Innovative Technology for Computer Professionals

Computer

February 2008

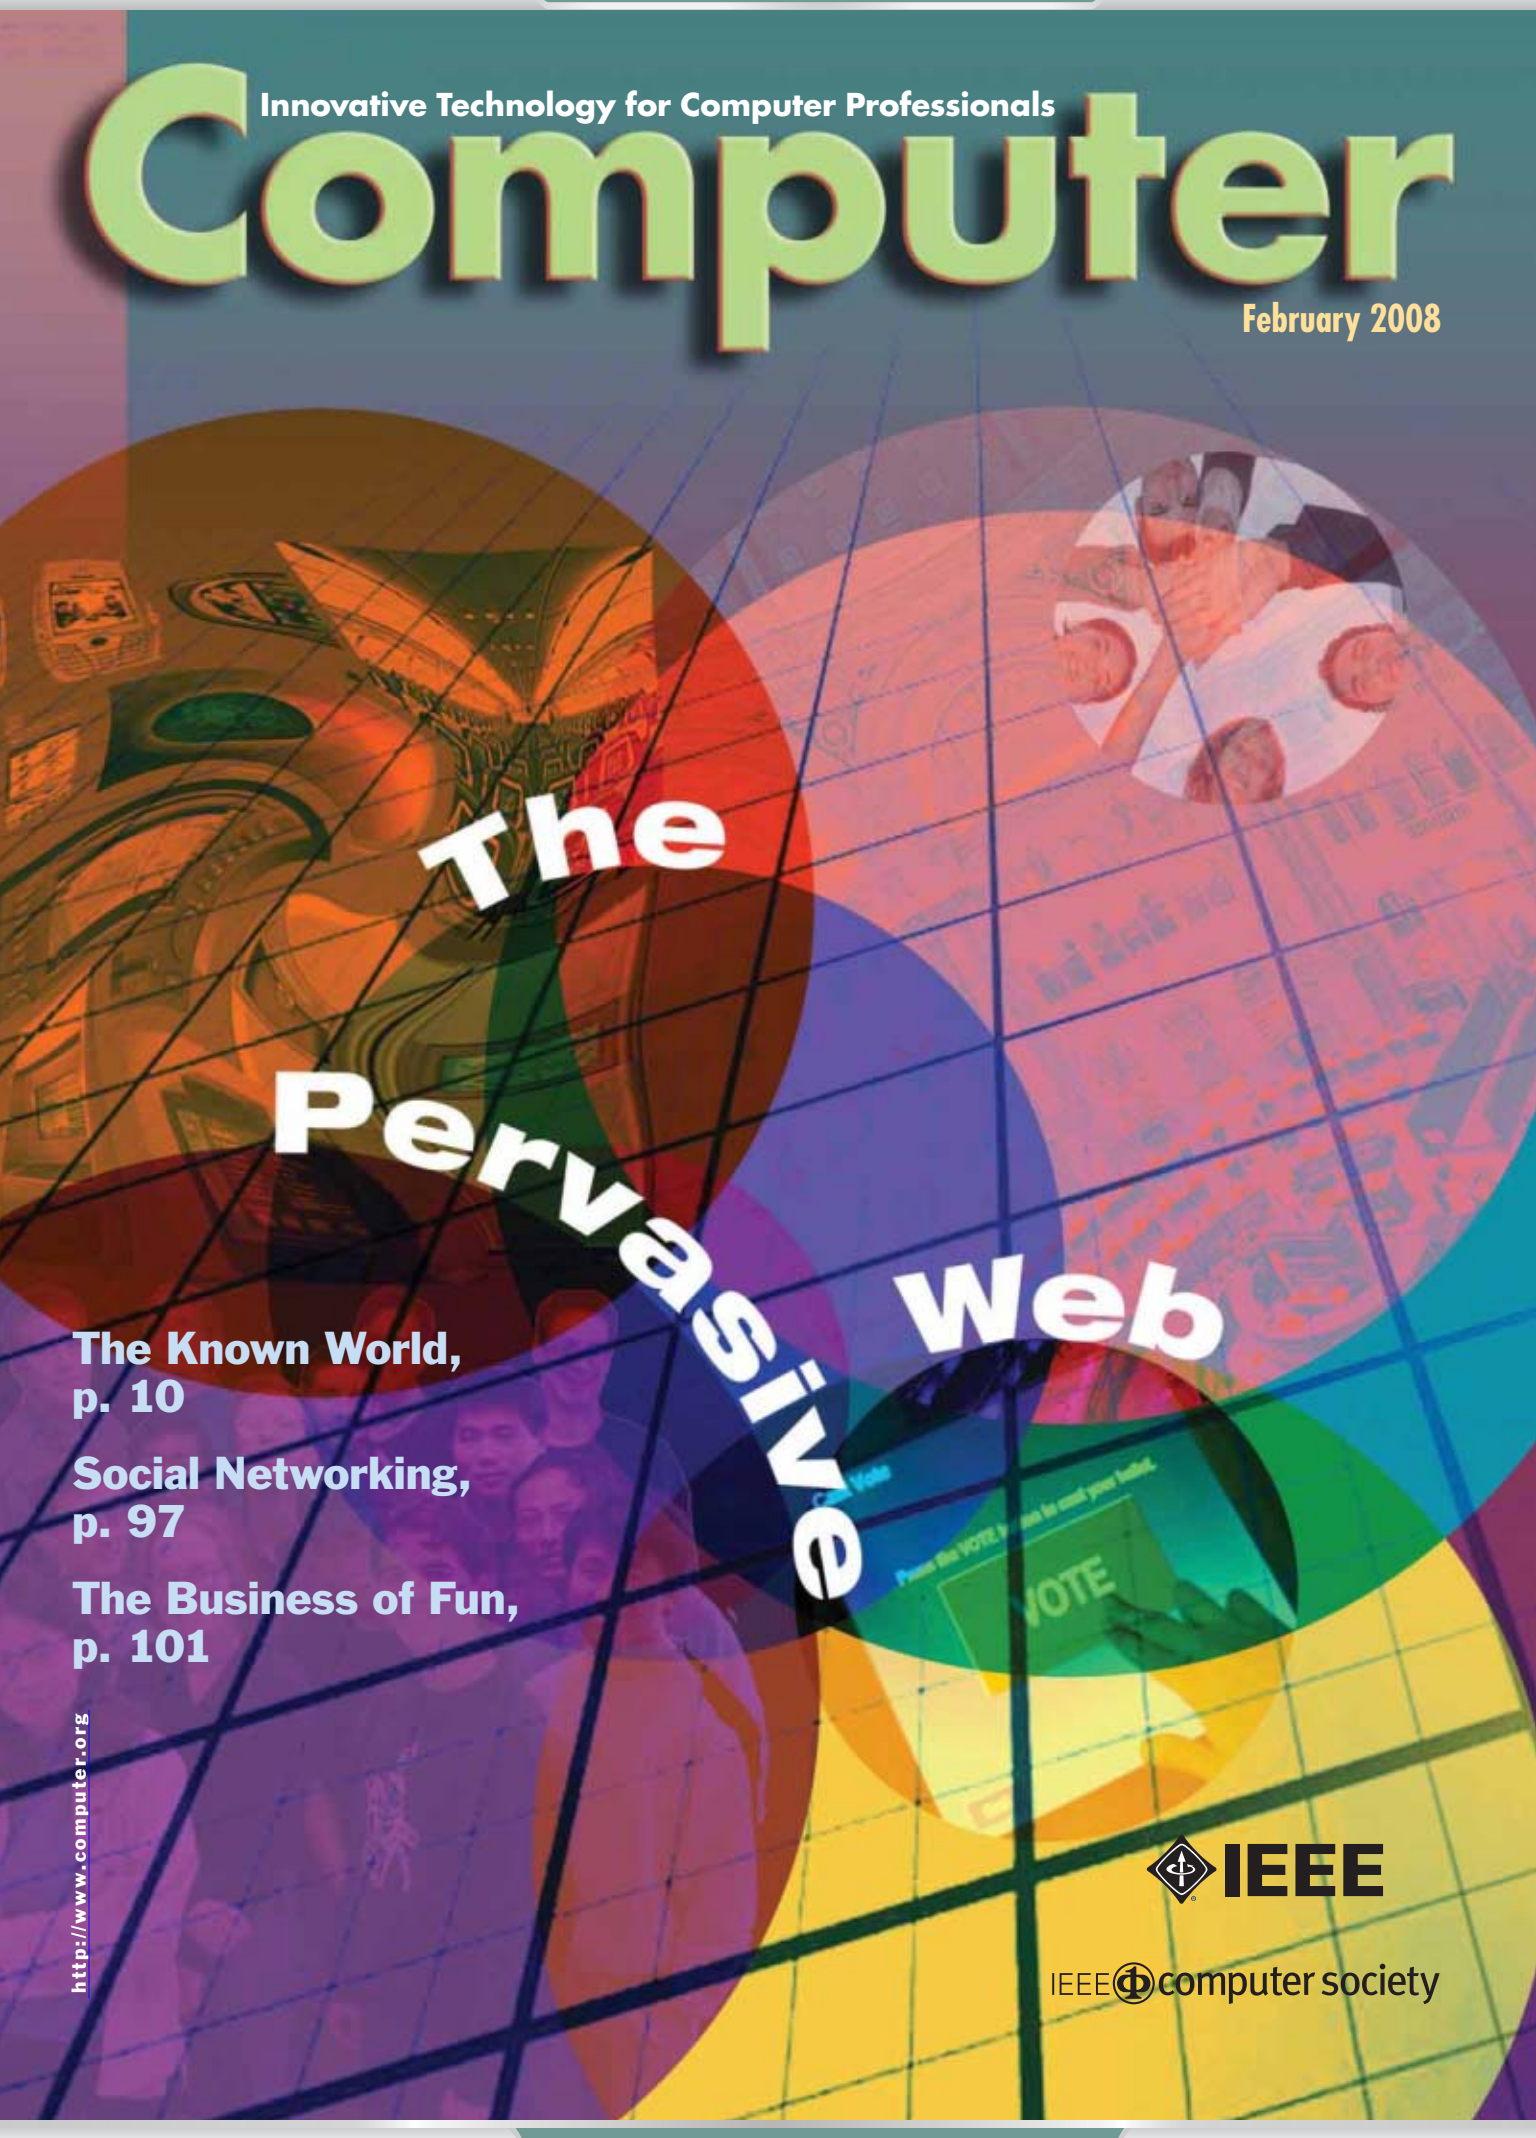

The Pervasive Web

http://www.computer.org

IEEE computer society

Newest Titles In Computing from Wiley & the IEEE Computer Society

Software Maintenance Management
Evaluation and Continuous Improvement

Alain April, Alain Abran

9780470147078 • March 2008 • Paper • 320pp • $65.00
Wiley-IEEE Computer Society Press

Written by experts in this area, *Software Maintenance Management* explores the domain of software maintenance in depth and describes the critical overlaps with software development. It presents a new way of capturing software maintenance activities in a model-based process-improvement approach. The focus of this book examines the uniqueness of software maintenance activities, teaches how to assess software maintenance using the Software Maintenance Maturity Model (S3m model), and explores paths for improvement.

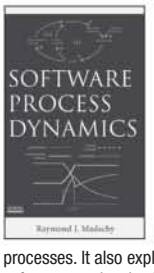

Software Process Dynamics

Raymond J. Madachy

9780471274551 • January 2008 • Cloth • 632pp • $84.95
Wiley-IEEE Press

Software Process Dynamics is an introduction to system dynamics and shows readers how to gather better information about interrelated technical and social factors to effectively improve the development process. This book reviews the field of software process modeling with system dynamics, and describes how others have used the principles to improve their processes. It also explains the modeling process (including calibration of models to software metrics data) and shows systems thinking-in-action by illustrating how to develop a deeper understanding of software process structures and behaviors.

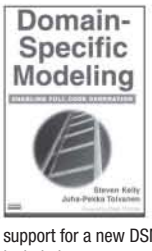

Domain-Specific Modeling

Steven Kelly, Juha-Pekka Tolvanen

9780470036662 • January 2008 • Paper • 448pp • $84.95
Wiley-IEEE Computer Society Press

This book illustrates examples from various fields of software product development and largely addresses the guidelines for implementing domain-specific modeling (DSM): how to identify the necessary language constructs, what options are available for code generation, and what tools are available to provide tool support for a new DSM language. The DSM example cases described in the book are included on an accompanying CD, paired with an evaluation copy of the MetaEdit+ tool to examine and use for modeling languages and code generators. Evaluation versions are included for Windows, Linux, and Mac OS X.

Emerging Methods, Technologies and Process Management in Software Engineering

Andrea De Lucia, Filomena Ferrucci, Genny Tortora, Maurizio Tucci

9780470085714 • February 2008 • Cloth • 296pp • $94.95
Wiley-IEEE Computer Society Press

Emerging Methods, Technologies and Process Management in Software Engineering covers emerging topics in the field of software engineering: requirement engineering, software system design, UML, software architectures, verification and validation, software configuration management, process management, empirical software engineering, and software evolution. Based on tutorials presented at the International School on Software Engineering (University of Salerno, Italy), the editors provide an overview of each section that places the material in perspective.

WILEY Publishers Since 1807 | IEEE computer society | IEEE

Scripting with Objects
A Comparative Presentation of Scripting with Perl and Python

Avinash C. Kak

9780470179239 • March 2008 • Paper • 1328pp • $89.95
Wiley-IEEE Computer Society Press

Scripting with Objects is based on the same overall philosophy as the author's previous title *Programming with Objects*—it demonstrates how a programming language is used through its applications to fully show its power. Designed for readers who want to acquire a more comprehensive and expansive perspective on scripting by exposure to two languages at the same time, this book takes a novel approach to a timely and important topic.

Software Development Rhythms
Harmonizing Agile Practices for Synergy

Kim M. Lui, Keith C. Chan

9780470073865 • March 2008 • Cloth • 320pp • $69.95 • Wiley

Software Development Rhythms: Harmonizing Agile Practices for Synergy builds upon the inherent flexibility of agile practices, focusing on understanding the "why and when" of the effective application practice-move-practice or activity-move-activity. The unique approach and technical quality of this work answers the key question of whether programmer productivity is impacted by the various agile practices, and provides a comprehensive and unbiased journey across software development rhythms.

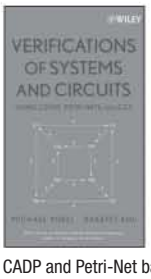

Verification of Systems and Circuits Using LOTOS, Petri Nets, and CCS

Michael Yoeli, Rakefet Kol

9780471704492 • March 2008 • Cloth • 160pp • $94.95 • Wiley

Part of the Wiley Series on Parallel and Distributed Computing, *Verification of Systems and Circuits Using LOTOS, Petri Nets, and CCS* provides computer science students and practicing logic design engineers with a step-by-step interactive introduction to formal verification of systems and circuits. This text makes use of two powerful analysis tool sets: LOTOS-based CADP and Petri-Net based PETRIFY. These systems cover alternating-bit protocol, arbiters, pipeline controllers, up-down counters, and phase converters.

FPGA Prototyping by VHDL Examples
Xilinx SpartanTM -3 Version

Pong P. Chu

9780470185315 • January 2008 • Cloth • 468pp • $84.95 • Wiley

This book uses a "learning by doing" approach to introduce the HDL (hardware description languages) and FPGA development process to designers through a series of hands-on experiments. A wide range of examples is included, from a simple gate-level circuit to an embedded system with an eight-bit soft-core microcontroller and customized I/O peripherals. All examples can be synthesized and physically tested on an actual FPGA prototyping board.

ORDER INFORMATION

1 (877) 762-2974 North America
+ 44 (0) 1243 843294 Rest of World
Log on to www.wiley.com
Enter Promotion Code CMPF8 to receive 20% off featured titles at checkout.

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, PO Box 3014, Los Alamitos, CA 90720-1314; voice +1 714 821 8380; fax +1 714 821 4010; **IEEE Computer Society Headquarters**,1730 Massachusetts Ave. NW, Washington, DC 20036-1903. 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.

Innovative Technology for Computer Professionals
Computer

February 2008, Volume 41, Number 2

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

Cover design and artwork by Dirk Hagner

ABOUT THIS ISSUE

Over the past decade, the Web has become all-pervasive, almost a physical manifestation in our social fabric. In this issue, we look at some maturing Web applications in education and knowledge management, entertainment and commerce, and politics. We also revisit the promise of reconfigurable computing.

COVER FEATURES

RESEARCH FEATURES

Flagship Publication of the IEEE Computer Society

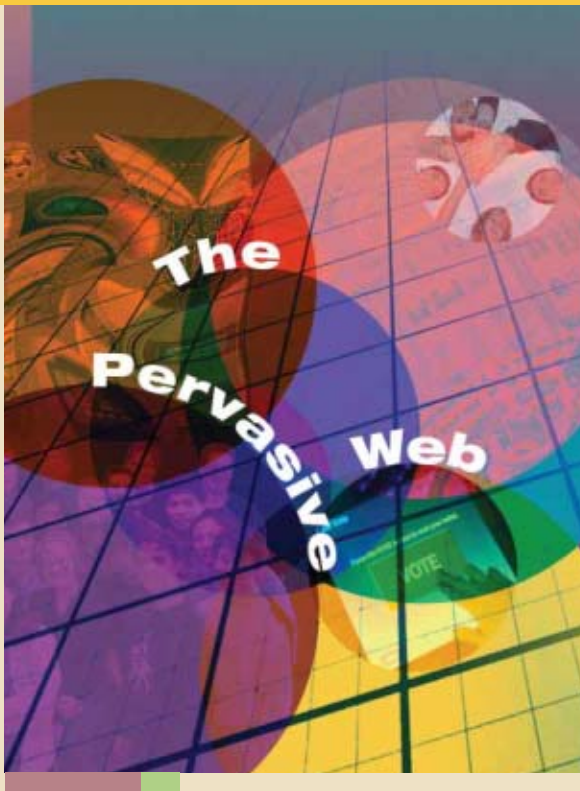

IEEE

IEEE computer society

ARTICLE SUMMARIES

National Science Digital Library: Shaping Education's Cyberinfrastructure

pp. 26-32

David McArthur

The National Science Digital Library's main portal, www.nsdl.org, has been online nearly five years. With this maturity, the National Science Foundation is now rethinking NSDL's status as a research program. NSDL has reached the point at which it must either change substantially or start winding down. Because NSF is primarily a research agency, investing further in NSDL would seem to run counter to NSF's policy of not supporting routine science and education operations.

Nonetheless, there are compelling arguments for NSF's continued investment in NSDL.

Wikis: 'From Each According to His Knowledge'

pp. 34-41

Daniel E. O'Leary

Given the explosive growth in wiki applications and the controversies surrounding the technology, it is useful to sort through the claims and criticisms to better understand what wikis are, how they are used, their advantages and limitations, and various issues surrounding their implementation.

Wiki is Hawaiian for quick and, as the term suggests, the technology's initial goal was to give users the ability to quickly put content on the Web. Wikis offer tremendous potential to capture knowledge from large groups of people.

EventWeb: Developing a Human-Centered Computing System

pp. 42-50

Ramesh Jain

Current interest in human-centered computing suggests new winds blowing in the computing community. HCC combines many powerful and independent approaches in different aspects of computing.

A combination of technological advances, a reduction in barriers to interactions among different parts of the world, and the quest for solving increasingly difficult problems has created a potential to impact the course of human civilization.

Secure and Easy Internet Voting

pp. 52-56

Giampiero E.G. Beroggi

Modern societies have thus far seemed hesitant to rely heavily on information and communication technology for democratic decision-making activities such as voting. One reason for the delay in implementing technologically sophisticated voting methods is the computer science community's almost unanimous wariness of Internet-based elections.

Fortunately, more countries are beginning to consider e-voting systems. However, three cantons in Switzerland—Zurich, Geneva, and Neuchatel—are already using an e-voting system.

Turning Teenagers into Stores

pp. 58-62

Srijith K. Nair, Bruno Crispo, Andrew S. Tanenbaum, and Ron Gerrits

Peer-to-peer file sharing has been immensely popular since 1999, when Napster began offering a central catalog of who had which songs so that people could directly copy them from the remote hard disks of people they didn't know, peer-to-peer file sharing has been immensely popular.

Initially, the music industry felt differently and viewed such file sharing as intellectual property theft. Now, music companies realize that digital music is their friend. This knowledge has led some music executives to dream of turning teenagers into stores.

Authorizing Card Payments with PINs

pp. 64-68

Václav (Vashek) Matyáš, Jan Krhovjak, Marek Kumpost, and, Dan Cvrcek

Chip and PIN technology for card-purchase authorization replaces card imprints or swiping cards with magnetic stripes through readers during face-to-face credit- or debit-card transactions. Signature verification secures both methods, with the clerk required to compare the customer's signature with the one on the card's back.

To resolve questions about this new authentication system's benefits, the authors conducted a two-part experiment to evaluate if it makes circumvention easier for thieves and disputing fraudulent transactions more difficult for customers.

The Promise of High-Performance Reconfigurable Computing

pp. 69-76

Tarek El-Ghazawi, Esam El-Araby, Miaoqing Huang, Kris Gaj, Volodymyr Kindratenko, and Duncan Buell

In the past few years, high-performance computing vendors introduced many systems containing both microprocessors and field-programmable gate arrays. In all these architectures, the main application executes on the microprocessors, while the FPGAs handle kernels with long execution times but that lend themselves to hardware implementations.

The authors' research revealed that HPRCs can achieve up to four orders of magnitude improvement in performance, up to three orders of magnitude reduction in power consumption, and a two orders of magnitude savings in cost and size requirements.

IEEE Symposium on Low-Power and High-Speed Chips

COOL Chips XI

CALL FOR PARTICIPATION

COOL Chips is an International Symposium initiated in 1998 to present advancement of low-power and high-speed chips. The symposium covers leading-edge technologies in all areas of microprocessors and their applications.
The COOL Chips XI is to be held in Yokohama on April 16-18, 2008, and is targeted at the architecture, design and implementation of chips with special emphasis on the areas listed below.

- **Low Power-High Performance Processors for Multimedia, Digital Consumer Electronics, Mobile, Graphics, Encryption, Robotics, Networking and Biometrics.**
- **Novel Architectures and Schemes for Single Core, Multi-Core, Embedded System, Reconfigurable Computing, Grid, Ubiquitous, Dependable Computing and Wireless.**
- **Cool Software including Binary Translations, Compiler Issues and Low Power Techniques.**

Dates and Location

April 16-18, 2008

Yokohama Joho Bunka Center, Yokohama, Japan

Yokohama Media & Communications Center, Yokohama, Japan

Keynote Presentations

- ***Satoru Ito***
  Chairman & CEO, Renesas Technology Corp.
- ***Hironori Kasahara***
  Department of Computer Science and Engineering, Waseda University

Invited Presentations

- ***Shorin Kyo***
  System IP Core Research Laboratories, NEC Corporation
- ***John Goodacre***
  Program Manager, Multiprocessing, ARM Ltd

Panel Discussion

- ***Josep Torrellas***
  Department of Computer Science, University Illinois, Urbana-Champaign
  Topics: Multi-core and Many-core (tentative)

Special Sessions
(invited lectures)

- ***Kevin Skadron***
  Department of Computer Science, University of Virginia
- ***Yan Solihin***
  North Carolina State University

For more information, please visit <**http://www.coolchips.org/**>

Sponsored by the Technical Committees on Microprocessors and Microcomputers and C Architecture of the IEEE Computer Society. cooperation with the IEICE Electronics Society, ACM SIGARCH and IPSJ.

Symposium Registration

In order to make a registration, please use COOL Chips XI REGISTRATION FORM on web site at

< http://www.coolchips.org/ >

REGISTRATION FEES

Registration Fees include a copy of the proceedings (copies of speakers' slides) of all plenary and technical sessions and special sessions presented on April 16-18, 2008.

	Early Registration by March 28, '08	Late Registration from March 29, '08
Member of any of IEEE IEICE, IPSJ or ACM	30,000 yen	36,000 yen
Non-Member	38,000 yen	46,000 yen
Student (Member)	10,000 yen	12,000 yen
Student (Non-Member)	13,000 yen	16,000 yen

PAYMENT AND REMITTANCE

For attendees from inside Japan:
- Fees should be prepaid with a bank transfer by April 11.

For attendees from outside Japan:
- On-site cash payment is highly recommended.
- Credit cards (only VISA) are necessarily acceptable.

Detailed information is available on the web site.

Notes:
1. Credit card charges will be billed in Yen.
2. Personal checks, bank drafts, and traveler's checks are not acceptable.
3. Registration is incomplete without payment of the registration fee.

Accommodations

HOTEL RESERVATIONS

A special group rate is available for COOL Chips XI attendees for Hotel Monterey Yokohama and Hotel New Grand.

CONTACT to

COOL Chips XI Organizing Committee Secretaries
E-mail:cool_XI@coolchips.org http://www.coolchips.org/

(As of January 10, 2008)

LETTERS

METADATA AND ONTOLOGIES

The article "Toward a Social Semantic Web" by Alexander Mikroyannidis (Web Technologies, Nov. 2007, pp. 113-115) displays a fundamental misunderstanding of the nature and limitations of both metadata and ontologies.

Cory Doctorow provided one of the best critiques of the limitations of metadata (www.well.com/~doctorow/metacrap.htm).

One limitation of an ontology is that it is a data model of entities over a domain (http://en.wikipedia.org/wiki/Ontology_(computer_science)).

Although the relational data model (E. Codd, "A Relational Model for Large Shared Databanks," *Comm. ACM*, June 1970, pp. 377-387) can be used to represent any model of data within a domain, it does not address the semantics of any database because

- defining an entity (relation) is arbitrary (W. Kent, *Data and Reality*, North Holland Publishing, 1978);
- partitioning an entity into a hierarchy (Codd's normalization) is arbitrary, and there is no a priori best hierarchy for this partitioning (W.S. Jevons, *The Principles of Science*, Dover Publications, 1874); and
- partitioning a set of sets (concept domain) into nonoverlapping subsets is an NP-complete problem and has no polynomial time-limited algorithmic solution. The best that can be done is to test a given partitioning to see whether it has overlapping subsets. If you do not require nonoverlapping subsets, then any arbitrary partitioning will do, but you will not be able to use it for reasoning about the domain.

This makes the creation of an ontology (from a folksonomy or any other source) an arbitrary exercise of the author, and it reflects all of the author's unstated assumptions and prejudices.

Rainer Schoenrank
rschoenrank@computer.org

The author responds:

Thanks for taking the time to read my article. I'd like to comment on two points:

You argue about the arbitrary nature of ontologies. As explained in the article, my proposal addresses this inherent drawback in ontologies by introducing tagging consensus into ontology construction.

You question the general contribution of ontologies and metadata in knowledge management. I need not argue against your views—many people have already done that for me. I suggest reading, for example, the following:

- J. Davies, D. Fensel, and F. van Harmelen, *Towards the Semantic Web: Ontology-Driven Knowledge Management*, John Wiley and Sons, 2003.
- M. Daconta, L. Obrst, and K. Smith, *The Semantic Web: A Guide to the Future of XML, Web Services, and Knowledge Management*, John Wiley and Sons, 2003.

Alexander Mikroyannidis
a.mikroyannidis@ieee.org

KILLER ROBOTS

Noel Sharkey might want to add the landmine to his collection of robotic killers (The Profession, "Automated Killers and the Computing Profession," Nov. 2007, pp. 124, 122-123). It's a very dumb robot and can't move about, but it will kill indiscriminantly—child, adult, soldier, farmer, friend, or foe.

The landmine models many of the problems Sharkey mentions pretty well, especially not having a human in the loop. Some suggested solutions include having a timer that deactivates the mine after some period of use. Maybe deactivation will be required for robots as well.

Charles J. Neuhauser
cneuhauser@earthlink.net

The author responds:

I agree that landmines are a type of robot, and their power to kill innocents long after wars have been fought is well-known and well-argued. The old antipersonnel mines will kill on contact and are like other reflex weapons such as the Navy CIWS. They fall foul of the notion of just war because they kill indiscriminately and there is difficulty in assigning responsibility for mishaps.

In 1997, 153 countries signed the Ottawa Mine Ban Treaty, but not the US, China, or Russia. President Clinton had planned to join in 2006, but George W. Bush abandoned the plan in 2004 because it would mean giving up a "needed military capability." The US policy was to move toward mines that self-destruct (*The Lancet*, 27 Aug. 2005; www.thelancet.com).

I did not include mines in the article because I wanted to focus tightly on creating discussions about new technological threats to humanity in the form of mobile autonomous weapons that will actually make decisions (so to speak) about who to kill. This is where my expertise may be of some use.

Unfortunately I have recently heard of a new breed of mine that is meant to "intelligently" determine friend from foe and fire torpedoes at the latter. These do come under my remit, and I am investigating further for future articles.

Noel Sharkey
noel@dcs.shef.ac.uk

DOING MORE WITH LESS?

Simone Santini's perspective on what we need for computers is painfully on target (The Profession, "Making Computers Do More with Less," Dec. 2007, pp. 124, 122-123). However, there is one fatal omission in his discussion of an ideal device.

Once you have a simple computer that does what you need, it is very

likely not to require replacement or need new software for many years. This is an economic disaster. Software vendors from Redmond to Rangoon and hardware vendors from Santa Clara to Beijing are addicted to regular fixes of money, which are inversely proportional to the life span of a given hardware-software platform.

When I get my One Laptop Per Child computer (see www.laptop.org), I may actually have that simple device. We will see.

Jim Isaak
j.isaak@snhu.edu

The author responds:

I assume that your definition of "economic disaster" is more than a bit tongue in cheek. Of course hardware and software vendors like the explosion of sales they provoke with every new release cycle, but they know (or they should know) that they are playing a risky game. They are generating such unreasonable get-rich-quick expectations in their stockholders that now if they meet expected revenues rather than exceeding them by a fat margin, their stock loses value.

What manufacturers don't like to hear is that after the initial transient in which every Tom, Dick, and Harry on the planet wants to buy a brand new product, they will have to accept a market plateau. The computer industry, it seems, is focusing on extending the transient beyond the limits of ridiculous rather than preparing for the plateau.

Well, what can I say? The day that software executives notice, sadly, that they can't afford a second private jet, I will shed a tear for them.

Simone Santini
simone.santini@uam.es

I read Simone Santini's article with great interest, and I agree with most of the issues he raised. In fact, an operating system that fits his description already exists. It's the IBM OS/2, which runs on a 25-MHz processor, can run with 8 Mbytes of RAM, and can install in 100 Mbytes of disk space.

The system was in active development from 1987 to 1996 (with participation of Microsoft until 1992, when it dropped out of the OS/2 effort to pursue its Windows product instead). At that time, processors were obviously much slower, and RAM was much more expensive. For this reason, most of the OS/2 kernel was written in assembler. But this does not prevent it from running on the latest dual-core Intel Core 2 and AMD Athlon64, and it runs *fast*.

I disagree with the comment about colors. At the time the windowing interface was being designed, IBM hired lots of psychologists to study visual perception effects on user interfaces. The outcome was the 1995 Common User Access standard in use today. CUA is based on a consistent user interface, not necessarily pretty, but that can use color icons and configuration notebook tabs.

The problem with application installation stems from the use of DLLs or shared libraries, which represent a major design decision by IBM because they help use RAM much more efficiently. As long as the system uses the same version of the DLLs for most applications, the DLLs stay in memory and are shared between processes. That's what makes the system lightweight and fast. Unfortunately, the price to pay for this is complexity of installation and updates. This was solved with the configuration-installation-distribution facility that is similar to package manager in AIX or RedHat Linux RPM, but predates it by a decade.

OS/2 still enjoys a small but devoted following, especially in Europe, but IBM stopped supporting it in 2003. Although the system has clear technical advantages over its rivals, it never was a commercial success. Today, there is an OEM distribution of OS/2 called eComStation (www.ecomstation.com).

Vadim Kavalerov
Vadim.Kavalerov@sig.com

The author responds:

From my point of view, the main point of this message is the statement that "the system was in active development from 1987 to 1996." At that time, I assume, things got out of hand (thanks, I suppose, to Windows 95), and the race to useless features and monster operating systems was under way.

I don't know OS/2 well enough to express a technical judgment on it but, if what you say is true (and I have no reason to believe it isn't), its demise is a good example of the trend I was criticizing in my article.

Now I would like to see a computer company with enough courage to produce a laptop with 300 Mbytes secondary memory and a 50-MHz CPU that is light, thin, and has a very long battery life running one of these stripped-down operating systems (and equally stripped-down programs). Even more, I would like to see a public with the culture, intelligence, and resistance to commercial pressure to make such a computer a success.

I partially buy the argument about color. Partially, because I don't think there is anything ergonomically significant that can't be done with 16 colors. Moreover, I stand by my opinion: If a black-and-white screen can buy me a couple of hours of battery life, I'll go for it!

Simone Santini outlined proposed guidelines for "a fast and efficient system." Are these guidelines realistic? Today, are computers really "meant to be for people who use a computer as a work instrument" or merely used to "write a mathematical paper?" A computer with a monochrome display and an operating system with less than 50 Mbytes storage is more like an early-stage terminal.

A computer is a complex piece of machinery consisting of many components, each of which is a separate invention. Within the past several decades, the speed and power of the computer have grown at an exponential rate. During that time, computers

LETTERS

Erratum

The correct e-mail address for Allen Stone, author of "Natural-Language Processing for Intrusion Detection" (Security, Dec. 2007, pp. 103-105), is astone@jasi.com.

have evolved from being primarily professional and business machines to become our primary entertainment and educational tools.

Computers have become the heartbeat of the modern world. They communicate. They act. They are our personal assistants. When we are surfing the Internet, participating in a videoconference with colleagues thousands of miles away, viewing with amazement 3D graphics for cars and architectural design, or exchanging e-mail messages, we can't imagine our lives without a computer being involved. Never in history has one invention had such an influence on humanity as a whole. However, without question, a "simple" computer can no longer fulfill our ever-changing demands.

The question is not whether we should make the operating system smaller but what exactly we should do to keep our computers running as fast as new. Performing regular tasks such as uninstalling old and unused software, performing disk cleanup, running hard-disk maintenance utilities, removing spyware/adware, and keeping the security software up to date are practical ways to keep systems running at peak performance.

Hong-Lok Li
lihl@ams.ubc.ca

The author responds:

As I understand it, your point is that computers today perform many functions, some of which require fast CPUs and—this point is more doubtful—large operating systems. Your underlying assumption seems to be that all computers should do all things.

Consider, as a parallel, motor vehicles. People who must transport a heavy load drive 18-wheeler trucks. This, of course, doesn't imply that every activity performed with a motor vehicle requires an 18-wheeler, or that everybody should buy one. Sometimes a small two-seater city car is the perfect solution for a given transportation problem.

The same applies to computers. Some people use them to watch videos or to "view with amazement 3D graphics" (should I suppose that if the people were not amazed, the requirements for the operating system would change?). Other people use computers to write reports and calculate simple spreadsheets. There is no reason why these two groups of people should use the same machines, the same operating systems, and the same programs.

The fact that a computer is a complex piece of machinery built of many components is utterly irrelevant: so is a skyscraper, and so is a car. But, as we have seen in the case of motor vehicles, the device's complexity doesn't mean that we should adopt a "one-size-fits-all" model.

Two statements in this message reveal a profound philosophical and attitudinal difference between the two of us. One is, "Computers have evolved from being primarily professional and business machines to become our primary entertainment and educational tools." There has, undoubtedly, been a change, but why is going from business applications to entertainment an "evolution" in the use of computers? It's a diversification, certainly, but considering it an evolution seems a trifle naïve.

Second, there is the somewhat triumphalistic observation: "Computers have become the heartbeat of the modern world. They communicate. They act." I assure you, they do not. We communicate, we act. Computers do not communicate any more than a telephone or a letter does. Computers are versatile instruments, useful for certain things, not so much for others.

I must confess that I am worried when I see such a triumphalistic attitude among academicians: We should value critical evaluation and detached analysis. There are already plenty of people out there who can write marketing brochures, and there is no need for us to join their ranks.

GREEN COMPUTING

I was pleased to see the articles in *Computer*'s December 2007 issue covering green computing in various forms. I look forward to the day when the IEEE gets real about the environment and makes its own contribution by offering totally paperless membership.

I have a comment on the practicality of the idea of recycling silicon (Oliver et al., "Life Cycle Aware Computing: Reusing Silicon Technology," pp. 56-61). I'm not sure if "recycling down the food chain," as the authors propose, is practical.

Moving to devices with lower computing requirements also often means moving to bigger markets. The example in this article suggests recycling a PDA processor in a GPS system, and later in a Nintendo DS. Numbers I dug up on the Net suggest that each cheaper device in this list has about a factor of four times the sales of the device a level above. What's more, the sales of the cheaper devices in this example are increasing faster. Also, a few years down the track, a lower-cost alternative for the cheaper device will probably be available.

This kind of recycling would also require factoring in the energy costs of dismantling the device (difficult with components designed to be used once), the shrinkage of damaged components from disassembly, and the higher failure rate of devices that have already seen significant use.

It might, however, be an option to ship off obsolete or recycled parts to poorer countries where labor costs are low and create a cottage industry in building low-end but functional computers with low power demands. Many PDAs, phones, and the like

easily have enough processing power to run a stripped-down free operating system like Linux.

This has more appeal to me than the One Laptop Per Child project, which is based on the misconception that owning a computer is in itself an advantage. If thousands of people in poor countries had the direct experience of building computers and massaging software to install on unusual configurations, the skills gained would be a huge boost to the local economy.

Philip Machanick
philip.machanick@gmail.com

The authors respond:

The letter writer raises some interesting issues relating to silicon reuse.

With respect to the volume of lower-end devices in a "food chain," it certainly might not be possible to supply enough recycled parts to meet demand. The goal, however, is to get more use out of the high-end devices and forestall their disposal in landfills.

Recycling costs are also definitely a concern, and our current research focus is on recycling entire systems (such as mobile handsets) instead of individual chips. The cottage industry in "poorer countries" has been suggested, but it should be noted that our industrial collaborators find this a sensitive issue and that, as the writer somewhat suggests, it is important that the recycled technology be an enabler for new applications.

Fred Chong
chong@cs.ucsb.edu

We welcome your letters. Send them to computer@computer.org.

THE KNOWN WORLD

Derek the Rocket Scientist

David Alan Grier
George Washington University

Trust is a difficult thing to engineer, as it involves history, character, and an ability to put aside your own goals for the good of the whole.

Many couples have testy discussions about golf. Or tennis. Or rock climbing. Or soccer. But my friends Derek and Sara are the only couple I know who have had an awkward conversation about rockets.

Sara and I have been friends since childhood. Our fathers worked together at the old Burroughs Corporation. When she told me that she was marrying a rocket scientist, I was both pleased and intrigued. I was happy to see her find a partner and interested to learn what sort of person a rocket scientist might be.

Like many an individual who has followed the American space program, I speculated that someone who sends machines beyond Earth's atmosphere might have a special view of the world, that earthly concerns might have a celestial tint. I quickly found that the opposite was true, that celestial activities were tied to those of common life.

"I'll be back in time for dinner" were the first words I heard Derek say.

The three of us had agreed to meet for coffee at a restaurant near Sara's apartment. I was visiting their city on business and had taken an extra day to meet Derek and to share a year's worth of news with Sara. Both Sara and Derek had arrived before I found the place, and the two were already in the midst of a vigorous discussion.

"It's two hours out there," said Sara, "and two hours back. You won't be home for dinner."

"I'll leave early," replied Derek, "I'll get back in plenty of time."

At this point, I was certain I understood the conversation between the two of them. The subject was clearly golf. Derek and his buddies had a tee time for noon at some distant course. Sara doubted that her fiancé could finish the game, drive home, take a shower, and arrive on time at the place where we were all to have dinner.

I was so confident of my judgment that I made a misguided attempt to defuse any anger. "Playing a round this afternoon?" I asked after we were introduced.

"No," replied Sara, "he's going over the mountains to fly rockets."

"Oh," I said, grasping the fact that Derek the rocket scientist was also Derek the rocket hobbyist.

"I finished this one last weekend," Derek interjected. "We think we can get it up to 35,000 feet."

There was not much more to say, though many words continued to be spoken. As a couple, Derek and Sara were still in a relatively early phase of their relationship. They had not yet combined households, merged their social goals, or even found a common way of discussing their needs.

Their conversation analyzed every aspect of Derek's schedule for the day, the potential problems, and the need to be back in the city at a certain time. It ended, as such things do, on an awkward note.

Derek was resolved to get his rocket flown. Sara was not convinced that he was going to be able to arrive on time for our reservation at the Bayside Restaurant at 6:30 p.m. We could only hope that all things would work for good and that the day's events would not be disrupted by poor planning, a shifting wind, or the happenstance of traffic congestion on I-95.

ENGINEERING TRUST

Derek the rocket scientist is technically a systems engineer. He oversees the design and construction of digital systems for large spacecraft projects. His job is to make the individual boxes that constitute the spacecraft electronics work as a whole. The process is tricky and fraught with problems.

He can specify the basic functions of a system in planning documents, but such documents still leave a great deal of discretion to the people who are building the basic components. "You can't design everything yourself," he explains. "You have to trust the other engineers to articulate what they can do, explain the shortcomings of their design, and suggest what can be done to make a system work."

Trust is a difficult thing to engineer, as it involves history, character, and an ability to put aside your own goals for the good of the whole. When are participants pursuing an unusual solution that will benefit the project, and when are they only

indulging their own curiosity at the group's expense?

Entering into a project and assuming that all parties automatically trust each other is like entering into a marriage without acknowledging the fact that you are about to share a bathroom with someone who has a very different idea of how such a facility should be used. It's all too easy to claim that the gender that sees little reason to return the toilet seat to the horizontal position has a deep character flaw or that those who believe that shower rods were invented as a place to dry underwear have a moral shortcoming. But such claims do little to help a marriage meet its design goals and do nothing to build trust.

EXTREME SILLINESS

Once, while I was sitting in Derek's office talking with him about nothing in particular, he picked up a cable from a side table and threw it across his desk. "This," he said, "is an example of extreme silliness. No one is willing to take responsibility for the cable and its problems. My component group says, 'That's a cable, and we don't take responsibility for cables,'" remarked Derek. "But the cable guys say, 'That cable contains a transformer, so it's a component and not our responsibility.'"

We often try to solve such problems with an overall plan, but such plans can have failings of their own. You can establish a budget for power, time, or weight and make the different units compete for their share, but in the end you might have a device that only fosters more problems.

Derek once worked on a large project that divided a satellite's outer skin and tried to make sure that each component group got the share it needed. The component groups fought hard to get prime real estate on the spacecraft. "In one design, two key devices were only 10 degrees apart," Derek said. It was an efficient use of the skin, but it posed a serious problem. "One device was a navigation sensor that had to find the sun. The other was a scientific experiment that would have been destroyed if it absorbed large amounts of solar radiation." Apparently, for a time, each group tried to blame the other for the mistake. Neither was willing to trust someone else to find a solution.

Silliness is, of course, in the eye of the beholder. Irresponsibility is also not easy to identify. When you're close to a design, you usually can't imagine that others don't see it as you do. In engineering organizations—indeed in many organizations—we try many different techniques to get people to articulate their needs, to listen to each other. Ultimately, we hope to understand and trust each other. We have team social gatherings, take our staff on retreat, play team-building games, give everybody imprinted sportswear with a common logo. Such things might not directly build trust, but they give team members a common experience, and common experiences can be used to build trust.

Most often, trust is built upon history and common experience.

REESTABLISHING TRUST

In the spring of a difficult war, the engineer and aviator Charles Lindbergh looked to common experiences to rebuild some trust in his marriage. Lindbergh was living apart from his wife so that he could work on military aircraft, but he was also feeling that his marriage was distant and strained. Their relationship had been rubbed raw by 15 years of intense public scrutiny. Every activity of Lindbergh and his wife had been reported in the press. After the Japanese attack on Pearl Harbor, that press was especially critical of the couple because of their isolationist politics, a stance that argued that the US should stay out of the war.

Writing from a hotel room outside a Michigan bomber factory, Lindbergh took a small step toward reconnecting with his wife. "The moon is at Bathurst takeoff height," he wrote to her. Such words could be taken as an attempt at poetry, an effort to charm his wife through pretty words, but they were actually a shorthand reference to a common experience, a time when they had learned to work together.

In the early years of their marriage, they had flown all over the world to survey air routes. They had gone to Europe, Asia, and Africa, where Bathurst could be found. Even though Lindbergh was the skilled engineer, he had split the work with his wife. In particular, she was responsible for navigation and communication. At Bathurst, they had a difficult departure and had to work together to get the plane airborne. Each had to shoulder responsibility. Each had to trust the other. No one else could help.

The dynamics of a marriage might be either easier or harder than those of an engineering team. A marriage can be sustained by social forces that are never quite captured by a contract or a design deadline. Occasionally, as in a national emergency, in the pursuit of a grand goal, or in the mere joy of doing something for the first time, an engineering team can find an extra bit of commitment that binds a group or team together, but most often, trust is built upon history and common experience.

DIFFERENT RESULTS

Derek the rocket scientist often struggled to find enough of those qualities to hold a project together. "We were once working on a fairly complicated spacecraft," he told me, "for which we contracted with two different firms and got two very different results." Both firms followed the specifications and both built instruments that ultimately worked. However, one firm devoted most of its resources to building their sensor

itself, while the second spent more time thinking about how their device would communicate with other satellite systems.

The first firm "created a 'roll your own' interface for its sensor," noted Derek. "They argued that they had no incentive to follow one of the military or civilian network standards, as they would never build enough of the sensors to recover their investment. It's not as if they were building a laser printer and could recover their costs over a production run of 100,000 units."

The firm delivered its device well before the deadline, but Derek's group could not make it talk with the other satellite systems. "We had a hard time making that thing work. We continued to debug that interface after we delivered the satellite to Cape Canaveral. We were working on it until liftoff."

"The second company spent less time on its instrument, but it connected its work to a standard interface. Its machine was late. The company had to deliver the sensor directly to the Cape, but we plugged it into the system, and it worked fine the first time."

"We weren't worried about it," Derek reported, "because it talked in a standard way. We could trust it."

WHAT'S IN IT FOR ME?

Not all group events are common experiences and not all of them build trust. Early in my career, I was part of a software group that went on a "team and trust-building" retreat. We listened to a talk about the company. We played a few games. We were encouraged to speak freely about our feelings. We all were asked to wear tennis shirts marked with the company logo and some heartening phrases, such as "Best software in the known universe."

At some point, I noticed that the women in the group were not especially enthusiastic about the day's events. When I asked how they were feeling, I got a bit of an earful in return. "Someone hasn't thought much about what we're doing," remarked one of my colleagues. "We're wearing men's clothing and playing men's games. I feel like I'm wearing a flour sack and being the sympathetic soul at a high school track meet. What's in it for me?"

On that day I met Derek the rocket scientist, I think he probably had a clear idea of what "was in it for him" and for Sara in their relationship. I don't remember if he arrived at the restaurant on time that night or if he had had a shower, or if the moon hung over the horizon, but I do recall it as a wonderful evening.

David Alan Grier, *an associate professor of International Science and Technology Policy at George Washington University, is the author of* When Computers Were Human *(Princeton University Press, 2005). Contact him at grier@gwu.edu.*

FEBRUARY 1976

INTERNATIONAL DATA TRANSMISSION (p. 3). "Western Union International has received approval from the Federal Communications Commission to furnish its International Digital Data Service to France, Italy, Spain, and Austria. The service is claimed to provide high quality and reliability, primarily by simultaneous transmission of data via undersea cable and satellite. The FCC approval is for one year of operation, within which IDDS must submit an analysis of its service. IDDS will be required to submit a tariff proposal before beginning service, based on transmission speed (from 50 bits to 9.6 kilobits per second), transmission volume, and other factors."

[Update: "WUI and ITT Datacomm Applications before FCC," p. 3.]

ATM ADOPTION (p. 3). "While the EFTS Commission continues to organize, a recent survey … indicates the extent and distribution of existing EFT systems. … 'The survey, which had a response rate of 97 percent of 4700 national banks, showed that fully 10 percent of the banks had at least one Automated Teller Machine. As expected, a high proportion of large banks have an EFT system—72.9 percent of billion dollar banks and 48.4 percent of those in the half billion to billion dollar range. However, more than half of all EFT systems are in banks with under $100 million in deposits. …"

[Update: "EFTS Survey Released," p. 3.]

COMPUTER CONSUMERS (p. 8). "In today's strong consumer market, the computer industry must be responsive to the message that is coming from the consumer sector and design products to meet their demands. Products needed include banking systems that are designed for the convenience of the customer and not only for the banker, and point-of-sales terminals that have digital displays on the side of the consumer as well as on the side of the cashier. … Machines should be easier to use. The need to have a computer expert at every application site must be diminished by the introduction of systems that are easier to use and do not require the complex error-prone programmer interventions to make them operate. The system of the future must be designed for the end user—i.e., the consumer."

[C. Adams, "Over the Horizon: A Report on the June 1975 Computer Elements Technical Committee Workshop in Vail, Colorado," pp. 8-11.]

DESIGN AUTOMATION (p. 12). "Design automation systems have made large-scale integrated electronics feasible for commercial production. Without computer-aided design, simulation, test data verification, and technology checking programs, manual LSI product-design cycles and error rates would exceed permissible product development times and costs. …

"The strategic goals of a DA system are to (a) provide a method for expeditious completion of a design; (b) formalize design methodologies to assure needed design discipline, integrity, and completeness; and (c) provide for orderly release-to-manufacturing procedures. …"

[W. Rosenbluth, "Design Automation Architecture and Applications," pp. 12-17.]

HIGH RELIABILITY (p. 18). "Just before the turn of this century, the newly invented Bell telephone began to receive widespread use. The early telephone network required many wires to carry conversations, and … these wires filled the sky at an alarming rate. An automatic dial telephone switch invented by an undertaker called Strowger further enhanced the proliferation of the telephone. Along with the invention of telephone switching came the corresponding invention (or curse) of devices to record subscriber billing information. Early billing accumulation devices consisted of electromechanical counters which were incremented when the subscriber used his telephone. Today, these counters are still widely used, but new semiconductor components for the first time make electronic replacement programs attractive."

[J.C. McDonald, "Testing for High Reliability: A Case Study," pp. 18-21.]

VIRTUAL MACHINES (p. 38). "The introduction of computers which have user alterable microprograms presents users with both an opportunity and a problem. To take full advantage of the opportunity to tailor the architecture of the computer to the application domain, the problem of microprogram development and testing must be solved. The technique outlined here provides a tool which furnishes the microprogrammer with a virtual machine which is microprogrammable. The function provided is similar to that provided by a simulator, but the technique is more efficient and allows multiple microprograms to be executed concurrently with regular programs on a single real machine."

[J.D. Bagley, "Microprogrammable Virtual Machines," pp. 38-42.]

UNSTRUCTURED PROGRAMMING (p. 47). "… My assertion is that machine code programs for commercial and administrative applications in the early days were sometimes astonishingly complex …; that such programs necessarily exhibited a high degree of structure; that, when theoreticians were later faced with designing programs whose complexity … matched that of earlier real-life programs, they naturally found themselves evolving the same structural disciplines; and that then, in their ignorance, they assumed that all programmers of the machine-code era had proceeded in the same undisciplined manner as they themselves had."

[J. Inglis, The Open Channel: "The True History of Unstructured Programming," pp. 47-50 (reprinted by courtesy of the Editor of the *Computer Bulletin*).]

