Innovative Technology for Computer Professionals

Computer

February 2004

Computer Lessons From Pop Songs

Location-Aware Computing

Smart Toys

ad hoc networks

IEEE

IEEE Computer Society

http://www.computer.org

New and Noteworthy

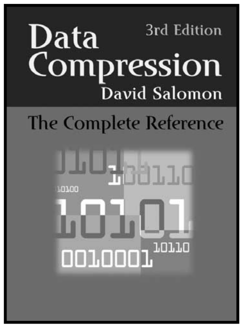

New Edition

Data Compression

The Complete Reference

Third Edition

DAVID SALOMON

From a review of a previous edition:

"A wonderful treasure chest of information; spanning a wide range of data compression methods, from simple text compression methods to the use of wavelets in image compression."

—COMPUTING REVIEWS

This new edition provides a comprehensive reference for the many different types and methods of compression. The presentation is organized into the main branches of the field of data compression: run length, encoding, statistical methods, dictionary-based methods, image compression, audio compression, and video compression

2004/882 PAGES/HARDCOVER/ISBN 0-387-40697-2/$69.95

Web Data Management

A Warehouse Approach

SOURAV BHOWMICK, SANJAY MADRIA, *and* **WEE KEONG NG**

This comprehensive resource presents a data model called WHOM (Warehouse Object Model) to represent HTML and XML documents in the warehouse. It defines a set of web-algebraic operators for building new web tables by extracting relevant data from the Web, as well as generating new tables from existing ones. Its web-warehouse approach incorporates modern and effective shared data-management concepts, methods, and models

2003/470 PAGES/HARDCOVER/ISBN 0-387-00175-1/$69.95

Guide to Biometrics

RUUD BOLLE, JONATHAN CONNELL, SHARATCHANDRA PANKANTI, NALINI RATHA, *and* **ANDREW SENIOR**

This complete, technical guide presents the core ideas underlying biometrics. It explains the definition and measurement of performance and examines the factors involved in choosing between different biometrics. It also delves into practical applications and covers a number of topics critical for successful system integration.

2004/335 PAGES/HARDCOVER/ISBN 0-387-40089-3/$49.95

Practical Biometrics

From Aspiration to Implementation

JULIAN ASHBOURN

This book concentrates on the practical implementation of biometric verification techniques, with specific regard to wide scale public applications. It details the associated issues around: scalability; interoperability, both from a technical as well as operational perspective; ethnicity and the implications for biometric verification checks; and failure to enroll and other error conditions. The author also highlights non-device specific issues such as human factors, privacy, and data protection.

2004/176 PAGES./HARDCOVER/ISBN 1-85233-774-5/$49.95

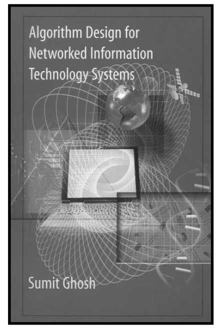

Algorithm Design for Networked Information Technology Systems

SUMIT GHOSH

This book presents a scientific theory of Networked information technology (NIT) systems and logically develops the fundamental principles to help synthesize control and coordination algorithms for these systems. The book explains through case studies the conception, development, experimental testing, validation, and rigorous performance analysis of practical asynchronous, distributed decision-making (ADDM) algorithms for real-world systems from a number of diverse disciplines.

2004/286 PAGES/HARDCOVER/ISBN 0-387-95544-5/$69.95

Verification of Reactive Systems

Formal Methods and Algorithms

KLAUS SCHNEIDER

This book is devoted to the foundation of the most popular formal methods for the specification and verification of reactive systems. In particular, the μ-calculus, omega-automata, and temporal logics are covered in full detail; their relationship and state-of-the-art verification procedures based on these formal approaches are presented. Furthermore, the advantages and disadvantages of the formalisms from particular points of view are analyzed.

2004/600 PAGES/HARDCOVER/ISBN 3-540-00296-0/$79.95

Evolvable Components

From Theory to Hardware Implementations

LUKAS SEKANINA

This book deals with adaptive evolvable systems operating in a changing environment. In particular, topics such as the computational power of evolvable systems, the design of complex human-competitive digital circuits and the implementation of a one-clock partial reconfiguration in a common field-programmable gate array are introduced, investigated and implemented. The application domain is adaptive image preprocessing.

2004/194 PAGES/HARDCOVER/ISBN 3-540-40377-9/$64.95

Introduction to Evolutionary Computing

A.E. EIBEN *and* **J.E. SMITH**

This book presents a complete overview of Evolutionary Computing (EC) aimed directly at lecturers and graduate and undergraduate students. It is also meant for those who wish to apply evolutionary computing to a particular problem or within a given application area. To this group the book is valuable because it presents EC as something to be used rather than just studied. The book also contains quick-reference information on the current state-of-the-art in a wide range of related topics.

2003/299 PAGES/HARDCOVER/ISBN 3-540-40184-9/$49.95

Easy Ways to Order:

CALL: Toll-Free 1-800-SPRINGER • **FAX**: 201-348-4505 • **WEB**: www.springer-ny.com • **E-MAIL**: orders@springer-ny.com • **WRITE**: Springer-Verlag New York, Inc., Dept. S5462, PO Box 2485, Secaucus, NJ 07096-2485 • **VISIT** your local scientific bookstore or urge your librarian to order.

Prices subject to change without notice.

Springer

www.springer-ny.com

2/04

PROMOTION #S8002

A Symposium on High-Performance Chips

HOT CHIPS 16

Stanford University, Palo Alto, California
August 22-24, 2004

CALL FOR CONTRIBUTIONS

For the last 15 years, Hot Chips has been the leading conference on high-performance microprocessors and related integrated circuits. The conference is held once a year in August on the Stanford University campus in the center of the world's electronics capital, Silicon Valley. The emphasis this year, as in previous years, is on real products and realizable technology. Topics of interest for this year's conference include but are not limited to:

- **Microprocessors**
- **Systems-on-chip**
- **Embedded processors**
- **Digital signal processors**
- **Application-specific processors**
- **Network/security processors**
- **Graphics/Multimedia/Game processors**
- **Communication/networking chips**
- **Wireless LAN/Wireless WAN chips**
- **Novel chips: quantum computing, microarray**
- **Low-power chips/Dynamic Power Management**
- **Reconfigurable chips/processors**
- **Reliability and design for test**
- **Compiler technology**
- **Operating system/chip interaction**
- **Advanced semiconductor process technology**
- **Advanced packaging technology**
- **Performance evaluation**

Presentations at HOT Chips are in the form of 30-minute talks. 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 must consist of a title, extended abstract (two pages maximum), and the presenter's contact information (name, affiliation, job title, address, phone, 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; we do our best to maintain confidentiality.

Submissions are evaluated by the Program Committee on the basis of the performance of the device (or devices), 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. Authors will be notified of the status of their submission by the end of April, 2004.

Don't miss this chance to present your work to an audience of engineers, computer architects, and computer system and device researchers. Submissions must be received no later than **March 15, 2004**.
Please make your submissions in plain ascii text (in the message, not as an attachment) to:
hotchips-submission@cva.stanford.edu
(Submissions containing figures may be submitted in pdf, but plain ascii is preferred.)
For more information, see the Hot Chips 16 Web site at: **http://www.hotchips.org**
Send questions to **hotchips@cva.stanford.edu**
or contact the co-program chairs:
Prof. Bill Dally at **hotchips@cva.stanford.edu**, or
Keith Diefendorff at **keithd@mac.com**

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

Program Committee Co-Chairs
Prof. Bill Dally, Stanford University
Keith Diefendorff, MemoryLogix

IEEE COMPUTER SOCIETY

***Check the HOT CHIPS 16 web page for updates:* http://www.hotchips.org**

IPDPS 2004

Monday, 26 April – Friday, 30 April 2004

Eldorado Hotel • Santa Fe, New Mexico USA

IPDPS 2004 Santa Fe, New Mexico International Parallel & Distributed Processing Symposium

Sponsored by the IEEE Computer Society Technical Committee on Parallel Processing

In cooperation with ACM SIGARCH, IEEE Computer Society Technical Committee on Computer Architecture (TCCA), and IEEE Computer Society Technical Committee on Distributed Processing (TCDP)

Hosted by University of New Mexico

IPDPS 2004 PROGRAM IS SET

From nearly 450 submissions, 142 papers have been accepted for presentation in 24 Technical Sessions. The scheduled program below includes 17 Workshops and 3 Tutorials as well as an Industrial Track on Wednesday and Birds-of-a-Feather meetings throughout the week. For more information on the advance program and registration, visit the IPDPS Web site at www.ipdps.org or send email to info@ipdps.org.

Monday, April 26

- Workshops

Heterogeneous Computing Workshop
Workshop on Parallel and Distributed Real-Time Systems
Reconfigurable Architectures Workshop
Workshop on High-Level Parallel Programming Models & Supportive Environments
Java for Parallel and Distributed Computing Workshop
Workshop on Nature Inspired Distributed Computing
Advances in Parallel and Distributed Computing Models
Communication Architecture for Clusters
Workshop on High Performance Computational Biology
NSF Next Generation Systems Program Meeting

- Tutorial 1: High Performance Computing & Networking Using Infiniband Technology – Hardware and Software Solutions
- Tutorial 2: An Introduction to Distributed Data Mining
- TCPP Reception Speaker: Srinidhi Varadarajan, Terascale Computing Facility, Virginia Tech - System X: The Virginia Tech Supercomputer

Tuesday, April 27

- Keynote: Ken Kennedy, Rice University - High Performance (and Grid) Computing Without a Degree in Computer Science
- Technical Sessions 1 through 8
- Panel 1: Bioinformatics and High Performance Computing - Moderator: David A. Bader, University of New Mexico

Wednesday, April 28

- Keynote: Michel Dubois, University of Southern California - Are We Entering the Golden Age of Parallel Processing? Finally?
- Technical Session 9: Best Papers
- Technical Sessions 10 through 12
- Industrial Track
- Panel 2: Internet Computing - Moderator: Thierry Priol, IRISA Paris
- Symposium Banquet Speaker: Christos Papadimitriou, University of California Berkeley - Networks and Games

Thursday, April 29

- Keynote: Reiner Hartenstein, University of Kaiserslautern, Germany - Software or Configware? About the Digital Divide of Parallel Computing
- Technical Sessions 13 through 24

Friday, April 30

- Workshops

Workshop on Fault-Tolerant Parallel and Distributed Systems
International Workshop on Wireless. Mobile, Ad Hoc and Sensor Networks
Workshop on Parallel and Distributed Scientific and Engineering Computing with Applications
Performance Modeling, Evaluation, and Optimization of Parallel and Distributed Systems
Workshop on Massively Parallel Processing
Parallel and Distributed Systems:Testing and Debugging
High-Performance Grid Computing

- Tutorial 3: Parallel Computing on Heterogeneous Networks

GENERAL CO-CHAIRS

David A. Bader, University of New Mexico
Michel Cosnard, Université de Nice & INRIA Sophia Antipolis, France

GENERAL VICE CHAIR

Charles Weems, University of Massachusetts at Amherst

PROGRAM CHAIR

Burkhard Monien, University of Paderborn, Germany

For more information visit the IPDPS Web site at www.ipdps.org or send email to info@ipdps.org.

Innovative Technology for Computer Professionals
Computer

February 2004, Volume 37, Number 2

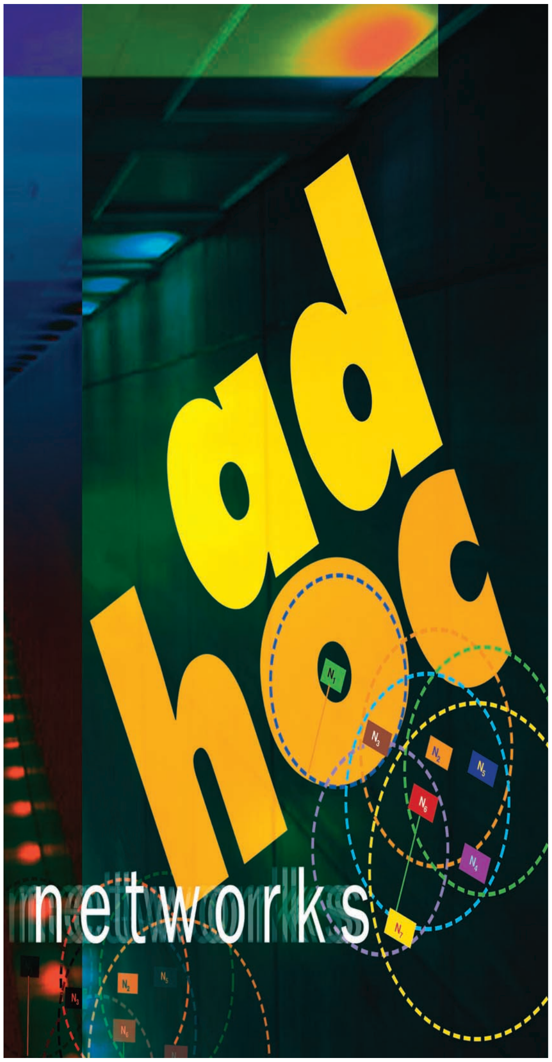

Cover design and artwork by Dirk Hagner

ABOUT THIS ISSUE

Ad hoc networking encompasses wireless multihop scenarios in which network nodes communicate via other network nodes such as in conference, healthcare, battlefield, rescue, and monitoring situations. This issue focuses on recent developments that offer potential solutions to problems encountered in ad hoc networks including topology control, data communication, and service access.

PERSPECTIVES

COVER FEATURES

Computer

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

OPINION

NEWS

MEMBERSHIP NEWS

COLUMNS

NEXT MONTH:
Hardware Speculation

DEPARTMENTS

Membership Magazine of the IEEE Computer Society

IEEE

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 $17 for a subscription to *Computer* magazine. Nonmember subscription rate available upon request. Single-copy prices: members $10.00; nonmembers $20.00.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Canadian GST #125634188. Canada Post Publications Mail (Canadian Distribution) Agreement Number 0487910. Printed in USA.

A Copper Bullet for Software Quality Improvement

pp. 21-25

Michael Blaha

Fred Brooks long ago observed that there is no silver bullet to improve software quality. However, there are copper bullets—lesser steps that improve quality over time. One of these is the notion of software engineering, the practice of thinking carefully before immersing yourself in the minutia of coding.

The author offers a summary of his experience indicating that database reverse engineering offers a quality improvement strategy that could benefit the entire software community.

Cooperative Cache-Based Data Access in Ad Hoc Networks

pp. 32-39

Guohong Cao, Liangzhong Yin, and Chita R. Das

Cooperative caching, in which multiple nodes share and coordinate cached data, is widely used to improve Web performance in wired networks. However, resource constraints and node mobility have limited the application of these techniques in ad hoc networks. The authors propose caching techniques that use the underlying routing protocols to overcome these constraints and further improve performance.

Energy-Efficient Area Monitoring for Sensor Networks

pp. 40-46

Jean Carle and David Simplot-Ryl

The nodes in sensor networks must self-organize to monitor the target area as long as possible. Researchers at the Fundamental Computer Science Laboratory of Lille are developing strategies for selecting and updating an energy-efficient connected active sensor set that extends the network lifetime. The authors report on their work to optimize energy consumption in three separate problems: area coverage, request spreading, and data aggregation.

Cross-Layering in Mobile Ad Hoc Network Design

pp. 48-51

Marco Conti, Gaia Maselli, Giovanni Turi, and Silvia Giordano

Mobile ad hoc network researchers face the challenge of achieving full functionality with good performance while linking the new technology to the rest of the Internet. A strict layered design is not flexible enough to cope with the dynamics of manet environments, however, and will prevent performance optimizations.

The MobileMan cross-layer architecture offers an alternative to the pure layered approach that promotes stricter local interaction among protocols in a manet node.

Group Communications in Mobile Ad Hoc Networks

pp. 52-59

Prasant Mohapatra, Chao Gui, and Jian Li

Efficient support of group communications is critical for most ad hoc network applications. However, manet group communications issues differ from those in wired environments because the wireless communications medium has variable and unpredictable characteristics, and the signal strength and propagation fluctuate with respect to time and environment.

The authors provide an overview of group communications issues such as protocol design, state maintenance, and performance. They also discuss representative protocols and examine related issues such as reliability, power conservation, quality of service, security, and areas requiring additional research.

Routing and Security in Mobile Ad Hoc Networks

pp. 61-65

Nikola Milanovic, Miroslaw Malek, Anthony Davidson, and Veljko Milutinovic

Mobile ad hoc networks remove the dependence on a fixed network infrastructure by treating every available mobile node as an intermediate switch, thereby extending the range of mobile nodes well beyond that of their base transceivers.

The authors present four manet routing algorithms along with a hybrid approach, discuss their advantages and disadvantages, and describe security problems inherent in such networks.

Prioritized Overlay Multicast in Mobile Ad Hoc Environments

pp. 67-74

Li Xiao, Abhishek Patil, Yunhao Liu, Lionel M. Ni, and Abdol-Hossein Esfahanian

Many proposed routing protocols for manets require nodes to maintain and update complicated route information, which incurs significant overhead when groups have different priorities.

To address this problem, some researchers have begun focusing on application-layer, or overlay, multicast in which an overlay network forms a virtual network consisting of only member nodes atop the physical infrastructure. The authors propose a prototype of prioritized overlay multicast for manets in which participating nodes can carry out multiple functions and thus be associated with more than one overlay tree.

The 2004 IEEE International Conference on Information Reuse and Integration

IEEE IRI-2004 Call for Papers

November 1-3, 2004
Luxor Hotel and Resort, Las Vegas, Nevada, USA
http://www.cs.fiu.edu/IRI04

IEEE IRI 2004

Sponsored by the IEEE Systems, Man and Cybernetics Society

Honorary General Chair
Lotfi Zadeh
University of California, USA
zadeh@cs.berkeley.edu

General Chairs
Stuart Rubin
SPAWAR Systems Center, USA
stuart.rubin@navy.mil

Shu-Ching Chen
Florida International University, USA
chens@cs.fiu.edu

Program Chairs
Du Zhang
California State University, USA
zhangd@ecs.csus.edu

Eric Grégoire
Université d'Artois, France
gregoire@cril.univ-artois.fr

Doug DeGroot
Leiden University, The Netherlands
degroot@acm.org

Local Arrangements Chair
Louellen McCoy
SPAWAR Systems Center, USA
louellen.mccoy@navy.mil

Finance Chair
David Barnett
Saint Louis University, USA
barnettd@slu.edu

Registration Chair
Willie Chang
California State University, USA
changw@ecs.csus.edu

Publications Chair
Atif Memon
University of Maryland at College Park, USA
atif@cs.umd.edu

Publicity Chair
Julie A. Adams
Vanderbilt University, USA
julie.a.adams@vanderbilt.edu

Student Activities Liaison
Mei-Ling Shyu
University of Miami, USA
shyu@miami.edu

Web Master
Na Zhao, Guo Chen
Florida International University, USA

European Liaison
Eric Grégoire

Asian Liaison
Doug DeGroot

This year's conference theme addresses all aspects of **Knowledge Management**. This theme was selected to reflect the pivotal role of reuse in the capture, maintenance, integration, validation, extrapolation, and application of domain knowledge. The IEEE International Conference on Information Reuse and Integration will feature contributed as well as invited papers. Theory, speculative, and application papers are all included in this call. The conference program will include tutorials and an open forum workshop. The focus of the conference includes, ***but is not limited to***, the areas listed below:

- Component-Based Design
- Data Mining and Knowledge Discovery
- Decision Support Systems
- HW & SW Engineering for Reuse
- Fuzzy and Neural Systems
- Large Scale Data Integration
- Human-Machine Information Systems
- Reuse in Software Engineering
- Intellectual Property
- Information Assurance
- Soft computing
- Knowledge Acquisition and Management
- Evolutionary Computing
- Case-Based Reasoning
- Model-Based Reasoning
- Multimedia Reuse & Integration
- Internet Computing
- Knowledge Management and E-Government
- Agent-Based Systems
- Sensory and Information Fusion
- Natural Language Understanding

Instructions for Authors:

Papers reporting original and unpublished research results pertaining to the above and related topics are solicited. Full paper manuscripts must be in English and should not exceed 6 pages (using the IEEE two-column template). Submissions should include the title, author(s), affiliation(s), e-mail address(es), tel/fax numbers, abstract, and postal address(es) on the first page. Papers should be submitted at the conference web site: http://www.cs.fiu.edu/IRI04. If web submission is not possible, manuscripts should be sent as an attachment via email to one of the Program Chairs listed below on or before the deadline date of June 1, 2004.

Dr. Du Zhang	Dr. Eric Grégoire	Dr. Doug DeGroot
Department of Computer Science	CRIL-CNRS	Digital Life Technologies
California State University	Université d'Artois	Dept of CS (LIACS)
6000 J Street	Rue Jean Souvraz	Leiden University
Sacramento, CA 95819-6021, USA	SP 18	Niels Bohrweg 1
zhangd@ecs.csus.edu	F-62307 Lens Cedex	2333 CA Leiden
1-916-278-7952 (office)	France	The Netherlands
1-916-278-6774 (fax)	gregoire@cril.univ-artois.fr	degroot@acm.org

The attachment must be in .pdf (preferred) or word.doc format. The subject of the email must be "IEEE IRI 2004 Submission." Papers will be selected based on their originality, timeliness, significance, relevance, and clarity of presentation. Authors should certify that their papers represent substantially new work and are previously unpublished. Organizers of prospective special sessions, panels, and tutorials are invited to submit proposals and should contact the Program Chair directly as soon as possible, but no later than April 15, 2004. Submission implies the intent of at least one of the authors to register and present the paper, if accepted. Authors of selected papers, presented at the conference, will be invited to submit an expanded version of their paper for review for publication in the appropriate IEEE SMC Transactions.

Important Dates:

Special Sessions, Panels, and Tutorials Deadline:	**April 15, 2004**
Full Paper Submission received by:	**June 1, 2004**
Notification of Acceptance:	**July 31, 2004**
Preregistration & Revised Camera-Ready Paper Due:	**September 14, 2004**

IMPROVING SOFTWARE PROGRAMMING

I can't help thinking that, for the most part, the author of "VIM: Taming Software with Hardware" (M. Halpern, Oct. 2003, pp. 21-25) is trying to solve a nonissue with a nonsolution.

First, I would like to see a specific memory-accessing model that allows me to access my data. This model must allow a programmer to differentiate elements within a collection and access any particular element reliably by a defined index of some sort. Then we can get into discussing the merits and drawbacks of this model.

Simply saying "virtually infinite memory" doesn't cut it because I need read/write access, not write-only access, to each identifiable subelement of my data. Arrays, therefore, are not expedients we resort to to conserve memory but a logical representation of the data plus a method to uniquely reference each element in that collection.

Second, I would like to see a suggested replacement for loops. The operations I need are:

- Perform the desired operation over a certain subset of my data—which I cannot calculate before that exact moment at runtime, and that I may not even be able to calculate during the operations. Current statements: for (i=a.begin();i!=a.end();i++) { operations; }; and while (! flag) { data operations; }.
- Repeatedly extract elements from an I/O device (file, socket, and so on) and process them until an EOT/EOF condition occurs: while (! eof(instream)) { data operations; }.

Third, programming is much different today than it was in 1980. Today, I can write a program that uses key-based array accesses to a "virtually infinite" collection. At work, I use MSVC,

so I use its built-in set, map, and list classes, but equivalent classes should either exist in most modern programming languages or be easily written from any handy book on data structures and operations.

In addition, most languages can dynamically create arrays of whatever size you need (up to available memory) using the new operator or the malloc() or an equivalent function. Programmers are no longer "forced" to allocate fixed-length arrays for buffer space; instead, they can fit the memory required to the problem at hand. Thus, the argument of array accesses being "bad" is mostly moot from this perspective—if you don't like fixed-length arrays, don't use them.

Thus,

- Throwing computer resources at the problem can ameliorate, but not eliminate, the effects of poor programming.
- We already have the tools we need to "solve" this "problem."

Michael J. Lewchuk
Edmonton, Alberta
michael.lewchuk@shaw.ca

Mark Halpern responds:

Michael Lewchuk writes that VIM is a nonsolution to a nonproblem. My first reaction is, hey, sounds like a plan! But the serious answer is that his objections are based on a misconception: VIM is not intended to change the way a programmer expresses his intentions; it is, in fact, directed to the way his intentions are carried out. It's the object program, not the source program, that VIM aims to change radically, and it's loops, not arrays, that I referred to as "simply expedients we resort to to conserve memory."

Using VIM, programmers still create arrays, and they refer to their elements by array names and subscripts—the "memory-accessing model" remains unchanged. What can differ significantly is the way VIM represents arrays internally.

As for dynamic arrays, a VIM system, like a conventional one, can allocate space only when requirements become known. The difference is that VIM can allocate *memory*—it does not need to juggle pages and interrupts and all the other nonproductive constructs of a memory-bound system. Lewchuk says that "...most languages can dynamically create arrays of whatever size you need (up to available memory)." They sure can; they can even create—conceptually—arrays that far exceed available memory. Such creation is carefree, but storing real data is another matter.

Finally, I made no claim that VIM would reduce the incidence of poor programming (although it would, at least, ensure that programmers would not botch memory management). And, while we may have the tools to resolve what Lewchuk refers to as a "nonproblem," VIM addresses a real problem—unless the burden of masses of unnecessary memory-management software is not considered a problem.

Mark Halpern
markhalpern@iname.com

We welcome your letters. Send them to computer@computer.org. Letters are subject to editing for style, clarity, and length.

Computer Lessons from Pop Songs

Bob Colwell

It has been two years since we last considered the startling and statistically significant correspondences between pop culture and our beloved information technology, in part because that's how long it has taken me to collect enough of these amazing oddities for another such column.

JOAQUIN IN L.A.

Consider the song "Words" from a 1980s album by the band Missing Persons called "Walking in L.A.":

What are words for when no one listens anymore
What are words for when no one listens
What are words for when no one listens it's no use talkin' at all

Clearly, in 1982 this band somehow anticipated the recent IEEE 802.11 and Bluetooth standards, both of which provide for a central access point that periodically broadcasts a "who's out there" beacon and then listens for responses. If there's no answer, the central access point gets somewhat depressed and lonely and polls less often. This is why it is important to always have at least one responder in any wireless network, so that the access points don't go from sadness and introversion to anger and open hostility—you know, the perennial state of Windows Millennium, with its ever-present Blue Screen of Terrible Sadness and Great Despair.

But here's the delicious part, redolent of self-referentiality. It turns out that Missing Persons were, er, enunciation-challenged. And with the same confident confusion with which many listeners thought they heard Jimi Hendrix sing "'Scuse me while I kiss this guy," the real lyrics to their song titled "Walking in L.A." were heard as "Joaquin in L.A." Wait, it gets worse. Some of the rest of the lyrics to this song went as follows:

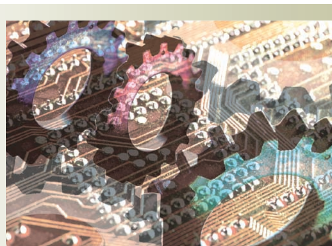

Pop songs often have computer design secrets buried deeply in their lyrics.

Look ahead as we pass, try and focus on it
I won't be fooled by a cheap cinematic trick
It must have been just a cardboard cut out of a man
Top forty cast off from a record stand

But what some people heard was this:

Whoever heard of the fab giant ultrasonic
I won't be fooled by the sheeps in a mattress
It must've been just a cardboard cannibal the man
Saw bowling Chester from the record stand

In effect, the message in this song sketches out the basic structure of today's communications standards, while also demonstrating the very point being made. It's a beautiful thing when mutually reinforcing communication occurs on multiple levels at once. The singer complains that no one is listening, but actually even those who are listening aren't getting the message. What message? The message that they aren't listening! Delightfully fractal, it seems to me.

The Rolling Stones' classic "(I Can't Get No) Satisfaction" also draws our attention to the plight of the lonely access point, but this song also points out that the number of retries equals four:

I can't get no satisfaction
I can't get no girl reaction
'Cause I try and I try and I try and I try
I can't get no, I can't get no

This excerpt could equally well describe our experience upon first power-up of Intel's Pentium Pro processor in 1995. The first three times we asserted the Reset pin, the chip wouldn't work, but the fourth time it did. However, I discount the possibility that the Rolling Stones were referring to our chip with this song because surely if they'd known the chip would do this, they would have told us the cause instead of making us figure it out.

MISINTERPRETATION

It's also intriguing how far from rationality the received message can wander before the receiver starts questioning whether the message was correctly decoded.

I think this constitutes proof that most pop music listeners don't expect to make sense out of most lyrics; therefore, when a quasi-random word sequence seems to emerge from their speakers, they cheerfully translate it into some other sequence without worrying that anything might have been lost.

Still, modern speech depends heavily on our ability to compensate with cognition for what is lost in the audio channel to noise, inattention, hearing impairments, and preconceptions—people hearing what they expect, independent of what is actually said. Doesn't that just drive you nuts sometimes?

I would expect that there is a threshold at which our cognition returns the result, "This song's lyrics are ludicrous babble." Then we act accordingly, either cheerfully singing along or looking up the lyrics on the Internet. But in any case, surely that threshold would have been set in such a way as to identify a potential problem with the received lyrics of "sheeps in a mattress"?

No wonder the Animals sang, "Oh Lord, please don't let me be misunderstood." Reminds me of a patent we wrote on a mechanism for handling page misses. The inventors called it the "page miss handler." When the first draft of the patent disclosure came back, the title of the patent said, "Page Mishandler." Spaces are important too.

In the 1960s, my parents teased me, a hapless de facto representative of my generation, for the apparently drug-induced weirdness of many of the lyrics of the time. I mounted the best defense I could, but with songs like "I Am the Walrus," the effort was doomed. Too bad nobody can understand what Bob Dylan sings, because many of his lyrics actually make sense. At least, I think they do—I could be hearing lyrics much different than what he intended.

Eventually, I fought back when I found my parents' old record on which there was a song with the lyrics, "mairsie-doats and doesie-doats and little lambsy-divie," and we called a truce. Nowadays, I tease my own children in a way they can't refute so easily. While some particularly vacuous drivel is emanating from our car's radio, I simply shake my head sadly, look at the kids, and say, "I'm sorry. We had the Beatles."

Today's songs might be telling us what the computers of 2023 will be like.

TIMING IS EVERYTHING

Alert design engineers (and what other kind are there?) will recognize what the Steve Miller hit "Fly Like an Eagle" is describing: "Time keeps on slippin', slippin', slippin', into the future...." That's right—he was trying to converge the timing on his latest mixed-signal SoC chip. Apparently he couldn't do it, so my guess is he had to relax the clock rate. Even then he probably had heat problems, which he all but admits to in another hit song, "Abracadabra," from the point of view of the fan sink on his chip:

I heat up, I can't cool down
You got me spinnin'
'Round and 'round
'Round and 'round and 'round it goes
Where it stops nobody knows

Of course, Miller could also have been singing about his pet gerbil on its little treadmill. Enigmatic ambiguity has spawned many a PhD thesis. Such is art.

Pink Floyd needs help, too. Normally, the aging effects of silicon don't have much of an impact on a clock-by-clock basis, but with its song "Time," Pink Floyd seems to be complaining about a speed path that progressively gets worse. Artsy, the way they use the periodicity of the sun to represent a single-phase clock in a synchronous system:

And you run and you run to catch up with the sun, but it's sinking
And racing around to come up behind you again
The sun is the same in a relative way, but you're older
Shorter of breath and one day closer to death

Not all rock stars have such problems with their designs. Mick Jagger brags about his success, probably just to irritate Steve Miller and Pink Floyd: "Time is on my side. Yes, it is. Time is on my side. Yes, it is." It is also possible that Jagger is referring to his botox injections, but if so, I question his conclusion, having seen some recent photos of him.

And then there's Devo, a strangely quantized quintet from the 1980s consisting of men wearing white shirts and bow ties with flowerpots on their heads, moving as if they had stepping motors instead of muscles—or maybe the video sampling rate was just far too slow. They understood pipelining, though. Every pipestage has work to be done, changes to be wrought on the data flowing through, and a deadline to meet (before Pink Floyd's sun comes up again):

Now whip it
Into shape
Shape it up
Get straight
Go forward
Move ahead
Try to detect it
It's not too late
To whip it
Whip it good

OUT-OF-ORDER POP SONGS

This one's positively eerie. Somehow, in her 1980s song, "Time After Time," Cyndi Lauper anticipated the challenges of designing an out-of-order engine like Intel's P6, where the front end that fetches, decodes, and renames is decoupled from the back end that performs all executions:

Sometimes you picture me
I'm walking too far ahead

You're calling to me, I can't hear
What you've said
Then you say go slow
I fall behind
The second hand unwinds

If you're lost, you can look and you will find me
Time after time
If you fall I will catch you I'll be waiting
Time after time

By the time the back end finds a branch that has been mispredicted, the front end has run off into the weeds. It could even be fetching and decoding data for all the back end knows. The back end then redirects the front end to the point where the error was made, and the front end restarts. But by now, the back end may be out of work to do. How did Cyndi Lauper see all this coming?

THE FUTURE

I'm an engineer. I may not understand all of the reasons why something works, but I can spot a pattern and extrapolate from it. If 1980s pop songs predicted the course of computer design 20 years ahead, then today's songs might be telling us what the computers of 2023 will be like. Let's consider some of today's lyrics to get a sneak preview.

The 12 Stones' song "Crash" predicts a continuation of the bane of today's designs: memory latency and snoopy bus protocols. I would have hoped that in two decades we would have gotten away from these problems:

My mind is slowly fading
So far away from me
Each time I start crawling
You're there watching me

With high probability (pun intended), Dave Matthews' "You Never Know" predicts quantum computing plus optical interconnect:

Out of the darkness comes light
Like a flash
You think you can you think you can
Sometimes that is the problem
Dream
Is it your dream
Dream

After extensive research, I can only conclude that the rest of today's pop songs have buried their computer design secrets too deeply in lyrics that consist mostly of profanity or inanity. Perhaps the pattern formed by the few normal words completely surrounded by the curse words is a coded message. I'll have to work on that some more. Meanwhile, all I can say is ... we had the Beatles. And the Walrus was Paul. ■

Bob Colwell *was Intel's chief IA32 architect through the Pentium II, III, and 4 microprocessors. He is now an independent consultant. Contact him at bob.colwell@comcast.net.*

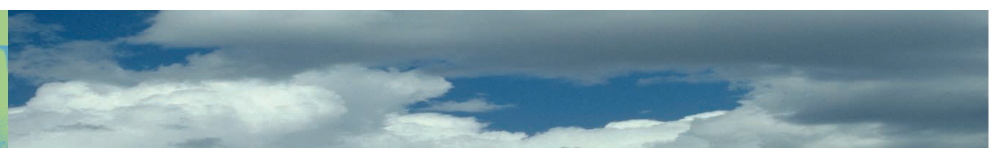

New Trends Revive Supercomputing Industry

Steven J. Vaughan-Nichols

Although grid computing—which links disparate machines so that they can function as a distributed supercomputer—has become an increasingly popular focus of high-performance-computing research, the traditional supercomputing industry has languished until recently. This occurred largely because the easing of the Cold War in the early 1990s reduced government use of and spending on supercomputer technology.

However, the industry is now reviving because of the development of low-cost supercomputer clusters that use commodity chips. Clusters of commodity computers linked by high-speed interconnect technologies, such as InfiniBand, have put supercomputers within reach of new users.

Meanwhile, longtime supercomputer vendors like Cray are making a comeback with systems based on *vector units*, powerful CPUs dedicated to floating-point and matrix calculations. Vector machines satisfy the demand by governments and large industries for supercomputers that can conduct complex tasks such as nuclear-weapons simulations, pharmaceutical drug modeling, mining of large data sets, and geological analysis to find oil deposits.

"Things are changing because of advancements in microprocessing and because we now know how to better use parallel processing," noted Jack Dongarra, professor and director of the Innovative Computing Laboratory at the University of Tennessee.

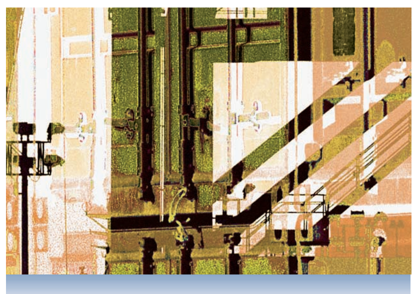

In addition, competition among the US, UK, German, Japanese, and other governments to sponsor the fastest supercomputers for scientific and military research has sparked development of increasingly powerful machines.

Attesting to supercomputing's new popularity, SC2003, last year's IEEE and ACM annual high-performance-computing conference and exhibition in Phoenix, Arizona, had the event's highest attendance ever with 7,641. Officials also had to turn away 50 companies because the exhibition hall ran out of space.

TRENDS

The Top500 (www.top500.org)—a list of the world's 500 most powerful supercomputers compiled every June and November by researchers at the University of Mannheim, the University of Tennessee, and Lawrence Berkeley National Laboratory—indicates some recent supercomputing trends.

For example, said Dongarra, "Supercomputing performance is doubling about every 14 months." The total performance of all systems on the Top500 has increased by a factor of 10 every four years. In 1993, the list's first year, all systems' total performance was about 1 Tflops. The total grew to 10 Tflops in 1997 and 100 Tflops in 2001, and is expected to reach 1 petaflops (1,000 Tflops) in 2005.

In the November 2003 Top500 list, there were 208 computer clusters, up from 149 in the previous listing. Clusters didn't even appear in the Top500 until 1998.

Of the top 10 supercomputers on the November 2003 list, two use processors that have never made the list before.

IBM's 2-GHz PowerPC 970 G5 processor, found in high-end Apple computers, was used in Virginia Polytechnic Institute and State University's third-ranked 10.28-Tflops Terascale Cluster. Srinidhi Varadarajan, director of Virginia Tech's Terascale Computing Facility, explained, "There is a lot of interest in [the chip] because everyone in this domain of scientific computing realizes that here you have a processor that can really do [floating-point calculations] … better than anybody out there."

Meanwhile, cluster-system vendor Linux Networx used Advanced Micro Devices' 64-bit, 2-GHz Opteron in a 2,816-processor, 8-Tflops, Linux-based cluster built for the Los Alamos National Laboratory. Opteron is fast, but it's the AMD-pioneered HyperTransport interchip communications technology (www.hypertransport.org/technology.html) that really attracts supercomputer designers.

HyperTransport's aggregate data throughput of 12.8 Gbytes per second is 40 times faster than legacy PCI buses and even more than the 12.3-Gbps internode data rate produced by proprietary technology in the Top500's fastest computer: Japan's 35.6-Tflops Earth Simulator, built by NEC for geosciences research. HyperTransport uses high-integrity signaling and high clock speeds to provide fast data rates.

Meanwhile, vendors are beginning to change the face of high-performance computing by building supercomputers that consume less power, as well as lower-cost supercomputers for general corporate use, as discussed in the "IBM's Energy-Efficient Supercomputer" and "Cray's Supercomputer for the Corporate Masses" sidebars.

VECTOR-SUPERCOMPUTING REVIVAL

As post-Cold-War US governments began spending less on supercomputer research, the then-new massively-parallel-processor model began to replace the previously dominant but very expensive vector model, made famous by the late Seymour Cray and his Cray Research, now called Cray Inc.

MPP computing uses many CPUs, each with its own memory, running in parallel and linked by high-speed interconnections to execute the various parts of a program. In traditional MPP computers, the CPUs are in the same box and are linked by buses. In cluster machines, the CPUs are in different boxes, linked by network interconnects.

Over time, MPP approaches showed that because the processors in the system weren't designed to work closely together, they were effective primarily for work that could be divided into subtasks that don't need to be tightly integrated. This would include work in which the results of one subtask don't affect the results of another, even though all the subtasks would be eventually aggregated for the purpose of yielding final results.

"In a vector computer, instructions, instead of performing a single operation at a time, act on arrays of elements [simultaneously]," explained Cray chief architect Steve Scott.

Vector machines use specialized CPUs and memory units to offer both high processing speeds and the ability to transfer data rapidly within the computer while calculations are taking place. The CPUs are designed to work with arrays of elements, while the memory units are fast, coupled tightly enough to work efficiently with one another, and accessible by multiple CPUs at the same time.

Vector supercomputers are particularly effective for tasks, such as the type of complex system modeling used in meteorology, that require tightly coupled processors and global shared memory.

By the late 1990s, many experts assumed MPP designs were supercomputing's future. Even Cray Inc. has built some MPP machines.

However, the Earth Simulator showed that vector computing still has life, particularly when combined with MPP techniques. In the Earth Simulator,

IBM's Energy-Efficient Supercomputer

IBM has introduced a small, high-performance prototype machine, Blue Gene/L, that uses relatively little energy, thereby saving money and mitigating the heating problems that plague typical power-hungry supercomputers. Many supercomputers run fast processors close together and thus generate so much heat that they can operate for only limited periods of time.

The dishwasher-sized Blue Gene/L prototype represents a departure from the ongoing development of increasingly fast supercomputers based on ever-speedier microprocessors. Instead, IBM has balanced performance and energy consumption by running large numbers of lower-powered processors.

Blue Gene/L performed at 1.4 Tflops, placing it 73rd in the Top500 list of the world's fastest supercomputers. It uses 512 700-MHz PowerPC 440 CPUs, which are deliberately run at only 500 MHz to minimize heat production. All of the CPUs have two PowerPC 440 cores, each with two floating-point units.

Because Blue Gene/L produces less heat, it can be air-cooled. Many supercomputers have to use water and refrigeration for cooling, which is expensive.

Blue Gene/L is a prototype for a system, ASCI Purple (ASCI is the US government's Accelerated Strategic Computing Initiative, now called the Advanced Simulation and Computing Program), that will have 128 times as many processors. IBM plans to complete the system—which will be used to model the folding of human proteins for medical research—for Lawrence Livermore National Laboratory in 2005.

Cray's Supercomputer for the Corporate Masses

Typically, supercomputers have been too expensive and complex for all but the largest corporations. However, by early 2005, Cray Inc. hopes to release a line of smaller, easier-to-use commercial products based on the Red Storm supercomputer that it's building for the US Department of Energy's Sandia National Laboratories. If Red Storm achieves 40 Tflops as planned, it would be the world's fastest computer.

The system will be fast in part because it will use HyperTransport technology, explained Bill Camp, head of Sandia's Red Storm project.

Red Storm's commercial spin-offs will also use HyperTransport and will work with Cray's MPP architecture, the Linux OS, and between 96 and 1,000 AMD 3.04-GHz Opteron processors, depending on customer needs. The processors can handle 32- and 64-bit software, making it easy to use existing 32-bit software and providing an upgrade path to 64-bit applications.

vector processors were used in an MPP architecture.

Demonstrating the renewed interest in vector supercomputing, the US Defense Advanced Research Projects Agency selected Cray, IBM, and Sun Microsystems for its High-Productivity Computing Systems program that aims to develop, by 2010, prototypes of vector/MPP supercomputers that can reach 1 petaflops.

CLUSTERS

According to Robert Pennington, interim director of the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign, "The biggest impact on supercomputing has been [made by] the advent of clustering. Vendors like Dell, IBM, and HP are now supplying cluster packages, so supercomputing has moved out of the realm that if you can't build it yourself, you can't have it. Now you can have a company supply and support it. Scientists can concentrate on science and not on computer science."

Clusters are attractive to many users because they are relatively inexpensive and can use commodity parts readily available from many suppliers. As with other types of MPP systems, clusters are effective for loosely coupled tasks.

Eric Pitcher, vice president of product marketing for Linux Networx, said, "The technology driver has been the increased power of the microprocessor and powerful processor-interconnect networks like InfiniBand and [like proprietary approaches from] Miracom or Quadrics."

InfiniBand, for example, specifies a serial bus, rather than the traditional parallel bus. The technology is faster because it carries multiple channels of data simultaneously in a multiplexing signal and supports multiple memory areas.

Virginia Tech's low-cost supercomputer

According to Jason Lockhart, associate director of Virginia Tech's Terascale Computing Facility, the school's Terascale Cluster was built in less than four months for only $5.2 million. Japan's Earth Simulator, on the other hand, cost an estimated $375 million and required about five years to build.

The Terascale Cluster was quick and relatively inexpensive to build because it uses commodity components extensively. For example, in addition to 1,100 off-the-shelf Apple computers, the cluster uses Cisco Systems' 4500 Gigabit switches for the management network and 24 Mellanox Technologies' 96-Port InfiniBand switches—each with 1.92 Tbits of bandwidth—to carry data between the component computers.

Explained Lockhart, "We were looking to produce a world-class supercomputer at a research-university price tag." The machine will be used for research on nanoelectronics, aerodynamics, molecular modeling, and computational acoustics.

Intel's Advanced Computing Program

Intel plans to spend $36 million on basic research to improve the performance of cluster-based supercomputers made from off-the-shelf parts, such as the company's own Xeon, Itanium 2, and Pentium 4 chips. The Advanced Computing Program is part of Intel's ongoing effort to gain a bigger presence in high-performance computing.

The three-year ACP—which eventually is slated to include several server manufacturers and universities—will work on supercomputer-friendly architectures and advanced system software, as well as software tools, languages, and libraries, explained Justin Rattner, Intel senior fellow and director of the Intel Systems Technology Lab.

Intel-based cluster systems already account for four of the 10 highest-ranking Top500 supercomputers and 189 of all machines on the list. A year ago, only 56 of the Top500 computers used Intel chips.

Using Linux in clusters

Early clusters typically used various versions of Unix. Now, though, Linux is becoming increasingly popular for several reasons. Cray's Scott said, "Linux is open, and people perceive that it provides high quality at a low price."

Linux Networx's Pitcher added, "You can treat a Linux cluster as a totally integrated system [rather than a collection of individual computers], thus making it easy to schedule software, monitor the cluster, and centrally operate it."

Moreover, the NCSA's Pennington said, with an open source operating system, you can adapt the OS as needed and thereby work on a common platform across different architectures and easily port applications from one system to another.

The Tungsten 9.8-Tflops, Linux-based cluster, which the NCSA makes available to scientific researchers, is currently the fourth-fastest supercomputer in the world. The system uses Red Hat Linux and includes 1,250 Dell PowerEdge 1750 servers, each running a pair of Intel 3.06-GHz Xeon processors, and the Myrinet 2000 low-latency, 4-Gbps interconnect fabric.

Building the optical clustering infrastructure

The US Department of Energy has awarded IBM and Corning $20 million to develop high-speed, optical connection hardware for supercomputers during the next 2.5 years. The faster optical equipment would replace today's electronic, copper-based supercomputing switches, such as those based on InfiniBand.

Craig Stunkel, IBM's manager of scalable server network and memory systems, said the goal is to improve interconnection technology's cost, latency, and throughput. The optical equipment would be more expensive than the electronic hardware. However, Stunkel noted, the optical approach would run cooler and eliminate the expensive cooling systems necessary with the electronic approach, saving money in the long run.

IBM and Corning hope to build a prototype optical system by the end of 2005. If the technology can be commercialized, Stunkel said, it would provide node-to-node interconnects.

The future direction of supercomputing may change after the completion of two US studies. The National Coordination Office for Information Technology Research and Development is developing a five-year roadmap for federal supercomputing investments. The Jasons, a group of scientific advisers who consult with the US government on classified research issues, are identifying how current supercomputing projects meet the needs of the government's nuclear-weapons program.

However, Pitcher said, he isn't sure that the government, which used to drive supercomputer sales, will continue to be that important. "Expensive proprietary designs have their place, and those will require a lot of upfront investment by the government to keep them in place for the small number of applications that require them," he said. "But with clustering that uses off-the-shelf technology making supercomputers less expensive, supercomputers will move into broader commercial accessibility for both large and mid-sized businesses." ■

Steven J. Vaughan-Nichols is a freelance technology writer based in Arden, North Carolina. Contact him at sjvn@vna1.com.

Editor: Lee Garber, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; l.garber@computer.org

Will New Standards Help Curb Spam?

David Geer

The percentage of all e-mail that is unsolicited has risen steadily during the past couple of years, as Figure 1 shows. This barrage of spam has caused numerous complaints from many sources. In fact, a recent survey by the Pew Internet and American Life Project—which funds research on the Internet's impact on individuals, families, and communities—showed that 75 percent of US users are concerned about spam, said Pew research fellow Deborah Fallows.

Some proponents say spam is a form of legitimate marketing. Others—such as Scott Richter, the president of OptinRealBig, an e-mail marketing firm—say revenue-generating spam is also the price users pay for having access to so much free material on the Internet.

However, European companies lose $9 billion a year because of spam, primarily from lost productivity due to employees reading or deleting unsolicited messages, trying to unsubscribe to mailings, or reporting spam-related problems to IT departments, said John Levine, cochair of the Internet Research Task Force's Anti-Spam Research Group (ASRG).

In light of these concerns, Ferris Research analyst Chris Williams noted, "In 2003, more antispam products came to market in the US than [any] other [type of] e-mail application or management tool."

He said US antispam product sales were about $130 million in 2003 and will increase to about $360 million this year. This diverts money that corporate

customers could spend on profit-making activities.

A principal antispam approach is filtering. In some cases, e-mail servers, clients, or both block messages recognized as potential spam based on headers, text, images, and code. Some filters compare messages against whitelists of approved sender addresses. Others block mail from addresses on blacklists of known spammers. However, filters aren't always easy to configure and can block legitimate e-mail.

Some governments are taking a legislative approach to fighting spam. For example, US officials recently adopted the Can-Spam (Controlling the Assault of Nonsolicited Pornography and Marketing) Act, and the European Union approved a directive making it illegal to send unsolicited e-mail without a preexisting business relationship. However, these laws are likely to have little effect on spammers who are from unregulated areas of the world or who can send their transmissions anonymously.

Thus, many industry observers say neither filtering nor regulation will effectively stop spam. They contend that more fundamental changes to e-mail technology—such as alterations to basic standards—are necessary.

With this in mind, various organizations, including the ASRG and the ePrivacy Group, are working on e-mail standards changes. These efforts are still in the early development stages and face a number of challenges, but they have already begun to garner support from various sources.

PROPOSED STANDARDS CHANGES

A key issue with spam is the ability of senders to mask both their identities and the servers from which they send their unsolicited mail. "The vast majority of return addresses on spam are generated randomly by spamming software, so they're usually bogus," explained Ray Everett-Church, the ePrivacy Group's chief privacy officer. This makes it difficult to identify, track, and stop spammers.

This, in turn, complicates efforts to block mail from known spammers and their mail servers. Thus, authenticating mail senders and their servers, thereby ensuring that e-mail messages actually come from their listed source, is a key approach to using standards and technology to curb spam.

Programs typically use the Simple Mail Transfer Protocol, which was adopted in 1982, to send e-mail. They use either Post Office Protocol 3 (POP3) or the Internet Message Access Protocol to receive messages sent to a local server.

SMTP was designed to be open, which helped make e-mail popular in the first place, explained Paul Hoffman, director of the Internet Mail Consortium, a clearinghouse that hosts mailing lists relating to the adoption of Internet mail standards.

However, as part of its being open, SMTP doesn't offer authorization. This lets spammers easily mask their identities by hacking unprotected mail servers or forging return addresses in a message's "mail from" command. "Most mail systems won't challenge it," explained Ken Schneider, chief technology officer of antispam vendor Brightmail.

This has led to calls for the protocol's revision or replacement.

The need for authentication

A key challenge in revising or replacing SMTP is authentication. With authentication, spam recipients have a way to identify senders, and senders realize they can be identified.

However, Hoffman said, experts still haven't figured out the best way to use authentication with e-mail. With small groups of users, he said, either direct or third-party authentication could work. On the other hand, he said, with large groups of e-mail users, neither approach will scale effectively. There are too many individuals and e-mail servers involved and too many issues in determining e-mail trustworthiness, he said.

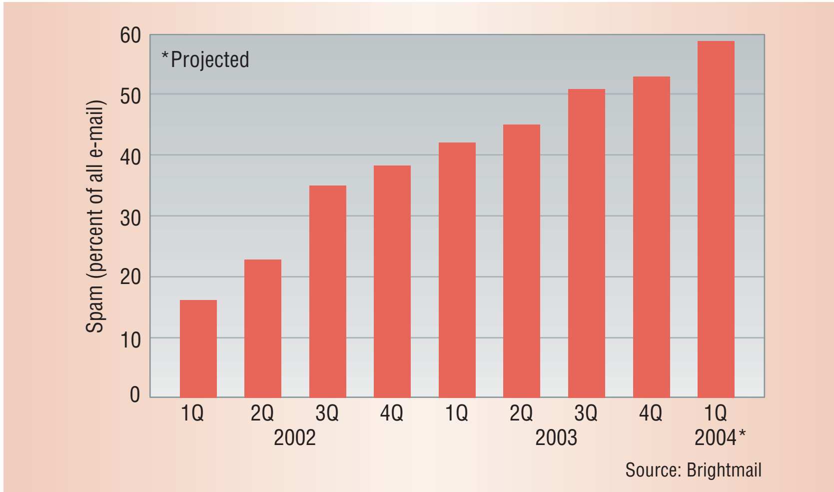

Figure 1. The percentage of US e-mail that is spam jumped from about 15 percent during the first quarter of 2002 to more than 50 percent during the last quarter of 2003, and is predicted to keep rising.

Anti-Spam Research Group

The Anti-Spam Research Group (http://asrg.sp.am/) is studying and considering adoption of three protocols that would curb spam by helping Internet service providers (ISPs) or organizations with their own e-mail systems determine whether message senders have spoofed their e-mail addresses.

The protocols would change the domain name system so that domains could publish all of the IP addresses legitimately associated with them—a type of reverse mail exchange (RMX) record—in their own DNS databases. If an e-mail sender transmits a message and spoofs his address to make it look like it came from a specific domain (such as computer.org), e-mail servers receiving the transmission could check the sender's real IP address, listed in the message's header, against addresses listed in the spoofed domain's RMX record. The server could then verify whether the message actually came from the domain.

None of the protocols do more than determine whether an e-mail address is forged. An ISP or other e-mail provider would have to decide whether to block e-mails with forged addresses or pass them on to intended recipients with warnings.

The ASRG initially intended to merge the three protocols but decided that would not yield the best result, Levine said. Instead, he explained, "We are trying to get a good idea of their relative technical merits to see if there is one that we want to pick."

Sender Permitted From. Meng Weng Wong, founder and chief technology officer of e-mail forwarding service Pobox, said he is finalizing development of SPF for submission to the Internet Engineering Task Force, via the ASRG, as a potential protocol.

According to Wong, SPF systems would check a domain's DNS to determine if an e-mail message has been spoofed before the bulk of the message is transmitted to its intended recipients, rather than waiting until the transmission is completed. This would save time and bandwidth.

Wong said he expects to finish a draft of SPF and begin the IETF approval process this month.

Designated Mailers Protocol. "[With DMP], you can still send spam in the name of your own domain, but then it can be tracked," said Gordon Fycek, DMP's developer and owner of Pan-Am Internet Services, a PC and Internet consultancy.

DMP is compatible with IPv4 and IPv6 and also has its own way of listing the findings of its DNS checks. Only Pan-Am Internet is using these DMP records now, though.

Fycek said he has been discussing DMP with the ASRG and is also currently working with the Internet Society. In addition, Pan-Am has already made the current version of its protocol available for use by organizations and companies, although the specifications might change during the standardization process.

Reverse Mail Exchange. RMX works much like SPF and DMP but would change the e-mail message headers themselves to include the results of DNS database checks. Recipients' e-mail clients could then use the information to appropriately handle the messages, based on user-specified rules.

"I am close to the end of the development process," said Hadmut Danisch, a security consultant with the SHD technology consultancy. Danisch has worked with the ASRG but said he is unsure exactly how he will proceed in trying to standardize RMX.

Trusted E-mail Open Standard

The ePrivacy Group, a privacy consultancy and product vendor, is developing TEOS to correct SMTP's lack of authentication. The proposed standard would also provide content descriptions, in which TEOS-compliant senders would list the type of message they are transmitting, to let recipients handle the e-mail as they wish.

TEOS works with existing third-party digital-certificate technologies, such as X.509, to authenticate the identity of e-mail senders, explained Everett-Church. This requires e-mail senders to use a certificate, an attachment to a message in which a trusted third party verifies the sender's identity. TEOS technology would also encrypt the certificate and associated information to prevent alteration by intruders.

The ePrivacy Group plans to submit a formal TEOS proposal to the ASRG in the near future.

In the long run, TEOS will need support from some of the major ISPs if it is ever to become successful, said Anne Mitchell, CEO of the Institute for Spam and Internet Public Policy.

CAUTIONS AND CONCERNS

The proposed new standards have raised several concerns. For example, some experts say that widely implementing SMTP changes will be difficult because of the protocol's huge global user base.

The Internet's size threatens the proposed antispam standard's effectiveness.

The Internet's size and openness threaten the effectiveness of the proposed antispam standards, contended the Internet Mail Consortium's Hoffman. There are so many e-mail servers, he explained, it may not be practical to authenticate each possible user combination. Even if the proposed protocols could provide adequate authentication, he said, they won't necessarily provide trust for authenticated senders with whom recipients are not familiar.

Because of this, said the ePrivacy Group's Everett-Church, the proposed ASRG standards won't address all of the issues that have let spam flourish.

"Even assuming a near-future agreement on and promulgation of new standards, getting widespread buy-in and adoption, let alone implementation, will take at a minimum [many] months, and that is being optimistic. It is no small project for any major ISP to revamp its e-mail sending and receiving infrastructure," said Mitchell.

"It's probably not pessimistic to think this all could take five to 10 years to unfold, if it unfolds at all," said Andrew Lockhart, director of product marketing for Postini, an e-mail security vendor.

Meanwhile, said Everett-Church, "Focusing on IP-address-based solutions opens up performance problems. IP-address-based solutions introduce the need for a back channel for verification of every e-mail message. This is the equivalent of doing an authorization phone call every time somebody uses a credit card."

OptinRealBig's Richter said the antispam protocols are unnecessary because the spammers he knows are engaging in a legitimate business activity and not looking to create problems or generate complaints. And when online, he added, users should be more careful in reading agreements and authorizations that permit e-mail solicitations.

Pobox's Wong was optimistic about the proposed protocols' future. "When enough organizations adopt protocols like SPF," he said, "spammers will have to use their real e-mail addresses, at which point technology could block them."

As for which protocol will become the most widely used, Brightmail's Schneider said that if a few large mail systems begin deploying one approach, others may follow and thereby create a de facto standard.

However, the ASRG's Levine said, if the standards community doesn't provide antispam protocols soon, big companies, acting individually or in concert, will make unilateral changes to how they handle mail, perhaps creating de facto standards in the process.

And in the long run, Mitchell said, using standards might curb some spam, but effective e-mail filtering and antispam laws will also be necessary. ■

***David Geer** is a freelance writer based in Ashtabula, Ohio. Contact him at geercom@alltel.net.*

Editor: Lee Garber, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; l.garber@computer.org

Researcher Develops Colorful Map of the Internet

A US-based security consultant is creating a graphical representation of Internet network routes, yielding striking, colorful images that depict those connections.

Barrett Lyon of Network Presence, an information-security consultancy, started the Opte project (www.opte.org/) to map the Internet. "My goal was actually to see if I could do it quickly and with less hardware than other people might need, to prove it is not this huge, difficult process," Lyon said. Similar efforts, such as the Lumeta Corp.'s Internet Mapping Project (http://research.lumeta.com/ches/map/), have operated on a larger scale.

In addition to appealing to general interest, Lyon said, Opte offers benefits such as modeling the Internet and analyzing wasted IP space.

For example, he explained, the project could show what happens to the Internet after a natural disaster or important geopolitical event. He said Opte could also show which major companies or organizations have huge blocks of IP addresses that are unused and that could be given to other entities that need them.

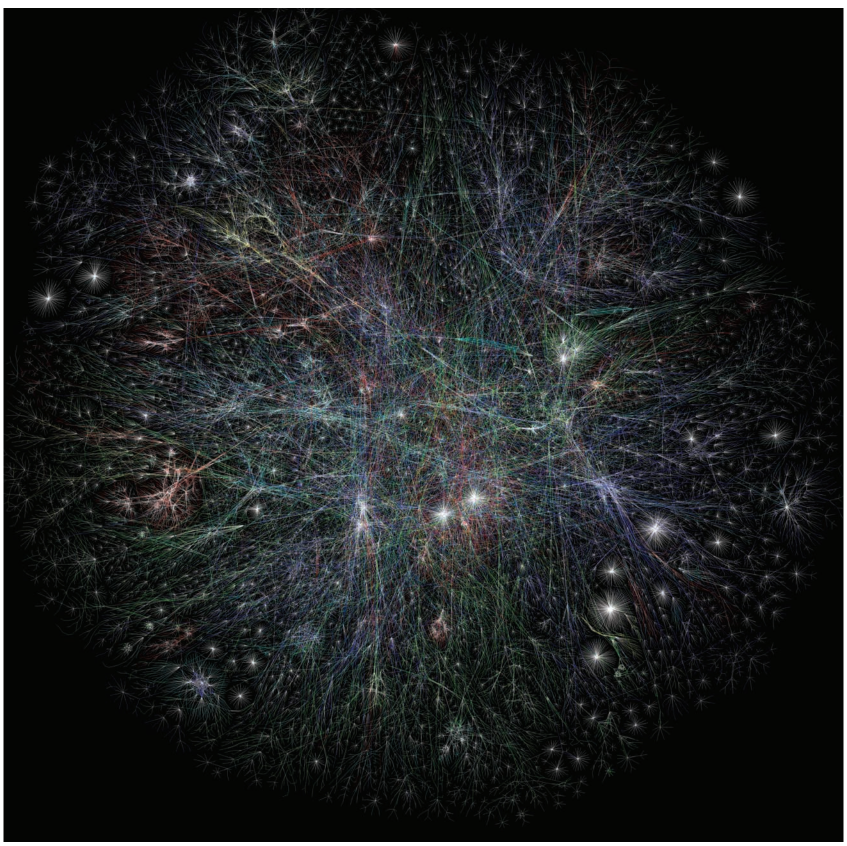

A US security consultant has begun a project to map Internet network routes. In this map, routes in red are based in the Asia-Pacific region; routes in green are based in Europe, the Middle East, Central Asia, or Africa; routes in blue are from North America; routes in yellow are from Latin America and the Caribbean; and routes whose origins are unknown are in white. Source: The Opte project.

Opte tracks Internet routes by using a traceroute program, which records the route a packet takes between two endpoints, indicated by the gateway computer at each hop. The scanrand2 utility performs traceroute searches involving thousands of packets per second.

A single computer—using a 1.7-GHz Intel Celeron processor, with 1 Gbyte of RAM, working with the Free BSD operating system and other open source applications—runs the Opte project's database, Web pages, image calculations, route collection, and distribution of scanrand2 requests.

Because mapping the Internet from a single computer could miss many Internet subnetworks with which the node may not connect during multiple transmissions, Opte collects information from four computers scattered throughout the Internet.

Opte maps only what used to be called the Internet's Class C networks. The Internet's original routing scheme in the 1970s defined three primary classes of IP address to be assigned to users: classes A (for very large networks), B (for mid-sized networks), and C (for smaller networks). Class D was for multicasts, and Class E was for experimental purposes.

Class C represents most Internet networks. Because of this, Lyon said, his project provides a good representation of the Internet. ■

—Linda Dailey Paulson

Using Plastic to Make High-Capacity Memory

Researchers are experimenting with a new type of compact, inexpensive, high-capacity electronic memory for small devices that uses a plastic material also sometimes employed as an antistatic coating for computer screens.

A team of Princeton University and Hewlett-Packard scientists has used polyethylenedioxythiophene, called Pedot, to develop the memory system for deployment in cameras, MP3 music players, cellular phones, and other small devices.

The new system is a type of WORM technology. "Write-once, read-many memory is very useful for large archival files such as digital photographs, music, etc.," commented Robert Street, a senior research fellow at the Palo Alto Research Center.

The new memory system consists of an amorphous silicon diode and Pedot printed on an inexpensive plastic substrate at each intersecting point on a grid of electrodes. The Pedot, which forms the device layer that stores binary data, is applied using the recently developed electronic imprint lithography process, which is less expensive than the traditional lithography used with silicon products, noted HP Labs researcher Craig Perlov.

The Pedot and diode are located between two electrical wires. To write data, researchers apply varying voltages across the grid. At low voltages, Pedot conducts electricity. At about 210 degrees centigrade, the material permanently loses conductivity, which is why the new storage device is not rewritable. The presence or lack of conductivity creates binary data's ones and zeros. To read data, current is run through one wire and the conductivity at the systems' various intersecting points is measured in the other.

Perlov said researchers want the new storage products to be the same size as current small storage devices so that they can work in slots already used for those devices.

Warren Jackson, a researcher at HP Labs, said the system's simplicity should make it inexpensive to manufacture. Further, the new storage devices would be produced within a sealed machine that would eliminate the need for the costly clean rooms used to make silicon-based products. Also, the new devices could be stacked, thereby increasing storage density.

When commercialized, Perlov noted, the new system will offer several gigabytes of storage, more than current flash-memory systems.

Researchers have only recently begun working with plastics for storage devices. The key has been making plastic materials conductive, explained Jackson. Plastics are attractive for storage devices because they are potentially simpler and less expensive to work with than silicon.

Over the years, Street noted, experimental plastic-based memory devices have failed to replace silicon-based systems, in large part because silicon and magnetic memory have developed so rapidly.

Nonetheless, he said, the new Pedot-based memory device has commercial potential, although he wondered whether it will work well and inexpensively enough on a large scale to compete with silicon-based systems.

And, Jackson said, it remains to be seen whether potential users will accept such a new type of memory system. ■

—Linda Dailey Paulson

Businesses Adopt Web Site's Satirical Company Names

It was supposed to be a joke. The Design Conspiracy—a UK graphic design, advertising, and digital media company—started the Whatbrandareyou.com Web site as a promotional effort that featured made-up company names as a spoof of the trend in which businesses adopt catchy-sounding but meaningless names. However, several companies have actually adopted the names.

"Someone alerted us to the fact that one name had been registered, and we then did a check on the Companies House Web site," said Ben Terrett, the Design Conspiracy's business and strategy director. Companies House is the UK government agency responsible for registering and recording the incorporation and dissolution of companies.

"Out of the 150 names we thought up, 20 have been registered," Terrett said. None of the companies, even those that may have gotten the idea through Whatbrandareyou.com, ever contacted The Design Conspiracy, he noted.

On the satirical Web site, visitors enter their name, their company's "core values" (such as dynamic, passionate, or innovative) and "main goal" (such as global leadership, client focus, or quality), and receive a recommended catchy name from a preprogrammed list of 150, including Aliquis, Emineo, Festinatio, and Sollicito.

"We thought them up one afternoon," said Terrett. "Some are vaguely Latin sounding, some are just silly. The site has had over 1.5 million visitors and still gets over 10,000 a week."

Computer magazine contacted Econoprint, a London firm that registered at least two of the names—Amplifico and Ovisovis—for comment, but the company did not respond. ■

—Linda Dailey Paulson

Researchers Develop Self-Assembling Chip Technique

IBM has developed prototype silicon-polymer flash-memory chips using a process based on self-assembling nanocrystals. The process uses a technique in which two polymers, when heated, self-organize into patterns that researchers can manipulate to make one of the stencils used to produce the chips.

This self-assembly process could enable manufacturers to produce smaller chips and reduce the high cost and complexity of traditional lithography, which uses expensive tools to precisely draw on a silicon substrate the places where every transistor, wire, and other element will go.

Kathryn Guarini, a lead researcher on IBM Research's project, said the self-assembly technique could be used to define particularly small elements in flash-memory devices, while conventional lithography would define the larger elements.

In the self-assembly technique, the polymers, also sometimes used in adhesives, are made with two types of molecules that are bonded together but that otherwise would repel each other.

When heated, the polymers crystallize and arrange themselves into patterns that can form a stencil of tiny memory cells 20 nanometers wide and 40 nanometers apart. Researchers can align the cells within the larger features created by conventional lithography. Silicon nanocrystals subsequently form the memory cells themselves.

The smaller memory cells are necessary to create the ever-smaller flash-memory chips that industry wants to fit within smaller devices, Guarini said. Size reduction is a particular challenge for flash-memory devices, which retain data when the host machine's power is shut down. As the memory devices get smaller, the program oxide becomes thinner, making it harder to avoid leakage of the electrical charge that represents the stored data.

With the self-assembly technique, the flash memory is made up of tiny nanocrystals, so any leakages generally occur in only one of the crystals and don't cause a large-scale electrical-charge loss.

In building flash-memory chips using self-assembly techniques, IBM worked with processes compatible with existing manufacturing tools. "The tools are exactly the same as the ones used to apply the polymers used in conventional lithography," said Chuck Black, IBM researcher and another lead researcher on the project. "All we have changed is the type of polymer used and the conditions used to apply and bake."

IBM says its research suggests that polymer-based self-assembly approaches could help build chips other than those used for flash memory. "Different types of polymers could be used to make different patterns [for different types of chips]," Guarini explained.

IBM plans to further explore self-assembly approaches—including which polymers are best for creating layouts for different types of chips, and how best to integrate the technique within today's fabrication processes—during the next five years. There are no current plans to commercialize the technology, according to Guarini. ■

—Linda Dailey Paulson

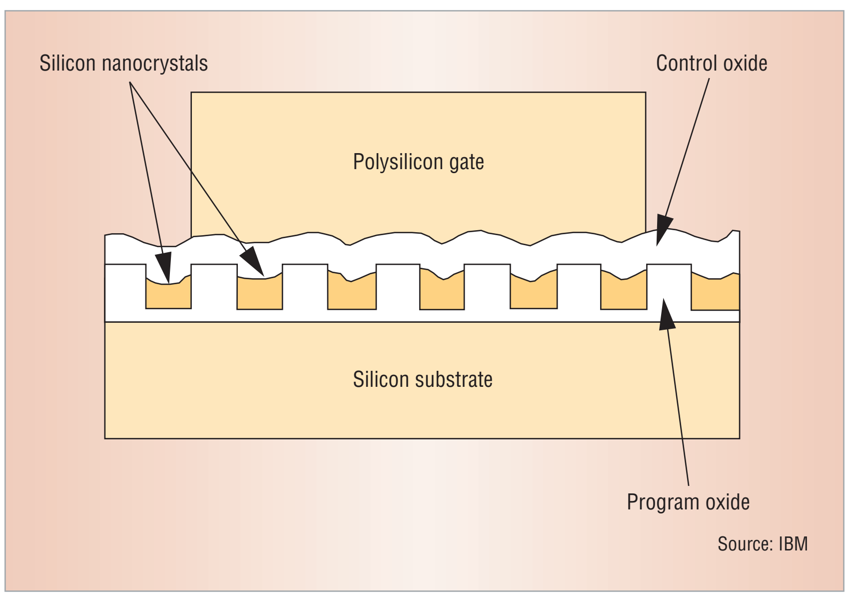

IBM has developed a self-assembly technique for making stencils used in manufacturing flash-memory chips. The technique uses two polymers that, when heated, self-organize into patterns that researchers can manipulate inexpensively into stencils for tiny memory cells. The process subsequently uses silicon nanocrystals to form the memory cells themselves.

Editor: Lee Garber, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; l.garber@computer.org

Call for Participation

VIS 2004

austin, texas

October 10-15, 2004

IEEE VISUALIZATION 2004

Vis 2004 is the premier forum for visualization advances in science and engineering for academia, government, and industry. This event brings together researchers and practitioners with a shared interest in techniques, tools, and technology. We invite you to participate by submitting your original research and joining us in beautiful Austin, Texas, nestled among the rolling hills and lakes of the hill country.

Co-located with Vis 2004 are two highly successful Symposia:

- **InfoVis 2004: IEEE Symposium on Information Visualization**
- **VolVis 2004: IEEE Symposium on Volume Visualization**

The combined Vis 2004 conference and symposia make this week in Austin the place to be to participate in this rapidly expanding field.

Submission deadline for conference and symposia: March 31, 2004

For more information, visit the Vis 2004 web site:

vis.computer.org/vis2004

Conference Chairs:
Kelly Gaither, *The University of Texas at Austin*
Jim Thomas, *Pacific Northwest National Laboratory*

Sponsored by the IEEE Computer Society Technical Committee on Visualization & Graphics.

IEEE

IEEE Computer Society

A Copper Bullet for Software Quality Improvement

Most software engineers agree that software quality improvement comes not from a silver bullet but from a combination of strategies. One of these copper bullets is to reverse-engineer a vendor's database as part of an evaluation to determine overall quality before buying the product.

Michael Blaha
Modelsoft Consulting Corp.

As Fred Brooks described in his famous "no silver bullet" paper, there is no single action the computing community can take to radically improve software quality. There are, however, "copper bullets"—lesser steps that improve quality over time. One such copper bullet is the notion of software engineering, the practice of thinking carefully before immersing yourself in the minutia of coding. Judiciously applied, software engineering should improve quality.

In theory, that makes sense, but for better or worse, software is becoming larger and more complex, which makes the benefits of software engineering less noticeable. As the software community faces an unprecedented number of project failures, researchers must continue the hunt for new copper bullets to offset the complexity.

A quality improvement strategy that the community has largely ignored is to use database reverse engineering to measure the quality of software that a company is looking to buy. Companies now routinely assess vendor software on the basis of cost, functionality, user interface, and vendor stability, but none of these dimensions addresses the software's intrinsic quality. Database quality, on the other hand, could be a litmus test for overall quality. If a product has a flawed database, it is likely to have other quality issues, such as messy programming. In contrast, the quality evident in a sound database is likely to be present in the software's other parts.

Over the past 11 years, my colleagues and I have been evaluating software using database quality as the basis for product grading. We have found that reverse-engineering a database can help a company deeply understand the associated product. Moreover, the time to do the evaluation (sometimes only a few person-weeks) is trivial compared to the millions it can cost to buy and deploy the application.

The benefits of this copper bullet are enormous. As reverse engineering pressures vendors to improve their offerings, vendors upgrade their software development practices to survive the scrutiny. Success then depends less on marketing prowess and more on technical merit.

Proficient vendors receive more notice, can negotiate more attractive prices, and can look forward to increased sales. Inferior vendors receive less revenue and are eventually forced from the marketplace as companies flock to their more proficient competitors. Openness is encouraged, since reverse engineering makes it plain what models and database designs vendors actually offer, regardless of what the vendor is willing to publish.

Thus, as more companies practice database reverse engineering, aggregate vendor quality should improve, benefiting the entire software community.

The suggestions offered here extend only to database assessment, not to the more general problem of reverse engineering to assess code. I have found that this is hardly a limitation, however, since large companies tend to buy mostly information systems, which

Table 1. Grading scale used in reverse-engineering databases.

Grade	Explanation	Design flaws	Model flaws
A	Clean	• No significant flaws	• Style is reasonable and uniform
B	Structural flaws, but they don't affect the application (can be repaired without much disruption)	• Data types and lengths not uniformly assigned • Not-null constraints not used to enforce required fields • Unique keys and enumerations not defined • Columns have cryptic names like *Cell123*	• Anonymous fields that application code must interpret
C	Major flaws that affect the application (bugs, low performance, difficult maintenance)	• Undefined primary keys • Propagated identity • Haphazard indexing • Foreign-key data type mismatches primary key • Parallel foreign keys	• Needless complexity • Excessive inheritance • Specific modeling errors
D	Severe flaws that compromise the application	• Much unnecessary, redundant data • Extensive binary data (compiled programming language data structures), subverting the declaration of data • Gross denormalization • Dangling foreign-key references	• Lack of crisp conceptualization • Many arbitrary restrictions
F	Appalling (the application won't run properly or runs only because of brute-force programming)	• Gross design errors	• Deep conceptual errors

location_address_1	location_address_2	location_address_3
456 Chicago Street	Decatur, IL xxxxx	
198 Broadway Dr.	Suite 201	Chicago, IL xxxxx
123 Main Street	Cairo, IL xxxxx	
Chicago, IL xxxxx		

Figure 1. Database design flaw—anonymous fields.

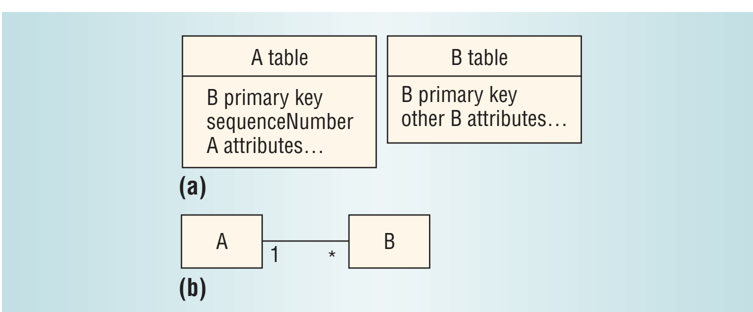

Figure 2. Database design flaw—using a sequence number for a one-to-many relationship: (a) tables as implemented and (b) the logical intent.

are built around a database. I've also found that vendors are more willing to provide their database structure than programming code, which makes database reengineering a more realistic improvement strategy.

WHY REVERSE ENGINEERING?

Forward engineering is the process of building software. Development flows forward from product conception, through analysis and design, and then finally to implementation. Most developers view this flow as more iterative than sequential, given the endless reviews and feedback loops. Reverse engineering begins at the end, with the application code, and works backward to deduce the requirements that spawned the software.[1]

Reverse engineering is certainly not new. Indeed, savvy developers use it to study existing applications and salvage useful ideas, data, and code. To my knowledge, however, few software engineers have applied reverse engineering to vendor assessment. I believe this is an oversight and that reverse engineering of vendor products should be a routine aspect of all software evaluations.[2]

CONNECTING DATABASE AND PRODUCT QUALITY

When we began evaluating databases in 1992, we decided to adopt the grading system in Table 1 to summarize the results of reverse engineering. We have found that business leaders readily understand the meaning of the grade, realize that we have the supporting technical details, and appreciate being allowed to study the details at their leisure. As the table shows, we assess the quality of both the database design and the conceptual model that underlies the database using A, B, C, D, or F, with A being the best grade and F the worst.

The first time we performed database reverse engineering, it was as an experiment. We were studying a vendor product and were perplexed by our experiences. The vendor had a great marketing story and clearly understood the business requirements. The company was both large and credible, so we expected high-quality software. When we encountered a number of problems with the product, we decided to look at the database and discovered that its poor quality was at the root of the problems.

This experience prompted us to look at additional databases, after which we decided that database reverse engineering was not an odd technology, but something we should routinely perform. We started keeping records of our experiences and have amassed 11 years of data in a grading table (available in its entirety at www.modelsoftcorp.com). I believe our results represent broad practice, given that a different team prepared each database. We evaluated databases only from developers we did not advise as part of our consultant work, and we included databases only for applications that a vendor actually completed.

The case studies include both vendor assessments and in-house reengineering and where possible indicate if the application succeeded or failed, with "success" defined as actual use. Of course, by definition, the case studies are biased toward success, given that we (as customers) did not see the products companies scrapped.

QUALITY PROBLEMS

Our assessments revealed many applications with flawed databases. More important, database quality has improved little in 11 years. At best, I can give only a flavor of the problems we encountered, but we did see many recurring design flaws. Several databases had anonymous address fields, for example, which satisfies database design theory, but is sloppy nonetheless.

Consider the data in Figure 1. To find a city, you must search multiple fields. Worse yet, it could be difficult to distinguish Chicago the city from Chicago the street. Furthermore, you might need to parse a field to separate city, state, and postal code. A better design would put address data in distinct fields that are clearly named.

Some database designs used a sequence number to resolve a one-to-many relationship so that it could be buried on the "one" side, as in Figure 2. The sequence number is a completely meaningless field, making it difficult to find data in the A table.

Another common flaw was overloaded foreign keys. In the tables of Figure 3, each address can be linked either to a person or a company (targetID) as indicated by switch.

Several databases had propagated identity, as in Figure 4. From reading the database textbooks, you might think that propagated identity is fine, but

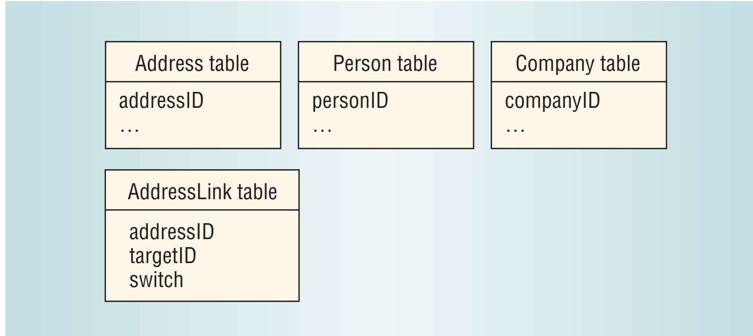

Figure 3. Database design flaw—overloaded foreign keys.

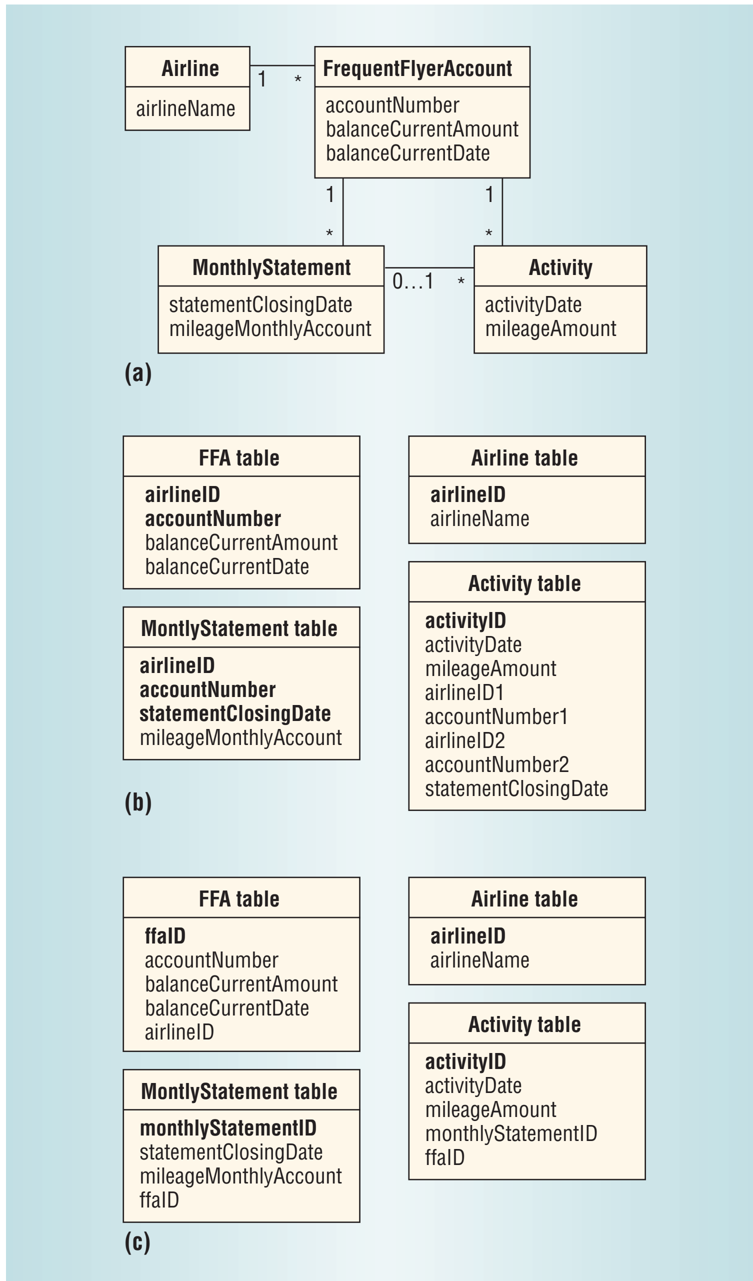

Figure 4. Database design flaw—propagated identity: (a) UML model; (b) implementation with propagated identity, which causes two sources to propagate a field; and (c) a better implementation with existence-based identity.

Company table

companyName	address	phoneNumber
Simplex	100 S. Main	555-1234
Acme	111 Broadway	555-3141
AT&T	201 Potomac	555-1812

(a)

Job table

jobName	estimated StartDate	estimated Duration	estimated StartDate
Simplex new PC	Jan 1, 1990	1 month	$5,000
Acme PC repair	Jul 3, 1986	2 weeks	$200
A.T.&T. LAN upgrade	Feb 9, 1987	3 months	$20,000
acme disk repair	Oct 5, 1986	2 weeks	$50

(b)

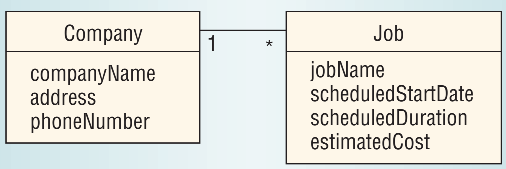

(c)

Figure 5. Database design flaw—informally linking tables: (a) company table, (b) job table, and (c) logical intent.

from a modeling perspective, it is clearly an inferior approach. Propagated identity leads to multiattribute primary keys and can give rise to the situation in which a field is propagated from two sources, but must be duplicated to enforce referential integrity. In the Activity table in Figure 4b, for example, airlineID1 and airlineID2 really represent the same field, but the design requires two separate fields—two copies—to satisfy referential integrity. Thus, AirlineID1 + accountNumber1 refers to the FFA table. AirlineID2 + accountNumber2 + statementClosingDate refers to the MonthlyStatement table. Unfortunately, with SQL, a field cannot refer to two tables, so a single airline field cannot refer to both FFA and MonthlyStatement without causing mathematical ambiguity. Figure 4c shows a correct implementation with existence-based identity, which eliminates the need for duplication.

One database informally linked tables by human inspection, rather than formally linking them via referential integrity. In Figure 5a and 5b, the company name is embedded in the job name. Each job pertains to a customer company; a company can have many jobs. The link may be easy for a human to detect, but it can be difficult for a machine. First, the link between the columns is not declared in the database. However, even if the link is known, the precise computational relationship varies and is arbitrary. In Figure 5b, company name is a substring of job name, but there are case differences (Acme vs. acme) and punctuation differences (AT&T vs. A.T.&T.).

INTERPRETING RESULTS

The data we collected in our reverse-engineering case studies proved to be a useful indicator of historical software quality. Table 2 summarizes the grades converted to points: A = 4.0, B = 3.0, C = 2.0, D = 1.0, and F = 0. Although there is some scatter in the data, the table shows clearly that database designs and models have improved little over the past 11 years. My colleagues and I suspect that the increased use of database design tools could explain the modest improvement in database design quality. Even so, the average database design and conceptual model remain mediocre.

The data essentially confirms our conclusion from anecdotal observation that modeling shows little improvement. Models baffle many developers, who do not appreciate the leverage that modeling can provide in building applications. Roger Box and Michael Whitelaw were quite accurate in observing that abstraction is the most difficult aspect of modeling.[3]

We believe that the root cause for the lack of improvement is that universities are not teaching students how to model software. Many universities teach the syntax of modeling, but they don't teach the art and thought processes. We presume that most professors are not teaching modeling because they don't know how to do it themselves.

We also found a correlation between vendor applications and software developed in-house. As Table 3 shows, the two have comparable quality, which means that software houses are not necessarily more professional than IT departments within corporations, as some might expect.

REVERSE ENGINEERING AND ETHICS

Many articles tend to give reverse engineering a sinister image, implying that developers typically use it to re-create a product.[4] In all our case studies, we made it clear that this was *not* our goal. Besides being unethical, reimplementation is usually uneconomical. Instead, we assured vendors that our focus was to assess the software's merit, to get past hidden assumptions and the sales claims, and to gain a deeper understanding of the product so that we could better communicate with the vendor and use the software more effectively. In short, when we assess products, we are merely trying to determine what the vendor is selling.

When we reverse-engineer a product, we openly ask vendors for their database structure and tell them why we want it. If they refuse, we tell them we will penalize them in the evaluation. We performed many of these reverse-engineering case studies under commission from large companies. Large companies emphatically do not want to devote their resources to re-create a product. Commercial software is important to them, but it is incidental to their primary business. Otherwise, they would be writing their own software, not purchasing it. The industrial mentality is to outsource work that is not a core competency, so rewriting commercial software is one of the last things these companies want to do. In light of that, most vendors acquiesce and settle for a nondisclosure agreement to protect them from competitors. We encourage our clients to agree to reasonable nondisclosure terms and to make it clear to the vendor that there is no intent to compromise its technology or reveal its secrets to competitors.

Some vendors might find reverse engineering threatening, but it should worry only the inept. Superb vendors should welcome the process because it makes their excellence visible in a much more credible way than words or an impressive sales ad.

Database quality is undeniably a good indicator of application quality. I have found that business leaders welcome the insights gained from database reverse engineering and use them to make more informed decisions about purchasing an application. More important, the benefits have the potential to ripple into the entire computing community. If vendors improve, in-house software development will follow suit. The same personnel, over time, move between vendors and in-house staff, and demand from vendors and customer companies will pressure universities to teach students better.

At this time, only pockets of people are assessing vendor software with database reverse engineering, and I would like to see the practice spread. If everyone would reward excellent vendors and penalize sloppy ones, we could improve overall software quality, not just in databases, but eventually in code. And that would be a copper bullet that benefits the entire computing profession. ■

Table 2. Grade average (converted to points) for database designs and models over time.

Statistic	Grade
Design average, first 21 case studies*	1.7
Design average, last 21 case studies	2.2
Model average, first 21 case studies	1.9
Model average, last 21 case studies	2.2

*The first 21 case studies are roughly from 1992 to 1997; the last 21 are roughly from 1997 to 2002.

Table 3. Relative average grades for vendor and in-house software.

Statistic	Grade
Design average, vendor	2.1
Design average, in-house	1.7
Model average, vendor	2.1
Model average, in-house	2.0

References

1. M. Blaha, *A Manager's Guide to Database Technology: Building and Purchasing Better Applications*, Prentice Hall, 2001.
2. M. Blaha, "On Reverse Engineering of Vendor Databases," *Proc. Working Conf. Reverse Eng.*, IEEE CS Press, 1998, pp. 183-190.
3. R. Box and M. Whitelaw, "Experiences When Migrating from Structured Analysis to Object-Oriented Modeling," *Proc. Australasian Computing Education Conf.*, ACM Press, 2000, pp. 12-18.
4. B.C. Behrens and R.R. Levary, "Practical Legal Aspects of Software Reverse Engineering," *Comm. ACM*, Feb. 1998, pp. 27-29.

Michael Blaha *is a partner of Modelsoft Consulting Corp., a consulting and training company based in Vancouver, BC, and is* Computer*'s area editor for databases and software. His research interests include object-oriented technology, modeling, system architecture, database design, enterprise integration, and reverse engineering. Blaha received a PhD in chemical engineering from Washington University and is a member of the IEEE Computer Society. Contact him at blaha@computer.org.*

2004 IEEE Computer Society
Professional Membership/Subscription Application

Membership and periodical subscriptions are annualized to and expire on 31 December 2004.
Pay full or half-year rate depending upon the date of receipt by the IEEE Computer Society as indicated below.

Membership Options*

All prices are quoted in U.S. dollars

	FULL YEAR Applications received 16 Aug 03 - 29 Feb 04	HALF YEAR Applications received 1 Mar 04 - 15 Aug 04
1 I do not belong to the IEEE, and I want to join **just** the Computer Society	$ 99 ❑	$50 ❑
2 I want to join **both** the Computer Society and the IEEE:		
I reside in the United States	$189 ❑	$95 ❑
I reside in Canada	$170 ❑	$85 ❑
I reside in Africa/Europe/Middle East	$166 ❑	$83 ❑
I reside in Latin America	$159 ❑	$80 ❑
I reside in Asia/Pacific	$160 ❑	$80 ❑
3 I already belong to the IEEE, and I want to join the Computer Society. (IEEE members need only furnish name, address, and IEEE number with payment.)	$ 42 ❑	$21 ❑

Are you now or were you ever a member of the IEEE?
Yes ❑ No ❑ If yes, provide member number if known: _______________

Add Periodicals**

	ISSUES PER YEAR	FULL YEAR PRINT	FULL YEAR ELECTRONIC	FULL YEAR COMBO	HALF YEAR PRINT	HALF YEAR ELECTRONIC	HALF YEAR COMBO
IEEE Computer Society Digital Library *BEST DEAL*	n/a	n/a	$109 ❑	n/a	n/a	$55 ❑	n/a
Computing in Science and Engineering	6	$42 ❑	$34 ❑	$55 ❑	$21 ❑	$17 ❑	$28 ❑
IEEE Computer Graphics and Applications	6	$39 ❑	$31 ❑	$51 ❑	$20 ❑	$16 ❑	$26 ❑
IEEE Design & Test of Computers	6	$37 ❑	$30 ❑	$48 ❑	$19 ❑	$15 ❑	$24 ❑
IEEE Intelligent Systems	6	$37 ❑	$30 ❑	$48 ❑	$19 ❑	$15 ❑	$24 ❑
IEEE Internet Computing	6	$39 ❑	$31 ❑	$51 ❑	$20 ❑	$16 ❑	$26 ❑
IT Professional	6	$40 ❑	$32 ❑	$52 ❑	$20 ❑	$16 ❑	$26 ❑
IEEE Micro	6	$37 ❑	$30 ❑	$48 ❑	$19 ❑	$15 ❑	$24 ❑
IEEE MultiMedia	4	$35 ❑	$28 ❑	$46 ❑	$18 ❑	$14 ❑	$23 ❑
IEEE Pervasive Computing	4	$41 ❑	$33 ❑	$53 ❑	$21 ❑	$17 ❑	$27 ❑
IEEE Security & Privacy	6	$41 ❑	$33 ❑	$53 ❑	$21 ❑	$17 ❑	$27 ❑
IEEE Software	6	$44 ❑	$35 ❑	$57 ❑	$22 ❑	$18 ❑	$29 ❑
IEEE/ACM Transactions on Computational Biology and Bioinformatics *NEW*	4	$35 ❑	$28 ❑	$46 ❑	$18 ❑	$14 ❑	$23 ❑
IEEE/ACM Transactions on Networking†	6	$44 ❑	$33 ❑	$55 ❑	$22 ❑	$17 ❑	$28 ❑
IEEE Transactions on:							
Computers	12	$41 ❑	$33 ❑	$53 ❑	$21 ❑	$17 ❑	$27 ❑
Dependable and Secure Computing *NEW*	4	$31 ❑	$25 ❑	$40 ❑	$16 ❑	$13 ❑	$20 ❑
Information Technology in Biomedicine†	4	$35 ❑	$25 ❑	$44 ❑	$18 ❑	n/a	$22 ❑
Knowledge and Data Engineering	12	$43 ❑	$34 ❑	$56 ❑	$22 ❑	$17 ❑	$28 ❑
Mobile Computing	4	$30 ❑	$24 ❑	$39 ❑	$15 ❑	$12 ❑	$20 ❑
Multimedia†	6	n/a	n/a	$38 ❑	n/a	n/a	n/a
NanoBioscience†	4	$30 ❑	$24 ❑	$38 ❑	$15 ❑	n/a	$19 ❑
Parallel and Distributed Systems	12	$40 ❑	$32 ❑	$52 ❑	$20 ❑	$16 ❑	$26 ❑
Pattern Analysis and Machine Intelligence	12	$44 ❑	$35 ❑	$57 ❑	$22 ❑	$18 ❑	$29 ❑
Software Engineering	12	$38 ❑	$30 ❑	$49 ❑	$19 ❑	$15 ❑	$25 ❑
Visualization and Computer Graphics	6	$34 ❑	$27 ❑	$44 ❑	$17 ❑	$14 ❑	$22 ❑
VLSI Systems†	12	n/a	n/a	$28 ❑	n/a	n/a	$14 ❑
IEEE Annals of the History of Computing	4	$31 ❑	$25 ❑	$40 ❑	$16 ❑	$13 ❑	$20 ❑

Choose PRINT for paper issues delivered via normal postal channels.
Choose ELECTRONIC for 2004 online access to all issues published from 1988 forward.
Choose COMBO for both print and electronic.

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

____
Expiration 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 15% HST or 7% GST to total. AL, AZ, CO, DC, NM, and WV add sales tax to all periodicals. GA, IN, KY, MD, and MO add sales tax to print and combo periodicals. NY add sales tax to electronic periodicals. European Union residents add VAT tax to electronic periodicals.
**** 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.
† Not part of the IEEE Computer Society Digital Library. Electronic access is through www.ieee.org/ieeexplore.

For fastest service, apply online at http://computer.org/join

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
PO Box 3014
Los Alamitos, CA 90720-1314 USA
Phone: +1 714 821 8380
Fax: +1 714 821 4641
E-mail: help@computer.org

Allow up to 8 weeks to complete application processing. Allow a minimum of 6 to 10 weeks for delivery of print periodicals.

IF4B

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

APPLICATION MUST BE SIGNED

NOTE: In order for us to process your application, you must complete and return both sides of this form.

IEEE

IEEE Computer Society

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, 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 ____________

Ad Hoc Networks

Recent developments offer potential solutions to problems encountered in ad hoc networks including topology control, data communication, and service access.

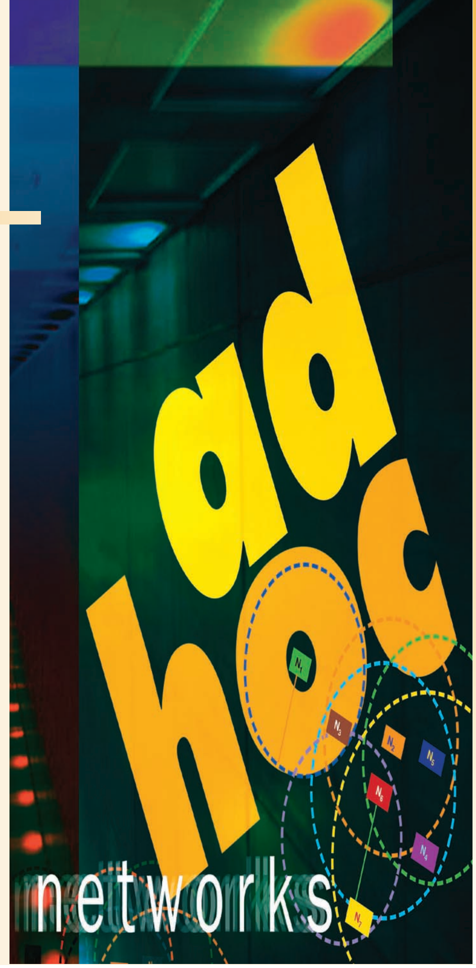

Jie Wu
Florida Atlantic University

Ivan Stojmenovic
University of Ottawa

Ad hoc networks are a key factor in the evolution of wireless communications. Self-organized ad hoc networks of PDAs or laptops are used in disaster relief, conference, and battlefield environments.

These networks inherit the traditional problems of wireless and mobile communications, such as bandwidth optimization, power control, and transmission-quality enhancement. In addition, their multihop nature and the possible lack of a fixed infrastructure introduce new research problems such as network configuration, device discovery, and topology maintenance, as well as ad hoc addressing and self-routing.

Various approaches and protocols have been proposed to address ad hoc networking problems, and multiple standardization efforts are under way within the Internet Engineering Task Force, as well as academic and industrial research projects.

NETWORK TOPOLOGY

In ad hoc networks, wireless hosts can communicate with each other in the absence of a fixed infrastructure.[1] These networks typically consist of equal nodes that communicate over wireless links without central control.

Sensor networks,[2] also called *hybrid ad hoc networks,* are linked to monitoring centers that collect data such as temperature, chemical detection, or movement. In recent years, government agencies in several countries have supported research on sensor networks. For example, the US National Science Foundation launched a multidisciplinary program on sensors and sensor network research in 2003 (www.nsf.gov/pubs/2003/nsf03512/nsf03512.htm).

Some ad hoc networks are linked to a fixed infrastructure via access points. For example, mesh or rooftop networks (www.sonic.net/sales/rooftop/faq.shtml) consist of antennas placed on top of buildings to provide wireless Internet access.

Vehicles on a highway can create an ad hoc network for use in disseminating traffic information. They can operate as a pure ad hoc network in which an individual vehicle detects traffic events and initiates a broadcast to other vehicles. Alternatively, cellular or Internet access points placed near the road can transmit the information.

Multihop cellular networks[3] have recently emerged as a communication alternative at events where huge numbers of users are concentrated in a small area such as a stadium.

Peer-to-peer networks are ad hoc networks in which an overlay network is built on the Internet. In a P2P network, two or more peers can use appropriate information and communication systems to collaborate spontaneously without requiring central coordination.

AD HOC NETWORK COMMUNICATION

Communication between two hosts in an ad hoc network is not always direct—it can proceed in a *multihop* fashion so that every host is also a router.

0018-9162/04/$20.00 © 2004 IEEE

Ad hoc networks are a key factor in the evolution of wireless networks.

Ad hoc network hosts can use protocols such as the IEEE 802.11 media-access control standard to communicate via the same frequency, or they can apply Bluetooth or other frequency-hopping technology.

Because power consumption is directly proportional to the distance between hosts, direct *single-hop* transmissions between two hosts can require significant power, causing interference with other such transmissions. To avoid this *routing problem*, two hosts can use multihop transmission to communicate via other hosts in the network.

With IEEE 802.11 technology, avoiding collisions—transmission interferences—is difficult because of the *hidden station problem*: Two hosts that do not communicate directly can transmit messages simultaneously to a common neighbor on the same frequency.

In addition to maintaining an ongoing routing task or facilitating route establishment, mobile networks also must support location management by keeping track of the host's location.

NETWORK LAYER PROBLEMS

The problems encountered in the network layer of ad hoc networks include topology control, data communication, and service access.