FEBRUARY 1992

VISION SYSTEMS (p. 9). "A typical vision system requires integrating algorithms from diverse areas such as image processing, numerical analysis, graph theory, artificial intelligence, and databases. There is no clear understanding and consensus on how to achieve this. Specific problems in integration can also be attributed to a lack of understanding of the vision process itself, even if the computations and parallelism of some individual components are well understood.

"Recent efforts in architectural design and development have embedded architectural components for each level of processing into one integrated architecture … Compared to the progress in architectural advances in general-purpose parallel processing for other scientific disciplines, however, architectural advances for vision systems are in their infancy."

ANIMATE VISION (pp. 12-13). "… *Animate vision* researchers, inspired by successful biological systems, seek to develop practical, deployable vision systems by discovering and exploiting principles that link perception and action. Animate systems use active vision and are structured as vertically integrated skills or behaviors, rather than as visual modules that try to reconstruct different aspects of the physical world.

"Despite the computational simplifications of the animate vision paradigm, a parallel implementation is necessary to achieve the required performance. Fortunately, many of the tasks in an animate vision system are inherently parallel. … Thus, finding parallelism in the application is easy. However, the type of parallelism we would like to exploit varies among tasks in the system, and no single model of parallel computation is likely to suffice for all tasks."

IMAGE PROCESSING (pp. 22-23). "Because image processing is such an important application area for parallel computers, it makes sense to provide [programming] languages for this area. Apply and Adapt are both languages that make it possible to write certain types of image-processing operations, while providing the highest possible level of architectural independence. …

"Apply is a simple, architecture-independent language for *local* image-processing operations. Local operations produce an output pixel based on a small window surrounding the corresponding input pixel; they include edge detection and smoothing. …

"Adapt allows the definition of local operations as well as *global* operations, in which an output pixel can depend on many or all input pixels. Histogram and other feature-extraction operations are global operations. Adapt is based on the split-and-merge programming model."

GEOMETRIC HASHING (p. 33). "Developing realistic vision systems that can recognize rigid objects from a database of hundreds of models is a continuing goal of vision researchers. A model-based vision system extracts features such as edges and points from digital imagery and compares them with a database of models to identify objects within a scene. Many model-based vision systems are based on hypothesizing matches between scene features and model features, predicting new matches, and verifying or changing the hypotheses through a search process. A new method, called *geometric hashing*, offers a different and more parallelizable paradigm …"

PARALLEL PROGRAMMING (p. 54). "The 'need for speed' has been the most influential factor in supercomputer design. In the past, technology fueled the development of faster computers through better semiconductor devices and very large scale integration (VLSI). Technology, as a source of speed for a single processor, is bounded by the speed of light and physical limitations on miniaturization. Consequently, it has become necessary to replicate hardware to allow concurrent execution to achieve the performance requirements of many of today's scientific and industrial applications. This concurrent execution, or parallel processing, has forced the reformulation of the most well-accepted sequential programs and even the mathematical rethinking of some problems. The parallel programmer needs to 'think parallel.'"

GRACE HOPPER (p. 84). "Rear Admiral Grace Murray Hopper, pioneer computer programmer and co-inventor of Cobol, died January 1 at her home in Arlington, Virginia. Known as the mother of computerized data automation in the US naval service, she retired from the Navy in 1986, having been the oldest military officer still on active duty. …"

WORK SCHEDULING (p. 94). "Microsystems Software has released four WAN versions of its CaLANdar workgroup-scheduling software. The package uses standard e-mail packages to distribute activity and availability information.

"DOS users can distribute CaLANdar activity through in-place interserver gateways in a multiserver environment. The package allows Microsoft Mail, cc:Mail, and Banyan Vine Mail administrators to maintain CaLANdar user lists within their e-mail naming system."

PDFs of the articles and departments in the February 1992 issue of Computer *are available through the Computer Society's website: www.computer.org/computer.*

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

INDUSTRY TRENDS

US Cell Phone Industry Faces an Open Future

George Lawton

The US wireless industry appears to be in the midst of a sea change in the way it does business. It is just beginning to move away from the traditional *walled-garden* approach, in which users can pick their service providers but the carriers determine the phones and applications that customers can use. This gives service providers control of the wireless experience.

The US is a much more tightly controlled and closed market than Asia or Europe, noted Phil Marshall, an analyst with the Yankee Group, a market research firm.

Now, the US cellular industry seems to be moving toward a future in which customers have many more choices of the wireless devices and software they use and thus more control over their mobile experience. In the process, service providers would no longer control activities such as the sale of phones for their systems.

Service providers have not accepted this easily. Carriers such as Verizon Wireless have maintained that controlling customers' use of phones and applications lets the company optimize the wireless experience and ensure that its network and equipment remain reliable.

However, several current and potential service providers, including Google, as well as a growing number of public interest groups—such as Free Press, a nonprofit media-reform organization—maintain that this practice limits choice and

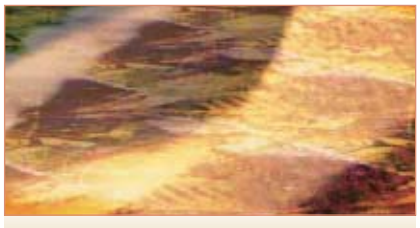

inhibits competition and innovation. They say this keeps the mobile online experience inferior to the PC-based online experience, in which users can choose their computers and applications.

The US move toward openness could take a significant step forward with the US Federal Communications Commission auction of a block of radio spectrum highly valued by service providers. For the first time, after pressure from Google, the FCC is requiring that purchasers of the block provide *open access*, meaning they can't restrict the types of commission-certified devices that use the spectrum.

Google has said it will bid on the FCC's open-access spectrum and develop an open mobile platform that would work with multiple devices and applications.

In the wake of these developments, Verizon has announced plans to open its network to new devices and applications.

OPEN ACCESS DRIVERS

Various factors are encouraging open access, including government initiatives such as the pending FCC spectrum auction; the success of Apple's iPhone, which the company says will soon have an open API that lets independent programmers freely develop applications for the device; and a push for openness by the influential and well-funded Google.

FCC auction

Starting 24 January 2008, the FCC officially began an auction of five blocks of radio spectrum in the 700-MHz frequency band—totaling 62 MHz in width—that providers can use for voice and data services. In the US, UHF TV stations will abandon this spectrum when they move from analog to all-digital broadcasts by February 2009, as federal law requires.

US cell phones currently don't use the spectrum, the last large amount available in the US for the foreseeable future, according to Mark Gibson, senior director at Comsearch, a spectrum-management consultancy.

The 700-MHz band is lower in frequency and thus has signals with a longer wavelength than the 800-, 1,900-, and 2,400-MHz bands that US cellular services currently use. Therefore, 700-MHz signals could travel farther in rural areas and penetrate buildings better in urban areas. This could translate into more bandwidth and fewer base stations, thereby reducing new providers' infrastructure costs.

Open access. After the FCC announced it would auction off parts of the 700-MHz spectrum in 2006, Google pressed the commission to require open access for two of the blocks.

Proponents say that wireless broadband is the next frontier of innovation in communications and computing but that the US is falling behind Asia and Europe in this area. They contend that open access has enabled competition and innovation in the wired Internet and thus is also necessary for the wireless Internet.

However, some service providers, such as Verizon, and industry orga-

nizations, such as CTIA-The Wireless Association, have opposed open access.

"We are concerned that a significant portion of this valuable spectrum will be encumbered with mandates that could reduce the number of interested bidders," said CTIA CEO and president Steve Largent. "We remain committed to the principle that wireless consumers and American taxpayers are best served when such a valuable commodity is auctioned … with no strings attached."

The FCC eventually mandated open access for the 22-MHz-wide C-block, which it will distribute via 12 regional licenses. The other blocks to be auctioned—including one for public-safety communications—will not require open access.

"A network that is more open to devices and applications can help foster innovation on the edges of the network," said FCC chair Kevin J. Martin recently. "And it will give consumers greater freedom to use the wireless devices and applications of their choice."

The 22-MHz-wide C-block is particularly desirable because it transmits more data at any one time than typical narrow blocks and thus increases the available bandwidth.

The FCC requires that providers winning the C-block licenses build enough infrastructure to reach 40 percent of the population in the areas they would serve by 2013. This could be costly, particularly for new operators with no local infrastructure, said Comsearch's Gibson. The failure to do so could result in fines.

For all blocks, the bidding will be closed, so no company will know what the others are offering. The minimum acceptable total of high bids for the open-access licenses will be $4.6 billion. The FCC has said that if it doesn't receive qualifying bids totaling at least this amount, it will auction the licenses again without the open-access requirement.

Verizon filed, but then withdrew, a lawsuit against the FCC's open-access decision.

Bidders. Two hundred and sixty-six companies submitted notices of the intention to bid for 700-MHz licenses, including Alltel Wireless, AT&T Mobility, Chevron, Cox Wireless, Google, Qualcomm, and Verizon.

The FCC doesn't allow bidders to publicly discuss the auction until it concludes and the commission announces the winners, probably in March of this year.

Openness could take away US carriers' control of the wireless experience.

A number of significant telecommunications-related companies—including Comcast, DirecTV, Nokia, Sprint Nextel, Time Warner, and T-Mobile USA—decided to sit out the auction.

New business models

New business models, such as the one Apple has used with its iPhone, will be required to promote openness in the US cellular industry, said the Yankee Group's Marshall.

The iPhone, released last year, is sold in the US only by AT&T Mobility. However, instead of allowing AT&T to provide all services, content, and applications, the iPhone lets users purchase and download content directly from, for example, Apple's iTunes Store via Wi-Fi.

Nokia, working with the Open Mobile Alliance, is creating an open ecosystem with application developers, content creators, phone vendors, and service providers to support its S60 smart phone platform, which is based on the Symbian mobile operating system.

GOOGLE MOVES IN

Google's foray into the US cellular market has already created the potential for industry openness.

Google's motivation

By entering the cellular market, Google could make money as a service provider. However, John Gauntt, a senior analyst with eMarketer, an industry analysis firm, speculates that Google's main goal is to extend the advertising business model it has used so successfully with its search engine during the past few years—with ads running next to search results on related topics—into the mobile online world. As Figure 1 shows, eMarketer expects mobile advertising revenue to grow substantially during the next few years.

According to Gauntt, Google could extend its ad-based model by directing wireless users to local businesses via geographically targeted advertisements.

Although Google is well-funded and has submitted its intention to bid for 700-MHz spectrum, winning the license may be expensive, particularly because the company's competitors also have a lot of money. In addition, unlike existing carriers, Google would have to design and build an infrastructure from scratch, noted the Yankee Group's Marshall.

Google thus may not make a serious bid if corporate executives think the wireless market is already opening up, predicted Paul Gallant, telecommunications analyst with the Stanford Group, a market research firm. This would give the company a chance to compete by providing content or services without having to become a carrier, he explained.

If not, he said, "I think Google will make a serious push to win the auction. Google doesn't want to cede the mobile Internet to Verizon and AT&T."

Android

In November 2007, Google announced formation of the Open Handset Alliance, consisting of 34 companies such as chip makers Intel and Qualcomm, phone manufacturers LG Electronics and Motorola, wireless carriers Sprint Nextel and T-Mobile, and eBay.

The OHA will develop and promote the Android mobile platform, an open source stack with all the free software needed to run a mobile phone: a Linux-based OS; a browser; middleware for delivering applications and services; and programs such as maps, e-mail, contact lists, and video-sharing and calendar-management tools.

The group will publish the Android specifications and make the platform available via an open source license to let developers create a variety of applications and services.

The basic specification calls for Wi-Fi support and multiple input approaches, including numeric keypads, traditional typewriter keyboards, and touch screens.

The OHA has released a toolkit that lets independent programmers, not just carriers, provide software for Android phones.

The OHA expects the first handsets to ship by the second half of this year.

Android—named for a company Google bought in 2005—will be available via the minimally restrictive Apache open source license. It would compete with mobile platforms such as Symbian, Windows Mobile, and those based on Linux.

"Google brings credibility [to Android]," noted analyst David Chamberlain with In-Stat, a market research firm, "which may attract developers and service providers."

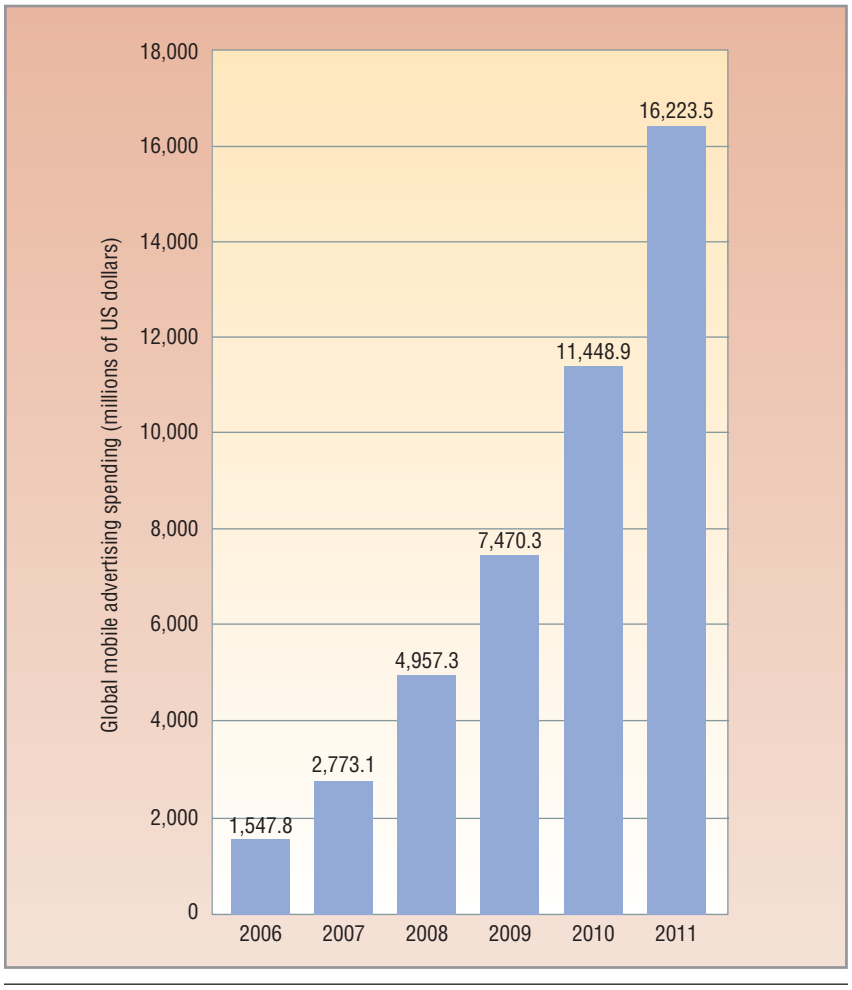

Figure 1. Market research firm eMarketer predicts a steady, substantial worldwide growth in spending on advertising to mobile devices. Experts speculate this opportunity is one reason that Google, which has made money via advertising on its search site, is interested in entering the mobile-phone market.

CURRENT CARRIERS OPEN UP

AT&T Mobility's and T-Mobile's networks use Global System for Mobile Communications technology, which lets subscribers use any GSM-compatible device.

This is not the case with US nationwide carriers Sprint Nextel and Verizon, which use code-division multiple-access cellular technology. Sprint and Verizon are now working on plans to open their networks.

Free Press policy director Ben Scott expressed skepticism about these plans, saying of the carriers, "Six months ago, they told us that open access meant the death of the wireless industry."

In addition to service providers, the Linux Phone Standards (LIPS) Forum—which includes major companies such as British Telecommunications, France Telecom, and Texas Instruments—is trying to foster openness in mobile and fixed telephony by standardizing Linux-based services and middleware APIs, said forum general manager Bill Weinberg.

This would promote Linux use; allow independent developers to write applications for LIPS devices; and enhance the development, deployment, and interoperability of software and services, enabling mobile phones to work with multiple programs and networks.

Verizon

Verizon, once a vocal opponent of open access, has announced it will start letting customers use almost any device or software on its network by the end of this year.

In the near future, Verizon plans to publish technical standards that the development community can use to design products that work with the company's network.

Verizon will test devices for compatibility and limit application use only for security reasons, noted company spokesperson Nancy Stark.

"This is a transformation point in the 20-year history of mass-market wireless devices, one we believe will set the table for the next level of innovation and growth," said

Verizon CEO and president Lowell McAdam.

Sprint's open WiMax network

Sprint has announced plans to spend $5 billion on Xohm, an open network the company hopes will reach 185 million subscribers within three years. Xohm would be based on Mobile WiMax, a long-range, broadband wireless technology.

"We will facilitate the use of application programming interfaces by developers to encourage device and service innovation [and] encourage the availability of a wide array of access devices," said Sprint spokesperson John Polivka.

The APIs would make it easy for developers to create applications and services that work with Sprint's infrastructure.

THE DOWNSIDE OF OPEN ACCESS

In-Stat's Chamberlain predicted that adding new, open approaches to the already multifaceted world of mobile telephony would further fragment the industry and thereby slow the move toward interoperability and openness.

Also, open wireless devices and applications could attract hackers and malware that could interfere with phone service, said Forrester Research analyst Charles Golvin.

"Carriers today largely prevent subscribers from putting third-party software on their devices, which helps prevent malicious activity. The key question is how much downloading freedom can carriers give subscribers without risking widespread network harm," noted the Stanford Group's Gallant.

AN OPEN FUTURE

Ross Rubin, director of industry analysis for the NPD Group, a market-research firm, predicted that openness will usher in a new age of creative networked devices and services.

He said it will open the doors for vendors, independent developers, and others to design and implement services that would, for example, let mobile media players stream music over the Internet, enable cameras and camcorders to automatically post pictures and videos online, and allow GPS devices to gather local traffic and neighborhood information.

The Yankee Group's Marshall predicted the US cellular industry will move from today's walled-garden approach to a more open model during the next five years or so.

If the major players like Verizon don't change, he said, they will lose business because consumers will be attracted to the lower-cost, more-customizable open networks and services.

"This change was inevitable," Forrester's Golvin said. "It is no great insight to draw an analogy to the fixed Web, where initially AOL and CompuServe [and their closed systems] were successful because they provided ways to organize information. But [eventually], there was more innovation toward an open environment, and consumers saw the opportunity to benefit. Closed environments like AOL and CompuServe were left in the dust. The same will eventually happen to mobile operators."

George Lawton is a freelance technology writer based in San Francisco. Contact him at glawton@glawton.com.

Editor: Lee Garber, *Computer*, l.garber@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 WEB SITE: 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 14 magazines, 12 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 over 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 and certification.

Next Board Meeting: 16 May 2008, Las Vegas, NV, USA

IEEE

EXECUTIVE COMMITTEE

President: Rangachar Kasturi*
President-Elect: Susan K. (Kathy) Land*
Past President: Michael R. Williams*
VP, Electronic Products & Services: George V. Cybenko (1st VP)*
Secretary: Michel Israel (2nd VP)*
VP, Chapters Activities: Antonio Doria†
VP, Educational Activities: Stephen B. Seidman†
VP, Publications: Sorel Reisman†
VP, Standards Activities: John W. Walz†
VP, Technical & Conference Activities: Joseph R. Bumblis†
Treasurer: Donald F. Shafer*
2008–2009 IEEE Division V Director: Deborah M. Cooper†
2007–2008 IEEE Division VIII Director: Thomas W. Williams†
2008 IEEE Division VIII Director-Elect: Stephen L. Diamond†
***Computer* Editor in Chief:** Carl K. Chang†
** voting member of the Board of Governors* *† nonvoting member of the Board of Governors*

BOARD OF GOVERNORS

Term Expiring 2008: Richard H. Eckhouse, James D. Isaak, James W. Moore, Gary McGraw, Robert H. Sloan, Makoto Takizawa, Stephanie M. White
Term Expiring 2009: Van L. Eden, Robert Dupuis, Frank E. Ferrante, Roger U. Fujii, Ann Q. Gates, Juan E. Gilbert, Don F. Shafer
Term Expiring 2010: André Ivanov, Phillip A. Laplante, Itaru Mimura, Jon G. Rokne, Christina M. Schober, Ann E.K. Sobel, Jeffrey M. Voas

EXECUTIVE STAFF

Executive Director: Angela R. Burgess
Associate Executive Director: Anne Marie Kelly
Associate Publisher: Dick Price
Director, Administration: Violet S. Doan
Director, Finance & Accounting: John Miller

COMPUTER SOCIETY OFFICES

Washington Office. 1828 L St. N.W., Suite 1202, Washington, D.C. 20036-5104
Phone: +1 202 371 0101 • Fax: +1 202 728 9614
Email: hq.ofc@computer.org

Los Alamitos Office. 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 Office. Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan
Phone: +81 3 3408 3118 • Fax: +81 3 3408 3553
Email: tokyo.ofc@computer.org

IEEE OFFICERS

President: Lewis M. Terman
President-Elect: John R. Vig
Past President: Leah H. Jamieson
Executive Director & COO: Jeffry W. Raynes
Secretary: Barry L. Shoop
Treasurer: David G. Green
VP, Educational Activities: Evangelia Micheli-Tzanakou
VP, Publication Services & Products: John Baillieul
VP, Membership & Geographic Activities: Joseph V. Lillie
VP, Standards Association Board of Governors: George W. Arnold
VP, Technical Activities: J. Roberto B. deMarca
IEEE Division V Director: Deboarh M. Cooper
IEEE Division VIII Director: Thomas W. Williams
President, IEEE-USA: Russell J. Lefevre

revised 15 Jan. 2008

TECHNOLOGY NEWS

A Storm (Worm) Is Brewing

Brad Smith

Last year marked a turning point in malicious software's evolution that has caused serious concern among security experts. Skilled hackers have designed a sophisticated type of malware that blends multiple techniques, hides and changes its code, and employs tricks to entice users to implement and spread it.

The malware is generally called the Storm worm, although it is also known by names such as Nuwar, Peacomm, Tibs, and Zhelatin. Security vendors believe the malware has been around since August 2006. However, it began gaining widespread attention on 19 January 2007, when it launched a series of attacks that generated an estimated 20 times the normal spam volume.

The Storm worm has created a massive network of remotely operated zombie computers that the person or people who control the malware have used to launch spam and distributed denial-of-service (DDoS) attacks.

A key to Storm's success has been its controllers' clever and creative use of social engineering to entice unsuspecting victims to open e-mail attachments or connect to harmful websites.

In addition, the controllers regularly change the malware's code and delivery mechanism.

Bruce Schneier, chief technology officer of security vendor and consultancy BT Counterpane, said the Storm worm is probably the wave of the future and is particularly dangerous because of its complex design and aggregation of existing malware tools.

Despite its name, the malware is not just a worm but also includes Trojans, botnets, rootkits, encryption, and peer-to-peer networking, noted Joe Stewart, researcher for SecureWorks, a managed-security service provider.

"This is top-quality, state-of-the-art malware," said David Perry, director of global education for security vendor Trend Micro.

INSIDE THE STORM WORM

The Storm worm has been definitively linked to several large waves of spam attacks in April, May, July, and August 2007, as Figure 1 shows. It is probably responsible for several others, according to security experts.

The malware's controllers have also used their zombie network to launch DDoS attacks on antispam and antivirus organizations, such as the Spamhaus Project (www.spamhaus.org), that have tried to stop their activities.

Storm targets PCs running Windows 2000, XP, and Server 2003, and security experts expect it also will target Vista.

Most experts think Storm's controllers are based in Russia or Eastern Europe. Much of the traffic from the botnet's controllers appears to come from that area, based on IP addresses traced to server farms there, noted Stewart.

The hosting servers also apparently use text written in Russian for communications, he added.

Show me the money

"The Storm worm is crimeware," noted Symantec director of security response Kevin Haley. "It's about making money in the underground economy."

For example, it has launched spam as part of "pump and dump" schemes, which promote low-value stocks that hackers have purchased. The promotion drives up the stocks' price so that the hackers can sell them for a profit.

Some evidence indicates that Storm's owners have divided the botnet into pieces they can rent to spammers, SecureWorks' Stewart said. In one case, he noted, it appears part of the Storm botnet was used to send out spam for a Canadian pharmaceutical outlet.

Social engineering

For Storm to spread, victims must voluntarily open e-mail attachments or click on links to infected websites. To convince people to open attachments, Storm-related messages use attractive subject lines about a fake or real news event, easy ways to make money, inexpensive products, or communications from a friend or family member.

The subject lines, which sometimes are grammatically incorrect, have read, for example, "230 dead as storm batters Europe," "A killer at 11, he's free at 21 and kill again!," "British Muslims Genocide," "Naked teens attack home director," "Re: Your text," "Russian missile shot down USA satellite," and "US Secretary of State

Published by the IEEE Computer Society

Condoleezza Rice has kicked German Chancellor Angela Merkel."

The person in command of the computer that controls the botnet generates the subject lines and sends them to the bots to use with the e-mail they send.

Storm worm elements

Trend Micro senior researcher Jaime Yaneza said the Storm worm infects computers with multiple payloads that contain several key elements.

The payload includes software that handles the forwarding of spam and the replacement of the core botnet code to disguise the malware's presence.

It also contains a worm that enables Storm to replicate and spread by sending copies of itself over a network and then infiltrating victims' systems by exploiting Web browser and other system vulnerabilities.

Storm also includes a Trojan horse that, once loaded on the infected bot, creates a backdoor that the malware's authors can use to issue attack commands. The Trojan—which appears to be a legitimate file such as a video clip of a news event—also carries a rootkit, which masks the malware. The rootkits add to or replace part of the kernel code, which makes them difficult for security software to detect.

The rootkit also removes evidence of Storm by replacing binaries that list a system's files and processes with a version that doesn't reveal the malicious code. It also intercepts API calls that provide a list of running processes and removes the ones related to Storm. And it deletes entire files or utilities that could indicate the malware's presence.

In addition, the rootkit can modify the code of legitimate existing drivers in the Windows registry so that they launch Storm every time Windows starts.

Storm's owners can divide the core botnet into subnetworks to, for example, rent to different groups of spammers. The owners can configure each of these subnetworks differently.

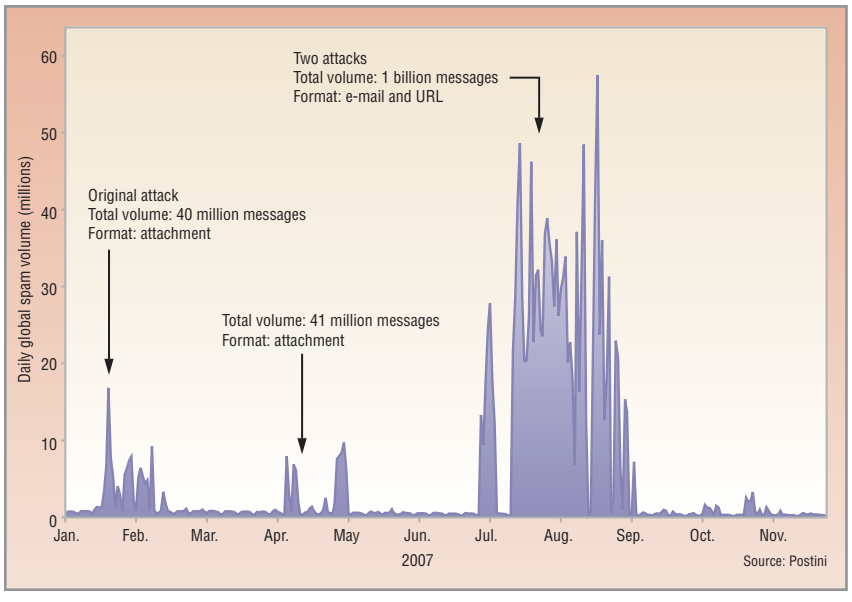

Figure 1. The Storm worm, which its controllers unleashed last year to distribute large amounts of unsolicited e-mail, caused huge spikes in global spam volumes during its major 2007 attacks.

Communications

The malware uses peer-to-peer networking—via the Overnet protocol, designed for decentralized networks such as Storm's botnet—and Internet Relay Chat to feed new code to the zombies.

Storm's human controllers typically use the ICQ instant-messaging platform—a favorite of Russian hackers—to send messages to each other or associates, perhaps to set up a spam attack or discuss new strategies, according to Stewart.

The infection process

Storm can spread via either opened e-mail attachments or visits to infected websites.

Its e-mail messages generally include no text but carry a link to an executable attachment—titled "Read More.exe," "FullClip.exe," "Full Story.exe," "FullVideo.exe," or "Video.exe"—that, if clicked on, downloads software that turns the victim's computer into a zombie.

The messages also sometimes link to infected websites. If unsuspecting users visit an infected site and click on an embedded hyperlink, they can download the Trojan.

Storm can infect webpages with a malicious iFrame (inline frame) that includes a piece of HTML code that, when clicked on, downloads the malware.

Once Storm infects a computer, it can take several additional actions. For example, the malware can upload keylogging software that reads a victim's keystrokes to capture information such as credit card numbers or passwords.

Botnets

According to BT Counterpane's Schneier, estimates of the Storm botnet's size frequently range from 1 million to 50 million computers.

Microsoft says there were probably about 500,000 Storm zombies as of September 2007, based on information from its Malicious Software Removal Tool installed on Windows-based PCs.

Storm worm was so big, it generated 20 million spam messages—20 times the normal volume—during an attack that occurred between 19 and 23 January 2007, noted Adam Swidler, product marketing manager for security vendor Postini.

During a six-week attack the following July and August, Storm worm generated 1 billion spam messages, including about 60 million in one day.

Security experts speculate that Storm worm was responsible for a massive DDoS attack against the Estonian government's cyberinfrastructure in May 2007.

Self-preservation

Storm's authors have changed the malware's delivery mechanism regularly and have used several other sophisticated techniques to make recognition by security products difficult.

Multiple delivery approaches. One of Storm's gambits has been to use spam that includes PDF attachments or electronic greeting cards—such as those distributed this past Christmas and New Year's holidays—with links that actually take users to infected websites. It has also worked via infected audio-file attachments and links in instant messages.

In addition, Storm spam has contained an attractive invitation and a supposed link to a YouTube video. The link actually takes users to a Storm distribution site with the YouTube logo and tells them to click on another link, which uploads malware onto their computer.

Decentralization. Unlike many botnets, Storm uses multiple zombies, rather than a single, central server, for command and control. There is thus no single computer that security experts can target to stop the malware's activities.

Using few bots at a time. Schneier said a small fraction of Storm's zombies spread the malware and an even smaller fraction act as command-and-control servers, while the rest wait for orders.

"By allowing only a small number of hosts to propagate the virus and act as command-and-control servers, Storm is resilient," he explained. Even if security experts shut down the active bots, he noted, the network remains largely intact and other zombies can take over.

Encryption. According to Trend Micro's Yaneza, Storm uses 40-bit encryption to prevent antivirus software from accurately reading its code and identifying it as malware. The system also takes each of its bots out of service for long periods, making them harder to detect.

Code replacement. Commands sent by Storm's controllers replace the malware's core code up to 10 times per hour, making identification by security software difficult, said Yaneza.

Fast fluxing. Storm also makes itself harder to detect via *fast fluxing*, which hides websites that can infect visitors behind an ever-changing network of compromised hosts acting as proxies. The compromised computers' public domain-name records change constantly, in some cases every few minutes. This makes it difficult to track their activities and shut them down.

To combat the Storm worm, security experts recommend basic computer hygiene by individuals and organizations. This includes exercising care when opening e-mail attachments, using intrusion- and rootkit-detection systems, and employing techniques such as the blocking of peer-to-peer communications.

Nonetheless, keeping up with Storm's many changes presents a challenge to security vendors, so the malware promises to continue causing problems.

For the time being, eliminating malware like Storm is going to be difficult because people will continue to open unsafe e-mail attachments, according to computing pioneer and Carnegie Mellon University professor David Farber.

Ultimately, he said, eliminating such malware will require authenticated e-mail because currently, there is no way of knowing whether a message has come from a trusted source or includes a malicious attachment.

In fact, he added, "I think we're going to have to redesign the protocols of the Internet anyway, and when we do that, we will have to pay attention to security."

"We designed the network with very little attention to security," Farber noted. "It wasn't a problem then."

Brad Smith is a freelance technology writer based in Castle Rock, Colorado. Contact him at pbradsmith@gmail.com.

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

NEWS BRIEFS

IBM System Is a Virtual Sign-Language Interpreter

UK-based IBM scientists have developed a prototype system incorporating avatars that translate speech into sign language.

The SiSi (say it, sign it) system uses speech-recognition technology to convert a conversation into text and then translation technology to generate commands that animate an avatar so that it makes the correct sign-language gestures.

The system—developed with assistance from sign-language users and the UK's Royal National Institute for Deaf People—is not designed to replace human interpreters. Instead, it could be used when no human interpreter is available or when confidentiality is important.

Possible applications include the teaching of sign language, as well as services that translate TV shows, voice mail, or public announcements into sign language, said Helen Bowyer, an IBM emerging-technology software engineer.

The SiSi prototype uses IBM's ViaVoice speech-to-text technology, noted Bowyer. The text input is then passed through a translation software module, which conducts syntactic parsing, lexical analysis, and other processes to convert the content into grammatically correct British Sign Language.

SiSi tags sentences with grammar markers and restructures them into BSL grammar. The system then sends commands—written in the Sign Gesture Markup Language, a format the UK's University of East Anglia developed for working with avatars—to the user's PC.

To animate a customizable, IBM-developed avatar, the system converts the commands into BSL gestures retrieved from a dictionary of signs and into the facial expressions that would be used to speak the translated words, Bowyer explained.

"The avatar has a built-in mechanism for smoothly transitioning from the end of one gesture to the start of the next one, to give fluidity to the movement on screen," she noted.

SiSi can't translate sign language into speech, which is a different problem than converting speech into gestures. Computer audio processing is difficult, but visual processing is harder and more expensive, requiring techniques such as motion capture and model-based analysis, she explained.

"IBM is not currently working in this area, although we recognize that it would be another important step in aiding two-way communication," she said.

SiSi works with Sign Supported English as well as BSL. IBM designed the system to be adaptable—via multiple translation modules that would include the appropriate signs, grammar, and syntax—for use with other languages, Bowyer noted.

Future improvements could enable SiSi to work with complex sentence structures or use an avatar that functions on a greater variety of devices such as digital TVs, said Bowyer. The researchers have already developed a prototype for viewing the avatars on cellular phones via streaming video.

The SiSi team will try to improve the accuracy of their system by gaining a deeper understanding of BSL syntax and grammar and improving the translation module to address more nuances, Bowyer noted.

IBM sees SiSi functioning commercially as either a stand-alone application or technology included in other products.

Currently, though, Bowyer added, "There are no plans to license or sell SiSi as it is still very much a prototype requiring further development." ■

News Briefs written by ***Linda Dailey Paulson****, a freelance technology writer based in Ventura, California. Contact her at ldpaulson@yahoo.com.*

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

IBM researchers have developed a prototype system, called SiSi (say it, sign it), that uses avatars to translate speech into sign language. (Source: IBM)

NEWS BRIEFS

Consumer Groups Advocate Do Not Track Program

US consumer and privacy groups have called for the creation of a Do Not Track program—similar to the Do Not Call program for telemarketers—that would let users avoid having companies track their Web surfing via cookies and other means.

Nine groups—the Center for Democracy and Technology, Consumer Action, the Consumer Federation of America, the Electronic Frontier Foundation (EFF), Privacy Activism, Public Information Research, the Privacy Journal, the Privacy Rights Clearinghouse, and the World Privacy Forum—have submitted this proposal to the US Federal Trade Commission.

Proponents say the plan is designed to protect people's privacy by allowing them to avoid having advertisers track their online activities without their permission.

"The main goal of the proposal is for advertisers to disclose whether they are engaged in behavioral tracking of any kind, not simply whether they are using cookies," said EFF senior staff attorney Lee Tein. "This is an attempt at an approach that is not intrusive, not heavily regulatory."

The FTC will accept public comments on the Do Not Track proposal until 22 February and then decide what to do next about the plan, noted Jessica Rich, assistant director of the commission's Division of Privacy and Identity Protection.

Advertisers frequently place cookies on PCs to identify users on future visits and track the websites they view. They often do this so that they can send ads related to visitors' interests as reflected by their online activities.

In response to this, Do Not Track proponents say they want to address cookies and other types of persistent identifiers.

"The FTC or someone else would maintain a list of who is engaged in behavioral tracking," Tein said. "Consumers would know by looking at that list who is and who is not engaged in these practices. It would allow you or me to see that a company is doing this, so we wouldn't visit their site [if we didn't want to be tracked]."

There might even be technology that would automate the avoidance of these sites, he added. Currently, the Do Not Track proponents don't specify the technology that would make their plan work.

"I suspect they understand that it is technologically difficult, if not impossible, to implement, and they just want the public exposure versus a real solution to benefit consumers," said Mike Zaneis, vice president of public policy for the Interactive Advertising Bureau (IAB), which represents about 300 online advertisers.

"The challenge of creating a universal software platform that would serve as judge and jury for which websites Americans are allowed to view is very daunting," he added.

Tein said the Do Not Track proposal was deliberately unspecific in many ways because this would let the concerned parties develop the best technical approaches for doing the job.

Online advertisers oppose the Do Not Track proposal. The IAB favors self-regulation and is reviewing best practices to disseminate to members.

"I couldn't begin to divine how they see such a system working effectively," stated Zaneis. "Cookies make e-commerce sites work, make customized websites possible, and often are used by publishers to deliver more relevant advertising."

Group Chooses to Run Fast Bluetooth over Wi-Fi and UWB

The Bluetooth Special Interest Group has changed its mind and has decided to base a fast version of its short-range wireless-connectivity technology—which would enable the streaming or transferring of video, audio, or other content to and from mobile devices—on both Wi-Fi and ultrawideband technology.

Bluetooth currently uses a radio approach designed specifically for the technology. However, proponents want to increase Bluetooth's speed by having it run over a faster Wi-Fi or UWB radio.

During their initial considerations in 2006, the SIG board of directors selected only UWB for fast Bluetooth because, at the time, it consumed less power and had a higher theoretical-maximum throughput than Wi-Fi—480 megabits per second compared to 54 Mbps, said Charles Golvin, an analyst with Forrester Research.

However, Golvin noted that ultrawideband's performance in practice has not met the Bluetooth SIG's expectations, perhaps because many UWB radios aren't using a sufficiently wide spectrum band for transmission. The technology is supposed to transmit signals over wide

frequency bands, thereby enabling high data rates.

Also, Golvin said, UWB radios cost more than the Bluetooth SIG expected, perhaps because the technology hasn't been implemented as widely as Wi-Fi and thus hasn't achieved the same economies of scale.

Now, the SIG is working on a protocol adaptation layer that will let a Bluetooth radio run over either Wi-Fi or ultrawideband, said Stephen Wood, Intel technology strategist and president of the WiMedia Alliance, a UWB industry organization. This would eliminate the need for the generally small Bluetooth devices to have both a UWB and a Wi-Fi radio, he explained.

Currently, Bluetooth typically offers transmission ranges up to 10 meters and speeds up to 3 Mbps.

UWB's next generation promises speeds of 960 Mbps and a transmission range of up to 10 meters.

Bluetooth SIG executive director Mike Foley said his organization is exploring the use of the IEEE 802.11n version of Wi-Fi, which offers a theoretical maximum data rate of 540 Mbps and a transmission range of up to 70 meters. However, this would force vendors to take into account Wi-Fi's higher energy usage during the design process.

Once the SIG develops the new fast Bluetooth specifications, Foley said, it will build prototypes, probably late this year, and then test them. The SIG expects its board to adopt the specification in the first half of 2009.

Because Wi-Fi has been so widely adopted, it might become Bluetooth's mainstream radio technology, while UWB might be used for niche applications such as linking devices like digital cameras with a PC, according to Golvin. ■

Researchers Design a Breath-Operated Computer Interface

Two Georgia Tech University researchers have developed a way for users to blow at their PC or laptop screens to control interactive applications.

The Blowable and Localized User Interaction interface that doctoral candidate Shwetak Patel and professor Gregory Abowd developed could help people work with computers when they can't use their hands because they are either busy with other tasks or have a disability or injury.

BLUI lets a user blow at the part of a screen on which an application control—such as a button, cursor, or even scroll bar—is located.

Users first install the BLUI engine on their computer, explained Patel. Computers currently need specially designed interfaces to work with BLUI. However, Patel added, designing the technology to work with the Windows interface would not be difficult.

BLUI works with the computer's microphone to determine where on the screen a person is blowing, based on the sound of the air as it reflects off the screen. With a PC, users would have to place the microphone near the screen.

To calibrate BLUI, the system displays boxes on different parts of the screen where a user should blow. During the calibration, the system calculates and then stores audio fingerprints for the sound of the blowing on each part of the screen, Patel noted.

He said researchers haven't yet determined the accuracy of BLUI's current version. Even with calibration, BLUI isn't as accurate as a mouse because a breath can cover a relatively large screen area, a single microphone doesn't always precisely capture data, and different microphones can produce different results. Software filtering and other techniques help improve accuracy.

BLUI offers a couple of advantages over speech interfaces, Patel said. BLUI users don't have to make audible noises when issuing commands and developers don't have to design multiple, intuitive verbal commands for different functions, he explained.

Patel said that he is releasing an open source toolkit to help users easily and effectively work with BLUI and that he will explore the technology's commercial use as a game controller.

Georgia Tech has obtained a provisional US patent for the technology and hopes to receive a full patent later this year.

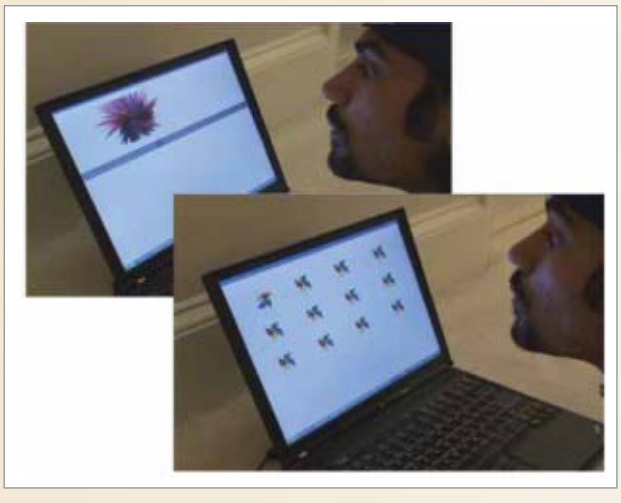

Two Georgia Tech University scientists have developed the Blowable and Localized User Interaction interface, which lets users blow at their PC or laptop screens to control interactive applications. BLUI could help people use computers when they can't work with their hands because they are busy with other tasks or because of disability or injury. (Source: Georgia Tech University)

COVER FEATURE

National Science Digital Library: Shaping Education's Cyberinfrastructure

David McArthur
GoH Corp.

As a National Science Foundation program, the NSDL is reaching maturity, but the library is already forging a strong link among research projects, which argues compellingly for continued NSF investment, although with new directions.

In summer 2007, the National Science Digital Library provided access to more than 2.5 million digital educational resources, covering science, technology, engineering, and mathematics from pre-K to postgraduate levels. At that time, the main portal, www.nsdl.org, had been online nearly five years, providing a single entry point to dozens of distributed collections and services, many of which were contributions of the more than 200 small projects that the NSDL program funded through its research tracks.

With this maturity, the National Science Foundation is now rethinking NSDL's status as a research program. In one sense, it remains a typical NSF program, operating through the traditional NSF project-award cycle of publishing a solicitation, receiving proposals from R&D teams, and awarding the best among them. In other ways, it is an atypical program because the goal of its projects is not simply to broaden the knowledge base of science education research and practice; it is also to build an integrated enterprise that will persist and be valuable to learners and teachers of all ages.

But whether typical or not, NSDL has reached the point at which it must either change substantially or start winding down. Many NSF programs come and go in less than a decade, often after accomplishing their primary goals and laying a foundation for a new research agenda. As a library, NSDL is becoming mature enough to be an operational center. Because NSF is primarily a research agency, investing further in NSDL would seem to run counter to NSF's policy of not supporting routine science and education operations.

Nonetheless, there are compelling arguments for NSF's continued investment in NSDL—but the nature of that support must change to match NSDL's new purposes. Generally speaking, NSF's policy is to "let a thousand flowers bloom," and to that end, it spawns programs that award distinct projects and that rely on conferences and publications (both traditional and electronic) to foster researcher crosstalk. This is an admirable goal, but I believe that NSDL gives NSF an opportunity to tighten the link among R&D projects: The library is poised to provide a standardized technical infrastructure that encourages—perhaps even requires—a much higher degree of project interaction.

In that mission, I see NSDL growing both as a platform for improving the productivity of educational resource development and transforming education research and also as a tool for creating and managing scientific knowledge about education and learning. More broadly, NSDL could be a key component in building a new cyberinfrastructure for education and education research.

NSF's continued investment in NSDL would have strong implications for how it funds education R&D and how it manages projects to foster effective partnerships among highly diverse and distributed groups of education researchers, developers, and practitioners. Having recently completed a rotator position at NSF, I have been able to acquire an understanding of what NSDL as a library is accomplishing and how NSDL as a program is run. The ideas set forth in this article draw from that dual understanding, but admittedly much of the thinking is speculative and reflects my own views, not necessarily those of the NSF.

Published by the IEEE Computer Society

NEW PURPOSES

NSDL can enhance its value in at least four ways: by serving as a repository and providing knowledge management for digital products from NSF's educational programs, by growing as a platform for improving the productivity of educational materials development and implementation, by continuing to be a testbed for interesting education research and information science questions, and by strengthening the foundation for a cyberinfrastructure that connects education research and practice at NSF.

Digital product repository

In its first years, NSDL had a collections track that supported researchers in their creation of new digital resources, which NSDL librarians then catalogued in collections for users. The program has recently dropped this track, however, in part because other NSF programs are now generating a sufficient supply of digital educational materials. Within the NSF's Division of Undergraduate Education, for example, the Course Curriculum and Laboratory Improvement (CCLI) program typically funds dozens of projects annually, each of which produces digital exercises, modules, simulations, and games for classrooms and laboratories. Given that hardware and development application costs are likely to keep dropping, such educational materials will continue to proliferate, with or without NSDL.

On the other hand, these resources often have limited use beyond the projects that created them. Even when developers would be happy to share materials, faculty who could put them to good use often don't know they exist, let alone how to find them. In this, NSDL can continue to make a difference. Among the Pathways projects, one of the four project tracks listed in the "NSDL Project Tracks" sidebar, ComPadre is assuming a stewardship role for the educational resources useful to broad communities in physics and astronomy. One of ComPadre's tasks is to review the digital products of recent CCLI projects in physics and astronomy and to add these materials to ComPadre's collections, thereby making them accessible to a broader community through NSDL.

Collections from NSF programs. The Internet Scout team leads another Pathways project, the Applied Mathematics and Science Education Repository (AMSER), which is providing a similar service by revamping NSF's Project Information Resource System. PIRS is an online database that records not only CCLI projects' digital products, but also information about the grant, project investigators, and the institution awarded. From the start, PIRS has been open to the public, but its use has been spotty. The Internet Scout team believes that by using NSDL as a foundation, they can implement a more powerful PIRS, providing a wealth of new search and browse facilities—all at a fraction of the original PIRS cost. The hope is that PIRS will become the CCLI wing of NSDL, not just including CCLI project materials, but also providing a community hub for education researchers and practitioners.

Such directions are only the beginning. CCLI is not the sole source of educational materials, even in the Division of Undergraduate Education. And projects in other divisions within the Education and Human Resources Directorate are also busy generating digital resources for K-12 learning and teaching and for graduate education. In addition, researchers supported through other NSF directorates—biology, computer and information science, engineering, geosciences, and mathematical and physical sciences, for example—are developing digital content that professionals in these fields can use in both research and teaching. If NSDL's AMSER and ComPadre projects are successful in finding and organizing CCLI resources, other projects could expand this work to different programs and disciplines, creating still more new wings of NSF educational content within NSDL.

Knowledge management and sharing. Such NSDL projects could help NSF manage its internal knowledge base. NSF funds many new projects each year, and it is difficult for NSF program directors—particularly rotators, who stay for only a year or two before returning to their home institution—to have information about relevant projects at their fingertips. Program directors hear hundreds of ideas for proposals each year.

To illustrate the type of aid that NSDL could provide, suppose a program director is listening to a prospective principal investigator describe a project over the phone. While the investigator is talking, the director should be able to search the new NSF wings in NSDL and instantly

NSDL Project Tracks

At present, NSDL has four main project tracks:

- *Core Integration,* a single large project, coordinates and manages the core library, develops the library's central portal and infrastructure, and engages and supports the other NSDL projects and community.
- *Services* projects develop tools that support users, collection providers, and the Core Integration effort. They enhance NSDL's efficiency and value.
- *Targeted Research* projects explore specific topics that have immediate applicability to collections, services, and other aspects of NSDL's development.
- *Pathways* projects provide stewardship for the collections and services that major learner communities require.

Information on specific projects in these tracks is available on the NSDL website (http://nsdl.org/about/index.php?pager=factsheet).

Inside NSDL

Users enter the National Science Digital Library through its main portal, where they can access hundreds of distributed collections in the NSDL Data Repository (NDR) using a wide range of services. Users can search by desired characteristics such as subject, grade level, and media format; services let users annotate resources and compose them into larger units. Recently added services include Expert Voices—a blog where subject experts communicate with students and teachers—and On Ramp—a system for distributed creation, editing, and dissemination of content from multiple groups.

Technical standards have been vital to the distributed development of both content and services. Using NSDL metadata standards, projects easily share resource descriptions through the NSDL data repository. This access, in turn, invites others to reuse, dissect, and recombine the resources in different ways and to return the results to NSDL.

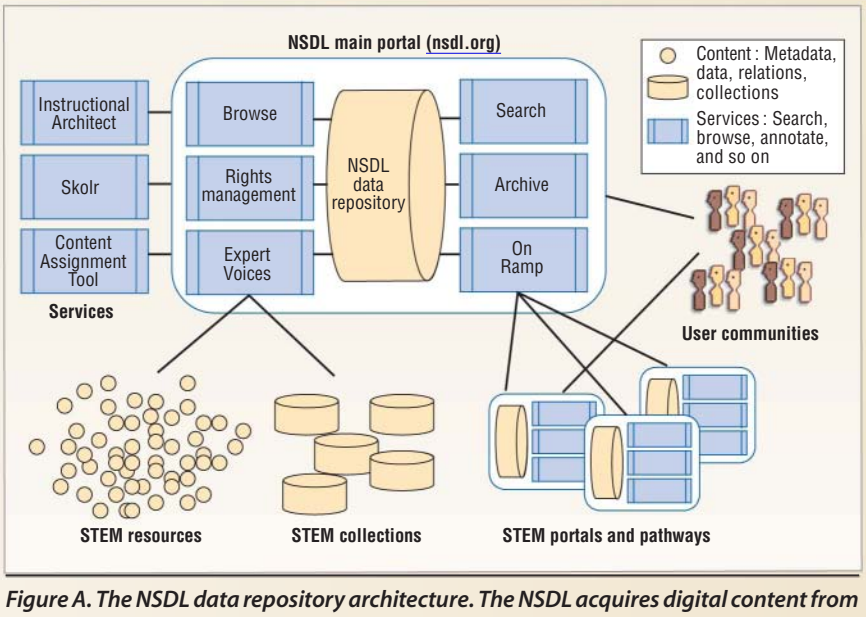

Figure A. The NSDL data repository architecture. The NSDL acquires digital content from science, technology, engineering, and mathematics (STEM) collections, including NSDL Pathways projects and external webportals. Also provided are a growing number of services, from search and annotation to tagging with learning standards.