Topology control problems include discovering neighbors, identifying position, determining transmission radius, establishing links to neighbors, scheduling node sleep and active periods, clustering, constructing the dominating set (each node either belongs to or has a neighbor from the dominating set), and maintaining the selected structure.

Data communication problems include

- *routing*—sending a message from a source to a destination node,
- *broadcasting*—flooding a message from a source to all other nodes in the network,
- *multicasting*—sending a message from a source to a set of desirable destinations,
- *geocasting*—sending a message from a source to all nodes inside a geographic region, and
- *location updating*—maintaining reasonably accurate information about the location of other nodes.

Service access problems include Internet access, cellular network access, data or service replication upon detection or expectation of network partition, and unique IP addressing in merge or split-network scenarios.

IN THIS ISSUE

The articles in this special issue review emerging ad hoc networking technologies, techniques, algorithms, and protocols, with emphasis on recent developments offering potential solutions to problems encountered in these networks.

In "Cooperative Cache-Based Data Access in Ad Hoc Networks," Guohong Cao, Liangzhong Yin, and Chita Das propose efficient solutions to the data-caching problem. In cooperative caching, some nodes in an ad hoc network replicate data from servers, using replicated files rather than original files to satisfy other nodes' access demands. This should reduce traffic in the network or even provide service if the server becomes disconnected in the meantime. The proposed solutions include caching data paths toward replicated copy, making another copy of data at the node, and using some novel hybrid methods.

An emerging area of research in sensor networks is area coverage and monitoring. In "Energy-Efficient Area Monitoring for Sensor Networks," Jean Carle and David Simplot-Ryl classify sensor data reporting into two categories: *event-driven* and *on-demand*. They propose dividing the area-monitoring problem into three subproblems, each of which requires an energy-efficient solution. These subproblems consist of constructing a broadcast tree (request propagation), selecting sensors for area coverage, and reporting sensor data with data aggregation. The protocols implement periodic changes in sensor roles to extend network life. The proposed solutions use dominating sets and localized minimal spanning trees.

In "Cross-Layering in Mobile Ad Hoc Network Design," Marco Conti and coauthors describe a European project that overcomes manet performance problems by allowing protocols belonging to different layers to cooperate, sharing network status information while still maintaining separate layers. The authors propose applying triggers to the Network Status so that it can send signals between layers. This lets each layer maintain network information and adapt its performance accordingly. This innovative cross-layering approach addresses, in particular, security and cooperation, energy management, and quality-of-service issues.

Many potential mobile ad hoc network applications involve collaboration among a group of nodes. Group communication models include one-to-many, one-to-any, many-to-many, and one-to-all patterns that facilitate collaboration among a group of nodes. In "Group Communications in Mobile Ad Hoc Networks," Prasant Mohapatra, Chao Gui, and Jian Li describe various techniques for group communications in ad hoc networks, including multicasting, broadcasting, anycasting, and geocasting and discuss representative protocols for each of these categories. They also provide an overview of related issues such as protocol design, state maintenance, and performance; examine issues such as reliability, power conservation, quality of service, and security; and comment on future research directions for group communications in ad hoc networks.

In "Routing and Security in Mobile Ad Hoc Networks," Nikola Milanovic and coauthors provide a survey of routing, flooding for routing, and security issues, based on current IETF drafts. The authors describe four nonposition-based routing algorithms: on-demand reactive dynamic source routing, ad hoc on-demand distance vector routing, proactive table-based optimized link-state routing, and topology broadcast based on reverse-path forwarding. They also discuss a recently proposed hybrid approach that combines the advantages of on-demand and optimized link-state routing for wireless sensor networks. To establish secure routing, the hybrid approach uses a mechanism similar to multipoint relays that applies threshold cryptography and attempts to find a fault-free path to each node only when needed.

Instead of using the traditional IP-based network layer to implement multicast routing protocols, in "Prioritized Overlay Multicast in Mobile Ad Hoc Environments," Li Xiao and coauthors propose a model that improves the efficiency and robustness of overlay multicast in manets by building multiple role-based prioritized trees, possibly with the help of location information about member nodes. Like P2P networks, POM forms a virtual network, consisting of only member nodes, on top of the physical infrastructure. Member nodes can form a short-term multicast group to perform certain important tasks. Overlay trees can have different levels of priority depending on the importance of the service they perform. This approach avoids the need to change the application layer tree when the underlying network changes.

We thank all the authors for their submissions to this special issue. We are indebted to the many referees who reviewed each paper promptly and professionally. We also extend special thanks to *Computer*'s staff for their professional guidance in editing the papers selected for publication. ■

References

1. S. Basagni et al., eds., *Mobile Ad Hoc Networking*, IEEE Press, 2003.
2. I.F. Akyildiz et al., "A Survey on Sensor Networks," *IEEE Comm.*, Aug. 2002, pp. 102-114.
3. Y-D. Lin and Y-C. Hsu, "Multihop Cellular: A New Architecture for Wireless Communications," *Proc. IEEE Infocom 2000*, IEEE Press, 2000, pp. 1273-1282.

***Jie Wu** is a professor in the Department of Computer Science and Engineering at Florida Atlantic University. His research interests are in the areas of wireless networks and mobile computing, routing protocols, fault-tolerant computing, and interconnection networks. Wu received a PhD in computer engineering from Florida Atlantic University. He is a member of the IEEE Computer Society and the ACM and is a senior member of the IEEE. Contact him at jie@cse.fau.edu.*

***Ivan Stojmenovic** is a professor of computer science in the School of Information Technology and Engineering and director of the Ottawa-Carleton Institute for Computer Science at the University of Ottawa. His research interests include wireless ad hoc, sensor, and cellular networks. Stojmenovic received a PhD in mathematics from the University of Zagreb, Croatia. Contact him at ivan@site.uottawa.ca.*

Cooperative Cache-Based Data Access in Ad Hoc Networks

A cooperative cache-based data access framework lets mobile nodes cache the data or the path to the data to reduce query delays and improve data accessibility.

Guohong Cao
Liangzhong Yin
Chita R. Das
Pennsylvania State University

Mobile ad hoc networks have potential applications in civilian and military environments such as disaster recovery efforts, group conferences, wireless offices, mobile infostations (in tourist centers, restaurants, and so on), and battlefield maneuvers, making them a focus of current research.

A battlefield ad hoc network might consist of several commanding officers and a group of soldiers. The soldiers could access officers' information centers for detailed geographic information, information about the enemy, new commands, and so on. Because neighboring soldiers tend to have similar missions and thus common interests, several soldiers might need to access the same data at different times. Having a nearby soldier serve later accesses to this data instead of the faraway information center saves battery power, bandwidth, and time.

In ad hoc networks, mobile nodes communicate with each other using multihop wireless links. Due to a lack of infrastructure support, each node acts as a router, forwarding data packets for other nodes. Most previous research in ad hoc networks focused on the development of dynamic routing protocols that can efficiently find routes between two communicating nodes. Although routing is an important issue, the ultimate goal of ad hoc networks is to provide mobile nodes with access to information.

If mobile users around infostations, which have limited coverage, form an ad hoc network, a mobile user who moves out of the range of a particular infostation can still access the data it contains. If one of the nodes along the path to the data source has a cached copy of the requested data, it can forward the data to the mobile user, saving bandwidth and power. Thus, if mobile nodes can work as request-forwarding routers, they can save bandwidth and power and reduce delays.

Cooperative caching, in which multiple nodes share and coordinate cached data, is widely used to improve Web performance in wired networks. The "Related Work in Cooperative Caching" sidebar provides additional information about recent research focusing on cooperative caching approaches for wired networks. However, resource constraints and node mobility have limited the application of these techniques in ad hoc networks.

Our proposed caching techniques—CachePath, CacheData, and HybridCache—use the underlying routing protocols to overcome these constraints and further improve performance by caching the data locally or caching the path to the data to save space.

To increase data accessibility, mobile nodes should cache different data items than their neighbors. Although this increases data accessibility, it also can increase query delays because the nodes might have to access some data from their neighbors instead of accessing it locally. In addition, replicating data from the server could create security problems.

As Figure 1 illustrates, the cooperative cache-

 Published by the IEEE Computer Society

Related Work in Cooperative Caching

Existing cooperative caching schemes for the Web environment can be classified as message-based, directory-based, hash-based, or router-based.

Duane Wessels and Kim Claffy introduced the standardized and widely used Internet cache protocol.[1] As a message-based protocol, ICP supports communication between caching proxies using a simple query-response dialog.

Directory-based protocols for cooperative caching—such as cache digests[2] and summary cache[3]—let caching proxies exchange information about cached content.

The cache array routing protocol is the most notable hash-based cooperative caching protocol. The rationale behind CARP constitutes load distribution by hash routing among Web proxy cache arrays.

As a router-based protocol, the Web cache coordination protocol transparently distributes requests among a cache array. Because these protocols usually assume fixed network topology and often require high computation and communication overhead, they might be unsuitable for ad hoc networks.

To tolerate network partitions and improve data accessibility, Takahiro Hara proposed several replica allocation methods for ad hoc networks.[4] In Hara's schemes, a node maintains replicas of data that is frequently requested. The data replicas are relocated periodically based on three criteria: access frequency, neighbor nodes' access frequency, or overall network topology. Later, Hara proposed schemes to deal with data updates. Although data replication can improve data accessibility, significant overhead is associated with maintaining and redistributing the replicas, especially in ad hoc networks.

Maria Papadopouli and Henning Schulzrinne[5] proposed a 7DS architecture similar to cooperative caching, which defines two protocols to share and disseminate data among users experiencing intermittent Internet connectivity. It operates on a prefetch mode to gather data for serving the user's future needs or on an on-demand mode to search for data on a single-hop multicast basis. The 7DS architecture focuses on data dissemination instead of cache management. Further, it focuses on a single-hop rather than a multihop environment.

References

1. D. Wessels and K. Claffy, "ICP and the Squid Web Cache," *IEEE J. Selected Areas in Comm.*, Mar. 1998, pp. 345-357.
2. A. Rousskov and D. Wessels, "Cache Digests," *Computer Networks and ISDN Systems*, vol. 30, 1998, pp. 2155-2168.
3. L. Fan et al., "Summary Cache: A Scalable Wide Area Web Cache Sharing Protocol," *Proc. ACM SIGCOMM*, ACM Press, 1998, pp. 254-265.
4. T. Hara, "Effective Replica Allocation in Ad Hoc Networks for Improving Data Accessibility," *Proc. IEEE Infocom*, IEEE CS Press, 2001, pp. 1568-1576.
5. M. Papadopouli and H. Schulzrinne, "Effects of Power Conservation, Wireless Coverage and Cooperation on Data Dissemination among Mobile Devices," *Proc. MobiHoc*, ACM Press, 2001, pp. 117-127.

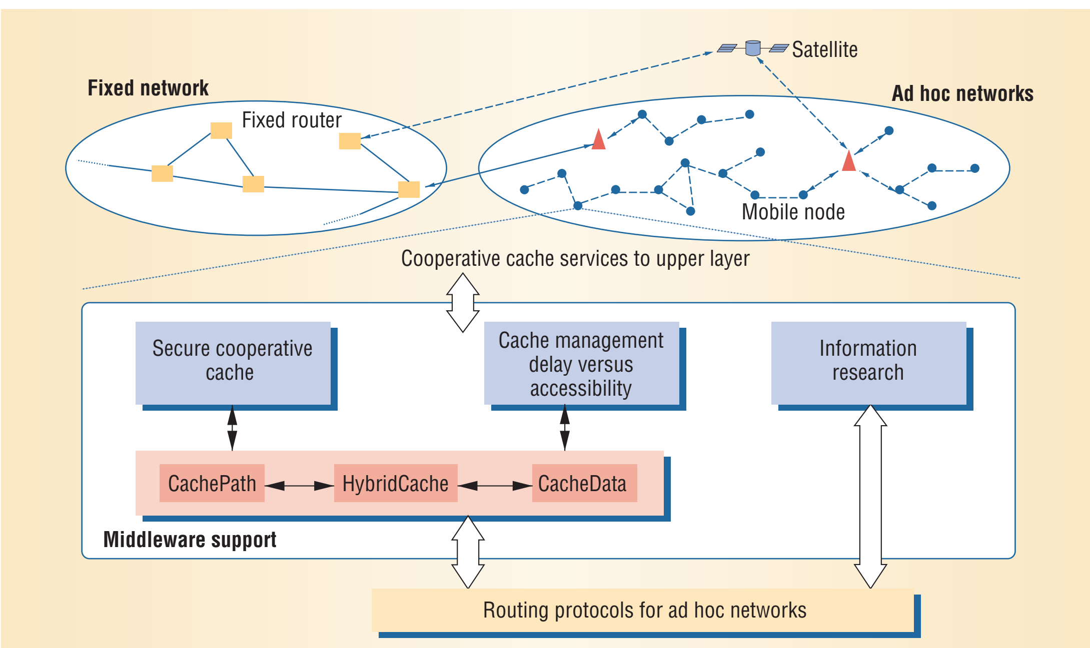

Figure 1. Cooperative caching in ad hoc networks. Middleware support mechanisms provide secure cooperative caching, cache management, and information search.

based framework stays on top of the routing protocols. It relies on several components, such as secure cooperative caching, cache management, and information search to provide services to the upper layer.

ROUTER SUPPORT FOR COOPERATIVE CACHING

Suppose that in the ad hoc network in Figure 2, N_{11} is a data source containing a database of n items d_1, d_2, …, and d_n. N_{11} might be a connecting

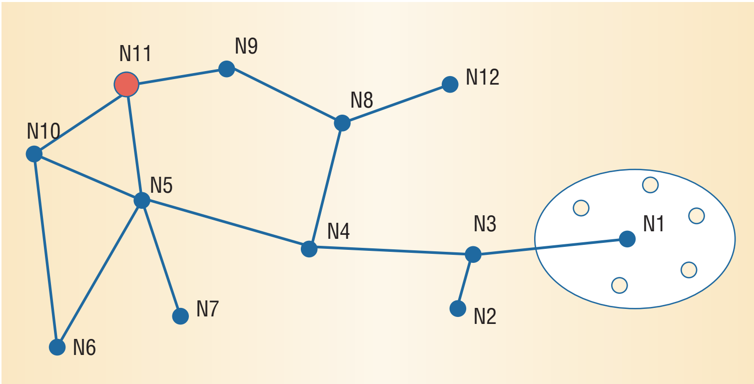

Figure 2. Ad hoc network. Node N_{11} is a data source and the blue nodes are router nodes. Node N_1 is a cluster head surrounded by mobile nodes.

node to the wired network with the database. The blue nodes are router nodes, which can be cluster heads if a cluster-based routing protocol is used; otherwise, they are just mobile nodes. Around each cluster head (as illustrated by node N_1) are several mobile nodes, or cluster members. To reduce bandwidth and power consumption, the number of hops between the data source and the requesting node should be as small as possible.

Routing protocols can help reduce bandwidth and power consumption to a limited degree. Our proposed caching techniques use the underlying routing protocols to further improve performance.

CachePath and CacheData concepts

Figure 2 illustrates the CachePath concept. Suppose node N_1 requests a data item d_i from N_{11}. When N_3 forwards d_i to N_1, N_3 knows that N_1 has a copy of the data. Later, if N_2 requests d_i, N_3 knows that the data source N_{11} is three hops away whereas N_1 is only one hop away. Thus, N_3 forwards the request to N_1 instead of N_4. Many routing algorithms provide the hop count information between the source and destination. Caching the data path for each data item reduces bandwidth and power because nodes can obtain the data using fewer hops. However, mapping data items and caching nodes increases routing overhead. We propose various optimization techniques to improve CachePath's performance.

In CachePath, a node need not record the path information of all passing data. Rather, it only records the data path when it's closer to the caching node than the data source. For example, when N_{11} forwards d_i to the destination node N_1 along the path $N_5 - N_4 - N_3$, N_4 and N_5 won't cache d_i's path information because they're closer to the data source than the caching node N_1.

In a mobile network, the node caching the data might move or it might replace the cached data because of cache size limitations. Consequently, the node modifying the route should reroute the request to the original data source after discovering that the node moved or replaced the data. Thus, the cached path might be unreliable, and using it can increase the overhead.

To deal with this issue, in our approach a node caches the data path only when the caching node is very close. The closeness can be defined as a function of the node's distance to the data source, its distance to the caching node, route stability, and the data update rate. Intuitively, if the network is relatively stable, the data update rate is low, and its distance to the caching node is much shorter than its distance to the data source, the routing node should cache the data path.

In CacheData, the router node caches the data instead of the path when it finds that the data is frequently accessed. For example, in Figure 2, if both N_6 and N_7 request d_i through N_5, N_5 might think that d_i is popular and cache it locally. N_5 can then serve N_4's future requests directly. Because the CacheData approach needs extra space to save the data, it should be used prudently.

Suppose N_3 forwards several requests for d_i to N_{11}. The nodes along the path—N_3, N_4, and N_5—might want to cache d_i as a frequently accessed item. However, they'll waste a large amount of cache space if they all cache d_i. To avoid this, CacheData enforces another rule: A node does not cache the data if all requests for the data are from the same node.

In this example, all the requests N_5 received were from N_4, and those requests in turn came from N_3. With the new rule, N_4 and N_5 won't cache d_i. If N_3 receives requests from different nodes, for example, N_1 and N_2, it caches the data. If the requests all come from N_1, N_3 won't cache the data but N_1 will, or the requesting node in N_1's cluster will cache the data if it's the only requesting node. Certainly, if N_5 later receives requests for d_i from N_6 and N_7, it can also cache the data.

Maintaining cache consistency. To handle cache consistency, CachePath and CacheData use a simple weak consistency model based on the time-to-live mechanism. In this model, a routing node considers a cached copy up-to-date if its TTL hasn't expired. If the TTL has expired, the node removes the map from its routing table (or removes the cached data). As a result, the routing node forwards future requests for this data to the data source. We optimize this model by allowing nodes to refresh a cached data item if a fresh copy of the same data passes by. If the fresh copy contains the same data but a newer TTL, the node updates only the cached data's TTL field. If the data center has updated the data item, the node replaces both the cached data item and its TTL with the fresh copy. When strong cache consistency is needed, we adopt techniques reported in earlier work.[1]

```
(A) When a data item d_i arrives:
        if d_i is the requested data by the current node
        then cache d_i
        else /* Data passing by */
        if there is a copy of d_i in the cache
        then update the cached copy if necessary
        else if s_i < T_s or TTL_i < T_TTL then
            cache data item d_i and the path;
        else if there is a cached path for d_i, then
            cache data item d_i;
            update the path for d_i;
        else
            cache the path of d_i;

(B) When a request for data item d_i arrives:
        if there is a valid copy in the cache
        then send d_i to the requester;
        else if there is a valid path for d_i in the cache then
            forward the request to the caching node;
        else
            forward the request to the data center;
```

Figure 3. HybridCache concept. The hybrid caching scheme caches the data or path based on criteria such as data item size and TTL time.

Locating cached data within a cluster. To save power, many routing protocols divide areas into clusters (or grids), with only one node in the cluster active while others sleep. To avoid network partitioning and maintain fairness, all nodes in the cluster alternate the role of the cluster head, or *coordinator*. Nodes can also leave or join the cluster. In such an environment, to ensure that all nodes can share the cached data, the cluster head needs to know which node caches which data.

A simple solution has each node send its cache data IDs to the cluster head. Most data IDs are very long, so sending the cached data IDs can consume a lot of bandwidth and power. *Cache digests*—a lossy compression of all cache keys with a lookup capability—offer a low overhead option that facilitates data searching in cluster-based ad hoc networks.[2] A node can check another node's digests to discover (with some uncertainty) whether it holds a given data item. When a node joins a cluster, it sends its cache digests to the cluster head. During the query reply phase, the cluster head calculates the cache digests based on the data ID and updates relevant information accordingly. If the current cluster head needs to sleep, it sends the cache digests to the new cluster head.

By caching the data or the data path, a nearby node can serve requests instead of the distant data center. This reduces query latency as well as bandwidth and power consumption because fewer nodes are involved in the query process. In addition, because the data center handles fewer requests, the workload is spread over the network, reducing the load on the data center.

The HybridCache approach

CachePath and CacheData can significantly improve system performance. Our analysis showed that CachePath performs better when the cache is small or the data update rate is low, while CacheData performs better in other situations.[3]

To further improve performance, we propose HybridCache, a hybrid scheme that exploits the strengths of CacheData and CachePath while avoiding their weaknesses. Specifically, when a mobile node forwards a data item, it caches the data or path based on some criteria. These criteria include the data item size s_i and the TTL time TTL_i.

For a data item d_i, we use the following heuristics to decide whether to cache data or the path:

- If s_i is small, CacheData is optimal because the data item only needs a small part of the available cache; otherwise, CachePath is preferable because it saves cache space. We denote the data size threshold as $\mathcal{T}_s$.
- If TTL_i is small, CacheData is preferable. Because the data item might soon be invalid, using CachePath can result in chasing the wrong path and having to resend the query to the data center. For large TTL_is, however, CachePath is acceptable. We denote the TTL threshold value as $\mathcal{T}_{TTL}$.

To achieve better performance, the threshold values used in these heuristics must be set carefully.[3] Figure 3 shows the algorithm applying these heuristics in the HybridCache scheme.

In CachePath, caching a data path only requires saving a node ID in the cache, which has a very small overhead. As a result, in HybridCache, when a node caches a data item d_i using CacheData, it also caches d_i's path. Later, if the cache replacement algorithm decides to remove d_i, it removes the cached data but keeps the path. Thus, CacheData degrades to CachePath for d_i. Similarly, CachePath can upgrade to CacheData again when data item d_i passes by.

When TTL expires, some cached data can be invalidated. Usually, the node removes such invalid data from the cache. However, invalid data items can be useful. For example, caching the data indicates the mobile node's interest in it.

When forwarding a data item, if a mobile node finds an invalid copy of that data in its cache, it caches the new copy for future use. To save space, when a cached data item expires, the mobile node removes the item from the cache, keeping the data's path in invalid state to indicate its interest. Because the mobile node's interest can change, it should not keep the expired path in the cache forever. In our design, if an expired path or data item has not been refreshed for the duration of its original TTL time

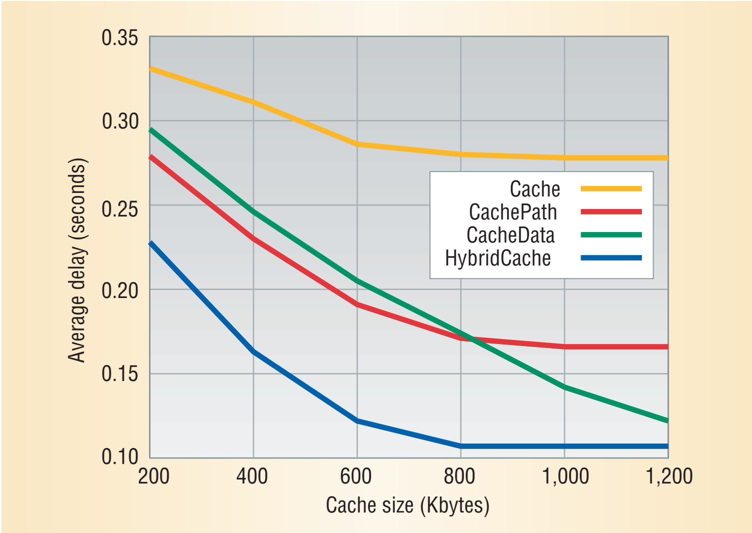

Figure 4. The average query delay as a function of the cache size. HybridCache performs better than both CacheData and CachePath because it combines the strengths of the two approaches.

(set by the data center), the node removes it from the cache.

Experimental results

To compare the performance of CachePath, CacheData, HybridCache, and traditional caching, in which a mobile node only caches the data it requests from the data center, we developed a simulation platform based on a simulator called ns-2[3].

Figure 4 shows the impact of cache size on average query delay. When the cache is small, CachePath performs better than CacheData. When the cache is more than 800 Kbytes, CacheData performs better because it uses more cache space to save passing data.

HybridCache performs better than either approach because it applies either CacheData or CachePath to different data items. HybridCache dynamically switches between CacheData and CachePath at the data item level, not at the database level.

PROACTIVE COOPERATIVE CACHING

With HybridCache, router nodes help other mobile nodes get the requested data quickly. A mobile node doesn't know whether the data source or some other nodes serve its request. If multiple data sources exist, or if the mobile node doesn't know where the data source is, HybridCache might not be a good option. In addition, caching nodes outside the path between the requesting node and the data source might not be able to share cache information with the requesting node.

Proactive cooperative caching, in which the requesting node actively searches for data from other nodes, is a possible solution. Rather than locating data within a cluster, the information search can go through multiple hops.

In proactive cooperative caching, the requesting node broadcasts a request to its neighbor nodes. If a node receiving the request has the data in its local cache, it sends an acknowledgment (ACK) to the requesting node; otherwise, it forwards the request to its neighbors. In this way, a request is flooded to other nodes and eventually acknowledged by the data source or a node with the cached copy.

The requesting node uses the arriving order of ACKs from the mobile nodes to select a path to the requested data. For example, if the requesting node gets an ACK from N_i earlier than it gets one from N_j, it requests the data from N_i. To improve performance, if the data size is small, it can be piggybacked with the ACK.

Flooding can create problems such as redundancy, contention, and collision—collectively referred to as the *broadcast storm* problem.[4] When a node receives multiple broadcast requests, many of its neighbors have already sent out broadcast requests, and an extra broadcast might not add much additional coverage. However, this extra broadcast can increase network contention and collision, thus the broadcast might be unnecessary. Techniques used to reduce the broadcast storm problem can be applied to information search to reduce flooding overhead. Other techniques such as dominant pruning[5] and dominating sets[6] can be used to further reduce collision, contention, and redundancy problems.

Because of flooding overhead, most information searches should be limited to some range instead of the entire network. One simple solution is based on the hop counter concept used in routing protocols. When the hop counter reaches some threshold, nodes stop flooding.

Setting up the hop counter is challenging. A hop counter that is too small reduces network traffic, but at the cost of not finding the data. A hop counter that is too large increases the probability of finding the data, but it also increases network traffic.

Intuitively, if the data size is very large or if the requesting node is closer to the caching node than the data source, the information search overhead might pay off. Otherwise, limiting the search range is preferable. Thus, the hop counter's value is related to the data size, flooding message size, distance to the caching node, and distance to the source node.

Information search is always valuable in network partitions, because the data may be found through information search when the data center is not accessible. However, because information search creates a large amount of network traffic, it should be used only when the data source is unreachable, using the basic cooperative cache scheme otherwise.

CACHE MANAGEMENT

Cache management is more complex in cooperative caching because deciding what to cache can also depend on the node's neighbors. Cooperative caching presents two problems: cache replacement and cache admission control.

Cache replacement algorithms

When the cache is full, cache replacement algorithms can find a suitable subset of data items to evict from the cache. Cache replacement algorithms have been extensively studied in operating systems, virtual memory management, and database buffer management. However, these algorithms might be unsuitable for ad hoc networks for several reasons:

- Because the data item size is not fixed in wireless environments, the least recently used policy must be extended to handle data items of varying sizes.
- The data item's transfer time might depend on the item's size and the distance between the requesting node and the data source (or cache). Consequently, the cache hit ratio might not be the most accurate measurement of a cache replacement algorithm's quality.
- The cache replacement algorithm should also consider cache consistency—that is, data items that tend to be inconsistent earlier should be replaced earlier. For example, one item is accessed 1 percent of the time at a particular client and is also updated 1 percent of the time.[7] A second item is accessed 0.5 percent of the time at the client, but updated only 0.1 percent of the time. In this example, the LRU algorithm would replace the second item and keep the first. However, keeping the second item might result in better performance.

Most cache replacement algorithms designed for the Web are function-based, employing factors such as time since last access, the item's entry time in the cache, transfer time cost, and item expiration time. Most of these cost functions are valid only for the traces used, however, and are not generic enough for insightful observations. These algorithms work relatively well under a certain goal—for example, response time. When the goal changes, they must produce another experience function.

Environments in which data accessibility is a concern require an enhanced cache replacement policy. Here, we consider two factors in selecting the data item (victim) to be replaced. The first factor is the distance (δ)—the number of hops away from the data source or caching node. Because δ is closely related to latency, if the node selects the data item with a higher δ as the victim, access latency will be high. Therefore, it selects the data item with the lowest δ.

The second factor is data access frequency. Node mobility means the network topology can constantly change. As the topology varies, the δ value becomes obsolete. Therefore, we use a parameter τ that captures the time elapsed since the last δ update. We obtain τ by $1/(t_{cur} - t_{update})$, where t_{cur} and t_{update} are δ's current and last updated times for the data item. If τ is close to 1, δ has recently been updated. If τ is close to 0, the update gap is long. Thus, we use τ as an indicator of δ to select a victim. Using these two factors, we can design different cache replacement algorithms. For example, the victim can be the item with the least value of $(\delta + \tau)$ or $(\delta \times \tau)$.

Node mobility means the network topology can constantly change.

Some preliminary results show that different values for these parameters affect the tradeoffs between data accessibility and query delay.[8]

Cache admission control

Cache admission control decides whether a data item should be brought into the cache. Inserting a data item into the cache might not always be favorable because it can lower the probability of cache hits. For example, replacing a data item that will be accessed soon with a data item that won't be accessed in the near future degrades performance. We can use the cache replacement value function to implement cache admission control simply by comparing the requested item's cost to the cached item with the highest cost.

To increase data accessibility, a node cannot cache data that some neighbors already cache; rather, it uses its local space to cache more data. For example, if the requesting node is less than *a* (a system parameter) hops away from a node that has cached the data, it won't cache the data. Thus, the same data item is cached at least *a* hops apart.

A tradeoff exists between query latency and data accessibility. With a small *a*, the number of replicas for each data item is high, and the access delay for this data is low. However, with a fixed amount of cache space, mobile nodes can cache a limited number of distinct data items.

If a network partition exists, many nodes might not be able to access this data. On the other hand, with a large *a*, each data item has a small number of replicas, and the access delay can be a little longer. On the positive side, mobile nodes can cache more distinct data items and still serve

The server can prevent nodes from caching some data or limit the number of nodes that can cache it.

requests when the data source is not accessible.

Depending on the application, *a* can take different values. Accessibility can outweigh access latency when network partitions occur frequently. Instead of waiting for the network topology to change, the nodes should maintain high probability of finding the requested data items. A large *a* lets a node distribute more distinct data items over the entire cache due to admission control, increasing the number of accessible data items and thereby improving overall data accessibility. Note that when *a* reduces to 1, the node can cache the data with minimal latency.

SECURE COOPERATIVE CACHE

Caching nodes can replicate data from the server. Because of security concerns, the owners of some sensitive data might want to restrict access to the data, preventing its duplication. We define different levels of data security with regard to duplication and storage in node caches. The data server can specify the level of security for each data item. Depending on the security level, the server can prevent nodes from caching some data or limit the number of nodes that can cache it. For most sensitive data, the data server sends the encrypted version to a few trusted nodes, which decrypt the data using a shared key.

Future research should focus on mechanisms that let the data owner control the caching scope without undermining its flexibility. Such mechanisms should maintain a balance between security strength and system performance because encrypting or limiting the distribution of the most sensitive data can reduce cooperative caching's benefits.

With cooperative caching, mobile nodes can return the cached data or modify the route and forward the request to the caching node; hence, the mobile nodes should not be able to modify the data maliciously. With data authentication, a receiver can ensure that the received data is *authentic*—that is, it originated from the source and was not modified on the way—even when none of the other data receivers is trusted.

Authenticating the data source is more complicated. Appending each packet with a message authentication code calculated using a shared key does not work because any receiver with the shared key can forge the data and impersonate the sender. Consequently, we use solutions based on asymmetric cryptography, namely digital signature schemes. After the data source signs the data with its private key, mobile nodes can verify the data's integrity using the data source's public key.

Because digital signatures have high overhead in terms of both time to sign and verify and bandwidth, we focus on reducing authentication overhead. For example, if the data has gone through nodes with good reputations, the receiver might not need to verify the signature. This trades some security strength for system performance. Periodically, the receiver might want to verify the signature and change a node's credit rating based on the verification results. A mobile node can verify the signature if the data is important or the node has enough computation power. This allows the user to choose the proper tradeoffs between security and performance—for example, trading security strength for performance if the data is not very important and requires less computation power.

Although ad hoc networks have attracted many researchers, most previous research in this area focuses on routing, with little work on data access. We anticipate that our work will stimulate further research on cooperative cache-based data access. For example, how can we reduce the broadcast overhead during information search? How can we maintain cache consistency and deal with attacks on cache consistency? Comparing cooperative cache-based data access with data replication techniques might also be interesting. We are currently setting up a prototype to test the proposed protocols and address the implementation issues. ■

Acknowledgments

This work was supported in part by the US National Science Foundation through career grant CCR-0092770 and grant ITR-0219711.

References

1. G. Cao, "A Scalable Low-Latency Cache Invalidation Strategy for Mobile Environments," *IEEE Trans. Knowledge and Data Eng.*, vol. 15, no. 5, 2003, pp. 1251-1265.
2. A. Rousskov and D. Wessels, "Cache Digests," *Computer Networks and ISDN Systems*, vol. 30, 1998, pp. 2155-2168.
3. L. Yin and G. Cao, "On Supporting Cooperative Cache in Ad Hoc Networks," to appear in *Proc. IEEE Infocom*, IEEE CS Press, 2004.
4. Y. Tseng, S. Ni, and E. Shih, "Adaptive Approaches to Relieving Broadcast Storms in a Wireless Multi-

hop Mobile Ad Hoc Network," *Proc. IEEE Int'l Conf. Distributed Computing Systems*, IEEE Press, 2001, pp. 481-488.
5. W. Lou and J. Wu, "On Reducing Broadcast Redundancy in Ad Hoc Wireless Networks," *IEEE Trans. Mobile Computing*, vol. 1, no. 2, 2002, pp. 111-123.
6. I. Stojmenovic and J. Wu, "Broadcasting and Activity Scheduling in Ad Hoc Networks," to appear in *Ad Hoc Networking*, S. Basagni et al., eds., IEEE Press, 2004.
7. G. Cao, "Proactive Power-Aware Cache Management for Mobile Computing Systems," *IEEE Trans. Computers*, vol. 51, no. 6, 2002, pp. 608-621.
8. S. Lim et al., "A Novel Caching Scheme for Internet-Based Mobile Ad Hoc Networks," *Proc. IEEE Int'l Conf. Computer Comm. and Networks* (ICCCN), IEEE Press, 2003, pp. 38-43.

Guohong Cao *is an assistant professor of computer science and engineering at Pennsylvania State University. His research interests are wireless networks, mobile computing, and distributed fault-tolerant computing. Cao received a PhD in computer science from Ohio State University. Contact him at gcao@cse.psu.edu.*

Liangzhong Yin *is a PhD candidate at Pennsylvania State University. His research interests include wireless/ad hoc networks and mobile computing. Yin received an ME in computer science and engineering from the Southeast University, Nanjing, China. Contact him at yin@cse.psu.edu.*

Chita R. Das *is a professor in the Department of Computer Science and Engineering at Pennsylvania State University. His research interests include parallel and distributed computing, cluster computing, mobile computing, performance evaluation, and fault-tolerant computing. Das received a PhD in computer science from the University of Louisiana, Lafayette. He is a Fellow of the IEEE. Contact him at dsag@cse.psu.edu.*

Energy-Efficient Area Monitoring for Sensor Networks

The nodes in sensor networks must self-organize to monitor the target area as long as possible. Optimizing energy consumption in area coverage, request broadcasting, and data aggregation can significantly extend network life.

Jean Carle
David Simplot-Ryl
University of Lille

Recent advances in microelectromechanical systems, digital electronics, and wireless communications have enabled the development of low-cost, low-power, multifunctional sensor devices. These devices can operate autonomously to gather, process, and communicate information about their environments.

When a large number of devices collaborate using wireless communications and an asymmetric, many-to-one data flow, they constitute a *sensor network*.[1] The sensor nodes usually send their data to a specific sink node or *monitoring station* for collection. If all the nodes communicated directly with the monitoring station, the communication load—especially over long distances—would quickly drain the network's power resources. Therefore, the sensors operate in a self-organized, decentralized manner that maintains the best connectivity as long as possible and communicates messages via multihop spreading.

Sensor networks are a special case of ad hoc networks. Figure 1 shows two kinds of applications. In *event-driven* applications like forest-fire detection, one or several sensor nodes detect an event and report it to a monitoring station. In *demand-driven* applications like inventory tracking in a factory warehouse, sensors remain silent until they receive a request from the monitoring station. In both cases, sensor networks generally deploy nodes densely, using hundreds or thousands of sensors—placed mostly at random—either very near or inside the phenomenon to be studied. The nodes are static and carry less battery and processing power than standard ad hoc networks.

A sensor's battery is not replaceable, so its energy is the most important system resource—especially when the network operates in hostile or remote environments. The best method for conserving energy is to put as many sensors to sleep as possible. At the same time, however, the network must maintain its functionality through a connected subnetwork that lets the monitoring station communicate with any of the network's active sensors.

Our group at the Fundamental Computer Science Laboratory of Lille (LIFL) is developing strategies for selecting and updating an energy-efficient connected active sensor set that extends the network lifetime. We report optimizing solutions to three separate problems:

- *area coverage*—maintaining full coverage of the monitoring area;
- *request spreading*—broadcasting from the monitoring station to the covering nodes; and
- *data aggregation*—transmitting information from nodes to the monitoring center.

Each sensor's monitoring area can be approximated as a disk around the sensor. We further assume that each sensor can measure or observe the physical parameter or event in its own monitoring area and can use radio-frequency technology to communicate with other sensors in its vicinity. The solutions we present here also assume that a sensor's monitoring area is exactly the same as its commu-

Published by the IEEE Computer Society

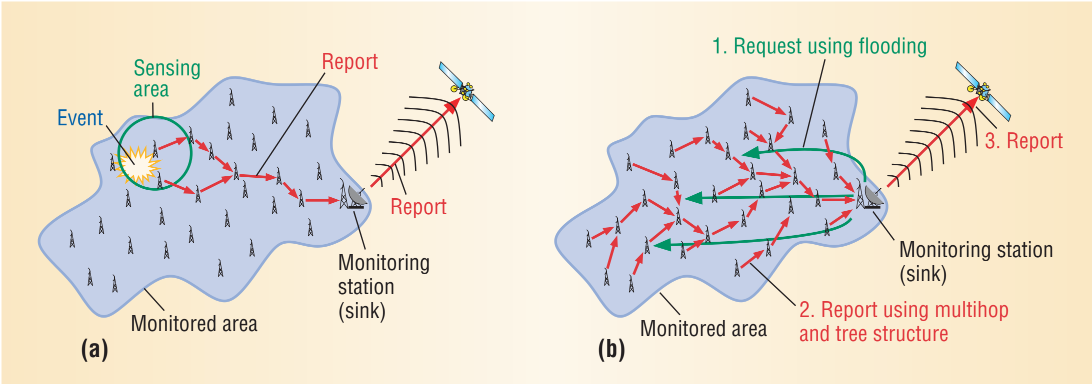

Figure 1. Sensor network applications. (a) In event-driven applications, one or several sensors detect an event and report it to a monitoring station. (b) In demand-driven applications, sensors remain silent until they receive a request from the monitoring station.

nication area—that is, the area in which nodes can receive communication from a transmitting node.

AREA COVERAGE

Consider a sensor node set dropped randomly on a target area that it must monitor. The monitored area is the union of all individual node monitoring areas. Assuming that the sensing radius of each node is the same and that the sensors can obtain their geographical position, solving the area-coverage problem requires finding the *area-dominating set*—that is, the smallest subset of sensor nodes that covers the monitored area. Nodes not belonging to this set do not participate in the monitoring—they sleep. The area-dominating set changes periodically both as a function of activity scheduling and to extend the network's monitoring capability.

Recently, Fan Ye and colleagues[2] proposed a simple localized protocol for dynamically selecting an area-dominating set. According to this protocol, a node sleeps for a while and then decides to be active if and only if there are no active nodes closer than a given threshold distance from it. When a node is active, it remains active until the end of its battery lifetime. Sleeping nodes periodically reevaluate their decision. With this protocol, the probability of having full coverage of a monitored area is close to 1 if the threshold is less than $1/(1+\sqrt{5})$ of the sensing area's radius. However, this protocol has limited usefulness because it is probabilistic and does not ensure full area coverage.

Di Tian and Nicolas D. Georganas[3] proposed a solution that requires every node to know all its neighbors' positions before making its monitoring decisions. Each node then selects a timeout interval. At the end of the interval, if a node sees that neighbors who have not yet sent any messages together cover its monitoring area, the node transmits a "withdrawal" message to all its neighbors and goes into sleep mode. Otherwise, the node remains active but does not transmit any message. The process repeats periodically to allow for changes in monitoring status.

The problem with this solution is that it requires a priori knowledge about all neighboring sensors, which could involve significant communication overhead once sensors start to die between activity periods.

Dominating node sets

We have developed solutions to the area-coverage problem that guarantee coverage—as long as the given sensor nodes cover the area—without requiring nodes to have prior knowledge of neighboring nodes.

Instead of covering an area, the goal is to select a *connected dominating set* of sensor nodes that "monitor" other sensors within their coverage range, or neighborhood. Researchers have studied this problem in the context of ad hoc network broadcasting.[4] A *dominating set* is a subset of network nodes in which each node is either in this subset or is a neighbor of a node in this subset. A dominating set is *connected* if any two nodes in the set can communicate, possibly through other nodes via multihop broadcasting. The broadcasting task is to send a message from one node to all network nodes using only nodes in a connected dominating set.

Among recently developed strategies for constructing small connected dominating sets, *localized protocols* offer the best prospect for achieving energy efficiency. In a localized protocol, each node makes decisions based solely on information about itself and its one-hop neighbor—if position information is also available—or its two-hop neighbors—if position information is not available. Moreover, each node makes decisions without communications between nodes beyond the message exchanges that nodes use to discover each other and establish neighborhood information. The local information must suffice for a node to decide whether or not it is in a connected dominating set; otherwise, the increased communication overhead could offset the energy savings.

Recently, Fai Dai and Jie Wu[5] proposed a distributed dominant-pruning algorithm to meet these solution criteria. This algorithm gives each node a priority, which can be simply its unique identifier or a combination of remaining battery life, number of neighbors, and identifier.

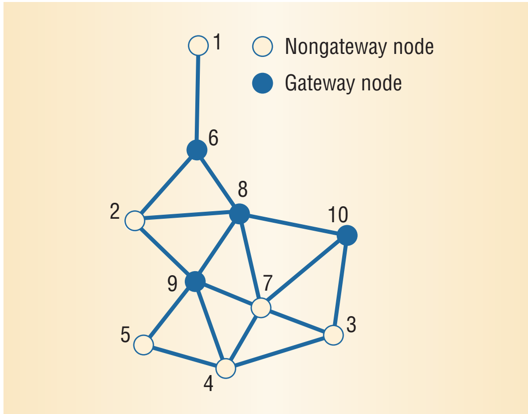

Figure 2. Dominant-pruning algorithm. A node belongs to the dominating set if and only if no subset fully covers it.

A node *u* is "fully covered" by a subset *S* of its neighboring nodes if and only if three conditions hold:

- the subset *S* is connected,
- any neighbor of *u* is a neighbor of at least one node from *S*, and
- all nodes in *S* have higher priority than *u*.

A node belongs to the dominating set if and only if no subset fully covers it. The nodes belonging to a dominating set are *gateway* nodes; other nodes are *nongateway* nodes.

For instance, Figure 2 shows a graph that uses the node identifier for priority. For node 1, *S* = {6}, so node 6 fully covers node 1; consequently, node 1 is a nongateway node. For node 6, *S* = {8}, but node 1 is not a neighbor of node 8; therefore, node 6 is a gateway node. For node 7, *S* = {8, 9, 10} and its additional neighbors, nodes 3, 4, and 5, are neighbors of nodes 10, 9, and 9, respectively; therefore, node 7 is covered by its set *S* and is a nongateway node. For node 10, the corresponding set *S* is empty; therefore, 10 is a gateway node because its neighbors have no node in *S* to satisfy the second condition.

Note that this definition allows each node to decide about its dominating node status without requiring a message exchange. The knowledge of either its two-hop neighbors or its one-hop neighbors with their geographic positions is sufficient. Each node can decide whether or not it is a gateway node by running the following procedure:

- collect information about neighborhood and neighbor priorities;
- compute subgraph of one-hop neighbors with higher priority;
- if this subgraph is connected and if each one-hop neighbor is either in this subgraph or the neighbor of at least one node in this subgraph, the node chooses nongateway status; otherwise, the node chooses gateway status.

Dai and Wu's original algorithm[5] defined priority by node identifiers, leaving the energy remaining in nongateway nodes available for extending network life.

A variation of the dominating-set protocol uses timeouts to transmit selected priorities to a vicinity. At the beginning of the process, each node selects a timeout that is inversely proportional to the node's priority (the timeout function may also depend on a random variable). This means that a node with high priority selects a short timeout and vice versa. At the end of this waiting period, the node can indicate its priority by broadcasting its identifier to its neighbors. The node can also transmit its gateway decision at the same time. Thus, at the end of the waiting period, each node knows which of its neighbors have higher priority and are gateway nodes.

Area-dominating sets

We can use this information to modify the dominating-set protocol to find area coverage rather than node coverage. In this modification, each node computes its timeout function based on its priority and listens to messages from other nodes before deciding its dominating status at the end of a timeout interval. A node choosing gateway status always transmits a message (*positive advertising*) to all its neighbors. A node choosing not to monitor its area has the option of transmitting this information to its neighbors (*negative advertising*) or not.

The protocol runs as follows:

- using a simple perimeter coverage scheme,[3] node *A* computes the area covered by each node that transmits either positive or negative advertising and includes the transmitting node in a subset;
- at the end of its timeout interval, node *A* computes a subgraph of its one-hop neighbors that sent advertisements (these are its neighbors with higher priority);
- if this subgraph is connected and if the subgraph nodes fully cover node A's area, node *A* opts for nongateway (sleeping) status; otherwise, the node chooses gateway (active) status.

Figure 3 shows the possible decision results at a central node *A* based on different information available through positive-only advertising and through positive and negative advertising.

The distributed dominant-pruning algorithm can prove that this area-dominating set is connected. In the case of positive-only advertising, each node simply ignores nodes that remain silent and inactive. In the positive-and-negative-advertising case, a node gives higher priority to all nodes with pre-

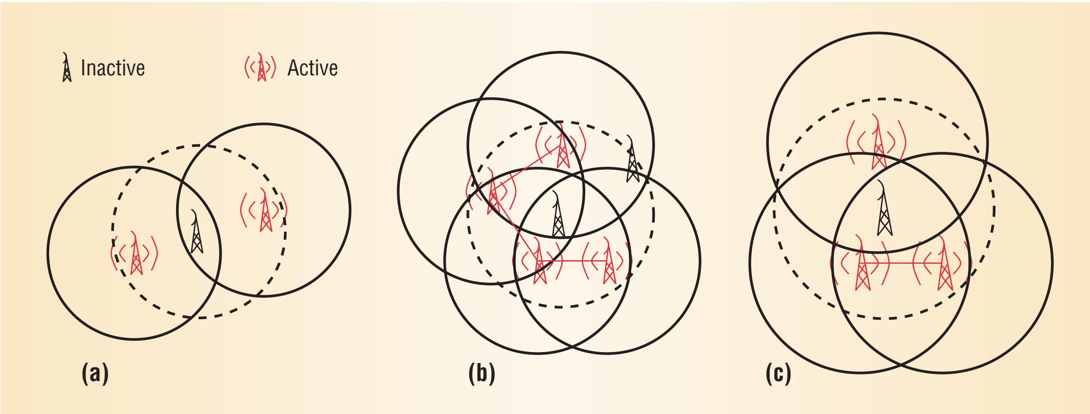

Figure 3. Examples of configurations in which the central node, which is in its default inactive state, makes its area-coverage decision. (a) Central node decides to be active because its active neighbors do not fully cover its monitoring area; (b) central node decides to remain inactive because its monitoring area is covered by active neighbors that are connected; and (c) central node decides to be active because the active neighbors that cover its monitoring area are not connected.

vious advertisements and treats them as part of a subset that must be connected. In both cases, the proof is similar to the connectivity proof given for the dominant-pruning algorithm. It suffices here to set the priority in Dai and Wu's algorithm to the remaining battery life.

Let *G* be the area-dominating set that the algorithm generates, and let *F* be the dominating set of *G* that the dominant-pruning protocol generates. If the monitoring area of a node *u* is covered by a connected subset of its neighbors with more battery life—that is, shorter timeout and therefore higher priority—the set of neighbors of *u* is "fully covered" by this same subset. In other words, if a node does not belong to *G,* it also does not belong to *F*. Hence, the area-dominating set *G* includes the constructed dominating set (of neighbors) *F*. Since *F* is proven to be connected, it follows that *G* is also connected.

This property of a connected area-dominating set is important for request propagation and collection of sensor replies.

Figure 4 shows how negative advertising can reduce the number of active sensors and therefore prolong network life. Numbers represent timeouts. Nodes 0, 1, 2, 3, and 4 announce their activation. Node 5 decides to be inactive since previously advertised nodes are connected and cover its monitoring area. If node 5 does not announce its deactivation, node 6 decides to be active because it does not know that the area *A* is covered.

If node 5 announces its status, node 6 decides to be inactive because the negative advertising also brings information about coverage—that other nodes with shorter timeouts cover node 5's sensing area.

The network can reselect covering nodes periodically to spread the sensing cost dynamically over all nodes in a fair manner. This method significantly extends the network's life. If the density is more than 30 nodes per unit area, the area-dominating graph is sparse, with nodes having on average three neighbors. In addition, the distance between its two neighboring nodes is typically two-thirds of the transmission radius. Hence, active nodes form a very simple network with a structure similar to regular hexagonal tiling.

REQUEST SPREADING

Wireless ad hoc networks commonly use broadcasting to find routes and to disseminate requests and data, and many research efforts have addressed the design of energy-efficient broadcast protocols.[4] The protocols differ for nodes with fixed transmission ranges and those that can adjust their transmission ranges.

Fixed transmission ranges

When all nodes have a fixed transmission radius, the basic broadcast protocol is *blind flooding*: A sequence number identifies each broadcast message, and a node receiving the message for the first time retransmits it to its neighbors. In general, blind flooding generates redundant transmissions and leads to many packet collisions in the media access control (MAC) network layer. But in the sparse and uniformly distributed networks of area-dominant nodes that we have defined for area coverage, blind flooding produces a satisfactory solution.

Researchers have proposed several protocols to minimize retransmissions while trying to guarantee that every node receives a broadcast message.[4] Dai and Wu's dominant-pruning algorithm[5] is one.