With NSDL's infrastructure now reimplemented using Fedora, an architecture for managing digital assets, projects can share in the development of interoperable services as well, which should greatly enhance NSDL's functional capabilities. Innovative services from projects (far left in Figure A) include Instructional Architect, a tool for authoring course modules using NSDL resources; Skolr, a personal collection service; and Content Assignment Tool, an application to align resources with educational standards.

find all the projects that resemble the one being proposed. NSF staff would also be able to summarize collections of recently supported projects, provide broad visualizations of funded areas, or even detect gaps in the topics that a program's project portfolio covers.

Giving prospective investigators access to the same kinds of tools would let them discover if other projects have already investigated the exciting idea they want to propose. If their idea turns out to be not so novel, they could gain enough information to decide how to build on the idea rather than reinvent it. Avoiding reinvention was, in fact, a motivation for the NSDL program, and it is a persuasive reason for NSF to continue investing in NSDL.

Materials development platform

Although NSDL can provide repository services for existing resources, it also can be a useful platform for developing new educational materials. Several NSDL projects are already collaborating to build resources and collections within NSDL, often by taking pieces from other projects and composing them in innovative ways. The "Inside NSDL" sidebar describes some of NSDL's development services.

ContentClips, a project in the Services track, is developing a system that lets teachers assemble standards-driven learning activities using templates that organize digital objects. The objects can come from any collection, but the project is relying heavily on content from existing NSDL collections to populate initial templates.

Even more ambitious is Instructional Architect, which lets teachers find shareable resources from the NSDL and other websites, and provides an authoring tool that assembles webpages and learning modules from the retrieved materials.

Combining and reusing content. The large Core Integration project is planning new services and tools that promise to dramatically improve NSDL's ability to combine materials in this way. Expert Voices invites professionals in exciting research areas, such as global warming or seismology, to engage in dialogue with students and teachers. These digitized raw discussion materials

or blog entries then become part of an NSDL collection and can be the basis for, say, building standards-aligned modules. The modules, in turn, could become NSDL resources linked to annotations or guides on how to use the modules in classrooms or laboratories. By providing tools to create new resources in standardized formats from existing ones and by enabling the composed products to rejoin the library, NSDL affords an array of opportunities for reusing library content.

Although all these tools are possible without NSDL, the library provides a platform that can make materials development much more productive. Computer-Assisted Content Standard Assignment and Alignment (CASAA; http://cnlp.org/research/project.asp?recid=48) is an NSDL project that underscores these potential productivity gains. NSDL now provides access to many high-quality educational resources, but teachers who use the library are often frustrated when their searches return hundreds of results. They usually want a few modules that meet the specific state or federal standards they are teaching to. Unfortunately, most of the materials submitted to NSDL do not come with such standards attached. CASAA aims to solve this problem by developing the Content Assignment Tool (CAT), which will at least partially automate the aligning of a resource with specific standards, as well as the process of putting the aligned resource back into NSDL. This is much more efficient than letting each teacher struggle to relate raw NSDL resources to standards.

The NSDL provides a platform that can make materials development much more productive.

The real productivity boost, however, comes next. Once a resource is aligned with a specific state standard, CAT can then immediately align it with different standards in other states through an automated process that correlates standards of one state to those of others. In addition, if state standards change—which they do all the time—CAT can take care of realigning NSDL resources to reflect the change, potentially saving thousands of hours of manual recataloguing.

Platform requirements. To ensure that NSDL is a productive platform for educational materials development, resources must adhere to established standards for describing educational content. With such standards, it is easier to combine diverse resources and create plug-and-play tools that recombine materials in innovative and valuable ways for educators and teachers. A less obvious requirement for a productive platform is the need for diverse developer communities and an appropriate division of labor among them.

A few like-minded software developers can find materials in NSDL and assemble a new resource from them. Such resources are of some value, but what teachers want most are modules that include not only the resource but also information on the educational standards that the resource satisfies and how others have used the resource in a classroom or laboratory. This knowledge often comes from different communities—software developers, academic researchers, educational theorists, content authors, standards experts, cataloguers, and teachers. In the past, these groups have rarely talked to one another, let alone worked together to build materials.

One of NSDL's great promises as a platform could very well be its ability to bring representatives of these distributed communities together, enabling the creation of much richer educational content than would be possible if each community worked alone.

Research testbed

In addition to its potential as a platform for productive materials development, NSDL can be a testbed for addressing education research questions about both library resources and their use.

Several early studies in the NSDL program looked specifically at library content and quality. One study[1] used a machine-learning approach to determine which of 16 indicators predicted quality ratings of digital resources and found that metadata currency—how recently authors had created or updated resource metadata—was the best indicator. In another study,[2] researchers applied a graphical analysis tool to evaluate the quality of NSDL metadata, much of which is generated by inexperienced cataloguers. The data inaccuracies and gaps uncovered had patterns that might be exploited to speed up metadata correction. The results of the study also could lead to a means for automating metadata repair or creation.

More recent research studies have addressed the educational use of library resources, including how teachers use the Instructional Architect to create lessons using digital library resources[3] and how middle-school students can use the Digital IdeaKeeper to make sense of the information they find, rather than just passively accepting it.[4]

For NSF teaching and learning. These studies are just the tip of the potential for using NSDL as a foundation for research. Many other questions are ripe for study, for example:

- How well do teachers use resources linked to standards?
- What differences does this make to student learning outcomes?
- Is an Amazon-style recommender system the best way to provide a context for using resources in an educational setting?
- Is it possible to create professional development communities in NSDL for teachers, and if so, what roles would such communities play?

Some NSDL projects are already investigating these questions. However, many focus not on digital library use, but on the nature of learning and the design of technology-augmented educational environments. Projects that investigate these broader questions might be funded not by NSDL, but through social, behavioral, and economic sciences programs at NSF, which would use NSDL as a testbed. Of course, NSDL could not be a testbed for *all* educational technology R&D, but it would provide the ability to implement a wide range of education research prototypes quickly and productively.

For social science projects. As NSDL grows from thousands to millions of resources, with archives of older materials and perhaps records of user interactions, it could support a range of information and social sciences research projects. Questions could move beyond classroom use to more sweeping issues, such as how a body of knowledge on new scientific topics grows, as measured by the time-indexed evolution of digital resource sets, or how communities of practice form, as reflected in resource use patterns.

The NSDL could support a range of information and social sciences research projects.

Yesternet, a project funded by the Next Generation Cyberinfrastructure Tools program and staffed with social and physical scientists from Cornell University, including one of NSDL's founders, is taking the first steps in this direction.[5] These researchers aim to use the 40-plus billion pages of the Internet Archive, which has captured snapshots of the Web for more than a decade, as a laboratory for social-science research. On top of this data set, the project will build intelligent front ends that make the Internet Archive data broadly accessible and let social scientists develop, test, and refine research tools.

The project will also investigate how scientists use these tools to examine social phenomena pertaining to how beliefs in online communities and organizations evolve. One hypothesis, for example, is that popular beliefs often acquire a bimodal distribution and tend to foster polarized camps of like-minded members. Unlike surveys and other conventional methods, new tools could track beliefs across time in the Internet Archive testbed, possibly revealing much richer insights about opinion dynamics.

Content-wise, NSDL is not as large as the Internet Archive; nor can it compare in sheer size with the vast collections from experiments in particle physics, global climate system models, satellite-sensed Earth imagery, or protein data banks. However, NSDL is one of the few social-science collections now represented in the Teragrid, NSF's project to build and deploy the world's largest distributed infrastructure for open scientific research data. In this context, NSDL affords interesting research opportunities not because it is a digital library, but because it offers social scientists a way to design and evaluate the tools they will need as their disciplines join with the physical sciences in using the evolving cyberinfrastructure.

Part of cyberinfrastructure

The prospect that NSDL might play several roles in the emerging cyberinfrastructure is yet another reason for NSF to support NSDL in the coming years. The seminal Atkins report[6] outlined many ways in which cyberinfrastructure in physical sciences can extend to environmental engineering, medicine, and atmospheric sciences. High-energy physicists have been among the earliest cyberinfrastructure users, in part because their experiments already generate petabytes of data that the network must store, analyze, and distribute to distant collaborators. But for the most part, early cyberinfrastructure reports did not consider education or the social sciences.

I believe that NSDL can play a significant role in the emerging cyberinfrastructure. The potential impact of cyberinfrastructure in education is at least as significant as that in other disciplines. New studies[7,8] suggest that future education data sets could easily be as rich and challenging to analyze and understand as ones from the physical sciences. An NSF-sponsored report on Cyberinfrastructure for Education and Learning for the Future (CELF)[9] envisions "lifelong learning chronicles" that would capture the sum of the learner's formal and informal education experiences. If cyberclassrooms log not just test results but also rich audio-video and neuroimaging records of student behavior, as well as the multimedia products that learners create, that data might easily be a petabyte per day—almost as much as the output of a high-energy physics experiment.

Beyond generating massive data sets, cyberinfrastructure applications might transform education every bit as dramatically as they are now changing practice in the sciences. Recent reports[10,11] envision that future teachers will tailor learning experiences for whole classes, just as tutors now deftly remediate misconceptions for individual learners. By accessing and analyzing authentic scientific data from remote labs, students might be able to learn concepts they could not grasp before and could hone their knowledge anytime, anywhere. Just-in-time evaluation tools could assess students' understanding and redirect lesson plans at a moment's notice.

In education reform, cyberinfrastructure's promise is twofold. As the previous examples show, it can greatly expand learning and teaching opportunities for students and teachers; yet, earlier generations of educational technologies could also say this. What earlier generations did not provide is the technical foundations for connect-

Connecting Research and Practice: A Vision

One novel approach for linking education research and practice comes from the Strategic Education Research Partnership,[1] which advocates partnerships between researchers, practitioners, developers, and policymakers. SERP's fundamental idea is to tightly couple education research with practice by conducting much research and intervention design in classrooms and other educational venues, rather than in laboratory settings. Instead of establishing a few prototype schools to try out their methods, SERP proposes recruiting hundreds of classrooms across the US to adopt interventions. The idea is to pilot-test interventions in a few classes, immediately suggest and implement improved designs, and then try the improved intervention in other sites. If initial test results look promising, experiments could quickly lead to changes in classrooms across the country. Cyberinfrastructure for Education and Learning for the Future (CELF) envisioned similar ways to connect research and practice, which they captured in the phrase, "Make America one big research school."[2]

Implementing this vision has been all but impossible in the past for both technical and organizational reasons. Now, however, highly distributed systems like NSDL provide a real chance to put it into operation. NSDL already offers access to a wealth of educational materials that research and practice teams can use to fashion experimental interventions for the big—and highly distributed—research school. Because all these resources are available at anyplace, anytime, changes to intervention designs based on results from initial trials would be immediately available to all other sites in the distributed network. In addition, the Core Integration project could integrate CELF's lifelong learning chronicles into this emerging NSDL network, since the creation of many digital tracks in such chronicles will be automatic as students and teachers interact online with materials available through NSDL.

NSDL's open architecture would also enable separate research and practice teams to conduct experiments in parallel and then jointly analyze their results. Indeed, some assessment teams might specialize in the secondary analysis of experimental data that other teams create—assuming, of course, that these results, like digital resources in NSDL, are shared, not proprietary.

Privacy concerns notwithstanding, if it were possible to extend NSDL's content and services along these lines, several other communities could also participate in educational partnerships. Curriculum and software developers might use the lifelong-learning chronicles to shape the designs of interventions and programs. Parents might monitor their children's learning progress. The students themselves could put pieces of their chronicles into performance portfolios. And policymakers and assessment experts should be able to use the data aggregated from a wide range of intervention studies to judge the effectiveness of educational programs and inform revisions.

References

1. M.S. Donovan, A.K. Wigdor, and C.E. Snow, eds., National Research Council, *Strategic Education Research Partnership*, Nat'l Academies Press, 2003; http://books.nap.edu/catalog/10670.html.
2. A. Bernat, J. Smith, and D. Rothschild, eds., *Cyberinfrastructure for Learning and Education for the Future: A Vision and Research Agenda*, Computing Research Assoc., 2005; www.cra.org/reports/cyberinfrastructure.pdf.

ing education research and practice communities. The sidebar "Connecting Research and Practice: A Vision" describes how an infrastructure based on NSDL can assist in forging these connections.

ORGANIZATIONAL CHALLENGES

I have advanced many reasons for continued investment in NSDL—from the immediate need to manage NSF's project knowledge to the future-oriented shaping of cyberinfrastructure for education. A common thread in these reasons is the importance of developing a social and technical platform that will support the shared use of resources and services among diverse communities.

Implementing the kind of community network that could transform the country's educational systems into one big research school is obviously a long-term goal with formidable obstacles that the Strategic Education Research Partnership, CELF, NSDL, and others committed to this vision must overcome. From a technical perspective, no cyberinfrastructure can yet support all the functions and services that education research and practice partnerships might require.

If NSDL ends, lessons learned might be incorporated into a successor program that does support them. On the other hand, NSDL could be part of an installed base that grows incrementally to eventually form cyberinfrastructure. The focus would have to be on building integrated layers of tools and services, rather than building disconnected pieces. This might be difficult for NSF, since it is accustomed to supporting science projects and experiments, not infrastructure.

But the toughest challenges would almost certainly be organizational not technical. As a CELF report notes,[9] this tighter link of research, educational design,

classroom practice, and assessment will have implications for how NSF supports learning and conducts education R&D. Currently, most projects that NSF programs support, including NSDL, are loosely coupled. Consequently, knowledge and results are often not as cumulative as they would need to be to realize cyberlearning's full potential. When projects do build on one another, it is often within single educational communities: Researchers sometimes talk with other researchers, but there is much less direct communication with practitioners and other key stakeholder groups, including assessment specialists and policymakers. To achieve the tighter coupling of research and practice, this would have to change.

Whether or not NSF continues to invest in NSDL, NSF educational programs will need to consider new ways of doing business to address these organizational issues. They might insist that projects form tight partnerships with different educational communities. Education programs could require projects to build on the ever-expanding cyberlearning infrastructure of shared digital resources and knowledge about their use. In a more sweeping change to encourage large-scale coordination, some programs might consider funding projects through cooperative agreements or contracts—now commonly used only with NSF's research facilities and major centers—rather than exclusively through grants. NSDL has already experimented with this by requiring new Pathways projects to agree to a memorandum of understanding that outlines the details of cooperation with the Core Integration project in several technical, marketing, and evaluation areas.

The potentially significant role of NSDL in developing cyberinfrastructure for education is one of the strongest arguments for continued NSF investment. Without having to make dramatic changes to their basic work practices, many existing communities in the physical sciences are already using cyberinfrastructure to address important problems that used to be far beyond their reach. Education, however, might never meet its Grand Challenges if the disparate groups that have a stake in teaching, learning, and educational reform continue to work alone. To use cyberinfrastructure to best advantage, education must not merely make existing practices faster and better; it must establish new and different practices. With the right stewardship, NSDL can be part of those differences and usher in a new era for education.

Acknowledgment

I thank Lee Zia, lead program director for NSDL, for the many conversations that benefitted this article.

References

1. M. Custard and T. Sumner, "Using Machine Learning to Support Quality Judgments," *D-Lib Magazine*, Oct. 2005; www.dlib.org/dlib/october05/custard/10custard.html.
2. N. Dushay and D. Hillman, "Analyzing Metadata for Effective Use and Re-Use," *Proc. DC-2003 Dublin Core Conf.: Supporting Communities of Discourse and Practice*, 2003; www.siderean.com/dc2003/501_Paper24.pdf.
3. M. Recker et al., "You Can Lead a Horse to Water: Teacher Development and Use of Digital Library Resources," *Proc. 5th ACM/IEEE-CS Joint Conf. Digital Libraries*, 2005; http://portal.acm.org/citation.cfm?id=1065385.1065387.
4. M. Zhang and C. Quintana, "Facilitating Middle School Students' Sense-Making Process in Digital Libraries," *Proc. 5th ACM/IEEE-CS Joint Conf. Digital Libraries*, 2005; http://doi.acm.org/10.1145/1065385.1065389.
5. W. Arms et al., "A Research Library Based on the Historical Collections of the Internet Archive," *D-Lib Magazine*, Feb. 2006; www.dlib.org/dlib/february06/arms/02arms.html.
6. D. Atkins, "Revolutionizing Science and Engineering through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure," 2003; www.nsf.gov/od/oci/reports/atkins.pdf.
7. T. Hey and A. Trefethen, "Cyberinfrastructure for e-Science," *Science*, 6 May 2005, pp. 817-821.
8. D. Atkins, "Cyberinfrastructure and the Next Wave of Collaboration," *Proc. Educause Australasia*, 2005; www.educause2005.auckland.ac.nz/interactive/presentations/Atkins.pdf.
9. A. Bernat, J. Smith, and D. Rothschild, eds., *Cyberinfrastructure for Learning and Education for the Future: A Vision and Research Agenda*, Computing Research Assoc., 2005; www.cra.org/reports/cyberinfrastructure.pdf.
10. F. Berman and H. Brady, "Final Report: NSF SBE-CISE Workshop on Cyberinfrastructure and the Social Sciences," 2005; vis.sdsc.edu/sbe/reports/SBE-CISE-FINAL.pdf.
11. American Council of Learned Societies, "Our Cultural Commonwealth: The Report of the ACLS Commission on Cyberinfrastructure for the Humanities and Social Sciences, Final Draft Report," July 2006; www.acls.org/cyberinfrastructure/acls.ci.report.pdf.

David McArthur *is a senior researcher at GoH and a consultant to the NSDL program. Until September 2006, he was a program director in NSF's Division of Undergraduate Education, where the NSDL program resides. His research interests include digital libraries, cyberinfrastructure and learning, and educational technologies. McArthur received a PhD in psychology from the University of Michigan. Contact him at dmcarthur1@nc.rr.com.*

A Symposium on High-Performance Chips

HOT CHIPS 20

Stanford University, Palo Alto, California
August 3-5, 2008

Call for Contributions

AUTHOR'S SCHEDULE

Submission due: March 24, 2008 Acceptance Notification: April 28, 2008
Final Version due: July 2, 2008

AREAS OF INTEREST:

- General Purpose Processor Chips
  - Low-Power
  - High-Performance
  - Multi-Core, Multi-Processor Technologies: interconnect, programming models, compiler, runtime systems
- Other Chips
  - Novel Technology: Quantum Computing, Nano-Structures, Micro-Arrays
  - Low-power chips/Dynamic Power Management
  - Communication/Networking
  - Chipsets
  - Wireless LAN/Wireless WAN
  - FPGAs and FPGA Based Systems
  - Display Technology
- Application-Specific / Embedded Processors
  - Systems-on-Chip
  - Mobile Phone
  - Digital Signal Processing
  - Network/Security
  - Graphics/Multimedia/Game
- Software:
  - Compiler technology
  - Operating System/Chip Interaction
  - Performance Evaluation
- Other Technologies
  - Advanced Packaging Technology
  - Reliability and Design for Test
  - Advanced Semiconductor Process Technology

AUTHOR INFORMATION and FORMAT

Presentations at HOT Chips are in the form of 30-minute talks using PowerPoint or .pdf. Presentation slides will be published in the HOT Chips Proceedings. Participants are not required to submit written papers, but a select group will be invited to submit a paper for inclusion in a special issue of IEEE Micro.

Submissions are evaluated by the Program Committee on the basis of performance of the device(s), degree of innovation, use of advanced technology, potential market significance, and anticipated interest to the audience. Research and software contributions will be evaluated with similar criteria.

Submissions consist of a title, extended abstract (two pages maximum), and the presenter's contact information (name, affiliation, job title, address, phone(s), fax, and email). Please indicate whether you have submitted, intend to submit, or have already presented or published a similar or overlapping submission to another conference or journal. Also indicate if you would like the submission to be held confidential.

Please submit your extended abstract in plain ASCII text by March 24, 2008 via:
https://www.softconf.com/starts/hotchips20
Submissions containing figures may be submitted in .pdf, but plain ASCII text is strongly preferred.

Send questions relating to the program to the program chairs at: program2008@hotchips.org, and send questions relating to conference operation or organization to the general chair, Don Draper, at:
info2008@hotchips.org

Sponsored by the Technical Committee on Microprocessors and Microcomputers of the IEEE Computer Society and the IEEE Solid State Circuits Society

IEEE SSCS

Program Committee Co-Chairs:

Christos Kozyrakis	Stanford
Jan-Willem van de Waerdt	NXP Semiconductors

IEEE computer society

Check the HOT CHIPS 20 web page for updates: http://www.hotchips.org

COVER FEATURE

Wikis: 'From Each According to His Knowledge'

Daniel E. O'Leary
University of Southern California

Wikis offer tremendous potential to capture knowledge from large groups of people, making tacit, hidden content explicit and widely available. They also efficiently connect those with information to those seeking it.

Much has been written about wikis in recent years by researchers, journalists, bloggers, and wiki software vendors. Not surprisingly, most of this information appears in wikis themselves. Given the explosive growth in wiki applications and the controversies surrounding the technology, it is useful to sort through the claims and criticisms to better understand what wikis are, how they are used, their advantages and limitations, and various issues surrounding their implementation.

WHAT IS A WIKI?

In 1994, Ward Cunningham implemented the first wiki, the WikiWikiWeb, to promote the exchange of ideas among fellow programmers on his consultancy's website (http://en.wikipedia.org/wiki/Ward_Cunningham). Shown in Figure 1, the WikiWikiWeb was written in Perl and based on a HyperCard stack Cunningham wrote in the late 1980s. Today, wiki software applications are based on numerous languages, including Java, Lisp, PHP, Smalltalk, Python, and Ruby (http://en.wikipedia.org/wiki/List_of_wiki_software).

"Wiki" is Hawaiian for quick, and, as the term suggests, the technology's initial goal was to give users the ability to quickly put content on the Web. Today, however, a wiki's purpose depends on who you ask and what kind of application is being developed. In general, wikis are designed to facilitate quick and easy content

- generation,
- collaboration, and
- distribution.

With wikis, multiple users can connect virtually in time or space—from private communities within enterprises to the general public—to create, update, and share knowledge with others.

Wikis typically allow users to

- add new content,
- link to other related content,
- edit existing content,
- organize and structure content,
- view content, and
- access a history of contributed content.

Most wiki contributions are written, but they can include media such as images, videos, and sound files. Web-based documents are created collaboratively in a simplified markup language, or "wikitext," using a Web browser over the Internet or an intranet. This enables nonprogrammers to create wiki applications and add new features without having to be familiar with the code base.

Wikis use various mechanisms to track the history of contributed content so that users can see who made what changes and when. Figure 2 provides an example of a wiki tool that compares versions of documents.

Knowledge management

Over the years, researchers have offered many proposals to facilitate knowledge management, particularly at the enterprise level.[1] However, the promise of various tools and applications to make tacit knowledge explicit remains largely unfulfilled—much tacit knowledge

Published by the IEEE Computer Society

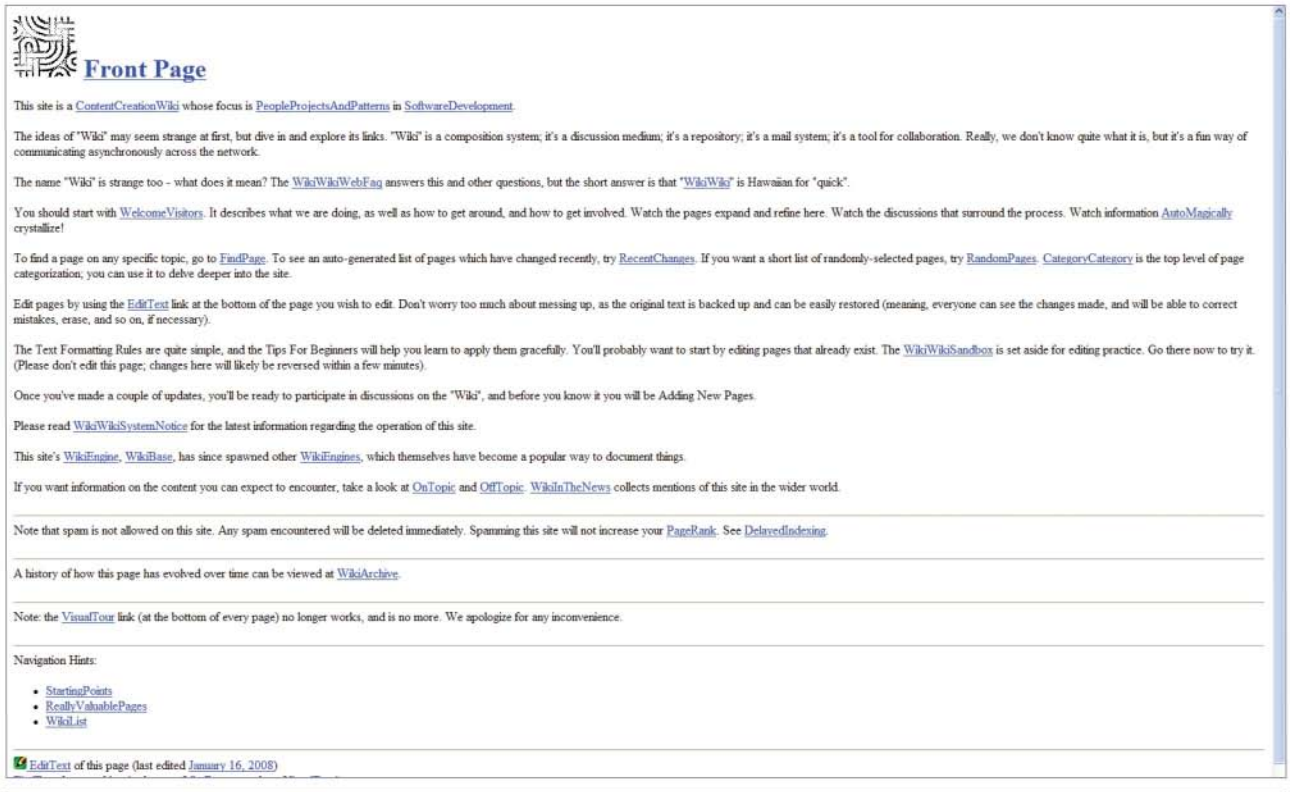

Front Page

This site is a ContentCreationWiki whose focus is PeopleProjectsAndPatterns in SoftwareDevelopment.

The ideas of "Wiki" may seem strange at first, but dive in and explore its links. "Wiki" is a composition system; it's a discussion medium; it's a repository; it's a mail system; it's a tool for collaboration. Really, we don't know quite what it is, but it's a fun way of communicating asynchronously across the network.

The name "Wiki" is strange too - what does it mean? The WikiWikiWebFaq answers this and other questions, but the short answer is that "WikiWiki" is Hawaiian for "quick".

You should start with WelcomeVisitors. It describes what we are doing, as well as how to get around, and how to get involved. Watch the pages expand and refine here. Watch the discussions that surround the process. Watch information AutoMagically crystallize!

To find a page on any specific topic, go to FindPage. To see an auto-generated list of pages which have changed recently, try RecentChanges. If you want a short list of randomly-selected pages, try RandomPages. CategoryCategory is the top level of page categorization; you can use it to delve deeper into the site.

Edit pages by using the EditText link at the bottom of the page you wish to edit. Don't worry too much about messing up, as the original text is backed up and can be easily restored (meaning, everyone can see the changes made, and will be able to correct mistakes, erase, and so on, if necessary).

The Text Formatting Rules are quite simple, and the Tips For Beginners will help you learn to apply them gracefully. You'll probably want to start by editing pages that already exist. The WikiWikiSandbox is set aside for editing practice. Go there now to try it. (Please don't edit this page; changes here will likely be reversed within a few minutes).

Once you've made a couple of updates, you'll be ready to participate in discussions on the "Wiki", and before you know it you will be Adding New Pages.

Please read WikiWikiSystemNotice for the latest information regarding the operation of this site.

This site's WikiEngine, WikiBase, has since spawned other WikiEngines, which themselves have become a popular way to document things.

If you want information on the content you can expect to encounter, take a look at OnTopic and OffTopic. WikiInTheNews collects mentions of this site in the wider world.

Note that spam is not allowed on this site. Any spam encountered will be deleted immediately. Spamming this site will not increase your PageRank. See DelayedIndexing.

A history of how this page has evolved over time can be viewed at WikiArchive.

Note: the VisualTour link (at the bottom of every page) no longer works, and is no more. We apologize for any inconvenience.

Navigation Hints:

- StartingPoints
- ReallyValuablePages
- WikiList

EditText of this page (last edited January 16, 2008)

Figure 1. WikiWikiWeb. Ward Cunningham implemented the first wiki in 1994 to promote the exchange of ideas among fellow programmers on his consultancy's website.

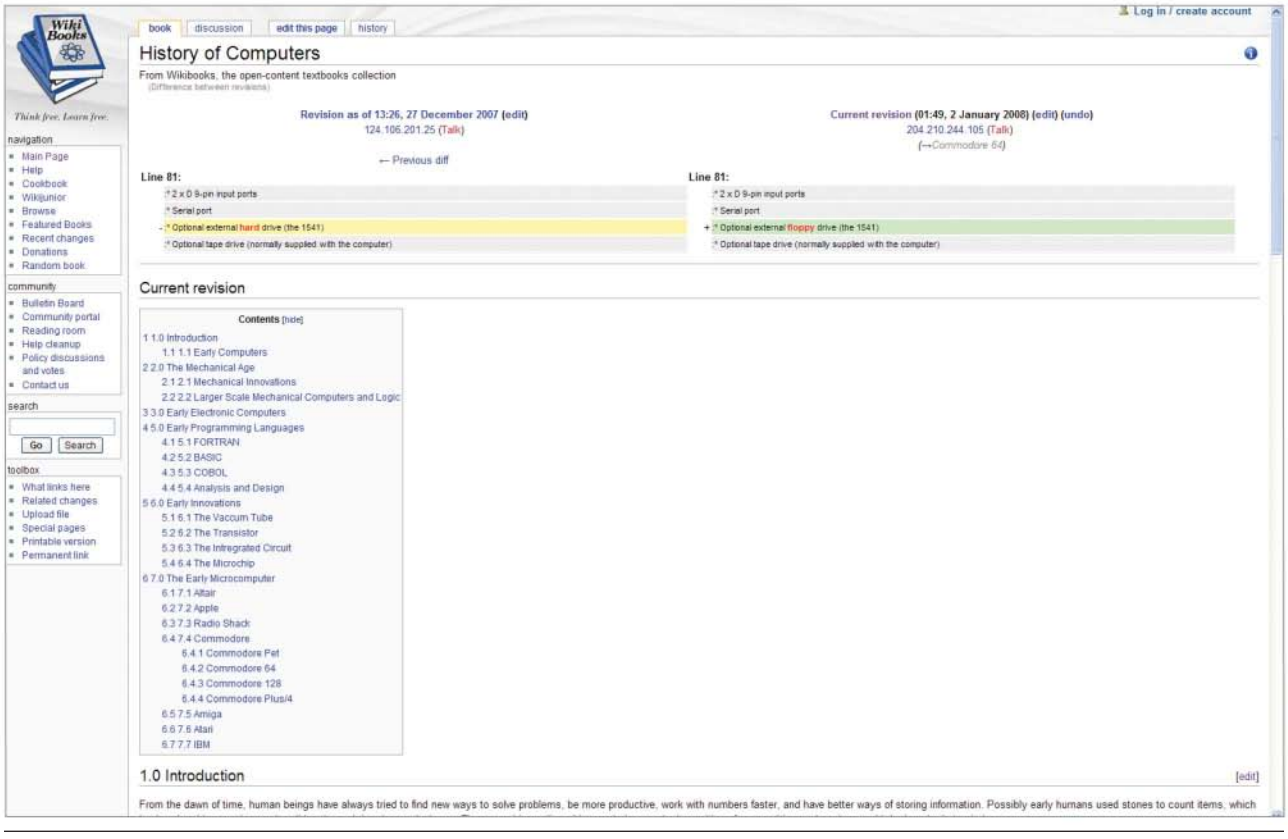

book | discussion | edit this page | history | Log in / create account

History of Computers

From Wikibooks, the open-content textbooks collection

Revision as of 13:26, 27 December 2007 (edit)
124.105.201.25 (Talk)

Current revision (01:49, 2 January 2008) (edit) (undo)
204.210.244.105 (Talk)
(→Commodore 64)

← Previous diff

Line 81:
* 2 x D 9-pin input ports
* Serial port
- * Optional external hard drive (the 1541)
* Optional tape drive (normally supplied with the computer)

Line 81:
* 2 x D 9-pin input ports
* Serial port
+ * Optional external floppy drive (the 1541)
* Optional tape drive (normally supplied with the computer)

Current revision

Contents [hide]
1 1.0 Introduction
1.1 1.1 Early Computers
2 2.0 The Mechanical Age
2.1 2.1 Mechanical Innovations
2.2 2.2 Larger Scale Mechanical Computers and Logic
3 3.0 Early Electronic Computers
4 5.0 Early Programming Languages
4.1 5.1 FORTRAN
4.2 5.2 BASIC
4.3 5.3 COBOL
4.4 5.4 Analysis and Design
5 6.0 Early Innovations
5.1 6.1 The Vacuum Tube
5.2 6.2 The Transistor
5.3 6.3 The Integrated Circuit
5.4 6.4 The Microchip
6 7.0 The Early Microcomputer
6.1 7.1 Altair
6.2 7.2 Apple
6.3 7.3 Radio Shack
6.4 7.4 Commodore
6.4.1 Commodore Pet
6.4.2 Commodore 64
6.4.3 Commodore 128
6.4.4 Commodore Plus/4
6.5 7.5 Amiga
6.6 7.6 Atari
6.7 7.7 IBM

1.0 Introduction [edit]

From the dawn of time, human beings have always tried to find new ways to solve problems, be more productive, work with numbers faster, and have better ways of storing information. Possibly early humans used stones to count items, which

Figure 2. Wikibooks "history" interface. Wiki software allows for version comparisons of documents so that contributors can see who made what changes and when.

remains inaccessible. Wikis have the potential to gather such knowledge from far-reaching sources.

Wikis satisfy four key knowledge management needs by

- capturing knowledge from those who have it,
- converting knowledge into an explicitly available format,
- connecting those who want knowledge with those who have it, and
- linking knowledge to knowledge.[2]

In classic knowledge management, acquisition experts are responsible for capturing knowledge from domain experts. Wikis offer a nonintrusive means of capturing information by removing the intermediary and letting people share knowledge directly. Wikis also make information or sources exclusively available to the contributor generally available; users thus directly influence the knowledge base's structure and content. In addition, by making available information about contributors, wikis facilitate connections between interested parties. Finally, through the use of hypertext, wikis let contributors link appropriate knowledge.

Wikis all rely on the community at large rather than an elite group to advance knowledge, education, and discussion.

Mass collaboration

Wikis are particularly effective in situations in which a large group of people want to leverage their collective knowledge to achieve some goal. For example, during the 2004 US presidential contest, one campaign used a wiki to compile political news stories for their candidate.[3] This approach enabled some 400 staffers to focus on different areas of coverage—for example, around a given periodical. The resulting database served as the basis of twice-daily briefing documents.

Within an enterprise, the choice of whether to implement a wiki depends on the nature of the information as well as the number of users. If a group wishes to keep information private, then wikis, unless tightly limited, are not appropriate as a means of fostering collaboration.

Transparency

To increase participation, content must be transparent; otherwise, multiple participants will not be able to provide coherent and related contributions. Wikis provide transparency by letting users see what others have contributed, thereby converting individual knowledge into communal knowledge.

Pull versus push

Wikis facilitate the connection between those who have information and those who need it. This "pull" mechanism is useful for organizations that want to continually draw on a dynamic, ordered information set. The alternative is to "push" static, unordered information directly to users, either individually or as a group. E-mail represents the most common form of this approach.

WIKI APPLICATIONS

A broad range of general and enterprise wiki applications is in use today.

General applications

The most well-known general wiki application is Wikipedia, the multilingual online encyclopedia that relies on volunteers from around the world to contribute and edit content on any given topic. Launched in January 2001 by Jimmy Wales, it is one of the 10 most popular websites and currently contains more than 9 million articles in 253 languages (http://en.wikipedia.org/wiki/Wikipedia).

The project's tremendous success spawned numerous siblings now operated, along with Wikipedia, by the nonprofit Wikimedia Foundation. These include Wiktionary, a dictionary of term meanings, synonyms, etymologies, and translations; Wikibooks, a collection of open source textbooks and other learning materials; Wikiquote, a compendium of quotations from prominent people and works; Wikisource, a library of public domain texts and other source documents; Wikimedia Commons, a repository of images, sounds, and video; and Wikinews, a source for reports by citizen journalists.

Tens of thousands of independent wiki applications have sprung up on the Web to serve communities interested in broad topics like computing, travel, and entertainment as well as niche subjects such as the online role-playing game *World of Warcraft*. For example, Wikia, a for-profit company cofounded by Jimmy Wales, alone hosts more than 4,700 wiki communities (www.wikia.com/wiki/About_Wikia).

Although some wikis impose restrictions on contributions, all rely on the community at large rather than an elite group to advance knowledge, education, and discussion. The power of wikis to reach a broad constituency has not been lost on technology-minded political candidates, who are beginning to incorporate them into their campaigns (http://vote.peteashdown.org/wiki/index.php/Main_Page).

Enterprise applications

Wikis have many applications within businesses and other organizations.

Wikipedia imitations. The high visibility of Wikipedia has led many companies to replicate this type of application internally.[3] These internal wikis are typically designed to support particular functions by letting employees input

information as appropriate in an encyclopedia-like setting. For example, a business might employ a wiki-type product directory to record changes and new offerings.

Meeting setup. Wikis can help mitigate information overload.[4] For example, they can facilitate meetings by gathering input in advance from attendees and making it generally available. This saves time, particularly in the case of multiday meetings with much to assimilate, by enabling participants to review what others have to contribute prior to the meeting so that they can concentrate on areas that need attention.

Project management. Companies can use wikis to capture information about projects. Participants can post documents and progress reports or generate and massage information related to a project on the wiki. For example, CommSecure, an Australian provider of e-billing and e-payment solutions, employs a wiki to help track the implementation status and related documentation of different projects.[5] This can facilitate buy-in by letting participants help construct key inputs and making constraints transparent.

Best practices. Employees can use wikis to describe best practices. For example, the wiki "Library Success" is a "one-stop shop for great ideas and information for all types of librarians" (www.libsuccess.org/index.php?title=Main_Page). Another wiki's expressed goal is to share best practices about the Common Base Event, a fundamental systems management standard (www.ibm.com/developerworks/wikis/display/CBEbestpractice).

Taxonomy development. Wikis can simplify taxonomy development within an enterprise, which generally requires the cooperation of multiple parties. Individual users can propose a portion of the taxonomy and its associated explanation, and others can point out their limitations and suggest changes.

Competitive intelligence. Wikis can be used to gather competitive intelligence, a function traditionally performed by a small group within the organization that acts in relative secrecy. SAP, one of the world's largest business software companies, employs a wiki to monitor how its pricing tactics and sales strategies are working in the field.[3] By making the process open and participatory, the company can get better and more timely collective intelligence and make it available to more people.

WIKI ADVANTAGES AND LIMITATIONS

In determining whether to implement a wiki, an enterprise or other organization must balance the advantages of the technology with its limitations as well as match the wiki's capabilities to the desired objectives. Figure 3 summarizes some of the pluses and minuses of wikis.

WIKI ADVANTAGES
- Wikis generate a network of knowledge by linking people and content
- Wikis can build consensus
- Wikis collect knowledge from multiple sources
- Wikis engage contributors
- Wikis can be as accurate as traditional published sources
- Wikis delegate control to contributors
- Wikis provide a forum to help users manage their behavior

WIKI LIMITATIONS
- Wikis often do not provide author information, raising questions about content accuracy
- Wikis typically lack referees or peer review, which provide some quality assurance
- Wikis can hinder as well as build consensus, focusing on contributors' conflicting opinions
- Contributors can easily introduce bias
- Wikis can compromise information security
- Wikis can encourage scope creep
- Contributions can decrease over time
- Wikis can expose an organization to legal problems
- Wikis are subject to vandalism
- Wikis can be contrived to look genuine but have an ulterior motive
- Wiki content is generally not available in a machine-processable format

Figure 3. Wikis have both advantages and limitations.

Advantages

Wikis offer numerous advantages.

Structure. At the highest level, wikis use a vocabulary or ontology to explicitly organize contributions. However, the use of hypertext to link related concepts and articles within the wiki embeds additional structure. Some wikis, such as Wikipedia, also contain references and external links to other subjects.

Consensus. Wikis can build consensus because many participants often "sign off" on the content. In fact, building consensus is Wikipedia's "fundamental model for editorial decision-making" (http://en.wikipedia.org/wiki/Wikipedia:Consensus). Wikis typically encourage a neutral point of view and have mechanisms to resolve disputes among contributors.

Collective wisdom. Because wikis are generally open, democratic environments, they harness the "wisdom of the crowd." Ideally, content draws on a wide range of contributors with varying perspectives and expertise. Everyone in the community has an opportunity to evaluate the quality of contributions, and those who have an interest in or are knowledgeable about a topic can add to or modify content.

User engagement. Wikis engage users by letting them express themselves freely and for all to see. Although most wikis have etiquette guidelines and codes of conduct prohibiting, for example, hateful content or personal attacks, individuals generally have tremendous flexibility in what they post. Users derive satisfaction from being part of a communal effort as well as seeing their creativity on display.

Accuracy. Contrary to the claims of some critics, wiki accuracy can be comparable to published sources. For example, one recent study found that Wikipedia had roughly four inaccuracies per entry, only one more than *Encyclopedia Britannica*.[6] Some wikis have verifiability guidelines that encourage contributors to cite reliable sources.

Delegation of control. Wikis delegate control of content to potential contributors. This is an advantage in organizations where management seeks bottom-up input on particular issues or processes.

User management. Wikis can help manage users as well as contributors by providing widespread access to equivalent standards for actions and behaviors, whether implicitly or explicitly.

Wiki limitations

Wikis also have several limitations.

Lack of authority. Users might want assurance that material they obtain online is backed by some authority or level of expertise. Unfortunately, in many cases there is limited information about authors of wiki material. For example, a Wikibooks contributor named "Psychofarm" has written books on both Mac OS and Asian honey chicken salad, while another has offered works on both physics and accounting. Such broad interests naturally raise doubts as to whether these authors have the necessary expertise.

No referees. Few wikis referee content to any appreciable extent, if at all, because that violates the open wiki spirit. Consequently, there is no guarantee that information in wikis is accurate or even reasonable. Wikipedia, for example, has had well-documented problems with users submitting invalid information.[7] In contrast, published research is typically peer-reviewed and edited, providing some quality assurance.

"Too many cooks in the kitchen." Wikis can hinder as well as build consensus. If multiple contributors express conflicting points of view or alternative solutions, the resulting content might be incoherent or focus on differences rather than similarities. Wikis can also misleadingly give the appearance of consensus if only one or a small group of contributors dominate the process early on, thereby thwarting further discussion.

Bias. Although many wikis have policies advocating a neutral point of view, their open nature makes it easy to introduce biased information. For example, a former MTV veejay and podcasting pioneer was caught anonymously editing the Wikipedia entry on podcasting to take credit for its development away from others and inflate his own role.[8]

Information insecurity. Wikis can compromise information security. Organizations often compartmentalize data, giving different pieces of information to different users, but wiki users could inadvertently share data that should not be available to all who have access to the wiki. For example, Microsoft purposely separates product and market information, and users able to intermingle data through a wiki could gain deep insights into the company's revenue stream.[9]

Scope creep. Because wiki contributors can range from amateurs to professionals, from beginners to experts, the resulting content might be too amorphous to be of use to any particular group. Scope creep is a common problem on complex projects, and wikis can encourage it by facilitating changes in team composition.

Decreased contributions. Wikis, particularly discretionary ones, can suffer a slow death. In some cases, contributions are initially heavy but subsequently decrease as participants turn to other activities. In other cases, contributions are light to begin with, increase as users familiarize themselves with the technology, and then decline as the uniqueness of the technology wears off. Unfortunately, both scenarios result in a similar outcome: decreased contributions over time.

Legal problems. Enterprise applications such as project management rely on contributors being frank and honest, but openness in company e-mail has led to expensive lawsuits—even in instances with only one recipient of a message. It is easy to imagine how a wiki could, by disseminating sensitive or private data to numerous people, expose an organization to all sorts of legal problems.

Vandalism. Wikis are only as good as their contributors, and these can include users who submit obscenities, personal attacks, and deliberate nonsense. Vandalism has actually forced some organizations to cancel wiki applications. For example, the *Los Angeles Times* closed down its "Wikitorial" feature because of contributors' repeated use of foul language.[10]

Contrived wikis. Because wikis facilitate consensus, some use them to try to generate consensus within an enterprise or the general public. Contrived wikis are implemented by some anonymous source to look like a standard wiki, with open contributions, but are actually not open and designed to influence public opinion.

Human consumption. In general, wikis are generated by and for humans. However, many knowledge management systems, such as rule-based systems, attempt to put information in a machine-consumable format, intermediary to human consumption. Such machine-based consumption is generally beyond the scope of wikis.

IMPLEMENTING WIKIS

Some organizations that implement a wiki might expect to simply "build it and they will come" (and use it). However, the open nature of wikis raises several issues that are often ignored.

Author information. Enterprise wikis usually keep data about wiki authors. As Figure 4a shows, capturing such information can be critical to achieving user

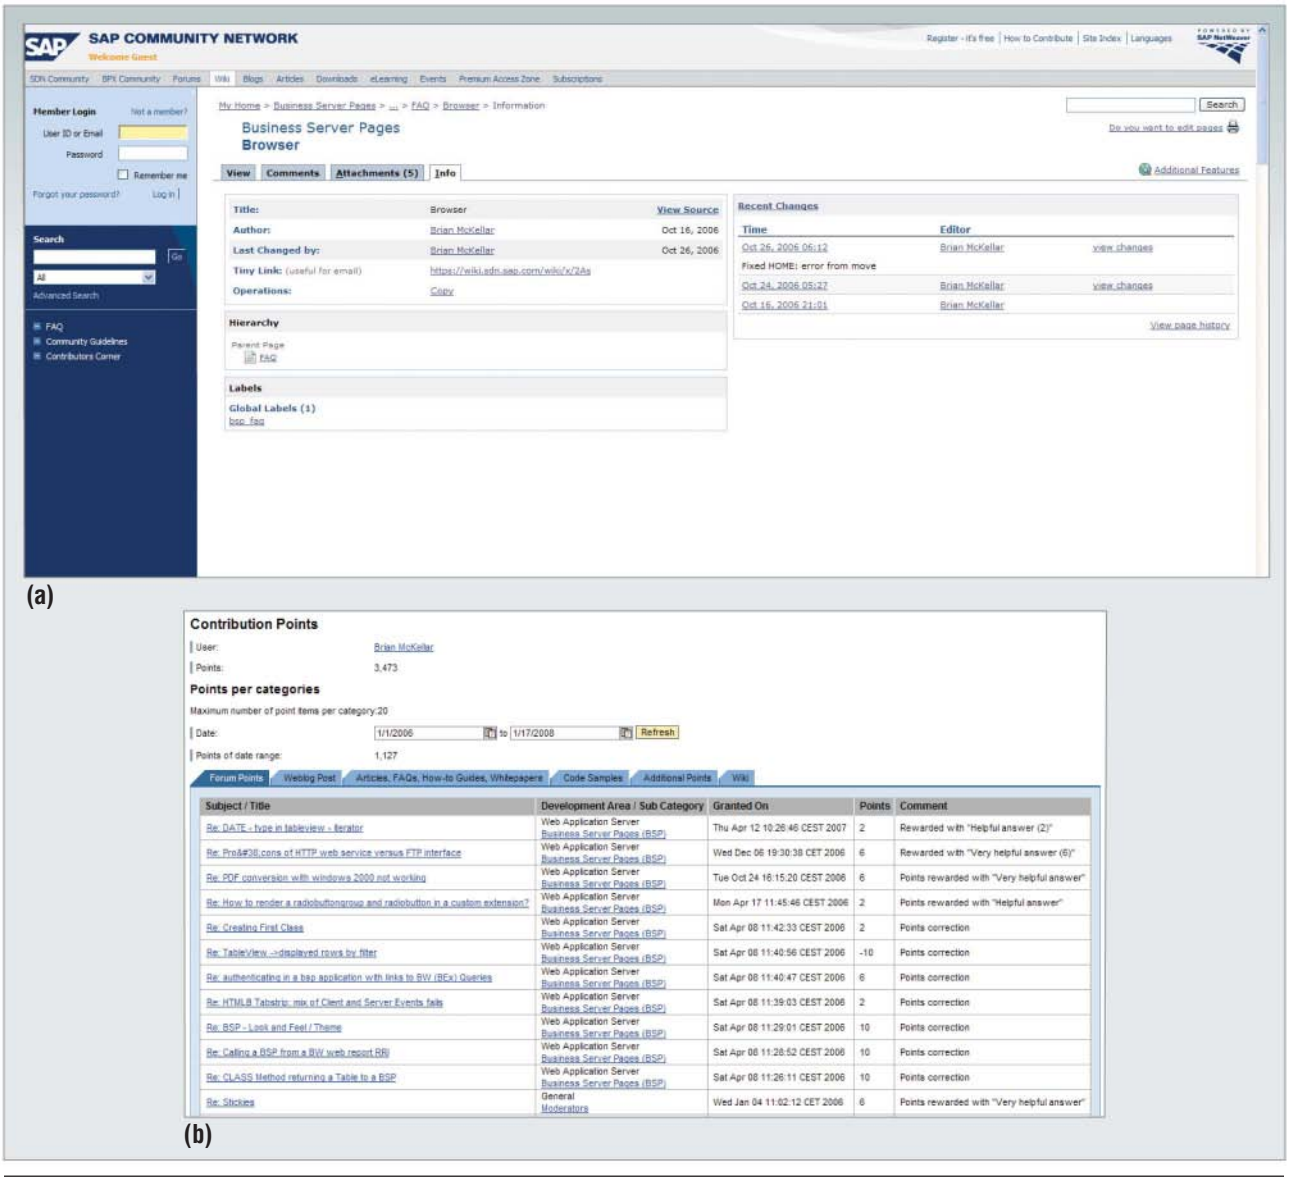

Figure 4. SAP wiki. (a) Capturing author information can be critical to achieving user acceptance of content. (b) Issuing "points" with associated benefits to employees for their efforts encourages participation.

acceptance of content. It can also foster connections between users and authors. Google recognizes this and is implementing its own competitor to Wikipedia, Knol, that prominently displays authors' names (http://blogoscoped.com/archive/2007-12-14-n19.html).

Incentives to participate. Wikis should provide potential contributors with incentives to participate. As Figure 4b shows, one way to do this within an enterprise is to issue "points" to employees, with some reward upon reaching a certain threshold, for their efforts. In addition to displaying authors' names, Google's Knol will let authors include advertisements and make money from their contributions.

Administration. Ideally, wikis should have an administrator who referees and manages the changing content. However, if the wiki is substantial, such as Wikipedia, no one person or even group can monitor all of the changes in real time.

Change alerts. One way to provide control over changes is to alert those who have indicated interest in a particular subject or whose previous contribution has been altered by another user. Participants who know that changes they make to existing content will be broadcast to the original author will likely be more discriminating, while those whose contributions have been edited will have a chance to quickly review the changes for inaccuracies or other issues.

Access and registration. The original philosophy of wikis was to let all users contribute and change content. However, such openness can lead to vandalism, tampering, compromised data, and other problems, particularly

in noncorporate settings where there is little recourse for destructive acts.

In corporate environments, it is important to determine whether wiki access should be open to outsiders or limited to employees, managers, or a select group of users within the company. Does the wiki contain information—for example, about product faults—or controversial content that, if made available to the wrong people, could negatively impact sales, compromise proprietary secrets, or lead to costly litigation?

Perhaps the least intrusive way to control access is to notify potential contributors that their IP address is being captured. Another method is to require that users register with a valid e-mail address and log in with a username and password. Although the effectiveness of these steps can be mitigated, they at least provide some potential control over users.

Contributor capabilities. One way to manage users is to categorize them according to their capabilities. Wikibooks distinguishes contributors according to their fluency in English and other languages. For example, User en-N connotes a native English speaker, while User en-0, -1, -2, and -3 represent users with zero, basic, intermediate, and advanced levels of English, respectively. To limit wiki access in enterprise settings, contributors can be assigned "roles" based on their responsibilities or level of expertise.

User practice. Although most wikis are relatively simple, they can be intimidating to first-time users. Many sites therefore provide a "sandbox" that lets contributors learn the wiki's various features and practice, thereby limiting potential mistakes. Sandboxes might also facilitate user buy-in.

Policies and guidelines. Wiki contributors should clearly understand what they can and cannot do. The site should therefore offer a list of mandatory policies and advisory guidelines, subject to community approval. For example, Wikipedia users must respect other contributors, respect copyrights, avoid bias, and include only verifiable information (http://en.wikipedia.org/wiki/Wikipedia:Key_policies_and_guidelines).

Copyrighted material. In many settings, wiki-based materials cannot be copyrighted. For example, Wikibooks considers all contributions to fall under the terms of the GNU Free Documentation License (http://www.gnu.org/licenses/fdl.html). Wikibooks warns potential violators that "the posting of copyrighted material without the express permission of the copyright holder(s) is possibly illegal and is a violation of our copyright policy" (http://en.wikibooks.org/wiki/Wikibooks:Copyrights).

Project completion estimates. Although wikis are typically open ended, some projects can have a completed format. In these cases, providing users with an estimate of how much work has been done can be helpful. For example, Wikibooks indicates whether text for any given project is "sparse" (0 percent), "developing" (25 percent), "maturing" (50 percent), "developed" (75 percent), or "comprehensive" (100 percent). Because there are likely to be multiple contributors, and completeness is in the eye of the beholder, estimates can be highly subjective.

Design for participation. Because wikis depend on contributors, any implementation should be designed to facilitate participation. Ross Mayfield, cofounder of Socialtext, the leading enterprise wiki company, suggests starting small with a pilot project that applies a wiki solution to a single process or application.[3] Once the project participants have evaluated the tool through a forum or discussion group, they can "take it public" by each inviting five others in the organization to use the wiki. This can be repeated with successively larger waves of contributors, gradually building a community, adding content, and evolving norms.

Researchers have used wikis to generate and maintain ontologies and taxonomies.

Personalization. Many wiki applications let users personalize some aspects such as privacy settings, link formats, image size, editing options, browser appearances, date format, and time zone.

EMERGING AI APPLICATIONS IN WIKIS

In simple terms, artificial intelligence aims to incorporate human intelligence into computer-based applications or analysis. There are numerous potential applications of AI in the area of wikis.

Wikis provide substantial structured material about particular subjects, and researchers have used them to generate and maintain ontologies[11] and taxonomies.[12] Similarly, group input could be used in a wiki to generate rule-based knowledge to capture insights and identify conceptual relationships. Systems designed to improve knowledge by intelligent questioning and answering could also leverage wiki content.

Just as electronic auction sites generate reliability or quality estimates about buyers and sellers, AI systems could search the Internet and other wikis to find out what particular authors have contributed on various topics and generate trust or expertise indices. Researchers also could develop intelligent agents to search multiple wikis and assemble material for a comprehensive article on a subject.

Wikis such as Wikipedia address the same topics in numerous languages. Researchers could use this multilingual data to disambiguate topics, terms, or words; generate translations; or analyze structure in a subject area.

Researchers also could use AI systems to help secure wikis, whose open nature makes them particularly vulnerable. For example, concept-based systems could identify vandalism and exclude such contributions prior to

posting, while intrusion-detection systems could leverage information gathered about contributors to unmask illegitimate users.

Wikis can be used as a training ground to search for knowledge obtained through machine-learning approaches. Further, annotating wikis with machine-readable content would make them both human and machine-friendly.

Wikis offer tremendous potential to capture knowledge from large groups of people, making tacit, hidden content explicit and widely available. They also efficiently connect those with information to those seeking it: "from each according to his knowledge, to each according to his need." Although wikis have inherent limitations that make them inappropriate in certain settings and for some applications, they are likely to replace existing processes and technologies, providing organizations with a wide range of additional capabilities.

Acknowledgment

The author thanks the anonymous referees for their comments on an earlier version and extends appreciation to the editors for their efforts to help make this a better article.

References

1. D.E. O'Leary, "Enterprise Knowledge Management," *Computer*, Mar. 1998, pp. 54-61.
2. D.E. O'Leary, "Knowledge-Management Systems: Converting and Connecting," *IEEE Intelligent Systems*, May/June 1998, pp. 30-33.
3. R. Mayfield, "How to Start a Wiki;" podcast, www.veotag.com/player/?u=rivtyskosd.
4. J. Spira, "Information Overload—A Growing Problem," *Collaboration Loop*, 1 Feb. 2007; www.collaborationloop.com/index.php?option=com_content&task=view&id=1945&Itemid=39.
5. L. Wood, "Blogs & Wikis: Technologies for Enterprise Applications?," *The Gilbane Report*, Mar. 2005, pp. 2-9; http://gilbane.com/artpdf/GR12.10.pdf.
6. J. Giles, "Internet Encyclopedias Go Head to Head," *Nature*, 15 Dec. 2005, pp. 900-901.
7. D. Terdiman, "Growing Pains for Wikipedia," CNET News.com, 5 Dec. 2005; http://news.com.com/Growing%20pains%20for%20Wikipedia/2100-1025_3-5981119.html?tag=st.prev.
8. R. Cadenhead, "Adam Curry Caught in Sticky Wiki," *Workbench*, 1 Dec. 2005; www.cadenhead.org/workbench/news/2818/adam-curry-caught-sticky-wiki.
9. D.E. O'Leary and M.L. Markus, "Microsoft's Management Reporting: SAP, Data Warehousing, and Reporting Tools," *J. Emerging Technologies*, vol. 3, no. 1, 2006, pp. 129-141.
10. P. Naughton, "Foul Language Forces LA Times to Pull the Plug on 'Wikitorial,'" *Times Online*, 21 June 2005; http://technology.timesonline.co.uk/tol/news/tech_and_web/article535749.ece.
11. M. Hepp, D. Bachlechner, and K. Siorpaes, "OntoWiki: Community-Driven Ontology Engineering and Ontology Usage Based on Wikis," *Proc. 2006 Int'l Symp. Wikis*, ACM Press, 2006, pp. 143-144; www.heppnetz.de/files/ontowikiDemo-short-camera-ready.pdf.
12. A.L. Burrow, "Negotiating Access within Wiki: A System to Construct and Maintain a Taxonomy of Access Rules," *Proc. 15th ACM Conf. Hypertext and Hypermedia*, ACM Press, 2004, pp. 77-86.

Daniel E. O'Leary *is a professor in the Marshall School of Business at the University of Southern California. His research focuses on information systems, including enterprise resource planning systems and knowledge management systems. O'Leary received a PhD in information systems and systems engineering from Case Western Reserve University. He is a member of the IEEE, the ACM, the Institute for Operations Research and the Management Sciences, and the Decision Sciences Institute. Contact him at oleary@usc.edu.*

COVER FEATURE

EventWeb: Developing a Human-Centered Computing System

Ramesh Jain, University of California, Irvine

Dealing with emerging applications of computing requires taking a fresh look at our tools. EventWeb is a human-centered computing system that will give users a compelling experience by combining quality content, carefully planned data organization and access mechanisms, and powerful presentation approaches.

In a sense, all computing is human centered. Once humans successfully developed machines to augment their mechanical strength, they focused on developing machinery to augment their analytical abilities. The first analytical ability they addressed was simple mathematical calculations. This resulted in different types of calculators, leading finally to electronic computers in the middle of the past century.

Augmentation of human analytical facilities remains the primary goal of computers today. Computing has evolved rapidly, with advances in processing, storage, communication, sensing, and related areas. Another driving force is the emergence of increasingly challenging sensory-data problems, including video and audio, that computers could solve.

Current interest in human-centered computing suggests new winds blowing in the computing community. HCC combines many powerful and independent approaches in different aspects of computing, ranging from human-computer interfaces (HCIs), computer vision, speech recognition, and pervasive computing to virtual reality systems. Most computing systems allow creation of powerful interfaces using audiovisual techniques. However, designing HCC systems that give users a compelling experience requires quality content, carefully planned data organization and access mechanisms, and powerful presentation approaches.

Content is ultimately what's interesting to users, so its quality is important. Content quality includes its credibility, depth, and timeliness. The challenges that HCC researchers face aren't limited to HCI or gestural approaches, but go deeper, into correctly organizing multimodal data from disparate sources, finding the best combination of multimedia sources to communicate the message or experience, presenting and distributing these sources for the best subjective quality of experience, and helping advance human knowledge and build stronger communities using these approaches.

Most current approaches in computer science and applications evolved from alphanumerical data—the dominant data in the early days of computing. Researchers have tried to extend these approaches to increasingly multimodal dynamic data. To deal with emerging applications of computing, ranging from biology to entertainment, and from security to business, we need to take a fresh look at our tools.

A combination of technological advances, a reduction in barriers to interactions among different parts of the world, and the quest for solving increasingly difficult problems has created a situation that's unique in its potential to impact the course of human civilization.

DISRUPTIVE INNOVATIONS IN COMPUTING

Computing has already gone through two major disruptive evolutionary stages, and is now on the verge of a third. Although the first two were revolutionary, the third stage will effect the most long-term fundamental changes in how computing influences human civilization.

Table 1 shows the basic features and characteristics of the three stages of this evolution. In addition to the nature of input and output, the nature of applications and user expectations has changed. Also, each phase

Published by the IEEE Computer Society

Table 1. Evolutionary stages in computing.

Feature	Stage 1	Stage 2	Stage 3
Input	Data	Documents	Multimodal data
Output	Results	Information	Experiences
Processing	Computations	Information extraction	Experience processing
Devices	Mainframe, workstations	Personal computers, Internet	Mobile phones
Applications	Computing	Information and communication	Insights and entertainment
User level	Trained	General population in the developed world	All humans
Interaction environment	Command languages	Windows	Experimental

builds on the previous phase's advances, so it subsumes and even enhances the previous phase's functionality.

Data and computation

The invention of electronic computers marked the first major event in computing. These computers processed data at an unimaginable speed and could perform calculations millions of times faster than humans. The ability to program computers opened avenues for many previously unimaginable applications.

This initial evolutionary phase focused on scientific and engineering computations. Soon, the business community and other organizations realized the potential of computing with alphanumeric data. This extended computing to businesses and large organizations. Early mainframe computers and workstations represented this style of computation. The terms "computer" and "computing" are the legacy of this phase.

Information and communication

The second major evolutionary stage brought personal computers, including laptops, and the Internet. In this phase, the emphasis shifted from data to information and communication technology (ICT). The rise of PCs and what-you-see-is-what-you-get (WYSIWYG) word-processing and spreadsheet programs brought computing from trained computer operators to the general population in the developed world.

The Internet and World Wide Web accelerated the ICT revolution. By networking computers to form a global pervasive network, users could connect alphanumeric data sources, including documents, and communicate data and information. The ICT revolution started unimaginable applications and affected human life in most developed and developing parts of the world.

Insights and experiences

The easy availability of multimodal sensory data and devices that can capture, play, store, and process this data is propelling the third major evolutionary stage. This phase will bring insights and experience to the forefront in the same way that the second phase focused on information and communication.

Insights refer to the perception of the true nature of things made possible when a human understands the relationships among a thing's different components. Insights help in deeper analysis and problem solving. We need insights to make decisions, and information to implement these decisions. Insights are closely correlated to experiences. Dictionary.com defines "experience" as the "active participation in events or activities leading to the accumulation of knowledge or skill" or as "knowledge or wisdom gained from what one has observed, encountered, or undergone."

People experience events and activities using their sensory tools: sight, sound, touch, smell, and taste. These observations constitute experience of the event resulting in insights related to the event and objects in the event. In the digital world, approaches that use sensing technology convert sight, sound, and touch to electronic form and then, after processing, convert it back to sensory data for human perception. Smell and taste remain relatively difficult to capture, convert to digital form, and then convert back to the original senses for presenting to humans.

It's common, however, to experience the sights and sounds of remote events, and touch is finding increasing use. Fortunately, people experience the world around them and create knowledge about their environment mainly through sight and sound. Thus, we can naturally extend digital experience to these dominant human senses.

This third stage is about experiencing events, saving experiences, gaining insights and knowledge from these experiences, and sharing these insights and experiences with others. It's no wonder that the past few years have seen the emergence of companies like Flickr, YouTube, and Facebook because these companies provide environments in which users can share experiences through photos, videos, and multimedia. Progress made in this stage could lead to revolutionary approaches for expanding access to the more than 80 percent of the Earth's population that computing has yet to reach.

APPLICATION ENVIRONMENTS

Given the emerging nature of data and computing and projecting emerging applications, several applica-

tion areas seem to be the natural extension of current popular computing application paradigms.

Immersive telepresence

Virtual reality systems use computing as a storytelling mechanism, with the user an integral part of the story. In these systems, the user is immersed in an environment. Developers are making immersive environments increasingly realistic by combining computer graphics technology with image processing, audio, and even tactile information. Videogames, for example, routinely use tactile processing in the form of force feedback and similar mechanisms.

In an immersive system,[1-3] a rendering engine creates a detailed model of the environment in response to user actions. On detecting a user action, the engine renders an appropriate segment of the model to keep the user immersed in the environment. The system's realism depends on its ability to render the model and adapt the environment to the user's actions.

In an immersive telepresence system, multiple sensors placed in the physical environment capture an event.

All these actions occur in an interactive setting. The system must have less latency than a person's perceptual limits. These systems usually combine a predefined environment model with known user behaviors to generate the situation model synthetically.

Suppose we modify such a system. First, we assume that the model is from a real, rather than synthetic, environment. Let's assume that we're in the US and the environment model is a football game taking place in India. Let's also assume that we've placed as many cameras and microphones as required for creating a realistic model of the sights and sounds of this event—including what's happening on the sidelines and what spectators are doing. And suppose this dynamic model is properly indexed using events. We'll call this a situation model to differentiate it from the static environment model. (A simpler version of this scenario is a live telecast of the game.)

But now let's assume that we can request what we want to see and from where on the field, and the system will render exactly that view using the model that exists at each moment. Effectively, we could be at any location on or off the field at any point in the game. Of course, in this scenario, we can't act, except to change our position. But everybody else—potentially millions of people—can enjoy the game from their own perspective and, if they desire, share their view with others.

Figure 1 shows this architecture. On the left side, different sensors and data sources feed the system through arrow A. Using the environment model, the system assimilates and indexes this information. It creates a situation model to represent the physical world of interest at every time moment. It also saves and indexes all sensory data, perhaps for later use. Arrow C is the interaction with different remote users. Each user requests to experience a different segment of the situation model, and the system generates that model using information from different sensor and data sources.

This is an immersive telepresence experience, based on reality that is "real" because the system's environment model comes from the real, rather than virtual, world.[4] Moreover, the situation model is prepared at every instant to represent the real world as captured using relevant sensors. This makes it different from the systems that use these models from artificial worlds. This model could prove equally effective in entertainment, videoconferencing, telemedicine, personal communications, scientific explorations, and education.

Information assimilation. Most current multimedia information systems deal with archived video, audio, and images. Continuous queries—that is, persistent queries that are issued once and then logically run continuously over live and unbounded streams, have recently become a major research area in data management. A media stream comes from a sensor device such as a video, audio, or motion sensor and produces a continuous or discrete signal, but typically a data-stream processor can't directly use it. To evaluate queries on media streams, a system needs to extract features continuously and assimilate them to form domain events.

When we place multiple sensors in a physical environment, their placement and the models of the events they're supposed to capture play important roles in information extraction and assimilation. My research focuses on knowledge about the location of sensors in a physical space and the role of spatially and temporally correlated information obtained from disparate sensors.[2,5,6] An environment model captures the physical placements and constraints on the information obtained from these sensors. My research in multimedia stream queries addresses issues related to live multimedia data.

In an immersive telepresence system, multiple sensors placed in the physical environment capture an event. In most cases, a significant amount of metadata associated with the environment, sensors, and event is available. The system should continuously process and assimilate this data to form a unified situation model. The unified model provides all information about the events and objects in the environment. In assimilating the data from sensors and other sources, the system considers each source as an observation source that contributes to the complete model.

Semantic indexing. Databases and search engines have traditionally used indexing to efficiently store and retrieve data. In addition to dealing with indexing at that level, an immersive telepresence system must deal

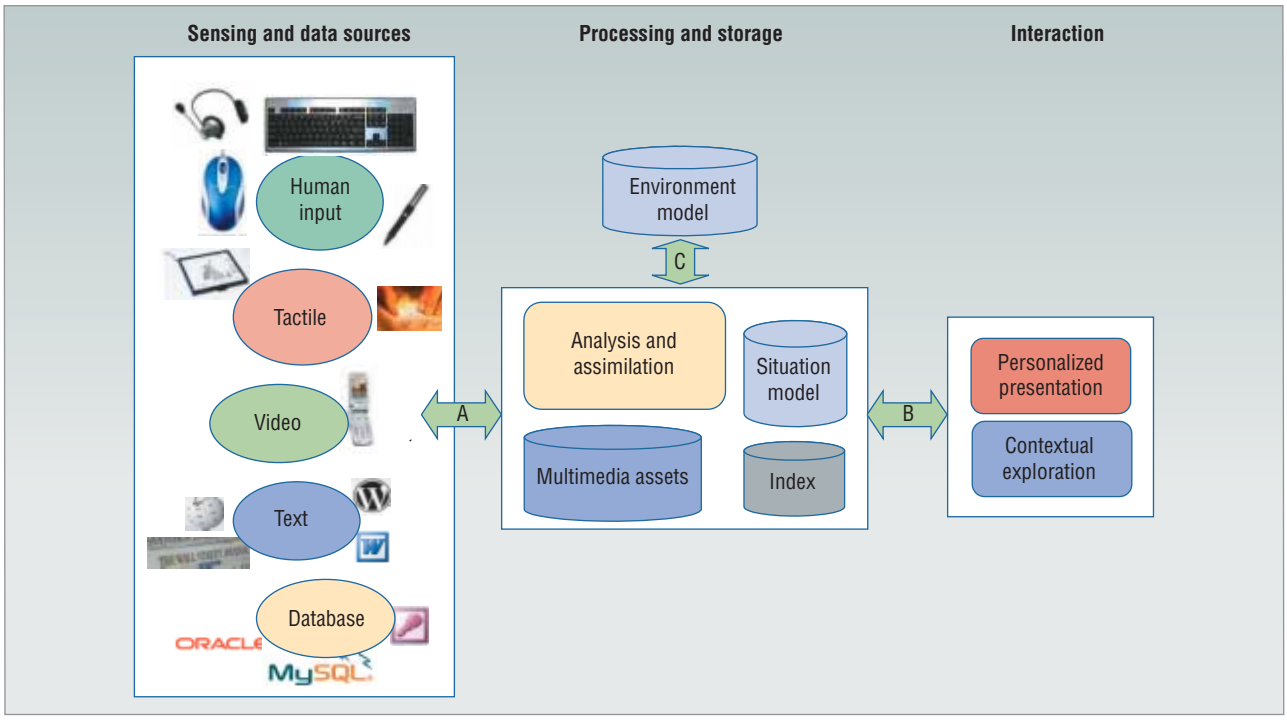

Figure 1. Immersive telepresence system architecture. On the left side, different sensors and data sources feed the system through arrow A. Arrow B is the interaction with different remote users.

with indexing at the semantic level. Current indexing techniques for different data types depend on metadata for that particular type. Metadata plays a key role in introducing semantics and determining how to use data. Schemas provide semantics in relational tables. XML is increasingly used to introduce semantics in strongly human-mediated environments.

Current information systems create data silos. They define and introduce the metadata for a particular data type, then index it and stash it neatly in a silo. Breaking down these silos can unify information. I proposed a unifying indexing system that introduces a layer on top of each data silo's metadata layer, or disparate data source.[7,8] The layer uses an event-based domain model and metadata to construct a new index that's independent of data type. This system can model an application domain in terms of events and objects.

An event ontology parses the data as it comes from the sensors and data sources and assimilates it to build a situation model that reflects knowledge about the event on the basis of information collected so far. An event index is essentially a list of spatiotemporal events as they occur. An event base stores the event's name, type, and all other relevant information.

We might not have access to the relevant information when the event is created. If this is the case, when it becomes available, the system attaches the information to the event. Thus, the event base is an organic database that grows as a result of many different processes running, in contrast to the current database form. The event base also stores links to original data sources, so the system can present the appropriate media in the context of a particular event.

Personalized distribution mechanisms. Different perspectives of the event might interest system users. The environment model contains all information used to generate different perspectives. In most cases, the system renders these perspectives by sending each user a different combination of sensor and other data streams. Moreover, because many different camera streams exist, the system can switch the streams at different times, depending on the user's request. In a way, the system performs dynamic semantic remixing for each user.

Current streaming systems are designed for fixed data streams. Some systems prepare the data stream at the receiver by combining multiple streams. Even in this case, the receiver renders a single stream. In dynamic switching systems, however, a user's actions or requests cause the system to switch the data stream—for example, from Camera 1's video to Camera 28's video. Because each user's requests are unique, the switching combination and timing differ for each user.

This dynamic switching ability will require new media-streaming techniques. This difficult real-time resource-management problem should consider quality-of-experience issues for different users.

EventWeb

The Web has revolutionized many aspects of human society in just over a decade. However, in the current

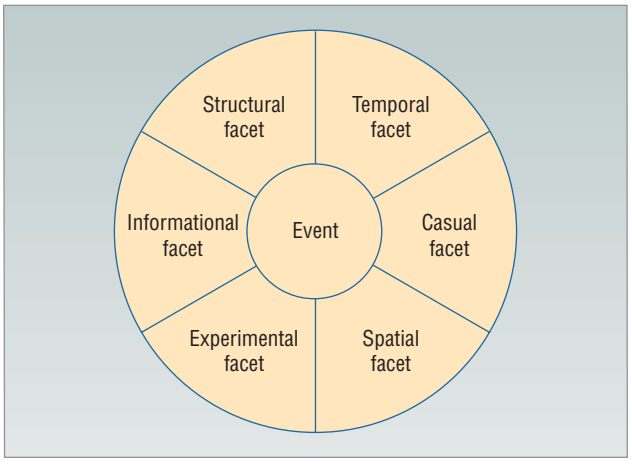

Figure 2. Different facets of an event. In EventWeb, each node (event) is represented by its basic properties—time, location, and type—as well as informational and experimental attributes.

Web each node is a document connected to other nodes using manually created referential links. Because of emerging digital-media devices and technology, we're now in a position to develop the EventWeb. Creating EventWeb will require developing technology to produce events using heterogeneous media elements, represent each event as a node, and then create explicit links between events and between events and information on the current Web.

EventWeb organizes data in terms of events and experiences and allows natural access from users' perspectives. For each event, EventWeb collects and organizes audio, visual, tactile, textual, and other data to provide people with an environment for experiencing the event from their perspective. EventWeb also easily reorganizes events to satisfy different viewpoints and naturally incorporates new data types—dynamic, temporal, and live. The current Web is document-centric hypertext. Unlike events, hypertext has no notion of time, space, or semantic structures other than often ad hoc hyperlinks.

Applications. EventWeb has significant applications, ranging from education to healthcare and from games to government. Events on this Web can offer multiple perspectives of important events, permitting remote personalized participation in live events such as meetings, lectures, concerts, and sports. Because users could conceivably take part in a meeting and change its course, they would find remote participation in a meeting different from remote participation in a concert or sports event.

Users would participate in a sporting or entertainment event only through observation from a particular perspective. In addition, they could archive these events to experience later, maybe from a different perspective each time. This could provide insights leading to valuable knowledge creation.

In EventWeb, each node (event) is represented by

- its basic properties—time, location, and type;
- informational attributes—participants and characteristics of the event; and
- experiential attributes—text or audio reports, photos, video, and other sensory information.

As Figure 2 shows, the structural and causal relationships among the events create the EventWeb. Event creators could specify these relationships explicitly, or different processes or people could discover them. Events captured in this Web could be as diverse as an application demands and at the desired granularity.

Dynamic structure. Anyone could produce an event node on EventWeb, whether it's the World Cup finals, a tsunami, a wedding, an accident, a bus arriving at a bus stop, a sale at a local store, my daughter's first music lesson, or me typing these words. You could use text descriptions, photos, audio, video, and haptic, smell, infrared, and other suitable sensor data to capture the important information and experiences related to the event.

EventWeb is fundamentally a dynamic Web structure that's linked to physical locations and uses familiar natural sensory characteristics. It uses text when needed. The EventWeb will link to the current Web, as content (for example, blogs and news sites) frequently describes events and experiences. Thus, EventWeb and the Web will work synergistically.

Realizing EventWeb will require technological advances in many areas. Some key areas are the same as for immersive telepresence. However, EventWeb also requires novel innovative concepts and tools from media processing, databases, Internet technologies, media creation and presentation, Web crawling, social networking, computer architecture, arts and architecture, and media search.

Archiving and indexing. People can use an Event Markup Language to post their events and related information and experiential data in the form of photos, audio, videos, and textual data. An EML will also provide an environment for expressing and creating relationships among events. Combining this language with event capture and a media-processing tool will help users identify events of interest in the EventWeb.

Events have an interesting life cycle. They're planned, they take place, then people store their experiences of the event in the form of experiential data and relate them to past and future events. Interestingly, past events play a more important role in our lives than current and future events.

For example, all sciences rely on the analysis of past events. Most paintings, novels, movies, and news reports are related to past experiences. An environment for capturing, archiving, and indexing events is therefore essential, as is the facility to continuously add new information and experiences and links to these events.

When billions of events (past, current, and future) exist, how do we find events of interest? We should somehow aggregate and develop an environment that lets users discover events of interest. We also need powerful and efficient indexing approaches for accessing all of these events. Current search-engine and relational-database techniques might not allow for event indexing. Instead, event indexing might require a combination of multidimensional and inverted file approaches.

Exploration and presentation environments. The keyword box on the current Web is practically unusable in EventWeb. We'll need novel approaches that combine navigation and search environments for finding appropriate events. Moreover, because photos and video are better than text for capturing and representing events, we'll need a novel presentation environment. Such a presentation environment will require a unique combination of ideas from visual arts and HCI to present event experiences to users.

Realizing immersive telepresence and the EventWeb will require advances in several research areas. Two areas in particular are essential to the development of such emerging applications and to HCC's advancement. These areas require a new perspective on current approaches and present some challenging issues. There's some research in these areas, but at best it's in its early infancy.

The concept of "event" can serve as the fundamental organizational principle for multimedia systems.

EVENT MODEL

Current information tools deal well with entities, objects, and keywords. However, information management in dynamic multimedia environments requires new concepts and techniques. Clearly, current concepts and tools are good for text-oriented and structured information systems that deal mostly with static information. But these tools aren't good for dealing with images, video, audio, and other sensory information. Consider, for example, major search engines' poor results for images and video. These search engines try to apply text-oriented search tools to the text associated with images and video, without processing images and video to extract meaningful indexing information from them, with surprisingly bad results.

Experiential systems

Current information tools evolved before the wave of mobile phones, digital cameras, and broadband systems changed the information system landscape. With all these advances, experiences are becoming an integral part of information systems. The recent flurry of activity in community-oriented systems such as YouTube, MySpace, and Facebook is a good example of experiential data's increasing popularity.

This data's popularity will likely increase even more rapidly in the developing world because of the abundance of mobile phones and the number of non-English-speaking people. Experiential systems deal with sensory data in sensory space without linguistic abstraction to bypass language and keyboard issues.

Event concept

The concept of "event" can serve as the fundamental organizational principle for multimedia systems. Strong conceptual, engineering, computational, and human-centered design principles support the use of event as a primary structure for organizing and accessing dynamic multimedia systems.

The definition of event depends on context and granularity. Dictionary.com defines an event as "something that occurs in a certain place during a particular interval of time." The term also refers to a significant occurrence or happening, or a social gathering or activity.

Events indeed have different contexts. A wedding is an event, as are the wedding reception and the cake-cutting. The bride and groom's first meeting is as much an event as the bride's birth, her parents' wedding, and so on. And, yes, the World Cup soccer final between Italy and France was an event, and so is my grandson's first soccer kick in his backyard. Theoretically, even moving my finger to a specific key is an event. So, events depend on context and occur at different granularities or resolutions.

Capturing and combining events

As Figure 3 shows, you can combine events in many ways to define other (compound) events. And you can again combine these combinations of events with other events to define yet another set of events. So this process of defining events continues. An application clearly determines these definitions. On the other side, an event is the result of one or more past events, which were in turn results of other events, and so on. Similarly, an event might result—maybe in combination with other events—in multiple events, which in turn might result in many other events. So this process of event creation has been ongoing and will continue into the future.

But, if all of these things are events, how can we capture them in our computing systems—or can we? At first, this situation appears confusing, but objects are equally confusing. Objects could be physical or conceptual. Objects can also exist at many resolutions. So, I'm an object, and so are the shirt I'm wearing and the buttons on the shirt.

Defining event aspects

Object-oriented programming and object-oriented design concepts have dominated computer science for

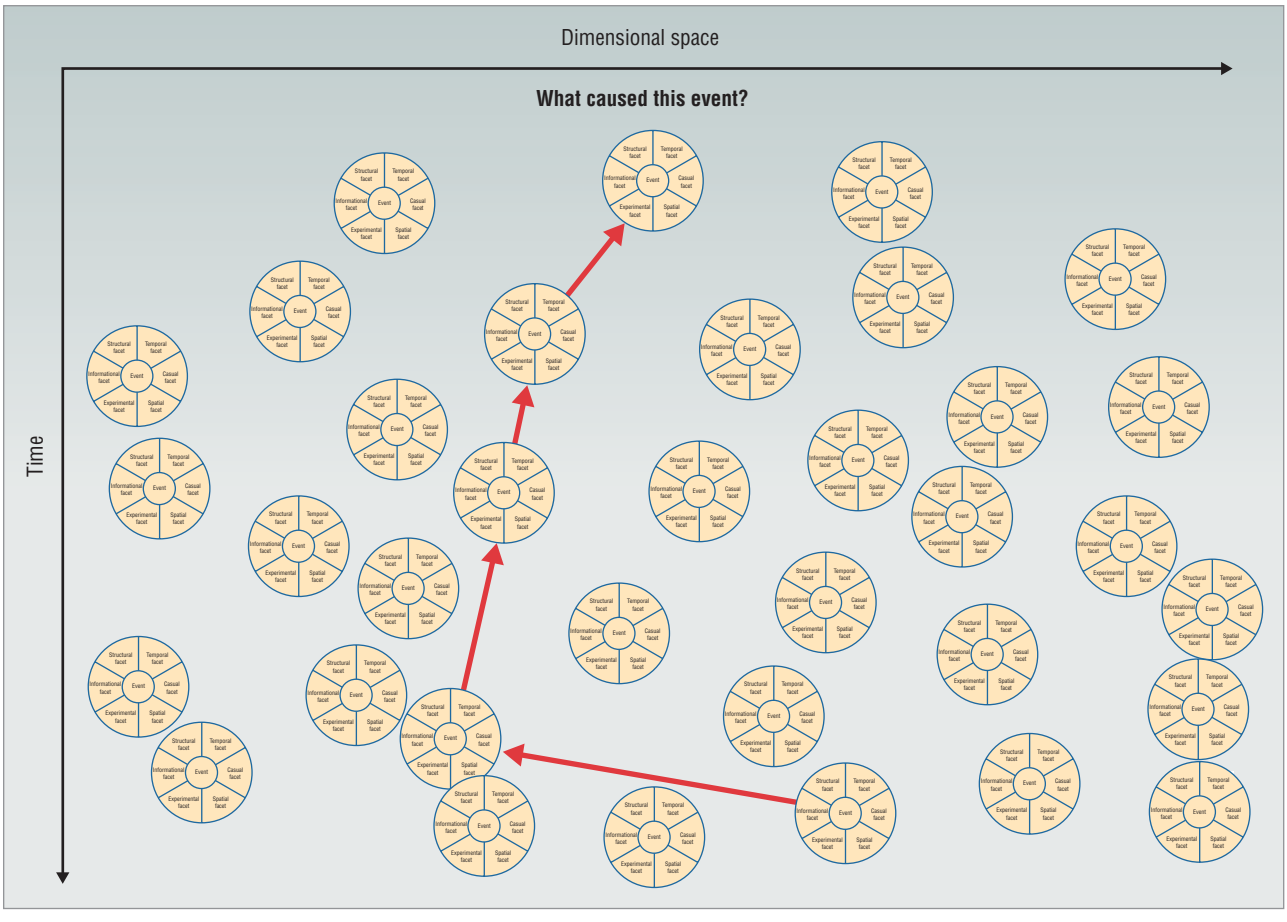

Figure 3. Time and space. We could consider a similar event occurring at a different time and space a different event.

more than a decade. This has been a powerful paradigm. In the context of computer science, "objects are a language mechanism for binding data with methods that operate on that data" (http://en.wikipedia.org/wiki/Object). This avoids any reference to physical, conceptual, or other types of objects that we use in our regular language, and it also provides an elegant new functional definition. Objects become a mechanism for binding data with methods that operate on that data.

Similarly, an event should explicitly define three important aspects:

- information about the event,
- experiences related to the event, and
- the event's structural and causal relationships with other events.

An event in computational form should represent data associated with these aspects as well as the processes necessary to acquire and present them. By providing flexible and expressive mechanisms to define these three components and associated methods, we could effectively define events. The event environment should provide tools for defining any event of interest from many disparate application domains. We could, for example, define event classes, and each event in the system could serve as an instance of a class.

An event's basic characteristics are its identification, time, and location, with the latter two becoming the event's fundamental defining characteristics. We could consider a similar event occurring at a different time and space a different event. In this sense, an event is defined in spatiotemporal space. Point events are just points in the spatiotemporal space, while interval events are regions in spatiotemporal space.

My team at the University of California, Irvine, is developing multiple applications using event models to validate our hypothesis that events can effectively capture multimedia semantics and help us build efficient systems to deal with multimedia information.[5,6]

EXPERIENTIAL ENVIRONMENTS

Assume we create a system containing events and all their related information and experiences. We'll then need an appropriate environment for interacting with such systems. The interaction environment should be human-centric and should work synergistically with human strengths and limitations.

Current information environments actually work against the human-machine synergy. Humans are effi-

cient in conceptual and perceptual analysis but relatively weak in mathematical and logical analysis, while computers are the opposite. In an experiential environment, users apply their senses to observe data and information of interest related to an event, and they interact naturally with the data based on their particular set of interests in the context of that event.

Support direct communication

People like to work in their natural environments using sensors. Unfamiliar metaphors and commands create confusion and make systems difficult to use. Similarly, a keyboard is an effective tool in some text-oriented situations, but most people would rather talk and gesture than type on a keyboard. An experiential environment should present a user with data that the human senses can easily and rapidly interpret. We should make user interactions with the data set as natural as possible. Gesture interfaces might play an important role here.

Same query and presentation spaces

WYSIWYG word processing and spreadsheets facilitated the personal computing revolution. WYSIWYG environments can provide quick feedback because they use the same space for queries and results. In fact, the notion of interface becomes just an environment for getting things done.

Different query and presentation spaces make most current information systems difficult to use. These spaces make users feel they're in a structured, rather than natural, environment. Popular search engines provide a box for entering keywords, and the system responds with a list of thousands of entries spanning hundreds of pages.

Most users never go beyond the first page. Contrast this to a spreadsheet, where users' actions result in a new sheet, showing new relationships. A WYSIWYG environment that merges query and presentation spaces would allow for easier interaction with spatiotemporal data.

User state and context

People feel comfortable in situations with static or gradually changing contexts and states. Computing systems should know the user's state and context and present information that's relevant to that state and context. Current information systems, including databases, were designed to provide scalability and efficiency, which are better achieved in stateless environments. This design was justified in the early days of computing, when computers were very expensive compared to human time. The situation has changed dramatically, however, and computing time is cheap compared to human time. So, we'll need to design systems that maintain context and state to maximize the efficiency and quality of users' experiences.

Perceptual analysis

Humans can use their perceptual facilities more rapidly and efficiently than machines, hence the popularity of visualization. Minimizing text-oriented displays and presenting information using appropriate sensory modalities can significantly improve performance and experience. Video's increasing popularity over any other medium is due to its ability to combine multiple modalities. Videogames and many simulation systems are so engaging because they provide a powerful visual environment, sound, and, in some cases, tactile inputs.

As speech recognition and computer vision become more sophisticated, we'll likely see more multimodal interfaces. Experiential environments must organize information as well as maintain user context and state. WYSIWYG applications are the first step in the direction of experiential environments.

It's no surprise that videogames are so successful. They're easy to learn and natural to use, and they provide a compelling experience by engaging our senses. General-computing-environment designers can learn a lot from the interaction environments in videogames.

How can we bring computing to the more than 5 billion of the world's 6 billion people it hasn't yet reached? Although many think selling a computer for less than US $100 would help achieve this, sound HCC practices would definitely help.[9]

Mobile phones—which are easy to use and outnumber computers by two to one—can bring the Internet and computing to the masses, even in remote parts of the world. Providing audiovisual-tactile interfaces in phones can help people create and access content. The iPhone and advanced phones like Nokia's N95 are more powerful than PCs were less than a decade ago, and they have built-in cameras, microphones, and other sensors. These devices could make content creation and access less dependent on language and education level than in current systems.

It's time to reduce our dependence on text by using all other data modalities and organizing and presenting information in ways that are more perceptual and cognitively meaningful to humans. By providing an environment for experiencing events to gain insights and share event experiences, technology will accelerate knowledge growth at an unprecedented rate. Moreover, the ability to effortlessly share event experiences with fellow humans will help identify those with similar interests and subsequently build communities across the globe.

References

1. P. Rander, "A Multi-Camera Method for 3D Digitization of Dynamic, Real-World Events," tech. report CMU-RI-TR-98-

12, Carnegie Mellon Univ., 1998; www.ri.cmu.edu/pub_files/pub1/rander_peter_1998_1/rander_peter_1998_1.pdf.
2. R. Jain, "Experiential Computing," *Comm. ACM*, July 2003, pp. 48-55.
3. S. Vedula et al., "Three-Dimensional Scene Flow," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Mar. 2005, pp. 475-480.
4. H. Rheingold, *Virtual Reality*, Simon & Schuster, 1992.
5. R. Jain, "Out of the Box Data Engineering: Events in Heterogeneous Data," *Proc. Int'l Conf. Data Eng.* (ICDE), CS Press, Mar. 2003, p. 8.
6. G.U. Westermann and R. Jain, "Events in Multimedia Electronic Chronicles (E-Chronicles)," *Int'l J. Semantic Web and Information Systems*, vol. 2, no. 2, 2006, pp. 1-27.
7. G. Utz Westermann and R. Jain, "Toward a Common Event Model for Multimedia Applications," *IEEE MultiMedia*, Jan. 2007, pp. 19-29.
8. A. Katkere et al., "Towards Video-Based Immersive Environments," special issue on multimedia and multisensory virtual worlds, *Multimedia Systems*, Mar. 1997, pp. 69-85.
9. R. Jain, "Folk Computing," *Comm. ACM*, Apr. 2003, pp. 27-29.

Ramesh Jain *is the Donald Bren Professor in Information and Computer Sciences in the Department of Computer Science at the University of California, Irvine. His research interests are multimedia information systems, image databases, and machine vision. Jain received a PhD in computer science from the Indian Institute of Technology. He is a Fellow of the IEEE and the ACM. Contact him at jain@ics.uci.edu.*

SPECTS 2008

2008 International Symposium on Performance Evaluation of Computer and Telecommunication Systems

http://atc.udg.edu/SPECTS2008/

June 16-18, 2008
Edinburgh, UK

This annual international conference is a forum for professionals involved in the performance evaluation of computer and telecommunication systems. Performance evaluation of computer systems and networks has progressed rapidly in the past decade and has begun to approach maturity. Significant progress has been made in analytic modeling, simulation and measurement approaches for performance evaluation of computer and telecommunication systems. Topics of interest include, but are not limited to:

Networking and Telecommunication Systems

- Internet Technology
  Quality of Service (QoS)
  DiffServ/IntServ
  MPLS/GMPLS
  TCP
  World Wide Web (WWW) Technology
- Networking Techniques
  Unicast and Multicast Routing
  Congestion Admission and Control
  Switching Techniques
  Tele-traffic
  Network Protocols
  Network Management and Control
  Network Capacity Planning
  Network Architecture Evaluation
  Service and QoS Pricing
  Security and Authentication
- Broadband Networks
  High-Speed Networking
  Optical Networks
- Wireless Systems and Networks
  Satellite Systems
  UMTS
  Mobile Networks/Computing
  Ad-hoc Networks
  Sensor Networks
- Multimedia Communications and Applications

Computer Systems

- Distributed Architecture
  Client/Server
  Distributed Systems and Agents
  Parallel and Distributed Computing
  Massively Parallel Systems
  Cluster Computing
  Grid computing
  Interconnection Networks
- Computer Architectures
  Microprocessors/Microcomputers
  Memory Systems
  High Performance I/O
  Real Time Systems
  Scheduling Schemes
- Software
  Software Performance, Evaluation, and Testing
  Parallel Algorithms and Languages
- Electronic Commerce
- Hardware and Software Monitors
- High Performance computing
- Information Assurance
- Reconfigurable Computing
- Scientific Computing Algorithms
- Workload and Traffic Characterization

Tools, Methodologies, and Applications

- Parallel and Distributed Simulation
- Verification and Validation
- Neural Networks and Fuzzy Logic Applications
- Performance Optimization, Bounds, and Models
- Queuing Systems and Networks
- Scalability Studies
- Integrated Modeling and Measurement
- On-Line Performance Adaptation and Tuning
- Process Algebra Based Models
- Mathematical Aspects and Integrated Design of Performance
- Case Studies

Paper Submission

Please submit your complete papers electronically to: http://www.softconf.com/scs/SPECTS08/ Instructions for authors will be posted on the paper submission web site. Submissions should not exceed 8 two-column, 8.5x11 inch pages (including figures, tables, and references) in 10 point fonts. Please include 5-10 keywords, complete postal and e-mail address, and fax and phone numbers of the corresponding author. If you have difficulties with electronic submission, please contact Technical Program Co-Chairs or the Local Arrangement Chair.

Extended Versions of Selected Accepted Papers in SPECTS 2008 will be considered for Possible Publication in Scholarly Journals.

Proposals for 1.5 hour tutorials and special sessions should be sent to Tutorials and Special Sessions Chair. Tutorial abstracts along with keynote speeches' abstracts will be included in the proceedings of the conference.

Sponsored by

The Society for Modeling and Simulation International, SCS

Technically Co-Sponsored by

IEEE Systems, Man, and Cybernetics (SMC) Society

IEEE Communications Society

Deadlines

Special sessions proposals	December 7, 2007
Submission of papers and tutorials proposals	February 25, 2008
Notification of acceptance	April 21, 2008
Submission of camera-ready papers	May 5, 2008

Call for Papers

General Chair
Mohammad S. Obaidat, Monmouth Univ., NJ, USA
obaidat@monmouth.edu

Program Chairs
- Jose Marzo, Univ. of Girona, Spain
  joseluis.marzo@udg.es
- Helena Szczerbicka, Univ. of Hannover, Germany
  hsz@sim.uni-hannover.de

Program Vice Chairs
- Pawel Gburzynski, University of Alberta, Canada
  pawel@cs.ualberta.ca
- José Luis Sevillano, University of Seville, Spain
  sevi@atc.us.es

Tutorials and Special Sessions Chair
S Dharmaraja, Indian Institute of Technology, Delhi, India
dharmar@maths.iitd.ernet.in

Awards Chair
Franco Davoli, Univ. of Genoa, Italy
franco@dist.unige.it

Publicity Chairs
- Essia Elhafsi (Chair), Univ. of California-Riverside, USA
  essia@cs.ucr.edu
- Jong Hyuk Park, Kyungnam Univ., Korea
  parkjonghyuk@gmail.com
- Farid Naït-Abdesselam, Univ. of Sciences & Technologies, France
  nait@lifl.fr
- Abdelmajid Khelil, Technical Univ. of Darmstadt, Germany
  khelil@informatik.tu-darmstadt.de

Local Arrangement Chair
Graham Shanks, Baesystems, UK
Graham.Shanks@baesystems.com

Web Masters
- Antonio Bueno, Univ. of Girona, Spain
  antonio.bueno@udg.edu
- Michael J. Chinni, US Army TACOM-ARDEC
  mchinni@pica.army.mil

Program Committee
Abdullah Abonamah, Zayed Univ., UAE
Tricha Anjali, Illinois Institute of Technology, USA
Noureddine Boudriga, Univ. of Tunis, Tunisia
Rafael F. S. Caldeirinha, Escola Superior de Tecnologia e Gestao de Leiria, Portugal
Maria C. Calzarossa, Univ. of Pavia, Italy
Fernando Cerdan, Technical Univ. of Cartagena, Spain
Tomaso de Cola, Univ. of Genoa, Italy
Floriano De Rango, D.E.I.S. Dept., Univ. of Calabria, Italy
S. Dharmaraja, Indian Institute of Technology, Delhi, India
Rachid El Abdouni Khayari, Univ. of the Armed Forces Munich, Germany
Matthew Faulkner, Lancaster Univ., UK
Geoffrey Fox, Indiana Univ., USA
John Fox, Foxband Consultancy, UK
Sebastia Galmes, Universitat de les Illes Balears, Spain
Daniel Garcia, Univ. of Oviedo, Spain
Pawel Gburzynski, Univ. of Alberta, Canada
Jarmo Harju, Tampere Univ. of Technology, Finland
Xavier Hesselbach-Serra, Univ. Politècnica de Catalunya, Barcelona, Spain
R. Jayaparvathy, Coimbatore Institute of Technology, India
Carlos Juiz, Univ. de les Illes Balears, Spain
Krishna Kant, Intel, USA
Helen Karatza, Aristotle Univ. of Thessalonici, Greece
Michalis E. Kounavis, Intel Research, USA
Veronica Lagrange M. Reis, HP Corp., USA
S. Kami Makki, Univ. of Toledo, USA
Petteri Mannersalo, VTT Technical Research Centre of Finland
Mario Marchese, Univ. of Genoa, Italy
Pascale Minet, INRIA, France
Ibrahim Onyuksel, Northern Illinois Univ., USA
Mohamed Ould-Khaoua, Univ. of Glasgow, UK
Elena Pagani, Univ. di Milano, Italy
Georgios I. Papadimitriou, Aristotle Univ., Greece
Krzysztof Pawlikowski, Univ. of Canterbury, New Zealand
Stephen Pink, Lancaster Univ., UK
Michal Pioro, Warsaw Univ. of Technology, Poland
Desanka Polajnar, Univ. of Northern British Columbia, Canada
Jernej Polajnar, Univ. of Northern British Columbia, Canada
Ramon Puigjaner, Univ. de les Illes Balears, Spain
Vicente Santonja, Technical Univ. of Valencia, Spain
Barbara Sorensen, Air Force Research Laboratory, USA
Tatsuya Suda, Univ. of California, Irvine, USA
Phuoc Tran-Gia, Univ. of Wuerzburg, Germany
Pere Vila Talleda, Univ. of Girona, Spain
Manuel Villen-Altamirano, Telefonica, Spain
Bernd E. Wolfinger, Hamburg Univ., Germany

COVER FEATURE

Secure and Easy Internet Voting

Giampiero E.G. Beroggi
Statistical Office, Canton Zurich

A Swiss e-voting system, operational since December 2004, is based on a service-oriented architecture that lets voters use Internet or mobile phones to cast votes. Two-step encryption and redundant storage systems keep votes authentic and confidential.

Although modern societies rely heavily on information and communication technology for business, work, and leisure time activities, they have thus far seemed hesitant to use ICT for democratic decision-making activities such as voting. Meanwhile, the lost and uncounted votes associated with current paper ballots could very well be contributing to biased political decisions.[1]

One reason for the delay in implementing more technologically sophisticated voting methods is the computer science community's almost unanimous wariness of Internet-based elections.[2] Many governments have simply dismissed e-voting as too risky. Others are not fully aware of e-voting's strong advantages over paper ballots: reliable and secure vote casting, precise vote counting, the option to conduct voting in a centralized and decentralized manner, and the rapid availability of results.

Fortunately, in light of these strong advantages, more countries are beginning to consider e-voting systems,[3] but most efforts are still in the conceptual or testing stage. In contrast, three cantons in Switzerland—Zurich, Geneva, and Neuchatel—are already using an e-voting system. The Zurich e-voting system (https://evoting.zh.ch), which has been in operation since December 2004, features a modular and service-oriented architecture that lets voters cast their votes through a range of digital media, including computers and mobile telephones (currently) and interactive TV and personal digital assistants (planned).

The system easily integrates into existing software solutions without loss of security and accommodates either centralized or decentralized operation. Both national and local authorities have embraced the system, particularly its smooth integration with traditional ballot-box voting. Offering both e-voting and paper ballots means that all citizens, regardless of their technology awareness, can vote, and there is no fear of a digital divide among the population. Zurich Minister of the Interior Markus Notter pronounced the system a "milestone in Swiss democracy [that] opens the ballot to today's information society."

The "Chronology of the Zurich E-Voting System" sidebar describes key points in developing and implementing the e-voting system. Annual operational costs are $400,000, which translates to approximately $0.50 per e-vote. Since the testing phase concluded in April 2006, three communities in Canton Zurich have started using the system. More would like to use it, but the Swiss government has mandated that only 10 percent of the electorate can use e-voting. As soon as the government lifts that restriction, however, all 171 Canton Zurich communities could begin using the system, thanks to the scalability of its service-oriented structure.

HOW THE SYSTEM WORKS

The Zurich e-voting system covers national votes on referenda, votes on citizen initiatives with counter referendum and contingency plans, majority elections, and proportional elections with predefined party lists.

Figure 1 illustrates the e-voting process. Six weeks before e-voting day, communities within Canton Zurich enter in the electronic ballot box the names of citizens

Published by the IEEE Computer Society

eligible to vote electronically. The electronic ballot box opens two weeks later. To vote, citizens use a special password that Canton Zurich's Statistical Office has mailed to them as part of their voting forms. E-voting then takes place during the next four weeks.

At present, voters can choose between using the Internet and mobile phones to cast their e-votes; other digital alternatives, such as interactive TV and personal digital assistant/Wireless Access Protocol (PDA/WAP) are technologically feasible but not yet active.

Voting process

To vote through the Internet, voters log onto the e-voting website using their identification numbers and follow the site's instructions for vote casting. Figure 2 gives a sample screen from the simulation software. After casting their votes, voters enter a personal identification number (PIN) and compare a security symbol with the one they received in the mail. If the two match, the system accepts the vote.

Two-step encryption protects voter confidentiality. The voter's client computer first encrypts the votes and identification and authentication characteristics, and the e-voting system then checks the incoming votes for their structure and integrity before once again encrypting them. Two redundant subsystems then store the cast votes in a database.