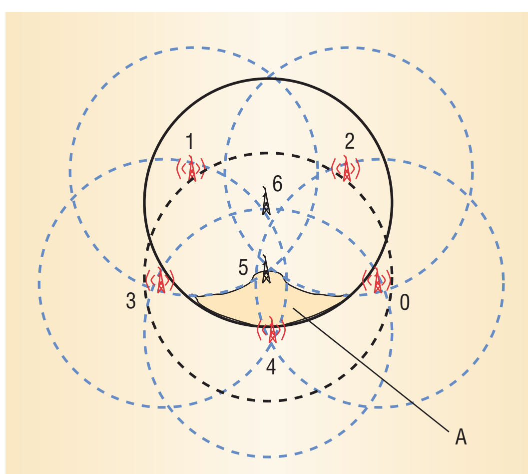

Figure 4. Negative advertising. Node 6 can exhibit different behaviors depending on whether or not node 5 advertises its decision to be inactive.

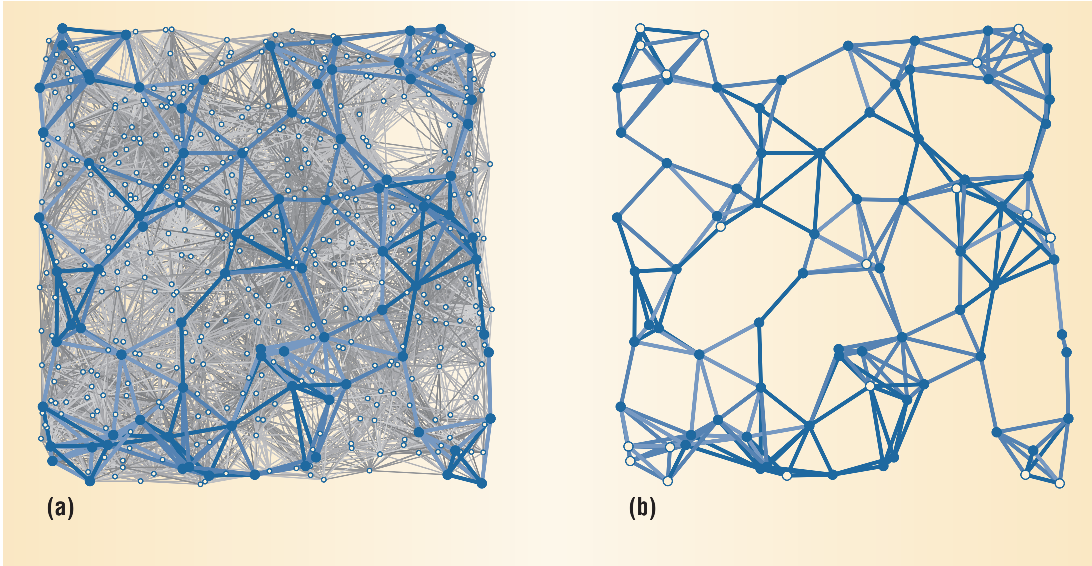

Figure 5. Sensor networks with dominating set applied over area-dominating set. (a) Original network with area-dominating set (black nodes), and (b) dominating set of neighbors applied over area-dominating set.

Applying it to the area-dominating set reduces the number of retransmissions with respect to blind flooding on the order of 20 percent, with most of the savings coming from sensors along the border of the monitored area.

The dominant-pruning method is easy to apply since the dominating-set information is already available from constructing the area-dominating set. Figure 5 illustrates a sensor network with its area-dominating set and its dominating set used for broadcasting.

Forwarding-neighbor protocols also can minimize retransmission requirements. In this approach, each network node has a relay subset composed of neighbor nodes. When a node transmits a broadcast packet, only nodes in its relay subset will consider forwarding the message. The multipoint relay protocol[6] is a deterministic method for reliable broadcasting in this context. The MPR algorithm selects a minimal set of one-hop neighbors that cover the same network as the complete set of neighbors. MPR is a greedy algorithm because computing the minimal set is an NP-complete problem. MPR is also an explicit broadcast protocol since each node finds its relay set by repeatedly adding the (one-hop) neighboring nodes until the relay subset constitutes the maximum possible neighbors not yet covered.

Broadcasting is source dependent and may exclude nodes that received the same message as the relay node from consideration for relay node decisions. The list of relay nodes is attached to the retransmitted packet. When applied on the area-dominating sets, MPR constructs relay subsets that contain nearly all nodes. Thus there is no advantage to applying MPR to area-dominating sets.

A proposed variant of MPR dedicated to sensor networks replaces the number of noncovered neighbors measured in the greedy algorithm.[7] Instead, it uses a "utility function" that multiplies the number of noncovered neighbors with a function of the remaining battery power.

Adjustable transmission ranges

Fixed-transmission-range protocols aim to save energy by reducing the number of sensors that participate in broadcasting. Adjustable-transmission-range protocols measure energy savings differently. The energy consumption for a unit message sent distance r is measured as $r^{\alpha} + c$, where α is a signal attenuation constant greater than 2 and c is a positive constant that represents c signal processing, minimum energy needed for successful reception, and c MAC-layer control messages.[8]

An energy-efficient solution using this method may require more nodes to reach every node than a fixed-transmission solution requires, but each node may expend less energy. Thus, control of the emitted transmission power can significantly reduce energy consumption and so increase the network's lifetime. However, adjusting the transmission signal strength generally implies topology alterations that lose connectivity. Hence, nodes must manage their transmission area while maintaining network connectivity.

To preserve connectivity and obtain full coverage of awake sensors, the main topology-control challenge is to design localized algorithms for deciding which edges are necessary for global connectivity. For instance, a relative neighborhood graph (RNG) is a locally defined subgraph that removes an edge between two nodes u and v if a node w is closer to u and v than the distance between u and v.

As Figure 6 shows, this RNG subgraph contains the original graph's minimum spanning tree, which is defined in a globalized manner. Ning Li, Jennifer C. Hou, and Lui Sha[9] recently proposed a local version of MST that preserves connectivity. To define the local MST, each node computes the MST over its one-hop neighborhood and retains only the

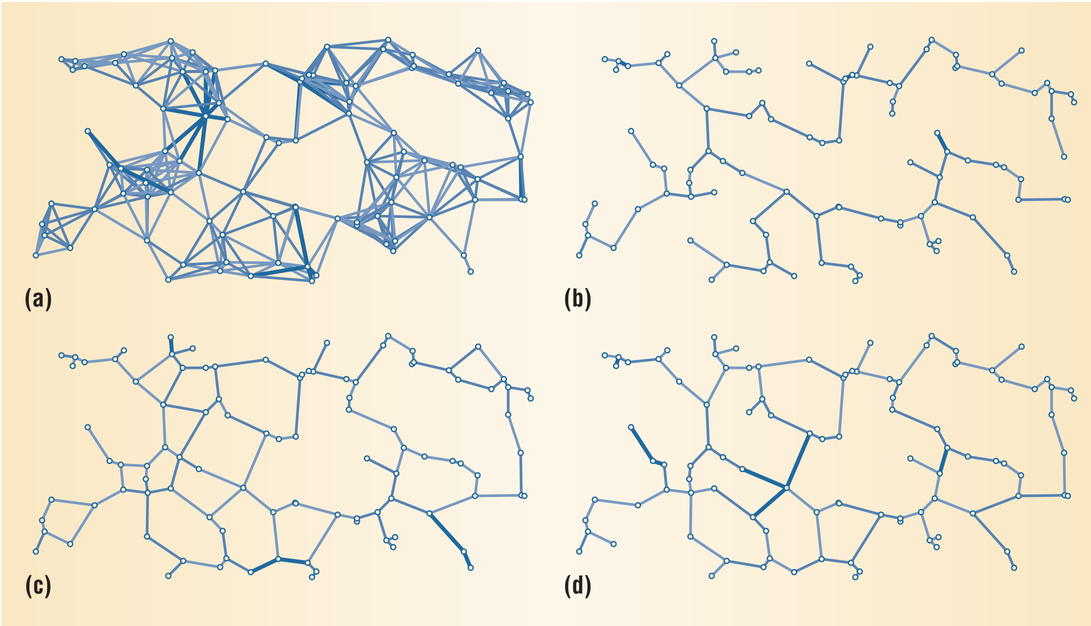

Figure 6. Proximity graphs. (a) Unit graph, (b) minimum spanning tree, (c) relative neighborhood graph, and (d) local MST (100 nodes with average degree of 14).

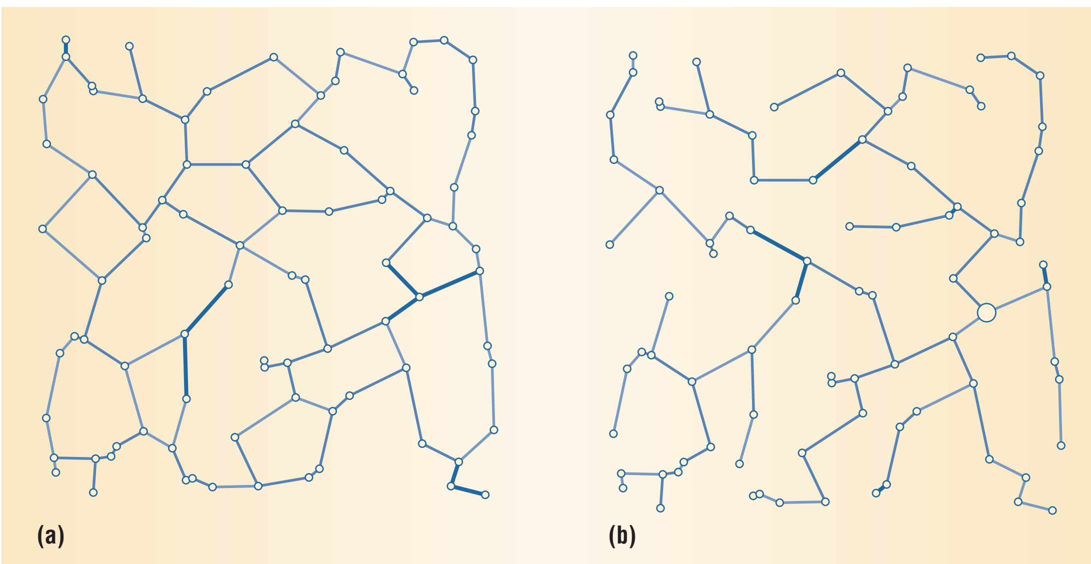

Figure 7. LMST algorithm. (a) Applying the LMST algorithm over an area-dominating set illustrates the simplicity of the obtained graph: Nodes only need to take care of two or three neighbors. (b) Spanning tree induced by flooding over an area-dominating set starting from the monitoring station. A node that receives the broadcast message for the first time considers the sender of the message as its parent in the distributed spanning tree.

neighbors in this subgraph. Thus, only edges that two endpoints choose remain in the graph.

MST is a subset of LMST—an edge that belongs to MST also belongs to the LMST graph. LMST is a subset of RNG.[10] Hence, LMST is the best-known local approximation of globalized MST. Since the MST subgraph is connected if the original graph is connected, LMST and RNG subgraphs that contain MST are also connected. Hence, a topology-control algorithm that preserves LMST edges guarantees network connectivity.

Other LIFL research has proposed adaptive protocols that use RNG or LMST subgraphs to ensure connectivity conservation.[10,11] In this work, each node's transmission radius is chosen to reach nonattained RNG or LMST neighbors. Although minimizing the transmission range is not always optimal (except for $c = 0$), it is possible to determine an optimal radius for power-efficient broadcasting either experimentally or theoretically.

In broadcasting requests over area-dominating sets, the regularity of the created topology limits transmission power choices. Indeed, as Figure 7a shows, the degree of an area-dominant node is about 3 and the number of LMST neighbors is around 2. When an area-dominant node receives a request message, usually only one LMST neighbor has not yet received it. This reduces the transmission radius by about 30 percent since the average distance between two area-dominant nodes is about two-thirds of maximal radius. For an energy model with $\alpha = 2$ and $c = 0$, the energy savings for a single broadcast is about 65 percent, given the reductions in communication power requirements and LMST leafs that must resend the request.

DATA AGGREGATION

After a sensor node receives a request, it must respond by reporting its measurements. Aggregating sensor measurements to report only important information, such as average or extreme values, can further reduce energy consumption. For instance, a surveillance application can request that the sensors count the number of sites that observe a temperature greater than a given threshold.

One technique for limiting the number of reply messages uses the spanning tree induced by flooding during request spreading.[7] Figure 7b shows this spanning tree for the same initial graph shown in Figure 5. Following construction of the tree, the parent nodes transmit data coming from multiple sensor nodes to the monitoring station via their own parent. For particular requests, such as inventory, a node can wait to have multiple replies from all its successors in the tree before replying with data fusion—sending the sum of its successors' replies instead of retransmitting all replies. This solves implosion and overlap problems.

In future work, we plan to study sensor networks in which the sensing and transmission radii are different. In fact, the existence of an optimal transmission radius in the request-spreading process suggests an advantage in having a transmission radius larger than the sensing radius because the sensing radius directly affects the average distance between area-dominant nodes. Moreover, enlarging the transmission radius can also benefit data-fusion schemes by allowing the construction of better-balanced trees. ■

References

1. I.F. Akyildiz et al., "Wireless Sensor Networks: A Survey," *Computer Networks*, vol. 38, 2002, pp. 393-422.
2. F. Ye et al., "PEAS: A Robust Energy Conserving Protocol for Long-Lived Sensor Networks," *Proc. IEEE Int'l Conf. Network Protocols* (ICNP 2002), IEEE CS Press, 2002, pp. 200-201.
3. D. Tian and N.D. Georganas, "A Coverage-Preserving Node Scheduling Scheme for Large Wireless Sensor Networks," *Proc. 1st ACM Workshop Wireless Sensor Networks and Applications*, ACM Press, 2002; pp. 32-41.
4. I. Stojmenovic and J. Wu, "Broadcasting and Activity-Scheduling in Ad Hoc Networks," to be published in *Ad Hoc Networking*, S. Basagni et al., eds., IEEE Press, 2004.
5. F. Dai and J. Wu, "Distributed Dominant Pruning in Ad Hoc Networks," *Proc. IEEE 2003 Int'l Conf. Communications* (ICC 2003), IEEE Press, 2003, pp. 353-357.
6. A. Qayyum, L. Viennot, and A. Laouiti, "Multipoint Relaying for Flooding Broadcast Messages in Mobile Wireless Networks," *Proc. 35th Ann. Hawaii Int'l Conf. System Sciences* (HICSS-35), IEEE CS Press, 2002, pp. 298-307.
7. J. Lipman et al., "Resource Aware Information Collection (RAIC) in Ad Hoc Networks," *Proc. 2nd Mediterranean Ad Hoc Networking Workshop* (MED-HOC-NET 2003), 2003, pp. 161-168.
8. L.M. Feeney, "An Energy-Consumption Model for Performance Analysis of Routing Protocols for Mobile Ad Hoc Networks," *ACM J. Mobile Networks and Applications*, vol. 3, no. 6, 2001, pp. 239-249.
9. N. Li, J.C. Hou, and L. Sha, "Design and Analysis of an MST-Based Topology Control Algorithm," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 1702-1712.
10. J. Cartigny et al., "Localized LMST and RNG-Based Minimum-Energy Broadcast Protocols in Ad-Hoc Networks," to be published in *Ad Hoc Networks*, 2004.
11. J. Cartigny, D. Simplot, and I. Stojmenovic, "Localized Minimum-Energy Broadcasting in Ad-Hoc Networks," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2210-2217.

Jean Carle *is an associate professor in the Fundamental Computer Science Laboratory of Lille (LIFL) at the University of Lille, France. His research interests include distributed computing, data communication, mobile ad hoc networks, and sensor networks. Carle received a PhD in computer science from the University of Amiens. Contact him at carle@lifl.fr.*

David Simplot-Ryl *is an associate professor in LIFL at the University of Lille. His research interests include ad hoc networks, distributed computing, embedded operating systems, and RFID technologies. Simplot-Ryl received a PhD in computer science from the University of Lille. Contact him at simplot@lifl.fr.*

COOL Chips VII

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. COOL Chips VII is to be held in Yokohama on April 14-16, 2004, 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.

Keynote Presentations

The following two keynote presentations will be provided;

- "Energy Optimization in Multi-Core Systems", John Cornish (ARM Limited)
- "Cool Chips Enable Ubiquitous Web Browsing", Daniel F. Zucker (ACCESS Co., Ltd.)

(Speakers and titles are as of December 19, 2003)

Symposium Registration

In order to make a registration, please use COOL Chips VII REGISTRATION FORM on web site at < http://www.coolchips.org >.

DATE AND LOCATION

The COOL Chips VII will be held at Yokohama Joho Bunka Center in Yokohama on April 14-16, 2004.

REGISTRATION FEES

Registration Fees include all plenary and technical sessions and special invited lectures presented 4/14/04 through 4/16/04, a copy of the proceedings (copies of speakers' slides).

	Early Registration by March 19, '04	Late Registration from March 20, '04
Member of any of IEICE, IEEE, IPSJ or ACM	30,000 yen	36,000 yen
Non-Member*	38,000 yen	46,000 yen
Student	10,000 yen	10,000 yen

PAYMENT AND REMITTANCE

Fees should be paid with a credit card or a bank transfer. Credit Cards (VISA, MASTERCARD and AMERICAN EXPRESS) are acceptable. Please indicate the account number of your credit card and expiration date on the Registration Form. Especially from inside Japan, bank transfer is also acceptable.

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.

CONFIRMATION AND RECEIPT

Upon receiving your Registration Form, the COOL Chips VII Secretariat will send your registration number for your confirmation. This confirmation should be retained and presented at the registration desk of the symposium site. Upon confirming your payment, the receipt of the registration fee would be issued at the symposium site.

CANCELLATION AND REFUND POLICIES

All registration transactions for the symposium are being handled by the secretariat of COOL Chips VII. Cancellation is available by March 19,2004 with 50% of registration fee refunded. We regret that no refunds for registration could be made from March 20, 2004. Please note that all refunds would be proceeded after the symposium.

CORRESPONDENCE

Please feel free to contact the following address, if you have any questions.

C/O: The IEICE (The Institute of the Electronics, Information and Communication Engineers)
Attn.: Mr. Y.Uchida, Mr. Y.Motoyama
Kikai-Shinko-Kaikan bldg., 202, 3-5-8,
Shibakoen, Minato-ku, Tokyo 105-0011, Japan
FAX: + 81-3-3433-6659,
E-mail:cool@coolchips.org
http://www.coolchips.org

Accommodations

HOW TO GET TO Yokohama Joho Bunka Center (Yokohama Media & Communication Center) in Yokohama.
Please visit <**http://www.idec.or.jp/jyobun/**>.

HOTEL RESERVATIONS

A special group rate is available for COOL Chips VII attendees for
THE YOKOHAMA NOVOTEL and
HOTEL NEW GRAND YOKOHAMA.
For more information, visit
<**http://www.coolchips.org/**>.

Cross-Layering in Mobile Ad Hoc Network Design

To overcome network performance problems, the MobileMan cross-layer design lets protocols that belong to different layers cooperate in sharing network-status information while still maintaining the layers' separation at the design level.

Marco Conti
Gaia Maselli
Giovanni Turi
IIT-CNR

Silvia Giordano
SUPSI

Mobile ad hoc network (manet) researchers face a major challenge: achieving full functionality with good performance while linking the new technology to the rest of the Internet. The IETF Manet Working Group addresses the latter issue by proposing a view that manets are an evolution of the Internet. The worldwide success of the Internet, mainly determined by a layered architecture, has promoted the adoption of a similar solution for manets. However, a strict layered design is not flexible enough to cope with the dynamics of manet environments, and will thus prevent performance optimizations.[1]

To what extent, then, must developers modify the pure layered approach by introducing stricter cooperation among protocols belonging to different layers? At one end, some solutions based on layer triggers use strict layering to maintain compatibility and solve interdependencies between protocols. A full cross-layer design represents the other extreme.

TRIGGERS VERSUS CROSS-LAYERING

Layer triggers—predefined signals to notify events such as data delivery failures between protocols—have been used extensively in both wired and wireless networks. Examples include

- the Explicit Congestion Notification mechanism, which intermediate routers use to notify the transmission control protocol layer about congestion; and
- L2 triggers, added between the link and Internet protocol layer to efficiently detect changes in the wireless links' status.

A full cross-layer design, on the other hand, introduces stackwide layer interdependencies to optimize overall network performance. In cross-layering, protocols use the state information flowing throughout the stack to adapt their behavior accordingly. For example, given current channel and network conditions, the physical layer can adapt rate, power, and coding to meet application requirements. Although the literature shows the advantages of this approach, previous work focused only on specific problems—such as data accessibility—and looked at the joint design of two to three layers only, such as the physical, media access control (MAC), and routing layers.[2]

In the ongoing cross-layer versus legacy-layer architecture debate, the ad hoc research community recognizes that cross-layering can provide significant performance benefits, but also observes that a layered design provides a key element in the Internet's success and proliferation.[3] Strict layering guarantees controlled interaction among layers because developing and maintaining single layers takes place independently of the rest of the stack.

On the other hand, an unbridled cross-layer design can produce spaghetti-like code that is impossible to maintain efficiently because every modification must be propagated across all protocols. Further, cross-layer designs can produce unintended interactions among protocols, such as adaptation loops, that result in performance degradation.

Published by the IEEE Computer Society

MOBILEMAN

We believe that developers must adopt a careful cross-layer design to overcome potential manet performance problems. Our approach introduces inside the layered architecture the possibility that protocols belonging to different layers can cooperate by sharing network-status information while still maintaining separation between the layers in protocol design.

The MobileMan project's primary aim is to exploit a manet protocol stack's full cross-layer design (http://cnd.iit.cnr.it/mobileMAN). We are not aware of an existing reference architecture that accomplishes this goal. MobileMan still implements protocols inside each layer, but offers the advantages of

- allowing for full compatibility with standards, as it does not modify each layer's core functions;
- providing a robust upgrade environment, which allows the addition or removal of protocols belonging to different layers from the stack without modifying the operations at the other layers; and
- maintaining the benefits of a modular architecture.

This reference architecture exploits the advantages of a full cross-layer design while still satisfying the layer-separation principle.

Figure 1 shows that some network functions, such as energy management, security, and cooperation, are cross-layer by nature. MobileMan seeks to extend cross-layering to all network functions through data sharing. The architecture presents a core component, *Network Status*, that functions as a repository for information that network protocols throughout the stack collect. Each protocol can access the Network Status to share its data with other protocols. This avoids duplicating efforts to collect internal state information and leads to a more efficient system design.

MobileMan achieves layer separation by standardizing access to the Network Status. This implies defining the way protocols can read and write the data from it. Interactions between protocols and the Network Status are placed beside normal-layer behavior, allowing optimization without compromising the expected normal functioning. Replacing a network-status-oriented protocol with its legacy counterpart will therefore let the whole stack keep working properly, although at the cost of penalizing functional optimizations.

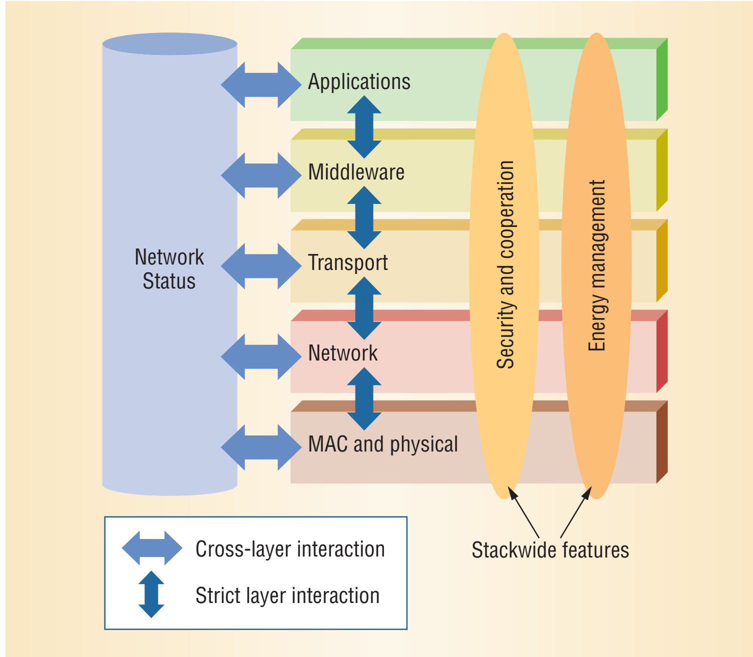

Figure 1. MobileMan architecture. Some network functions, such as energy management and security and cooperation, are cross-layer by nature. MobileMan seeks to extend cross-layering to all network functions through data sharing.

For example, using the legacy TCP implies that cross-layer optimizations will not occur at this layer and that the transport protocol will not provide any information to the Network Status component. However, the overall protocol stack will still operate correctly, although with degraded performance.

Performance gains

We believe that the MobileMan reference architecture offers the following performance advantages in ad hoc network design:

- *Cross-layer optimization for all network functions.* Cross-layering is a must for functions such as energy management, but provides benefits for all network functions.
- *Improved local and global adaptation.* Developers can use MobileMan to adapt the system to highly variable ad hoc network conditions and to better control system performance. For example, developers can exploit a cross-layering design to perform both local and global adaptations to network congestion. Specifically, the MAC layer reacts locally to congestion by exponential back-off. When congestion is high, this response is insufficient, requiring dual-option compensation: Either the forwarding mechanism can reroute traffic to avoid the bottleneck or, if alternate routes do not exist, the optimization can use transport protocol mechanisms to freeze traffic transmissions.
- *Full context awareness at all layers.* Developers can design protocols to be aware of the Network Status, energy level, and other factors. Cross-layering facilitates achieving context awareness at higher layers—such as

The MobileMan project is researching a physical implementation of an enhanced IEEE 802.11 wireless technology.

middleware and application layers.

- *Reduced overhead.* Collecting Network Status information avoids data duplication at different layers.

Protocol redesign

The only way to gain these benefits, however, is to redesign protocols. To fully exploit cross-layering and measure its impact on ad hoc network performance, developers are currently redesigning the full protocol stack even when they can still integrate legacy protocols to form a mixed architecture. MobileMan introduces the following modifications at each layer.

Wi-Fi. The MobileMan project uses IEEE 802.11 as its reference technology. The project is conducting ongoing research into a physical implementation of an enhanced IEEE 802.11 wireless technology[4] to fix performance problems. Specifically, by exploiting interactions between Wi-Fi and the network layer, through data sharing, MobileMan enhances the 802.11 MAC as follows:

- *Enhanced back-off.* Research shows that the binary exponential back-off scheme performs inadequately in ad hoc networks by causing channel underutilization and unfair sharing.[5] To fix these problems, the MobileMan project is working on an enhanced card, which implements the asymptotically optimal back-off algorithm, extending the standard protocol to achieve a theoretical optimum performance.[6] In addition, MobileMan can fix problems stemming from exposed and hidden station phenomena by exploiting Network Status information collected at the MAC and network layers. Taking into consideration that a large physical-carrier sensing range[7] causes nodes within two to three hops to conflict with each other when accessing the shared channel, the MAC layer exploits the topology information that the routing protocol provides to achieve fair channel scheduling.
- *Routing performed at the MAC layer.* By exploiting at the MAC layer the topology information collected by the routing protocol layer, we are developing an efficient packet-forwarding scheme inside the MAC card.[2]

Routing and forwarding. MobileMan considers routing according to the cross-layering principle, so that other layers also can use routing data. A proactive protocol meets this requirement because it collects topology information even when it is not required to perform packet forwarding. The protocol uses this apparently unnecessary information to facilitate other layers' tasks. Indeed, the MobileMan project is investigating the use of a link-state protocol[8] in which a node propagates link-state information to other nodes in the network, limiting flooding of updates in space and time to reach scalability.

The primary consequence of such a link-state routing protocol is a hazy, node-centered knowledge of network topology. A node has a precise knowledge of the neighborhood that nodes two to three hops away from it form. As the distance from this node increases, the neighborhood data becomes less precise. This knowledge can be used to implement multipath reliable forwarding mechanisms[9] to deliver data over existing paths, provided by routing, according to performance and reliability criteria that a performability index summarizes.

A performability estimator classifies the reliability, performance, and cooperation along used routes and computes the performability index for each route. This index summarizes path behavior, taking into account factors such as congestion, link quality, and selfish nodes—all of which can influence system performance. Every time a node sends a packet to a destination, the protocol updates the performability index for the relative route, according to the delivery outcome. The outcome can be inferred by looking at transport ACKs notified through the Network Status. The forwarding function uses the performability index to select among alternative paths to achieve traffic load balancing.

Transport protocol. This protocol seeks mainly to provide the upper layers with a reliable and connection-oriented service. It minimizes useless data retransmissions by analyzing and reacting appropriately to different events occurring at lower layers, such as route failures, route changes, and congestions.[10] The efficient implementation of a reliable transport protocol in ad hoc networks requires strict cooperation with lower layers.[11] Therefore, the MobileMan transport protocol exploits information reported by the routing and Wi-Fi layers in the Network Status component.

Middleware. The middleware layer generally provides abstractions that hide complex details from application programmers. In a manet environment, this trend must be reversed to provide context awareness.[12] In MobileMan, the Network Status contains the network context, which provides context awareness in the cross-layer architecture. For example, we are currently investigating how ad hoc networks could efficiently implement peer-to-peer routing substrates to exploit the information that the Network Status exports. By directly using the topology information the routing protocol collects, a peer-to-peer substrate can locally identify which peers are participating in a specified service, reducing the overhead of building overlay structures.

The MobileMan cross-layer architecture promotes stricter local interaction among protocols in a manet node. The Network Status uniformly manages the cross-layer interaction, and respects the principle of dividing functionalities and responsibilities in layers. This approach aims to optimize overall network performance by increasing local interaction among protocols, decreasing remote communications, and consequently saving network bandwidth. Engineering the Network Status component presents the greatest challenge. This component should be general enough to represent a vertical layer whose changes do not affect the overall system. ■

Acknowledgment

This work was partially funded by the Information Society Technologies program of the European Commission, Future and Emerging Technologies under the IST-2001-38113 MobileMan project.

References

1. A.J. Goldsmith and S.B. Wicker, "Design Challenges for Energy-Constrained Ad Hoc Wireless Networks," *IEEE Wireless Comm.*, vol. 9, no. 4, 2002, pp. 8-27.
2. MobileMan project, "Architecture, Protocols and Services," Deliverable D5; http://cnd.iit.cnr.it/mobileMAN.
3. V. Kawadia and P.R. Kumar, "A Cautionary Perspective on Cross Layer Design," submitted for publication in *IEEE Wireless Comm.*, 2004.
4. R. Bernasconi et al., "An Enhanced MAC Architecture for Multi-Hop Wireless Networks," *Proc. Personal Wireless Comm. Conf. 2003* (PWC 2003), LNCS 2775, Springer, pp. 811-816.
5. X. Shugong and T. Saadawi, "Does the IEEE 802.11 MAC Protocol Work Well in Multihop Wireless Ad Hoc Networks?" *IEEE Comm.*, June 2001, pp. 130-137.
6. L. Bononi, M. Conti, and E. Gregori, "Run-Time Optimization of IEEE 802.11 Wireless LANs Performance," *IEEE Trans. Parallel and Distributed Systems*, Dec. 2003, pp. 66-80.
7. S. Basagni et al., *Mobile Ad Hoc Networking*, IEEE Press and John Wiley & Sons, 2003.
8. C.A. Santivanez et al., "On the Scalability of Ad Hoc Routing Protocols," *Proc. Infocom 2002*, IEEE Press, vol. 3, 2002, pp. 1688-1697.
9. M. Conti, E. Gregori, G. Maselli, "Towards Reliable Forwarding for Ad Hoc Networks," *Proc. Personal Wireless Comm. Conf. 2003* (PWC 2003), LNCS 2775, Springer, 2003, pp. 790-804.
10. I. Chlamtac, M. Conti, and J. Liu, "Mobile Ad Hoc Networking: Imperatives and Challenges," *Ad Hoc Networks*, vol. 1, no. 1, 2003, pp. 13-64.
11. G. Anastasi and A. Passarella, "Towards a Novel Transport Protocol for Ad Hoc Networks," *Proc. Personal Wireless Comm. Conf. 2003* (PWC 2003), LNCS 2775, Springer, 2003, pp. 805-810.
12. C. Mascolo, L. Capra, and W. Emmerich, "Middleware for Mobile Computing (A Survey)," *Advanced Lectures on Networking*, LNCS 2497, E. Gregori, G. Anastasi, and S. Basagni, eds., Springer, 2002, pp. 20-58.

Marco Conti *is a senior researcher at the Istituto di Informatica e Telematica, Consiglio Nazionale delle Ricerche (IIT-CNR). His research interests include mobile computing and wireless networks. Conti received a Laurea degree in computer science from the University of Pisa, Italy. Contact him at marco.conti@iit.cnr.it.*

Gaia Maselli *is a PhD candidate at IIT-CNR. Her research interests include architectural and networking aspects of mobile ad hoc networks. Maselli received an MS in computer science from the University of Pisa, Italy. Contact her at gaia.maselli@iit.cnr.it.*

Giovanni Turi *is a PhD candidate at IIT-CNR. His research interests include architectural and middleware aspects of mobile computing. Turi received an MS in computer science from the University of Pisa, Italy. Contact him at giovanni.turi@iit.cnr.it.*

Silvia Giordano *is a senior researcher at the Department of Technology and Information, Scuola Universitaria Professionale della Svizzera Italiana. Her research interests include wireless networking and traffic control. Giordano received a PhD in engineering science from EPFL, Switzerland. Contact her at silvia.giordano@supsi.ch.*

Group Communications in Mobile Ad Hoc Networks

This survey of approaches to group communications in mobile ad hoc networks explores several potential solutions to the unique problems of wireless mobile communications, which have variable and unpredictable characteristics due to mobility and signal strength fluctuations with respect to time and environment.

Prasant Mohapatra
Chao Gui
Jian Li
University of California, Davis

A mobile ad hoc network (manet) comprises a set of wireless devices that can move around freely and cooperate in relaying packets on behalf of one another. A manet does not require a fixed infrastructure or centralized administration. Because mobile nodes have limited transmission range, distant nodes communicate through multihop paths.

Their ease of deployment makes manets an attractive choice for a variety of applications. Examples include battleground communications, disaster recovery efforts, communication among a group of islands or ships, conferencing without the support of a wired infrastructure, and interactive information sharing. Unlike typical Internet applications, most applications of manets involve one-to-many and many-to-many communication patterns.

Efficient support of group communications is critical for most ad hoc network applications. However, manet group communications issues differ from those in wired environments for the following reasons: The wireless communications medium has variable and unpredictable characteristics and the signal strength and propagation fluctuate with respect to time and environment. Further, node mobility creates a continuously changing communication topology in which routing paths break and new ones form dynamically.

Because manets have limited bandwidth availability and battery power, their algorithms and protocols must conserve both bandwidth and energy. Wireless devices usually use computing components—processors, memory, and I/O devices—that have low capacity and limited processing power. Thus, their communications protocols should have lightweight computational and information storage needs.

MULTICASTING

The multicasting communications model can facilitate effective and collaborative communication among groups. *Flooding* and *tree-based routing* represent two ends of the multicasting spectrum. Flooding is a simple approach that offers the lowest control overheads at the expense of generating very high data traffic in the wireless environment. The tree-based approach, on the other hand, generates minimal data traffic in the network, but tree maintenance and updates require many control-traffic exchanges. Both flooding and tree-based approaches scale poorly.

Multicast routing protocols for manets vary in terms of route topology, state maintenance, reliance on unicast routing, and other attributes. Instead of using a taxonomic approach to previously proposed multicasting protocols, our approach emphasizes the schemes' salient features.

Most proposed multicasting protocols primarily exploit one or more specific characteristics of the manet environment. These characteristics include variable topology, soft-state and state aggregations, knowledge of location, and communication pattern randomness. For example, mesh-based protocols exploit variable topology; stateless multicasting

Published by the IEEE Computer Society

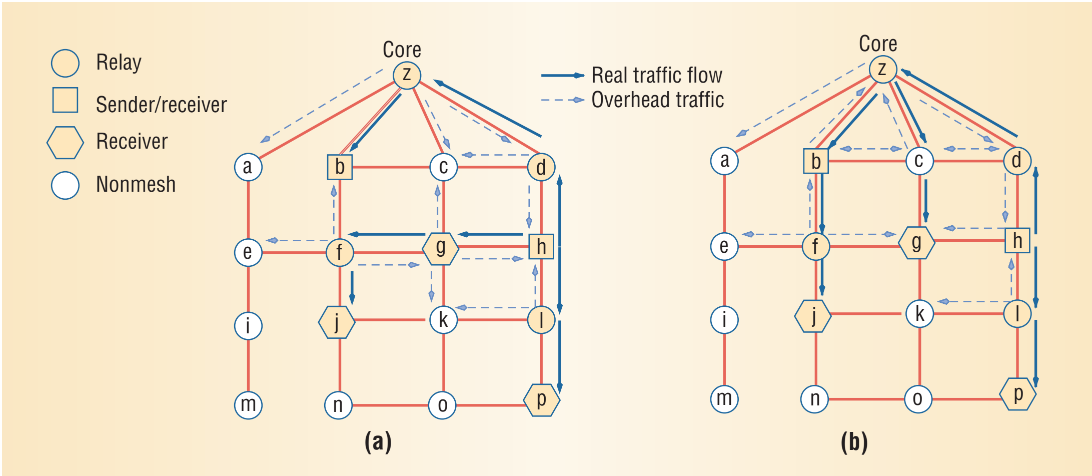

Figure 1. CAMP multicast routing. (a) Data-packet forwarding from node h using the core-assisted mesh protocol. (b) An equivalent shared tree, which uses a receiver-initiated router method. The solid arrows indicate the flow of actual traffic and the dashed arrows indicate the broadcast traffic due to the broadcast nature of wireless links.

exploits soft-state maintenance; location-aided multicasting exploits knowledge of location; and gossip-based multicasting exploits randomness in communication and mobility.

Mesh-based protocols

The addition of redundant paths between on-tree nodes converts a multicast tree into a mesh topology. The availability of alternative paths lets nodes deliver multicast packets regardless of link breakages. Mesh-based protocols thus achieve higher robustness against node mobility.

Core-assisted mesh protocol. CAMP[1] uses a shared mesh structure to support multicast routing in dynamic ad hoc networks. This structure ensures that the mesh includes the *reverse shortest paths,* the shortest paths from all receivers to the source.

Figure 1 shows how the protocol forwards data packets from node *h* to the rest of the group. To prevent packet replication or looping in the mesh, each node maintains a cache to keep track of recently forwarded packets. Periodically, a receiver node reviews its packet cache to determine whether it is receiving data packets from those neighbors not on the reverse shortest path to the source. When such situations arise, the node sends a heartbeat message to its successor in its reverse shortest path to the source. When the successor is not a mesh member, the heartbeat message triggers a *push join* message, which includes all nodes along any reverse shortest path in the mesh.

CAMP uses cores to limit the control traffic needed to create multicast meshes. Unlike the core-based tree protocol, CAMP does not require that all traffic flow through the core nodes. CAMP uses a receiver-initiated method for routers to join a multicast group. If a node wishing to join such a group finds it has neighbors that belong to the group, it simply updates its multicast routing table and uses a standard update procedure to announce its membership. When none of its neighbors are mesh members, the node either sends a join request toward a core or attempts to reach a group member using an expanding-ring search process. Any mesh member can respond to a join request with a join ACK, which propagates back to the request originator.

On-demand multicast routing protocol. Based on a sender-advertised approach to building a mesh, ODMRP[2] uses the *forwarding group* concept, in which a set of nodes forwards multicast data along the shortest paths between any member pairs. In ODMRP, each source establishes and updates a group membership and a multicast mesh *on demand*. By flooding a member-advertising packet, a source node starts building a forwarding mesh for the multicast group, collecting membership information at the same time.

When a node receives a nonduplicate message requesting admission to the multicast group, it stores the upstream node identity and rebroadcasts the packet. When this request message packet reaches a multicast receiver, the receiver creates or updates the source entry in the *member table*. The system then uses the member table to prepare periodic control packets and broadcasts them via the receiver node. The nodes relay the packets back toward the source along the reverse path that the member-advertising packet traverses. This process constructs or updates the routes from sources to receivers and builds a mesh of nodes, called the *forwarding group*. Multicast sources send the member-advertising packet periodically to refresh the membership information and update the routes. A soft-state approach maintains the multicast group and the mesh.

State maintenance

The techniques for maintaining multicast protocol states can be classified as stateless, constrained, or unconstrained. Stateless multicasting protocol nodes do not maintain any state information. Constrained state protocols reduce the state maintenance overhead through abstraction via application-layer multicasting or by aggregation via hierarchical multicasting. In unconstrained state protocols, both group members and nonmembers must maintain the protocol states to support a multicast group.

Broadcasting can provide a building block for route discovery in on-demand ad hoc routing protocols.

For multicasting in manets, a wider spread of protocol states restrict robustness and scalability. Changing network states requires more updates and exchange of control messages. If the routing tree or mesh involves fast-moving nonmember or member nodes, the multicast session will be severely hampered for unconstrained-state protocols. However, zero-state and constrained-state protocols are usually less affected by host mobility.

Location-aided multicasting

In networks that can access the Global Positioning System (GPS), the network provides each node with location and mobility information. Multicast protocols can also use this information to improve protocol robustness and performance. With GPS support, ODMRP can adapt to node movements and can use location and mobility information to estimate route expiration time, while receivers select the path that will remain valid longest. Sources can reconstruct routes in anticipation of route breaks, thereby making the protocol more resilient to node mobility.

Martin Mauve and colleagues[3] proposed a *position-based multicasting* (PBM) technique that does not require flooding to maintain a tree or mesh structure. In PBM, a multicast source node finds a set of neighboring next-hop nodes and assigns each packet destination to one next-hop node. The next-hop nodes, in turn, repeat the process. Thus, no global distribution structure is necessary.

Researchers have proposed two forwarding techniques for PBM. In *greedy multicast forwarding*, the next hop is selected based on the position of the forwarding node, its neighbors, and the destination. The distance toward the destination node is reduced at each hop.

When the greedy forwarding approach fails, the system adopts a recovery process using *perimeter forwarding*, in which it forwards the packet by traversing the network boundary gaps until it can resume greedy forwarding.

Gossip-based multicasting

Some multicasting protocols use gossip as a form of probabilistically controlled flooding to solve several problems, including network news dissemination. The basic idea of applying gossip to multicasting involves having each member node periodically talk to a random subset of other members. After each round of talk, the gossipers can recover their missed multicast packets from each other. In contrast to deterministic approaches, a probabilistic scheme will better survive a highly dynamic ad hoc network because it functions independently of network topology and its nondeterministic property matches the network's characteristics.

Anonymous gossip. Manet designers can apply the anonymous gossip[4] multicast performance enhancement technique atop any tree-based or mesh-based protocol with minimal overhead. This technique does not require a group member to have any knowledge of the other group members.

An anonymous gossip multicast protocol proceeds in two phases. In the first phase, a protocol multicasts packets to the group. In the second phase, periodic anonymous gossip takes place in the background as each group member recovers any lost data packet from other members of the group that might have received it.

Route-driven gossip. The RDG protocol[5] relies on a unicast protocol such as dynamic source routing to provide routing information for guiding the gossip process. Each node maintains two data structures for a multicast group: a *data buffer* that stores received data packets and a *view*, which lists all other group member nodes known to this node. The view at each node consists of two parts: the *active view*, which contains the IDs of known members to which at least one routing path is known, and the *passive view*, which contains the IDs of known members to which no routing path is currently available.

A node seeking to join a group floods the network with a Group-Request message. All members receiving the message update their active view. They also return a Group-Reply to the request initiator with a certain probability. The initiator updates its active view after receiving the Group-Reply message. Each member node periodically generates a gossip message and sends it to a set of other nodes randomly chosen from its active view. The message includes a selected subset of the data buffer and the sequence number of the most recent missing data packets. A group member receiving a gossip message will update its view of other group members and update its data buffer with newly received data.

BROADCASTING

Network-wide broadcasting, which attempts to deliver packets from a source node to all other nodes in the network, serves an important function in manets. Broadcasting often provides a building block for route discovery in on-demand ad hoc routing protocols. When designing broadcast protocols for ad hoc networks, developers seek to

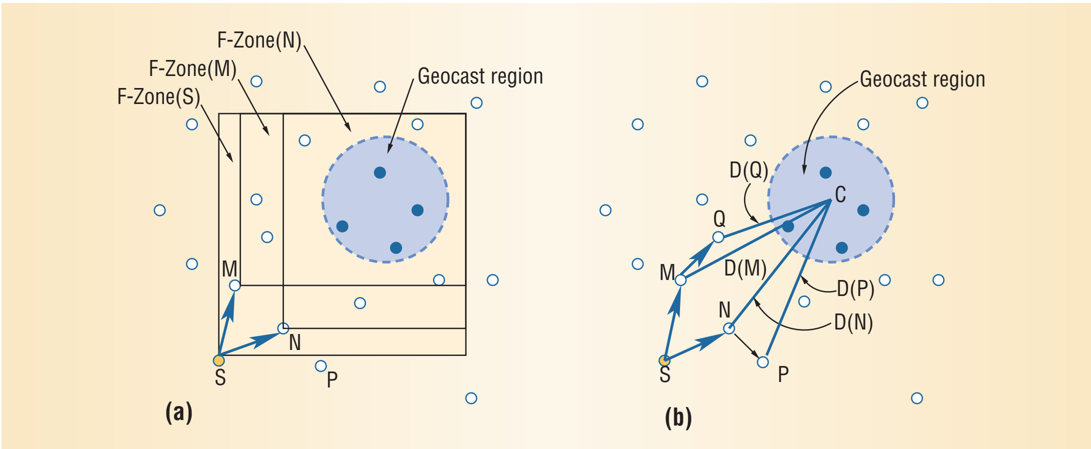

Figure 2. Location-based multicast forwarding schemes. (a) The first LBM scheme defines a rectangular forwarding zone, one corner of which lies at the source node and spans the geocast region. (b) The second LBM scheme uses a distance-based heuristic in which source node S defines the center point, C, of the geocast region in the geocast packets. Each intermediate node decides whether to forward a geocast packet by comparing its distance to the packet sender's distance.

reduce the overhead—such as collision and retransmission or redundant retransmission—while reaching all the network's nodes.

Although a wireless signal broadcast causes more contention and collisions in the shared wireless channel, it also allows a single transmission to reach multiple neighboring nodes. One comparison of existing techniques categorizes manet broadcast protocols into four types: simple flooding, probability-based, area-based, and neighbor-knowledge-based.[6]

In a more recent work,[7] researchers proposed a general framework for self-pruning-based broadcast-redundancy-reduction techniques in ad hoc networks. Upon receiving a packet, intermediate nodes use the two proposed neighborhood coverage conditions to determine whether or not they should rebroadcast it. These coverage conditions depend on neighbor connectivity and the history of visited nodes. Since global network information is costly, the manet can use a distributed and local pruning process to select the forwarding node set based on local information such as the *k*-hop neighbor. Researchers have used this framework to propose new algorithms that combine features of previous work and show better performance.

GEOCASTING AND ANYCASTING

Applications that need to deliver messages of interest to every node in a specific geographical area can adopt *geocasting*, which is either flooding- or route-based. Each node's position with regard to the specified geocast region implicitly defines group membership. Each node is required to know its own physical location, which it can identify using GPS. This does not require any explicit join and leave procedures. The group members tend to be clustered both geographically and topologically.

The IPv6 specification includes *anycast*, a similar Internet-based network service. Several servers, which jointly support a particular service, receive an anycast address. When a host sends its packets to this address, the network delivers the packets to at least one and preferably only one of the servers in the anycast group. Although little work has been proposed for using anycasting in manets, researchers have used it in other applications, especially in battlefield or disaster-recovery communications.

Flooding-based geocasting

Flooding is the simplest way to deliver a message to all nodes in the network. Although expensive and inefficient, a simple flooding algorithm achieves the geocasting goal. Some flooding-based geocast protocols use the *forwarding zone technique* to constrain the flooding scope. A forwarding zone is a geographic area that extends from the source node to cover the geocast region. The source node defines a forwarding zone in the header of a geocast data packet. Upon receiving a geocast packet, other nodes will forward it only if its location is within the forwarding zone.

A geocast protocol's accuracy is defined as the probability that the transmission delivers a geocast packet to each geocast group member. Enlarging the forwarding zone can increase the accuracy. Given that the protocol overhead increases dramatically with an increase in the forwarding zone's size, a geocast protocol must achieve a workable tradeoff between the two factors.

Young-Bae Ko and Nitin H. Vaidya[8] have proposed two flooding-based geocast protocols, both termed *location-based multicast* (LBM) schemes. Figure 2 shows these two LBM forwarding schemes. As Figure 2a shows, the first scheme defines a rectangular forwarding zone. One corner of the zone is at the source node and extends across the full geocast region. An adaptive-forwarding-zone technique ensures that each intermediate forwarding node redefines the forwarding zone by its location relative to the geocast region.

The second scheme uses a distance-based heuristic. As Figure 2b shows, instead of a forwarding zone, the source node S defines the center point, C, of the geocast region in the geocast packets. Each intermediate node decides whether to forward a geocast packet by comparing its distance to that of

In manets, an anycast protocol can simplify the access management and building process of a network distributed service.

the packet's sender. Thus, nodes M, N, and Q will forward the packet. However, node P decides not to forward the packet because the node also receives a geocast packet from node N, whose distance to center point C is shorter.

Route-based geocasting

Route-based geocast protocols use a two-step method for packet delivery. First, the protocol performs an anycast that delivers a geocast packet to any node within the geocast region. Thus, the source node builds a route to one or a few selected nodes in the intended region. Upon receiving a geocast packet, the selected nodes use a localized flooding method to further deliver the packet to all reachable nodes within the geocast region.

The GeoTORA[9] protocol extends the unicast *temporally ordered routing algorithm* for geocast routing. TORA potentially builds multiple routes from any source to a desired destination. The routing procedure assigns a height value to each network node, then uses their heights to determine the logical direction of a link between two nodes, working always from the higher to lower node. Thus, the destination node always has zero height, which the routing procedure uses to derive a destination-oriented *directed acyclic graph* (DAG). Any node that intends to send or forward a packet to the given destination simply follows the logical direction of the adjacent links.

GeoTORA adopts a similar method, building and maintaining a DAG for each geocast group. All nodes that belong to the geocast region have a zero-height link between a pair of nodes. If both end nodes have zero height, the system does not assign a direction. If a node wants to geocast to a region, it forwards the packet to a single node in that region—which, in turn, floods the packet within the region to reach all members.

Anycasting

Anycasting is defined as a point-to-point flow of packets between a single source and the "nearest" destination server identified by an anycast address.

In manets, an anycast protocol can simplify the access management and building process of a network distributed service. The protocol provides a route to the nearest server for a distributed service. It also maintains this route in the face of node mobility and can switch the connection to another server if needed. Rather than designing a completely separate anycasting protocol, extensions of several different classes of unicast routing protocols—such as link state and distance vector—can efficiently construct and maintain anycast routes.