To vote through a mobile phone, voters enter codes to a dedicated phone number using the short message system (SMS). Citizens enter codes for personal identification (g3387y55, for example), the name of the referendum (such as sg1), and the actual yes or no vote (er2 for yes, for example). The SMS message for a user voting yes on referendum sg1 would thus be g3387y55 sg1 er2. The system replies by asking the voter to enter a PIN (separate from the identification access code) and birth date (such as 14031968 for 14 March 1968) in a second SMS message. The citizen receives then a confirmation that the e-vote was entered in the e-voting ballot box.

On voting day, the communities enter the results from the regular ballot box into the vote registration software. As soon as the regular voting ballot box is closed, the e-voting system transfers the e-votes to the computer system that handles the regular votes. An overview of the total results—regular votes and e-votes—is available immediately.

Vote transmission

Because the e-voting system is based on the IT Infrastructure Library, it can accommodate a range of formats—the Extensible Markup Language (XML), the Electronic Markup Language (EML), open database connectivity, the comma-separated value format, and the Simple Object Access Protocol (SOAP)—as well as direct database access. To meet a Swiss government requirement, all formats are convertible to EML for

Chronology of the Zurich E-Voting System

With approximately 1.2 million people, Zurich has the largest population of the 26 Swiss cantons. The Statistical Office of the Canton Zurich (www.statistik.zh.ch), which belongs to the Ministry of Justice and Interior, is responsible for planning and conducting federal and local elections and referenda. As part of its responsibility, the office must provide the technological means for citizens and local authorities to conduct and participate in elections and referenda.

In 2001, the office introduced a fully computerized election and referendum system that connected all 171 communities within the canton, allowing real-time progress monitoring and community assistance on voting days. The e-voting pilot project began in 2003 and successfully completed in spring 2006. The total project cost was $3.7 million—$1.9 million for planning and $1.8 million for implementation.

- **February 1998:** Swiss government defines as part of its ICT strategy the need to test the use of ICT for democratic decision-making processes.
- **August 2000:** Swiss government mandates Federal Chancellery to study the feasibility of e-voting.
- **June 2002:** Swiss Parliament creates legal basis for e-voting pilot study.
- **February 2002:** Federal Chancellery signs contract with Ministry of the Interior of Canton Zurich to participate in the e-voting pilot study.
- **October 2003:** Unisys wins the bid to design the Zurich e-voting system and starts development.
- **December 2004:** First e-voting in Canton Zurich through Internet and mobile phone to elect 70 student board members at the University of Zurich. Voting participation was 93 percent; of the 1,767 people participating in the election, 1,582 used the Internet and 205 used mobile phones. Only one person used the traditional ballot box.
- **October 2005:** First e-voting election in the city of Bulach with 37 percent participating in e-voting.
- **November 2005:** First e-voting for federal and regional offices in three communities. E-voting participation was 37 percent.
- **April 2006:** First e-voting through Internet and mobile phone for proportional election system. E-voting participation was 20 percent.
- **July 2006:** End of pilot project and start of e-voting for any upcoming elections and referenda. Currently, Canton Zurich is waiting for the government to lift its restriction so that all communities can use e-voting.

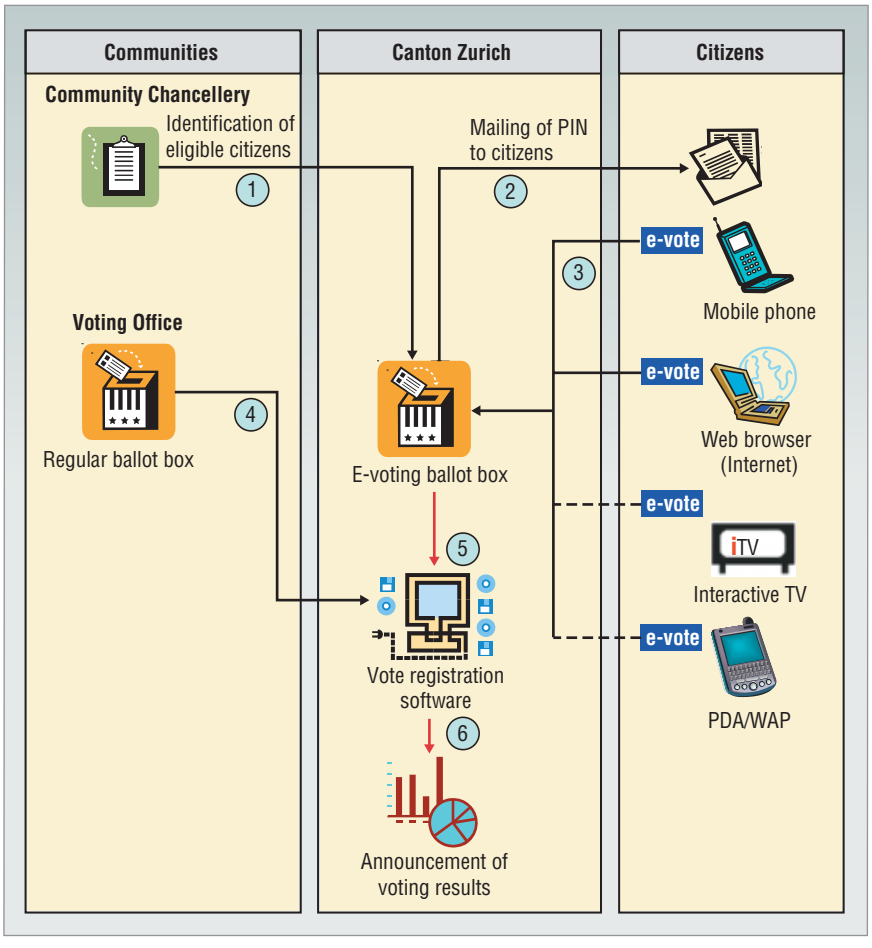

Figure 1. E-voting in Canton Zurich. The e-voting process has six main steps: (1) Communities send a list of those eligible to vote electronically. (2) Voters receive a list of system codes for identifying themselves and the referenda or candidates they are voting for, as well as codes for entering "yes" and "no" responses. (3) Voters cast their vote using their preferred digital medium. At present only Internet and mobile phone options are active, although the system is designed to handle iTV and PDA/WAP as well. The e-voting ballot box closes 24 hours before the regular ballot box. (4) The communities send the paper ballot results to the vote registration software. (5) Finally, Canton Zurich's Statistical Office counts all votes (electronic and paper) and (6) produces the final vote count.

import. Each community and organization has field mapping and tracing options at all user levels. Swiss standards for e-government dictate how data fields and records are integrated.

To enable voting on a mobile phone, two of the three mobile phone companies in Switzerland use a virtual private network (VPN) communications network to link directly into the e-voting system. The third company uses an IP VPN communications network to link into the Canton Zurich secure network (LeuNet), which in turn links directly to the e-voting system.

SYSTEM DEVELOPMENT

The e-voting system had to ensure voting rights and secrecy, capture votes accurately, and prevent abuses such as multiple votes from the same individual. The Swiss government was adamant that any alternative to traditional ballot-box voting not compromise the Swiss Federal Law for Political Rights, which protects the fundamental right of citizens to express their free political will without any technological, psychological, or other restriction or bias. The government also wanted an e-voting system that would encourage more citizens to participate in public-policy decision making. Finally, e-voting had to meet the same high security standards as traditional voting approaches.

To meet these requirements, the overarching design goals were to provide more flexibility and security without additional restrictions or controls and to offer a superior service for citizens and communities responsible for elections and referenda. To meet the superior service goal, all current IT systems had to integrate within the e-voting system. The aim was to require only minor changes to the communities' election and referendum processes.

Another requirement was the ability to operate in the decentralized manner of the Swiss voting structure. Thus, each of the 171 communities within Canton Zurich had to be able to manage its own voting register. The e-voting system also had to account for features of the Swiss elections and referenda rules. For majority elections, this could involve a predefined list of candidates or the entire citizenry. The system had to allow each community to define when the electronic ballot boxes would be open. The election officers would receive the decoding keys with all the passwords to decode the votes on voting day.

Finally, to prevent citizens from abuses such as casting their vote multiple times, the system had to have several safety features, which either the communities could activate individually or the canton could activate centrally.

Testing

The e-voting system had to undergo scientific test monitoring as well as technological testing. The three cantons chosen for these tests, including Canton Zurich, signed contracts with the Federal Chancellery to adhere to four rules during testing:

- No one can intercept, change, or reroute electronically cast votes.
- No third party can obtain knowledge of the cast vote.
- Only registered citizens can vote.
- Every registered person can vote only once.

Testing used an algorithm developed to simulate vote casting, vote counting, and results reporting. The aim of these tests was to reveal gaps that might not be detectable during regular applications. From 2004 to 2006, the e-voting system was tested during real elections and referenda. Swisscom Solutions, Switzerland's leading telecommunications company, conducted the system and internal security audits. The Federal Chancellery also conducted a separate security audit and suggested changes in the architecture, user interface, and password structure. Designers considered these suggestions in improving the e-voting system during the testing phase.

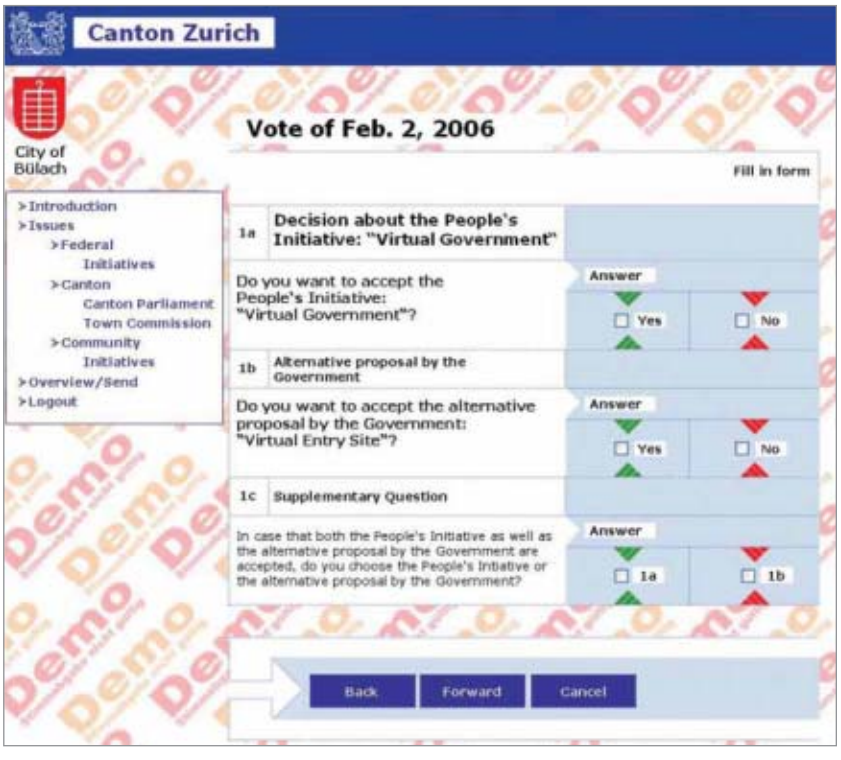

Figure 2. Sample e-voting screen. A menu indicates which issues are up for e-voting (left). In this case, 1a is the referendum, 1b is the alternative government proposal, and 1c is the supplementary question. Voters click on yes or no three times and then click forward to go to the next screen.

Security

The e-voting system's security requirements are based on the Information Security Management System (BS 7799). Both the Swiss government and the Federal Chancellery assess security annually. External parties perform security audits, one of which involves attempting to hack into the e-voting system (so far, all attempts have failed). The hardware and its physical security environment are in compliance with the US Department of Defense's protection class B2 or lower. The security concept definition complies with both ISO/IEC 17799 and BS 7799.

Data exchange between the communities and the e-voting system is based on the Secure Data Exchange Platform (SeDAP), which is based on the Online Services Computer Interface (OSCI) standard, which in turn is based on SOAP. All entries into the e-voting system—voter identification and authentication as well as voter rights—occur through a secure entry server, which ensures that only registered voters can vote.

Both the citizens' votes through the Internet and the files containing the names of citizens eligible to vote are transmitted using the Secure Sockets Layer (SSL) protocol. The confidentiality of voter access codes and passwords is of utmost importance, so Canton Zurich uses three independent companies to print these. After the system identifies the access codes and the voters cast their votes, the system immediately asks them to validate their vote by entering their birth date and a six-digit numerical identification code. The system accepts their votes only after validation.

Encryption occurs in two steps. The voter's client computer first encrypts the votes and identification and authentication characteristics through an SSL channel (1,024-bit encryption). The e-voting system then checks the incoming votes for their structure and integrity before once again encrypting them (1,024-bit encryption) and passing them to the high-security zone (second firewall). Two redundant systems store the votes on a write-once, read-multiple-times database.

For every election and referenda event, Canton Zurich's Statistical Office uses a virtual community to cast votes and then checks that the e-system properly recorded them. It also analyzes the citizens' votes, making sure that the sum of the validated codes during e-voting equals the sum of received electronic votes. These two plausibility checks must match perfectly—have zero tolerance—for the e-voting to be trustworthy. The separate encryption and storage of cast votes and names of citizens eligible to vote ensures that vote counts are accurate and keeps voting rights from being corrupted.

The literature on e-voting emphasizes the danger of making source code available as a way to build trust in the system,[4] since attackers with such access could

modify voting and auditing records.[5] For these reasons, the Zurich e-voting system does not make source code available. Rather, it relies on the ACM Statement on Voting Systems,[6] which recommends that e-voting systems "embody careful engineering, strong safeguards, and rigorous testing in both design and operation." The Federal Chancellery supervises the decoding of e-votes, which takes place only after physical balloting closes. As a further precaution, the e-voting hardware itself is in a steel cage with physical access control mechanisms such as fingerprint identification and appropriate safety precautions, such as fire detection and break-in alert.

The ACM statement also recommends that each voter be able to inspect a physical record to verify the accuracy of that vote. Obviously, e-voting does not lend itself to a reproducible recording of each voter's actions, but the codes provide an audit trail of sorts. This trail is still subject to attack and will never fully replace the physical trail, but a paper trail is equally dangerous in that it provides a visible receipt. Such a receipt could subject voters to bribery from those seeking to sell or buy votes.

Scalability and portability

Because of its service-oriented architecture and modularity, the e-voting system is fully scalable and portable. Cantons can define any number of voting districts, and communities can define their own electorate districts, entering district-specific data and information. Because the e-voting procedure is based on EML, any additional voting device will integrate with the e-voting system. Because voting transfer is independent of the user interface, users can integrate new applications and input devices quite easily. Thus, it is possible to analyze voting results independently of the media used to cast the vote.

ADOPTION RESULTS

Perhaps the main contributor to the e-voting system's favorable reception is its modularity and service-oriented architecture. Both national and local authorities have embraced the system because of its extreme flexibility in accommodating both centralized and decentralized operation and the full range of voting concepts, as well as its ability to integrate into existing infrastructure without compromising system security.

Adopting the e-voting system has already heightened voter participation. In response to the high participation in e-voting during the system's testing phase, the board of the University of Zurich decided to abolish traditional ballot-box voting. Consequently, the 2006 student board elections were, for the first time, based solely on e-voting. The result was higher efficiency and lower cost with no compromise in the approximately 24,000 students expressing their political preferences.

The Swiss ICT Society awarded the Zurich e-voting system the prize for Best Software in 2005, citing "its flexible compliance with complex elections and referenda concepts, its modular structure allowing for extension, and its remarkably high security standard." In 2007, the system won the 2007 United Nations Public Service Award for "fostering participation in policy-making decisions through innovative mechanisms." These awards, as well as lessons from the testing phase and first year of general use, are evidence that the e-voting system will successfully handle all Canton Zurich's 171 communities as well as port to other cantons or to any organization desiring to enjoy e-voting's compelling advantages.

References

1. A.S. Belenky and R.C. Larson, "To Queue or Not to Queue?," *OR/MS Today*, June 2006, pp. 30-34.
2. D.W. Jones and P.G. Neumann, "Interview: A Conversation with Douglas W. Jones and Peter G. Neumann," *ACM Queue*, Nov. 2006, pp. 16-23.
3. R. Krimmer, ed., *Proc. 2nd Int'l Workshop Electronic Voting*, Gesellschaft für Informatik, Bonn, Köllen Druck+Verlag GmbH, Bonn, 2006 (in German).
4. J. Kitcat, "Source Availability and e-Voting: An Advocate Recants," *Comm. ACM*, Oct. 2004, pp. 65-67.
5. B. Simons, "Electronic Voting Systems: The Good, the Bad, and the Stupid," *ACM Queue*, Oct. 2004, pp. 20-26.
6. J. Grove, "ACM Statement of Voting Systems," *Comm. ACM*, Oct. 2004, pp. 69-70.

Giampiero E.G. Beroggi *is director of the Statistical Office of Canton Zurich. He is also a professor at the Zurich School of Business Administration, where he specializes in decision support systems. Beroggi received a PhD from Rensselaer Polytechnic Institute. He is a senior member of the IEEE and a member of the IEEE Computer Society. Contact him at giampiero@beroggi.net.*

COVER FEATURE

Turning Teenagers into Stores

Srijith K. Nair, Bruno Crispo, and Andrew S. Tanenbaum, Vrije Universiteit
Ron Gerrits, Inovia

Paradiso is a prototype of a system that lets consumers contact content providers to buy songs and videos—and to buy optional content-resale rights. In essence, the scheme would turn customers into content distributors, provide wider reach, and free up content providers' bandwidth. However, such an architecture requires strict security precautions and interoperable digital rights management standards among player manufacturers and content providers.

The intersection of computers, the Internet, music, and teenagers has uprooted the music industry. For more than a century after Thomas Edison's 1877 invention of the phonograph, the industry sold singles or albums recorded on wax or plastic media to consumers in record and department stores. New media were introduced, including the 78-rpm wax single, the 45-rpm single, the 33-rpm LP record, and finally the Philips CD, but the business model stayed the same. With the invention and 1991 standardization of the MP3 psychoacoustic compression algorithm by engineers working on the European Union's Eureka project 147, the era of downloadable digital music was launched.

After the Fraunhofer Institute released the first MP3 encoder in 1994, many young music fans began to encode their audio CDs in MP3 format and store them on their computers' hard disks. Before the invention of MP3, storing music on a PC's hard disk wasn't practical because a single CD could take up to 650 Mbytes, and hard disks were smaller than 1 Gbyte at this time. But with tenfold compression possible with little quality loss, storing and playing music on computers skyrocketed.

It didn't take long before friends began exchanging music files over the Internet. Napster debuted in 1999, offering a central catalog of who had which songs, so people could directly copy songs from the remote hard disks of people they didn't know. Napster users thought of it as a wonderful new invention: peer-to-peer file sharing.

Unfortunately for them, people in the music industry didn't see it that way. They saw it as theft of their intellectual property, and they responded by suing Napster and closing it down. Decentralized services such as Kazaa and Grokster soon replaced Napster, and they were sued with mixed results. Then the music industry began suing individual teenagers for copyright violation, seeking maximum publicity when they settled out of court for thousands of dollars.

Eventually, it dawned on them that suing their own customers (especially children) wasn't a good business model. This led to the development of online music stores that let customers legally buy and download songs from the store's central server. The first major online-music seller was Apple with its hugely successful iTunes store (www.apple.com/itunes) and iPod player.

iTunes uses a completely centralized digital rights management (DRM) system called FairPlay, with users contacting an Apple server to buy and download music and authorize their usage. When Microsoft released its Zune player (www.zune.net) and online store in November 2006, it added a new feature lacking in iTunes: a limited ability for a user to transmit a song to a friend's Zune player offline, without having to contact the central Zune server. However, a user can only transmit a

Published by the IEEE Computer Society

song three times and store it for three days. If the friend likes the song, he must contact Zune's server to buy it. Figure 1 shows the two models.

By now, the music companies have come to realize that digital music is their friend (just as the movie studios eventually stopped suing VCR manufacturers and began releasing movies for rent). They also realize that many teenagers become aware of songs when they plug into a friend's music player and listen to music that way, a practice now known as "jack sharing." This knowledge has led some music executives to dream of turning teenagers into stores, legally reselling songs they've bought, a concept more prosaically called *superdistribution*.[1]

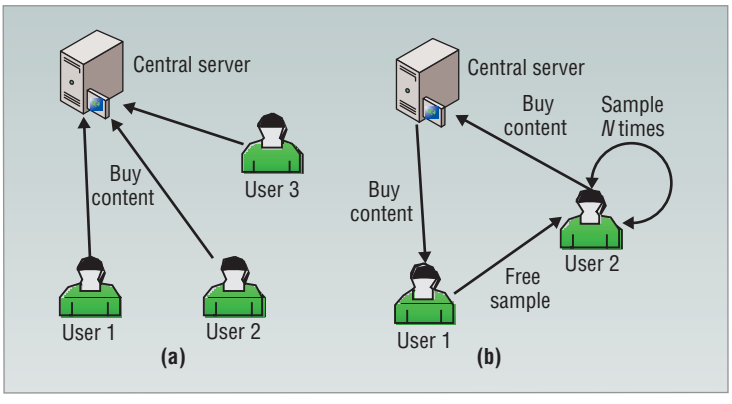

Figure 1. Music distribution models. (a) Apple uses a central store model to sell its iTunes. (b) Microsoft's Zune allows limited sharing of content.

A MODEL FOR RESELLING MUSIC

What we need is a scheme that turns willing customers into full-fledged resellers. Amsterdam's Vrije Universiteit has developed a system that could serve as a prototype.

Consider this scenario: Bob visits an online content provider like iTunes and buys a song for 99 cents. Having an inkling that the song will also be a hit with his friends, he buys the right to resell the song to nine friends for a total of $8.91, getting a 10 percent discount for buying 10 units. He pays the $8.91 in advance by credit card. Bob then hooks up with his friend Mark and tells him about the cool song he just got. After hearing the song using Bob's player, Mark decides to buy a copy. Bob sells Mark the copy (using the wireless link) for 95 cents, making a 6-cent profit.

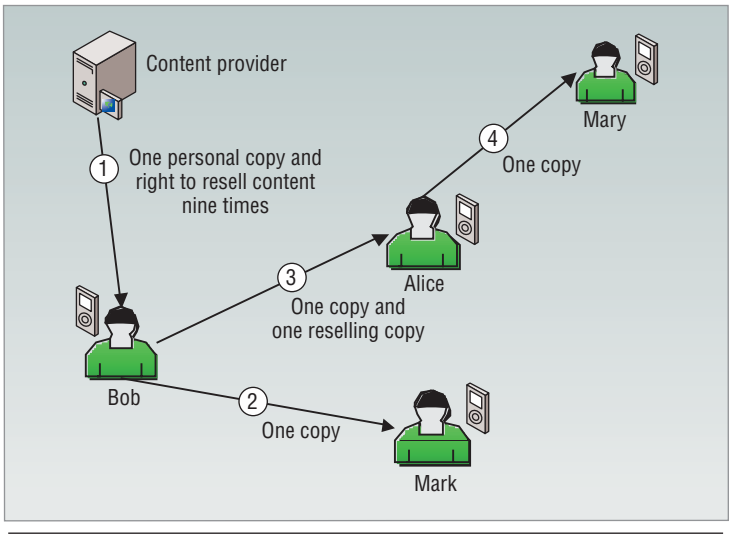

Figure 2. Content buying and reselling process flow.

On his way home, Bob meets up with Alice and sells a copy of the song to her. Alice tells Bob that her friend Mary might also be interested in the song, so she buys it and the right to resell it once, paying Bob $1.90. When Alice runs into Mary, she sells Mary the song for 97 cents. Figure 2 illustrates these transactions.

From the point of view of consumers like Bob and Alice, the benefit is evident. By acting as a reseller on behalf of the content owner, the consumer earns a profit per song sold. Mark and Mary also benefit by getting the song immediately and for less than the retail price. The content owner can reach a wider range of potential consumers, and, more importantly, the network formed between the consumers is more taste-targeted than any marketing campaign. Furthermore, by transacting many of the sales without involving the central distribution server, the content vendor's server and bandwidth requirements are greatly reduced, cutting associated costs.

Of course, the trick is to have the technology to enable the above scenario in such a way that content doesn't become freely available.

THE TECHNOLOGY

Achieving interoperability between different manufacturers' players and providers' content requires IEEE or other standards. Player manufacturers would need to design according to these specifications and undergo compliance testing and accreditation from a certification authority. Upon passing the accreditation tests, the CA would certify the manufacturer by signing its public key. The manufacturer would in turn sign the unique public key of each player it produces, thus producing a chain of trust to identify all specification-compliant players. Each player would contain its own certified public and private keys, the manufacturer's certified public key, and the CA's public key.

The manufacturer needs to store the player's private key in secure hardware to prevent direct unmediated access. All private-key operations must be performed within the secure hardware in a controlled manner. Among other functions, the secure hardware must perform asymmetric and symmetric key encryption and decryption and collision-resistant hashing.

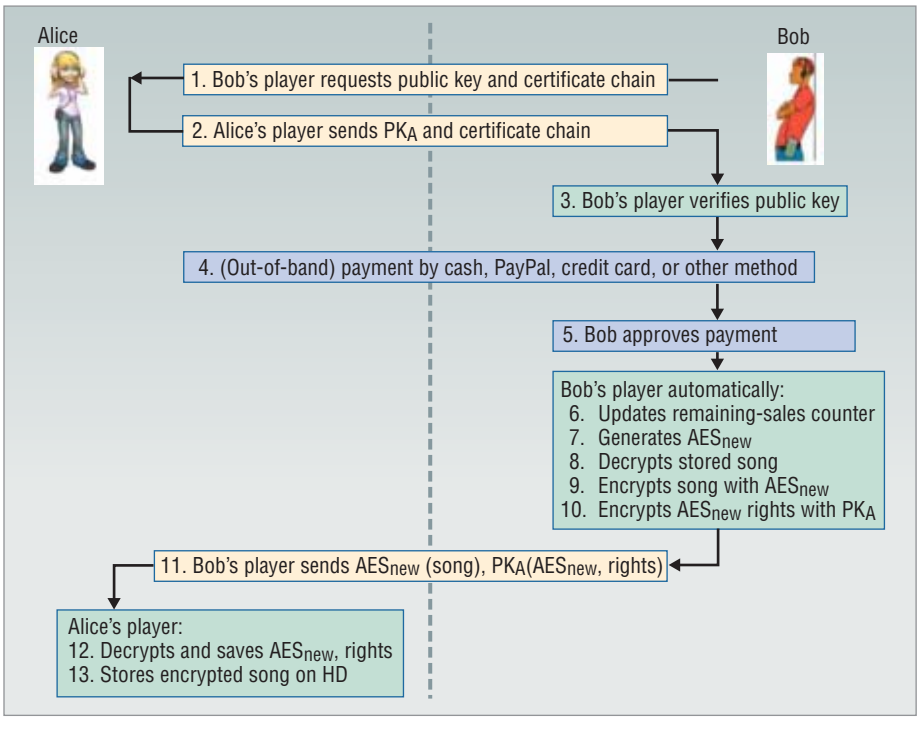

Figure 3. How Bob sells Alice a song. The yellow boxes are messages, the green ones are steps the cryptochips perform, and the blue ones are manual steps.

Since an attacker or misbehaving consumer might misuse content, it's stored encrypted. In addition, a cryptochip (preferably soldered to the player's motherboard or to a PCI plug-in card) performs sensitive operations. The cryptochip must contain a CPU, nonvolatile memory for key storage, and some working RAM. In this way, the player can issue a command to the cryptochip saying, "Fetch the encrypted song at memory address 0x122400, decrypt it with symmetric key #4 in your internal list, and generate audio on your output pins."

In this way, the plaintext music is never released outside the cryptochip. A chip like the trusted platform module (www.trustedcomputinggroup.org/groups/tpm) already provides some of these features. We feel that technology is advanced enough to expect implementation of such features, given enough financial incentives. With the content never appearing in plaintext outside the cryptochip, the security requirements on the rest of the software become much less stringent.

SYSTEM ARCHITECTURE

In our architecture,[2,3] a consumer contacts a provider to buy a song, video, or other content and optionally the right to resell it *N* times. The request message contains the player's public key. Once the consumer has paid for the content and the rights, the provider encrypts the content on the fly with a newly generated symmetric Advanced Encryption Standard (AES) key and sends the encrypted content to the consumer's (now called the reseller's) player.

The content provider also sends the symmetric key encrypted with the player's public key and the rights the reseller bought, expressed in a suitable language. It signs the rights with a secret key to prove they're valid. The secure hardware on the consumer's player checks the integrity of the content and the rights. If valid, they're stored on an insecure hard disk or flash memory. Since the rights information is signed by the content provider and its hash is stored in a secure memory, the owner can't tamper with it.

Process steps

Each time a reseller such as Bob wishes to resell the content, the cryptochip first checks to see if the maximum number of sales the license defines has been reached. If not, it goes through the following steps, which Figure 3 illustrates. In Step 1, Bob asks a buyer such as Alice to send her public key, PK_A, and a certificate chain rooted at the CA over the wireless link to Bob's player. In Step 2, Alice's player sends the PK_A and certificate chain. In Step 3, Bob's cryptochip verifies that Alice's claimed public key is in her certificate, that the player's manufacturer signed the certificate, and that the CA approved the manufacturer.

If all goes well, Bob's player now knows that Alice's player has been certified as compliant. After all, although Alice can easily generate a private-public-key pair, she can't produce a certificate chain back to the CA guaranteeing that the key is authentic, and without this authenticity, Bob's cryptochip won't allow the transaction to proceed.

In Step 4, Alice pays Bob using cash, PayPal, credit card, or another agreed-upon means. This step is out of band and not part of the protocol. In Step 5, when Bob is satisfied with the payment, he pushes a button on his player to approve the sale. In Step 6, the cryptochip first updates the number of sales remaining and keeps this counter in its secure internal memory. Then, in Step 7, it generates a fresh symmetric key, AES_{new}.

Using the existing stored per-song symmetric key, in Step 8, the cryptochip reads the song from main memory and decrypts it, and in Step 9, reencrypts it with AES_{new}, and puts the newly encrypted song elsewhere in insecure RAM, leaving the original intact.

In Step 10, Bob's cryptochip encrypts the new symmetric key, AES_{new}, with Alice's valid public key and puts it

in RAM as well. Bob can steal the song and the key from RAM, but it won't do him any good as they're encrypted, the song with a symmetric key he doesn't know and the song key with Alice's public key. Furthermore, he can't get at the "remaining-sales" counter, which is kept safely in the cryptochip's internal nonvolatile memory.

Next, in Step 11, Bob's player sends Alice the AES-encrypted song and encrypted AES key for this song. In Step 12, upon receipt of the message, Alice's cryptochip decrypts it and saves the song key, AES_{new}, internal to itself; in Step 13, it saves the encrypted song on the (insecure) hard disk or flash memory.

Critical events

Steps 6-10 are performed as a single atomic transaction, but if Bob's player is switched off between steps 6 and 11, maliciously or otherwise, he loses one resale right and must deal with an unhappy customer who didn't get the song she paid for. Other critical events are an accidental communication breakdown between the reseller and the consumer while the transaction is in progress or when the reseller cheats the consumer by delivering a bogus song. To resolve all these situations, a "recovery subprotocol"[3] lets the consumer contact the content provider directly to resolve the issue.

It's important to note that the resale is offline. Neither Bob nor Alice has to contact the content provider since Bob has already paid for Alice's copy of the music (as well as the eight unsold copies) in advance. By using teenagers as salespeople, the content provider saves on computing power and bandwidth costs.

What happens if Bob can't find eight more friends who want the song? The publishers of books, magazines, and newspapers have precisely the same problem, and they generally allow their sales outlets to return unsold stock for credit to encourage them to have an ample supply on hand. Of course, the publisher can rescind any quantity discount granted initially when the vendor returns the unsold copies. Following this tradition, music vendors are likely to follow suit, but that's their business decision to make.

When Alice wants to listen to her newly purchased song, the cryptochip in her player extracts its symmetric key, AES_{new}, stored in its internal memory and fetches, decrypts, and plays the song one block at a time. In this way, the bulk data—the songs—are stored in the large cheap memory, with each song encrypted with a unique symmetric (AES) key.

SECURITY CONSIDERATIONS

However secure and foolproof we assume a system to be, experience has shown that all it takes is a single weak link to compromise its security. Our scheme assumes that certified players behave in the stipulated manner and that they follow the protocols correctly. However, it might be possible to crack a player using out-of-band methods, such as using an electron microscope to read the keys in the cryptochip's internal nonvolatile memory.

Figure 4. Prototype implementation using Neuros OSD boards.

Watermarking and traitor-tracing techniques[4] can be used to identify such compromised players. Once identified, the compromised player's identity (public key) is added into a player revocation list. The system can push this list to each consumer's player the next time it connects to a content provider. Other researchers have proposed various ways to minimize the size of such lists.[5] As an enhancement, the players could also exchange revocation lists when they exchange content. A compliant player is designed to refuse communication with any player listed in the revocation list.

Of course, our system also suffers from the "analog hole" problem. An attacker can always record the content with a microphone while it's being played and redistribute it in an uncontrolled manner. There's no definitive solution for this problem; however, the degradation in the quality of the copy obtained through analog recording could be an attack deterrent.

It's important to note that our system doesn't introduce any new vulnerability. These attacks also apply to current players that don't allow controlled peer-to-peer distribution.

PROTOTYPE IMPLEMENTATION

We've implemented Paradiso (www.few.vu.nl/~srijith/paradiso), a system prototype, using a $230 Neuros development board (http://wiki.neurostechnology.com/index.php/OSD_Beta), representative of what's found in mobile music players. Shown in Figure 4, this board has a TI 200-MHz ARM926, 120-MHz C54x DSP processor specifically developed for multimedia applications, 64 Mbytes of SDRAM and 10/100 Mbps Ethernet port, among others. The board runs a modified version of the Linux 2.6 kernel. We used OpenSSL libraries for cryptographic support and software techniques for atomic actions.[6]

The developer boards acted as compliant players. However, since we couldn't obtain a developer board with a suitable cryptochip and secure store, we used a software layer to emulate the hardware security layer. We believe that once the interface and protocols have been defined, implementing them on another (secure) processor wouldn't be difficult.

Experiments performed with our prototype show that it takes around 10 seconds to perform steps 6-13 for a 5-Mbyte file. Performance measurements also show that the music file's quality doesn't suffer from the lag due to the decryption steps. While the prototype implements the cryptographic steps in the software, a production unit will implement them in the hardware, thus we can expect a speedup and better performance.

Microsoft's Zune took the first baby step toward implementing our proposed system by letting users forward songs to friends. However, the similarity ends there. The recipient still must contact the content provider to purchase the content and associated license. Zune's existence is an indication of the digital medium's potential, as well as content owners' and player manufacturers' receptiveness to explore new avenues to widen their reach.

Although we designed our prototype to generate revenue, DRM technology can easily be extended to serve the needs of consumer-produced digital content. For example, a band could produce and release a song under one of the "noncommercial" Creative Commons licenses and upload it to a content provider as a way to promote its new album. The trusted player, on noticing the song's license, would let the song be exchanged for free.

Similarly, Bob could use the same technology to share the latest video clip he's shot. One of Zune's perceived shortcomings is that irrespective of the origin and license of the content a user exchanges with another, the content is deleted after three plays or days. It's evident that designers incorporated such limitations to prevent using Zune as a new illegal peer-to-peer medium. However, in this age of consumer-generated content, a DRM scheme shouldn't deny copyright owners the right to give away content for free if they so choose. Just imagine the fuss if all computers automatically deleted all free software after three days.

The goals of a Paradiso-like system, however, aren't realizable without some mind-set change. As of now, every player manufacturer uses DRM technology that's not interoperable with other manufacturers. A Paradiso-like architecture would require major manufacturers and content owners to use interoperable DRM technology standards. The success of industry-wide specifications like the mobile industry's Open Mobile Alliance indicates that such an alliance is possible, given strong enough incentives.

Acknowledgments

The work described in this article has been supported by NWO project Account 612.060.319. We thank Bogdan C. Popsecu, Chandana Gamage, Mohammad T. Dashti, and Hugo Jonker for their assistance.

References

1. R. Mori and M. Kawahara, "Superdistribution: The Concept and the Architecture," *IEICE Trans.*, E73, no. 7, 1990, pp. 1133-1146.
2. S.K. Nair et al., "Enabling DRM-Preserving Digital Content Redistribution," *Proc. 7th IEEE Int'l Conf. E-Commerce Technology*, IEEE Press, 2005, pp. 150-158.
3. M.T. Dashti, S.K. Nair, and H.L. Jonker, "Nuovo DRM Paradiso: Towards a Verified Fair DRM Scheme," *Proc. Int'l Symp. Fundamentals of Software Eng.* (FSE 07), Springer-Verlag, 2007, pp. 33-48.
4. B. Chor, A. Fiat, and M. Naor, "Tracing Traitors," *Proc. Advances in Cryptology*, (CRYPTO 94), LNCS 839, Springer-Verlag, 1994, pp. 257-270.
5. D.A. Cooper, "A More Efficient Use of Delta-CRLs," *Proc. 2000 IEEE Symp. Security and Privacy*, IEEE Press, 2000, pp. 190-202.
6. R. Gerrits, "Implementing a DRM-Preserving Digital Content Redistribution System," master's thesis, Vrije Universiteit, 2006.

Srijith K. Nair *is a PhD student in the Department of Computer Science at Vrije Universiteit, Amsterdam. His research interests are information flow control, policy enforcement, system security, and privacy. He received an MSc in computer science from the National University of Singapore. Nair is a member of the IEEE and the ACM. Contact him at srijith@few.vu.nl.*

Ron Gerrits *is the owner of the Web-development firm Inovia. He obtained an MSc in Internet and Web technology from Vrije Universiteit. Contact him at msc.r.gerrits@inovia.nl.*

Bruno Crispo *is an associate professor in the Department of Computer Science at Vrije Universiteit and at the University of Trento, Italy. His research interests are networks, distributed systems, cryptography, and security protocols. Crispo received a PhD in computer science from the University of Cambridge, United Kingdom. He is a member of the IEEE. Contact him at crispo@cs.vu.nl.*

Andrew S. Tanenbaum *is a professor in the Department of Computer Science at Vrije Universiteit. His research interests are reliable operating systems and security. He received a PhD from the University of California, Berkeley. Tanenbaum is a Fellow of the IEEE and the ACM and a member of the Royal Netherlands Academy of Arts and Sciences. Contact him at ast@cs.vu.nl.*

Get the building blocks you need.

Take your career to the next level in software development, systems design, and engineering with:

- Article collections from the IEEE Computer Society
- Materials from Harvard Business School Publishing
- Computer discounts
- Online courses and certifications

Our experts. Your future.

www.computer.org/buildyourcareer

RESEARCH FEATURE

Authorizing Card Payments with PINs

Václav (Vashek) Matyáš, Jan Krhovják, and Marek Kumpost, Masaryk University
Dan Cvrcek, Brno University of Technology

Chip and PIN technology was introduced as a means of decreasing payment-card fraud. However, according to results of a two-phase experiment, the technology makes it easier for thieves to obtain PINs and more difficult for customers to defend against counterfeiting.

The introduction of Chip and PIN technology for card-purchase authorization has led to vigorous discussions within the computer-security community. An implementation of the EMV (Eurocard, MasterCard, and VISA) standard for secure payments, Chip and PIN is intended as a replacement for taking card imprints or swiping cards with magnetic stripes through readers during face-to-face credit- or debit-card transactions. Signature verification is the main means of security for both methods, with the clerk required to compare the customer's signature with the signature on the back of the card. Both systems have proven reasonably effective; however, thieves can still circumvent them by stealing cards from the mail and forgers can mimic card signatures.

Chip and PIN features a smart card containing an embedded chip that's placed in a PIN-pad terminal or modified swipe-card reader. Once the system verifies the card as authentic, the customer enters a four-digit PIN that's matched against the PIN on the card. The system was launched in the United Kingdom in 2004, and banks have since started issuing cards with both magnetic stripes and chips (http://en.wikipedia.org/wiki/Chip_and_PIN).

To resolve questions about this new authentication system's benefits, we conducted a two-part experiment to evaluate whether it makes circumvention easier for thieves and disputing fraudulent transactions more difficult for customers.

PROBLEMS WITH PINS

Chip and PIN technology can potentially increase the cost of card counterfeiting as well as complicate abuse of stolen cards. However, it might also have an adverse effect on Chip and PIN card users dealing with card-loss detection and reporting.

The first problem with PIN authorization lies in repudiation. Customers can fight the loss of a magnetic-stripe card reasonably well by tracking poorly faked signatures. However, losses with Chip and PIN cards are more difficult to dispute, since there's no tracking in the purchase or authorization documents after entry of the correct PIN. Furthermore, customers might have a difficult time obtaining relevant merchants' camera recordings due to limited retention time or merchants' disinterest in investigating disputed purchases. Furthermore, card-activation procedures only partly protect against mail theft to obtain PINs.

The second drawback is that current Chip and PIN systems don't deploy different security mechanisms or settings for various threat environments. Consequently, a thief obtaining a PIN by spying on a low-level transaction can use that PIN for a high-level transaction. While signatures have more entropy than PINs, merchants often don't check them, as many investigators (www.zug.com/pranks/credit) and our experiment verified.

Some indirectly related issues include the following:[1]

- We must allow for both Chip and PIN and signature

Published by the IEEE Computer Society

authorizations and so combine the drawbacks of both.
- The technology change lets some participants shift the parameters of risk exposure.
- Thieves can read information from the card chip at any merchant's reader and use it for other types of fraud.

Critics have uncovered many problems with using Chip and PIN for card-payment authorization (www.chipandspin.co.uk). The technological change from magnetic stripe and signature to Chip and PIN surprisingly makes low-tech attacks on cards less difficult. It's now easier to spy on customers entering their PINs, since PINs are used more frequently, often in overcrowded stores. The need to insert a PIN card in the reader during the entire transaction also reduces merchants' opportunity to check card details and compare signature-purchase authorization.

PIN VERSUS SIGNATURE

We hoped that results from our experiment[2] would help us and the broader community understand the security implications of the new technology, which represents a major security-related computer deployment. In particular, we wanted to find out

- whether PIN-entry spying is easier than signature falsification;
- under what conditions the above holds true or false; and
- alternatives for purchase authorization using chip-equipped payment cards.

Figure 1. The first phase of the experiment was carried out in a university bookstore in near-realistic conditions.

Figure 2. PIN pads. The PIN pads used during the experiment's first phase had varying security shields.

For customers to behave naturally during the experiment, we presented them with a cover story stating a slightly different purpose for our experiment, an approach our university ethically cleared. We based our results on the assumption of an honest merchant and an outside attacker. We also disregarded the possibility that the merchant could use CCTV or modified PIN pads to eavesdrop PINs,[3] although this type of attack might be easy and would have a high success potential. A thief might also deploy a miniature camera for an attack.

We initiated the experiment's first phase in 2005 in near-realistic conditions in our university bookstore with staff and students as participants, as Figure 1 shows. During this phase, we examined the success rates of PIN observation for two PIN pads (one with no security and the second with a very robust protective shield, shown in Figure 2), and of signature falsification, with the signature thoroughly verified.

The second phase took place in 2006 in one of the region's largest supermarkets, with slightly modified settings.

FIRST PHASE

We undertook the first phase in two rounds, focusing on the success rates of observing a customer entering the PIN and falsifying someone else's signature without

the store assistant's detecting the deception. Thirty-two customers participated in the first round, along with four observers, two store assistants, and three experiment supervisors. Three or four bystanders were also present at any given time.

First-phase setup

We set up the experiment so that after shopping, participants in a given round didn't exchange information with participants in subsequent rounds. The shop area plus two separated rooms prevented customers from interfacing. This phase involved simulated payments with nonbranded smart cards.

In our cover story, we told merchants and customers that we were surveying the pros and cons of two different methods of payment authorization. At the time, PIN-based authorization was used for few card payments in the Czech Republic, so it was known that signatures still overshadowed the authorization method.

We measured user comfort and acceptance and the time it took for all related operations. We used the two different types of PIN pads and split customers into two groups. We also asked the customers and merchants to follow their normal security and logistic provisions and fill out an opinion survey.

A person posing as a researcher from the university's School of Social Studies presented the cover story, and members of our research group posed as technology consultants. After the experiment was over, we informed participants about its real purpose.

Round one. In this round, we took all 32 customers into one room and gave each of them a purchase card with a randomly generated PIN. We split the cards into two groups of 17 and 15 participants, with each group using a PIN pad with a different security level. Customers were sent into the bookstore individually, and each picked a random item and approached the counter (or queue at the counter).

We started timing the operation once the customer handed the selected item to the merchant and entered the PIN (correct at the first attempt), then set a 10-second delay for the purchase authorization and receipt printing. We read the final time at the moment the merchant handed the item and a receipt to the customer. Once the customer left the bookstore, we recorded the observers' guesses of the PIN and called in the next customer.

Round two. We used all 32 participants from the first round to create two groups of 15 and 17 members for the second round. We issued the first group cards with their own signatures on the backs, and we gave the second group cards containing other people's signatures. We gave participants from the second group 20 minutes (and, at a special request, 30 minutes) to practice the given signatures.

We told the merchants that some customers would falsify other people's signatures, but didn't indicate how many would do that. The customers again entered the bookstore one at a time, and, as in the first round, we measured their purchase times. The merchants had to determine the signature's validity at the time of the purchase. We let them ask the customers to repeat the signatures if they were in doubt.

First-phase results

We divided the results into four parts. The first part contained interesting data from the opinion survey. That was followed by results obtained for the PIN pads with and without security. The final part covered signatures.

Opinion survey. While we asked participants to fill out the opinion surveys just to strengthen their belief that the experiment was about the user-friendliness of customer-authorization technologies, the results are worth mentioning.

Twenty-five of the 32 participants had used magnetic-stripe cards for payments, and about half had used Chip and PIN payment cards. On a scale where 1 is the best and 5 the worst, participants gave magnetic-stripe card payments an overall satisfaction level of 3.4, compared with 2.5 for Chip and PIN card payments.

Given the options of 10, 20, 30, 40, and 50 seconds as the maximum acceptable time for the entire payment operation, the participants agreed on an average of 21 seconds. And finally, participants experienced no problems with the card payment in 89 percent of their transactions. In about 7.5 percent of instances, they experienced a minor nuisance, in 2 percent a major problem, and in less than 2 percent of cases, they couldn't pay with their cards.

PIN pad with security. For the 17 customers who made their card payments using PINs, six observers (or 35.5 percent) successfully guessed the PIN. In five of those six instances, the thieves guessed the PIN on their first attempt. For three PINs, two observers got the PIN right, and for two PINs, one observer learned the PIN. In the sixth case, the observers built the PIN using their shared knowledge. Viewing the recordings from another point of view, the observers correctly reported 75 digits (48 percent) in 39 tips of the four-digit PINs (for 156 digits altogether).

PIN pad without security. Results for this PIN pad were shockingly different. Participants lost their PINs to the observers in 12 of 15 cases (or 80 percent). Observers correctly guessed 10 PINs at the first attempt. All four observers saw two PINs, three observers saw one pin, two observers saw four PINS, and one observer saw three PINs. The remaining two correct guesses were built from the shared knowledge.

Using the alternative view, from the 46 tips of four-digit PINs (184 digits) the observers provided, they guessed 129 digits (or 70.1 percent) correctly.

We instructed the four observers to make sure one observer interacted with bystanders instead of watching customers. In all four cases, the observers reported results. The fourth observer noted that "he couldn't have

helped" seeing the ease of PIN-entry observation even at a distance while talking to the bystanders.

Signatures. The merchants correctly identified 12 of the 17 cheaters (five of them bypassed the merchant's control). The merchants pointed out eight of the 12 fraudsters right after their first attempt, and the remaining four after their second signature.

In the group of 20 customers (five cheaters and 15 signing their own signatures), only four had to sign twice to convince the merchants. Both the participants and the supervisors for this experiment believed that the merchant performed a thorough check. We attribute this to the signature verifier working in a luxury jewelry store where signature checks are conducted more thoroughly than in a supermarket or bookstore.

SECOND PHASE

While we initially weren't certain we'd run the second phase of the experiment, results from the first phase reassured us that we had to undertake the second phase to determine whether PIN-entry observation was as easy in a supermarket as a bookstore and to establish the level of signature-forgery detection in a less strict environment.

Second-phase setup

We decided to run the second phase using genuine payment cards in one of the region's largest supermarkets. We opened new accounts in our names for this experiment and set appropriate legal and contractual arrangements between the formal account owners, cardholders, and actual card users so that cardholders wouldn't deem the card usage within our experiment as unauthorized.

To obtain unbiased data, we needed as few people as possible to know about the experiment, including cash-register assistants and ground-floor security. Therefore, the only people aware of the experiment were the university experiment supervisors; a supermarket chain manager; and the supermarket's director, head of security, and chief camera supervisor.

We performed two rounds using goals similar to the first-phase rounds. Twenty relatives and friends who didn't know anything about the experiment's first phase took part as customers in round one of the second phase. In the cover story, we told participants that we were examining the security of payment processes and the supermarket's internal procedures. We told them we'd be using our PhD students' cards with the students' consent.

Essentially, we shared all practical aspects of the experiment. However, we didn't tell them about the three teams observing our customers entering the PINs at the cash registers. We again asked participants to fill out a questionnaire to reassure them about the pretended purposes of our experiment. We revealed its real purpose after the experiment was over.

Of the 20 customers who took part in this round, 13 used tills equipped with security-shielded PIN pads, and the remaining seven used tills equipped with unsecured PIN pads, as Figure 3 shows.

Figure 3. The supermarket PIN pad without shielding.

Three groups of observers operated on the ground floor, reporting the results of their individual observations to three experiment supervisors. Each customer was always observed only by one dedicated group of observers, and a supervisory team member made sure observers followed the right person and was available to handle incidents on the ground floor. We took care to have equal representation of sexes and ages in the customer group.

Second-phase results

Observation results were initially surprising. Observers would correctly put together four PINs within the typical three-attempt limit, yet only observed one customer entry into an unsecured PIN pad (out of seven such observations). The remaining three came from those 13 observations of payments using secured PIN pads. The other three observations would lead to the right PIN within 10 attempts, and customers could guess three more PINs in less than 222 attempts. Using the alternative view, from the tips of four-digit PINs the observers provided, they guessed 42 percent correctly.

A more thorough analysis of results that groups of observers obtained led to a reasonable explanation of the results. Of the three groups of observers, the first correctly reported individual PINs in one-quarter of the cases. The second correctly reported PINs in 27 percent of the cases, while the third had a 68 percent success rate. Not surprisingly, the third group of observers discovered the four correct PINs within three attempts. Particularly, one member of the third group was rather

"assertive" and thus able to observe customers from very convenient spots.

Then we asked volunteers to spend 10 to 30 minutes practicing someone else's signature. When they visited the shop later, all 17 signatures were accepted without the volunteers being asked to sign a second time or show their IDs. Several volunteers reported that the clerk barely looked at their signatures. The result demonstrates that it might not be a good idea to require payment-card signature authentication with little or no human intervention.

The experiment's second phase results are unpleasant for card users, considering that signature forgers and PIN-entry observers were new to their tasks. All signature forgeries went undetected, and PIN observations yielded a success rate over 35 percent. However, we acknowledge that signature verification will likely be more thorough in a luxury jewelry store than a supermarket.

Our experiment indicated that the risk of PIN compromise is high, and eliminating paper records with signatures significantly reduces customers' ability to defend themselves during disputes over unauthorized transactions. While the proportion of correctly observed PIN digits in our experiments was roughly half (60 percent in the first and 42 percent in the second phase, respectively), we've seen partial evidence of skilled observers achieving a success rate of more than two-thirds in both the supermarket and the bookstore.

In addition, there's considerable room for improvement in the signature-verification process. Figures from both stages of our experiment support the view that PIN pads need robust security shielding. Our experiment suggests that we're replacing a weak biometric with an equally weak or possibly even weaker means of customer payment authorization. The risk of PIN-entry observation is clearly different at an ATM and a PIN pad in a crowded store. Biometric authentication might be a long-term remedy for the problem, yet the path to successful deployment will likely take years. We see PINs and verified signatures and different PINs for low- and high-level transactions as short-term solutions.

From the global security point of view, we can improve crime detection. However, as it's currently deployed, Chip and PIN makes detection much harder. Chip and PIN reduces some credit-card related crime, but it also shifts liability from banks to customers even further. In general, payment systems should acknowledge possibilities of unauthorized transactions and the fact that many customers are strongly disadvantaged against professional thieves. Currently, customer-protection mechanisms are the best solution, including careful control of card possession and the card issuer bearing contractual responsibility for card losses. Dedicated customer devices, blocking some types or levels of transactions, and secondary PINs might serve as solutions in the future.

Acknowledgments

We thank the Future of Identity in the Information Society consortium and our university management for supporting experiment, and Monet+ for providing the PIN pads; Pavel Mareček for providing access to his bookstore; and Dalibor Hanák for verifying customer signatures. Thanks also to Jozef Vyskoč and Tomáš Plos for assisting us with arrangements for the experiment.

References

1. B. Adida et al., "Phish and Chips—Traditional and New Recipes for Attacking EMV," *Proc. 2006 Int'l Workshop Security Protocols*, to appear in LNCS, Springer-Verlag.
2. D. Cvrcek, J. Krhovják, and V. Matyáš "PIN (and Chip) or Signature—Beating the Cheating?" *Proc. 2005 Int'l Workshop Security Protocols*, LNCS 4631, Springer-Verlag, 2007, pp. 69-75.
3. R. Anderson and M. Bond, "The Man-in-the-Middle Defence," *Proc. 2006 Int'l Workshop Security Protocols*, to appear in LNCS, Springer-Verlag; www.cl.cam.ac.uk/~rja14/Papers/Man-in-the-Middle-Defence.pdf.

Václav (Vashek) Matyáš *is an associate professor and chair of the Department of Computer Systems and Communications at Masaryk University, Czech Republic. His research interests are applied cryptography and security. He received a PhD in mathematics and informatics from Masaryk University. Contact him at matyas@fi.muni.cz.*

Dan Cvrcek *is an associate professor at the Brno University of Technology, Czech Republic, and a research associate at the University of Cambridge in the United Kingdom. His research interests are IT security and applied cryptography. Dan received a PhD in informatics from the Brno University of Technology. Contact him at cvrcek@fit.vutbr.cz.*

Jan Krhovják *is a PhD student in Masaryk University's Faculty of Informatics. His research interests are authentication issues and generating key material from passwords and random numbers. He received an MS in informatics. Contact him at xkrhovj@fi.muni.cz.*

Marek Kumpost *is a PhD student in Masaryk University's Faculty of Informatics and a research assistant at the Brno University of Technology. His research interests are privacy modeling issues and wireless network security. He received an MS in informatics from Masaryk University. Contact him at xkumpost@fi.muni.cz.*

RESEARCH FEATURE

The Promise of High-Performance Reconfigurable Computing

Tarek El-Ghazawi, Esam El-Araby, and Miaoqing Huang, George Washington University
Kris Gaj, George Mason University
Volodymyr Kindratenko, University of Illinois at Urbana-Champaign
Duncan Buell, University of South Carolina

Several high-performance computers now use field-programmable gate arrays as reconfigurable coprocessors. The authors describe the two major contemporary HPRC architectures and explore the pros and cons of each using representative applications from remote sensing, molecular dynamics, bioinformatics, and cryptanalysis.

In the past few years, high-performance computing vendors have introduced many systems containing both microprocessors and field-programmable gate arrays. Three such systems—the Cray XD1, the SRC-6, and the SGI Altix/RASC—are parallel computers that resemble modern HPC architectures, with added FPGA chips. Two of these machines, the Cray XD1 and SGI Altix, also function as traditional HPCs without the reconfigurable chips. In addition, several Beowulf cluster installations contain one or more FPGA cards per node, such as HPTi's reconfigurable cluster from the Air Force Research Laboratory.

In all of these architectures, the FPGAs serve as coprocessors to the microprocessors. The main application executes on the microprocessors, while the FPGAs handle kernels that have a long execution time but lend themselves to hardware implementations. Such kernels are typically data-parallel overlapped computations that can be efficiently implemented as fine-grained architectures, such as single-instruction, multiple-data (SIMD) engines, pipelines, or systolic arrays, to name a few.

Figure 1 shows that a transfer of control can occur during execution of the application on the microprocessor, in which case the system invokes an appropriate architecture in a reconfigurable processor to execute the target operation. To do so, the reconfigurable processor can configure or reconfigure the FPGA "on the fly," while the system's other processors perform computations. This feature is usually referred to as runtime reconfiguration.[1]

From an application development perspective, developers can create the hardware kernel using hardware description languages such as VHDL and Verilog. Other systems allow the use of high-level languages such as SRC Computers' Carte C and Carte Fortran, Impulse Accelerated Technologies' Impulse C, Mitrion C from Mitrionics, and Celoxica's Handel-C. There are also high-level graphical programming development tools such as Annapolis Micro Systems' CoreFire, Starbridge Systems' Viva, Xilinx System Generator, and DSPlogic's Reconfigurable Computing Toolbox.

Readers should consult *Computer*'s March 2007 special issue on high-performance reconfigurable computing for a good overview of modern HPRC systems, application-development tools and frameworks, and applications.

HPRC ARCHITECTURAL TAXONOMY

Many early HPRC systems, such as the SRC-6E and the Starbridge Hypercomputer, can be seen as attached processors. These systems were designed around one node of microprocessors and another of FPGAs. The

0018-9162/08/$25.00 © 2008 IEEE Published by the IEEE Computer Society

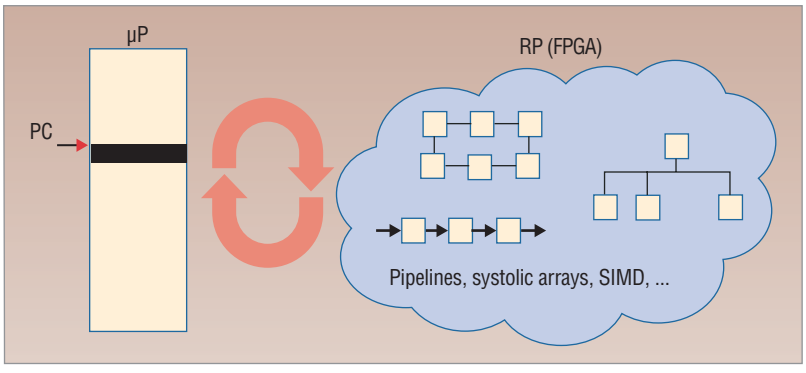

Figure 1. In high-performance reconfigurable computers, field-programmable gate arrays serve as coprocessors to the microprocessors. During execution of the application on the microprocessor, the system invokes an appropriate architecture in the FPGA to execute the target operation.

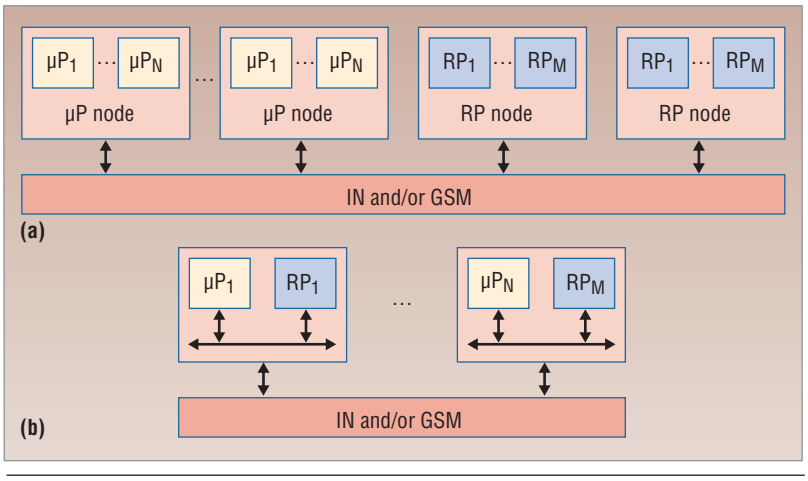

Figure 2. Modern HPRCs can be grouped into two major classes: (a) uniform node nonuniform systems (UNNSs) and (b) nonuniform node uniform systems (NNUSs).

two nodes were connected directly, without a scalable interconnection mechanism.

Here we do not address these early attached processor systems but focus instead on scalable parallel systems such as the Cray XD1, SRC-6, and SGI Altix/RASC as well as reconfigurable Beowulf clusters. These architectures can generally be distinguished by whether each node in the system is homogeneous (uniform) or heterogeneous (nonuniform).[2] A uniform node in this context contains one type of processing element—for example, only microprocessors or FPGAs. Based on this distinction, modern HPRCs can be grouped into two major classes: uniform node nonuniform systems and nonuniform node uniform systems.

Uniform node nonuniform systems

In UNNSs, shown in Figure 2a, nodes strictly have either FPGAs or microprocessors and are linked via an interconnection network to globally shared memory (GSM). Examples of such systems include the SRC-6 and the Altix/RASC. The major advantage of UNNSs is that vendors can vary the ratio of reconfigurable nodes to microprocessor nodes to meet the different demands of customers' applications. This is highly desirable from an economic perspective given the cost difference between FPGAs and microprocessors, and it is particularly suitable for special-purpose systems.

On the downside, having the reconfigurable node and the microprocessor node interact over the shared interconnection network makes them compete for overall bandwidth, and it also increases the latency between the nodes. In addition, code portability could become an issue even within the same type of machine if there is a change in the ratio between the microprocessor nodes and the FPGA nodes.

A representative example of the UNNS is the SRC-6/SRC-7, which consists of one or more general-purpose microprocessor subsystems, one or more MAP reconfigurable subsystems, and global common memory (GCM) nodes of shared memory space. These subsystems are interconnected through a Hi-Bar switch communication layer. The microprocessor boards each include two 2.8-GHz Intel Xeon microprocessors and are connected to the Hi-Bar switch through a SNAP interface. The SNAP card plugs into the dual in-line memory module slot on the microprocessor motherboard to provide higher data transfer rates between the boards than the less efficient but common peripheral component interconnect (PCI) solution. The sustained transfer rate between a microprocessor board and the MAP processors is 1,400 Mbytes per second.

The MAP Series C processor consists of one control FPGA and two user FPGAs, all Xilinx Virtex II-6000-4s. Additionally, each MAP unit contains six interleaved banks of onboard memory (OBM) with a total capacity of 24 Mbytes. The maximum aggregate data transfer rate among all FPGAs and OBM is 4,800 MBps. The user FPGAs are configured such that one is in master mode and the other is in slave mode. A bridge port directly connects a MAP's two FPGAs. Further, MAP processors can be connected via a chain port to create an FPGA array.

Nonuniform node uniform systems

NNUSs, shown in Figure 2b, use only one type of node, thus the system level is uniform. However, each node contains both types of resources, and the FPGAs are connected directly to the microprocessors inside the node.

Examples of such systems are the Cray XD1 and reconfigurable clusters. NNUSs' main drawback is their fixed ratio of FPGAs to microprocessors, which might not suit the traditional vendor-buyer economic model. However, they cater in a straightforward way to the single-program, multiple-data (SPMD) model that most parallel programming paradigms embrace. Further, the latency between the microprocessor and its FPGA coprocessor can be low, and the bandwidth between them will be dedicated—this can mean high performance for many data-intensive applications.

A representative example of the NNUS is the Cray XD1, whose direct-connected processor (DCP) architecture harnesses multiple processors into a single, unified system. The base unit is a chassis, with up to 12 chassis per cabinet. One chassis houses six compute cards, each of which contains two 2.4-GHz AMD Opteron microprocessors and one or two RapidArray Processors (RAPs) that handle communication. The two Opteron microprocessors are connected via AMD's HyperTransport technology with a bandwidth of 3.2 GBps forming a two-way symmetric multiprocessing (SMP) cluster. Each XD1 chassis can be configured with six application-acceleration processors based on Xilinx Virtex-II Pro or Virtex-4 FPGAs. With two RAPs per board, a bandwidth of 8 GBps (4 GBps bidirectional) between boards is available via a RapidArray switch. Half of this switch's 48 links connect to the RAPs on the compute boards within the chassis, while the others can connect to other chassis.

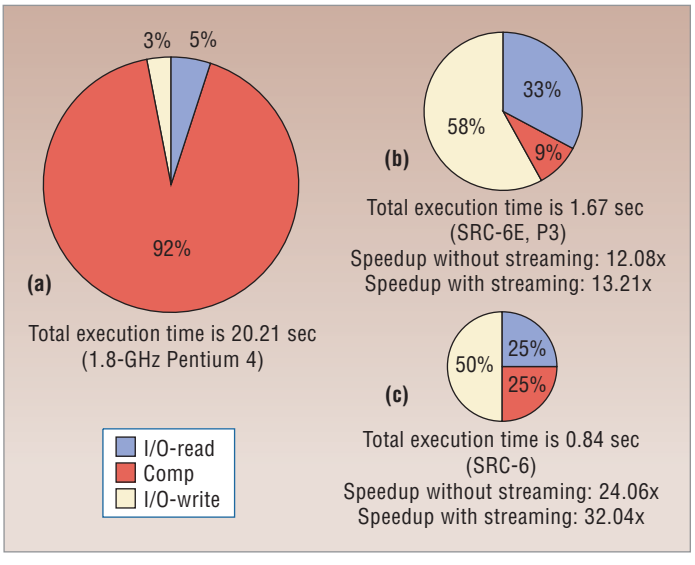

Figure 3. Execution profiles of hyperspectral dimension reduction. (a) Total execution time on 1.8-GHz Pentium 4 microprocessor. (b) Total execution time on SRC-6E. (c) Total execution time on SRC-6.

NODE-LEVEL ISSUES

We have used the SRC-6E and SRC-6 systems to investigate node-level performance of HPRC architectures in processing remote sensing[3] and molecular dynamics[4] applications. These studies included the use of optimization techniques such as pipelining and data transfer overlapping with computation to exploit the inherent temporal and spatial parallelism of such applications.

Remote sensing

Hyperspectral dimension reduction[3] is representative of remote sensing applications with respect to node performance. With FPGAs as coprocessors for the microprocessor, substantial data in this data-intensive application must move back and forth between the microprocessor memory and the FPGA onboard memory. While the bandwidth for such transfers is on the order of GBps, the transfers are an added overhead and represent a challenge on the SRC-6 given the finite size of its OBM.

This overhead can be avoided altogether through the sharing of memory banks, or the bandwidth can be increased to take advantage of FPGAs' outstanding processing speed. Overlapping memory transfers—that is, streaming—between these two processing elements and the computations also can help. As Figure 3a shows, such transfers (I/O read and write operations) take only 8 percent of the application execution time on a 1.8-GHz Pentium 4 microprocessor, while the remaining 92 percent is spent on computations.

As Figure 3b shows, the first-generation SRC-6E achieves a significant speedup over the microprocessor: 12.08× without streaming and 13.21× with streaming. However, the computation time is now only 9 percent of the overall execution time. In the follow-up SRC-6, the bandwidth between the microprocessor and FPGA increases from 380 MBps (sustained) to 1.4 GBps (sustained). As Figure 3c shows, this system achieves a 24.06× speedup (without streaming) and a 32.04× speedup (with streaming) over the microprocessor.

These results clearly demonstrate that bandwidth between the microprocessor and the FPGA must be increased to support more data-intensive applications—an area the third-generation SRC-7 is likely to address. It should be noted, however, that in most HPRCs today, transfers between the microprocessor and FPGA are explicit, further complicating programming models. These two memory subsystems should either be fused into one or integrated into a hierarchy with the objective of reducing or eliminating this overhead and making the transfers transparent.

Molecular dynamics

Nanoscale molecular dynamics (NAMD)[4] is representative of floating-point applications with respect to node performance. A recent case study revealed that when porting such highly optimized code, a sensible approach is to use several design iterations, starting with

the simplest, most straightforward implementation and gradually adding to it until achieving the best solution or running out of FPGA resources.[5]

The study's final dual-FPGA-based implementation was only three times faster than the original code execution. These results, however, are data dependent. For a larger cutoff radius, the original CPU code executes in more than 800 seconds while the FPGA execution time is unchanged, which would constitute a 260× speedup. The need to translate data between the C++ data storage mechanisms and the system-defined MAP/FPGA data storage architecture required considerable development effort. When creating code from scratch to run on an FPGA architecture, a programmer would implement the data storage mechanisms compatible between the CPU and FPGA from the beginning, but this is rarely the case for existing code and adds to the amount of work required to port the code.

Although the "official benchmark" kernel employs double-precision floating-point arithmetic, the NAMD researchers applied algorithmic optimization techniques and implemented their kernel using single-precision floating-point arithmetic for atom locations and 32-bit integer arithmetic for forces. Consequently, the final design occupies most available slices (97 percent), yet utilization of on-chip memory banks (40 percent) and hardware multipliers (28 percent) is low. The fact that the slice limit was reached before any other resource limits suggests that it might be necessary to restructure code to better utilize other available resources. One possible solution is to overlap calculations with data transfer for the next data set to use more available on-chip memory.

Despite the relatively modest speedup achieved, the NAMD study clearly illustrates the potential of HPRC technology. FPGA code development traditionally begins with writing code that implements a textbook algorithm, with little or no optimization. When porting such unoptimized code to an HPRC platform and taking care to optimize the FPGA design, it is easy to obtain a 10×-100× speedup. In contrast, we began with decade-old code optimized to run on the CPU-based platform; such code successfully competes with its FPGA-ported counterpart. It is important to keep in mind that the study's 100-MHz FPGA achieved a 3× application performance improvement over a 2.8-GHz CPU, and FPGAs are on a faster technology growth curve than CPUs.[6]

Lessons learned

Optimization techniques such as overlapping data transfers between the microprocessors and FPGAs with computations are useful for data-intensive, memory-bound applications. However, such applications, including hyperspectral dimension reduction and NAMD, can only achieve good performance when the underlying HPRC architecture supports features such as streaming or overlapping. Streaming can be enabled by architectures that are characterized by high I/O bandwidth and/or tight coupling of FPGAs with associated microprocessors. New promising examples of these are DCP architectures such as AMD's Torrenza initiative for HyperTransport links as well as Intel's QuickAssist technology supporting front side bus (FSB) systems. Large enough memory bandwidth is another equally important feature.

By memory bandwidth we mean that the memory system has sufficient multiplicity as well as speed, width, or depth/size. In other words, because FPGAs can produce and consume data at a high degree of parallelism, the associated memory system should also have an equal degree of multiplicity. Simply put, a large multiple of memory banks with narrow word length of local FPGA memory can be more useful to memory-bound applications on HPRCs than larger and wider memories with fewer parallel banks.

Vendor-provided transparent transfers can enhance performance by guaranteeing the most efficient transfer modes for the underlying platform.

In addition, further node architecture developments are clearly necessary to support programming models with transparent transfers of data between FPGAs and microprocessors by integrating the microprocessor memory and the FPGA memory into the same hierarchy. Vendor-provided transparent transfers can enhance performance by guaranteeing the most efficient transfer modes for the underlying platform. This will let the user focus on algorithmic optimizations that can benefit the application under investigation rather than data transfers or distribution. It also can improve productivity.

SYSTEM-LEVEL ISSUES

We have used the SRC-6 and Cray XD1 systems to investigate system-level performance of HPRC architectures in bioinformatics[7] and cryptanalysis[8-10] applications. These applications provide a near-practical upper bound on HPRC potential performance as well as insight into system-level programmability and performance issues apart from those associated with general high-performance computers. They use integer arithmetic, an area where HPRCs excel, are compute-intensive with lots of computations and not much data transfer between the FPGAs and microprocessors, and inherit both spatial and temporal parallelism.

We distributed the workload of both types of applications over all nodes using the message passing interface (MPI). In the case of DNA and protein analysis, we broadcast a database of reference sequences and scatter sequence queries. The application identified matching scores locally and then gathered them together. Each

				Expected		Measured	
				Throughput (GCUPS)	Speedup	Throughput (GCUPS)	Speedup
FASTA (ssearch34)	Opteron 2.4 GHz	DNA		NA	NA	0.065	1
		Protein		NA	NA	0.130	1
SRC-6 100 MHz (32x1)		DNA	1 Engine/chip	3.2	49.2×	3.19 → 12.2 1 → 4 chips	49 → 188 1 → 4 chips
			4 Engines/chip	12.8	197×	12.4 → 42.7 1 → 4 chips	191 → 656 1 → 4 chips
			8 Engines/chip	25.6	394×	24.1 → 74 1 → 4 chips	371 → 1,138 1 → 4 chips
		Protein		3.2	24.6×	3.12 → 11.7 1 → 4 chips	24 → 90 1 → 4 chips
XD1 200 MHz (32x1)		DNA	1 Engine/chip	6.4	98×	5.9 → 32 1 → 6 chips	91 → 492 1 → 6 chips
			4 Engines/chip	25.6	394×	23.3 → 120.7 1 → 6 chips	359 → 1,857 1 → 6 chips
			8 Engines/chip	51.2	788×	45.2 → 181.6 1 → 6 chips	695 → 2,794 1 → 6 chips
		Protein		6.4	49×	5.9 → 34 1 → 6 chips	45 → 262 1 → 6 chips

Figure 4. DNA and protein sequencing on the SRC-6 and Cray XD1 versus the open source FASTA program. An FPGA with one engine produced a 91× speedup, while eight cores on the same chip collectively achieved a 695× speedup.

FPGA had as many hardware kernels for the basic operation as possible. In the case of cryptanalysis, we broadcast the ciphertext as well as the corresponding plaintext; upon finding the key, a worker node sent it back to the master to terminate the search.

Bioinformatics

Figure 4 compares DNA and protein sequencing on the SRC-6 and Cray XD1 with the open source FASTA program running on a 2.4-GHz Opteron microprocessor. We used giga cell updates per second (GCUPS) as the throughput metric as well as to compute speedup over the Opteron. With its FPGA chips running at 200 MHz, the XD1 had an advantage over the SRC-6, which could run its FPGAs at only 100 MHz.

By packing eight kernels on each FPGA chip, the Cray XD1 achieved a 2,794× speedup using one chassis with six FPGAs. An FPGA with one engine produced a 91× speedup instead of the expected 98× speedup due to associated overhead such as pipeline latency, resulting in 93 percent efficiency. On the other hand, eight cores on the same chip collectively achieved a 695× speedup instead of the expected 788× speedup due to intranode communication and I/O overhead. The achieved speedup for eight engines/chip was 2,794× instead of the estimated (ideal) of 4,728× due to MPI internode communications overhead, resulting in 59 percent efficiency.

These results demonstrate that, with FPGAs' remarkable speed, overhead such as internode and intranode communication must be at much lower levels in HPRCs than what is accepted in conventional high-performance computers. However, given the speed of HPRCs, very large configurations might not be needed.

Cryptanalysis

The cryptanalysis results, shown in Tables 1 and 2, are even more encouraging, especially since this application has even lower overhead. With the Data Encryption Standard (DES) cipher, the SRC-6 achieved a 6,757× speedup over the microprocessor—again, a 2.4-GHz Opteron—while the Cray XD1 achieved a 12,162× speedup. The application's scalability is almost ideal.

In the case of the Cray XD1, straightforward MPI application resulted in using all nodes. However, it made sense for the node program to run on only one microprocessor and its FPGA; the other microprocessors on each node were not used. On the SRC-6, MPI processes had to run on the microprocessors, and the system had to establish an association between each microprocessor and a MAP processor. Because the SRC-6 was limited to two network interface cards that could not be shared efficiently, two MPI processes were sufficient. This meant the program could only run on one microprocessor and one MAP processor.

Lessons learned

Heterogeneity at the system level—namely, UNNS architectures—can be challenging to most accepted

Table 1. Secret-key cipher cryptanalysis on SRC-6.

Application	Hardware		Software		Speedup
	Number of search engines	Throughput (keys/s)	Number of search engines	Throughput (keys/s)	
Data Encryption Standard (DES) breaking	40	4,000 M	1	0.592 M	6,757×
International Data Encryption Algorithm (IDEA) breaking	16	1,600 M	1	2.498 M	641×
RC5-32/12/16 breaking	4	400 M	1	0.351 M	1,140×
RC5-32/8/8 breaking	8	800 M	1	0.517 M	1,547×

Table 2. Secret-key cipher cryptanalysis on Cray XD1.

Application	Hardware		Software		Speedup
	Number of search engines	Throughput (keys/s)	Number of search engines	Throughput (keys/s)	
Data Encryption Standard (DES) breaking	36	7,200 M	1	0.592 M	12,162×
International Data Encryption Algorithm (IDEA) breaking	30	6,000 M	1	2.498 M	2,402×
RC5-32/8/8 breaking	6	1,200 M	1	0.517 M	2,321×

SPMD programming paradigms. This occurs because current technology utilizes the reconfigurable processors as coprocessors to the main host processor through a single unshared communication channel. In particular, when the ratio of microprocessors, reconfigurable processors, and their communication channels differs from unity, SPMD programs, which generally assume a unity ratio, might underutilize some of the microprocessors. On the other hand, heterogeneity at the node level does not present a problem for such programs.

Heterogeneity at the system level is driven by nontechnological factors such as cost savings, which developers can achieve by tailoring systems to customers using homogeneous node architectures. However, this is at least partly offset by the increased difficulty in code portability. NNUS architectures are more privileged in this respect than their UNNS counterparts.

HPRC PERFORMANCE IMPROVEMENT

To assess the potential of HPRC technology, we exploited the maximum hardware parallelism in the previously cited studies' testbeds at both the chip and system levels. For each application, we filled the chip with as many hardware cores as possible that can run in parallel. We obtained additional system-level parallelism via parallel programming techniques, using the MPI to break the overall problem across all available nodes in order to decrease execution time. After estimating the size of a computer cluster capable of the same level of speedup, we derived the corresponding cost, power, and size savings that can be achieved by an SRC-6, Cray XD1, and SGI Altix 4700 with an RC100 RASC module compared with a conventional high-performance PC cluster.

As Tables 3-5 show, the improvements are many orders of magnitude larger. In this analysis, a 100× speedup indicates that the HPRC's cost, power, and size are compared to those of a 100-processor Beowulf cluster. The estimates are very conservative, because when parallel efficiency is considered, a 100-processor cluster will likely produce a speedup much less than 100×—in other words, we assumed the competing cluster to be 100 percent efficient. We also assumed that one cluster node consumes about 220 watts, and that 100 cluster nodes have a footprint of 6 square feet. Based on actual prices, we estimated the cost ratio to be 1:200 in the case of the SRC-6 and 1:100 in the case of the Cray XD1. The cost reduction is actually much larger than the tables indicate when considering the systems' associated power and size.

These dramatic improvements can be viewed as realistic upper bounds on the promise of HPRC technology because the selected applications are all compute-intensive integer applications, a class at which HPRCs clearly excel. However, with additional FPGA chip improvements in the areas of size and floating-point support, and with improved data-transfer bandwidths between FPGAs and their external local memory as well as between the microprocessor and the FPGA, a much wider range of applications can harness similar levels of

benefits. For example, in the hyperspectral dimension reduction study, data transfer improvements between the SRC-6E and SRC-6, while using the same FPGA chips, almost doubled the speedup.

Our research revealed that HPRCs can achieve up to four orders of magnitude improvement in performance, up to three orders of magnitude reduction in power consumption, and two orders of magnitude savings in cost and size requirements compared with contemporary microprocessors when running compute-intensive applications based on integer arithmetic.

In general, these systems were less successful in processing applications based on floating-point arithmetic, especially double precision, whose high usage of FPGA resources constitutes an upper bound on fine-grained parallelism for application cores. However, they can achieve as high performance on embarrassingly parallel floating-point applications, subject to area constraints, as integer arithmetic applications. FPGA chips will likely become larger and have more integrated cores that can better support floating-point operations.

Our future work will include a comprehensive study of software programming tools and languages and their impact on HPRC productivity, as well as multitasking/multiuser support on HPRCs. Because porting applications from one machine to another, or even to the same machine after a hardware upgrade, is nontrivial, hardware architectural virtualization and runtime systems support for application portability is another good research candidate.

Table 3. Performance improvement of SRC-6 compared with a Beowulf cluster.

Application	Speedup	Savings		
		Cost	Power	Size
DNA and protein sequencing	1,138×	6×	313×	34×
DES breaking	6,757×	34×	856×	203×
IDEA breaking	641×	3×	176×	19×
RC5 breaking	1,140×	6×	313×	34×

Table 4. Performance improvement of Cray XD1 compared with a Beowulf cluster.

Application	Speedup	Savings		
		Cost	Power	Size
DNA and protein sequencing	2,794×	28×	148×	29×
DES breaking	12,162×	122×	608×	127×
IDEA breaking	2,402×	24×	120×	25×
RC5 breaking	2,321×	23×	116×	24×

Table 5. Performance improvement of SGI Altix 4700 with RC100 RASC module compared with a Beowulf cluster.

Application	Speedup	Savings		
		Cost	Power	Size
DNA and protein sequencing	8,723×	22×	779×	253×
DES breaking	28,514×	96×	3,439×	1,116×
IDEA breaking	961×	2×	86×	28×
RC5 breaking	6,838×	17×	610×	198×

References

1. M. Taher and T. El-Ghazawi, "A Segmentation Model for Partial Run-Time Reconfiguration," *Proc. IEEE Int'l Conf. Field Programmable Logic and Applications*, IEEE Press, 2006, pp. 1-4.
2. T. El-Ghazawi, "Experience with Early Reconfigurable High-Performance Computers," 2006; http://hpcl.seas.gwu.edu/talks/Tarek_DATE2006.ppt.
3. S. Kaewpijit, J. Le Moigne, and T. El-Ghazawi, "Automatic Reduction of Hyperspectral Imagery Using Wavelet Spectral Analysis," *IEEE Trans. Geoscience and Remote Sensing*, vol. 41, no. 4, 2003, pp. 863-871.
4. J.C. Phillips et al., "Scalable Molecular Dynamics with NAMD," *J. Computational Chemistry*, vol. 26, no. 16, 2005, pp. 1781-1802.
5. V. Kindratenko and D. Pointer, "A Case Study in Porting a Production Scientific Supercomputing Application to a Reconfigurable Computer," *Proc. 14th Ann. IEEE Symp. Field-Programmable Custom Computing Machines*, IEEE CS Press, 2006, pp. 13-22.
6. K. Underwood, "FPGAs vs. CPUs: Trends in Peak Floating-Point Performance," *Proc. 12th ACM/SIGDA Int'l Symp. Field Programmable Gate Arrays*, ACM Press, 2004, pp. 171-180.
7. D.W. Mount, *Bioinformatics: Sequence and Genome Analysis*, 2nd ed., Cold Spring Harbor Laboratory Press, 2004.
8. O.D. Fidanci et al., "Implementation Trade-Offs of Triple DES in the SRC-6E Reconfigurable Computing Environment," *Proc. 5th Ann. Int'l Conf. Military and Aerospace Programmable Logic Devices*, 2002; www.gwu.edu/~hpc/rcm/publications/MAPLD2002.pdf.
9. R.L. Rivest, "The RC5 Encryption Algorithm," revised version, MIT Laboratory for Computer Science, Cambridge, Mass., 20 Mar. 1997; http://people.csail.mit.edu/rivest/Rivest-rc5rev.pdf.
10. A. Michalski, K. Gaj, and T. El-Ghazawi, "An Implementation Comparison of an IDEA Encryption Cryptosystem on Two General-Purpose Reconfigurable Computers," *Proc. 13th Ann. Conf. Field-Programmable Logic and Applications*, LNCS 2778, Springer, 2003, pp. 204-219.

Tarek El-Ghazawi *is a professor in the Department of Computer and Electrical Engineering, a founder of the High-Performance Computing Lab (HPCL) at George Washington University, and cofounder of the NSF Center for High-Performance Reconfigurable Computing (CHREC). His research interests include high-performance computing, parallel computer architectures, high-performance I/O, reconfigurable computing, experimental performance evaluations, computer vision, and remote sensing. El-Ghazawi received a PhD in electrical and computer engineering from New Mexico State University. He is a senior member of the IEEE and a member of the ACM. Contact him at tarek@gwu.edu.*

Esam El-Araby *is a doctoral student in the Department of Computer and Electrical Engineering and a research assistant in the HPCL at George Washington University. His research interests include reconfigurable computing, hybrid architectures, evolvable hardware, performance evaluation, digital signal/image processing, and hyperspectral remote sensing. El-Araby received an MSc in computer engineering from the George Washington University. Contact him at esam@gwu.edu.*

Miaoqing Huang *is a doctoral student in the Department of Computer and Electrical Engineering and a research assistant in the HPCL at George Washington University. His research interests include reconfigurable computing, high-performance computing architectures, cryptography, image processing, and computer arithmetic. Huang received a BS in electronics and information systems from Fudan University, Shanghai. Contact him at mqhuang@gwu.edu.*

Kris Gaj *is an associate professor in the Department of Electrical and Computer Engineering and leads the Cryptographic Engineering Lab at George Mason University. His research interests include reconfigurable computing, cryptography, computer arithmetic, hardware description languages, and software-hardware codesign. Gaj received a PhD in electrical and computer engineering from Warsaw University of Technology, Warsaw. He is a member of the International Association for Cryptologic Research. Contact him at kgaj@gmu.edu.*

Volodymyr Kindratenko *is a senior research scientist in the Innovative Systems Laboratory at the National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign. His research interests include high-performance, reconfigurable, and scientific computing. Kindratenko received a DSc in analytical chemistry from the University of Antwerp, Antwerp, Belgium. He is a senior member of the IEEE and the ACM. Contact him at kindr@ncsa.uiuc.edu.*

Duncan Buell *is a professor and chair of the Department of Computer Science and Engineering at the University of South Carolina. His research interests include high-performance computing applications, parallel algorithms and architectures, computer security, computational number theory, information retrieval, and algorithm analysis. Buell received a PhD in mathematics from the University of Illinois at Chicago. He is a senior member of the IEEE Computer Society and a member of the ACM and the American Mathematical Society. Contact him at buell@engr.sc.edu.*

FEBRUARY 2008

Advertisers	Page Number
Cool Chips Symposium 2008	5
ETH Zurich	79, 80
Hot Chips Symposium 2008	33
IEEE Computer Society Membership	94-96
IRI 2008	9
ISVCS 2008	Cover 3
John Wiley & Sons, Inc.	Cover 2
Kuwait University	79
Microsoft	12
MobiQuitous 2008	Cover 3
Nanyang Technological University	80
Philips Electronics North America	78
Seapine Software	Cover 4
SPECTS 2008	51
Stony Brook University, SUNY	79
Classified Advertising	*78-80*

Boldface denotes advertisements in this issue.

Computer
IEEE Computer Society
10662 Los Vaqueros Circle
Los Alamitos, California 90720-1314
USA
Phone: +1 714 821 8380
Fax: +1 714 821 4010
http://www.computer.org
advertising@computer.org

Advertising Sales Representatives

Mid Atlantic (product/recruitment)
Dawn Becker
Phone: +1 732 772 0160
Fax: +1 732 772 0164
Email: db.ieeemedia@ieee.org

New England (product)
Jody Estabrook
Phone: +1 978 244 0192
Fax: +1 978 244 0103
Email: je.ieeemedia@ieee.org

New England (recruitment)
John Restchack
Phone: +1 212 419 7578
Fax: +1 212 419 7589
Email: j.restchack@ieee.org

Northwest (product)
Lori Kehoe
Phone: +1 650-458-3051
Fax: +1 650 458 3052
Email: l.kehoe@ieee.org

Southeast (recruitment)
Thomas M. Flynn
Phone: +1 770 645 2944
Fax: +1 770 993 4423
Email: flynntom@mindspring.com

Midwest (product)
Dave Jones
Phone: +1 708 442 5633
Fax: +1 708 442 7620
Email: dj.ieeemedia@ieee.org

Will Hamilton
Phone: +1 269 381 2156
Fax: +1 269 381 2556
Email: wh.ieeemedia@ieee.org

Joe DiNardo
Phone: +1 440 248 2456
Fax: +1 440 248 2594
Email: jd.ieeemedia@ieee.org

Midwest/Southwest (recruitment)
Darcy Giovingo
Phone: +1 847 498 4520
Fax: +1 847 498 5911
Email: dg.ieeemedia@ieee.org

Southwest (product)
Steve Loerch
Phone: +1 847 498 4520
Fax: +1 847 498 5911
Email: steve@didierandbroderick.com

Connecticut (product)
Stan Greenfield
Phone: +1 203 938 2418
Fax: +1 203 938 3211
Email: greenco@optonline.net

Southern CA (product)
Marshall Rubin
Phone: +1 818 888 2407
Fax: +1 818 888 4907
Email: mr.ieeemedia@ieee.org

Northwest/Southern CA (recruitment)
Tim Matteson
Phone: +1 310 836 4064
Fax: +1 310 836 4067
Email: tm.ieeemedia@ieee.org

Southeast (product)
Bill Holland
Phone: +1 770 435 6549
Fax: +1 770 435 0243
Email: hollandwfh@yahoo.com

Japan
Tim Matteson
Phone: +1 310 836 4064
Fax: +1 310 836 4067
Email: tm.ieeemedia@ieee.org

Europe (product/recruitment)
Hillary Turnbull
Phone: +44 (0) 1875 825700
Fax: +44 (0) 1875 825701
Email: impress@impressmedia.com

Advertising Personnel

Marion Delaney
IEEE Media, Advertising Director
Phone: +1 415 863 4717
Email: md.ieeemedia@ieee.org

Marian Anderson
Advertising Coordinator
Phone: +1 714 821 8380
Fax: +1 714 821 4010
Email: manderson@computer.org

Sandy Brown
IEEE Computer Society,
Business Development Manager
Phone: +1 714 821 8380
Fax: +1 714 821 4010
Email: sb.ieeemedia@ieee.org

THE HONG KONG POLYTECHNIC UNIVERSITY, Department of Computing. The Department invites applications for Assistant Professors in most areas of Computing, including but not limited to Software Engineering / Biometrics / Digital Entertainment / MIS and Pervasive Computing. Applicants should have a PhD degree in Computing or closely related fields, a strong commitment to excellence in teaching and research as well as a good research publication record. Initial appointment will be made on a fixed-term gratuity-bearing contract. Re-engagement thereafter is subject to mutual agreement. Remuneration package will be highly competitive. Applicants should state their current and expected salary in the application. Please submit your application via email to hrstaff@polyu.edu.hk. Application forms can be downloaded from http://www.polyu.edu.hk/hro/job.htm. Recruitment will continue until the positions are filled. Details of the University's Personal Information Collection Statement for recruitment can be found at http://www.polyu.edu.hk/hro/jobpics.htm.

UNIVERSITY OF CALGARY, Assistant and Associate Professors, Department of Computer Science. The Department of Computer Science at the University of Calgary seeks outstanding candidates for several tenure-track positions at the Assistant and Associate Professor levels. Of particular interest are applicants from information security, theory, computer games and information visualization or HCI. Applicants must possess a PhD in Computer Science or related discipline, and have strong potential to develop an excellent research record. Details for each position appear at: www.cpsc.ucalgary.ca/department/employ. The Department is one of Canada's leaders, as evidenced by our commitment to excellence in research and teaching. It has an expansive graduate program and extensive state-of-the-art computing facilities. Further information about the Department is available at www.cpsc.ucalgary.ca. Calgary is a multicultural city and the fastest growing city in Canada. Located beside the natural beauty of the Rocky Mountains, Calgary enjoys a moderate climate and outstanding year-round recreational opportunities. Interested applicants should send a CV, a concise description of their research area and program, a statement of teaching philosophy, and arrange to have at least three reference letters sent to: Dr. Ken Barker, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, T2N 1N4 or via email to: search@cpsc.ucalgary.ca. Applications will be reviewed immediately and will continue until the position is filled. All qualified candidates are encouraged to apply; however, Canadians and permanent residents will be given priority. The University of Calgary respects, appreciates, and encourages diversity. For more information on the University of Calgary and the city, please visit http://www.ucalgary.ca/hr/careers.

UNIVERSITY OF PUERTO RICO AT MAYAGUEZ, College of Engineering, Department of Electrical and Computer Engineering Doctoral Program in CISE. The Department of Electrical and Computer Engineering (ECE) of the University of Puerto Rico at Mayaguez (UPRM) invites applications for tenure-track positions in Computer Sciences and Engineering. The Department is interested in attracting faculty in the area of software engineering, databases or networking. Candidates are expected to make important contributions to research activities being conducted under the doctoral program in Computing and Information Sciences and Engineering (CISE), either by enhancing and strengthening current projects or starting new research tracks. Applicants must possess a PhD degree in computer science or computer engineering or closely related field, and demonstrate strong potential for excellence in research. If applicant's native language is not Spanish, applicant should be able to communicate fluently in English, and by the end of the tenure-track appointment be able to communicate effectively in Spanish. The ECE Department offers a Baccalaureate and Masters in Computer Engineering and the PhD in CISE jointly with the UPRM Department of Mathematics. For further information please visit www.ece.uprm.edu and www.phd.cise.uprm.edu or call 1-787-833-3338. Applications with a curriculum vitae, three reference letters, and MS and PhD original transcripts, should be sent no later than February 29, 2008, to: CISE Search Committee, University of Puerto Rico at Mayaguez, P.O. Box 5028, Mayagüez, PR 00681-5028. Or in pdf format to cisephd@ece.uprm.edu. Appointments will be made by July 1st, 2008. UPRM is an equal opportunity affirmative action employer.

SIEMENS PLM SOFTWARE INC., & its subsidiaries including Siemens Product Lifecycle Management Software II (US) Inc. have positions in software/implementation engineering, sales engineering, PLM and technical/software marketing, Applied Specialist, and Solutions Architect in various locations including Ames, IA, Plano, TX, Richardson, TX, Cypress, CA, San Diego, CA, San Jose, CA, State College, PA, Milford, OH, Shoreview, MN, Bloomington, IN, Ann Arbor, MI, Detroit, MI metro area and Seattle, WA metro area. Ph.D., Master's or Bachelor's degree w/ experience required bas ed on position. Send resumes to PLMCareers@ugs.com & list location of interest. Job code J08 must be referenced in email subject line. EOE.

HEWLETT – PACKARD COMPANY has an opportunity for the following position in Cupertino, California. Technology Consultant II Reqs. BS in Computer Science, Electrical Engineering or related and 2 yrs related exp. Reqs. Windows, Unix, Siebel, CRM and Outlook. Send resume referencing #CUPCCH. Please send resumes with reference number to Hewlett-Packard Company, 19483 Pruneridge Ave., MS 4206, Cupertino, CA 95014. No phone calls please. Must be legally

Engineering

The Philips organization has the following job opportunities available (various levels/types):

PHILIPS ELECTRONICS NORTH AMERICA
ANDOVER, MA
- Software Engineers (SWE-PENAC-MA)
- Test Engineers (TE-PENAC-MA)

ATLANTA, GA
- Software Engineers (SWE-PENAC-GA)
- Clinical Support Engineers (CSE-PENAC-GA)
- Technical Support Engineers (TSE-PENAC-GA)

BOTHELL, WA
- Software Engineers (SWE-PENAC-WA)
- Programmer Analysts (PA-PENAC-WA)

FOSTER CITY, CA
- Software Engineers (SWE-PENAC-CA)

ROSEMONT, IL
- Design Engineers (DE-PENAC-IL)

SAN JUAN, PUERTO RICO
- Field Service Engineers (FSE-PENAC-PR)

PHILIPS ORAL HEALTHCARE
SNOQUALMIE, WA
- Software Engineers (SWEPOHC-WA))
- Commodity Manager (CM-POHC-WA)
- Process Engineer (PE-POHC-WA)

PHILIPS ULTRASOUND
BOTHELL and SEATTLE, WA
- Software Engineers (SWE-PU-WA)
- Design Engineers (DE-PU-WA)
- Test Engineers (TE-PU-WA)
- Programmer Analysts (PA-PU-WA)

PHILIPS MEDICAL SYSTEMS CLEVELAND
HIGHLAND HEIGHTS, OH
- Software Engineers (SWE-PMSC-OH)
- Design Engineers (DE-PMSC-OH)
- Engineer (E-PMSC-OH)

PHILIPS NUCLEAR MEDICINE
MILPITAS, CA
- Software Engineer (SWE-PNM-CA)

Some positions may require travel. Submit resume by mail to PO Box 4104, Santa Clara, CA 95056-4104. **Must reference job title and job code (i.e. SWE-PENAC-CA) in order to be considered**. EOE.

SUBMISSION DETAILS: Rates are $299.00 per column inch ($320 minimum). Eight lines per column inch and average five typeset words per line. Send copy at least one month prior to publication date to: Marian Anderson, Classified Advertising, *Computer* Magazine, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; (714) 821-8380; fax (714) 821-4010. *Email: manderson@computer.org.*

authorized to work in the U.S. without sponsorship. EOE.

HEWLETT – PACKARD COMPANY has an opportunity for the following position in New York, NY. Business Consulting IV. Reqs. exp. with SQL, Sales & Marketing, Profit & Loss, ROI, Net Present Value Analysis, & Risk Mgt. Reqs. incl. Master's degree or foreign equiv. in MIS, Business or related & 1 year of related exp. Send resume & refer to job #NYCSHA. Please send resumes with job number to Hewlett-Packard Company, 19483 Pruneridge Ave., MS 4206, Cupertino, CA 95014. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

NETWORK AND COMPUTER SYSTEMS ADMINISTRATOR: Valley Cottage, NY. Bach deg in Comp Science (or equiv), ext Windows 2000/SP/2003 install, TCP/IP network, SQL server 2000 & MS 2000, Ciscos switches, PIX firewalls, & routers admin; ARCServe & NT backup. Fax resumes to CreditRiskMonitor, Attn L Fensterstock at (845) 267-4110.

OPERATION SYSTEM ANALYST. NY, NY. Analyze IT development of company: operation/system & infrastructure dev. & integration, system analysis & program-

KUWAIT UNIVERSITY
DEPARTMENT OF INFORMATION SCIENCE
FACULTY POSITION ANNOUNCEMENT

The Department of Information Science at College for Women, Kuwait University, invites applications for faculty positions starting February 2008 in Information Science & Technology, Information Systems, Database Systems, Telecommunication Networks & Security, Data Mining and Web Development & Multimedia or related areas.

Preferences will be given to those applying for Associate and Full Professor ranks. Applications for short time visiting professor appointment will also be considered. The medium of instruction is English. Responsibilities include teaching undergraduate courses, conducting scholarly research, and carrying administrative duties. Both **male and female** candidate are invited to apply.

Qualifications include an earned **Ph.D.** from a reputed western university in the area of specialization or related fields. The candidate must also demonstrate evidence of quality teaching and research and have full command of English. To be considered for an Associate or professor level, the candidate must have a strong publication record in refereed international journals.

The College for Women, *www.cfw.kuniv.edu*, is part of Kuwait University which is one of the leading public institutions of higher education in the Gulf region. Kuwait University offers a generous benefit package that includes competitive tax-free salary, annual air tickets, tuition allowance for children schooling, a one-time settling-in allowance, housing allowance, free national health care, paid mid-year holidays and summer vacation, and an end-of-contract gratuity. The University offers an excellent academic environment and financial research support.

To apply, please submit a completed Application Package. Detailed information can be found at http://www.cfw.kuniv.edu/about/positions.htm. All communication should be addressed to:

ISC Department Head, College for Women, Kuwait University, P.O. Box 5969, Safat 13060, Kuwait, Phone:(+965) 498 3095, Fax: +(965) 251-4252, e-mail: ISChead@cfw.kuniv.edu

Faculty Positions

Stony Brook University's Department of Electrical and Computer Engineering is seeking outstanding candidates for tenure-track faculty positions for Fall 2008. Strong candidates in all areas will be considered, but we are particularly interested in receiving applications from highly qualified candidates who can drive new research initiatives by supplementing our existing strengths in the area of VLSI, computer systems, and wireless communications and networking, as well as in interdisciplinary areas or whose research bridges traditional academic boundaries. The successful candidate must have demonstrated excellence in research and the ability to develop an independent research program, and have a strong interest in teaching at both the undergraduate and graduate levels. Detailed information on the research activities of these groups can be found on the Department Web site at *http://www.ece.sunysb.edu*. The Department is in a stage of significant expansion along with the New York State Center of Excellence in Wireless and Information Technology (CEWIT), which is currently under construction at the University R&D Park. The Department is affiliated with the University's newly established clusters focusing on multi-disciplinary studies in wireless and information technology, high-performance computing, computational neuroscience, and digital media. The Department is also associated with the Computational Science Center at the neighboring Brookhaven National Laboratory and the NY Sensor CAT. Stony Brook enjoys close proximity to both New York City and Long Island's majestic ocean beaches. Its school districts are highly ranked nationally. Opportunities for industrial collaborations abound with many high-profile IT companies close by. Moreover, the Department of Electrical and Computer Engineering offers a congenial and productive working environment. **Required:** Applicants must have a Ph.D. in Electrical and/or Computer Engineering, or a related discipline. Review of applications will begin soon and will continue until the positions are filled.

To apply online visit *www.stonybrook.edu/jobs* or send a detailed résumé, the names of at least three references, and three publications to: Sangjin Hong, Chair of Faculty Recruiting Committee
Department of Electrical and Computer Engineering
Light Engineering Building, Room 273
Stony Brook University, SUNY
Stony Brook, NY 11794-2350

Equal Opportunity/Affirmative Action Employer.

STONY BROOK

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Professor for Computer Science / Informatics (Computer Systems)

ETH Zurich invites applications for a full professor position in Computer Science/Informatics (Computer Systems). The successful applicant is someone who bridges the gap between software systems and hardware systems. Applications are solicited from candidates with internationally recognized research credentials and proven teaching ability. He or she should be able to teach computer architecture, system design, performance evaluation, digital design, and system programming in the B.S. and M.S. program. Courses at Master level may be taught in English. We expect a well established track record in research. The candidate's research focus can be on Performance Evaluation, Embedded and Mobile Systems, Multiprocessor/Multicore Systems, Dependable Systems, Fault Tolerant Computing, Real-time Systems, Custom Computing, Support for Streaming and Continuous Media, or System Architecture.

Please submit your application together with a curriculum vitae, a list of publications, the names of at least three referees, and a short overview of the research interests to the **President of ETH Zurich, Prof. Dr. R. Eichler, Raemistrasse 101, ETH Zurich, 8092 Zurich, Switzerland, no later than April 30, 2008**. For further information, candidates may contact the Head of the Department, Prof. J. Gutknecht (gutknecht@inf.ethz.ch). With a view toward increasing the number of female professors, ETH Zurich specifically encourages qualified female candidates to apply.

ming, tech support & telecommunications. BS/MS in related field or foreign equiv & related exp. RES: Consulate Hotel Associates, 200 W 55th St, Ste 42, NY, NY 10019.