COMMON ISSUES IN GROUP COMMUNICATIONS PROTOCOLS

In addition to performance, some common issues that researchers have considered while designing most group communications protocols include energy conservation, reliability, security, and QoS support. However, the techniques for achieving these goals can differ significantly.

Reliability

Given the dynamic nature of manets, reliable group communications presents a challenging task. When node mobility is high, flooding becomes a viable approach. We assume that mobility is not so high that flooding, or even its more persistent variations, become the only choice for reliable multicasting and broadcasting. Given node mobility and network dynamics, more efficient and flexible alternatives are available for reliable group communications in manets.

A broadcast protocol based on a clustering technique assumes that an underlying clustering protocol manages construction of a clustered architecture that covers the entire population of network hosts.[10] The clustering protocol distributes the broadcast packets to form a forwarding tree consisting of cluster head nodes. This approach achieves reliability at the cost of maintaining the cluster structure proactively, even in the absence of traffic, and by using acknowledgments that travel backward along the path to the source node. Its efficiency also relies on the accuracy of the forwarding tree and underlying cluster structure, which poses a challenge when nodes move swiftly.

Obtaining a hard guarantee of manet reliability becomes extremely difficult when network size and mobility increase. The RDG protocol[5] adopted a practical probabilistic specification that achieves high reliability without relying on any inherent multicast primitive. In RDG, each node has only a random partial view of the group, which results from the randomness of the routing information that any given node has.

RDG uses a pure gossip scheme in the sense that it gossips uniformly about multicast packets, negative acknowledgments, and membership information. A *gossiper push* mainly propels the spread of information, with each group member forwarding a multicast packet to a random subset of the group. This technique is complemented by a *gossiper pull* in which the multicast packets piggyback any negative acknowledgments that the forwarding group members may generate. Given its nondeterministic characteristics, the notion of probabilistic reliability seems quite fitting in the dynamic manet environment.

A few efforts have focused on developing *medium-access-control* support for reliable group communications in manets. A new wireless ad hoc MAC protocol proposes a *broadcast medium window*.[11] BMW strives to ensure the reliable round-robin transmission of each packet to its neighbor. The protocol borrows some concepts from IEEE 802.11 and attempts to achieve reliable broadcast support at the MAC layer when the traffic load is manageable. If reliable transmission is counterproductive, BMW reverts to the unreliable delivery of IEEE 802.11. BMW's round-robin approach does not take advantage of the wireless signal's broadcast nature, so it can incur significant overhead by unicasting packets to each neighbor.

Energy efficiency

Since a limited battery source typically drives nodes in manets, designing energy-conserving protocols becomes essential. Even when energy is not a stringent constraint, reducing power consumption can result in less interference and better throughput.

Researchers can use various techniques to build power-aware and energy-efficient broadcast and multicast infrastructures in manets. Wireless transmission provides the greatest contributor to energy consumption in ad hoc networks, so reducing the number of nodes that participate in transmissions can reduce the total energy for a broadcast and multicast process. Many protocols thus strive to minimize the forwarding node set.

Several proposed techniques for energy-efficient broadcasting and multicasting share a common feature: combining a minimum or reduced forward-node set with power-level selection. The *broadcast incremental power* protocol[12] adds new nodes to the multicast tree one at a time, starting from the source node. BIP decides which specific node to add at each step based on the minimum transmission energy each addition would require. A leaf or parent node with increased transmission power also can reach this new node. A greedy heuristic, BIP requires global network information, but it might not generate the minimum-cost tree. Another proposed localized algorithm requires only neighborhood information and attempts to take advantage of wireless transmission's broadcast nature.[13]

Energy consumption from retransmission at the data-link layer when computing the minimum-energy-cost tree should also be considered when designing protocols. Although many efforts have been made to design energy-efficient broadcast and multicast protocols, issues such as how to address energy efficiency in highly mobile manets and how to factor in traffic conditions when using contention-based MAC protocols still present wide-open challenges.

Since a limited battery source typically drives nodes in manets, designing energy-conserving protocols becomes essential.

The reception and idle-listening process provides another power-consumption source. To reduce power consumption during idle listening, the power-aware PAMAS MAC protocol[14] selectively turns off some network nodes for certain durations. PAMAS has a separate signaling channel for manets. PAMAS overhears exchanges of node request-to-send and clear-to-send messages and uses this information about traffic demand and neighboring node status to determine when a mobile node should sleep, for how long, and what to do if a neighboring destination node is asleep.

The *wakeup mechanism* also plays a critical role in designing power-efficient protocols. Existing wakeup mechanisms fall into three categories:[15]

- *On-demand wakeup* typically uses a wakeup radio to awaken a neighboring node.
- *Scheduled rendezvous* requires that sleeping nodes wake up at the same time periodically to meet the potential demand for intercommunication. This approach is unsuitable for a multihop environment because it requires time synchronization among all nodes.
- *Asynchronous wakeup mechanisms* do not require time synchronization among different nodes. The sleep and wakeup schedules are designed so that any two neighboring nodes will have overlapped active time within a specified number of cycles.

Quality of service

QoS is usually defined as a set of service requirements that the network must meet while transporting a packet stream from a source to its destination. The network is expected to guarantee the performance of a set of measurable prespecified service attributes such as delay, bandwidth, probability of packet loss, and delay variance. Two other QoS attributes, power consumption and service coverage area, are more specific to manets.

With the increase in QoS needs in evolving applications, supporting QoS-aware group communi-

With advances in support for group communications, the use of multimedia objects will proliferate in manet application domains.

cations in manets is also desirable. Resource limitations and variability further add to the need for QoS in such networks.[16] However, the characteristics of these networks make QoS support a complex process; thus, QoS-aware group communications remains an open problem.

A proposed QoS-aware core migration protocol for the multicast protocols uses a group-shared multicast tree.[17] This protocol seeks to construct a tree in which the leaves achieve the multicast application's desired qualities. To reduce the communications cost, the protocol conducts the core selection algorithm only on the current core node. If, for example, delay is the standard QoS metric, the core records the history of delays to group members in terms of the relative time difference between sending the packets to the core and receiving the corresponding acknowledgments from the respective subtree branches. If the averaged delay exceeds the QoS requirement by a given threshold, the core selects a better core candidate from the members close by. Thus, core migration occurs incrementally, which is more suitable to manet dynamics.

Security

Because security is an essential manet requirement, researchers have proposed secure routing protocols for unicast routing in ad hoc networks. Ariadne[18] is a secure on-demand unicast routing protocol that prevents attackers or compromised nodes from tampering with uncompromised nodes. It uses symmetric cryptography and an efficient broadcast authentication scheme to prevent denial-of-service attacks.

Group communications amplify security concerns because they involve more nodes. However, research in this area is just beginning. Ad hoc networks have created additional challenges for implementing security services for wireless communications. The wireless broadcast media is more prone to both passive and active attacks. MAC layer solutions to group-key management and source-authentication proposals for wireline networks must be modified and enhanced for use in a wireless environment. Compared with other wireless communications such as cellular networks, ad hoc networks require even more sophisticated, efficient, and lightweight security mechanisms to achieve the same goals.

Once again, the dynamic characteristics of manets cause these extra challenges. First, ad hoc networks lack trusted centralized infrastructure, which previous security proposals for wireline networks often required. Threshold-based and quorum-based approaches have been investigated to address this problem.

Second, the wireless links between nodes in a manet form and dissolve unpredictably, resulting in ephemeral relationships between nodes. These relationships make it more difficult to build trust based on direct reciprocity.

Third, proposed ad hoc group communication schemes differ markedly from those proposed for wireline networks. In some applications, especially in hostile environments such as battlefield communications, attackers can capture and compromise individual mobile nodes, posing a severe threat to the entire ad hoc network.

Finally, given the stringent nodal budgets in many applications, any solutions proposed for manets must view overhead as a key concern. These applications need strong security mechanisms, yet the solutions must be lightweight in terms of message overhead and computational cost.

For most mobile ad hoc network applications, group communications will be as frequent as unicast communications. Considering the wireless medium's broadcast nature, manets require effective and efficient support for group communications. Although much work has been reported on unicast communications, efforts to improve group communications in manets have not kept pace. Important aspects that researchers must pursue more aggressively include efficient MAC layer support for multicasting and broadcasting and providing reliable and efficient transport layer support.

With advances in support for group communications, the use of multimedia objects, such as video, audio, or images from various sites, will proliferate in manet application domains. Thus, developers must design and integrate QoS support into group communications protocols. The wireless medium's nature, as well as that of its potential applications, warrants the integration of security aspects in all group communications protocols. ■

Acknowledgment

This work was supported in part by the National Science Foundation under grants ANI-0240502 and ANI-0296034, and a generous gift from Hewlett-Packard Corporation.

References

1. J.J. Garcia-Luna-Aceves and E.L. Madruga, "The Core-Assisted Mesh Protocol," *IEEE J. Selected Areas in Comm.*, Aug. 1999, pp. 1380-1394.
2. S.J. Lee, W. Su, and M. Gerla, "On-Demand Multicast Routing Protocol in Multihop Wireless Mobile Networks," *Mobile Networks and Applications,* ACM/Kluwer, vol. 7, no. 6, 2002, pp. 441-453.
3. M. Mauve et al., "Position-Based Multicast Routing for Mobile Ad Hoc Networks," poster paper, MobiHoc 2003.
4. R. Chandra, V. Ramasubramanian, and K. Birman, "Anonymous Gossip: Improving Multicast Reliability in Mobile Ad-Hoc Networks," *Proc. IEEE Int'l Conf. Distributed Computing Systems*, IEEE CS Press, 2001, pp. 275-283.
5. J. Luo, P.T. Eugster, and J-P. Hubaux, "Route Driven Gossip: Probabilistic Reliable Multicast in Ad Hoc Networks," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2229-2239.
6. B. Williams and T. Camp, "Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks," *Proc. ACM Int'l Symp. Mobile Ad Hoc Networking and Computing* (MobiCom 2002), ACM Press, 2002, pp. 194-205.
7. J. Wu and F. Dai, "Broadcasting in Ad Hoc Networks Based on Self-Pruning," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2240-2250.
8. Y-B. Ko and N.H. Vaidya, "Flooding-Based Geocasting Protocols for Mobile Ad Hoc Networks," *Proc. Mobile Networks and Applications*, Kluwer Academic, vol. 7, no. 6, 2002, pp. 471-480.
9. Y-B. Ko and N.H. Vaidya, "GeoTORA: A Protocol for Geocasting in Mobile Ad Hoc Networks," *Proc. 8th Int'l Conf. Network Protocols*, IEEE Press, 2000, pp. 240-250.
10. E. Pagani and G.P. Rossi, "Providing Reliable and Fault-Tolerant Broadcast Delivery in Mobile Ad-Hoc Networks," *Mobile Networks and Applications*, vol. 4, 1999, pp. 175-192.
11. K. Tang and M. Gerla, "MAC Reliable Broadcast in Ad Hoc Networks," *Proc. MilCom 2001*, IEEE Press, 2001, pp. 1008-1013.
12. J. Wieselthier, G. Nguyen, and A. Ephremides, "On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks," *Proc. IEEE Infocom 2000*, IEEE Press, 2000, pp. 585-594.
13. J. Cartigny, D. Simplot, and I. Stojmenovic, "Localized Minimum-Energy Broadcasting in Ad Hoc Networks," *Proc. IEEE Infocom 2003*, IEEE Press, 2003, pp. 2210-2217.
14. C.S. Raghavendra and S. Singh, "PAMAS: Power Aware Multi-Access Protocol with Signaling for Ad Hoc Networks," *ACM Computer Comm. Rev.*, July 1998, pp. 5-26.
15. R. Zheng, J.C. Hou, and L. Sha, "Asynchronous Wakeup for Ad Hoc Networks: Theory and Protocol Design," *Proc. 9th Ann. Int'l Conf. Mobile Computing and Networking* (MobiCom 2003), ACM Press, 2003, pp. 35-45.
16 P. Mohapatra, J. Li, and C. Gui, "QoS in Mobile Ad Hoc Networks," *IEEE Wireless Comm.*, June 2003, pp. 44-53.
17. M. Kochhal et al., "An Efficient Core Migration Protocol for QoS in Mobile Ad Hoc Networks," *Proc. IEEE Int'l Performance Computing and Comm. Conf.*, IEEE Press, 2002, pp. 387-391.
18. Y. Hu, A. Perrig, and D.B. Johnson, "Ariadne: A Secure On-Demand Routing Protocol for Ad Hoc Networks," *Proc. 8th Ann. Int'l Conf. Mobile Computing and Networking* (MobiCom 2002), ACM Press, 2002, pp. 12-23.

Prasant Mohapatra *is a professor in the Department of Computer Science at the University of California, Davis. His research interests include wireless networks, sensor networks, and quality-of-service issues in Internet and Web servers. Mohapatra received a PhD in computer engineering from Pennsylvania State University. Contact him at prasant@cs.ucdavis.edu.*

Chao Gui *is a PhD candidate in the Department of Computer Science at the University of California, Davis. His research interests include wireless networking and mobile computing. Gui received an MS in computer science from the University of Central Florida. Contact him at guic@cs.ucdavis.edu.*

Jian Li *is a PhD candidate in the Department of Computer Science at the University of California, Davis. His research interests include mobile ad hoc and sensor networks. Li received an M.Eng. in pattern recognition and intelligent systems from Tsinghua University, Beijing, China. Contact him at lijian@cs.ucdavis.edu.*

ADVERTISER / PRODUCT INDEX - FEBRUARY 2004

Boldface denotes advertisements in this issue.

Routing and Security in Mobile Ad Hoc Networks

Manets offer a promising new wireless communications paradigm, but researchers must develop efficient routing algorithms and address security concerns before such networks can be extensively deployed.

Nikola Milanovic
Miroslaw Malek
Humboldt University

Anthony Davidson
New York University

Veljko Milutinovic
University of Belgrade

Wireless technologies such as General Packet Radio Service, Wi-Fi, HomeRF, and Bluetooth make it possible to access the Web from mobile phones, print documents from PDAs, and synchronize data among various office devices. However, such applications rely at some point on mobility support routers or base stations, and it is often necessary to establish communication when the wired infrastructure is inaccessible, overloaded, damaged, or destroyed.

Mobile ad hoc networks remove this dependence on a fixed network infrastructure by treating every available mobile node as an intermediate switch, thereby extending the range of mobile nodes well beyond that of their base transceivers. Other advantages of manets include easy installation and upgrade, low cost and maintenance, more flexibility, and the ability to employ new and efficient routing protocols for wireless communication.

We present four manet routing algorithms along with a hybrid approach, discuss their advantages and disadvantages, and describe security problems inherent in such networks.

MANETS

Suppose that we want to easily and efficiently connect two office floors using short-range wireless communication devices. Every employee has one of these mobile devices, and some fixed devices—computers, printers, and so on—have the same capability.

We could connect these devices to the existing wired infrastructure using access points, but this option offers limited mobility, adds load on the wired network, and relies on existing protocols for wired communications. Another possibility is to build a network of dedicated and mutually connected base stations that enable cellular communication, but this is expensive with respect to time, installation, and maintenance.

The best solution is to create a mobile ad hoc network using surrounding electronic devices as intermediate switches when they are idle and if they are capable of performing this task. For example, the packet from one device can hop to the mobile phone of a person passing through the corridor in front of the office, then from the mobile phone to the shared laser printer in the next office, then to someone's digital wristwatch on the floor below, then from the wristwatch to the coffee machine, and, finally, from the coffee machine to its ultimate destination—say, another colleague's device or computer.

Manets are also useful for disaster management. A communications infrastructure is designed to survive common short-term problems, such as overloading, but not to sustain major physical damage. In most cases, the collapse of a single system will cause many dependent devices to fail. If a fire, earthquake, or other natural catastrophe disables a subset of base stations, every mobile phone within range of those stations automatically becomes unreachable. In such situations, rescue workers can use the nodes in manets to create a network "on the fly."

Small-scale manets are also effective for emergency search and rescue, battlefield surveillance, and other communication applications in haz-

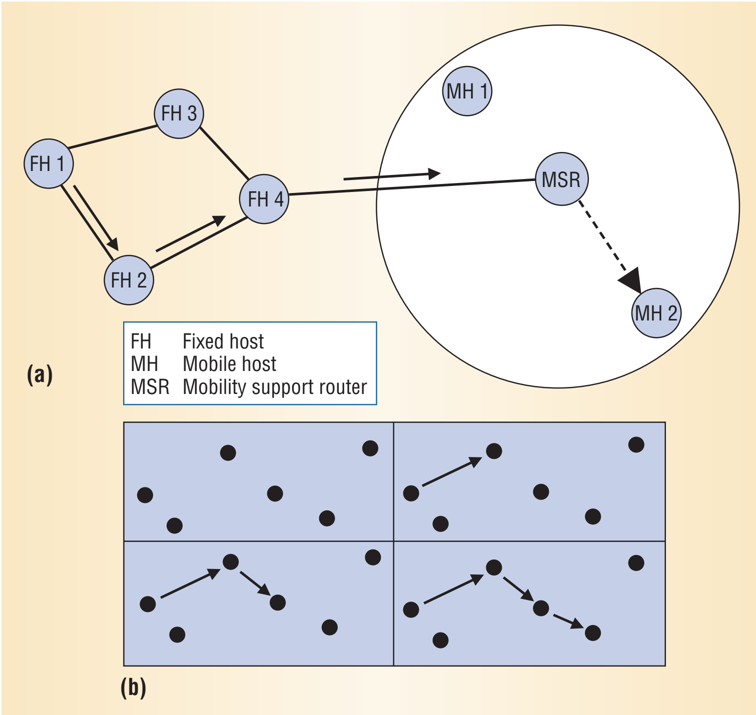

Figure 1. Routing in manets. (a) A classic cellular topology routes each packet in only one hop. (b) Manets route packets in multiple hops.

ardous environments. For example, robots or autonomous sensors deployed in an area inaccessible to humans could use simple manet routing protocols to transmit data to a control center. Even if many robots or sensors are disabled or destroyed, the remaining ones would be able to reconfigure themselves and continue transmitting information.

ROUTING IN MANETS

Efficient routing of packets is a primary manet challenge. Conventional networks typically rely on distance-vector or link-state algorithms, which depend on periodic broadcast advertisements of all routers to keep routing tables up-to-date. In some cases, manets also use these algorithms, which ensure that the route to every host is always known. However, this approach presents several problems:

- periodically updating the network topology increases bandwidth overhead;
- repeatedly awakening hosts to receive and send information quickly exhausts batteries;
- the propagation of routing information, which depends on the number of existing hosts, causes overloading, thereby reducing scalability;
- redundant routes accumulate needlessly; and
- communication systems often cannot respond to dynamic changes in the network topology quickly enough.

Manets use multihop rather than single-hop routing to deliver packets to their destination. As Figure 1a shows, a standard cellular topology routes each packet in only one hop, from the base station to the mobile host. However, as Figure 1b shows, manets can route packets in multiple hops, enabling direct communication between mobile hosts without the need for mobility support router mediation.

ON-DEMAND ROUTING ALGORITHMS

Rather than relying on periodical broadcasts of available routes, algorithms such as dynamic source routing (DSR) and ad hoc on-demand distance vector routing (AODVR) discover routes as needed. Because the route to every mobile node is not known at any given time, these algorithms must build and maintain routes.

Dynamic source routing

DSR[1] is a fairly simple algorithm based on the concept of *source routing*, in which a sending node must provide the sequence of all nodes through which a packet will travel. Each node maintains its own *route cache,* essentially a routing table, of these addresses. Source nodes determine routes dynamically and only as needed; there are no periodic broadcasts from routers.

Figure 2 illustrates the DSR algorithm's route discovery/route reply cycle. A source node that wants to send a packet first checks its route cache. If there is a valid entry for the destination, the node sends the packet using that route; if no valid route is available in the route cache, the source node initiates the route discovery process by sending a special route request (RREQ) packet to all neighboring nodes.

The RREQ propagates through the network, collecting the addresses of all nodes visited, until it reaches the destination node or an intermediate node with a valid route to the destination node. This node in turn initiates the route reply process by sending a special route reply (RREP) packet to the originating node announcing the newly discovered route. The destination node can accomplish this using inverse routing or by initiating the route discovery process backwards.

The DSR algorithm also includes a route maintenance feature implemented via a hop-to-hop or end-to-end acknowledgment mechanism; the former includes error checking at each hop, while the latter checks for errors only on the sending and receiving sides. When the host encounters a broken link, it sends a route error (RERR) packet.

Dynamic source routing is easy to implement, can work with asymmetric links, and involves no overhead when there are no changes in the network. The protocol can also easily be improved to support multiple routes to the same destination.

DSR's main drawback is the large bandwidth overhead inherent in source routing. Because each route cache collects the addresses of all visited nodes, RREQ packets can become huge as they propagate through the network. Routing information can also increase enough to exceed the accompanying message's usefulness. These problems limit the network's acceptable diameter and therefore its scalability.

Ad hoc on-demand distance vector routing

With AODVR,[2] a source node that wants to send a message to a destination for which it does not have a route broadcasts an RREQ packet across the network. All nodes receiving this packet update their information for the source node. Thus, unlike DSR, this approach does not use route caching. Instead, each node maintains only the next hop's address in a routing table, and these routing tables are updated all the way along the RREQ propagation path.

The RREQ contains the source node's address, broadcast ID, and current sequence number as well as the destination node's most recent sequence number. Nodes use these sequence numbers to detect active routes. A node that receives an RREQ can send an RREP if it either is the destination or has a route to the destination with a corresponding sequence number greater than or equal to the sequence number the RREQ contains. In the latter case, the node returns an RREP to the source with an updated sequence number for that destination; otherwise, it rebroadcasts the RREQ.

Nodes keep track of the RREQ source address and broadcast ID, discarding any RREQ they have already processed. As the RREP propagates back to the source, nodes set up entries to the destination in their routing tables. The route is established once the source node receives the RREP.

This algorithm also includes route maintenance facilities. For every route in a routing table, a host maintains a list of neighboring nodes using that route and informs them about potential link breakages with RERR messages. Each node also records individual routing table entries and deletes those not used recently.

AODVR offers several key advantages compared to DSR[3]:

- it supports multicast by constructing trees connecting all the multicast members along with the required nodes;
- smaller control and message packets result in less network bandwidth overhead; and
- the need for only two addresses when routing—destination and next hop—rather than the entire sequence ensures good scalability because packet size does not depend on network diameter.

However, AODVR only works with symmetric links, and because it does not allow for multipath routing, new routes must be discovered when a link breaks down.

LINK-STATE ROUTING ALGORITHMS

Link-state routing algorithms exploit the periodic exchange of control messages between routers, ensuring that the route to every host is always known and immediately providing required routes as needed. However, this proactivity comes at the cost of high bandwidth overhead. Ad hoc link-state routing algorithms attempt to conserve bandwidth by reducing the size and number of control messages.

Optimized link-state routing

Classic link-state algorithms declare all links with neighboring nodes and flood the entire network with routing messages. Optimized link-state routing[4] compacts control packet size by declaring only *multipoint relay selectors*, a subset of neighboring links. To further reduce traffic, OLSR uses only the selected nodes, called *multipoint relays* (MPRs), to flood the network with routing messages.

Each node selects a set of neighboring nodes as MPRs, and these nodes rebroadcast packets received from the originating node. Thus, unlike ordinary broadcast, not every node forwards routing messages. Each node maintains a table of MPR selectors and rebroadcasts every message coming from those selectors. In this way, the network distributes only partial link-state information, which OLSR can use to calculate an optimal route in terms of number of hops.

Each node periodically broadcasts hello messages containing information about its neighbors and a link status. Nodes select the minimal subset of MPRs among one-hop neighbors to cover all nodes two hops away. Thus, every node in the two-hop neighborhood must have a symmetric link to a given node's MPR set.

Because OLSR significantly reduces the number of broadcast retransmissions, this algorithm is most effective in networks with dense node distribution and frequent communication.

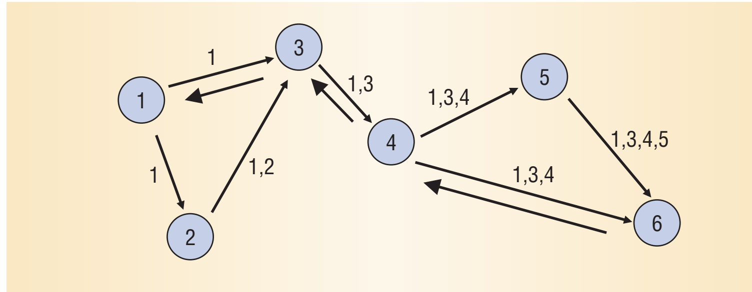

Figure 2. Dynamic source routing. A source node (1) sends a special route request packet to all neighboring nodes, and it propagates through the network. Upon receiving the RREQ, the destination node (6) sends a special route reply packet to the originating node announcing the newly discovered route.

> **A hybrid approach captures the advantages of on-demand and optimized link-state routing for wireless sensor networks.**

Topology broadcast based on reverse-path forwarding

TBRPF[5] broadcasts link-state updates via *source trees* that provide paths to all reachable nodes. It computes these source trees with partial topology information using a modification of Dijkstra's algorithm. Similar to OLSR, each node declares only part of its source tree to neighbors.

TBRPF uses both periodic broadcasts and differential updates to report updates, but each node can declare a full tree, leading to the entire topology's link-state behavior. Each route update travels along a single path to every node on a source tree; leaves do not forward updates. Nodes discover neighbors using differential hello messages that only report changes in the neighborhood, which makes the messages smaller than those in OLSR.

This algorithm is useful in dense mobile networks. Unlike OLSR, it is not limited to two-hop trees, which eliminates redundancy while delivering routing information. Also, while OLSR computes only routes with a minimum number of hops, TBRPF can use arbitrary link metrics if the links are symmetric.

HYBRID APPROACH

A recently proposed hybrid approach[6] captures the advantages of on-demand and optimized link-state routing for wireless sensor networks. This algorithm discovers the route to each node only when it is needed. However, route discovery does not occur through simple flooding but through a mechanism similar to multipoint relays.

The algorithm defines three types of nodes: master, gateway, and plain. A group of nodes selects a master to form a piconet and then synchronizes and maintains the neighbor list. A node can be a master in only one piconet, but it can be a plain member in any number of piconets. Gateway nodes belong to two or more piconets. Only masters and gateways forward routing information; plain nodes receive and process this information, but they do not forward it.

Simulation shows that this algorithm works best when the piconets are densely populated; otherwise, it degrades to simple network flooding. Future research should focus on using some well-defined and accepted metrics, such as power consumption, to compare various ad hoc routing approaches.[7]

SECURITY IN MANETS

The use of wireless links makes manets susceptible to attack. Eavesdroppers can access secret information, violating network confidentiality. Hackers can directly attack the network to delete messages, inject erroneous messages, or impersonate a node, which violates availability, integrity, authentication, and nonrepudiation. Compromised nodes also can launch attacks from within a network.

On-demand and link-state routing algorithms do not specify a scheme to protect data or sensitive routing information. Because any centralized entity could lead to significant vulnerability in manets, a security solution must be based on the principle of distributed trust.

This is similar to the dilemma posed by the classic Byzantine generals problem,[8] in which a general commands each division of the army, and some of the generals, who communicate via messenger, are traitors. All loyal generals must decide upon the same plan of action—that is, a small number of traitors cannot cause the loyal generals to adopt a bad plan. The same holds for manets: A number of compromised nodes cannot cause the network to fail.

Although no single node in a manet is trustworthy, threshold cryptography can distribute trust to an aggregation of nodes.[9] This scheme lets n parties share the ability to perform a cryptographic operation such that any t parties can do it together, while up to $t - 1$ parties cannot perform the operation. However, dividing a private key into n shares and constructing t partial signatures is nontrivial given that traditional key distribution schemes either do not apply to the ad hoc scenario or are not efficient for resource-constrained devices.

Combining identity-based techniques with threshold cryptography can achieve flexible and efficient key distribution.[10] After distribution, a combiner can verify the t signatures and compute the final signature for the certificate. In this way, up to $t - 1$ compromised nodes cannot generate a valid certificate by themselves.

If a large number of nodes are compromised, attributing fault to a specific malicious node is impossible. A proposed algorithm[11] addresses this problem by limiting the possible fault location to the link between two adjacent nodes; as long as a fault-free path exists between two nodes, they can establish a secure communication link even if most nodes in the network are compromised. In addition, this algorithm can detect selfish nodes that refuse to cooperate with other nodes. If their behavior is the result of a denial-of-service attack rather than power-savings activity, the algorithm can isolate the selfish nodes.

Wireless research today primarily focuses on the functional aspect of manets—improving the delivery of packets from one node to another. However, as technology matures, nonfunctional properties such as semantics and security will play the leading role. The challenge lies in managing these two layers, which are orthogonal to each other. If ad hoc communication is to be the foundation for pervasive computing, we must be able to seamlessly interconnect different platforms and devices, offer services on demand, and make it all secure and trusted. ■

References

1. D.B. Johnson, D.A. Maltz, and Y-C. Hu, "The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR)," IETF Mobile Ad Hoc Networks Working Group, Internet Draft, work in progress, 15 Apr. 2003.
2. C.E. Perkins, E.M. Belding-Royer, and S.R. Das, "Ad Hoc On-Demand Distance Vector (AODV) Routing," IETF Mobile Ad Hoc Networks Working Group, Internet Draft, work in progress, 17 Feb. 2003.
3. S.R. Das, C.E. Perkins, and E.M. Belding-Royer, "Performance Comparison of Two On-Demand Routing Protocols for Ad Hoc Networks," *Proc. IEEE Infocom 2000,* vol. 1, IEEE Press, 2000, pp. 3-12.
4. P. Jacquet et al., "Optimized Link State Routing Protocol for Ad Hoc Networks," *Proc. IEEE Int'l Multi Topic Conf., 2001*, IEEE Press, 2001, pp. 62-68.
5. R. Ogier, F. Templin, and M. Lewis, "Topology Dissemination Based on Reverse-Path Forwarding (TBRPF)," IETF Mobile Ad Hoc Networks Working Group, Internet Draft, work in progress, 14 Oct. 2003.
6. N. Milanovic et al., "Bluetooth Ad-Hoc Sensor Network," *Proc. 2002 Int'l Conf. Advances in Infrastructure for e-Business, e-Education, e-Science, and e-Medicine on the Internet,* Scuola Superiore G. Reiss Romoli, 2002; www.informatik.hu-berlin.de/~milanovi/bt_adhoc_sensor.pdf.
7. I. Stojmenovic and X. Lin, "Power-Aware Localized Routing in Wireless Networks," *IEEE Trans. Parallel and Distributed Systems,* vol. 12, no. 11, 2001, pp. 1122-1133.
8. L. Lamport, R.E. Shostak, and M. Pease, "The Byzantine Generals Problem," *ACM Trans. Programming Languages and Systems,* vol. 4, no. 3, 1982, pp. 382-401.
9. Y. Desmedt, "Some Recent Research Aspects of Threshold Cryptography," *Proc. 1st Ann. Workshop Information Security,* LNCS 1396, Springer-Verlag, 1997, pp. 158-173.
10. A. Khalili, J. Katz, and W.A. Arbaugh, "Toward Secure Key Distribution in Truly Ad-Hoc Networks," *2003 Symp. Applications and the Internet Workshops (SAINT 03 Workshops),* IEEE CS Press, 2003, pp. 342-346.
11. B. Awerbuch et al., "An On-Demand Secure Routing Protocol Resilent to Byzantine Failures," *Proc. ACM Workshop Wireless Security,* ACM Press, 2002, pp. 21-30.

Nikola Milanovic *is a PhD student at the Institute for Informatics, Humboldt University, Berlin. His research interests include wireless communications, ad hoc networking, ubiquitous computing, and component- and service-based environments. Milanovic received a Dipl.-Ing. in electrical engineering from the University of Belgrade. He is a member of the IEEE. Contact him at milanovi@informatik.hu-berlin.de.*

Miroslaw Malek *is a professor and chair of computer architecture and communication at Humboldt University. His research focus is on high-performance responsive computing, including parallel architectures, real-time systems, networks, and fault tolerance. Malek received a PhD in computer science from the Technical University of Wroclaw, Poland. He is a member of the ACM. Contact him at malek@informatik.hu-berlin.de.*

Anthony Davidson *is a professor and director of graduate programs in management and systems at New York University's School of Continuing and Professional Studies. His research interests include systems science, enterprise management, strategic planning, and operations. Davidson received a PhD in management and systems science from the City University of London. He is a member of the Association for the Advancement of Computing in Education. Contact him at anthony.davidson@nyu.edu.*

Veljko Milutinovic *is a professor in the Department of Computer Engineering at the University of Belgrade School of Electrical Engineering. His research interests include infrastructure for e-business on the Internet, shared memory multiprocessors, distributed shared memory, and microprocessor architecture and design. Milutinovic received a PhD in electrical engineering from the University of Belgrade. He is a Fellow of the IEEE. Contact him at vm@etf.bg.ac.yu.*

Prioritized Overlay Multicast in Mobile Ad Hoc Environments

The authors propose a model to improve the efficiency and robustness of overlay multicast in manets by building multiple role-based prioritized trees, possibly with the help of location information about member nodes.

Li Xiao
Abhishek Patil
Yunhao Liu
Michigan State University

Lionel M. Ni
Hong Kong University of Science and Technology

Abdol-Hossein Esfahanian
Michigan State University

An increasing number of multicast applications are being developed for mobile ad hoc networks. However, available multicast routing protocols for manets are not as efficient and robust as those used in unicast networks.[1] Many network-layer, or IP-based, multicast routing protocols have been proposed for manets[2] to respond to both network and group dynamics. These protocols require member and nonmember nodes of a multicast group to maintain and update route information, which is very complicated and incurs significant overhead when groups have different priorities.

Recently, many researchers have begun focusing on application-layer, or overlay, multicast. In this approach, participating member nodes perform multicast functions, and an overlay network forms a virtual network consisting of only member nodes atop the physical infrastructure. This eliminates the need to change the application-layer tree when the underlying network changes and enables the overlay network to survive in environments where nonmember nodes do not support multicast functionality. An overlay protocol monitors group dynamics, while underlying unicast protocols track network dynamics, resulting in more stable protocol operation and low control overhead even in a highly dynamic environment.

Although not as efficient as IP-based multicast, overlay multicast is flexible and easy to implement. In some applications, participating nodes can be members of more than one overlay tree or they can build a temporary tree to perform certain important tasks. For such applications to be successful, however, nodes belonging to more than one tree must be smart enough to ignore incoming messages from members in low-priority trees while listening to members from a higher-priority tree.

We propose a prototype of prioritized overlay multicast (POM) for manets in which participating nodes can carry out several different functions and thus be associated with more than one overlay tree. At times some member nodes can form a short-term multicast group to perform certain important tasks. Various overlay trees can have different levels of priority depending on the importance of the service they perform.

OVERLAY MULTICAST IN MANETS

Researchers have proposed numerous protocols to improve the efficiency and reduce the latency of overlay multicast in manets. These include the ad hoc multicast routing protocol, progressively adaptive subtree in dynamic mesh, and the location-guided tree construction scheme.

AMRoute

The ad hoc multicast routing protocol[3] builds a robust multicast network out of user-multicast trees and dynamic logical cores. AMRoute first constructs per-group multicast distribution trees using unicast tunnels between group members. It then creates bidirectional tunnels between neighbors in the multicast tree to form a virtual mesh. From this mesh, the protocol uses a subset of links to generate a shared multicast distribution tree. Packets

0018-9162/04/$20.00 © 2004 IEEE

POM is generally applicable to situations requiring the setup of a communications network without an infrastructure.

physically pass between neighboring nodes via a unicast tunnel and can go through several intermediate nodes. Unicast tunnel paths can change with the network topology without affecting the user-multicast trees.

AMRoute maintains a logical core in every tree responsible for mesh and tree creation. Noncore members cannot perform these actions and act only as passive responding agents. The core can migrate dynamically depending on group membership or network connectivity; thus, core loss will not disrupt data flow. However, AMRoute is inefficient because it uses a static virtual mesh to build the shared multicast distribution tree.

PAST-DM

Progressively adaptive subtree in dynamic mesh[4] is an overlay multicast protocol defined for manets that tries to eliminate redundant physical links and thereby reduce the multicast session's overall bandwidth consumption. Unlike AMRoute, PAST-DM's virtual topology constantly adapts to changes in the underlying network topology.

With PAST-DM, each node implements the expanded ring search algorithm[5] to become aware of neighboring member nodes. Nodes periodically exchange the link-state table with their neighbors in a nonflooding manner such that, after several exchanges, a given node's link state reaches distant nodes. Thus, by looking at each node's link state, a node can view the entire network.

PAST-DM uses this information to build source-based trees, which are more efficient for data delivery than shared trees. Because link-state information is more accurate and up-to-date for nodes closer to the source, the virtual link closer to the source node wins any tie between links of the same cost during tree construction.

LGT

The location-guided tree[6] construction scheme includes two position-based multicast protocols for groups of nodes modeled by complete unit graphs, in which the source of multicast messages and all destination nodes are within transmission radius of one another and aware of the geographic position of any other node in the group.

In the location-guided *k*-ary (LGK) algorithm, the sender node selects *k* nearest destinations as children nodes, groups the rest of the *k* children according to close geometric proximity, and forwards a copy of the packet to each of the *k* children with its corresponding subtree as destinations. The process continues recursively with these children as new source nodes.

The location-guided Steiner tree (LGS) algorithm uses an incremental approach to generate a Steiner tree. Initially, the tree contains only the source node. At each iteration, the algorithm finds the nearest—in terms of geographic distance—unconnected destination node to the partially constructed tree and adds it to the tree; this node receives the message from the tree node to which it connects. Only group nodes are used as tree nodes.

Simulations show that the bandwidth cost for LGS is less than that for LGK when the nodes' location information is up-to-date. However, when the information is obsolete, LGK performs better due to its lower computational complexity. The protocols are efficient solutions for the complete graph environment only in wireless networks in which each node's receiver frequency is known to all other nodes.

Two techniques have been proposed to improve the LGT scheme. In the *optimal paths* method, every node receiving a multicast message for a group of nodes forwards it to each neighbor that is closest to one of the group members. In the *aggregate paths* method, each node counts the closest destinations and then applies a covering algorithm to choose a neighbor that covers the maximum number of destinations. These destinations are eliminated from the list, and then another neighbor is chosen that covers the maximum number of remaining destinations, and so on. As in the optimal-paths method, this algorithm changes a multicast group's forwarding list.

PRIORITIZED OVERLAY MULTICAST

In contrast to these approaches, POM builds priority trees with certain nodes carrying important tasks in overlay networks and rearranges low-priority trees whenever some nodes temporarily move to a high-priority network. The model is generally applicable to situations requiring the setup of a communications network without an infrastructure, such as at a large sports venue.

Policing a stadium or arena with thousands of spectators has always been a challenge. The most effective approach is to scatter security personnel among the crowd to monitor any suspicious activity and, if necessary, request nearby assistance. However, security guards typically convey requests for help as well as information such as physical descriptions of suspects via walkie-talkies, which have static noise. Crowd noise can also make it difficult to hear messages.

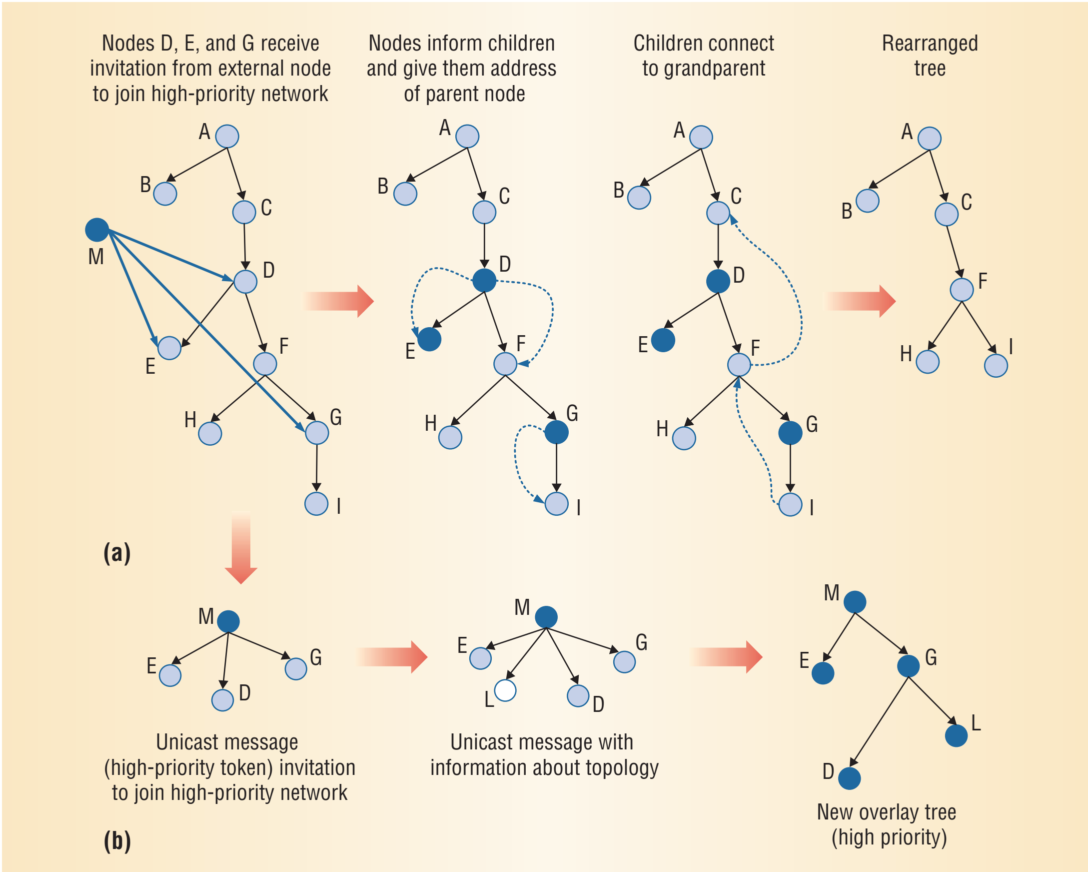

Figure 1. Priority tree formation in POM. (a) Upon receiving an invitation to join a high-priority network, nodes D, E, and G inform their children and provide their parent node's address; the children then connect to their grandparent, resulting in a new tree. (b) Node M sends a unicast high-priority token to each desired node and exchanges information about the formation of the new overlay tree topology.

POM would enable security personnel at such events to exchange multimedia data, such as images or video clips of suspects, via wireless handheld devices using an overlay multicast network. Because this data does not rely on an individual's perception or point of view, it is far more accurate than simple audio descriptions. In addition, a group of security guards could form a temporary network that would give higher priority to messages from group members while ignoring messages from other, lower-priority networks such as those that event operators and managers use.

Priority tree formation

Multiple priority trees can be built in the same environment to offer different services of varying importance. Nodes belonging to different trees switch between networks depending on what functionality they provide. At any given time, a node associates itself only with the highest priority tree in its set, ignoring incoming messages from members in lower-priority trees.

A node that initiates formation of a new priority tree supplies *priority tokens* whose value determines the tree's priority. Thus, a node that is currently a member of priority tree i would not listen to data from member nodes belonging to $i - 1$ or a lower-priority tree. Upon dissolution of i, the member nodes downshift to the next-highest priority tree in the set.

As Figure 1a shows, a node that decides to form a high-priority tree or receives an invitation to join a high-priority network could leave behind several orphans. Because these nodes must connect to another node in the original network to receive data from the source node, a departing node sends a control message to its children informing them that it is leaving the network. The departing node also provides its parent address, enabling the children to contact their grandparent node and receive data from it.

Member nodes use multihop means to communicate with one another; nonmember nodes can act as intermediate nodes. Thus, although nodes F and I are not physically close to each other, they can be neighbors in the logical topology. If the grandparent cannot support the new nodes, it will pass on the connection request to the source node. In location-aware trees, the source node has location information for the entire topology, and it should be able to redirect the orphan's connection to a suitable node.

Figure 1b illustrates the new high-priority tree's formation. External node M contacts nodes D, E, and G of the original tree and another external node, L, to form a high-priority network. M asserts its priority by sending a unicast token message to each of the desired nodes. It then exchanges information about the tree topology's formation, which is based on the nodes' location information with

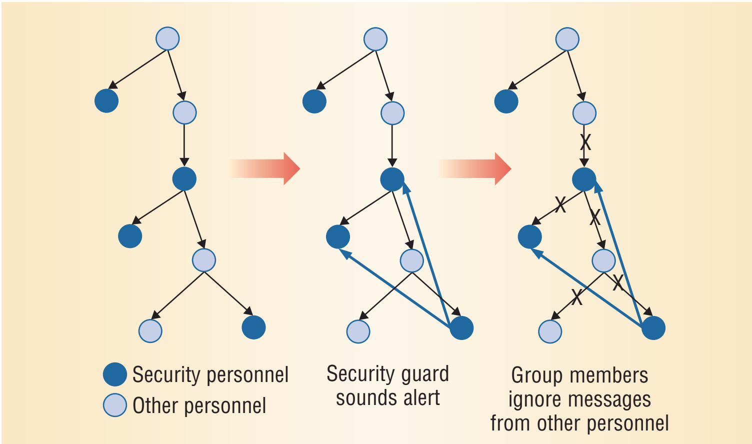

Figure 2. Role-based partitioning of overlay network. Because the network is partitioned at the application layer, nodes in the new higher-priority tree would not recognize members of lower-priority trees.

respect to one another. The number of steps in the tree's final formation varies depending on the implementation or algorithm used. In our approach, the source node implements location discovery and informs the other member nodes.

Location discovery

POM also uses location information to build overlay trees. In the security scenario, security personnel and event organizers would carry mobile devices that are tuned to communicate over the same channel using the same service set identifier for the wireless network. Under normal conditions, the nodes would all belong to the same tree. However, in an emergency situation, a security node can initiate formation of a new high-priority network consisting of only security nodes.

The security guards would carry devices that implement some form of neighbor discovery protocol to identify mobile devices that other nearby guards carry. This protocol can proactively conduct periodic checks for neighboring devices, or the user can initiate the search. When the participating security nodes receive an alert message, they will ignore messages from nonsecurity nodes, which are now considered as belonging to a lower-priority network.

The application layer at the mobile receivers can filter messages according to source or message/group priority. For example, security personnel would carry devices that are smart enough to ignore messages from an event organizer while they are receiving a video clip from another security guard or when they are assigned the task of frequently reporting on a suspect's activities. As Figure 2 shows, the security nodes would be part of a new higher-priority tree; thus, they would not recognize lower-priority tree members.

IMPLEMENTING POM

To assess POM's feasibility, we have identified a suitable unicast ad hoc routing protocol, explored the use of location information to build more efficient overlay trees, and studied how wireless node density and packet size impact system performance.

Simulation methodology

We used the CMU extension for ad hoc networks to simulate POM with an *ns* 2.1b7a discrete event simulator on a Dell Precision Workstation 330 running RedHat Linux 7.3. Currently, *ns* does not have an extension for simulating overlay multicast, although it does have an extension to simulate two network-layer multicast routing protocols: on-demand multicast routing and adaptive demand-driven multicast routing.

With the help of *bash* scripting, we modified the traffic pattern that CMU's *cbrgen* utility generates. We used the extension's *setdest* utility to generate different node positions and movement patterns. The parameters included the number of nodes, pause time, speed, and time and area of simulation. In our simulation, the total number of nodes is 25, of which 12 are member nodes, and the simulation area is 800 × 800 square meters.

The nodes move according to the random waypoint model.[8] When the simulation starts, each node is stationary at a particular location in the specified area for a time equal to the pause time. After the pause time expires, the nodes select a random destination within the given area and start moving with a maximum speed specified during creation of the scenario file. After reaching the destination, the nodes remain stationary for a period equal to the pause time, then they select another destination and proceed toward it in the same manner. Given the high sensitivity of protocol performance to the movement pattern, we carried out simulations for 10 different patterns for every parameter combination.

Unicast protocol identification

Because an overlay network forms a logical network of multicast member nodes, the underlying network regards the data exchange between such nodes as a unicast communication. This communication can use dynamic source routing (DSR),[8] ad hoc on-demand distance vector routing (AODVR),[5] destination-sequenced distance-vector routing (DSDVR),[9] or the temporally ordered routing algorithm (TORA).[10]

To identify an efficient ad hoc routing protocol with low latency, a low drop rate, and minimal overhead for POM, we analyzed simulation results for a general overlay tree. Because it performed poorly compared with the other three protocols, TORA was eliminated after the initial simulation rounds.

The speeds considered in the simulation were 1 meter per second (human walking) and 5 meters per second (human running). We measured the average time to complete the transfer of a 100-Kbyte file

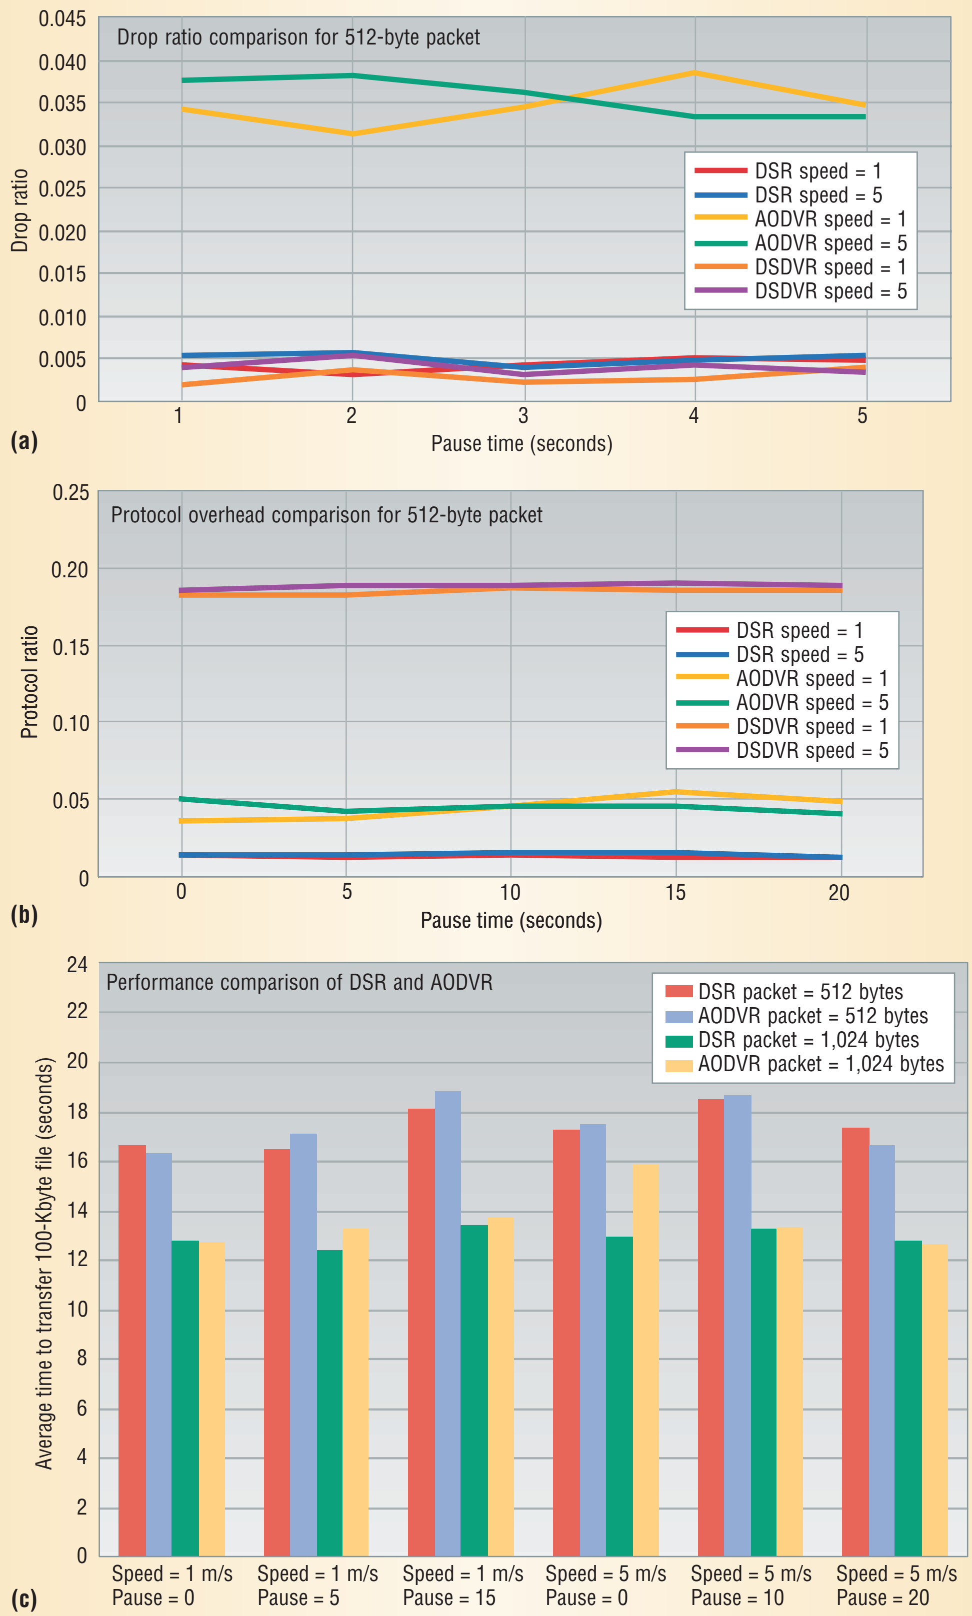

Figure 3. Performance comparison of unicast protocols in POM in terms of (a) drop ratio, (b) protocol overhead, and (c) average completion time. DSR = dynamic source routing; AODVR = ad hoc on-demand distance vector routing; DSDVR = destination-sequenced distance-vector routing.

to all member nodes, the average drop ratio (the ratio of total number of packets dropped to the total number of packets sent), and the average protocol ratio (the ratio of total number of protocol message packets to total number of packets sent).

Figure 3 shows that AODVR has a high drop ratio compared with the other two protocols, while DSDVR has a high protocol overhead. The figure

Figure 4. Impact of location information and node density on performance. (a) Location-aware overlay trees generally have lower latency than trees built without such information. (b) Increasing node density facilitates multihop forwarding, resulting in lower latency.

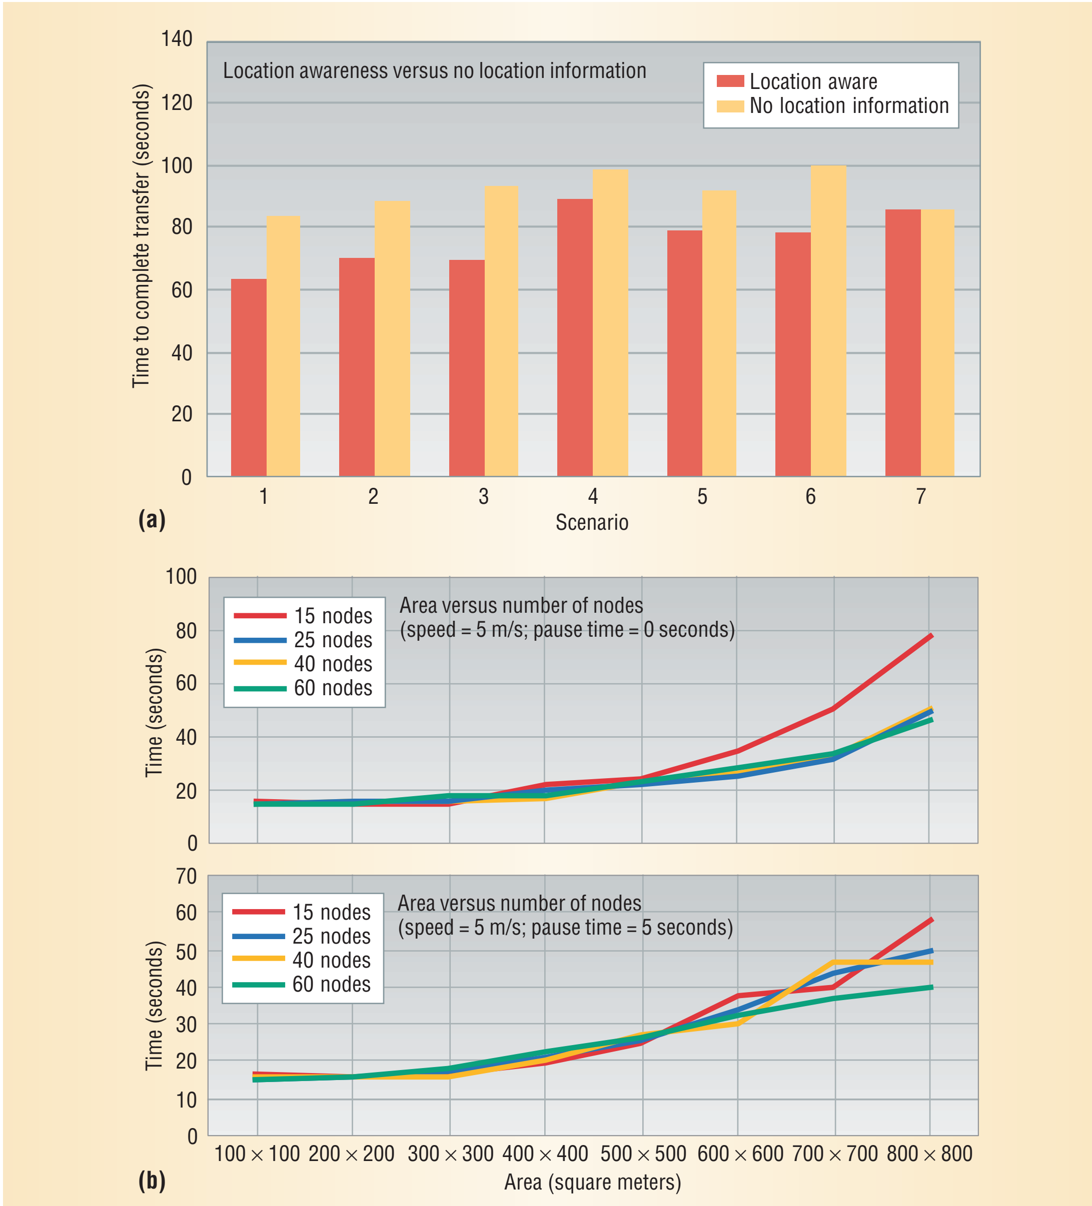

does not show DSDVR's average completion time because it is much higher than that of AODVR and DSR. Completion times for 1,024-byte packets are less than those for 512-byte packets, indicating that larger packet sizes can result in better performance. DSR and AODVR are comparable in terms of transfer time. However, because AODVR's drop ratio increases with packet size, it is high for 1,024-byte packets. Increasing packet size also reduces transfer time.

DSR benefits from source routing, in which packets carry route information to the destination. Consequently, aside from initial route discovery, DSR does not exhibit a high protocol overhead. With AODVR, each node participating between the source and destination must maintain information about the route. Also, to maintain routes, AODVR normally requires periodical transmission of a hello message with a default rate of once per second.

With these results in mind, we settled on a 1,024-byte packet size and DSR as the underlying unicast routing protocol for location information and node density simulations.

Location-based multicast trees

Improvements in location sensing techniques now make it possible to inexpensively locate an object's position within 1 meter in an indoor environment.[11] At the same time, differential Global Positioning System technologies have greatly increased outdoor location-positioning accuracy;[12] CompactFlash cards are available that can be plugged into handheld devices to enable GPS capability. Given the dynamically changing topology of manets, location-sensing technologies could be helpful in identifying nearby member nodes during tree formation—for example, using geometric distance as a heuristic.[6]

Figure 4a compares the performance of overlay trees built with and without location information in seven different movement scenarios. As shown, location-aware overlay trees have lower latency than trees built without such information. However, in extreme cases such as scenario 7, in which one or more member nodes is isolated or does not have overlapping coverage with other nodes, even location information brings little improvement.

Mobile node density

Because each mobile node has limited coverage, node density greatly influences network performance. When there is a high density of nodes in a given area, more nodes are available to perform multihop forwarding. We tested 15, 25, 40, and 60 nodes in areas ranging from 100 × 100 to 800 × 800 square meters. As Figure 4b shows, in smaller areas, the number of nodes has little impact on the network, but as the area increases, nodes scatter and there is little overlap in coverage. Thus, as the number of hops from source to destination increases, latency increases.

POM can be applied to a wide spectrum of mobile communications applications in which setting up an infrastructure-based system is difficult and the organizers desire a role-based partition in their network. For example, at a hospital with an overlay network for medical personnel certain doctors, nurses, and attendants could form their own temporary network to respond to a particular emergency. We continue to explore alternative ways to build a location-aware tree in POM that efficiently balances the tree's effectiveness with the overhead involved in building it. ■

References

1. C. Perkins, "Mobile Ad Hoc Networking Terminology," IETF Mobile Ad Hoc Networks Working Group, Internet Draft, work in progress, 30 Oct. 1997.
2. C.M. Cordeiro, H. Gossain, and D.P. Agrawal, "Multicast over Wireless Mobile Ad Hoc Networks: Present and Future Directions," *IEEE Network,* vol. 17, no. 1, 2003, pp. 52-59.
3. E. Bommaiah et al., "AMRoute: Adhoc Multicast Routing Protocol," IETF Internet Draft, work in progress, 6 Aug. 1998.
4. C. Gui and P. Mohapatra, "Efficient Overlay Multicast for Mobile Ad Hoc Networks," *Proc. 2003 IEEE Wireless Comm. and Networking Conf.,* vol. 2, IEEE Press, 2003, pp. 1118-1123.
5. C.E. Perkins, E.M. Belding-Royer, and I. D. Chakeres, "Ad hoc On-Demand Distance Vector (AODV) Routing," IETF Mobile Ad Hoc Networks Working Group, Internet Draft, work in progress, 19 Oct. 2003.
6. K. Chen and K. Nahrstedt, "Effective Location-Guided Tree Construction Algorithms for Small Group Multicast in MANET," *Proc. 21st Ann. Joint Conf. IEEE Computer and Comm. Societies,* vol. 3, IEEE Press, 2002, pp. 1180-1189.
7. M. Mauve et al., *Position-Based Multicast Routing for Mobile Ad-Hoc Networks,* tech. report TR-03-004, Computer Science Dept., Univ. of Mannheim, 2003.
8. D.B. Johnson and D.A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks," *Mobile Computing,* T. Imielinski and H. Korth, eds., Kluwer Academic, 1996, pp. 153-181.
9. C.E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," *ACM SIGCOMM Computer Comm. Rev.,* vol. 24, no. 4, 1994, pp. 234-244.
10. V. Park and M. Corson, "Temporally-Ordered Routing Algorithm (TORA) Version 1 Functional Specification," IETF Mobile Ad Hoc Networks Working Group, Internet Draft, work in progress, 20 July 2001.
11. L.M. Ni et al., "LANDMARC: Indoor Location Sensing Using Active RFID," *Proc. 1st IEEE Int'l Conf. Pervasive Computing and Communication,* IEEE CS Press, 2003, pp. 407-415.
12. G. Dommety and R. Jain, *Potential Networking Applications of Global Positioning Systems (GPS),* tech. report TR-24, Computer Science Dept., Ohio State Univ., 1996.

***Li Xiao** is an assistant professor in the Department of Computer Science and Engineering at Michigan State University. Her research interests include parallel/distributed and Internet systems, system resource management, and the design and implementation of experimental algorithms. Xiao received a PhD in computer science from the College of William and Mary. She is a member of the IEEE Computer Society and the ACM. Contact her at lxiao@cse.msu.edu.*

***Abhishek Patil** is a PhD candidate in the Department of Computer Science and Engineering at Michigan State University. His research interests include wireless networking, location-aware technologies, and Web technologies. Patil received an ME in electrical and computer engineering from Michigan State University. He is a student member of the IEEE Computer Society. Contact him at patilabh@cse.msu.edu.*

Yunhao Liu *is a PhD candidate in the Department of Computer Science and Engineering at Michigan State University. His research interests include peer-to-peer systems, overlay networking, pervasive computing, distributed systems, network security, and Internet and e-commerce technologies. Liu received an MS in computer science from Michigan State University. He is a student member of the IEEE Computer Society. Contact him at liuyunha@cse.msu.edu.*

Lionel M. Ni *is a professor and head of the Department of Computer Science at Hong Kong University of Science and Technology. His research interests include location-aware computing, high-speed networking, wireless networking, high-performance computer architectures, Internet and Web technologies, parallel and distributed systems, and network security. Ni received a PhD in electrical and computer engineering from Purdue University. He is an IEEE Fellow. Contact him at ni@cs.ust.hk.*

Abdol-Hossein Esfahanian *is an associate professor and graduate program director in the Department of Computer Science and Engineering at Michigan State University. His research interests include graph theory applications in computer science, fault tolerance in computer networks, and the design and analysis of algorithms and database applications. Esfahanian received a PhD in computer science from Northwestern University. He is a senior member of the IEEE and a member of the ACM and the Society for Industrial and Applied Mathematics. Contact him at esfahanian@cse.msu.edu.*

IEEE Names Fellows for 2004

Each year, a select few IEEE members are honored with the title of Fellow. The title is conferred by the IEEE Board of Directors upon a person with an extraordinary record of achievements in any of the IEEE fields of interest. A brief citation notes the accomplishments of each new Fellow.

The IEEE Board of Governors voted in 2003 to confer the title of Fellow upon 260 IEEE senior members, of whom 65 were members of the Computer Society. Senior membership status itself recognizes demonstrated achievement in engineering. As a matter of IEEE policy, the total number of

Fellows selected in any one year may not exceed one-tenth of one percent of the IEEE's total voting membership. In some cases, a Computer Society member has been named as a Fellow based upon contributions to a field other than computing. The name of the IEEE evaluating society (if other than the Computer Society) appears after each new Fellow's citation.

Two IEEE members with no society affiliation were selected as 2004 Fellows for their contributions to computing. They are **Vijay K. Garg,** University of Texas at Austin, for contributions to distributed computing systems and discrete event systems, and **Vishvjit Singh Nalwa,** FullView Inc., for contributions to high-resolution electronically steerable video.

The Computer Society is proud to recognize among its members the following new IEEE Fellows:

B

Miroslav Miodrag Begovic, Georgia Institute of Technology, for leadership in developing analysis tools and protection techniques for electric power transmission systems and renewable generation. (Power Engineering)

Jon Atli Benediktsson, University of Iceland, for contributions to pattern recognition and data fusion in remote sensing. (Geoscience and Remote Sensing)

Jerome John Blair, Bechtel, for contributions to the design and testing of analog-to-digital converters. (Instrumentation and Measurement)

Piero P. Bonissone, General Electric Global Research Center, for leadership in the development of artificial and computational intelligence technologies and their applications to real-world problems. (Neural Networks)

Kim Boyer, Ohio State University, for contributions to computer vision.

C

Shih-Fu Chang, Columbia University, for contributions to digital video and multimedia technologies. (Circuits and Systems)

Ming-Syan Chen, National Taiwan University, for contributions to algorithms for query processing and data management in parallel and distributed systems.

Philip A. Chou, Microsoft, for contributions to variational methods in information theory, signal processing, and compression. (Signal Processing)

D

Frederica Darema, National Science Foundation, for contributions to the programming of parallel and distributed computers.

Serge N. Demidenko, Massey University, for contributions to electronic testing. (Instrumentation and Measurement)

Bart L.R. De Moor, Catholic University Leuven, for contributions to algebraic and numerical methods for systems and control. (Control Systems)

Atam P. Dhawan, New Jersey Institute of Technology, for contributions to optical imaging of skin lesions and multimodality medical image analysis. (Engineering in Medicine and Biology)

John McG Dobbs, Analogic Corp., for contributions to CT scanners and biomedical instrumentation. (Instrumentation and Measurement)

Edmund Durfee, University of Michigan, for contributions to distributed artificial intelligence, multiagent systems, and real-time intelligent control.

E

Joel S. Emer, Intel, for contributions to computer architecture and quantitative analysis of processor performance.

F

James Farmer, Wave 7 Optics, for technical leadership in the cable television industry. (Consumer Electronics)

Aly E. Fathy, University of Tennessee, for contributions to advanced antenna concepts and implementations. (Microwave Theory and Techniques)

Hiromu Fujioka, Osaka University, for contributions to electron beam testing of semiconductor devices and circuits. (Electron Devices)

G

Forouzan Golshani, Arizona State University, for contributions to the field of multimedia information systems.

William Eric Leifur Grimson, Massachusetts Institute of Technology, for contributions to surface reconstruction, object recognition, image database indexing, and medical applications.

Rajesh K. Gupta, University of California, San Diego, for contributions to high-level synthesis and computer-aided design of digital circuits and systems.

H

Graham Reginald Hellestrand, VaST Systems, for contributions to computer system architecture simulations. (Circuits and Systems)

William Evan Higgins, Pennsylvania State University, for contributions to three-dimensional medical imaging and processing. (Engineering in Medicine and Biology)

I

Hajime Ishikawa, Fujitsu Laboratories, for technical leadership in the development of high-performance Si and GaAs devices and circuits. (Electron Devices)

J

Pankaj Jalote, Indian Institute of Technology, Kanpur, for contributions to software process improvement, fault-tolerant computing, and software engineering education.

K

Dilip Dinkar Kandlur, IBM T.J. Watson Research Center, for contributions to the development of quality of service in networks and network servers.

Daniel Koditschek, University of Michigan, for contributions to the theory and practice of robotics and intelligent systems. (Robotics and Automation)

L

Jack N. Little, Math Works, for leadership in the development of engineering software for technical computing. (Control Systems)

Hans-Andrea Loeliger, Swiss Federal Institute of Technology, for contributions to group codes, iterative decoding, and analog implementation of decoders. (Information Theory)

Michael Rung-Tsong Lyu, Chinese University of Hong Kong, for contributions to software reliability engineering and software fault tolerance.

M

Koso Murakami, Osaka University, for contributions to switching technologies and systems for broadband communications networks. (Communications)

N

Yukihiro Nakamura, Kyoto University, for contributions to very large-scale integration synthesis methodologies. (Circuits and Systems)

O

Bjorn Erik Ottersten, Royal Institute of Technology, Sweden, for contributions to antenna signal processing and wireless communications. (Signal Processing)

P

Krishna Palem, Georgia Institute of Technology, for contributions to embedded computing.

Raymond Paul, US Department of Defense, for contributions to metrics-guided testing and evaluation of software systems.

Sanjoy Paul, Bell Laboratories, for contributions to the design and development of communication network protocols. (Communications)

Ron H. Perrott, The Queen's University of Belfast, for contributions to the design and implementation of programming languages of parallel and distributed computers.

R

S. Ramadorai, Tata Consultancy Services, for leadership in the development of multidisciplinary software solutions. (Engineering Management)

Robert Cornelius Rassa, Raytheon, for contributions to automated system testing. (Instrumentation and Measurement)

S

Peter Irvin Scheuermann, Northwestern University, for contributions to logical and physical database design.

M. Ibrahim Sezan, Sharp Laboratories, for technical leadership in digital image and video processing. (Signal Processing)

Ghavam G. Shahidi, IBM Microelectronics, for contributions to silicon-on-insulator technology products. (Electron Devices)

Gurindar Singh Sohi, University of Wisconsin, for contributions to thread-level speculation in computer architecture.

Alfred Spector, IBM T.J. Watson Research Center, for leadership in reliable, scalable distributed computer systems.

Jaideep Srivastava, University of Minnesota, for contributions to the development of models and metrics for multimedia information processing.

Leon Stok, IBM T. J. Watson Research Center, for the development and application of high-level and logic synthesis algorithms. (Circuits and Systems)

Douglas Strain, for leadership in the development of automated test and calibration systems. (Instrumentation and Measurement)

Michael G. Strintzis, University of Thessaloniki, for contributions to digital filtering, image processing, and coding. (Circuits and Systems)

Christer M. Svensson, Linkoping University, for contributions to single-phase clocking and high-speed CMOS circuits. (Solid-State Circuits)

Wim Sweldens, Bell Laboratories, for contributions to multiresolution methods for image and 3D geometry compression.

T

Tieniu Tan, Chinese Academy of Sciences, for contributions to pattern recognition research and applications.

Yuan Yan Tang, Hong Kong Baptist University, for contributions to wavelet analysis to pattern recognition and document analysis. (Systems, Man, and Cybernetics)

David L. Tennenhouse, Intel, for leadership in the development of active networks.

Stuart K. Tewksbury, Stevens Institute of Technology, for contributions to telecommunications and interconnections in high-performance digital systems. (Components, Packaging, and Manufacturing Technology)

Osamu Tomisawa, Mitsubishi, for contributions to low-power, high-speed integrated circuits. (Solid-State Circuits)

Josep Torrellas, University of Illinois at Urbana-Champaign, for contributions to shared-memory multiprocessors.

W

Michael Waidner, IBM Zurich Research Laboratory, for contributions to the theory and practice of information security, privacy, and cryptography.

Steven J. Wallach, Chiaro Networks, for contributions to high-performance computing.

DeLiang Wang, Ohio State University, for contributions to advancing oscillatory correlation theory and its application to auditory and visual scene analysis. (Neural Networks)

Mark Wegman, IBM T.J. Watson Research Center, for contributions to the design, implementation, and analysis of algorithms and compiler technology.

David Wood, University of Wisconsin-Madison, for contributions to the design and evaluation of shared-memory multiprocessors.

Cheng-Wen Wu, National Tsing Hua University, for contributions to design and test of array structures.

Y

Daniel So Yeung, Hong Kong Polytechnic University, for contributions to sensitivity analysis of neural networks and fuzzy expert systems. (Systems, Man, and Cybernetics)

Z

Hong-Jiang Zhang, Microsoft, for contributions to media computing and leadership in content-based visual media analysis, retrieval, and browsing. (Circuits and Systems)

Albert Y. Zomaya, University of Sydney, for contributions to the solution of scheduling problems in parallel computing systems.

Computing in Science & Engineering Seeks Editor in Chief

The IEEE Computer Society and the American Institute of Physics are soliciting applications for a volunteer to serve as editor in chief of *Computing in Science & Engineering*. *CiSE* aims to reach across disciplines using the lingua franca of computation.

The *CiSE* Search Committee seeks nominations from the science and engineering community for candidates suitable for the post. Each applicant should submit a package that includes a letter of employer support, a vision statement that details a proposed leadership plan for the magazine, and a resume that lists publications and editorial experience.

The two-year editor-in-chief term will begin in January 2005. All application materials must be received by **15 March**.

Direct all nominations, inquiries, and application materials to Norman Chonacky, Chair, *CiSE* EIC Search Committee, Department of Earth and Environmental Engineering, Columbia University, Room 918 Mudd—Mail code 4711, 500 West 120th St., New York, NY 10027; chonacky@columbia.edu.

Further application requirements are detailed at www.computer.org/pr/Jan04/cise_eic. htm.

IEEE Foundation Solicits Grant Proposals for 2004

The IEEE Foundation, the philanthropic arm of the IEEE, is soliciting proposals for grants to be awarded in mid-2004. Proposals should be for projects that

- improve education in mathematics, science, and technology from pre-college through continuing education;
- preserve, study, or promote the history of IEEE-associated technologies;
- recognize major contributions to these technologies; or
- promise a major contribution to communities served by the IEEE.

At its November meeting, the IEEE Foundation awarded $138,500 in new grants, bringing the total for 2003 to more than $400,000. This included a $130,000 grant to the IEEE History Center for the operation of the IEEE Virtual Museum during 2004 and 2005. The IEEE Virtual Museum is devoted to helping pre-college students appreciate technology from a social and historical perspective.

The IEEE Foundation also awarded $26,035 to Indiana University-Purdue University in Indianapolis for a pre-college program, *Introducing Robotics to K12*, in which students learn problem-solving skills by using Lego robotic toys to create working prototypes.

To view a full list of the grants awarded in 2003, visit www.ieee.org/organizations/foundation/html/2003grants.html. For information and guidelines on IEEE Foundation grants, visit www.ieee.org/foundation.

Early 2004 grant proposal abstracts are due **12 March**, and full proposals are due by **23 April**. For consideration later in 2004, abstracts are due by **6 August**, and proposals by **17 September**.

IEEE Fellow Nominations Due 15 March

James T. Cain, Chair, 2004 Computer Society Fellows Nomination Committee

The grade of Fellow recognizes unusual distinction in IEEE designated fields, and election to Fellow is a significant honor. This level of membership can only be achieved through nomination and not by a self-generated application. Nominees receive careful scrutiny by the IEEE and an IEEE society.

To nominate a candidate for IEEE Fellow recognition, begin the process by visiting http://elektra.ieee.org/Fellows/FellowNo.nsf. The new Electronic Fellow Nomination Process is detailed at www.ieee.org/about/ awards/fellows/request.htm. The deadline for Fellow nominations is **15 March**. In the event that the online nomination

2004 Meetings and Election Schedule Now Available

The IEEE Computer Society has released its official administrative schedule for 2004. The 2004 calendar includes dates for the three weeklong administrative meetings of the Society's governing boards, including the Board of Governors and subsidiary bodies such as the Technical Activities Board and the Conferences and Tutorials Board.

2004 Election

The 2004 calendar also includes significant dates in the 2004 election process. The 4 October election will name the 2005 first and second vice presidents; the 2005 president-elect, who serves as president in 2006; and seven members of the Board of Governors, who each serve a three-year term. Officers elected in 2004 begin their terms on 1 January 2005.

Nomination recommendations for candidates in this year's election must be received by the Nominations Committee no later than **22 May**. 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 Stephen L. Diamond, Nominations Committee Chair, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1992; voice +1 202 371 0101; fax +1 202 296 6896; NomCom2004@computer.org.

2004 Schedule

The following calendar shows the Board of Governors meeting schedule for 2004, as well as important dates in the Society's election process this year.

28 February: Computer Society Board of Governors Meeting, Savannah, Georgia. Concludes weeklong administrative meetings series for Society governing boards.

15 May: The Nominations Committee sends its slate of officer and board candidates to the Board of Governors.

22 May: Recommendations from membership for board/officer nominees mailed to Nominations Committee.

2 June: Last day to send candidates' petitions, signed by members of the 2004 Board of Governors, to Oscar N. Garcia, Society Secretary, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1992; voice +1 202 371 0101; fax +1 202 296 6896; o.garcia@computer.org.

12 June: Computer Society Board of Governors Meeting, Long Beach, California. Concludes weeklong administrative meetings series for Society governing boards.

14 June: Last day to submit 2005 IEEE delegate-director-elect petition candidates to the IEEE.

30 June: Position statements, photos, and biographies of those candidates approved by the Board of Governors are due at the Society's publications office in Los Alamitos, California, for publication in the September issue of *Computer*.

July: *Computer* publishes the Board-approved slate of candidates and a call for petition candidates for the same officer and Board positions.

30 July: Member petitions and petition candidates' position statements, biographies, and photos due to Society Secretary Oscar N. Garcia at the address above.

August: *Computer* publishes schedule and call for 2006 IEEE delegate-director-elect recommendations to Nominations Committee.

13 August: Ballots are mailed to all members who are eligible to vote.

September: *Computer* publishes position statements, photos, and biographies of the candidates.

4 October: Ballots from members are received and tabulated.

8 October: The Nominations Committee makes recommendations to the Board of Governors for 2006 IEEE delegate-director-elect.

5 November: Computer Society Board of Governors Meeting, New Orleans, Louisiana. Concludes weeklong administrative meetings series for Society governing boards.

5 November: The IEEE delegate-director-elect slate is approved by the Board of Governors.

December: *Computer* publishes election results.

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. Printable forms are available at www.ieee.org/about/awards/fellows/forms.htm. Nominators should avoid sending the forms via fax.

Nominees

A nominee must be a senior member at the time of nomination and must have been an IEEE member in any grade for the previous five years. This includes exchange, student, associate, member, senior, and honorary member, as well as the life category of membership. It excludes affiliate, however, because this category does not include IEEE members. The five-year requirement must be satisfied at the date of election, 1 January 2005; thus, a nominee must have been in any member grade continuously since 31 December 1999. 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 Fellow Committee, or
- the technical society or council's Fellow 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. For the IEEE Computer Society, the point of contact is Lynne Harris, whose address appears at the end of this article. Computer Society members can also obtain 2004 nomination kits by calling the Society headquarters at +1 202 371 0101.

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 may be useful when these entities or individuals are in the best position to provide credible statements.

References

The nominator should select references who are familiar with the nominee's contributions and can provide insights into these achievements. For nominees in the US and Canada, these references must be IEEE Fellows; outside the US and Canada, senior members may serve as references if necessary. References cannot be from the IEEE staff or members of

- the IEEE Board of Directors,
- the Fellow Committee, or
- the technical society or council's Fellow evaluation committee.

While a minimum of five references is 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 an engineer or scientist, technical leader, or educator;
- technical evaluation by one 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 references 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, highly qualified individuals may 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 2004 nomination forms by **15 March**. The staff secretary must also receive at least five Fellow-grade reference letters directly from the references by that date. In addition, the evaluating society or council must receive a copy of the nomination by **15 March**. The deadline will be strictly enforced. If the Computer Society is to conduct the evaluation, send a copy to Lynne Harris, IEEE Computer Society, 1730 Massachusetts Ave. NW, Washington, DC 20036-1992; voice +1 202 371 0101; fax +1 202 728 9614; l.harris@computer.org.

Nominators should review the nomination requirements as described at www.ieee.org/about/awards/fellows/request.htm. ■

Editor: Bob Ward, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; bnward@computer.org

CALLS FOR IEEE CS PUBLICATIONS

Some security practitioners believe that the only way to know how to protect a system against attack is to know how attacks really work. Such practitioners advocate teaching about attacks when building security expertise, carrying out attacks as part of testing, and thinking and writing creatively about attacks. Others feel that discussing, publishing, and teaching attacks is irresponsible.

The July/August special issue of *IEEE Security and Privacy* is devoted to the idea of attacking systems in order to better understand how to defend against them. Suitable topics include exploiting software; using worms, viruses, and malicious code as attack vectors; attacking modern extensible systems such as Java and .NET; and ethical hacking, red-teaming, and penetration testing.

Submissions are due **4 April**. The complete call is available at www.computer.org/security/cfp.htm.

Conventional wisdom holds that security and usability are two antagonistic goals in system design. A classic example is passwords: systems without passwords are thought to be usable, but not very secure, while systems with long passwords that must be frequently changed are thought to be secure, but not very usable.

An alternative view holds that the expanded use of computers by the general public has turned the traditional tradeoff of security-for-usability on its head: Unless designers create systems that are both secure and usable, they will build systems that are neither. Systems that are secure but not usable may fail to gain market acceptance, while systems that are usable but not secure may be hacked and rendered useless.

The September/October issue of *IEEE Security and Privacy* will be devoted to security and usability. Suitable topics include techniques for increasing security and usability, metrics for quantifying user interfaces, user studies involving security and security-related topics, and visualization tools for security and privacy.

Submissions are due **3 May**. The complete call is available at www.computer.org/security/cfp.htm.

OTHER CALLS

UML 2004, 7th Int'l Conf. on Unified Modeling Language, 11-15 Oct., Lisbon, Portugal. Abstracts due **21 Mar.** www.umlconference.org/

SRDS 2004, 23rd Symp. on Reliable Distributed Systems, 18-20 Oct., Florianópolis, Brazil. Submissions due **2 Apr.** www.srds2004.ufsc.br/

ISSRE 2004, 15th Int'l Symp. on Software Reliability Eng., 2-5 Nov., Saint-Malo, France. Abstracts due **2 Apr.,** submissions due **18 Apr.** www.issre.org/2004/

ATS 2004, 13th Asian Test Symp., 15-17 Nov., Kenting, Taiwan. Papers due **15 Apr.** http://ats04.ee.nthu.edu.tw/~ats04/

LCN 2004, 29th IEEE Conf. on Local Computer Networks, 16-18 Nov., Tampa, Fla. Papers due **21 May.** www.ieeelcn.org/

Submission Instructions

The Call and Calendar section lists conferences, symposia, and workshops that the IEEE Computer Society sponsors or cooperates in presenting. Complete instructions for submitting conference or call listings are available at www.computer.org/conferences/submission.htm.

A more complete listing of upcoming computer-related conferences is available at www.computer.org/conferences/.

CALENDAR

MARCH 2004

1-3 Mar: CSEE&T 2004, 17th Conf. on Software Eng. Education & Training, Norfolk, Va. www.cs.virginia.edu/~cseet04/

3-5 Mar: PRDC 2004, 10th Pacific Rim Int'l Symp. on Dependable Computing, Papeete, Tahiti. www.laas.fr/PRDC10/

8-10 Mar: LATW 2004, 5th IEEE Latin American Test Workshop, Cartagena, Colombia. www.latw.net/

10-12 Mar: ISPASS 2004, IEEE Int'l Symp. on Performance Analysis of Systems & Software, Austin, Texas. www.ispass.org/ispass2004/

14-17 Mar: PerCom 2004, IEEE Int'l Conf. on Pervasive Computing & Comm., Kissimmee, Fla. www.percom.org/

20-24 Mar: CGO 2004, Int'l Symp. on Code Generation & Optimization, San Jose, Calif. www.cgo.org/

23-25 Mar: WMTE 2004, 2nd IEEE Int'l Workshop on Wireless & Mobile Technologies in Education, Taoyuan, Taiwan. http://lttf.ieee.org/wmte2003/

23-26 Mar: ICDCS 2004, 24th Int'l Conf. on Distributed Computing Systems, Tokyo. www.cis.ohio-state.edu/icdcs04/

25-26 Mar: HASE 2004, IEEE 8th Int'l Symp. on High-Assurance Systems Eng., Tampa, Fla. http://hasrc.csee.wvu.edu/hase04/

27-31 Mar: VR 2004, IEEE Virtual Reality 2004 Conf. (with Haptics 2004), Chicago. www.VR2004.org/

28-29 Mar: RIDE WS-ECEG 2004, 14th Int'l Workshop on Research Issues on Data Eng. (with ICDE 2004), Boston. www.nvc.cs.vt.edu/ride04/

28-30 Mar: SSIAI 2004, IEEE Southwest Symp. on Image Analysis & Interpretation, Lake Tahoe, Nev. www.ee.ttu.edu/Conferences/SSIAI2004/

28-31 Mar: EEE 2004, IEEE Int'l Conf. on e-Technology, e-Commerce, & e-Service, Taipei, Taiwan. http://bikmrdc.lm.fju.edu.tw/eee04/

29-31 Mar: AINA 2004, 18th Int'l Conf. on Advanced Information Networking & Applications, Fukuoka, Japan. www.takilab.k.dendai.ac.jp/conf/aina/2004/

30 Mar.-2 Apr: ICDE 2004, 20th Int'l Conf. on Data Eng., Boston. www.cse.uconn.edu/icde04/

APRIL 2004

4-7 Apr: ITSW 2004, 11th IEEE Int'l Test Synthesis Workshop, Santa Barbara, Calif. www.tttc-itsw.org/

5-7 Apr: ITCC 2004, 5th Int'l Conf. on IT, Las Vegas. www.itcc.info/

8-9 Apr: IWIA 2004, 2nd IEEE Int'l Information Assurance Workshop, Charlotte, N.C. www.iwia.org/2004/

14-16 Apr: COOL CHIPS VII, Int'l Symp. on Low-Power & High-Speed Chips, Yokohama, Japan. www.coolchips.org/

14-16 Apr: ICECCS 2004, 9th IEEE Int'l Conf. on Eng. Complex Computer Systems, Florence, Italy. www.dsi.unifi.it/iceccs04/

18-21 Apr: DDECS 2004, 7th IEEE Workshop on Design & Diagnostics of Electronics Circuits & Systems Workshop, Tatranská Lomnica, Slovakia. www.ui.savba.sk/DDECS2004/

19-22 Apr: CCGRID 2004, 4th IEEE/ACM Int'l Symp. on Cluster Computing & the Grid, Chicago. www-fp.mcs.anl.gov/ccgrid2004/

19-23 Apr: ASYNC 2004, 10th Int'l Symp. on Asynchronous Circuits & Systems, Hersonissos, Crete. www.async04.gr/

25-27 Apr: EDP 2004, 10th IEEE/DATC Electronics Design Processes Workshop, Monterey, Calif. www.eda.org/edps/edp04/

25-29 Apr: VTS 2004, 22nd VLSI Test Symp., Napa, Calif. www.tttc-vts.org/

26-30 Apr: IPDPS 2004, 18th Int'l Parallel & Distributed Processing Symp., Santa Fe, N.M. www.ipdps.org

MAY 2004

9-12 May: IEEE Symposium on Security and Privacy, Oakland, Calif. www.ieee-security.org/TC/SP-Index.html

10-13 May: ISEE 2004, Int'l Symp. on Electronics & the Environment, Scottsdale, Ariz. www.iseesummit.org

13-14 May: NATW 2004, IEEE 13th North Atlantic Test Workshop, Essex Junction, Vt. www.ee.duke.edu/NATW/

16-19 May: PADS 2004, 18th Workshop on Parallel & Distributed Simulation, Kufstein, Austria. www.pads-workshop.org/pads2004

Call for Articles for *Computer*

Computer seeks articles for a special issue on Internet data centers, to appear in November 2004. Guest editors are Krishna Kant from Intel and Prasant Mohapatra from the University of California, Davis.

Internet data centers form the backbone of most Internet-based services, including e-commerce, IP-based telecom services, hosting services, and the like. As the reach of the Internet widens and more business-critical services are offered, the demands on IDCs grow along multiple dimensions, including responsiveness, service differentiation, security, and availability. Many other forces are likely to affect how the data centers of the future are designed, provisioned, and operated.

Computer's special issue will focus on research issues in identifying and implementing new strategies for optimizing IDCs: application services, protocol enhancements, performance evaluations, provisions for adequate security, protection and isolation, and ensuring an adequate quality of service. *Computer* invites high-quality papers from academia and industry that highlight various problems and solutions and provide a vision for future work in this area.

Topics of particular interest include system architecture and converged data centers; symmetric multiprocessors versus clustered systems; scalability, reliability, and fault tolerance; performance evaluation and workload characterization; operations, control, and autonomic management; power management issues; exploitation of new hardware/software technologies; and issues of security, protection, and isolation.

The deadline for papers is **1 April**. Submission guidelines are available at www.computer.org/computer/author.htm. **Submit manuscripts at http://cs-ieee.manuscriptcentral.com/.**

Send inquiries to the guest editors at krishna.kant@intel.com and prasant@cs.ucdavis.edu.

17-18 May: ICAC 2004, Int'l Conf. on Autonomic Computing (with WWW 2004), New York. www.autonomic-conference.org/

19-21 May: VisSym 2004, Joint Eurographics/IEEE TCVG Symp. on Visualization, Konstanz, Germany. www.inf.uni-konstanz.de/cgip/vissym04/index.shtml

19-22 May: ISMVL 2004, 34th Int'l Symp. on Multiple-Valued Logic, Toronto. www.eecg.utoronto.ca/ ~ismvl2004/

23-28 May: ICSE 2004, 26th Int'l Conf. on Software Eng., Edinburgh, UK. http://conferences.iee.org/icse2004/

24-27 May: ECBS 2004, 11th IEEE Int'l Conf. & Workshop on the Eng. of Computer-Based Systems, Brno, Czech Republic. www.fit.vutbr.cz/events/ECBS2004/

26-27 May: SDD 2004, IEEE Int'l Workshop on Silicon Debug & Diagnosis, Ajaccio, France. Contact Mike Ricchetti, miker@intellitech.com.

JUNE 2004

2-4 June: PBG 2004, Symp. on Point-Based Graphics, Zurich, Switzerland. www.graphics.ethz.ch/PBG/

2-4 June: IWLS 2004, 13th Int'l Workshop on Logic & Synthesis, Temecula, Calif. www.iwls.org/

6-9 June: SWTW 2004, Southwest Test Workshop, San Diego, Calif. www.swtest.org/

7 June: CLADE 2004, Workshop on Challenges of Large Applications in Distributed Environments, Honolulu. www.caip.rutgers.edu/clade2004/

7-9 June: POLICY 2004, IEEE 5th Int'l Workshop on Policies for Distributed Systems & Networks, Yorktown Heights, N.Y. www.research.ibm.com/policy2004/index.html

12-15 June: WICSA 2004, 4th IEEE/IFIP Working Conf. on Software Architecture, Oslo, Norway. http://wicsa4.cs.rug.nl/

21-24 June: CCC 2004, 19th Ann. IEEE Conf. on Computational Complexity, Amherst, Mass. http://facweb.cs.depaul.edu/jrogers/complexity/

23-25 June: MEMOCODE 2004, 2nd ACM/IEEE Conf. on Formal Methods & Programming Models for Codesign, San Diego, Calif. www.irisa.fr/MEMOCODE

23-26 June: IMSTW 2004, 10th IEEE Int'l Mixed-Signal Test Workshop, Portland, Ore. www.ece.pdx.edu/imstw04/

24-25 June: CBMS 2004, 17th IEEE Symp. on Computer-Based Medical Systems, Bethesda, Md. www.cvial.ttu.edu/ Conferences/cbms2004/cbms2004.html

24-26 June: IWPC 2004, 12th Int'l Workshop on Program Comprehension, Bari, Italy. http://iwpc2004.di.uniba.it/

27-30 June: ICME 2004, Int'l Conf. on Multimedia & Expo, Taipei, Taiwan. www.icme2004.org/

27 June-2 July: CVPR 2004, IEEE Computer Soc. Conf. on Computer Vision & Pattern Recognition, Washington, D.C. http://cvl.umiacs.umd.edu/conferences/cvpr2004/

28 June-1 July: DSN 2004, Int'l Conf. on Dependable Systems & Networks, Florence, Italy. www.dsn.org/

JULY 2004

6-9 July: ICWS 2004, IEEE Int'l Conf. on Web Services, San Diego, Calif. http://conferences.computer.org/icws/

6-9 July: CEC 2004, IEEE Conf. on E-Commerce, San Diego, Calif. http://tab.computer.org/tfec/cec04/

7-9 July: ICPADS 2004, 10th Int'l Conf. on Parallel & Distributed Systems, Newport Beach, Calif. www.cacs.louisiana.edu/icpads2004/

ISWC 2004 to Feature Student Wearable Computing Design Contest

ISWC 2004, the eighth annual IEEE International Symposium on Wearable Computers, will feature a student wearable computing design contest. Teams will select a problem from a predetermined list, then design and present a solution or prototype in a poster session at the symposium in November.

A committee of researchers and industry experts from the wearable computing community has solicited real-world problems from business, government, and industry. The contest goal is to increase industry involvement in the development and application of wearable computers. Winning teams will receive a trophy and a certificate for individual members. Other prizes may be added as the contest date approaches.

ISWC 2004 brings together researchers, product vendors, fashion designers, textile manufacturers, and others interested in wearable computing. Visit the ISWC 2004 Web site at www.cc.gatech.edu/ccg/iswc04/ for more information on the symposium and the student design contest. ISWC 2004 takes place from 31 October through 3 November in Washington, D.C.

Industry's First Wireless Broadband Video Phone

D-Link has released a wireless version of the company's D-Link i2eye VideoPhone. According to the company, the wireless i2eye can turn any TV into a broadband video phone, delivering IP videoconferencing without the need to run a network cable from the Internet connection to the television. The new D-Link wireless video phone works with any television and broadband connection to capture and display real-time streaming video images at up to 30 frames per second. The D-Link AirPlus Wireless i2eye VideoPhone costs $249; www.dlink.com.

American Arium Upgrades SourcePoint Debugger

Embedded software tool developer American Arium has released a major upgrade of the company's flagship SourcePoint debugging software to version 6.0. Designed for use with ARM-architecture targets operating in a Microsoft Windows or Linux host environment, SourcePoint 6.0 represents Arium's latest support for ARM7/ARM9, Intel XScale, and TI OMAP processors. Key among the new features is improved trace depth, performance analysis capability, SourcePoint IDE, and debug functionality on Linux; www.arium.com.

OMAP Processor Uses TI's 90-nm Technology

Texas Instruments announced its newest applications processor for mobile phones and other mobile devices. The industry's first multimedia applications processor developed using TI's advanced 90-nm process technology, the OMAP1710 device offers up to 40 percent performance improvement while consuming as little as half the power of current TI application processors. The OMAP1710 leverages the same software environment as all OMAP chips, enabling application developers and mobile phone designers to reuse and build upon existing software; www.ti.com.

Please send new product announcements to products@computer.org.

World's First 10-Gbps Programmable Ethernet Adapter

Napatech and Xilinx announced what they are calling the industry's first 10-Gbps programmable Ethernet server adapter. The NT10F is an advanced network interface card that supports Ethernet connectivity at 10 Gbps. Developers can add custom features like packet tracing, data compression, protocol offload, and encryption-decryption algorithms to the adapter hardware through Xilinx Virtex-II Pro FPGAs and onboard processors. The NT10F 10-Gbps Ethernet Server Adapter costs $10,000; www.xilinx.com.

360-Degree Panoramic Image Capture System

Tokyo-based AOS Technologies has granted an exclusive license to Xybernaut, a US developer of wearable/mobile computing systems, to integrate and market its proprietary Location View System technology. The LVS is a unique multiheaded and 360-degree panoramic image-capture-and-mapping system. Mounted on a variety of transport systems ranging from hand-propelled carts to armored robotic platforms and conventional vehicles, the LVS produces geocoded location view imagery that lets viewers see the scene as if they were in the tracking vehicle; www.xybernaut.com.

RFID Capabilities Added to Wireless Middleware

Responding to the growing demand for wireless enterprise applications that use radio-frequency identification (RFID) technology, Defywire has released a new version of its flagship product. Mobility Suite is a wireless middleware system for giving mobile devices access to data in front office, back office, supply chain, and other enterprise information systems. The new version includes built-in capabilities to share information captured by RFID systems; www.defywire.com.

Free Preview of JDeveloper 10g Available

Oracle is giving developers early access to the upcoming release of JDeveloper 10g. Featuring the Oracle application development framework, JDeveloper 10g is designed to simplify application development by providing a "productivity layer" that lets developers create J2EE applications and Web services optimized for deployment on enterprise grids. Developers can use the new software to focus on defining higher-level business logic and process flows; www.oracle.com.

Sybase Adaptive Server Enterprise on Solaris

Sybase and Sun Microsystems announced the general availability of Sybase Adaptive Server Enterprise (ASE) for the Solaris OS x86 platform. Additionally, Sybase ASE 12.5.1, Developer Edition, is now copackaged with Solaris OS Sparc and x86 Platform Editions, letting customers develop and evaluate applications of Sybase's relational database management system on the Sun platform. Prices start at $1,495; www.sybase.com.

Featuring an adjustable tilt camera lens, the new D-Link wireless i2eye can stream video at up to 30 frames per second to any TV, along with full-duplex audio to the television speakers so that users can speak naturally with one another through the unit's built-in microphone.

STATE UNIVERSITY OF NEW YORK AT BINGHAMTON, Department of Computer Science, The Thomas J. Watson School of Engineering and Applied Science, http://www.cs.binghamton.edu. Applications are invited for two or more anticipated tenure-track positions at the Assistant/Associate Professor level beginning in Fall 2004. Salary and startup packages are competitive. We are especially interested in candidates with specialization in Software Engineering, Programming Languages/Compilers, Information Security, Web-based Systems, Data Mining and areas related to systems development. Applicants must have a Ph.D. in Computer Science or a closely related discipline by the time of appointment. Strong evidence of research capabilities and commitment to teaching are essential. We offer a significantly reduced teaching load for junior tenure track faculty for at least the first three years. Binghamton is one of the four Ph.D. granting University Centers within the SUNY system and is nationally recognized for its academic excellence. The Department has well-established Ph.D. and M.S. programs, an accredited B.S. program and is on a successful and aggressive recruitment plan. Local high technology companies such as IBM, Lockheed-Martin, BAE and Universal Instruments provide opportunities for collaboration. Binghamton borders the scenic Finger Lakes region of New York. Send a resume and the names of three references to Professor Kanad Ghose, Department of Computer Science, State University of New York at Binghamton, P.O. Box 6000, Binghamton, New York 13902-6000. For candidates applying for Fall 04, first consideration will be given to applications that are received by February 27, 2004. Applications will be considered until the positions are filled. Binghamton University is an equal opportunity/affirmative action employer.

PROGRAMMER ANALYST. Masters Degree in MIS or equiv req. Send ad w/resume to: OSI Consulting, 5950 Canoga Ave., #300, Woodland Hills, CA 91367.

THE UNIVERSITY OF TEXAS AT ARLINGTON, Computer Science and Engineering Department, Faculty Openings for Fall 2004. The University of Texas at Arlington (UTA), Computer Science and Engineering (CSE) Department - CSE@UTA invites applications for multiple tenure-track faculty positions at all levels. However, preference will be given to positions at assistant or associate professor levels. All areas of computer science will be considered, including: computer security; bio-informatics, software engineering; pervasive computing; multimedia and video processing; intelligent systems; networks and telecommunications; database and data mining; and applied theory. UTA, part of The University of Texas System, is located in the heart of the rapidly growing Dallas/Fort Worth area, one of the nation's largest high-technology regions, with a flourishing industrial base and excellent opportunities for industry/university collaboration. We at CSE@UTA are committed to excellence in research, teaching, and service. We are in the Fourth year of our "Top 25 Initiative" plan to reach a national top 25 ranking within 10 years. The initiative is strongly supported by all CSE@UTA stakeholders including the UTA administration, faculty, students and alumni, and industry partners. Since 2000, we have added 8 new tenure track faculty and 5 new full-time non-tenure track faculty to our roster of 38 full-time faculty. The number of our PhD students has more than doubled since 2000 and the actively funded research awards exceeded $5.8 million in 2002-03 academic year. Applicants must have an earned doctorate in computer science, computer engineering, or closely related fields and a commitment to teaching and scholarly research. Applicants are expected to have an excellent record of professional accomplishments, commensurate with their level of experience. The faculty opening is anticipated for September 2004. Screening of applications will begin immediately and will continue until all positions are filled. Interested persons should submit a letter of application, a resume, and reference letters online at: http://www.cse.uta.edu/application/. Please note that we do not accept hardcopy submissions. For additional information, please contact: Dr. David Kung, Chair of Search Committee, Department of Computer Science and Engineering, The University of Texas at Arlington, Phone: 817-272-3605, FAX: 817-272-3070, Email: search@cse.uta.edu; http://www.cse.uta.edu. The University of Texas at Arlington is an Equal Opportunity/Affirmative Action Employer.