NOKIA SIEMENS NETWORKS US LLC has the following exp/degree position at its Irving, Texas location. Travel to unanticipated U.S. worksites may be required. Equal Opportunity Employer. IRVING, TEXAS: *Specialist: Perform Wide Code Division Multiple Access Base Transceiver Station system testing, troubleshooting and reporting using test tools and equipment. ID# NSN-TX-SPEC. Mail resume to: NSN Recruiter, Nokia Siemens Networks, 6000 Connection Dr., 4E-338, Irving, TX 75039. MUST REFERENCE ID #.

CISCO SYSTEMS, INC is accepting resumes for the following positions: CALIFORNIA, San Jose/Milpitas/Santa Clara: Finance Business Manager, Treasury (Ref# SJ54IC), Finance Business Manager (Ref# SJ28IC), Manager, Technical Marketing (Ref# SJ56IC), Network Engineer (Ref# SJ57IC), Manufacturing Quality Engineer (Ref# SJ58IC), Marketing Manager (Ref# SJ59IC), Web Program/Specialist (Ref# SJ60IC), Program Manager (Ref# SJ22IC). ILLINOIS, Chicago: Business Development Manager (Ref# CHI2IC). NEW YORK, New York: Channel Systems Engineer (Ref# NY6IC). OHIO, Richfield: Hardware Engineer (Ref# OH2IC). TEXAS, Austin: Software Engineer (Ref# AUS2IC). Please mail resumes with job reference number to Cisco Systems, Inc., Attn: Jasbir Walsh, 170 W. Tasman Drive, Mail Stop: SJC 5/1/4, San Jose, CA 95134. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE. www.cisco.com

COMPUTER & INFORMATION SYSTEMS MANAGER (NYC) Manage backup, security & user help systems. Consult w/users, mgmt, vendors & technicians to assess computing need & system reqmts. Make modifications/corrections to s/ware. BS/MS in related field & related exp. Res: Neil A. Weinrib & Associates, 305 Broadway, Ste 1002, NY, NY 10007.

HEWLETT – PACKARD COMPANY has an opportunity for the following position in Cupertino, CA. Software Designer II. Reqs. exp. with Java, J2ee, Junit, SQL, PL/SQL, XML, Ant, Eclipse, Webservers, Application Servers, Software Development Life Cycle for an enterprise application, Oracle, industry standard Java Software Design Patterns, working w/ at least 1 bug-tracking system; exp. Integration w/ 3rd party libraries pref. Open source (Apache/Source Forge). Reqs. incl. Master's degree or foreign equiv. in CS, CE, EE or related. Send resume & refer to job #CUPSPA. Please send resumes with job number to Hewlett-Packard Company, 19483 Pruneridge Ave., MS 4206, Cupertino, CA 95014. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

HEWLETT – PACKARD COMPANY has an opportunity for the following position in Cupertino, CA. Software Designer. Reqs. exp. with Java and Python, programming languages; networking technology & protocols incl. TCP/IP, HTTP, SSL, XML/RPC, LDAP, DHCP & Dykstra's routing algorithms; UNIX file systems & shell programming; Database technology incl. JDBC; SQL programming & performance tuning; security concepts incl. RBAC, authentication authorization protocols; and ability to design clean, usable APIs. Reqs. incl. Bachelor's degree in CS/Engineering or related & 6 yrs. exp. in job offered or related. Send resume & refer to job #SUNKGR. Please send resumes with job number to Hewlett-Packard Company, 19483 Pruneridge Ave., MS 4206, Cupertino, CA 95014. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

NANYANG TECHNOLOGICAL UNIVERSITY

NUS National University of Singapore

LEE KUAN YEW POSTDOCTORAL FELLOWSHIP

Applications are invited from young and outstanding academics for the prestigious Lee Kuan Yew Postdoctoral Fellowship (LKY PDF) in the National University of Singapore (NUS) and Nanyang Technological University (NTU).

The Fellowship is tenable for up to 3 years, with possible extension for 2 further years. LKY PDFs can apply for academic positions following the Fellowship.

Gross annual salary ranges from S$72,000 to S$144,000 (approx US$50,000– US$100,125) with commencing salary depending on qualifications and experience. Leave and medical benefits will be provided. For details of other benefits offered and application procedure, please visit the websites of the respective University*.

APPLICATION

Interested candidates should send their complete application package, comprising all documents listed below, to the respective University *

- NTU Application or NUS Personal Particulars Form (downloadable from website)
- Detailed Curriculum Vitae, List of Publications & Educational Certificates
- 3 International Referee Reports (including contact details)
- Statement of Research Intent (details of proposed research plan)

Closing date: 5 March 2008

Successful candidates will be notified in June 2008

* For application and contact details, please see
NUS: http://www.nus.edu.sg/ore/fellowships/fellowship_lky.htm
NTU: http://www.ntu.edu.sg/hr/recruit/research/LKY2008.htm

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Assistant Professor of Theoretical Computer Science

ETH Zurich invites applications by candidates with internationally recognized research credentials for an assistant professor position in Theoretical Computer Science within the Department of Computer Science (Informatics). Duties of the new professor include research and teaching in Theoretical Computer Science and its applications. The successful candidate should preferably have an expertise in an area that complements those currently present at ETH Zurich, like e.g. formal methods, semantics, alternative computing paradigms, or computer algebra. Assistant professors have the opportunity of performing independent research in an attractive environment and with their own research group. All professors are expected to participate in teaching introductory computer science courses.

This assistant professorship has been established to promote the careers of younger scientists. Initial appointment is for four years, with the possibility of renewal for an additional two-year period.

Please submit your application together with a curriculum vitae, a list of publications, the names of at least three referees, and a short overview of the research interests to the **President of ETH Zurich, Prof. Dr. R. Eichler, Raemistrasse 101, ETH Zurich, 8092 Zurich, Switzerland, no later than April 30, 2008**. For further information, candidates may contact the Head of the Department, Prof. J. Gutknecht (gutknecht@inf.ethz.ch). With a view toward increasing the number of female professors, ETH Zurich specifically encourages qualified female candidates to apply.

BOOKSHELF

CMMI and Six Sigma: Partners in Process Improvement, Jeannine M. Siviy, M. Lynn Penn, and Robert W. Stoddard. This book focuses on the synergistic, rather than competitive, implementation of CMMI and Six Sigma. Topics range from formation of the value proposition to specific implementation tactics. The authors show how not taking advantage of what both initiatives offer puts an organization at risk of sinking time, energy, and money into inventing a solution that already exists. Along the way they debunk a few myths about Six Sigma applications in software.

Although the authors concentrate on the interoperability of Six Sigma and CMMI, they also recognize that organizations rarely implement only these two initiatives. Accordingly, the discussion turns to the emerging realm of multimodel process improvement, as well as strategies and tactics that transcend models to help organizations effectively knit together a single unified internal process standard.

Addison-Wesley Professional; informit.com; 0-321-51608-7; 368 pp.

Algorithmic Game Theory, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, eds. In the past few years game theory has substantially influenced computer science, especially Internet- and e-commerce-related issues. More than 40 top researchers in this field have contributed chapters to this book that go from game theory's foundations to its state-of-the-art applications.

Early chapters cover algorithmic methods for equilibria, mechanism design, and combinatorial auctions, while later chapters address incentives and pricing, cost sharing, information markets, and cryptography and security.

Cambridge University Press; www.cambridge.org; 978-0-521-87282-9; 776 pp.

Advances in Biometrics: Sensors, Algorithms and Systems, Nalini K. Ratha and Venu Govindaraju, eds. Biometrics technology continues to progress, with its wider acceptance and the need for various new security facets in modern society. From simply logging on to a laptop to crossing a national border, biometrics is being called upon to meet the growing challenges of identity management.

With contributions by leading international authorities in the field, this book presents a comprehensive treatment of biometrics that covers the entire gamut of topics in the field, including data acquisition, pattern-matching algorithms, and system-level issues such as standards, security, networks, and databases. Organized into three sections, the book covers: sensors, advances in biometric matching algorithms, and topics that deal with issues at the systems level.

Springer; www.springer.com; 978-1-84628-920-0; 508 pp.

Networking with Microsoft Windows Vista: Your Guide to Easy and Secure Windows Vista Networking, Paul McFedries. The author provides a beginner's guide to creating, configuring, administering, and using a small network with Windows Vista computers. Coverage spans networking hardware, including Ethernet hardware (from NICs to cables to switches to routers) and wireless hardware (from wireless NICs to access points to range extenders).

The book includes buyer's guides that help readers make smart choices when purchasing network hardware, then shows readers how to put everything together, including configuring a router, laying cable, and connecting the devices. Mastering Windows Vista's networking features involves using the Network and Sharing Center, managing wired and wireless connections, accessing shared network resources, sharing local resources on the network, and working with network files offline.

The book also provides extensive coverage of security issues that affect anyone connected to the Internet. The author shows how to secure each computer, secure global networking settings, and batten down wireless connections. The book also includes intermediate networking tasks such as making remote connections, monitoring the network, troubleshooting network problems, and setting up Vista's built-in Web and FTP servers.

QUE; www.informit.com/que; 0-7897-3777-9; 552 pp.

Conversational Informatics: An Engineering Approach, Toyoaki Nishida, ed. This book investigates human behavior with a view to designing conversational artifacts capable of interacting with humans. It spans an array of topics, including linguistics, psychology, and human-computer interaction. Until recently, research in such areas has been carried out in isolation, with no attempt made to connect the various disciplines. Advancements in science and technology have changed this.

This book provides an interdisciplinary introduction to conversational informatics and places emphasis on the integration of scientific approaches to achieve engineering goals and advance further understanding of conversation. A collection of surveys explores four prominent research areas: conversational artifacts; conversational contents; conversation environment design; and conversation measurement, analysis, and modeling.

Wiley; www.wiley.com; 978-0-470-02699-1; 430 pp.

Send book announcements to newbooks@computer.org.

COMPUTER SOCIETY CONNECTION

Top Researchers Win Technical Achievement Award

The IEEE Computer Society sponsors an active and prestigious awards program as part of its mission to promote the free exchange of ideas among computer professionals around the world and to recognize its members for their outstanding accomplishments. The awards honor technical achievements as well as service to the computer profession and to the Society.

TECHNICAL ACHIEVEMENT AWARD

Each year, the IEEE Computer Society selects several individuals to receive its Technical Achievement Award. This award recognizes outstanding and innovative contributions to the fields of computer and information science and engineering or computer technology, usually within the past 10, and not more than 15, years. Each winner receives a certificate and a $2,000 honorarium.

2006 Winners

The following four award recipients were selected for their contributions to the technical advancement of the computer science and engineering fields.

Hsinchun Chen, of the University of Arizona, was chosen "for innovative contributions to digital libraries, medical informatics, and intelligence and security informatics."

Michael T. Goodrich, of the University of California, Irvine, was cited "for outstanding contributions to the design of parallel and distributed algorithms for fundamental combinatorial and geometric problems."

Shashi Shekhar, of the University of Minnesota, gained honors "for technical achievements in spatial databases, spatial data mining, and geographic information systems."

Roberto Tamassia, of Brown University, was recognized "for pioneering the field of graph drawing and for outstanding contributions to the design of graph and geometric algorithms."

For more information about these and other IEEE Computer Society awards, including nomination forms and guidelines, visit www.computer.org/awards.

Society Debuts Career-Development Initiative

The IEEE Computer Society has launched a career site intended to help professionals in the computing field navigate the rapid technology advances, globalization, shifting demographics, and new business approaches that will dramatically change the workforce over the next decade.

Build Your Career (www.computer.org/buildyourcareer) is designed to give technologists practical, affordable guidance to benefit their job prospects. The site is envisioned as a one-stop shop for those who are either entering the field, looking for a new job, or interested in advancing with their current employer. Users can get quickly up to speed on a broad spectrum of topics with TechSets, article packages that provide comprehensive knowledge on subjects in the fields of software and programming, security and privacy, networking, and wireless, Web, and management technologies.

TechSets are compiled by such well-known Computer Society experts as Cisco engineering manager Wes Chou, MITRE Corp.'s Susan (Kathy) Land, CSDP, the Society's 2008 president-elect; Pennsylvania State University software engineering professor Phillip Laplante, an author and editor of 22 books; and University of British Columbia professor Philippe Kruchten, an expert on the Rational Unified Process.

Besides career-related technical articles, the site features online technical courses, training aids, jobs boards, career news, and columns that address current industry issues. In partnership with the Computer Society, Harvard Business School Publishing is making its *Harvard Business Review* articles, Harvard Business School Press books, and other content available to Build Your Career visitors.

IEEE Names 2008 Fellows

The IEEE Board of Directors recently conferred the title of Fellow upon 295 senior members of the IEEE, including 62 Computer Society members, who have demonstrated outstanding achievement in engineering. The original 1912 constitution of the American Institute of Electrical Engineers, a forerunner of the IEEE, outlined a procedure for naming Fellows. Today, Fellow status recognizes a person who has established an extraordinary record of achievements in any of the IEEE fields of interest.

The honorees are selected from among the more than 350,000 IEEE members. Since IEEE policy limits the number of Fellows selected each year to no more than 0.10 percent of the IEEE's total voting membership, this year's cohort of 295 new Fellows is an especially select group of outstanding individuals.

The names below include both new Fellows who are Computer Society members and other IEEE members who the Computer Society recommended for elevation to Fellow status. All are now IEEE Fellows, effective 1 January. An accompanying citation details the accomplishments of each new Fellow.

A

Ishfaq Ahmad, University of Texas at Arlington, for contributions to scheduling techniques in parallel and distributed computing systems.

Rajeev Alur, University of Pennsylvania, for contributions to automata, logics, and verification techniques for real-time and hybrid systems.

Cleon Anderson, L-3 Communications, for contributions to servo control systems.

B

Victor (Paramvir) Bahl, Microsoft Research, for contributions to the design of wireless networks and systems and leadership in mobile computing and communications.

John Bay, Air Force Research Laboratory, for leadership in model-based design and integration methods for large-scale embedded systems.

Bernd Becker, University of Freiburg, for contributions to the development of algorithms and data structures for testing and verification of integrated circuits.

Andrew Blake, Microsoft, for contributions to the foundations of segmentation and tracking, and innovation in vision applications.

Gunilla Borgefors, Swedish University, for contributions to discrete geometry and image analysis.

Ronald Brachman, Yahoo!, for leadership in knowledge representation and reasoning in computer science and artificial intelligence.

Joe Brewer, IEEE, for contributions to nonvolatile memory integrated circuit technology and digital signal processor architecture.

Michael Bushnell, Rutgers University, for contributions to testing methods for digital and mixed-signal VLSI circuits.

C

Manuel Castro, Spanish University for Distance Education, for contributions to distance learning in electrical and computer engineering education.

Krishnendu Chakrabarty, Duke University, for contributions to the testing of core-based system-on-chip integrated circuits.

Peter Chen, University of Michigan, for contributions to fault-tolerant storage systems.

George Chiu, IBM, for leadership in supercomputer technology.

Jen-Yao Chung, IBM, for contributions to the application of electronic business and Web-based information systems.

Pau-Choo Chung, National Cheng Kung University, for contributions to neural network models for biomedical image analyses.

Peter Corke, Commonwealth Scientific Industrial Research Organization, for contributions to visual-based robot control and its applications to field robotics.

D

Nikil Dutt, University of California, Irvine, for contributions to architecture description languages for the design and exploration of customized processors.

G

Guang Gao, University of Delaware, for contributions to architecture and compiler technology of parallel computers.

Daniel Gamota, Motorola, for leadership in nanotechnology-based printed electronic products.

Michael Gschwind, IBM, for contributions to high-performance computer architecture and compilation technology.

Rajiv Gupta, University of Arizona, for contributions to computer architecture and optimizing compilers.

H

Dong Ha, Virginia Polytechnic Institute and State University, for leadership in VLSI design and test.

Kazuo Hagimoto, Nippon Telegraph and Telephone, for contributions to very large capacity optical transmission systems.

Jennifer Hou, University of Illinois at Urbana-Champaign, for contributions to protocol design and analysis of wireless communications networks.

J

Ravi Jain, Google, for contributions to wireless networks and standard programmable interfaces for converged networks.

Christian Jensen, Aalborg University, for contributions to temporal, spatiotemporal, and mobile data management.

Christian Jutten, University Joseph Fourier, for contributions in source separation and independent component analysis.

K

Hisao Kameda, University of Tsukuba, for contributions to performance-optimization methods for information processing systems.

L

Phillip Laplante, Pennsylvania State University, for leadership in engineering education and for contributions to software and systems education.

Jorg Liebeherr, University of Toronto, for contributions to the design and analysis of computer networks and their protocols.

Steven Low, California Institute of Technology, for contributions to Internet congestion control.

M

Benoit Macq, Université Catholique de Louvain, for contributions to visual communication technologies.

Vijay Madisetti, Georgia Institute of Technology, for contributions to embedded computing systems.

Roy Maxion, Carnegie Mellon University, for contributions to real-time monitoring and analysis of computer systems for trend analysis, fault prediction, and anomaly detection.

Samiha Mourad, Santa Clara University, for contributions to fault modeling in digital circuits and systems.

N

Klara Nahrstedt, University of Illinois at Urbana-Champaign, for contributions to end-to-end, quality-of-service management of multimedia systems.

Chandrasekhar Narayanaswami, IBM, for contributions to pervasive computing systems.

Sani Nassif, IBM, for contributions to semiconductor manufacturing processes.

David Notkin, University of Washington, for contributions to software engineering and software evolution research.

O

Oyekunle Olukotun, Stanford University, for contributions to multiprocessors on a chip and multithreaded processor design.

P

Dhabaleswar Panda, Ohio State University, for contributions to high-performance and scalable communication in parallel and high-end computing systems.

Karen Panetta, Tufts University, for leadership in engineering education and curriculum development to attract, retain, and advance women in engineering

Fernando Pereira, Instituto Superior Técnico, for contributions to object-based digital video representation technologies and standards.

R

Sanguthevar Rajasekaran, University of Connecticut, for contributions to sequential, parallel, and randomized algorithms and to bioinformatics.

Raghu Ramakrishnan, Yahoo!, for contributions to deductive databases, data mining, and query optimization.

Nageswara Rao, Oak Ridge National Laboratory, for contributions to algorithms for reliable communication in distributed sensor networks.

Ronny Ronen, Intel, for leadership in microarchitecture, low-power design, and compilers for high-performance superscalar microprocessors.

S

Yvon Savaria, University of Montreal, for contributions to the development of long interconnect VLSI signal-processing architectures.

Jyuo-Min Shyu, Industrial Technology Research Institute, for leadership in the microelectronics industry.

Mani Srivastava, University of California, Los Angeles, for contributions to energy-aware wireless communications and sensor networking.

Ivan Stojmenovic, University of Ottawa, for contributions to data communication algorithms and protocols for wireless sensor and ad hoc networks.

T

Sergios Theodoridis, University of Athens, for contributions to the design of adaptive signal processing systems.

Anand Tripathi, University of Minnesota, for contributions to distributed system software architectures and programming frameworks

U

Lalita Udpa, Michigan State University, for contributions to development of forward and inverse electromagnetic nondestructive evaluation methodologies.

V

Paulo Verissimo, University of Lisbon, for contributions to dependable and secure distributed computing.

W

Laung-Terng Wang, SynTest Technologies, for leadership in practical design for test of integrated circuits.

Jacob White, Massachusetts Institute of Technology, for contributions to simulation tools for RF circuits, electrical interconnects, and micromachined devices.

Ja-Ling Wu, National Taiwan University, for contributions to image and video analysis, coding, digital watermarking, and rights management.

Y

Rajendra Yavatkar, Intel, for contributions to network protocols and multiprocessor systems-on-a-chip for wire-speed packet processing.

Hoi-Jun Yoo, Korea Advanced Institute of Science and Technology, for contributions to low-power and high-speed VLSI design.

Z

Alexander Zelinsky, Commonwealth Scientific Industrial Research Organization, for contributions to vision-based robotics.

For more information on IEEE Fellows, see the related story on IEEE Fellow nominations in this issue of *Computer*, or visit www.ieee.org/fellows.

Obituary: James P. Anderson Jr.

On 18 November 2007, noted computer pioneer James P. Anderson, Jr., 77, died at his home in Pennsylvania. Anderson first introduced the notion of intrusion detection in 1980 with his seminal paper, "Computer Security Threat Monitoring and Surveillance" (US Air Force, 1972). Widely known as "The Anderson Report," that paper defined the research agenda in information security for well over a decade.

After graduating from Pennsylvania State University with a degree in meteorology, Anderson served in the US Navy as a gunnery officer and as a radio officer. Later, he worked at Univac under noted computer pioneer John Mauchly. Anderson subsequently joined Burroughs, where he explored issues of compilation, parallel computing, and computer security. At Burroughs, he received a patent for one of the earliest multiprocessor systems, the D-825. From the late 1960s until 2007, Anderson maintained an independent consulting firm.

In 1968, he served on the Defense Science Board Task Force on Computer Security, which produced the "Ware Report," defining the technical challenges of computer security. In 1990, Anderson received the prestigious NIST/NCSC National Computer Systems Security Award. He also figured prominently in the development of more than 200 other seminal standards, policies, and reports including Blacker and the Trusted Computer System Evaluation Criteria or "Orange Book" (US National Security Agency, 1983.)

IEEE/ACM TCBB Seeks Editor in Chief for 2009-2010 Term

The IEEE Computer Society seeks applicants for the position of editor in chief of *IEEE/ACM Transactions on Computational Biology and Bioinformatics* for a two-year term starting **1 January 2009**.

REQUIREMENTS AND QUALIFICATIONS

Candidates for any Computer Society editor in chief position should possess a good understanding of industry, academic, and government aspects of the specific publication's field. In addition, candidates must demonstrate the managerial skills necessary to process manuscripts through the editorial cycle in a timely fashion. An editor in chief must be able to attract respected experts to his or her editorial board.

Applicants, with clear employer support, must possess recognized expertise in the computer science and engineering community, have editorial experience, and be able to work effectively with technical and publishing professionals.

CANDIDATE SEARCH

Prospective candidates are asked to provide, by **15 March**, a complete curriculum vitae, a brief plan for the publication's future, and a letter of support from their institution or employer. Materials should be sent as PDF files to staff liaison Alicia Stickley at astickley@computer.org.

NEW TRANSACTIONS EDITORS IN CHIEF

Two other IEEE Computer Society transactions have editors in chief who are beginning an initial two-year term in 2008.

Wolfgang Nejdl, of the University of Hannover, now directs the recently launched *IEEE Transactions on Learning Technologies*. Mani Srivastava, of the University of California, Los Angeles, now heads *IEEE Transactions on Mobile Computing*.

IEEE Fellow Nominations Due 1 March

André Ivanov, Chair, 2008 Computer Society Fellows Committee

The IEEE and its member societies cooperate each year to select a small group of outstanding professionals for recognition as IEEE Fellows. A senior IEEE member who has achieved distinction in his or her field can be named an IEEE Fellow only after being nominated for the honor. All such nominations undergo rigorous review before the IEEE Board of Governors votes to bestow the prestigious rank of Fellow.

For information regarding nominating a candidate for IEEE Fellow recognition, visit www.ieee.org/fellows. The Electronic Fellow Nomination Process is detailed at www.ieee.org/web/membership/fellows/index.html.

The deadline for Fellow nominations is **1 March**. In the event that the online nomination process is unsuitable, paper nomination materials can be obtained from the IEEE Fellow Committee, 445 Hoes Lane, PO Box 1331, Piscataway, NJ 08855-1331; voice +1 732 562 3840; fax +1 732 981 9019. Hard copies can also be obtained by request from fellow-kit@ieee.org. Nominators may not submit the forms via fax or email.

Nominees

A nominee must be a senior member at the time of nomination and must have been an IEEE member at any level for the previous five years. This includes exchange, student, associate, senior, honorary, and life membership levels. It excludes affiliates, however, because this category does not comprise IEEE members. The five-year requirement must be satisfied at the date of election, 1 January 2009; thus, a nominee must have been a member at any level continuously since 31 December 2003. The five-year membership requirement may be waived in the case of nominees in Regions 8, 9, and 10. Fellows are never named posthumously.

Nominators

A nominator need not be an IEEE member. However, nominators cannot be IEEE staff or members of the IEEE Board of Directors, the Fellows Committee, the technical society, or council evaluation committee.

Preparing a nomination

Essential to a successful nomination is a concise account of a nominee's accomplishments, with emphasis on the most significant contribution. The nominator should identify the IEEE society or council that can best evaluate the nominee's work and must send the nomination form to the point of contact for that group.

Careful preparation is important. Endorsements from IEEE entities such as sections, chapters, and committees, and from non-IEEE entities and non-IEEE individuals are optional but might be useful when these entities or individuals are in the best position to provide credible statements.

References

The nominator should select referrers who are familiar with the nominee's contributions and can provide insights into these achievements. For nominees in the US and Canada, references must be from IEEE Fellows; outside the US and Canada, senior members can provide references if necessary. References cannot come from IEEE staff or from members of the IEEE Board of Directors, the Fellows Committee, a technical society, or a council evaluation committee. While a minimum of five references are needed, it is strongly recommended that the maximum of eight be sought.

Evaluation of nominees

In evaluating nominations, the IEEE Fellow Committee considers the following criteria:

- individual contributions as a research engineer or scientist, application engineer or practitioner, technical leader, or educator;
- technical evaluation by an IEEE society or council;
- tangible and verifiable evidence of technical accomplishment, such as technical publications, patents, reports, published product descriptions or services, as listed on the nomination form;
- confidential opinions of referrers who can attest to the nominee's work;
- IEEE and non-IEEE professional activities, including awards, services, and offices held, committee memberships, and the like; and
- total years in the profession.

Resubmission of nominations

Typically, less than half of the nominations each year are successful. Therefore, even highly qualified individuals might not succeed the first time. Because reconsideration of a nominee is not automatic, nominators are encouraged to update and resubmit nominations for unsuccessful candidates. To resubmit these materials, ensure that the nomination forms are current. The deadline for resubmission is the same as for new nominations.

Nomination deadline

The IEEE Fellow Committee must receive 2008 nomination forms by **1 March**. The staff secretary, Pamela Kemper (pkemper@computer.org), must also receive at least five Fellow-grade reference letters directly from the referrers by that date. The deadline will be strictly enforced.

Computer Society Board Announces 2008 Meetings and Election Schedule

The IEEE Computer Society has released its official 2008 administrative schedule. Highlights include the annual Board of Governors administrative meeting series that serves as a fixed point around which many other deadlines are scheduled. In a move to control costs, Computer Society administrative meetings have been scaled back from three weeks to one week during the year. The balance of administrative deliberations throughout the year will take place via teleconference.

The 2008 calendar includes significant dates in the 2008 election and governance cycle. The 6 October election will name the 2009 first and second vice presidents; the 2009 president-elect, who will serve as president in 2010; and seven members of the Board of Governors, who serve three-year terms. Officers selected in the 2008 elections begin their terms on 1 January 2009.

Nominations

The Nominations Committee must receive recommendations for candidates in this year's election no later than **4 April**. Recommendations must be accompanied by the nominee's biographical information, which should include facts about past and present participation in Society activities. Nomination materials should be sent to Michael Williams, Nominations Committee Chair, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1992; voice +1 202 371 0101; fax +1 202 296 6896; m.williams@computer.org.

2008 Schedule

Member participation and volunteer involvement are welcomed throughout the year. The following calendar lists dates of note for both Computer Society and IEEE election materials in 2008.

- **8 February:** Board of Governors teleconference
- **4 April:** Recommendations from membership for board/officer nominees due to Nominations Committee
- **18 April:** Nominations Committee slate of officer and board candidates due to Board of Governors
- **6 May:** Last day for board/officer petition candidates to be submitted to Board secretary
- **12-16 May:** First Board of Governors meeting, Rio Hotel, Las Vegas
- **16 May:** Periodical page budgets, prices, and board/officer candidates approved at Board of Governors meeting, Rio Hotel, Las Vegas
- **13 June:** Last day for 2009 IEEE Division V Delegate/Director-Elect petition candidates to be submitted to IEEE
- **27 June:** Board-selected candidate statements, biographies, and pictures due in the Publications Office
- **July:** Board-approved slate and call for petition candidates published in *Computer*
- **31 July:** Last day for members to submit board/officer petition candidates (with statements, biographies, and pictures)
- **August:** Schedule and call for 2010 IEEE Division VIII Delegate-Director-Elect recommendations to Nominations Committee
- **4 August:** Executive Committee teleconference
- **5 August:** Board of Governors teleconference
- **5 August:** Ballots mailed
- **September:** Candidate statements/biographies published in *Computer*
- **6 October:** Ballots returned and tabulated
- **21 October:** Nominations Committee makes recommendations to Board of Governors for 2010 IEEE Division VIII Delegate-Director-Elect
- **17-18 November:** Second Board of Governors meeting, Hyatt Regency, New Brunswick, New Jersey
- **18 November:** 2010 IEEE Division VIII Delegate/Director-Elect slate approved at Board of Governors meeting
- **December:** Election results published in *Computer*

CALLS FOR ARTICLES FOR IEEE CS PUBLICATIONS

IEEE Intelligent Systems seeks articles for a September 2008 special issue on natural-language processing and the Web. Articles should focus on innovative uses of the Web as a large-scale, distributed, hyperlinked, and multilingual corpus and on building state-of-the-art natural-language interfaces to search engines.

Submissions are due by **5 March**. Visit www.computer.org/intelligent to view the complete call for papers.

IEEE Pervasive Computing seeks articles for a December 2008 issue on environmental sustainability. The creation, use, and disposal of large quantities of pervasive technologies such as sensors and mobile devices have strong implications for resource consumption and waste production. Submissions should address design for technology reuse, repurposing, or lifetime extension; sensor network applications that support the efficient use or protection of natural resources; novel systems, devices, or interfaces that support stewardship of the natural environment; and resource-efficient system design, among other topics.

Articles are due by **23 June**. Visit www.computer.org/pervasive to view detailed author instructions and the complete call for papers.

CALLS FOR PAPERS

ICWS 2008, IEEE Conf. on Web Services, 23-26 Sept., Beijing; Submissions due **7 Apr.;** http://conferences.computer.org/icws/2008/call-for-papers.html

Music And Multimedia 2008, The Use of Symbols to Represent Music and Multimedia Objects, 8 Oct., Lugano, Switzerland; Submissions due **31 May;** http://conferences.computer.org/icws/2008/call-for-papers.html

WI-IAT 2008, IEEE/WIC/ACM Int'l Conf. on Web Intelligence & IEEE/WIC/ACM Int'l Conf. on Intelligent Agent Technology, 9-12 Dec., Sydney; Submissions due **10 July;** http://datamining.it.uts.edu.au/wi08/html/wi/?index=cfp

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

FEBRUARY 2008

18-21 Feb: WICSA 2008, Working IEEE/IFIP Conf. on Software Architecture, Vancouver, Canada; www.wicsa.net

MARCH 2008

3-7 Mar: SimuTools 2008, 1st Int'l Conf. on Simulation Tools and Techniques for Communications, Networks, and Systems, Vancouver, Canada; www.simutools.org

8-12 Mar: VR 2008, IEEE Virtual Reality, Reno, Nevada; http://conferences.computer.org/vr/2008

24-26 Mar: SSIAI 2008, IEEE Southwest Symp. on Image Analysis and Interpretation, Santa Fe, New Mexico; www.ssiai.org

25-28 Mar: AINA 2008, 22nd IEEE Int'l Conf. on Advanced Information Networking and Applications, Okinawa, Japan; www.aina-conference.org/2008

25-28 Mar: SOCNE 2008, 3rd IEEE Workshop on Service-Oriented Architectures in Converging Networked Environments (with AINA), Okinawa, Japan; www.c-lab.de/RLS/SOCNE08

31 Mar. - 4 Apr.: ECBS 2008, 15th IEEE Int'l Conf. on Eng. of Computer-Based Systems, Belfast, Northern Ireland; www.compeng.ulster.ac.uk/events/ecbs2008

APRIL 2008

7-12 Apr: ICDE 2008, 24th IEEE Int'l Conf. on Data Eng., Cancun, Mexico; www.icde2008.org

8-12 Apr: MCN 2008, 2nd IEEE Workshop on Mission-Critical Networking (with InfoCom), Phoenix; www.criticalnet.org

14 Apr: HiCOMB 2008, 7th IEEE Int'l Workshop on High-Performance Computational Biology (with IPDPS), Miami; www.hicomb.org

14-15 Apr: RAW 2008, 15th Reconfigurable Architectures Workshop (with IPDPS), Miami; www.ece.lsu.edu/vaidy/raw

14-17 Apr: CSEET 2008, 21st IEEE Conf. on Software Eng. Education and Training, Charleston, South Carolina; www.csc2.ncsu.edu/conferences/cseet

14-18 Apr: IPDPS 2008, 22nd IEEE Int'l Parallel and Distributed Processing Symp., Miami; www.ipdps.org

15-17 Apr: InfoCom 2008, 27th IEEE Conf. on Computer Communications, Phoenix; www.ieee-infocom.org

18 Apr: Hot-P2P 2008, 5th Int'l Workshop on Hot Topics in Peer-to-Peer Systems (with IPDPS), Miami; www.disi.unige.it/hotp2p/2008/index.php

18 Apr: PCGrid 2008, 2nd Workshop on Desktop Grids and Volunteer Computing Systems (with IPDPS), Miami; http://pcgrid.lri.fr

18 Apr: RoSOC-M 2008, Int'l Workshop on the Role of Services, Ontologies, and Context in Mobile Environments (with MDM), Beijing; http://events.deri.at/RoSOC-M

18 Apr: SSN 2008, 4th Int'l Workshop on Security in Systems and Networks (with IPDPS), Miami; www.cse.buffalo.edu/~fwu2/ssn08

27-30 Apr: MDM 2008, 9th Int'l Conf. on Mobile Data Management, Beijing; http://idke.ruc.edu.cn/mdm2008

MAY 2008

4-8 May: VLSI 2008, 28th IEEE VLSI Test Symp., San Diego; www.tttc-vts.org

5-7 May: ISORC 2008, 11th IEEE Int'l Symp. on Object/Component/Service-Oriented Real-Time Distributed Computing, Orlando, Florida; http://ise.gmu.edu/isorc08

7-9 May: EDCC 2008, 7th European Dependable Computing Conf., Kaunas, Lithuania; http://edcc.dependability.org

10-18 May: ICSE 2008, 30th Int'l Conference on Software Eng., Leipzig, Germany; http://icse08.upb.de

12-13 May: HST 2008, 8th IEEE Int'l Conf. on Technologies for Homeland Security, Waltham, Massachusetts; www.ieeehomelandsecurityconference.org

14-16 May: ICIS 2008, 7th IEEE Int'l Conf. on Computer and Information Science, Portland, Oregon; http://acis.cps.cmich.edu:8080/ICIS2008

19-22 May: CCGrid 2008, 8th IEEE Int'l Symp. on Cluster Computing and the Grid, Lyon, France; http://ccgrid2008.ens-lyon.fr

22-24 May: ISMVL 2008, 38th Int'l Symp. on Multiple-Valued Logic, Dallas; http://engr.smu.edu/ismvl08

Events in 2008

MARCH

3-7	SimuTools 2008
8-12	VR 2008
24-26	SSIAI 2008
25-28	AINA 2008
25-28	SOCNE 2008
31 Mar-4 Apr	ECBS 2008

APRIL

7-12	ICDE 2008
8-12	MCN 2008
14	HiCOMB 2008
14-15	RAW 2008
14-17	CSEET 2008
14-18	IPDPS 2008
15-17	InfoCom 2008
18	Hot-P2P 2008
18	PCGrid 2008
18	RoSOC-M 2008
18	SSN 2008
27-30	MDM 2008

MAY

4-8	VLSI 2008
5-7	ISORC 2008
7-9	EDCC 2008
10-18	ICSE 2008
12-13	HST 2008
14-16	ICIS 2008
19-22	CCGrid 2008
22-24	ISMVL 2008
24	ULSI 2008
25-28	GPC 2008
25-28	WaGe 2008
25-28	WMCS 2008

24 May: ULSI 2008, 17th Int'l Workshop on Post-Binary ULSI Systems (with ISMVL), Dallas; http://engr.smu.edu/ismvl08

25-28 May: GPC 2008, 3rd Int'l Conf. on Grid and Pervasive Computing, Kunming, China; http://grid.hust.edu.cn/gpc2008

25-28 May: WaGe 2008, 3rd Int'l Workshop on Workflow Management and Applications in Grid Environments (with GPC), Kunming, China; www.swinflow.org/confs/WaGe08/WaGe08.htm

25-28 May: WMCS 2008, 4th Int'l Workshop on Mobile Commerce and Services (with GPC), Kunming, China; www.engr.sjsu.edu/wmcs

CCGrid 2008

Grid computing started as a generalization of cluster computing, promising to deliver large-scale levels of parallelism to high-performance applications by crossing administrative boundaries. Today, the use of computational and data resources in high-performance applications, undertaken over grid infrastructure, has become a reality.

The IEEE International Symposium on Cluster Computing and the Grid provides researchers and practitioners with an opportunity to share their research and experience in cluster and grid technology. Topics to be addressed at the conference include grid architectures and systems, service composition and orchestration, middleware for clusters and grids, scheduling and load balancing, and parallel and wide-area file systems.

CCGrid 2008, which takes place in Lyon, France from **19-22 May**, attracts top computer engineers and scientists from around the world.

The four-day program includes contributed papers, tutorials, a poster exhibition, and a doctoral symposium, complemented by workshops held throughout the week.

For further details, visit the CCGrid website at http://ccgrid2008.ens-lyon.fr.

JUNE 2008

4-6 June: SMI 2008, IEEE Int'l Conf. on Shape Modeling and Applications, Stony Brook, New York; www.cs.sunysb.edu/smi08

10-13 June: ICPC 2008, 16th IEEE Int'l Conf. on Program Comprehension, Amsterdam; www.cs.vu.nl/icpc2008

11-13 June: SIES 2008, IEEE 3rd Symp. on Industrial Embedded Systems, La Grande Motte, France; http://www.lirmm.fr/SIES2008

11-13 June: SUTC 2008, IEEE Int'l Conf. on Sensor Networks, Ubiquitous, and Trustworthy Computing, Taichung, Taiwan; http://sutc2008.csie.ncu.edu.tw

23-25 June: CSF 2008, 21st IEEE Computer Security Foundations Symp. (with LICS), Pittsburgh; www.cylab.cmu.edu/CSF2008

23-25 June: WETICE 2008, 17th IEEE Int'l Workshop on Enabling Technologies: Infrastructures for Collaborative Enterprises, Rome; www.sel.uniroma2.it/wetice08/venue.htm

23-26 June: ICITA 2008, 5th Int'l Conf. on Information Technology and Applications, Cairns, Australia; www.icita.org

24-27 June: LICS 2008, IEEE Symp. on Logic in Computer Science, Pittsburgh; www2.informatik.hu-berlin.de/lics/lics08

JULY 2008

8-11 July: Services 2008, IEEE Congress on Services, Hawai'i;; http://conferences.computer.org/services/2008

8-11 July: CIT 2008, IEEE Int'l Conf. on Computer and Information Technology, Sydney; http://attend.it.uts.edu.au/cit2008

8-11 July: SCC 2008, IEEE Int'l Conf. on Services Computing, Honolulu, Hawai'i; http://conferences.computer.org/scc/2008

11-12 July: NCA 2008, 7th IEEE Int'l Symp. on Network Computing and Applications, Cambridge, Massachusetts; www.ieee-nca.org

28 July–1 Aug: SAINT 2008, IEEE/IPSJ Symp. on Applications and the Internet, Turku, Finland; www.saintconference.org

AUGUST 2008

4-7 Aug: ICSC 2008, 2nd IEEE Int'l Conf. on Semantic Computing, Santa Clara, California; http://icsc.eecs.uci.edu

SEPTEMBER 2008

1-3 Sept: AVSS 2008, 5th IEEE Int'l Conf. on Advanced Video and Signal-Based Surveillance, Santa Fe, New Mexico; www.cpl.uh.edu/avss2008

23-26 Sept: ICWS 2008, IEEE Int'l Conf. on Web Services, Beijing; http://conferences.computer.org/icws/2008

OCTOBER 2008

8 Oct: Music and Multimedia 2008, The Use of Symbols to Represent Music and Multimedia Objects, Lugano, Switzerland; www.cm.supsi.ch

22-24 Oct: ICEBE 2008, IEEE Int'l Conf. on e-Business Eng., Xi'an, China; http://conferences.computer.org/icebe

SOFTWARE TECHNOLOGIES

A Product-Focused Approach to Software Certification

Tom Maibaum and Alan Wassyng
McMaster University

Process alone is insufficient to guarantee necessary product attributes for certification.

History tells us that self-regulation of critical products doesn't work—certification by overseeing bodies is necessary. As software invades more areas of everyday life, certification of systems containing software is increasingly important for governments, industry, and consumers alike. Even if an organization isn't worried about safety, it must consider the consequences of using mission-critical software that isn't certified or qualified as fit for purposes. The US Sarbanes-Oxley Act of 2002, for example, imposes stringent requirements on companies' financial IT systems.

Many standards bodies and licensing authorities describe attributes of processes by which software should be developed to meet certain standards or certification criteria. However, a good process on its own doesn't necessarily result in high-quality software.

Standards and certification processes should be primarily product-focused rather than process-based to raise the certainty in evaluation of software reliability. Evaluations should be based on direct evidence about the product's attributes, not circumstantial evidence about the process.

MISSING THE POINT

Licensing organizations currently aim to establish a common understanding between software producers and certifiers. For example, the US Food and Drug Administration has published several guidance documents concerning the validation of medical software. However, such recommendations aren't explicit or precise. For example, the FDA doesn't adequately specify the

- objects subject to assessment,
- measurable attributes that characterize these objects, or
- criteria on which the FDA staff will base its decision to approve or reject the medical software and, thus, the measurement procedures to ascertain the values of the relevant attributes of the objects being assessed.

Instead, the FDA's approach focuses on the characteristics of a development process that is likely to produce satisfactory software. It shares this approach with most certification authorities' requirements as well as those of standards based on maturity, such as the Capability Maturity Model (CMM).

Although this focus on process has had an important and positive effect on software development, it misses the point of certification—namely, to ascertain whether the product for which a certificate is being sought has appropriate characteristics.

A MEASUREMENT-BASED ACTIVITY

Certification should assess a product in terms of measurable attributes using an agreed-upon objective function, combining these measures as a basis for decision making. This function is itself subjectively defined, but once specified—by, for example, a certifying authority (CA), government regulatory body, or company procuring the software for itself—its use is objective, repeatable, and, perhaps most importantly for the economic well-being of all parties concerned, predictable.

Using a process likely to deliver a sound product (for example, CMM level 5) offers no guarantee that it will not deliver a faulty product that could cause an avoidable disaster. A process simply can't provide the kind of assurance required for certification if it doesn't relate directly to relevant qualities of the product being certified. Even if the process provides correctness by construction, this doesn't ensure the product's acceptability. For example, the specification on which the correctness assertion is based might be flawed.

Process-oriented standards and certification criteria alone will never prove to be satisfactory ways

of guaranteeing software properties and providing a basis for licensing. What's needed is a scientific—that is, evidence-based—and product-focused approach to certification.

FAKING IT

Process-based certification imposes many procedures designed exclusively to manage the process's imperfections. For example, the FDA recommends developing and implementing a configuration-management plan for medical software. But what does configuration management have to do with the product's properties? Similar observations can be made about bug-reporting and bug-fixing mechanisms.

David Parnas and Paul Clements argued that, while the actual instances of a development process are likely to be imperfect, it's nevertheless useful for the manufacturer to develop documentation that fakes the ideal execution of the project, eliminating backtracking, fixes, workarounds, and so on ("A Rational Design Process: How and Why to Fake It," *IEEE Trans. Software Eng.*, Feb. 1986, pp. 251-257).

Intermediate products are needed as part of the evidence being assessed for certification of the final product—for example, a requirements specification, a design specification, a document describing validation of the design against the requirements, documents relating to testing, and documents proving correctness. These various products can be organized in terms of an idealized development process, such as a simplified waterfall model, in which every stage must be completed before moving on to the next.

Whatever actual process it follows, the onus is on the organization seeking certification for its product to map evidence onto the idealized process. Thus, the CA can measure both the actual product and associated products. The CA shouldn't mandate any particular development process or the necessity of having, say, a configuration-management plan.

Software manufacturers can define their own internal process as long as they can effectively map their products onto the ones that the much simpler, faked process requires. They can then manage their own internal process without having to undertake the difficult, time-consuming job of "proving" conformance to a mandated process.

Of course, the CA must decide what evidence is required and design the idealized process to deliver this evidence. Striking the fine balance between the level of detail in the idealized process and the weight of evidence necessary to make the certification decision is probably the most difficult job the CA has to perform.

The onus is on the organization seeking certification for its product to map evidence onto the idealized process.

An example of this idea in action is the process model used to redevelop the safety system of the Darlington Nuclear Generating Station on Lake Ontario in Canada. Figure 1 shows a simplified version of the actual process, which the certifiers approved but didn't mandate.

MEASUREMENT FRAMEWORK

Many models attempt to capture an existing or yet-to-be-built process and its associated products, which can be technical or administrative or both (M. Myers and A. Kaposi, *A First Systems Book: Technology and Management*, 2nd ed., Imperial College Press, 2004). A process constitutes behavior that occurs over time, while a product is an instantaneous entity; both have measurable attributes.

Processes and products

Norman Fenton and Shari Pfleeger defined processes, in the context of software engineering, as "collections of software-related activities" (*Software Metrics: A Rigorous and Practical Approach*, 2nd ed., PWS Publishing Co., 1998). They used the term "internal process attributes" to describe those attributes that can be measured directly by examining the process definition on its own. We refer to these attributes as *static*. A software certifier should, after identifying these attributes, evaluate the procedures used to measure each of them. The evaluation can occur at predefined times or continue over a time interval.

In contrast, *dynamic* attributes such as quality and stability can only be measured with regard to the way the process relates to its environment—that is, the process behavior rather than the process definition is the focus of the measurement activity. The values of these attributes, which can depend on some values of static process attributes, aren't meaningful outside their operating environment.

Fenton and Pfleeger defined products as "any artifacts, deliverables or documents that result from a process activity." Examples include software requirements specifications, software design specifications, source code, or any other outcomes of the development process. Products are atemporal in that their attributes can be measured at any time instant, though the measured values might differ from one instant to another.

The same notion of static and dynamic attributes applies to products. The software's version number is a static attribute related to the product itself, whereas reliability is a dynamic attribute related to the way the product behaves in its operating environment.

Specifying product attributes

To facilitate the formal evaluation of software products for validation or certification, their attributes must be specified. This involves defining a measurement

- scale for the attribute—that is, its type; and
- procedure for ascertaining the specific value in the type.

An acceptability criterion for each attribute must also be established that defines the acceptable values from the evaluation's point of view. Often it's not the value of a specific attribute that determines product acceptability but the result of some objective function applied to some or all of the product's attribute measures.

Tools used in software production must also be evaluated. These are considered entities with attributes, exactly like products, and thus their attributes must be specified and measured.

The FDA's guidelines for medical software validation lack such objective criteria. Because the FDA doesn't explicitly define what processes or products are to be measured, software developers and FDA evaluators have no deep, common understanding about what evidence to inspect, what attributes to measure, and what values are acceptable.

The Common Criteria (CC) for Information Technology Security Evaluation attempts to identify products of development processes and determine what the evaluator must do, though it still lacks adequate specificity. Interestingly, the CC classifies certification into different levels that reflect security features' criticality to the organization. The top-most levels require formal specification and verification products. For each level, various kinds of evidence are required to satisfy the CA—a common feature in regulatory environments.

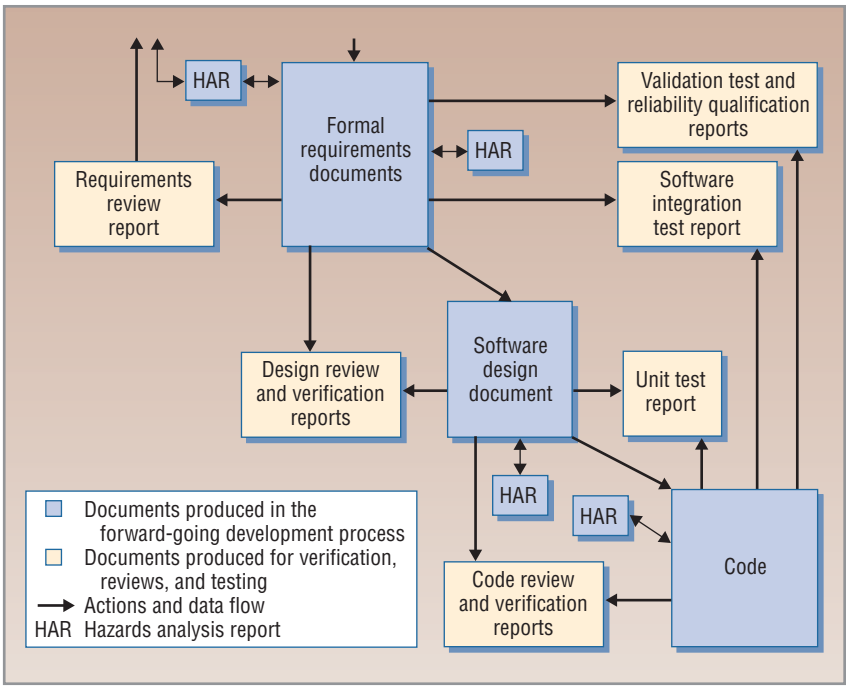

Figure 1. Idealized version of the process used to redevelop the safety system of the Darlington Nuclear Generating Station. (Figure adapted from A. Wassyng and M. Lawford, "Software Tools for Safety-Critical Software Development,"* Int'l J. Software Tools for Technology Transfer, *vol. 8, nos. 4-5, 2006, p. 338.)

Regulations designed to produce reliable software generally rely on process-based approaches, but there's little evidence that particular processes produce software with the quality required for critical applications. Without a fundamental understanding of the relationship between software attributes and quality, process-oriented approaches will remain hit-and-miss aids. However, if we develop this understanding, we can build audit processes to support predictable, objective, product-focused certification methods.

During the past year, McMaster University's Software Quality Research Laboratory started an initiative on software certification under the aegis of the Software Certification Consortium, which includes North American representatives of academia, industry, national agencies interested in software, and regulators. The SCC is identifying research-oriented activities that will facilitate and promote product-oriented, evidence-based approaches to certification. It is also providing Grand Challenge problems as testbeds for both engineering methods and potential certification regimes.

The first of these is the Pacemaker Formal Methods Challenge (www.cas.mcmaster.ca/sqrl/pacemaker.htm). Boston Scientific has made public a typical requirements specification for a pacemaker manufactured 10 years ago. This application is of a size and complexity not seen in past grand challenge problems and provides a unique opportunity to apply a product-oriented, evidence-based certification evaluation. The SCC will soon add analogous challenges in other domains.

Tom Maibaum *is a professor and holds the Canada Research Chair in the Foundations of Software Engineering in the Department of Computing and Software at McMaster University, Hamilton, Ontario, Canada. Contact him at tom@maibaum.org.*

Alan Wassyng *is an associate professor in the Department of Computing and Software and acting director of the Software Quality Research Laboratory at McMaster University. Contact him at wassyng@mcmaster.ca.*

Editor: Michael G. Hinchey, Loyola College in Maryland; mhinchey@loyola.edu

2008 RATES for IEEE COMPUTER SOCIETY

Membership Dues and Subscriptions

Membership and periodical subscriptions are annualized to and expire on 31 December 2008.
Pay full or half-year rate depending upon the date of receipt by the IEEE Computer Society as noted below.

Membership Options*

All prices are quoted in U.S. dollars.

	FULL YEAR Applications received 17 Aug 07 – 29 Feb 08	HALF YEAR Applications received 1 Mar 08 – 15 Aug 08
I do not belong to the IEEE and I want to join only the Computer Society:	☐ $113.00	☐ $57.00
I want to join **both** the Computer Society and the IEEE:		
I reside in the USA	☐ $220.00	☐ $110.00
I reside in Canada	☐ $195.00	☐ $98.00
I reside in Africa/Europe/Middle East	☐ $187.00	☐ $94.00
I reside in Latin America	☐ $180.00	☐ $90.00
I reside in Asia/Pacific	☐ $181.00	☐ $91.00
I already belong to the IEEE, and I want to join the Computer Society:	☐ $50.00	☐ $25.00

Are you now or were you ever a member of the IEEE?
☐ Yes ☐ No If yes, please provide member # if known: _____

Add Periodicals**

	ISSUES PER YEAR	FULL YEAR Applications received 16 Aug 07 – 29 Feb 08 PRINT + ONLINE	HALF YEAR Applications received 1 Mar 08 – 15 Aug 08 PRINT + ONLINE
BEST VALUE!			
IEEE Computer Society Digital Library (online only)	n/a	☐ $121	☐ $61
ARTIFICIAL INTELLIGENCE			
IEEE Intelligent Systems	6	☐ $43	☐ $22
IEEE Transactions on Learning Technologies†	4	☐ $25	☐ $13
IEEE Transactions on Pattern Analysis and Machine Intelligence	12	☐ $52	☐ $26
BIOTECHNOLOGY			
IEEE/ACM Transactions on Computational Biology and Bioinformatics	4	☐ $36	☐ $18
COMPUTATION			
Computing in Science & Engineering	6	☐ $45	☐ $23
COMPUTER HARDWARE			
IEEE Computer Architecture Letters	2	☐ $29	☐ $15
IEEE Micro	6	☐ $41	☐ $21
IEEE Design & Test of Computers	6	☐ $40	☐ $20
IEEE Transactions on Computers	12	☐ $47	☐ $24
GRAPHICS & MULTIMEDIA			
IEEE Computer Graphics and Applications	6	☐ $43	☐ $22
IEEE MultiMedia	4	☐ $38	☐ $19
IEEE Transactions on Haptics	2	☐ $31	☐ $16
IEEE Transactions on Visualization and Computer Graphics	6	☐ $43	☐ $22
HISTORY OF COMPUTING			
IEEE Annals of the History of Computing	4	☐ $34	☐ $17
INTERNET & DATA TECHNOLOGIES			
IEEE Internet Computing	6	☐ $43	☐ $22
IEEE Transactions on Knowledge and Data Engineering	12	☐ $49	☐ $25
IEEE Transactions on Services Computing†	4	☐ $25	☐ $13
IT & SECURITY			
IT Professional	6	☐ $42	☐ $21
IEEE Security & Privacy	6	☐ $24	☐ $12
IEEE Transactions on Dependable and Secure Computing	4	☐ $33	☐ $17
MOBILE COMPUTING			
IEEE Pervasive Computing	4	☐ $43	☐ $21
IEEE Transactions on Mobile Computing	12	☐ $43	☐ $22
NETWORKING			
IEEE Transactions on Parallel and Distributed Systems	12	☐ $47	☐ $24
SOFTWARE			
IEEE Software	6	☐ $49	☐ $25
IEEE Transactions on Software Engineering	6	☐ $38	☐ $19

* Member dues include $17 for a 12-month subscription to *Computer* magazine.
** Periodicals purchased at member prices are for the member's personal use only.
† Online issues only .

Payment Information

Payment required with application

Membership fee
$

Periodicals total
$

Applicable sales tax***
$

TOTAL:
$

Enclosed:
☐ Check/Money Order****

Charge my:
☐ MasterCard
☐ VISA
☐ American Express
☐ Diner's Club

Card Number

Exp Date (month/year)

Signature

USA Only include
5-digit billing zip code

* Member dues include $17 for a 12-month subscription to *Computer*.

** Periodicals purchased at member prices are for the member's personal use only.

*** Canadian residents add 14% HST or 6% GST to total. AL, AZ, CO, DC, GA, IN, KY, MD, MO, NM, and WV add sales tax to periodical subscriptions. European Union residents add VAT tax to IEEE Computer Society Digital Library subscription.

**** Payable to the IEEE in U.S. dollars drawn on a U.S. bank account. Please include member name and number (if known) on your check.

Allow up to 8 weeks for application processing. Allow a minimum of 6 to 10 weeks for delivery of print periodicals.

Please complete both sides of this form.

For fastest service, apply online at www.computer.org/join

Personal Information

Enter your name as you want it to appear on correspondence. As a key identifier in our database, circle your last/surname.

☐ Male ☐ Female

Date of birth (Day/Month/Year)

Title | First name | Middle

Last/Surname

Home address

City | State/Province

Postal code | Country

Home telephone

Home facsimile

Preferred e-mail

Send mail to: ☐ Home address ☐ Business address

Educational Information

First professional degree completed | Month/Year degree received

Program major/course of study

College/University | State/Province | Country

Highest technical degree received | Program/Course of study

Month/Year received

College/University | State/Province | Country

Business/Professional Information

Title/Position

Years in current position | Years of practice since graduation

Employer name

Department/Division

Street address

City | State/Province

Postal code | Country

Office phone

Office facsimile

I hereby make application for Computer Society and/or IEEE membership and agree to be governed by IEEE's Constitution, Bylaws, Statements of Policies and Procedures, and Code of Ethics. I authorize release of information related to this application to determine my qualifications for membership.

Signature | Date

IF8B

NOTE: In order for us to process your application, you must complete and return BOTH sides of this form to the office nearest you:

Asia/Pacific Office
IEEE Computer Society
Watanabe Bldg.
1-4-2 Minami-Aoyama
Minato-ku, Tokyo 107-0062 Japan
Phone: +81 3 3408 3118
Fax: +81 3 3408 3553
E-mail: tokyo.ofc@computer.org

Publications Office
IEEE Computer Society
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314 USA
Phone: +1 800 272 6657 (USA and Canada)
Phone: +1 714 821 8380 (worldwide)
Fax: +1 714 821 4641
E-mail: help@computer.org

BPA Information

This information is used by society magazines to verify their annual circulation. Please refer to the audit codes and indicate your selections in the box provided.

A. Primary line of business

1. Computers
2. Computer peripheral equipment
3. Software
4. Office and business machines
5. Test, measurement, and instrumentation equipment
6. Communications systems and equipment
7. Navigation and guidance systems and equipment
8. Consumer electronics/appliances
9. Industrial equipment, controls, and systems
10. ICs and microprocessors
11. Semiconductors, components, sub-assemblies, materials, and supplies
12. Aircraft, missiles, space, and ground support equipment
13. Oceanography and support equipment
14. Medical electronic equipment
15. OEM incorporating electronics in their end product (not elsewhere classified)
16. Independent and university research, test and design laboratories, and consultants (not connected with a manufacturing company)
17. Government agencies and armed forces
18. Companies using and/or incorporating any electronic products in their manufacturing, processing, research, or development activities
19. Telecommunications services, and telephone (including cellular)
20. Broadcast services (TV, cable, radio)
21. Transportation services (airlines, railroads, etc.)
22. Computer and communications and data processing services
23. Power production, generation, transmission, and distribution
24. Other commercial users of electrical, electronic equipment, and services (not elsewhere classified)
25. Distributor (reseller, wholesaler, retailer)
26. University, college/other education institutions, libraries
27. Retired
28. Others (allied to this field) _____

B. Principal job function

1. General and corporate management
2. Engineering management
3. Project engineering management
4. Research and development management
5. Design engineering management — analog
6. Design engineering management — digital
7. Research and development engineering
8. Design/development engineering — analog
9. Design/development engineering — digital
10. Hardware engineering
11. Software design/development
12. Computer science
13. Science/physics/mathematics
14. Engineering (not elsewhere classified)
15. Marketing/sales/purchasing
16. Consulting
17. Education/teaching
18. Retired
19. Other _____

C. Principal responsibility

1. Engineering or scientific management
2. Management other than engineering
3. Engineering design
4. Engineering
5. Software: science/management/engineering
6. Education/teaching
7. Consulting
8. Retired
9. Other _____

D. Title

1. Chairman of the Board/President/CEO
2. Owner/Partner
3. General Manager
4. V.P. Operations
5. V.P. Engineering/Director Engineering
6. Chief Engineer/Chief Scientist
7. Engineering Manager
8. Scientific Manager
9. Member of Technical Staff
10. Design Engineering Manager
11. Design Engineer
12. Hardware Engineer
13. Software Engineer
14. Computer Scientist
15. Dean/Professor/Instructor
16. Consultant
17. Retired
18. Other Professional/Technical _____

Social Networking

Alfred C. Weaver and Benjamin B. Morrison
University of Virginia

The mass adoption of social-networking websites points to an evolution in human social interaction.

Social networking is a concept that has been around much longer than the Internet or even mass communication. People have always been social creatures; our ability to work together in groups, creating value that is greater than the sum of its parts, is one of our greatest assets.

At its bare essentials, a social network consists of three or more entities communicating and sharing information. This could take the form of a research coalition, a Girl Scout troop, a church, a university, or any number of other socially constructed relationships.

Since the explosion of the Internet age, more than 1 billion people have become connected to the World Wide Web, creating seemingly limitless opportunities for communication and collaboration. In the context of today's electronic media, social networking has come to mean individuals using the Internet and Web applications to communicate in previously impossible ways. This is largely the result of a culture-wide paradigm shift in the uses and possibilities of the Internet itself.

The current Web is a much different entity than the Web of a decade ago. This new focus creates a riper breeding ground for social networking and collaboration. In an abstract sense, social networking is about everyone.

The model has changed from top-down to bottom-up creation of information and interaction, made possible by new Web applications that give power to users. While in the past there was a top-down paradigm of a few large media corporations creating content for the consumers to access, the production model has shifted so that individual users now create content that everyone can share.

The social-networking trend is causing a major shift in the Internet's function and design. While we previously thought of the Internet as an information repository, the advent of social networks is turning it into a tool for connecting people. The mass adoption of social-networking websites of all shapes and sizes points to a larger movement, an evolution in human social interaction.

AGE DISTINCTION AMONG SOCIAL NETWORKING USERS

A social-networking site typically allows users to post their profiles and create personal networks for exchanging information with other users. According to a Pew Internet and American Life Project report (www.pewInternet.org/pdfs/PIP_SNS_Data_Memo_Jan_2007.pdf), more than half of Americans aged 12 to 17 use online social-networking sites, 55 percent of teens have created a personal profile online, and 55 percent have used social-networking sites such as Facebook or MySpace. Regardless of whether the information they exchange is serious or frivolous, the underlying Web technology is the same.

The mass murder at Virginia Tech on 16 April 2007 provides a poignant lesson concerning the functions of social networking. As this tragic event unfolded during the early-morning hours, university officials communicated via conventional police radios and telephones. As the scope of the tragedy became apparent, students and peers notified each other via e-mail. But students at other schools used Facebook to ascertain the status of their friends attending Virginia Tech, aided by VT students annotating their personal pages with comments such as "I'm all right," "I'm safe," and "I'm coming home. See you soon." This is a perfect example of using the Internet as a way of keeping updated with a network of friends, a function that the younger generation is readily adopting.

But will everyone adopt this conceptualization of the Internet as a social platform? The ubiquity of social networking among college-age individuals remains largely unappreciated by their parents' demographic. This is a generational issue that will in turn affect how well the modern workplace accepts and adopts social-networking concepts. Information sharing, collegiality, quick group formation, and a sense of shared purpose and mission characterize the various Web-based social-networking applications that today's young professionals commonly use.

Table 1. Social networking websites.

Website	URL	Focus	Registered users*
Broadcaster	broadcaster.com	Video sharing and webcam chat	26,000,000
Classmates	classmates.com	School, college, work, and military groups	40,000,000
Facebook	facebook.com	Upload photos, post videos, get news, tag friends	58,000,000
Flickr	flickr.com	Photo sharing	4,000,000
Fotolog	fotolog.com	338 million photos from around the world	12,695,007
Friendster	friendster.com	Search for and connect with friends and classmates	50,000,000
MySpace	myspace.com	Videos, movies, IM, news, blogs, chat	217,000,000
Windows Live Spaces	spaces.live.com	Blogging	40,000,000

* The counts of registered users are taken from http://en.wikipedia.org/wiki/List_of_social_networking_websites.

SOCIAL-NETWORKING SUCCESSES

A large or rapidly growing user population characterizes a successful social network. The most popular social networks grow their membership through viral marketing—the natural human behavior that causes people to tell others about products or services that are particularly good or bad. The value proposition is that the user must see enough return on investment of some measure of involvement—for example, time, energy, or money—to continue using the service over a long time. Table 1 lists some currently popular social-networking websites.

The service provider must see sufficient return on its investment—for example, staff, equipment, advertising, and website maintenance—to continue to provide a stable service to its users.

"Successful" is not necessarily synonymous with "profitable." While many social-network sites are (or at least hope to be) profitable, others, such as Wikipedia, the free online encyclopedia, are socially motivated, not-for-profit ventures. Wikipedia, funded entirely by donations, seeks to serve society by making the vast store of human knowledge available to everyone, for free, in the world's most common languages.

CASE STUDIES

Different social networks emphasize different aspects of human interaction.

Four case studies show the range of socialization that can occur within social networks focused around the core purpose of connecting people online. All four of these sites can be characterized as social networks. They each give tools that let users express themselves and interact with others in different ways, choosing to focus on different areas of social interaction. Although each site expresses sociality in a different way, they all are built around the core idea of connecting people and enabling them to interact.

MySpace

MySpace (www.myspace.com) is a peer- and media-based social network in which members can create their own miniature websites containing pictures, profile information such as age and interests, and media. MySpace allows users to embed media snippets such as music or video in their sites. They also can customize the look and feel of their MySpace websites by changing the webpage files.

Initiated in 2003 in Los Angeles, MySpace was originally a music-oriented site designed to help independent musicians keep in touch with their fans. Bands can create MySpace pages and post samples of their music as well as tour dates and show locations. Thus, the concept of sharing media is at the core of MySpace, and the idea of sharing music naturally expanded to sharing video.

Now the largest online social network, MySpace claims more than 217 million registered users worldwide. It has expanded far beyond its initial focus on music to saturate the financially valuable 16-to-34 age group. The distribution of males and females is almost even (within 1 percent), and the site is also incredibly popular with younger teens and preteens.

MySpace differs from other social-networking sites (such as Facebook) mainly in its informal social-networking contracts. To sign up, MySpace requires little more than an e-mail address. A MySpace member's page has its own URL, and anyone can view the page. An Internet user doesn't even need to be a MySpace member to browse the websites and view members' pages. Few privacy settings are available to determine who can or can't see pages.

MySpace experienced explosive growth due to its viral design of inviting friends to join and also because it encourages the practice of "friending" people the users don't necessarily know well in real life so that they can share media.

MySpace also emphasizes that its members retain proprietary ownership of the media they upload to the site. This helped reassure independent bands and groups who were concerned that their material might be stolen, and instead lets members focus on connecting as many people to their "spaces" as possible.

While MySpace didn't have first-mover advantage—a similar site

called Friendster was launched years before MySpace, but focused mainly on dating—it does allow users more control over things they deem important, such as uploading their favorite media and customizing their websites. Being flexible and responsive to members' demands and desires has fueled the site's growth.

Facebook

Launched by a Harvard sophomore in February 2004, Facebook (www.facebook.com) was initially restricted to Harvard students. Over the next two years, it was expanded to include all universities, then high schools, and now to anyone age 13 or older.

Facebook is a peer-relationship-based social network that allows users to create personal profiles describing their real-world selves and then establish connections with other users. In addition to basics such as name, age, and school, profiles also include information such as favorite books, movies, quotes, and photos. Users can further customize their accounts by uploading other media such as images or movies, and interact with other users by commenting on their profiles or media.

Even though social networking is no longer a new phenomenon, Facebook has created innovations that have allowed it to keep its users interested. Facebook is a unique social network that views itself as a platform. This means that it has made extensive tools, documentation, and an application programming interface available for third-party developers to use in creating "applications" they can seamlessly embed into Facebook. Using server-side software from other companies and partners, as well as from lone programmers sitting in their bedrooms, users can fill the profiles and pages they habitually use. Opening up both the front and back end of Facebook to others lets third-party developers create their own Web applications that utilize the Facebook network.

Figure 1. Wikipedia is a multilingual, Web-based, free content encyclopedia project, written collaboratively by authors around the world. Wikipedia has more than 75,000 active contributors working on more than 9 million articles in more than 250 languages. Source: www.wikipedia.org.

As with any online business, Facebook experienced many growing pains, particularly with regard to privacy. Cases of school administrators and legal authorities using Facebook's user-uploaded images as evidence of illegal behavior such as underage drinking have engendered controversy and have made users more aware of their responsibility with regard to their online presence. Facebook also has followed user feedback by implementing privacy settings that let users select what information will be available to certain parties online.

As Facebook has learned, it's important to recognize that while social networks are focused on connecting others, users must follow certain rules and etiquette for the interactions to be viewed as legitimate. Otherwise, privacy is violated, making users feel uncomfortable and prompting them to leave the site.

Wikipedia

Wikipedia (www.wikipedia.org) is a collaborative online encyclopedia project in which anyone can create and edit the content displayed on the site. As Figure 1 shows, to date, Wikipedia hosts 9.25 million articles in 253 languages (2 million articles in English), edited by tens of thousands of users every day. The encyclopedia is organized as a system of articles that are interconnected using hyperlinks, allowing readers to look up subject areas and then click links to explore related areas. The thousands of volunteer users and editors create new topics and pages.

Owned by the Wikimedia Foundation, a nonprofit organization seeking to develop and maintain open content that allows anyone to contribute, Wikipedia does not display advertisements on its pages and generates no revenue. It is supported through donations from both individuals and corporate sponsors (such as Yahoo!). Wikipedia's founders view the site as a public service and have no interest in making a profit from it.

Wikipedia's users organize the editing of the site behind the scenes. "Discussion" and "history" tabs allow users to look "under the hood" and see the work that goes into each article. Contributors discuss topics such as an article's scope and the information's reliability and bias. The editors discuss the best ways to adhere to Wikipedia's policies of fairness and quality. Editors can view changes made to the article and revert to an earlier version if they detect tampering or invalid editing.

Since users can change the site without having an account, Wikipedia

has been criticized for its susceptibility to vandalism. Whether Wikipedia is a reliable resource is still open to debate. *Encyclopedia Britannica* has released reports claiming that Wikipedia contains more errors than a traditional encyclopedia that experts compile (www.post-gazette.com/pg/06083/676130.stm). However, other reports have claimed that Wikipedia is just as reliable as a conventional encyclopedia (http://news.com.com/2100-1038-5997332.html).

YouTube

An online video-sharing network, YouTube (www.youtube.com) allows users to upload videos and share movie clips. Users can browse other users' videos as well as store their own videos on the site for free. They also can recommend popular videos to their friends via e-mail, or embed simple code that lets them display YouTube videos on their own websites or blogs.

Users mainly govern the YouTube community, and there are several ways that viewers can interact with the videos. Users can leave comments or recommendations of other videos or they can "flag" or report videos that violate YouTube's usage agreements (such as pornography or offensive content). YouTube's staff then investigates these flagged videos.

CREATING A STANDARD

In March 2007, Google unveiled an ambitious plan to standardize the movement to convert the Internet into a social platform. Google's open source framework, OpenSocial, is intended for use by the dozens of already popular social-networking websites such as MySpace, LinkedIn, Beebo, and others.

With OpenSocial, developers will be able to use a common HTML language and JavaScript to take advantage of the user connections within all social networks, not just a single site. Instead of writing one application for MySpace and another for Facebook, OpenSocial hopes to let developers create a single application that can spread across all platforms seamlessly.

In essence, Google is attempting to standardize the code base for the new social Internet. If most social-network sites adopt the OpenSocial standard, new ways of human interaction will proliferate, just as new sources of information and webpages did in the early days of the Web. With all social networks connected, the Internet will truly mature in its transformation to a social platform. Only time will tell if others will follow Google's lead.

So what is the future of the Internet? Before the advent of the Web, the Internet was primarily a repository of information and data—a giant encyclopedia. It was the place where users went to find what they needed to know.

In this decade, with the growth of social networks and user-generated content, the Internet is becoming a hub of socialization, a social utility. The Internet is now where users go to interact and connect with others. Far beyond e-mail, the Internet is becoming a means of connecting people to one another, across distances and time, allowing an order of socialization and culture never before seen. Social networking is the logical extension of our human tendencies toward togetherness, whether that socialization is down the hall or across the world.

Alfred C. Weaver *is a professor of computer science at the University of Virginia. One of his research projects studies social networking to learn how to build better Web-based collaboration environments. Contact him at weaver@virginia.edu.*

Benjamin B. Morrison *is a senior computer science major and a media studies minor at the University of Virginia. In addition to researching the social implications of emerging technology, Morrison is the founder of UVa's Student Game Developers, an organization that enables students to create entertainment software. Contact him at bbm6f@virginia.edu.*

Computer welcomes your submissions to this bimonthly column. For additional information, or to suggest topics that you would like to see explained, contact column editor Alf Weaver at weaver@cs.virginia.edu.

ENTERTAINMENT COMPUTING

The Business of Fun

Michael van Lent
USC Institute for Creative Technologies

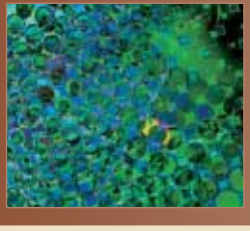

Videogames promote flash and features, but compelling gameplay wins the day.

According to the NPD Group (www.npd.com), US videogame retail sales approached $12.5 billion in 2006, an increase of almost 20 percent over the 2005 total of $10.5 billion. While the frequently cited claim that "the game industry is bigger than the movie industry" does not hold up to scrutiny, it is true that US videogame retail sales exceeded the domestic box office movie gross of $9.2 billion in 2006, according to Box Office Mojo (www.boxofficemojo.com). Further, game sales are now going toe-to-toe with more established entertainment forms.

GAMING HOLLYWOOD

Microsoft's *Halo 3* raked in an estimated $170 million on its first day of release, topping the previous motion picture release-day champion, *Spider-Man 3*, which grossed only $151 million. Despite Microsoft's claim, *Halo 3* probably doesn't beat the launch-day record for the best-selling book—*Harry Potter and the Deathly Hallows* holds that title, having sold 11 million copies worldwide in the first 24 hours, according to BBC News.

These are big numbers. Hit games like *Halo 3* and *World of Warcraft* are, without a doubt, highly profitable. *Halo 3* probably cost less than $50 million to develop and market, an investment Microsoft made back three times over in its first-day sales. Statistics like these for hit games can give the impression that money comes easily in the videogame industry.

The truth is less rosy. The cost of developing a modern videogame for either the PC or a leading console (Xbox 360, PlayStation 3, Nintendo Wii) is increasing rapidly, but the return on investment has not kept pace.

To take an extreme example of two hit games, in 1982 Atari spent $100,000 to develop the Video Computer System version of its *Pacman* title, which went on to sell 10 million copies at $30 each. In 2004, Microsoft spent an estimated $40 million to develop *Halo 2*, which sold 8 million copies at $50 each (http://arstechnica.com/articles/paedia/hardware/crossplatform.ars/2).

This staggering escalation in development costs poses serious challenges for game developers and publishers seeking to maintain and even increase their profit margins.

ART IS KING

The major costs of developing and launching a videogame title involve artwork creation, programming, console fees, and marketing—not necessarily in that order. According to a *Forbes* magazine analysis (www.forbes.com/technology/2006/12/19/ps3-xbox360-costs-tech-cx_rr_game06_1219expensivegames.html), about 30 percent of development costs go for artwork creation and design, 25 percent for programming, and 15 percent for console fees—Microsoft, Sony, and Nintendo take a cut of every game sold for their consoles. Another 15 percent goes for marketing. Not surprisingly, a major factor in the increasing cost of game development is the growing budgets for artwork creation and programming.

As game consoles become more powerful, the images they can render become much more detailed and therefore more costly. For example, the shift from the PlayStation to the PlayStation 2 increased polygon counts by a factor of 50 and increased the resolution of the rendered image two and a half times. This forced artists to spend more time developing much more detailed models, which in turn led to larger teams of artists and correspondingly larger art budgets.

Accordingly, the composition of game development teams has shifted from mostly programmers to mostly artists. In 2006, Bing Gordon, Electronic Arts' chief creative officer, estimated that its current titles had more individual art files than lines of code. The current-generation Xbox 360 and PlayStation 3 both support high-definition (HD) video, which raises the development bar even higher.

UPPING THE ANTE

The increased computational power that renders all this detailed and expensive artwork brings with it higher programming costs as well. Current-generation game consoles are complex multiprocessor machines that pose significant technical challenges even for such legendary game programmers as

id Software's John Carmack. As a result, development teams require more experienced and better-educated programmers who command correspondingly higher salaries. Although plausible in the industry's early days, the proposition that two kids could develop a modern console game in their parents' garage today seems ludicrous.

While some might worry that spiraling costs will bankrupt the videogame industry, game developers and publishers still find creative ways to stay out of the red. Some game developers work to decrease the cost of developing games while others cut out the middle man—in this case game publishers and retailers.

For example, *Spore*, the next brainchild of master game designer Will Wright, decreases the cost of art production by automating some content creation, such as animation, that would typically be created painstakingly by a human artist. Valve's Steam digital distribution platform lets consumers buy games such as *Half-Life 2* online. This approach directly eliminates manufacturing costs, retail markups, and publisher overhead.

MILKING THE COW

However, some of the game business's most recent developments tread the opposite path and seek to extract *more* cash from gamers. The most obvious approach to increasing revenue from game sales is to raise the game's price, the price of the platform it runs on, or both. While games and consoles have both become costlier, the industry perceives that there's a limit to what people will pay to play videogames. This is especially true in the case of game consoles, which generally must be sold below cost to hit the $300-$500 range that seems to be the consumer's sweet spot.

Business Week (www.businessweek.com/technology/content/nov2005/tc20051122_410710.htm) estimates that the Xbox 360, which sells for $399, actually costs $525 to manufacture. Console makers depend on those already mentioned console fees, 15 percent of the game's development cost, to make up the difference.

Even so, making game consoles requires deep pockets and a long-term view. Amazingly, Microsoft's Entertainment and Devices Division, which launched the original Xbox in 2001, didn't have a profitable year until 2007, when *Halo 3* proved such a massive success. Since game developers can't simply raise prices to fatten their bottom lines, they must devise other approaches to enhance revenues at their customers' expense.

The industry perceives that there's a limit to what people will pay to play videogames.

MASSIVELY PROFITABLE

To date, massively multiplayer online role-playing games have led the charge to find new revenue streams. The main attraction of MMORPGs like *EverQuest* and *World of Warcraft* is the shared, persistent virtual world populated by tens of thousands of paying customers. To play, these gamers must log into one node of the vast server array that hosts their virtual world. Each node, in turn, can host thousands of players.

MMORPG publishers not only charge each player for a copy of the boxed game (usually costing $40 to $60), the players must also pay a monthly subscription fee to access the servers. While maintaining these subscriptions undoubtedly incurs an expense, the fee still represents a significant second stream of revenue for the game company—and its profits can be substantial.

In July 2007, Blizzard Entertainment announced that *World of Warcraft* had already recruited 9 million subscribers, each paying between $13 and $15 a month in subscription fees. This amounts to a staggering $126 million a month—or more than $1.5 billion annually—in subscription fees. In addition, because it is impossible to play without a subscription, MMORPGs can ignore most software piracy and hacking concerns.

Similarly, Microsoft's Xbox platform offers a subscription-based Xbox Live service that, while not technically an MMO, lets large numbers of players chat and compete online; download additional games, demos, and optional content; and interact with the rest of the Xbox Live community.

DANCE, MONKEY BOY

While each MMORPG charges a subscription fee for access to its imaginary world, other games charge extra for specialized peripherals that let the player interact with the game in engaging new ways. *Dance Dance Revolution* provided the groundbreaker in this trend when it replaced the joystick-like game controller with a floormat. To dance, the player steps on different parts of this mat, mimicking the patterns displayed on the game screen. *Dance Dance Revolution* has kicked off an entire genre of rhythm games with a variety of increasingly realistic peripherals.

The *Guitar Hero* series, for example, gives players a guitar-shaped controller with fret buttons on the neck and a strum bar on the body. The game *Rock Band* takes this a step further by supporting controllers for the entire band, including two guitar controllers, a set of drum controllers, and even a microphone controller that tracks pitch and individual vowels and consonants. The publisher sells the game disk and peripherals packaged together and at a higher price point than a traditional videogame: *Guitar Hero III*, with a Les Paul-style controller, retails for $100; *Rock Band*, with one guitar, a drum set, and a microphone, retails

for $170. In addition to creating an entire new product line, these music-oriented rhythm games appeal to many casual gamers and thus have the potential to expand the existing customer base.

Game companies have begun exploring a variety of revenue-enhancing models. Several, including Linden Lab (*Second Life*) and Sony (*EverQuest*), have started selling virtual goods for real-world dollars. Others sell small chunks of additional game content for micropayments of a few dollars each.

As always in the game industry, however, nothing matters as much as creating fun-to-play games. Nintendo's Wii lacks the processing power of the other game consoles but, thanks to its lower system specifications, it alone of this generation's consoles sells for more than it costs to make. On the downside, the Wii is also the only console that doesn't support HD video. However, the intuitive motion-based controller—which uses accelerometers and infrared to track the player's gestures in several simple but engaging games—has become the surprise winner of this generation's consoles.

Michael van Lent, *the Entertainment Computing column editor, is a research scientist at the USC Institute for Creative Technologies. Contact him at vanlent@ict.usc.edu.*

INVISIBLE COMPUTING

Doing the Mobile Mash

Jonathan Trevor, Yahoo!

The Pipes system makes it easy to create mobile-friendly mashups.

Mashups are designed to collect data that can't be found in any one place on the Web, or to present online information in a more usable or visually suitable manner.

Suppose you're looking for a new apartment. You could refine your search by criteria commonly supported by rental websites such as price range and the number of bedrooms, but what if you want to know something that often isn't included, such as how close an apartment is to a park?

You could go to each apartment listing in turn and then enter the address on a map, but that's both error-prone and extremely time-consuming. In fact, all the data needed to solve this task is already available on the Web, but in different places. For example, Craigslist provides comprehensive apartment listings available through RSS (Really Simple Syndication), while Yahoo! Local offers a REST (Representational State Transfer)-style API to locate attractions and businesses in a given area.

A technically savvy developer could handcraft a mashup, like HousingMaps or chicagocrime.org's crime maps, that fetches the apartment listings from various rental websites, parses the contents, uses another online service to geocode the locations, and finally plots these side-by-side on a map (with the park locations fetched from Yahoo! Local). However, there needs to be an easier way to fetch and manipulate this data.

PIPES

Yahoo! Pipes (http://pipes.yahoo.com) is a service platform for processing well-structured data formats such as XML, RSS, iCal (iCalendar), JSON (JavaScript Object Notation), or CSV (comma-separated values). Developers of all capabilities can use Pipes' Web-based visual programming environment to combine data sources and user input into mashups without having to host or write any code. These mashups, somewhat analogous to Unix pipes, can power badges on personal publishing sites, provide core functionality for Web applications, or serve as reusable components within the Pipes platform itself.

Modules and wires

Figure 1 shows the Pipes visual editor. Developers drag *modules* from a toolbox on the left onto the Pipes canvas. Modules provide high-level data-processing functions such as geocoding locations in feed items or translating from one language to another. Modules are configured directly on the canvas, and data flows through them via *wires*. Connecting a wire is as simple as dragging the output of one module into the input of another. Developers can inspect the data being processed at any point of the Pipe simply by clicking on the debugger module at the bottom of the page. No plug-ins or downloads are required for the editor, which opens seamlessly in a Web browser.

Using Pipes, you could solve the apartment search problem in four easy steps:

- obtain the apartment listings using a "fetch RSS feed" module;
- connect the fetch output to a location extractor module to geocode the apartment locations;
- loop over each of the locations and search Yahoo! Local for the nearest park; and
- sort the results by the distance the apartment is from the park.

This Pipe, hosted on the Pipes website, can provide an RSS or JSON feed that updates as Craigslist posts new apartments within a certain distance of a park. Pipes can send the feed to you via your favorite feed reader or SMS (short message service) alert when new listings come up or your application of choice can fetch it.

Mashup 'buckets'

While Pipes is a developer- rather than consumer-oriented tool, a central design goal was to attract a broad range of users, from experienced developers to those with limited or no development skills at all. Along with the easy-to-use visual data editor, the system features a "view source" mechanism inspired by that in HTML that lets you

inspect any public Pipe on the site (to see how it works) and copy it as your own (to modify at your leisure) with just a single click.

This accessibility has enabled numerous developers of all abilities to create Pipes spanning a wide range of topics including religion, politics, news, weather, photo sharing, and deal hunting. While a mashup's nature prevents clean classification, these Pipes can broadly fit into four "buckets."

Feed aggregation with filtering. The "Hot Deals Search" Pipe aggregates 10 different website feeds related to bargain hunting on the Web into one place, allowing users to filter the list with keywords.

Two-source mashups. The Pipe "New York Times thru Flickr" combines a primary source of data (news headlines) with an additional stream of information (pictures about that item).

Data transformation and geocoding. The "Hurricane Katrina with pressure and wind speed" Pipe converts a CSV file into an RSS feed that can be plotted on a map. "Geo-Annotate Reuters News" geocodes news data that can be used to display stories on a world map. Babbler is a Second Life plug-in that translates instant messages between languages by running various translation Pipes. These types of Pipes are used by applications but often aren't visible to the user.

Complex mashups using REST APIs. Flickr has no automatic way of grabbing your friends' favorite pictures, but the "Friend Flickr Photos" Pipe fetches the last five favorite pictures of your friend using a combination of Flickr API calls, with no code, ready for your favorite RSS reader.

System implementation

Pipes has a cluster of back-end engines responsible for executing pipe definitions stored in a relational database. The engines fetch the data that a Pipe requires through an HTTP cache to ensure high performance and prevent overloading of remote sources. The system returns a Pipe "run" in a language-independent format via another layer of HTTP caches to the front-end Web servers.

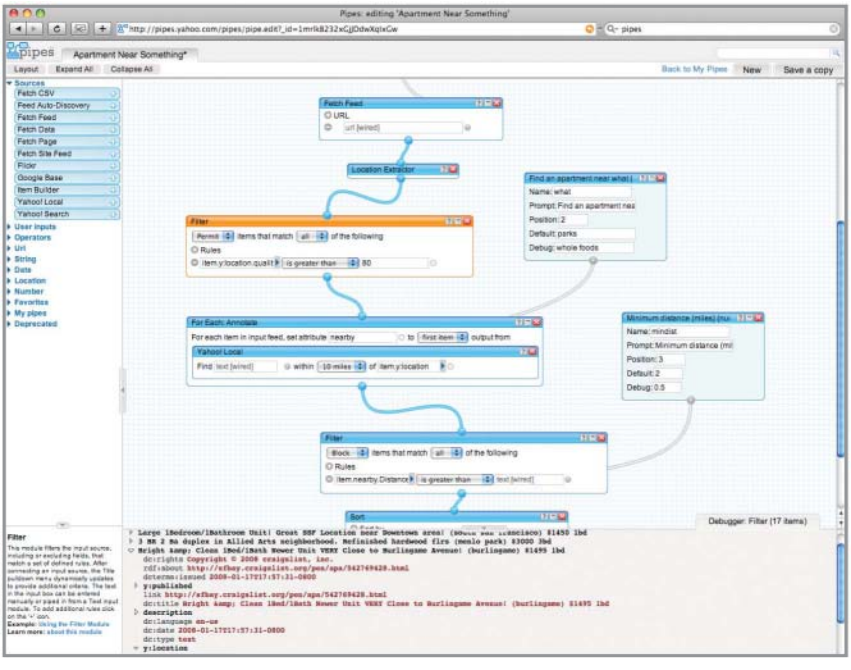

Figure 1. Yahoo! Pipes visual editor. Developers drag modules from a toolbox on the left onto the Pipes canvas. Modules are configured directly on the canvas, and data flows through them via wires.

In addition to providing the various pages that make up the Pipes website and visual editor, the front-end servers can transform the run data into numerous formats.

All caches extend the default HTTP semantics with two additional capabilities: *stale content serving*, to return a cache hit even when it's stale but continues to run the request asynchronously to update the entry; and *request collapsing*, to forward one request for any number of identical requests arriving at the cache at once.

GOING MOBILE

Mobile devices' capabilities have dramatically increased in recent years, offering bigger displays, multimodal input, and faster network connections. However, providing the right information at the right time continues to be a challenge for mobile applications.

Because mashups often are conceived with the goal of collecting data from various sources for easier consumption, it's not surprising that many mashups are useful to mobile users. Indeed, geocoded output, for displaying items on a map, remains one of the most popular components of Pipes (and mashups in general).

Enter the iPhone

Apple's iPhone, with its easy-to-use built-in mapping application, provided Yahoo! developers with an ideal platform to move Pipes mashups from the desktop to the mobile domain.

The Pipes website provides iPhone users with a list-centric interface (http://iphone.pipes.yahoo.com) consistent with native applications that run in the device's browser. As Figure 2 shows, any Pipe—a user's own or someone else's—can be run simply by selecting it from the list. The system initially presents the results in a similar compact list format. Pipes that geocode their results, via a single module in the visual editor, can be directly displayed on the built-in mapping application by touching the Map button.

Figure 2. Using Pipes on the iPhone. (a) Pressing the "Restrooms nearby" entry will prompt a query for the user's address. (b) After the user inputs this information, Pipes generates a list of businesses ordered by distance that are likely to have public restrooms; selecting an item from the list opens a webpage with information about that business. (c) Pressing the Map button on the bottom left plots the same results on the iPhone's built-in map.

Mobile Pipes

The Pipes system makes it easy to create numerous mobile applications that have previously been limited to one-off solutions. For example: The "Yahoo! Local by Rating" Pipe searches Yahoo! Local to find the "best" restaurants; "GasBuddy.com Gas Prices" locates the cheapest gas nearby; and "Live Traffic Results" plots detailed information about traffic problems on a map.

Not all useful mobile Pipes are map-related. The "Price Comparison" Pipe lets you determine whether a particular item you're looking at is a good value by searching for its price at nearby stores, or whether you could buy it used locally.

Pipes can also be used on many Internet-enabled mobile devices. It produces outputs in two of the most common geocoded formats—GeoRSS (Geographically Encoded Objects for RSS) and KML (Keyhole Markup Language)—so mobile mapping applications that support either can also display any Pipe results.

Pipes lets both novice and expert developers grab Web data sources from multiple formats, manipulate that data, mash it up with other data sources or services using a visual editor, and host the results on its webpage. Once a Pipe is built, anyone can use its data output however they want, wherever they want.

Mobile application development continues to be challenging given the enormous number of devices and capabilities. Services like Pipes enable developers to focus on one of the keystones of such applications: getting the right data. They also make list- and map-based visualizations easy to both build and consume.

A current shortcoming of many mobile Pipes is the lack of an automatic way to provide user location. However, as location brokerage services like Yahoo! Fire Eagle (http://fireeagle.research.yahoo.com) become available, this information becomes just another source of data to be mashed up in Pipes.

Jonathan Trevor is an architect/developer at Yahoo!. Contact him at jtrevor@yahoo-inc.com.

Editor: Bill N. Schilit, Google; bill.schilit@computer.org

Continued from page 108

contacted, the Russians replied simply that they used a pencil.

When I look at the history of science, many instances in which outsiders have brought novel ideas to various disciplines come to mind.

The Frenchman Louis de Broglie predicted the wave nature of electrons and created the field of wave mechanics. Before him, scientists knew that waves could behave like particles, as Albert Einstein described in his famous theory of the photoelectric effect. On the other hand, nobody, including Einstein, had imagined that the dual phenomenon of particles behaving like waves could also exist. De Broglie predicted this in 1924, and Clinton Davisson and Lester Halbert Germer confirmed it experimentally in 1927 by observing electron diffraction with crystals. Surprisingly, de Broglie was not a physicist by training, but a historian who had graduated in literary studies. Physics was simply a hobby he learned much later in life.

Several others independently discovered the conservation of energy law. One of this field's earliest researchers, Julius Robert von Mayer, was not a physicist but a surgeon. Von Mayer observed that the color of his Dutch East Indian patients' blood was a deeper shade compared to his European patients. From this he conjectured that they absorbed less oxygen because maintaining body temperature in a hotter climate requires less energy. Von Mayer conjectured that heat and mechanical work were both forms of energy, and later, after improving his knowledge of physics, calculated a quantitative relationship between them, leading to his discovery of the energy conservation law.

It would seem that such discoveries might be less likely in developing countries. Not so, as another example shows, this time from the field of economics and finance. In the past few years, the notion of micro loans and micro credits has become one of the brightest ideas to emerge in the world of banking and finance. Their inventor, Muhammed Yunus, is from Bangladesh, one of the world's poorest countries. Yet Yunus's ideas have proven so effective in fighting poverty that he has been appointed to the World Bank's advisory committee and received a Nobel Prize.

Georges Clemenceau once observed that "war is too important a matter to be left to the generals." Similarly,

It is not enough to talk about our new ideas; it is equally important that we listen to those of others.

parallel computing is too important to be left to the computer scientists. Those with a background in another field—such as physics, mathematics, or chemistry—can still make a worthwhile contribution to research on parallel computing. New ideas often emerge from novel applications.

LISTENING IN CONFERENCES

It is not enough to talk about our new ideas; it is equally important that we listen to those of others. Most of us are fond of talking. If allowed to stand in one place for two days, I could keep talking without tiring. But when I sit in the audience at a conference and listen, I tire within 10 minutes.

This reminds me of an incident I heard about from a friend. The traffic police in various cities noticed that accidents often occurred involving public transport buses serving the airports. Investigators traced a common reason for these accidents to drivers being distracted by passengers, particularly tourists, asking them for directions or other information.

Consequently, the bus companies decided to post signs on their vehicles to prevent this. All the cities, devised signs with similar messages. For instance, on one city's buses the signs read, "Please do not speak to the driver." Another town's buses had signs that read, "You are kindly requested to refrain from speaking to the driver." In yet another town the signs read, "It is strictly forbidden to speak to the driver." However, in one particular town, the bus signs read, "Please do not answer the driver," because the drivers in that locale would start speaking to the passengers first, as they boarded the bus.

Talking is easy, listening far more difficult. Not surprisingly, then, most of us lack the quality of listening, but we can develop this skill by attending conferences. As teachers, we can also engage in a kind of role reversal, experiencing how our students feel while listening to us lecturing to them on subjects with which they are unfamiliar.

Quite often we find low attendance at conference sessions for contributed papers—people attend only the sessions with their own talks, then leave. Many attendees tell me they feel sleepy while listening to the lectures. I tell them it is perfectly fine to sleep during some of the talks—an energizing siesta such as this can fortify attendees and improve their alertness. For example, Peter Medawar, who received a Nobel Prize in medicine for his studies on organ transplantation, writes in one essay that lectures provide him with his most refreshing sleep.

I have the same experience: Give me a five-star hotel room, a comfortable bed, and a long night of sleep, and I would still wake up feeling tired. But if I can sleep for a few minutes during a lecture, I wake up feeling completely refreshed and ready to focus on the remaining lectures.

***Raja Natarajan** is a researcher at the Tata Institute of Fundamental Research, Mumbai, India. Contact him at raja@tifr.res.in.*

Editor: Neville Holmes, School of Computing and Information Systems, University of Tasmania; neville.holmes@utas.edu.au.

THE PROFESSION

On Attending Conferences

Raja Natarajan
Tata Institute of Fundamental Research

Talking is easy, listening far more difficult.

As a researcher and academician, I get many opportunities to give technical lectures on various aspects of computer science. On the other hand, a formal occasion to share other important but nontechnical ideas rarely occurs. Recently, I received one such opportunity quite unexpectedly.

I was attending a two-day conference on parallel and distributed computing at which surprisingly many of the participants were teachers from various undergraduate colleges. Although the organizers had initially invited me to give a technical lecture, they later requested that I deliver the inaugural address as well.

I approached the dais with great trepidation because I had little experience in giving nontechnical lectures. Fortunately, by the time I completed my preamble with a quick overview of the various planned technical sessions, and congratulated the organizers for their good work in planning the conference program so well, I had made up my mind that my talk's theme would be to motivate the audience to meaningfully participate in more conferences. What follows is an annotated transcript of my lecture.

WHY ATTEND CONFERENCES?

As teachers and academicians, we must keep ourselves informed about the latest developments in our field and related areas. Attending conferences and meeting others can help us learn about recent ideas and advances. As the *Rig Veda* says, "Let noble thoughts come to us from all directions." We can make this happen by attending conferences and meeting other people, which exposes us to ideas from all directions.

Some might think they can keep themselves informed by reading books and journals. But learning from these sources offers greater limitations compared to how much can be learned from listening to others. As the organizers for a conference have taken the initiative to arrange the event, and likely have done their job well, the participants have the opportunity to take an active part in making the conference a success.

SPEAKING IN CONFERENCES

Conference attendees should arrive prepared. They should try to contribute at least one lecture to each conference they attend. Such lectures require providing material beyond that available in standard textbooks. Presenters must instead lecture on what they have actively thought about and worked on—their research projects. This raises the more basic question of why we should do research at all.

Some instructors might say that their job is to teach, and as long as they do that, they needn't bother doing research at all. Granted, teaching is a necessary part of our job, but it isn't sufficient to just teach, nor is it enough to just keep reading about what others have already discovered or invented. We must seek to bring our own new perspective to things.

As John Milton noted in *Paradise Regained*, every person who reads without engaging in imaginative thinking simply becomes "deep-versed in books and shallow in himself." There is a great danger of becoming shallow if we abstain from research. Researching an advanced field at the edge of conventional knowledge, such as parallel computing, can prove daunting, yet often complete newcomers can bring novel ideas to the discipline.

OUTSIDERS HAVE ORACLES

An anecdote from the early years of the NASA space program shows how unconventional thinking can lead to radically different solutions. Soon after NASA's early space flights, its engineers encountered a problem: In a time before pocket digital diaries, the astronauts could not write properly during their space flights because none of the existing pens functioned properly in outer space. Under zero gravity and low-pressure conditions, ink flow in the pens would be uneven, making smooth writing impossible.

Two years of research and several projects costing many million of dollars later, NASA's engineers finally succeeded in making a pen that could write in outer space. At this point, the American engineers wondered if the Russians had yet solved the same problem. When

Continued on page 107

Important Dates

Workshop Proposals:
March 1, 2008
Paper Registration:
March 8, 2008
Paper submission:
March 15, 2008
Notification of acceptance:
May 4, 2008
Camera-ready version due:
May 18, 2008

General Chair
Vinny Cahill
Trinity College Dublin, Ireland

Steering Committee Chair
Imrich Chlamtac
Create-Net

Program Chair
Liviu Iftode
Rutgers University, USA

Workshop Chairs
Christian Becker
Universitat Mannheim, Germany
Fabian E. Bustamante
Northwestern University, USA

Posters Chair
Dean Eckles
Nokia Research, USA

Demonstrations Chairs
Tatsuo Nakajima
Waseda University, Japan
Alan Messer
Samsung Electronics, USA

MobiQuitous 2008

July 21st – 25th, 2008, Dublin, Ireland

The Fifth Annual International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services

http://www.mobiquitous.org

Sponsored by ICST
Technically-sponsored by Create-Net

The Fifth Annual International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous 2008) will provide a forum where practitioners and researchers coming from many areas involved in ubiquitous solutions design and deployment will be able to interact and exchange experiences needed to build successful ubiquitous systems. Areas addressed by the conference include systems, applications, service-oriented computing, middleware, networking, agents, data management and services, all with special focus on mobility and ubiquitous computing.

PAPERS: Technical papers describing original, previously unpublished research, not currently under review by another conference or journal, are solicited. Technical papers clearly identifying how the specific contributions fit to an overall working system are particularly of interest. Topics include, but are not limited to, the following:

- **Ubiquitous architectures, systems and applications**
- **Wearable computing and personal area networks**
- **Wireless technologies (Bluetooth, ZigBee, 802.15.x, WiFi, WiMAX)**
- **Wireless Internet access in ubiquitous systems**
- **Ad hoc and sensor network support for ubiquitous computing**
- **Reconfigurability and personalization of wireless networks**
- **Wireless/mobile service management and delivery**
- **Security, privacy and social issues of mobile and ubiquitous systems**
- **Service and knowledge discovery, matching and composition mechanisms**
- **Location-based services and tracking in ubiquitous environments**
- **Context- and location-aware applications**
- **Agent technologies in ubiquitous, wearable, and mobile systems**
- **Context modeling, services and frameworks**
- **Toolkits, testbeds, development environments, and languages for ubiquitous computing**
- **Rapid prototyping of ubiquitous applications**
- **Ontologies for mobile and ubiquitous computing**
- **Mobile and ubiquitous data management and processing**
- **Data replication, migration and dissemination in ubiquitous environments**
- **Queries, transactions and workflows in mobile and ubiquitous environments**
- **Multimodal interfaces (speech, video kinetic, tactile)**

Important Dates

Paper Registration:
March 8, 2008
Paper submission:
March 15, 2008
Notification of acceptance:
May 4, 2008
Camera-ready version due:
May 18, 2008

General Chair
René Maier
Trinity College Dublin, Ireland
Eylem Ekici
Ohio State University, USA

Program Chair
Raj Rajkumar
Carnegie Mellon University, USA

Steering Committee Chair
Imrich Chlamtac
Create-Net

Posters/Demos Co-Chairs
Raja Sengupta
UC Berkeley, USA
Sam Reisenfeld
UTS, Australia

Industry Chair
Fan Bai
General Motors, USA

ISVCS 2008

July 22nd – 24th, 2008, Dublin, Ireland

The Fifth Annual International Symposium on Vehicular Computing Systems

http://www.isvcs.eu

Sponsored by ICST
Technically-sponsored by Create-Net

The ISVCS is an annual event that seeks to bring together people from academia, industry and government interested in building and using vehicular computing systems, middleware, protocols, services and applications. The symposium is primarily interested in papers reporting innovative research resulting in real implementations and working prototypes. Work in progress and new ideas will also be accepted as well as presentations of new products.

Topics of interest include:

- **Design and implementation of computer systems for network-connected vehicles**
- **Middleware and service architectures for safety, road sensing, route planning, in-car networking and other vehicle-centric applications**
- **Protocols for data collection and dissemination over network-connected cars.**
- **Security and privacy issues in vehicular networks, systems, services and applications**
- **Location-aware computing models for vehicular systems**
- **Driver-to-computer interfaces**
- **Data management systems for road sensing and traffic monitoring**
- **Vehicle-to-vehicle and vehicle-to-infrastructure communication protocols**
- **Inter-vehicle distributed systems for entertainment and gaming**
- **Pervasive computing applications using vehicular networks**
- **Tools and methodology for vehicular computing systems verification and evaluation**
- **Fault-tolerance solutions for vehicular computing systems**
- **Maintenance for vehicular computing software**
- **Experience reports of testing vehicular computing systems on real-world conditions**
- **Policies, laws and regulations for adopting vehicular computing technologies**
- **Social and psychological implications of vehicular computing applications.**