UNIVERSITY OF FLORIDA, Department of Electrical and Computer Engineering. The Department of Electrical and Computer Engineering invites applications for two tenure-track or tenured faculty positions in the general area of Computer Engineering. Strong candidates in all areas of Computer Engineering are encouraged to apply for tenure-track positions at all ranks. The department is particularly interested in candidates with strong backgrounds and interests in one or more of the following areas: distributed computing, computer architecture, hardware/software systems design, high-performance computing, nanocomputing, adaptive computing, mobile computing, biologically-inspired computing and embedded systems. Applicants for tenured or senior positions must have distinguished records. Candidates should be committed to excellence in teaching and research. All applicants should hold a PhD degree. Salary and support are competitive and commensurate with experience and background. Successful candidates will have the opportunity to work with highly qualified faculty, excellent students and well-equipped laboratories from both the Department of Electrical and Computer Engineering and the Department of Computer and Information Science and Engineering of the University of Florida. More information is available at http://www.ece.ufl.edu and http://www.cise.ufl.edu. Candidates should send their resume, research and teaching plan, cover letter stating the desired position, and names and addresses of at least three references to Dr. José Fortes, Professor and BellSouth Eminent Scholar, C/O Computer Engineering Faculty Search Committee, Dept. of Electrical and Computer Engineering, P.O. Box 116200, 339 Larsen Hall, University of Florida, Gainesville, Florida 32611-6200. Email applications are encouraged. They should be sent in pdf format to cesearch@acis.ufl.edu. The committee will begin reviewing applications on January 28th, 2004 and will continue to receive applications until the positions are filled. The University of Florida is an Affirmative Action Employer and women and minorities are encouraged to apply. According to Florida law, applications and meetings regarding applications are open to the public on request.

TEXAS STATE UNIVERSITY / SAN MARCOS, Department of Computer Science. Applications are invited for two assistant professor level tenure-track positions to begin Fall 2004. Applicants must have an earned doctorate in Computer Science or Computer Engineering or an equivalent degree. A commitment to excellence in teaching, research and service is essential. The Department of Computer Science has 19 faculty members, 485 undergraduate majors and 249 graduate students in Computer Science and Software Engineering. The Bachelor of Science program in Computer Science is accredited by the ABET. We have a thriving master's program and the Ph.D. program proposal is in process. Review of applications will begin on receipt and continue until the positions are filled. For more information and to apply, please see www.cs.txstate.edu/recruitment. Texas State is an Equal Opportunity, Affirmative Action employer. Texas State is committed to increasing the number of women and minorities in faculty and administrative positions.

LOUISIANA STATE UNIVERSITY, ASSISTANT/ASSOCIATE/FULL PROFESSOR, (Computer Engineering/Tenure-track/One or more positions), Electrical and Computer Engineering. The Department of Electrical and Computer Engineering at Louisiana State University invites applications for one or more tenure-track positions in Computer Engineering at all levels available August 2004 or until positions are filled. Though applicants from all areas of Computer Engineering will be considered, of particular interest are applicants with research interests in computer architecture (including microarchitecture, memory systems and special-purpose architectures), design automation (including system-on-chip design and behavioral synthesis), and networking (including traffic engineering and network security). Required Qualifications: Ph.D. or equivalent degree in Electrical or Computer Engineering or related field; potential for excellence in teaching and research. The positions involve teaching graduate and undergraduate courses in electrical or computer engineering and research in areas of individual interests. Salary is competitive and commensurate with qualifications and experience. Release time and resources are provided in order to enhance the development of a quality research program. Opportunities for summer support are available. In addition to the resources in the department, the faculty has access to a 1024-processor cluster. Application deadline is March 1, 2004, or until candidate is selected. Please send your resume (including e-mail address), names of at least three references, and a statement of teaching and research interest, in electronic form to kemin@ece.lsu.edu (preferred) or mail your application package to: Dr. Kemin Zhou, Interim Chair, Electrical and Computer Engineering, Louisiana State University, Ref: Log# 0518, Baton Rouge, LA 70803. LSU IS AN EQUAL OPPORTUNITY/EQUAL ACCESS EMPLOYER.

WASHINGTON STATE UNIVERSITY, Computer Science. Looking for a career in a research university with a very high quality of living? Washington State University in Pullman offers a great opportunity to live in a small town with a strong sense of community, excellent schools, and abundant outdoor recreation nearby, while working with a team of excellent researchers. The School of Electrical Engineering and Computer Science (EECS) at Washington State University is inviting applications and nominations for a faculty position to be filled in Software Engineering. Candidates at all ranks will be considered. Senior applicants must have a strong record of achievement in academia and/or industry. Junior applicants must have earned a PhD in Computer Science by August 16, 2004. All candidates must have both a strong interest in conducting publishable research and a strong commitment to teaching. A record of publication in peer-reviewed journals and conference proceedings is required. The School offers junior faculty a reduced teaching load for the first three years of their appointment. The successful candidate will be expected to teach, effectively communicate and interact with students and colleagues, conduct funded research, publish, and direct MS and PhD student research programs. For appropriately qualified senior candidates in software engineering, the School of Electrical Engineering and Computer Science (EECS) at Washington State University is soliciting applications and nominations for the position of Boeing Endowed Professorship in Software Engineering. The Boeing Professorship has an endowment of $600,000, and the annual proceeds are available to support the research and teaching program of the Chair holder. Candidates must 1) possess a PhD degree in computer science or a related discipline, 2) have a demonstrated record of outstanding accomplishments, and 3) be committed to building an active research program. The School of EECS is the largest of six academic departments in the College of Engineering and Architecture at Washington State University. The School presently has forty faculty, and awards B.S., B.A., M.S., and Ph.D. degrees in computer science, B.S., M.S., and Ph.D. degrees in electrical engineering, and the B.S. degree in computer engineering. Major areas of research emphasis include software engineering, embedded systems, algorithmics, distributed systems, databases, and computer networks. The School has six endowed chairs and distinguished professorships in computer science, computer engineering, and electrical engi-

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Assistant Professor in Computer Science (Interactive Multimedia)

Candidates should have expertise in interactive multimedia technology, human computer interaction, augmented and virtual reality systems, computer graphics, or related areas and contribute to the development of innovative algorithms or systems in these fields.

The professor is expected to establish and lead a research group within the Department of Computer Science at ETH Zurich. Moreover, he or she shall supervise graduate students, teach courses in his or her own field of research, and participate in core courses of computer science. Collaboration with existing groups, in particular with the Computer Graphics Laboratory, is desired.

Applicants with internationally recognized research credentials are asked to send their curriculum vitae, list of publications, names of at least three references, and a short overview of their research interests **to the President of ETH Zurich, Prof. Dr. O. Kübler, ETH Zentrum, CH-8092 Zurich, no later than March 1, 2004.** ETH Zurich specifically encourages female candidates to apply with a view towards increasing the proportion of female professors.

General information about the Department of Computer Science at ETH Zurich can be found at «http://www.inf.ethz.ch/». Questions regarding this position can be addressed to Prof. Dr. P. Widmayer, Head Department of Computer Science, ETH Zentrum, CH-8092 Zurich (email: widmayer@inf.ethz.ch).

INFORMATICS iCORE CIRCLE OF RESEARCH EXCELLENCE

iCORE Research Chair

Distributed High Performance Computing

The Alberta Informatics Circle of Research Excellence (iCORE), in conjunction with Alberta universities, is seeking a Chair to lead a research program in the distributed systems, high performance computing, or grid computing areas. The Chair is expected to be an internationally recognized, exceptional researcher who will develop a research team in one or more of these areas.

Western Canada, through the WestGrid project (http://www.westgrid.ca), has invested $44 million in high performance computing research infrastructure that spans Calgary, Edmonton and Vancouver. These grid-enabled HPC resources are connected by a leading-edge dedicated optical network.

iCORE Chairs can be held at any of the three research universities in Alberta: University of Alberta, University of Calgary and the University of Lethbridge.

UNIVERSITY OF CALGARY University of Lethbridge UNIVERSITY OF ALBERTA

Successful applicants will be appointed with tenure and have substantial research funding for an initial period of five years, renewable once for a second five years.

Salaries and research funding associated with iCORE Chairs are highly competitive.

If you are interested, or know of someone who may be interested, please contact:

Lynn Sutherland, Vice President, Programs
(403) 210-5335 sutherland@icore .ca

www.icore.ca

SINGAPORE MANAGEMENT UNIVERSITY

School of Information Systems

Openings for Faculty

Applications for tenure-track and practice-track are invited at all levels.

The Singapore Management University (SMU) was officially incorporated in January 2000. While it is a public funded institution, it is chartered in a unique way to provide the flexibility and operating characteristics of an American-style private university. SMU's mission is to generate leading-edge business and business-technology research with global impact, and to produce creative and entrepreneurial leaders for the knowledge-based economy.

SMU – Carnegie Mellon Partnership
SMU and Carnegie Mellon University (Pittsburgh, USA) have entered into a close partnership to jointly establish the SMU School of Information Systems (SIS). Carnegie Mellon faculty are actively participating in SIS faculty selection, mentoring and development, and in the design of the SIS undergraduate curriculum, research center, and post-graduate and professional programmes.

SIS Research themes include:

1 e-business technology and management
- Business process integration
- Collaborative work
- Data management
- Internet distributed computing (including web services, grid, peer-to-peer)
- Mobile, embedded and pervasive computing

2 Information security technology and management
- Available and secure computing services and systems
- Available and secure network services and systems
- Secure and legal access to devices and data
- Privacy and trust

3 Architecture and software engineering

4 Information systems management
- Economic and risk analysis of information systems
- Management of design and development projects
- Management of ongoing IT operations

A major research center focusing in interdisciplinary work in cybersecurity and trusted e-business will be launched in collaboration with Carnegie Mellon in 2004.

SIS research and educational projects will demonstrate innovative IT applications and economic value propositions in the following industry sectors: financial services, supply chain & logistics services, manufacturing, health & medical services, and the public sector.

SMU is committed to innovative pedagogy. Candidates must be capable of designing and delivering one of the following undergraduate courses: Object oriented systems, Data management, Networking, Software engineering, Enterprise systems and integration, Security, Architectural analysis, or related electives.

Tenure-track applicants must have a PhD from an internationally recognized university in the areas of Information Systems, Information Technology, Computer Science or related disciplines and an outstanding record of academic research and journal publishing that is commensurate with their desired rank. Tenure-track faculty must also demonstrate a strong interest in innovative research oriented applications in the targeted industry sectors.

Practice-track faculty applicants must also have a PhD in the related IT disciplines from an internationally recognized university, an outstanding record of participating in leading-edge applications that impact business practice, and a record of professionally relevant publications in applied magazines or conferences.

Qualified candidates should submit a cover letter, curriculum vitae, and at least three letters of recommendation and samples of published work. All candidates please submit electronically or hardcopy to:

Dr Steven Miller, Dean, SIS
c/- Office of Faculty Administration
Singapore Management University
469 Bukit Timah Road
Singapore 259756
Telephone: +65 6822 0385
Email: siscv@smu.edu.sg
Website: www.sis.smu.edu.sg

Selected candidates will be asked to interview at Carnegie Mellon University

Singapore Management University
469 Bukit Timah Road
Singapore 259756

www.smu.edu.sg

neering. Two NSF IUCRC centers are in EECS. To learn more about WSU, the School, and faculty research interests, please see http://www.eecs.wsu.edu/. Screening of applications will begin immediately upon receipt of curriculum vitae. Maximum consideration will be given to applications received by January 31, 2004. Applicants should send a cover letter that includes a summary of their research and teaching interests, a curriculum vitae, and the names and addresses of three references qualified to comment on the applicant's research and teaching qualifications to: Chair, Computer Science Search Committee, School of Electrical Engineering and Computer Science, Washington State University, PO Box 642752, Pullman, WA 99164-2752. WSU is an EO/AA educator and employer.

THE UNIVERSITY OF TEXAS AT ARLINGTON, Associate Dean for Research, College of Engineering. The College of Engineering at The University of Texas at Arlington (UTA) invites applications and nominations for the position of associate dean for research. We seek a dynamic and experienced researcher with a record of scholarship and funding that qualifies the person for appointment as a tenured professor in an appropriate department. Candidates must hold a doctorate in engineering, or a closely related field, and must have the communication, management, and leadership skills necessary to promote the diverse research activities in the college and to foster interdisciplinary research programs. The College of Engineering offers thirty-one degree programs including bachelors, masters, and doctoral degrees in aerospace engineering, civil engineering, computer science, computer science and engineering, electrical engineering, industrial engineering, and mechanical engineering. Bachelors and masters degrees in software engineering, masters and doctoral degrees in biomedical engineering, environmental science and engineering, and materials science and engineering, and masters degrees in engineering management and logistics. The College currently has 128 full-time faculty members (108 tenured or on tenure track) and external funding exceeding $15 million. The associate dean will report to the dean and will have lead responsibility, in collaboration with the dean and department chairs, for increasing research funding to our short-term goal of $25 million. The University of Texas at Arlington is located in the heart of the Dallas/Fort Worth/Arlington metropolitan area. UTA is the second largest component of The University of Texas System and is currently the fastest growing university in Texas. The College of Engineering is the third largest in the state and the largest and most comprehensive in the region. DFW, with a population of more than 5 million, is the home to numerous high-

FACULTY POSITION IN COMPUTER ENGINEERING
Computer Sciences Department
The City College of the City University of New York

The City College invites applications for a tenure-track faculty position in Computer Engineering at the Assistant or Associate Professor level, starting with the Fall semester of 2004. Applicants must have a PhD in computer engineering, computer science or a closely related field. Appointment at the Assistant Professor level requires teaching experience and clear evidence of research potential; appointment as Associate Professor requires demonstrated excellence in research and teaching, and a track record of research funding.

Research areas of interest include but are not limited to: next-generation computer architectures, configurable computing technology and FPGA, system engineering of real-time systems and middleware, computer assisted software/hardware engineering, simulation, model driven computing, machine learning, natural language processing, human-computer interfaces, and secure system technology and engineering.

City College is the oldest college of the City University system; Computer Science resides within the School of Engineering. The Department offers the BS jointly with the Department of Electrical Engineering. It offers as well BS and MS degrees in Computer Science; and the PhD in collaboration with the Graduate School of the City University of New York.

The Department's active research reflects its position in the School of Engineering, and encompasses diverse areas of Computer Science and Engineering. Our primary research concentrations are in computational geometry and vision, data systems and informational retrieval, distributed computing systems and algorithms, combinatorial mathematics and optimization, web-based human-computer interaction, information management and E-commerce, multimedia networks and digital libraries, computational methods for image, speech, and language processing, remote sensing, programming systems and language paradigms, and symbolic computation. Further information can be obtained from the Department's web site at http://www-cs.engr.ccny.cuny.edu. Information on the doctoral program in computer science can be obtained at http://web.gc.cuny.edu/Computerscience.

Rank and salary are commensurate with experience and qualifications. Applications with vitae and names and contact information for three (3) references must be postmarked by February 27, 2004, to **Professor Douglas Troeger, Chair (PVN# FY 8882), Computer Sciences Department, R8/206, The City College of New York, Convent Avenue at 138th St., NY, NY 10031.**

the City College of New York

The City College/CUNY is an EEO/AA/IRCA/ADA Employer

NUS National University of Singapore

NATIONAL UNIVERSITY OF SINGAPORE
LEE KUAN YEW POSTDOCTORAL FELLOWSHIP

Applications are invited from young and outstanding academics for prestigious Lee Kuan Yew Postdoctoral Fellowship (LKY PDF) positions with the School of Computing and the Faculties of Engineering, Medicine and Science at the National University of Singapore (NUS).

ELIGIBILITY
- New or recently graduated PhD
- Below 35 yrs of age
- Research interests that fit NUS' strategic directions
- Strong research training
- Strong motivation for an academic research career
- Fluent in spoken & written English

The Fellowship will be tenable for up to 3 years in the first instance, with possible extension for 2 further years. PDFs can apply for academic positions following the Fellowship.

REMUNERATION. Gross annual salary ranges from **S$72,000 to S$144,000** with commencing salary depending on qualifications and experience.

Leave and medical benefits will be provided. Other benefits may include: provident fund benefits or an end-of-contract gratuity, a settling-in allowance, subsidized housing, passage assistance and baggage allowance on appointment.

LKY PDFs will be entitled to the same eligibility as normal academic staff for research support, including research grants.

APPLICATIONS. Applicants are to send the following documents to the relevant NUS Faculty/School (Contact details at **www.nus.edu.sg/ore/resource/fac_resrc.htm**):
- NUS Personal Particulars Form
- Detailed Curriculum Vitae & Educational Certificates
- 3 International Referee Reports (including contact details)
- Statement of Research Intent
- List of Publications

Closing date: 10 Apr 2004

Further information: **www.nus.edu.sg/ore/funding/fellow/lky_pdf.htm**
General Enquiries: **oremyn@nus.edu.sg**

D. E. Shaw Research and Development

Research on Algorithms and Architectures for Computational Biochemistry

Extraordinarily gifted computer scientists, systems architects, electrical engineers and systems software professionals are sought to join a rapidly growing New York–based research group pursuing an ambitious, long-term project aimed at achieving major scientific advances in the field of biochemistry and fundamentally transforming the process of drug discovery. This research effort is being financed by the D. E. Shaw group, an investment and technology development firm with approximately $6 billion in aggregate capital, and operates under the direct scientific leadership of its founder, Dr. David E. Shaw.

Among the group's current research activities is a project aimed at developing a massively parallel special-purpose supercomputer and innovative mathematical and computational techniques to direct unprecedented computational power toward the solution of key scientific and technical problems in the fields of molecular simulation and molecular design. Successful candidates will be working closely with a number of the world's leading computational chemists and biologists, and will have the opportunity not only to participate in an exciting entrepreneurial venture with considerable economic potential, but to make fundamental contributions within the fields of biology, chemistry and medicine.

Serious candidates will have an exceptionally distinguished history of academic and/or industrial accomplishment in computer science, electrical engineering, applied mathematics, or a related area. Particularly relevant areas of expertise might include parallel computation, high-speed interconnection networks, scientific computing, numerical analysis, optimization, the analysis of algorithms, operating systems, digital systems simulation, reconfigurable computing, and ASIC design, but specific knowledge of any of these areas is less critical than exceptional intellectual ability and a demonstrated track record of achievement. We are prepared to reward exceptionally well-qualified individuals with above-market compensation.

Please send your curriculum vitae (including list of publications, thesis topic, and advisor, if applicable) to Research.Development7@deshaw.com.

D. E. Shaw Research and Development, L.L.C. does not discriminate in employment matters on the basis of race, color, religion, gender, national origin, age, military service eligibility, veteran status, sexual orientation, marital status, disability, or any other protected class.

D E Shaw & Co

tech, manufacturing, and infrastructure companies and to a wide-range of cultural, entertainment, and sport venues. Applications and nominations should be submitted by e-mail to shirazi@cse.uta.edu (Behrooz Shirazi, Search Committee Chair, Department of Computer Science and Engineering, The University of Texas at Arlington, PO Box 19015, Arlington, TX 76019-0015.) Applications must include a curriculum vita; a statement of the candidate's academic and administrative philosophy; and the names, mailing and email addresses, and phone numbers of five references. Review of applications will begin January 15, 2004 and will continue until the position is filled. Final offer of employment is contingent on completion of a satisfactory criminal background investigation for security sensitive positions. The University of Texas at Arlington is an Equal Opportunity and Affirmative Action Employer.

WEST VIRGINIA UNIVERSITY, Lane Department Chair and Professorship, Lane Department of Computer Science & Electrical Engineering, College of Engineering & Mineral Resources, Morgantown, West Virginia. West Virginia University is conducting a national search to fill the position of Chairperson and Lane Professor for the Lane Department of Computer Science and Electrical Engineering. The University seeks a dynamic leader with a strong record of scholarly and/or professional accomplishments. Candidates must have an earned doctoral degree in CpE, CS, EE or a closely related field and must have credentials suitable for tenure at the level of full professor in the department, with visibility and leadership in teaching and research and the interpersonal skills necessary to serve as a mentor, counselor, and leader of faculty and staff. Successful candidates will have either demonstrated ability or strong potential to work well with constituencies both internal and external to the university including alumni and the corporate community, work collaboratively with faculty and administration to further develop the department's growing research program, and effectively obtain external development funds and gifts in support of continued department growth. The Chair reports to the Dean of the College and is responsible for administration of all programs and budgets of the Department. The department has 32 tenure track faculty, 9 research faculty and/or research associates, and 7 support staff. It has an enrollment of over 500 undergraduate students and over 300 graduate students. The Lane department has degree programs in Computer Science, Computer Engineering, Electrical Engineering, Software Engineering, and Biometric Systems. Major research initiatives include: Biometrics, Software Engineering, Electric Power Engineering, Integrated Devices, and Sensors. Growing research areas include: Computer Forensics, Communications and Nanotechnology. Departmental research funding has seen considerable growth, including funding from corporate and government sources. WVU is a Carnegie Doctoral/Research-Extensive institution and a public land grant university. On its main campus in Morgantown, WVU enrolls over 24,500 undergraduate, graduate and professional students. Morgantown is a safe vibrant college community in a beautiful rural region and was recently rated the No. 1 Small City in America and the Best Small City in the East. The Morgantown area is home to several national facilities and laboratories including: the FBI Automated Fingerprint Identification System (AFIS) Center, the NASA Independent Verification and Validation (IV&V) Facility, DoE National Energy Technology Laboratory (NETL), National Institute for Occupational Safety and Health (NIOSH) and DoD Biometric Fusion Center (BFC). Interested candidates should submit a letter of application, a description of leadership style and department vision, a current resume and a list of three references to: Warren R. Myers, Associate Dean, Lane Department Chair and Professorship Search Committee, College of Engineering and Mineral Resources, West Virginia University, P.O. Box 6070, Morgantown, WV 26506-6070. For best consideration completed applications should be received by March 1, 2004, but the position will remain open until filled. WEST VIRGINIA UNIVERSITY IS AN EQUAL OPPORTUNITY/AFFIRMATIVE ACTION EMPLOYER.

SECURITY SYSTEMS SOFTWARE PROGRAMMER. Bach Deg in Electronics Engr or equiv req. Send ad w/resume to: Security Frameworks, 1875 Century Park East, #1185, Los Angeles, CA 90067.

SIMON FRASER UNIVERSITY, Canada Research Chair (CRC) – Tier I. The School of Computing Science at Simon Fraser University is seeking candidates for a prestigious Canada Research Chair (CRC) Tier I senior-level faculty position. Applicants are expected to have a truly outstanding research program with a strong record of publication, research funding, and graduate student supervision. Applicants also must have a strong commitment to excellence in teaching. The School is interested in candidates in systems, application and foundational areas. We also have a special interest in interdisciplinary areas. However, higher priority will be given to the overall innovation and achievement of the candidate than to any specific area. Additional information about the CRC chairs can be found at www.chairs.gc.ca. Simon Fraser University is consistently one of the top-ranked public universities in Canada. The School of Computing Science currently has over 170 Ph.D. and M.Sc. students, 700 undergraduate majors, and over 48 faculty members. The School is at the outset of a phase of strong growth. As part of the "Double the Opportunity" program of the government, both the number of faculty and students will significantly increase over the next three years. The School of Computing Science highly encourages interdisciplinary research building upon the strengths of Simon Fraser University. The main campus of Simon Fraser University is situated on Burnaby Mountain in Greater Vancouver. Vancouver thrives as one of the most scenic waterfront cities, located just minutes away from the mountains and a wide range of outdoor activities. It has the mildest climate in Canada. Vancouver's cultural events, skiing, boating, hiking, favorable climate, and clean, safe environment consistently make it one of the most desirable places in the world to live and work. All qualified candidates are encouraged to apply; however, Canadians and Canadian Permanent Residents will be given priority. Simon Fraser University is committed to employment equity and encourages applications from all qualified women and men, including visible minorities, aboriginal people and people with disabilities. Applications will be accepted until the position is filled. For updated information see www.cs.sfu.ca. Applicants should send a CV, reprints of representative publications and names of six references to: CRC Search, School of Computing Science, Simon Fraser Uni-

Mississippi State UNIVERSITY

Mississippi State University invites applications for tenure-track faculty positions in the Department of Electrical and Computer Engineering; all ranks will be considered with an endowed chair available to qualified candidates. Areas of particular focus include: embedded and real-time computer systems, parallel architectures, asynchronous circuits, microelectronics, hardware/software co-design, speech and information processing, visualization, and telecommunications. An earned doctorate in Electrical Engineering, Computer Engineering, or a related field and a clear potential for gaining national prominence through funded research and teaching are required. Industrial or federal laboratory experience is highly desirable. Applications will be accepted until positions are filled. To apply, please visit **http://www.ece.msstate.edu/** employment and include a resume and list of three references. Inquiries can be made to **employment@ece.msstate.edu** or 662-325-3912.

MSU is an AA/EOE.

versity, 8888 University Drive, Burnaby, B.C., Canada V5A 1S6, Email: crc-search@cs.sfu.ca.

VIRGINIA TECH, Department of Computer Science, Bioinformatics Faculty Positions. The Department of Computer Science at Virginia Tech seeks applications for several tenure-track positions in the Department of Computer Science from individuals desiring to make fundamental contributions to both computer science and the life sciences in bioinformatics, as broadly defined. Special funding from the Commonwealth of Virginia provides competitive salaries and startup funding for tenure-track faculty positions at all ranks. CS faculty in bioinformatics have access to the 2200-processor Terascale Computing Facility recently established by the university, as well as other multiprocessor clusters within the Department. Excellent opportunities for collaborative research exist with researchers in life science departments and at the Virginia Bioinformatics Institute, which is located on campus. Applicants for a senior position must have a significant track record of grant funding. All applicants must have a PhD in Computer Science or an allied area, a demonstrated record of publications in computer science or computational science, and a commitment to addressing significant life science problems. Ability to collaborate with researchers within the Department and in the life sciences is required. A demonstrated record of accomplishments in bioinformatics is preferred. Additional information is available at http://www.cs.vt.edu/FacultySearch. Applicants should send a curriculum vitae, a 1-2 page statement of research goals in both computer science and life science, and at least three letters of reference, as separate PDF files, by email to facultysearch@cs.vt.edu. Review of candidates will begin January 5, 2004 and continue until the positions are filled. Virginia Tech is an equal opportunity employer.

PURDUE UNIVERSITY, Department of Industrial Technology, Department of Computer Sciences. The Department of Computer Sciences and the Department of Industrial Technology at Purdue University invite applications for an academic year tenure-track Assistant Professor position to begin fall semester 2004. This shared position will require research in the area of biometrics, teaching biometrics within the Department of Industrial Technology, and teaching undergraduate courses each semester in the Department of Computer Sciences. Further information about the program in biometrics is available at http://www.tech.purdue.edu/it/resources/biometrics/. The home department will be the Department of Industrial Technology. Both departments offer a stimulating and nurturing academic environment. Forty faculty members in

UNIVERSITY *of* VIRGINIA

CHAIRPERSON - Department of Computer Science

Department Chairperson, Department of Computer Science School of Engineering and Applied Science University of Virginia The University of Virginia School of Engineering and Applied Science invites applications for the chairperson of the Department of Computer Science. The University of Virginia has approximately 18,000 students pursuing degrees in nine colleges and professional schools, including the School of Engineering and Applied Science (SEAS). The University of Virginia is consistently ranked as one of the top public universities in the country.

The School of Engineering and Applied Science has approximately 2,000 undergraduate students, 625 graduate students, 150 tenure and tenure track faculty members in nine departments, and conducts approximately $45 million per year in externally funded research programs. The Department of Computer Science has more than 300 undergraduate students, 110 graduate students, and 27 (23 tenured or tenure-track and four teaching) faculty members.The department offers B.S., M.S., M.C.S., and Ph.D. degrees in computer science and B.S., M.S., M.E., and Ph.D degrees in computer engineering in conjunction with the Department of Electrical and Computer Engineering.The department has undergone significant growth and programmatic evolution and aspires to become one of the top departments in experimental systems research and the leading institution in undergraduate CS education. The department has a novel undergraduate curriculum, is actively expanding its graduate research program, and is engaged in multi-disciplinary programs with other departments, schools, and universities.

Active research programs exist in the following areas: algorithms, architecture, compilers and languages, computer-aided design, computer graphics, computer security, distributed systems and databases, mobile computing, computer and wireless sensor networks, operating systems, parallel and grid computing, programming environments, real-time and embedded systems, and software engineering.

We seek a chairperson with a strong research reputation who can provide leadership for a department committed to excellence in teaching, research, and service, guide the growth of new promising directions at a university that is committed to enhancing the general area of information technology, and visibly enhance the department's national stature. Candidates must have an earned doctorate in an appropriate discipline and an outstanding record of scholarship and research achievement.Candidates with exceptional credentials will be considered for an endowed professorship. Applications will be reviewed as they are received and until the position is filled.

Applicants are encouraged to visit the departmental website for more information at: **http://www.cs.virginia.edu/**. Please submit a letter of application to:

Professor James H. Aylor, Search Committee Chair
School of Engineering and Applied Science
University of Virginia, 351 McCormick Road, P. O. Box 400232
Charlottesville, VA 22904-4743

Email: jha@virginia.edu

The University of Virginia is an Equal Opportunity/Affirmative Action Employer. Women and members of minority groups are strongly encouraged to apply.

DIRECTOR OF COMPUTER ENGINEERING PROGRAM

Senior Faculty Position
Departments of Computer Sciences and Electrical Engineering
The City College of the City University of New York

The Departments of Computer Science and Electrical Engineering of The City College seek to hire an outstanding faculty member as Director of the Computer Engineering Program. The Computer Engineering Program has been operating as a degree-granting program since Fall 2000 and is jointly administered by the Computer Sciences and Electrical Engineering Departments. More than 300 students are currently enrolled in the program.

Qualifications for the position of Professor - in Computer Sciences or Electrical Engineering, depending on the candidate's background - and Director of the Computer Engineering Program include: an earned PhD in a Computer Engineering-related field; proven ability to leverage day-to-day operations with strategic plans; experience in academic program development and leadership; evidence of ability to attract sources of external funding; and demonstrated commitment to excellence in undergraduate and graduate education. Research areas of interest include but are not limited to: next-generation computer architectures, configurable computing technology and FPGA, system engineering of real-time systems and middleware, computer assisted software/hardware engineering, simulation, model driven computing, machine learning, natural language processing, human-computer interfaces, and secure system technology and engineering.

Day-to-day duties of the Director include the following: conduct research in Computer Engineering; provide academic advising to Computer Engineering students; foster the growth of the Computer Engineering Program; help build a research environment that attracts sustained funding, external partnerships and opportunities for student involvement; create and further advance opportunities for collaborative relationships between Computer Engineering faculty and industry/community partners; sustain the faculty's innovative research, teaching, and curriculum development at the interdisciplinary intersections of Software Engineering, Electrical Engineering, Computer Engineering, and Computer Science; work with the Chairs of CSs and EE departments and the Dean of SOE to provide academic leadership; create a supportive environment in which faculty, particularly junior faculty, can thrive; coordinate/manage ABET needs and requirements as far as Computer Engineering accreditation is concerned and coordinate efforts with Associate Deans to ensure that assessment of CpE students is performed on a continual basis; teach (at a reduced level) within the program.

Appointment is tenure track, at a senior rank; salary is competitive. There is a possibility of a recruitment award. Exceptional candidates may be considered for appointment at the rank of Distinguished Professor.

Further information about the Departments may be obtained from our web sites at: http://www-cs.engr.ccny.cuny.edu and http://www-ee.ccny.cuny.edu.

the City College of New York

Applications with curriculum vitae, detailed letter of application, and the names and contact information for three (3) references must be postmarked by February 27, 2004, to **Director of Computer Engineering Search Committee (PVN# FY 8878), Steinman Hall, Room 142, The City College of New York, Convent Avenue at 138th St., NY, NY 10031.**

The City College/CUNY is an EEO/AA/IRCA/ADA Employer.

Computer Science direct research programs in analysis of algorithms, bioinformatics, compilers, databases, distributed and parallel computing, graphics and visualization, information security, networking and operating systems, programming languages and compilers, scientific computing, and software engineering. Further information about the department is available at http://www.cs.purdue.edu. The Department of Industrial Technology, with 10 full-time faculty members, is one of eight departments in the School of Technology. Master's and Ph.D. degrees in Technology with an area of specialization in Industrial Technology are offered through the department. Further information on the Department of Industrial Technology is available at http://www.tech.purdue.edu/it. Applicants must hold a Ph.D. in Computer Science, Computer Technology, or a closely related discipline; be committed to excellence in teaching; and have demonstrated strong potential for excellence in research. Salary and benefits are highly competitive. Special departmental and University initiatives are available for junior faculty. Applicants are encouraged to apply by sending a letter of application, a curriculum vita, a statement of career objectives, and names and contact information of at least three references to Dr. Stephen J. Elliott, Biometrics Faculty Search Committee Chair, 401 N. Grant Street, Knoy Hall, Purdue University, West Lafayette, IN 47907-2021, phone (765) 496-2474, email biometrics-fac-search@tech.purdue.edu. Applications are now being accepted electronically in PDF format and will be considered until the position is filled. Purdue University is an Equal Opportunity/Equal Access/Affirmative Action employer and is committed to building a diverse faculty of excellence.

LOTUS NOTES DEVELOPER, PROJECT MANAGER. Devel Lotus Notes doc mgmt & work flow apps w/ Lotus Script, Lotus Formula lang; test & devel databases w/ LEI & Notetrix integration tools; perf server maintenance; monitor server; diagnose & correct prob w/ Notes Software; resolve prob in a WAN envrnmnt; assist in req anal for h/ware & s/ware selection for Notes network; assist in plan & design Notes expan & imprvmnts; maintain docs on install & config of s/ware, upgrades, systems admin, & statistical reports; manage app devel projects through project life cycle. Henkel Loctite, Bay Point, CA, F/T, 40 hrs./wk, 9-5, sal comm w/ exp REQs: Bach's deg, or equiv in educ, training, & exp, in Comp Sci, Info Sys & 2 yrs in job offrd or 2 yrs exp as Consultant, Programmer, Analyst. Exp must incl 2 yrs Lotus Notes app devel, Lotus Script, Lotus Formula lang, LEI, Notetrix & prob resol in WAN envrnmt; R5CLP cert. Send resume to IEEE Computer Society, 10662 Los Vaqueros Circle, Box# COM2, Los Alamitos, CA 90720.

SOFTWARE ENGINEER. Design, develop and implement Software components by using Telecommunication Management Network (TMN) Tools, Windows Application (Gateways) and ANSI, ITU Standards. Use C, C++, JAVA, SQL languages, VC++ and OSI Stack on various operating systems such as UNIX, LINUX and Windows. Must have knowledge of seven layers protocol stacks includes ES-IS, IS-IS Routing Protocol, rfc1006, TCP/IP, X.25 & Ethernet. Must have knowledge of DSET's TMN Tools (DSG, GDMO AgentToolkit, Manager-Toolkit, ASN.C, ASN.C++, CMISE, ACSE tools development) and Applications. Must have knowledge of ISO standards in software lifecycle process. Must be conversant with the version controlling systems such as Clearcase. Must have a Bachelor's degree in Computer Information System and 18 mths exp in the job offered. 40hrs/wk, 9-6. Send your resume to NE Technologies, Inc. at 5085 Avalon Ridge Parkway, Ste 100, Norcross, GA 30071. Reference No. 234.

UNIVERSITY OF TENNESSEE, Assistant/Associate Professor in Mathematics or Computer Science. The Joint Institute for Computational Sciences at the University of Tennessee and the Oak Ridge National Laboratory seeks to fill a tenure-track position in the area of mathematics or computer science to begin August 2004. Applications from individuals whose research focuses on the use of high-end computers to advance the frontiers of science and engineering will be considered. Required qualifications for the position include a Ph.D. and relevant postdoctoral experience, evidence of significant scientific productivity, and a commitment to an integrated program of teaching and research. The successful candidate will be qualified for appointment in the Department of Computer Science or Mathematics. Interested candidates should submit a resume and a description of proposed research program. Applicants should also arrange for at least three letters of reference to be submitted. Review of applications will begin upon receipt and will continue until the position is filled. All application materials should be sent to the address below or to mathcompscisearch@jics.utk.edu. Mathematics/Computer Science Search Committee, Joint Institute for Computational Sciences, University of Tennessee-Oak Ridge National Laboratory, P. O. Box 2008, Oak Ridge National Laboratory, Oak Ridge, TN 37831-6008. The University of Tennessee is an EEO/AA/Title VI/Title IX/Section 504/ADA/ADEA institution in the provision of its education and employment programs and services.

SENIOR WEB ANALYST/PROGRAMMER (Unisys/Chicago). Analyze & evaluate majorweb application project requirements. Review user requirements & providedirection in identification of

ASU
ARIZONA STATE UNIVERSITY

Department of Computer Science and Engineering

The Computer Science and Engineering Department in the Ira Fulton School of Engineering at Arizona State University has one faculty position, in the area of Bioinformatics, open at the tenure-track assistant professor level and invite outstanding candidates to apply. The expected starting date for this position is August 16, 2004. Applicants are required to have completed their Ph.D. in computer science, computer engineering, computational biology, bioinformatics, or a closely related field by the appointment date. Applicants must show exceptional promise in bioinformatics research and teaching, evidence of interdisciplinary research potential, and a commitment to quality teaching. Desired areas of interest for the position include genome-sequence (DNA or protein) based informatics, algorithm design, theory, data management and data mining, machine learning, statistical analysis, and related areas. The successful candidate will be expected to establish an extramurally funded research program and participate in the interdisciplinary research, teaching, and training initiatives in the Department of Computer Science and Engineering (http://cse.asu.edu) and the Center for Evolutionary Functional Genomics (http://www.azbio.org/efg).

ASU is a major research university widely recognized as a rapidly emerging educational institution in the US. The main campus is located in the city of Tempe, in the metropolitan Phoenix area. This year the School of Engineering received a $50 million gift, which will provide funding for scholarships, fellowships, and research programs. The Computer Science and Engineering Department provides a stimulating and fast-growing environment of research and teaching, with ample opportunities for partnerships with high-technology industry and with emphasis on quality, leading-edge graduate and undergraduate education. The department is in active collaboration with various interdisciplinary research centers, such as AZBioDesign institute, a hub for biotechnical and biodesign research in central Arizona, building collaboration networks among scientists and clinical researchers from leading industries and institutions. AzBioDesign Institute aims to promote multidisciplinary investigations in basic science and engineering required to design the critical biotechnology solutions of the 21st century. The successful candidate will join a thriving local genomics research community that includes the Translational Genomics Research Institute (TGen) and the International Genomics Consortium.

Initial closing date is March 1, 2004 and applications received by this date will receive full consideration, if not filled be reviewed bi-weekly after this date until the search is filled. Application packages must include a cover letter, detailed curriculum vitae, research and teaching statements, hard copies of the most important publications, and the names and addresses of four references.

These packages must be sent by regular post to

Chair of Faculty Search Committee
Department of Computer Science & Engineering
Ira A. Fulton School of Engineering
Arizona State University
Tempe, AZ 85287-8809

ASU is an equal opportunity, affirmative action employer.

problem & potential resolution. Provide analytical & tech support in conceptualization, design, development & implementation of complex, multiple interlinked web applications. Define system/application objectives & prepare design specs to meet user requirements & satisfy interface problems, involved in a relational d/base mgmnt systems & on-line transaction processing. Recommend corrections in technical applications & analysis. Contribute to the development of new technical principles & concepts. Reqs. Bachelor Degree or equiv in comp sci +3 yrs exp in job offered or 3 yrs exp as a Sr. Programmer/Web Analyst. Exp in deploying full life cycle client/server internet/intranet applications. Exp must also incl web server admin, maintenance & security. Must also have exp using ASP. Must have strong knowl in multi-d/base environmentincl Oracle & MS SQL server. 40 hrs/wk (9-5). Salary commensurate w/exp. Send resume to IEEE Computer Society, 10662 Los Vaqueros Circle, Box# COM1, Los Alamitos, CA 90720.

GJØVIK UNIVERSITY COLLEGE, Norway. Professor/Associate Professor, Information Security, Department of Computer Science and Media Technology, http://nislab.hig.no/People/Jobs/.

UNIVERSITY OF CENTRAL FLORIDA, Director of the School of Computer Science (http://www.cs.ucf.edu/). The University of Central Florida (UCF) invites applications and nominations for the Director of the School of Computer Science. The University is strongly committed to computer science, designating it a key discipline and allotting the School $2.5 million per year in special recurring funds. In 2004, UCF will break ground on a new state-of-the-art 75,000 sq. foot building to maintain and expand the excellence of the School of Computer Science. The Director has a unique opportunity to continue the School's rapid growth in ranking and reputation through recruitment of faculty members at all levels, and by helping to develop and execute the School's growth plan. The University of Central Florida has over 42,000 students and is among the nation's fastest growing universities. We are located in Orlando, Florida at the center of the Florida High Tech Corridor with a thriving industrial base in telecommunications, computer systems, semiconductors, defense and space, lasers, simulation and software, and the world renowned entertainment industry. The University offers $2 million per year in matching funds for joint research projects with regional high-tech industry and government organizations. In addition, there is a University-sponsored incubator to foster high-tech entrepreneurship. The Orlando metropolitan area enjoys an exceptional climate with rapid access to the Atlantic seashore and the Gulf of Mexico and benefits from a major airport with numerous direct international and national connections. The School of Computer Science has 39 faculty members covering all traditional areas in computer science as well as newer areas such as digital media and bioinformatics. We have over 270 graduate students (approximately 130 Ph.D. students) and 860 undergraduate CS majors. In addition, we offer a B.S. degree in Information Technology (IT), with over 400 undergraduate majors. We have a rapidly growing research program with over $2.8 million in annual expenditures from grants and contracts. Our Federal research sponsors include NSF, ARO, ONR, NASA, PEOSTRI, RDECOM, the Department of Transportation, and other agencies of the Department of Defense. We also enjoy the support of numerous industrial sponsors including established companies as well as local hi-tech startups. Please see our web page at http://www.cs.ucf.edu/ for more information. A Director candidate must have (1) an earned doctorate in computer science or related discipline; (2) a record of scholarly activity, including refereed publications and external funding, commensurate with the rank of tenured full professor. Preference will be given to candidates who meet these qualifications and further demonstrate: (1) leadership and vision in their field; (2) experience partnering with industry and other external constituencies; (3) strong administrative, planning, and decision-making capabilities; and (4) an ability to foster an environment of collegiality while motivating the School toward continued excellence. Candidates should be accomplished scholars in their field with a firm commitment to academic excellence. We seek outstanding applicants from academia, industry, and government. Salary will be highly competitive. Applications should include a cover letter addressing the required and preferred criteria, a detailed CV including publications, research grants, patents, and students supervised, and names and addresses of at least three references. Screening begins upon receipt of application. Applications will be accepted until the position is filled. All materials should be sent to: Dr. M.J. Soileau, Chair, Director of School of Computer Science's Search Committee, UCF Office of Research, 12443 Research Parkway, Suite 302, Orlando, FL 32826. (407) 823-5538. email: director-search@mail.ucf.edu. The University of Central Florida is an Equal Opportunity/Affirmative Action employer. Women and minorities are particularly encouraged to apply. As an agency of the State of Florida, UCF makes all application materials and selection procedures available for public review.

WRIGHT STATE UNIVERSITY, Lecturer Positions. The Department of Computer Science and Engineering at Wright State University seeks applicants for non-tenure track positions as Lecturer. The Department resides in the College of Engineering and Computer Science and offers B.S., M.S. and Ph.D. degrees both in Computer Science and Computer Engineering. Candidates for these positions are expected to have an earned Ph.D. in computer science, computer engineering, or a closely related area. We will also consider candidates with a graduate degree in CS or CE, who have extensive experience in a closely related field. Successful candidates will be expected to participate fully in the life of the university through teaching and service. Teaching experience and the ability to contribute to the research mission of the department would be desirable. All high quality applicants will be considered regardless of their field of specialization. Wright State University is an AA/EEOC and has a strong institutional commitment to diversity. Therefore, we are particularly interested in receiving applications from a broad spectrum of people, including underrepresented minorities, women, persons with disabilities and veterans. Applicants should provide a brief statement of their teaching experiences and research interests. They should include a complete vita with names, addresses, telephone numbers and e-mail addresses of at least three references, documentation of teaching abilities, plus any additional supporting information. Salaries are highly competitive. Address applications and supporting information to: Chair,

Memorial University of Newfoundland

MEMORIAL UNIVERSITY OF NEWFOUNDLAND

Department of Computer Science

2 Tenure-Track Faculty Positions

The Department of Computer Science at Memorial University of Newfoundland invites applications for two tenure-track positions at the Assistant or Associate Professor level. Candidates with a strong background in interdisciplinary work and from all areas of application of Computer Science are invited to apply.

Information on the application procedure and required qualifications may be found at **http://www.cs.mun.ca/news/**

Review of applications will begin March 31, 2004, and continue until suitable candidates have been identified. Memorial University is committed to employment equity and encourages applications from qualified women and men, visible minorities, aboriginal people and persons with disabilities. All qualified candidates are encouraged to apply; however, Canadians and permanent residents will be given priority.

Lecturer Search Committee, Department of Computer Science & Engineering, 3640 Col. Glenn Hwy, Wright State University, Dayton, OH 45435. Consideration of candidates starts April 1, 2004 and continues each month until June 1, 2005, or until the positions are filled. For details and information you may call (937) 775-5134 or contact Forouzan Golshani, NCR Distinguished Professor and Chair, at golshani@cs.wright.edu. Wright State University is an equal opportunity/affirmative action employer.

NETWORK ADMINISTRATOR/DESKTOP SUPPORT SPECIALIST. Qualified prof'l needed to admin. Network & computer systems for NYC based Co. Must install, configure & support LAN and WAN; maintain file and WEB servers and clients; monitor network & systems to ensure viability; perform req'd. maint. to support network availability; and provide desktop support to office users. Must have B.S. degree or equivalent in computer science or related field, MCSE certification, plus 4 yrs experience. Send resume to Weinstein & Holtzman, Inc., 96 9th Street, Brooklyn, NY 11215.

UNIVERSITY OF CALIFORNIA AT DAVIS, Edward Teller Professor of Applied Science, Department of Applied Science. The Department of Applied Science at the University of California, Davis, invites applications from, and nominations of, outstanding individuals as candidates for the Edward Teller Professor of Applied Science. The department is seeking a senior distinguished scholar with an outstanding record in physical or applied science and a strong commitment to the education and teaching mission of the department. Apart from two undergraduate majors, Optical Science and Engineering and Computational Applied Science, the department has a strong diverse graduate program in applied physics, computational science, optics and lasers, plasma physics, and materials science, reflecting the academic backgrounds of the faculty. The graduate program has traditionally interacted closely with the University of California managed National Laboratories, particularly the Lawrence Livermore and Lawrence Berkeley National Laboratories. According to the endowment of the Professorship, the Edward Teller Professor must also be the Chair of the department. Thus, the Edward Teller Professor will relinquish the title of the endowment, becoming a regular faculty member, when stepping down as department Chair. The service of department Chair is usually for a period of 3-5 years. The position is open until filled. For full consideration, applications should be received on or before March 20, 2004. Applications, nominations, and inquiries should be directed to Prof. Niels Grønbech Jensen, Chair of the Edward Teller Professor Search Committee, Department of Applied Science, University of California, One Shields Avenue, Davis, CA 95616-8254, U.S.A. A complete application consists of a Curriculum Vitae with a list of at least 5 professional references, a statement of intended professional directions and goals for how the Edward Teller Professorship can be used to enhance the Department. The University of California is an affirmative action/equal opportunity employer.

THE OHIO STATE UNIVERSITY, Columbus, Chair, Department of Electrical and Computer Engineering. The Department of Electrical and Computer Engineering at The Ohio State University invites applications and nominations for the position of Chair, which becomes available July 2004. Consistently ranked in the top 10 percent nationally, the Department of Electrical and Computer Engineering at The Ohio State University enjoys the tradition of an excellent reputation among electrical and computer engineering programs internationally. The department consists of 49 faculty members, including three members of the National Academy of Engineering, an Oldenberger Medal winner, and 16 Fellows of the Institute of Electrical and Electronics Engineers (IEEE). We also have several faculty members who are Fellows of other learned societies. Our faculty members have received the Presidential Early Career Award for Scientists and Engineers, the NASA Distinguished Public Service Medal, Presidential Young Investigator Awards, Office of Naval Research Young Investigator Award, and National Science Foundation CAREER Awards. OSU faculty have achieved international recognition for research in analog-digital integrated circuit design, communications, computer engineering, computer networks, computer vision, control, electromagnetics, electronic materials, high performance computing, optics, power engineering, robotics, signal processing, transportation, and wireless systems. Our strategic plan sets an ambitious but realistic agenda for the future with thrusts in advanced electronic materials and nanoscale devices, computational and distributed intelligence, and biotechnologies. Central to this plan are initiatives in multidisciplinary research; the department aggressively pursues these opportunities with, for example, the OSU Center for Materials Research, the College of Medicine, the Ohio Supercomputer Center, the Department of Biomedical Informatics and the Biomedical Engineering Center, and the Center for Automotive Research. It is expected that the next Chair will work to support and enhance interdisciplinary activities such as these. The successful applicant will present leadership skills and an extensive record of accomplishment in research and scholarship. The Ohio State University is committed to excellence in undergraduate and graduate education, and to diversity. Applications should include a curriculum vita, a statement of the candidate's vision for the future of ECE research and education, a self-assessment of leadership qualities and style, and the names of five professional references, all in the form of Word or pdf attachments (only) to the search committee chair, Prof. Kim L. Boyer, at: chairapp@ee.eng.ohio-state.edu. Review of applications will begin immediately and continue until the position is filled. The Ohio State University is an Equal Opportunity/Affirmative Action Employer; applications from women and other underrepresented groups are expressly encouraged.

THE UNIVERSITY OF TENNESSEE, The Imaging, Robotics, and Intelligent Systems (IRIS) Laboratory. The IRIS Lab invites applicants for multi-year Research Assistant/Associate Professorships and Ph.D. Fellowships. The IRIS Lab's emphasis is in the fields of Three-dimensional Imaging, Data Fusion, and Visualization. For 2004, the IRIS Lab is expected to have a staff of 50 and an annual budget over $3.5Million. Interested persons should contact: Mongi Abidi, Professor and Associate Department Head, Department of Electrical and Computer Engineering, 328 Ferris Hall, Knoxville, TN 37996-2100. Web: http://imaging.utk.edu/opportunities/opportunities.htm, Phone: 865-974-5454, Fax: 865-974-5459, E-Mail: abidi@utk.edu. UTK is an EE/AA/Title VI/Title IX/Section 504/ADA/ADEA Employer.

GRAPHIC ART CO. seeks animator for motion pictures, TV commercials, interactive CD-ROMS & architectural renderings. Knowledge of 3d Studio Max, Maya, PhotoShop, CorelDraw, Illustrator,Freehand,QuarkExpress & PageMaker program reqd. BA in Computer Animation. Comp. salary. Mail/fax resume to Terragenesis LLC, 3550 Biscayne Blvd. Ste 307, Miami, FL 33137; fax 305-576-9080.

DREXEL UNIVERSITY, Head, Department of Computer Science, College of Engineering. The Department of Computer Science seeks nominations and applications for the position of Department Head. The Department offers an ABET accredited undergraduate program in computer science, as well as a new program in software engineering. It is one of six departments within the College of Engineering, which is well known for its commitment to graduate research and cooperative education and for its leadership in developing innovative approaches to computer science education. Drexel's Computer Science Department has rapidly expanding graduate research and education programs in Software Engineering, Graphics, AI, Networks, Information Assurance and Security, HCI, and Scientific Computing. We are seeking a dynamic individual with outstanding leadership and administrative skills, proven abilities in research and teaching, and a doctorate in computer science or a closely related field. The candidate should have a strong interest in leading development activities, and in interacting with government, industry, and professional societies. Further, the candidate should be able to work with a diverse group of faculty, staff, students and administrators to achieve common goals and to maintain rapport with alumni and industrial representatives. Consideration and review of applications begins immediately and is expected to continue until the position is filled. To be considered, please send cover letter, CV along with names and contact information for at least six references to: Department Chair Search Committee, College of Engineering, Drexel University, Philadelphia, PA 19104. Email: search_chair@coe.drexel.edu. Electronic submissions to search_chair@coe.drexel.edu in PS/PDF format are preferred. Drexel University provides affirmative action and equal opportunity in employment for all qualified persons regardless of race, color, religion, national origin, age, sex, disability or veteran status. Each student, faculty or staff, interacting with any part of the University will be treated equally, and will be given the same opportunities as provided to any other person in similar circumstances. Drexel University is an Equal Opportunity / Affirmative Action Employer. For more information: Department of Computer Science http://www.cs.drexel.edu, College of Engineering http://www.coe.drexel.edu, Drexel University http://www.drexel.edu.

UNIVERSITY OF NEBRASKA AT OMAHA, Faculty Position in Bioinformatics. The University of Nebraska at Omaha (UNO) announces a tenure-track faculty position in Bioinformatics starting August 2004. The position is offered jointly by the Department of Biology and the Department of Computer Science. A Ph D. degree in a Bioinformatics or a related field is required. The successful candidate is expected to develop an active research program and to teach courses in Bioinformatics-related areas. Substantial experimental and computational resources are available in the two departments. These resources have been greatly upgraded after the development of the Peter Kiewit Institute, a new 193,000 sq. foot, $37.5M state-of-the-art facility for teaching and research. Bachelors, masters and Ph.D. degrees in areas related in Bioinformatics are presently available at UNO. The university has a strong commitment to achieving diversity among faculty and staff. We are particularly interested in receiving applications from members of under-represented groups and strongly encourage women, disabled, minorities and Vietnam-era veterans to apply for this position. The University of Nebraska at Omaha is an Affirmative Action/Equal Opportunity Employer. Review of applications will begin immediately and will continue until the position is filled. Send CV, statements of research and teaching interests and objectives, and 3 letters of recommendation to: Hesham H. Ali, Bioinformatics Search Committee, Department of Computer Science, PKI 172C, University of Nebraska at Omaha, 1110 South 67th Street, Omaha, NE, 68182-0116. Electronic applications can be sent to hesham@unomaha.edu. For more information related to the position, check the web site: http://www.ist.unomaha.edu/employment/bioinfoposition.php

WRIGHT STATE UNIVERSITY, Assistant & Associate Tenure Track Faculty Positions. The Department of Computer Science and Engineering at Wright State University seeks applicants for tenure track Assistant or Associate Professor positions. The Department resides in the College of Engineering and Computer Science and offers B.S., M.S. and Ph.D. degrees both in Computer Science and Computer Engineering. Candidates for these positions are expected to have an earned Ph.D. in computer science, computer engineering, or a closely related field and evidence of scholarship in that field appropriate to the position. Applicants for associate professor positions should have a distinguished record in computing, demonstrating strong leadership in both research and teaching, commensurate with the rank. Successful candidates will be expected to participate fully in the life of the university through research, teaching and service. There is specific interest in faculty specializing in distributed and embedded computing, hardware-software co-design, software engineering, and bioinformatics; however, all high quality applicants will be considered regardless of their field of specialization. The Computer Science and Engineering Department is one of four departments in the College of Engineering and Computer Science. The Department currently has 24 faculty members, 580 undergraduate majors, and more than 150 M.S. and 35 Ph.D. students. It is housed in an attractive engineering building with fully networked Unix and Windows environments, an Origin 2000 supercomputer, an NCR4850 Teradata machine, and excellent research laboratories. Wright State University, an institution of 16,000 students, is located on a spacious campus with a significant area of protected green space in a growing high-technology suburban community. It is surrounded by commercial (NCR, Lexis-Nexis, Reynolds & Reynolds, Mead-Westvaco, etc.) and government (WPAFB) research and development facilities. The University is proactively committed to industrial and government partnerships for research and economic development ventures and has established the Information Technology Research Institute closely aligned with the CS&E Department. In addition, the university has just won an $11M award from the State of Ohio to establish a Wright Center of Innovation for Advanced Data Management and Analysis which will be housed in a new building adjacent to the College of Engineering and Computer Science. The Department also receives special support for enhancement of its graduate program from WSU, and the Ohio Board of Regents. Last year, the Department received $2.7M in externally sponsored research funding. A variety of affordable and pleasant living environments, schools and parks, attractive to professionals and families, are conveniently located close to the campus. Wright State University is an AA/EEOC and has a strong institutional commitment to diversity. Therefore, we are particularly interested in receiving applications from a broad spectrum of people, including underrepresented minorities, women, persons with disabilities and Veterans.Applicants should clearly indicate the rank for which they are applying, and should provide a brief statement of their research and teaching interests and goals. They should include a complete vita with names, addresses, telephone numbers and e-mail addresses of at least three references, documentation of teaching abilities, plus any additional supporting information. Salaries are highly competitive. Address applications and supporting information to: Chair, Tenure Track Faculty Search Committee, Department of Computer Science & Engineering, 3640 Col. Glenn Hwy, Wright State University, Dayton, OH 45435. Consideration of candidates starts April 1, 2004 and continues each month until June 1, 2005, or until the positions are filled. For details and information you may call (937) 775-5134 or contact Forouzan Golshani, NCR Distinguished Professor and Chair, golshani@cs.wright.edu. Wright State University is an equal opportunity/ affirmative action employer.

NATIONAL TAIPEI UNIVERSITY, TAIWAN, College of Science and Technology has several faculty positions at the levels of professors, associate professors, and assistant professors in the beginning academic year 2004. We will have additional openings for each of the following three years. Candidates for Department of Computer Science and Institute of Communications Engineering must have a Ph.D. degree in related areas from an accredited university before the appointment is in effect. Candidates with specialization in the Computer Science and Communication areas are welcome to submit their curriculum vitae and 2 recommendation letters by March 1, 2004. Please send all information to: Professor Tong-Ying Juang, Searching Committee for College of Science and Technology, National Taipei University, No 67, Section 3, Ming-Sheng East Road, Taipei, Taiwan 104, Fax: +886-2-25015974, Email: juang@mail.ntpu.edu.tw.

BOOKSHELF

Analyzing Neural Development Mathematically

Modeling Neural Development, Arjan van Ooyen, editor. This collection looks at neural development using computational and mathematical modeling, which provides a tool for precise expression that augments other research modes. The book shows how models can be used to study the nervous system's development at different levels of organization and different phases of development, from molecule to system and from neurulation to cognition.

The book loosely follows the chronology of neural development, starting with the very early development of the nervous system, including gene networks, cell differentiation, and neural tube development. Next, the text examines neuronal morphogenesis and neurite outgrowth as well as different aspects of the self-organization of neurons into networks.

The book's final sections cover refinement of connectivity and the development of specific connectivity patterns, then focus on some of the functional implications of morphology and development.

MIT Press; mitpress.mit.edu; 0-262-22066-0; 336 pp.; $60.

WRITING THE RIGHT REQUIREMENTS

Writing Better Requirements, Ian F. Alexander and Richard Stevens. Experience has shown that investment in the requirements process saves time, money, and effort. Yet, development efforts still consistently charge ahead without investing sufficiently in the requirements process. Developers intent on building technical solutions are unwilling to take the time and effort to understand and meet real customer needs.

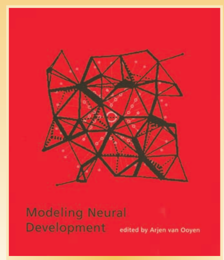

Readers involved in the systems engineering process in any industry can learn how to write requirements that guarantee the final system will have the functionality they need. This book shows how to write simple, clear requirements; organize those requirements into scenarios; and review the requirements to ensure they accurately reflect true user needs.

Addison-Wesley; http://www.awprofessional.com; 0-321-13163-0; 176 pp.; $36.99.

ACCELERATING BLUETOOTH DEVELOPMENT WITH JAVA

Bluetooth Application Programming with Java APIs, C. Bala Kumar, Paul J. Kline, and Timothy J. Thompson. Adoption of Bluetooth wireless technology has made great strides recently. This book explains in detail a major leap forward—standardization of Java APIs for Bluetooth wireless technology.

The authors describe how to write Bluetooth applications using Java APIs to exploit the power of both technologies. The book provides explanations and concrete examples for implementing Java APIs in Bluetooth applications. It also covers the programming areas necessary to successfully design and build these applications, including RFCOMM, OBEX, device discovery, service discovery, and L2CAP.

Morgan Kaufmann; www.mkp.com; 1-55860-934-2; 500 pp.; $44.95.

EVALUATING WEB SITES AUTOMATICALLY

Automated Web Site Evaluation: Researchers and Practitioners' Perspectives, Melody Y. Ivory. Automated evaluation has gained prominence as a way to help Web site builders produce more usable and accessible sites. But how effective are automated evaluation approaches and tools? Where do they fit into Web practitioners' work practices?

This book offers an in-depth look at automated Web site evaluation methodologies and existing software tools from both the researcher and practitioner perspectives. The author documents the field's current state and highlights promising research and tool developments. She summarizes surveys and analyses of automated Web site evaluation methodologies and tools; describes Web Tango, an automated evaluation approach she developed; and presents results from automated evaluation tools studies and the Web site design process itself.

Kluwer Academic/Plenum Publishers, New York; www.wkap.nl/; 1-4020-1672-7; 912 pp.; $94.

Editor: Michael J. Lutz, Rochester Institute of Technology, Rochester, NY; mikelutz@mail.rit.edu. Send press releases and new books to *Computer*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; fax +1 714 821 4010; newbooks@computer.org.

Location-Aware Computing Comes of Age

Mike Hazas, Lancaster University
James Scott, Intel Research
John Krumm, Microsoft Research

At the core of invisible computing is *context awareness*, the concept of sensing and reacting to dynamic environments and activities. Location is a crucial component of context, and much research in the past decade has focused on location-sensing technologies, location-aware application support, and location-based applications. With numerous factors driving deployment of sensing technologies, location-aware computing may soon become a part of everyday life.

LOCATION-SENSING TECHNOLOGIES

A central problem in location-aware computing is the determination of physical location. Researchers in academia and industry have created numerous location-sensing systems that differ with respect to accuracy, coverage, frequency of location updates, and cost of installation and maintenance.

Coarse-grained systems

For applications in open, outdoor areas, the Global Positioning System is a common choice. A GPS receiver estimates position by measuring satellite signals' time difference of arrival. Although GPS offers near-worldwide coverage, its performance degrades indoors and in high-rise urban areas, and receivers have a relatively long start-up time and high cost.

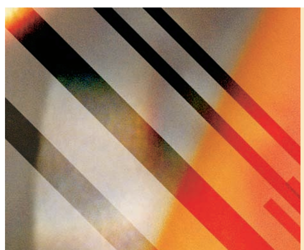

The widespread deployment of sensing technologies will make location-aware applications part of everyday life.

In 1989, Roy Want, Andy Hopper, and others pioneered the study of indoor location sensing with their infrared-based Active Badge system. This provides room-grained location using wall-mounted sensors that pick up an infrared ID broadcast by tags worn by the building's occupants.

Many of the location-sensing systems developed since then are based on radio. By using base station visibility and signal strength, it is possible to locate Wi-Fi-enabled devices with accuracies from several meters to tens of meters. Bluetooth technology, which offers a shorter range than Wi-Fi, can give more accurate positioning, but at the expense of requiring more fixed base stations to provide coverage. Inexpensive radio-frequency identification tags can be used for location determination as well by placing RFID readers at doorways and other strategic points to detect the passage of people or objects.

Location information can also be derived from other types of RF infrastructures, including those for mobile phones and TV broadcasts. These can be deployed over a wide area with relative ease, in contrast to technologies such as RFID that have limited transmission range. With mobile phones, Cambridge Positioning Systems has demonstrated location accuracies of 20 meters, while Rosum has achieved accuracies from 3 to 25 meters with digital TV signals.

Fine-grained systems

Many of the above systems are based on technologies that were not developed with location sensing in mind. Perhaps as a consequence they exhibit modest accuracy, generally measured in meters. However, at least three types of systems have been designed specifically to provide fine-grained location sensing, achieving accuracies on the order of centimeters.

Ultrasound can be used to determine distances between mobile tags and known points in the environment. A process akin to triangulation can then be employed to derive a location estimate for the tag. One type of ultrasonic ranging device is the Cricket indoor location system developed at MIT, which is set to become available for purchase from Crossbow this year.

Some computer vision-based systems are appealing because they do not require users to wear any sort of tag. However, such systems have difficulty identifying and simultaneously tracking many subjects. Vision-based systems using barcode-like tags tend to be more robust.

Ubisense, a company that builds real-time local positioning systems,

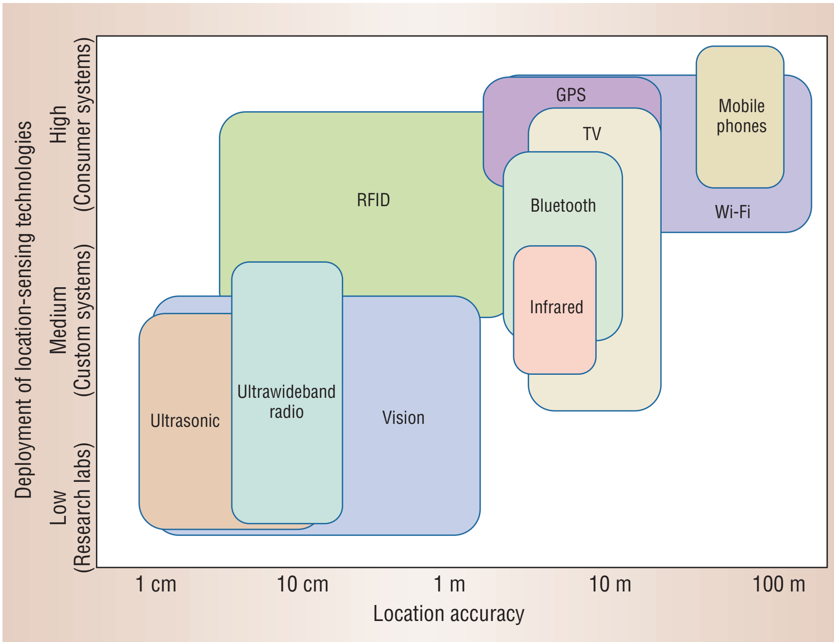

Figure 1. Location-sensing technologies. Each box's horizontal span shows the range of accuracies the technology covers; the bottom boundary represents current deployment, while the top boundary shows predicted deployment over the next several years.

recently demonstrated a fine-grained tracking system that uses ultrawideband radio signals. Unlike conventional radio signals, these signals can have pulse durations short enough to allow accurate time-of-arrival and angle-of-arrival measurement, with an accuracy of about 15 centimeters. In addition, ultrawideband technology does not require a direct line of sight between tags and sensors.

DEPLOYMENT

Figure 1 shows the current and predicted deployment of location-sensing technologies within the next two to three years. The widest existing deployments are based on GPS, which is particularly suited for outdoor applications. These include servicing applications centered on vehicle location such as route planning and fleet tracking, as well as applications integrated into handheld GPS units.

Other current deployments are found in vertically integrated solutions and comprise a specific location-aware application, appropriate location-sensing hardware, and a custom software platform. A handful of firms offer these systems in targeted application areas such as military training, human-body motion capture, supply chain management, and asset tracking.

Looking ahead, numerous factors are accelerating the adoption of coarse-grained location-sensing technologies. To begin with, the recent explosion of Wi-Fi, Bluetooth, and other wireless networking technologies has led to many end-user devices being equipped with RF hardware that can be used for location sensing.

In addition, the Enhanced 911 requirement—which mandates that US wireless carriers provide location accuracies of 50 to 100 meters for emergency 911 calls by the end of 2005—is driving incorporation of location-sensing systems into mobile phones using GPS, base-station triangulation methods, and a combination of these technologies known as Assisted GPS. Similar requirements exist in the European Union.

Another important factor is a directive by US retail giant Wal-Mart that requires its top 100 suppliers to include RFID tags in their products by January 2005. Current plans call for tagging items by the case, but this may be extended to the tracking of individual items in the future.

In addition to the deployment of sensing technologies, increased hardware and software support is available for location-aware applications. For example, mobile phones have emerged as a ubiquitously deployed computing platform capable of downloading and running location-aware applications. Also, Microsoft recently revealed that its next PC operating system, code-named Longhorn, will include location-aware software components.

ABSTRACTING LOCATION

To provide support for a variety of location-aware applications, researchers are working on techniques for fusing data from multiple sensors, on methods for representing location data, and on drawing high-level contextual information from location data.

Sensor fusion

Vertically integrated location-aware systems typically use one type of sensor for a single application. However, in the near future, many kinds of location sensors may be available to a particular client system. The task of making sense of this vast amount of sometimes contradictory information, known as *sensor fusion,* presents a major challenge.

Borrowing from the field of robotics, location researchers have settled on Bayesian inferencing as the preferred method for processing data from disparate location sensors. Using Kalman filters, hidden Markov models, dynamic Bayes nets, and particle filters, they have developed principled methods of incorporating sensor uncertainty as well as limits on speed and travel paths.

The result is a location measurement derived from multiple sensors and constraints that uses a probability distribution rather than a single value to describe the inherent uncertainty. For example, researchers at the University of Washington have demonstrated an indoor location-measuring system that processes data from multiple sources, including infrared and ultrasonic sensors, using a particle filter. In addition, the system learns typical walking paths through the building to aid in location estimation.

Representation

Finding a way to represent location readings that facilitates storage, communication, and interpretation is also a challenge. Unfortunately, no single representation is useful in all circumstances.

Sensors tend to report locations as numerical coordinates such as (x, y, z), but semantic representation is more effective for application-level reasoning about locations. Semantic representations usually include a hierarchy of locations such as (building, floor, room) or (country, state, city), and, like numerical representations, also must include a notion of uncertainty.

Some applications, such as travel planners, require both types of representations: Map-like coordinates are needed to perform Euclidian distance calculations, while information about terrain, traffic patterns, customs procedures, and other factors is important to achieve accurate travel-time projections. Due to the varying types of measurements from different sensors and the varying requirements of different applications, no dominant representation for location has emerged.

Sensed locations are useless without a *location authority* that gives a space of possible locations and can respond to queries about distances, routes, and proximity. Example location authorities include WGS84 data for latitude/longitude and maps or floor plans in computational form.

Place and context

Simply knowing that a person is "at home," "in my office," or "in my car" is often sufficient for applications to carry out predetermined actions in a given situation, such as turning off a cell phone's ringer during a film or concert. In these cases, the person's relationship or interaction with a place is more important than the physical location.

Toward this end, location data can be used to infer higher-level contextual information. For example, University of Washington researchers have used a time sequence of GPS data to infer a person's mode of transportation—foot, bus, or car. Others working at MIT have developed a location-aware augmented-reality museum guide that classifies a visitor's movement patterns as "greedy," "busy," or "selective" and tailors its content accordingly.

Location-aware applications have been developed for a number of everyday scenarios.

APPLICATIONS

In addition to specialized applications such as military training and asset tracking, location-aware applications have been developed for a number of everyday scenarios:

- office applications such as nearest-printer services and mobile desktop control can increase workplace productivity;
- tour and museum guides can help people navigate an unfamiliar space;
- "locate my friends" utilities can be linked with instant messaging for social or business purposes;
- conference aids can track presentation attendance and facilitate note taking and discussion;
- medical facilities can track staff and monitor patients for emergency response; and
- home applications can help with household management and home entertainment, as well as aid the aged and disabled in performing everyday tasks.

While coarse-grained location-sensing deployments will enable some of these applications, others will require fine-grained location-sensing systems. The deployment of this additional infrastructure will rely on an adequate return on investment. Unlike some technologies, fine-grained location-aware computing does not appear to have a "killer app" to provide the required economic incentive. However, the combination of benefits that many applications provide might justify widespread adoption.

Advances in location-sensing technologies and factors promoting wide-scale deployment will soon make coarse-grained location information widely available. Recent achievements in sensor fusion techniques, location representation, and software support will facilitate the development of applications that can use this multitude of sensors. In time, fine-grained location systems and applications will become more economically viable.

For more information on location-aware computing, visit www.cambridge.intel-research.net/lac/links.php. Many of the research examples discussed in this article were presented at the 2003 Workshop on Location-Aware Computing, held as part of the 5th International Conference on Ubiquitous Computing in October 2003. ■

***Mike Hazas** is a research associate in the Computing Department at Lancaster University. Contact him by visiting www.comp.lancs.ac.uk/~hazas.*

***James Scott** is a researcher at Intel Research in Cambridge, United Kingdom. Contact him at james.w.scott@intel.com.*

***John Krumm** is a researcher at Microsoft Research. Contact him at jckrumm@microsoft.com.*

Editor: Bill Schilit, Intel Research Seattle; bill.schilit@intel.com

Interactivity, Interfaces, and Smart Toys

Lydia Plowman, University of Stirling
Rosemary Luckin, University of Sussex

Although computers can represent a medium for children's social and intellectual development, some researchers believe that using computers before age seven subtracts from important developmental tasks and other types of learning (J.M. Healy, *Failure to Connect: How Computers Affect Our Children's Minds—for Better or Worse*, Simon & Schuster, 1998). Those opposed to computers believe that computer-based activities are less effective in developing understanding and skills than are artifacts that young children can handle.

These anxieties extend to technologies such as smart toys. Some researchers (D.E. Levin and B. Rosenquest, "The Increased Role of Electronic Toys in the Lives of Infants and Toddlers: Should We Be Concerned?" *Contemporary Issues in Early Childhood*, vol. 2, 2001, pp. 242-247), for example, claim that electronic toys produce limited and repetitive interactions that might inhibit the healthy play and development of young children. They contend that playing with dolls, trucks, blocks, and similar toys encourages children to be the "creators and controllers of their play" and helps "parents play in imaginative, give-and-take ways with their infants and toddlers."

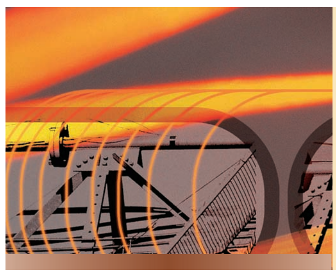

Despite some technical hurdles, using smart toys and software for educational play shows promise.

CACHET

Our recently completed research project, Computers and Children's Electronic Toys (www.ioe.stir.ac.uk/cachet/), investigated how children use smart toys. Cachet combines recent interest in mobile learning, tangible interfaces, and the home use of technologies. This research aimed mainly to explore interactivity and interfaces in the context of smart toys that children could use alone or in conjunction with a computer.

Toys

We based these toys on Arthur and his sister D.W., two aardvark characters from the Marc Brown stories and cartoon series (http://pbskids.org/arthur/) familiar to more than 75 percent of the children in our study. The toys have a plush finish, resemble traditional soft toys, and—unlike some smart dolls—their sensors and batteries do not make them heavy and unwieldy. Arthur and D.W. stand 60 cm tall and have a vocabulary of about 4,000 words, motors to provide movement, and electronic chips to recognize inputs. Because the toys cannot respond intelligently to spoken input, they depend on gestural interaction. If a child squeezes a toy's hand or wristwatch, the toy will ask questions. If a child squeezes the toy's toe, it will suggest a game.

Games include estimating a time—5, 10, 15, or 20 seconds—by squeezing the toy's hand when the time is up, saying the alphabet backwards and forwards, and tongue twisters. In addition to using them by themselves, children also can use the toys in conjunction with specially encoded CD-ROMs that feature language and number games.

Playing with the toy and the software simultaneously requires an accessory: a radio transmitter that looks like a modem and connects to the computer's game port. A PC pack add-on increases the toy's vocabulary to 10,000 words, letting the toy "talk" to children, comment on their interaction with the software, and offer advice and encouragement. Children can elicit help and information from the toy by squeezing its ear. In this mode, the child interacts with both the computer and the toy, while the toy interacts with the computer and mediates the child's actions.

If two children play with the toy, the interaction possibilities multiply. If the toy is not present, a clickable onscreen icon of Arthur or D.W. provides help and information. If children have difficulty with a game, the toy or icon reminds them of this help. We focused on these toys because they deliver the same help content through different mechanisms.

Conducting the study

We developed a multidisciplinary

approach to provide multiple response-and-behavior indices appropriate to the social and cultural context of using the toy while also allowing for investigation of detailed interactions. Given that children can play with Arthur and D.W. by themselves or in conjunction with the compatible CD-ROM, we looked at toy use alone, software use alone, and using the two together. Using a common core of data collection methods across all sites, we looked at toy and computer use in children's homes, after-school clubs, and classrooms.

Twelve children aged five to six participated in the home studies and were visited three times over approximately two weeks. Half the children received the toy first and were given the CD-ROM during the midway visit, while we reversed this order for the other half. In all cases, the children kept both items for the second week and could thus play with the toy and PC together.

Classroom use was more controlled, with detailed, dual-source video analysis of 32 children aged four to five. Twenty-two children in the four after-school clubs participated in the study. All children received verbal instructions on using the software and viewed a demonstration of how to access the help facility. In the classrooms and at the after-school clubs, the researchers remained in the background but prompted activity when children became stuck.

Analyzing the data

To capture the multiple interactions experienced by pairs of children using both the toy and software and either one or the other, we used the following main categories when transcribing the videos: action and gesture, expression and gaze, and dialogue. We further divided these categories into between children; between toy and children; between toy, onscreen software characters, and children; and between children and researcher. We used these structured transcripts to code categories of interaction and explore a range of questions about children's help-seeking preferences and social behaviors.

We compiled all data relating to the individual children in the homes and individuals or pairs in the after-school clubs to produce separate case studies. These studies contain personal details, field notes, complete video session transcripts, interviews, the diaries, test results, and the rating scale. This information let us analyze data for individuals and across children while building a detailed picture that includes the context of use, individual differences, patterns of play, and toy and software use.

Half the parents claimed the toys provided some educational value and said they would buy one at an affordable price.

RESULTS

Although the study focused on one particular type of toy, our approach has proven valuable in laying the groundwork for future research.

Toy characteristics

Interview questions revealed a considerable diversity of views about the extent to which children attributed human characteristics to the toys. Some knew batteries powered the toy, but younger children tended to think the toy had feelings and could think and talk by itself.

Invited to suggest ways in which the toys could be improved, several children mentioned that they would like it to be able to walk, perhaps because this would make the toy more lifelike. Although some children seemed to imbue the toy with sentient qualities, parents' comments suggest that this did not translate to a greater degree of dialogue or other forms of interaction. Most children viewed the toy as just that and some preferred to play with it switched off, taking the toy to the dinner table or making a bed for it next to the child's own.

At the first visit, parents tended to find the toys scary or feared they would stifle imaginative play. Yet by the end of the loan period, half the parents claimed the toys provided some educational value and said they would buy one at an affordable price. Although parents found the ways in which the toy and software interacted impressive, children seemed to take it for granted.

In homes, the time spent playing with the toy or software decreased as the novelty diminished—from about 45 minutes per day to less than 15 minutes, with interest in the toy wearing off faster than interest in the software. Those who received the toy after receiving the software played with it less frequently, if at all, partly because the help features were not generally needed—the software had already been explored. However, in this sequence we also found that children viewed the toy mostly as an adjunct to the software and rarely played with it away from the PC.

Children sometimes placed the toy on a table, next to the computer screen, even though the radio transmitter has a four-meter range. Although this made the toy seem more like a peripheral than an interaction partner, placing the toy on the child's lap could interfere with mouse manipulation. Paired children avoided this problem because one would typically control the mouse and the other the toy.

Interacting with toys and software

Although the toys can verbally interact with the children at a basic level, the spoken interaction could detract from the possibility of extended child-toy interaction and role-playing. Initially intriguing, the toy's vocabulary presents only an illusion of reciprocity and seems too limited to imply real personality. Most children found the toy's talking monotonous or irritating and eventually switched it off.

We found no evidence that these toys make either a beneficial or detrimental difference to the children's ability to engage in child-led imaginative play.

The children enjoyed the tactile nature of the toys more than the toy's interactivity, which challenges the view that technological toys are psychologically damaging.

Detailed video analysis revealed that young children can make the connection between two different interfaces and coordinate the experience received through their convergence. Our evidence suggests that children as young as age four are not disconcerted when faced with feedback and interaction possibilities from different artifacts.

Many children required assistance from the researcher or a peer to elicit help from the toy or onscreen icon. Some children also ignored the help that Arthur or D.W. provided. Even when they took notice of the help prompts, however, they did not necessarily interpret them correctly.

Help features

We analyzed the toy's role in supporting the child's learning in terms of *scaffolding*: how a more knowledgeable partner can assist the cognitive development of a less able one and gradually foster the development of successful independent task performance. Considerable research has explored this concept in relation to software, but with emphasis on desktop computers. We extended this research by examining how children requested and used assistance from the toy, the accompanying software, or other people.

To understand how children use the help available to them, we combined the descriptive results across contexts and with detailed activity analysis from the school studies. Initially, children were more likely to seek help from human companions. Although they often failed to notice the unsolicited clues the toy or onscreen icon was giving, when prompted by their human companion, they became competent at using the toy to elicit hints and encouragement. The children were discerning users and recognized the help content's questionable value.

Girls' attraction to the toy could help redress an imbalance in the greater appeal of technology to boys than girls.

The toy's presence showed a statistically significant increase in incidences of help being successfully implemented, with children actively seeking help and adults actively engaged in prompting or assisting the children in their software use. However, the toy and no-toy conditions made little difference in how often children refused or ignored the offered help.

These findings suggest that system developers interested in software scaffolding might reap benefits from considering tangible interfaces instead of screen-based ones. The onscreen icon and the toy provided the same feedback content, with only the method for invoking and delivering this feedback differing.

While all the children in our study enjoyed using the software, interest in the toy appeared to be age-related, with parents reporting much more interest from younger children. Their poor feedback makes these toys unimpressive interaction partners. Nevertheless, the technology has potential, and generalizing from findings relating to the specific toys in the study suggests several areas for future development.

Using existing work on software scaffolding could improve the type and mode of feedback by linking it to the children's performance. Because the youngsters were more likely to interact with each other or with the researcher when the toy was present, tangible interfaces offer promise for improving interaction between peers.

Children or their families failed to discover some of the toy's functionality, such as the alarm clock. This shows the tradeoff between a toy being complex enough to maintain interest and simple enough for use by very young children.

In preschool settings and the early years of primary school, the computer is usually used as a free play activity without the benefit of adult mediation. There may be value in developing this technology for such circumstances, but this would require close analysis of the contingent help adults can provide. Likewise, the plausibility issue must be overcome. Finally, girls' attraction to the toy could provide an avenue for redressing an imbalance in the greater appeal of technology to boys than girls. ■

Lydia Plowman is Reader in Education and Director of Postgraduate Research in the Institute of Education, University of Stirling. Contact her at lydia.plowman@stir.ac.uk.

Rosemary Luckin is director of the Human Centred Technology Research Group at Sussex and the Interactive Digital Educational Applications Research Lab. Contact her at rosel@sussex.ac.uk.

Editor: Michael Macedonia, Georgia Tech Research Institute, Atlanta; macedonia@computer.org

Simplifying Public Key Management

Peter Gutmann, University of Aukland

Many security protocols in use today were designed under the assumption that some form of global distributed public-key infrastructure would eventually emerge to address key management problems. These protocols go back to the early 1990s, when a universal PKI was thought to be just around the corner. Ten years later, it's still just around the corner, and it probably always will be.

Consequently, existing protocols originally designed to rely on a global PKI must either employ ad hoc solutions or use any public key that turns up, because the only alternative is not to use any keys at all. In the absence of a PKI, system administrators can incorporate alternative approaches that are easy to use, transparent to end users, and have a low unit cost.

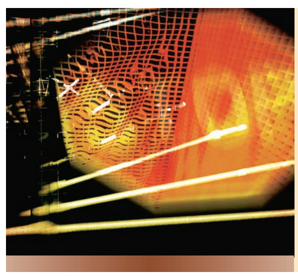

Alternative approaches to public key management offer ease of use, transparency to end users, and low unit cost.

KEY CONTINUITY

A store close to where I work assumes a new identity every year or two. Presumably this is done for accounting or tax purposes, but it's the same store, run by the same people, and that's what matters to customers—not what's on the sign out front.

The same principle applies to key management. If I used a given public key to access a file server, mail server, online store, or bank last week, all I need to know is that I'm dealing with the same entity this week. I don't need third-party attestation for this; I can verify it myself by comparing last week's key to this week's one. Key continuity is thus the basis for public-key management: Once you have a known good key, you can confirm a remote party's identity by verifying that it's still using the same key.

MITM attacks

The problem with this approach is that it's vulnerable to man-in-the-middle attacks by adversaries who can intercept or even modify the traffic passing between two communicating parties. A weak tie between a key and its owner invites MITM attacks, which succeed if the system can't distinguish communications with an intended recipient from those with the intervening attacker.

That's the theory. In practice, as Ian Grigg notes, the magnitude of the MITM threat is usually greatly exaggerated (www.iang.org/ssl/). For example, not a single instance of an MITM interception of a credit card sent over an HTTP link has ever been reported. Yet that is the threat the secure sockets layer (SSL) protocol's key management mechanism was designed to address.

On the other hand, attackers have harvested much credit card data from merchant databases after a transfer took place—what the failed secure electronic transactions (SET) protocol was designed to prevent. The success of the far-simpler SSL was one of the many nails in SET's coffin.

Advantages of key continuity

Some consider SSL a solution to the MITM threat; others regard the protocol as a placebo to allay unfounded fears. In any case, key continuity mechanisms can address a large percentage of key management issues, with the small remainder left to the unwieldy, heavyweight, and usually very expensive PKI mechanism (P. Gutmann, "PKI: It's Not Dead, Just Resting," *Computer*, Aug. 2002, pp.41-49).

Because it is effectively free, users can implement and operate key continuity management automatically, making it invisible to the user except when it warns of a potential problem. The importance of such ease of use should not be underestimated. As Ravi Sandhu has observed, "Cumbersome technology will be deployed and operated incorrectly and insecurely, or perhaps not at all" ("Good-Enough Security: Toward a Pragmatic Business-Driven Discipline," *IEEE Internet Computing*, Jan./Feb. 2003, pp. 66-68).

SSH

The secure shell (SSH) protocol was the first widely used security application to employ continuity as its primary form of key management. When a user initially connects to an SSH server, the client application warns that it has received a new public key and asks the user whether he wants to continue. When the user clicks "OK," the client application remembers the key that was used and stores it for comparison to future keys from the server. If the client application repeatedly receives the same key, this *known-hosts mechanism* identifies the source as most likely the same server.

SSH also lets a user verify a key through its *fingerprint*: a universal key identifier consisting of the hash of the key components or of the certificate if the key is in that form.

However, SSH doesn't provide complete key continuity. When the server software is reinstalled, unless the system administrator has remembered to archive the configuration and keying information, the system loses its connection to its previous key. Because SSH is typically used to secure access to frequently upgraded open source Unix systems, this loss of continuity can occur more often than would first appear.

The ease with which system administrators can change SSH keys and thus break key continuity introduces new problems. In PKI environments, where changing keys is arduous and frequently expensive, system administrators often choose instead to reuse the same key in perpetuity. This ensures key continuity, but recycling the same key over several years will compromise all data that the key protected during that time if the key is ever compromised.

In contrast, an SSH key can be replaced quickly and easily, limiting its exposure to attack but breaking the chain of continuity. A solution to this problem would be to have the server automatically generate and certify key $n + 1$ when key n is used, and then to save key $n + 1$ to offline media such as a floppy disk or a Universal Serial Bus memory token for future use.

This approach maintains continuity to the previous known server key when the system administrator reinstalls or replaces the system software or SSH server. A small change to the SSH protocol would allow the system to send an old-with-new key exchange message after the changeover.

CERTIFICATE APPROACHES

Unlike SSH, the SSL/transport layer security (SSL/TLS) and IPsec protocols were designed to rely on an external key management infrastructure. If necessary, both can function without a PKI by using shared keys—typically, passwords.

The real test of a security mechanism's effectiveness is its user acceptability.

Hard-coded certificate-authority certificates

SSL/TLS and IPsec compensate for the lack of a PKI in two ways. First, in SSL/TLS, the client—particularly in its most widespread form, the Web browser—contains more than 100 hard-coded certificate-authority certificates that in turn are trusted to issue SSL server certificates.

However, CA certificates can exhibit numerous vulnerabilities. Many hard-coded CAs are completely unknown; they often follow dubious practices such as using weak 512-bit keys or keys with 40-year lifetimes; many appear moribund; and in some cases, CA keys have been sold to third parties when the original owners went out of business.

Because all CAs are assigned the same level of trust, the entire system is only as secure as the least secure CA. Subverting any one CA compromises the entire collection. This lackadaisical approach to certificate security may explain why a recent SecuritySpace.com survey found that about three-quarters of all certificates in use by Web servers were invalid (www.securityspace.com/s_survey/sdata/200312/certca.html).

Self-issued certificates

Alternatively, both SSL/TLS and IPsec can act as their own certificate authorities, issuing themselves certificates for installation on client/peer machines.

In both protocols, the security these certificates provide is little better than that for SSH keys unless the client is careful to disable all but one or two trusted CA certificates. However, this process requires about 700 mouse clicks in the latest version of Internet Explorer.

With most hard-coded CA certificates disabled, the client software begins to display warning dialogs when an attempt is made to connect to a server with a certificate from a now-untrusted CA. However, the SecuritySpace.com survey indicates that browsers must either be accepting these certificates on many sites anyway or ignoring the problem of certificate validity for fear of scaring users.

STARTTLS

StartTLS is an extension to the simple mail transfer protocol (SMTP) that provides opportunistic TLS-based encryption for mail transfers. Similar facilities exist for other mail protocols such as POP and IMAP. Because of their use of mechanisms like StartTLS, applications like mail transfer agents have to deal with self-issued and similar informal certificates more frequently than other applications such as Web servers.

StartTLS provides a means of authenticating legitimate users and preventing misuse by spammers. Because mail servers are configured by system administrators rather than commercial organizations worried about adverse user reactions to browser warnings, they typically use self-issued certificates to avoid paying a CA.

Key continuity management among StartTLS implementations remains somewhat haphazard. Because StartTLS is intended to be a completely transparent solution, the ideal setup would automatically generate a certificate on the server side when the software is installed and use standard SSH-style key continuity management on the client side, with optional out-of-band verification available through a key or certificate fingerprint.

Some implementations—typically open source ones—support this concept fully; others support various aspects of it by, for example, requiring tedious manual operations for certificate generation or key or certificate verification. Some commercial implementations require using commercial CA certificates—an even more tedious and expensive manual operation.

Because of their mostly automated, transparent nature, mechanisms such as StartTLS have been extremely successful. Just as SSH has almost completely displaced the insecure telnet and rcp protocols that it was designed to supplement, StartTLS could replace or at least augment SMTP. A year or so after appearing, StartTLS is protecting more e-mail than all other e-mail security mechanisms combined, despite their 10-15 year lead in this area. For StartTLS use to continue to grow, all MTAs and mail user agents would need to fully automate the process by using automatic certificate generation on the server side and key continuity management on the client side.

S/MIME

Like SSL, the S/MIME protocol has a built-in mechanism for dealing with the absence of a PKI. Applications using S/MIME include any certificates that might be needed as part of each message sent. Signed e-mail, for example, includes all the certificates needed for verification along with the message. The S/MIME application remembers these certificates for future use, implementing a form of lazy-update PKI that distributes certificates on an on-demand basis.

S/MIME gateway implementers and some S/MIME clients have added two steps to this process. First, when it encounters a new certificate from a remote user, the S/MIME application can return a signed challenge message to the remote user. This challenge message is encrypted with the public key in the new certificate and includes the local S/MIME user's certificate.

The remote user decrypts the challenge and returns it re-encrypted with the key in the local user's certificate or performs some equivalent variant thereof, depending on the implementer's preferences. This provides a copy of the local user's certificate to the remote user and offers mutual proof of possession for local and remote user keys and certificates.

Second, the S/MIME application can automatically and invisibly generate a self-issued certificate the first time a new user sends e-mail, thereby avoiding the arduous and often expensive task of having users obtain certificates themselves.

Although no formal standard governs this use of S/MIME, most who are involved in deploying an S/MIME gateway have independently hit upon the process over the years. Specifying the challenge-response mechanism would help standardize this process.

The real test of a security mechanism's effectiveness is its user acceptability. Users will invariably regard an intrusive and awkward security mechanism as an impediment to bypass or ignore.

The rapid—and completely voluntary—adoption of SSH and StartTLS shows that opportunistic key generation and key continuity management make it possible to deploy security mechanisms at a scale and level of effectiveness that years of effort with more complex alternatives have failed to achieve. ■

Peter Gutmann is an honorary researcher in the Department of Computer Science at the University of Auckland, New Zealand. Contact him at pgut001@cs.auckland.ac.nz.

Editor: William A. Arbaugh, Dept. of Computer Science, University of Maryland at College Park; waa@cs.umd.edu

Mobile Agents in Wireless Devices

Luminita Vasiu, Middlesex University
Qusay H. Mahmoud, University of Guelph

The networks that connect handheld wireless devices such as cell phones and PDAs suffer from low bandwidth and a high incidence of network errors. Mobility can also result in the loss or degradation of wireless connections.

By employing mobile agents, such devices could provide a reliable technology for message transport over wireless links. Mobile agents are inherently distributed software entities that reduce the load on the network when they move. In addition, they support disconnected operations because they continue to execute after they move, even if they lose network connectivity with their dispatchers.

Mobile agents can be employed in wireless handheld devices in two ways: An agent platform could be installed on the device, enabling mobile agents to run directly on it, or devices could access and use remote mobile agents running on wired networks.

Each approach is viable and has its own advantages and domain-specific applications. Some high-end devices would benefit from running a mobile agent platform that lets agents run locally, but this would not be beneficial to others because of processing power and memory constraints or for security reasons.

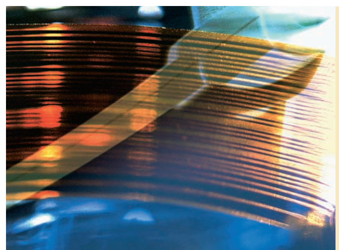

Mobile agents offer a potential solution to network problems associated with handhelds.

DEVICES FOR AGENTS

Researchers are developing a number of mobile agent platforms for handheld devices. For example, the Lightweight Extensible Agent Platform (http://leap.crm-paris.com) can execute on small devices, is scalable in both size and functionality, is "operating system agnostic," and complies with Foundation for Intelligent Physical Agents (www.fipa.org/) standards. Initiated by a number of companies including Motorola, Siemens, and British Telecom, LEAP is currently implemented as a lightweight extensible kernel for the Java Agent Development Framework (http://sharon.cselt.it/projects/jade) and will offer knowledge and travel management services.

Other proposed mobile agent platforms contain several common features in existing platforms such as agent messaging, events, and an agent execution environment. Primarily designed to offer mobile commerce services, such platforms are embedded directly into a K virtual machine (http://java.sun.com/products/cldc) rather than define additional classes on top of the KVM.

This approach has several disadvantages. As a practical matter, many handheld devices do not currently have sufficient processing power and memory to run an agent platform. In addition, running agents on wireless handhelds makes the devices and their data vulnerable to malicious code. Such platforms tend to concentrate on mobile agent features rather than security or quality.

Interoperability is also a problem. For example, a proprietary KVM is incompatible with the standard KVM. Because KVMs do not support object serialization, other mechanisms such as XML must be used to transmit mobile agents to and from the wireless device. Users will thus have to install new software components as well as install and configure the agent platform itself.

AGENTS FOR DEVICES

Instead of running an agent platform and executing agents locally, wireless handheld devices can include an interface that accesses and executes remote mobile agents. We have developed MobiAgent, a prototype agent-based framework and software infrastructure that minimizes the load on the wireless link and supports disconnected operations.

System components

As Figure 1 shows, the MobiAgent system consists of four major components: clients, the communication manager, the agent gateway, and MobiAgent services. Both wired and wireless links connect these components.

Clients are wireless applications or Java 2 Micro Edition MIDlets that run on the handheld device to provide the user with an interface for creating profiles and accessing MobiAgent services. Like applets, applications downloaded to the device run in a *sandbox*, a restricted environment with no access to the underlying resources. Clients request a service from the communication manager.

The communication manager mediates between the wireless device and the wired network. It manages all system connections and communicates with the agent gateway. A client using MobiAgent services for the first time downloads a wireless application or MIDlet from the communication manager that provides facilities for creating a user profile. Java servlets and common gateway interface (CGI) scripts communicate the user profile's identification and addressing properties to the agent gateway.

The agent gateway provides an interface to the communication manager and MobiAgent services to create and update user profiles, which are stored in a database. For security reasons, clients cannot directly access the user profiles database. The agent gateway can also run a mobile agent platform, although the agent platform can run on another host as well. MobiAgent runs Voyager, but it could use any agent platform.

MobiAgent services perform actions on behalf of clients that range from accessing network resources—such as Web sites containing information about books, cars, or travel information—to controlling remote applications. MobiAgent services have unique identifiers and accept requests from the agent gateway.

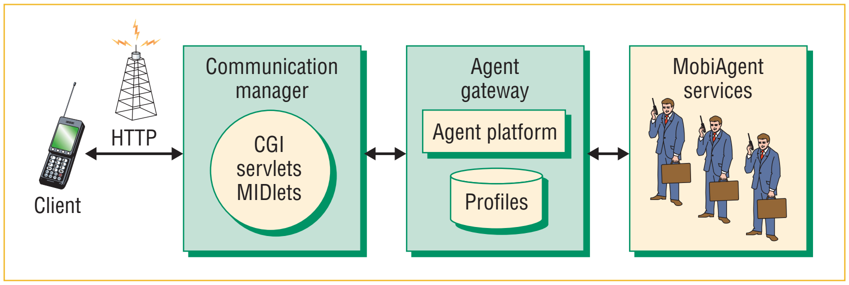

Figure 1. MobiAgent system components. Both wired and wireless links connect clients, the communication manager, the agent gateway, and MobiAgent services.

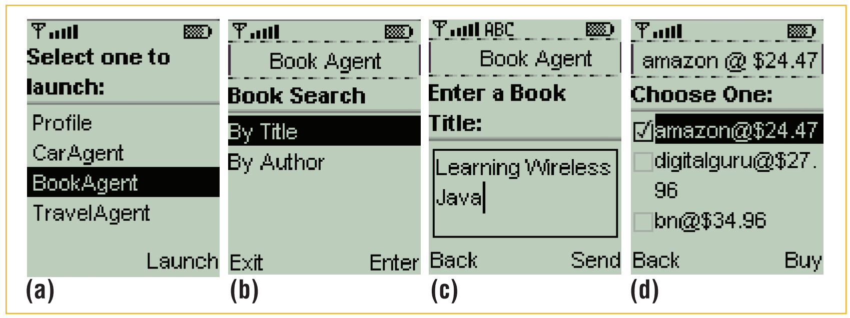

Figure 2. BookAgent service. (a) The user requests the service and (b) is prompted to search by title or author. (c) The user enters the book's title, and the mobile agent begins its task. (d) The agent returns to its host environment and sends the results to the agent gateway, which in turn forwards them to the user.

Operating scenario

Figure 2 shows how a cell phone user might apply the MobiAgent system to locate price information about a particular book. When using the system, the user can disconnect from the network while waiting for results to become available. If the user remains disconnected when the agent has finished its task, the agent gateway will inform the user that the results are available via Short Message Service. After reconnecting, the user can download the results.

Advantages

The major advantage of the MobiAgent system is that the user does not have to install or configure a mobile agent platform on the device. The system also reduces communications over wireless links to overcome low bandwidth and network disconnection. In addition, MobiAgent enhances service functionality by operating without constant user input. Finally, the system is platform independent.

Mobile agents must have an important role in wireless computing if software applications for handheld devices are to be truly useful and intelligent. However, security concerns present a major obstacle. For agents to collaborate, they must be able to meet somewhere and then return home with results for their dispatchers. One possible way to get around this problem is to let agents run on handheld devices but exclude foreign agents. ■

***Luminita Vasiu** heads the Wireless IT Research Centre, School of Computing Science, Middlesex University, and is Chair of the British Computer Society Wireless IT Specialist Group. Contact her at l.vasiu@mdx.ac.uk.*

***Qusay H. Mahmoud** is an assistant professor in the University of Guelph's Department of Computing and Information Science and associate chair of the Distributed Computing and Communications Systems Technology program at the University of Guelph-Humber. Contact him at qmahmoud@cis.uoguelph.ca.*

Editor: Upkar Varshney, Department of CIS, Georgia State University, Atlanta GA 30002-4015; uvarshney@gsu.edu

Making Computer Programming Fun and Accessible

Qusay H. Mahmoud, Wlodek Dobosiewicz, and David Swayne,
University of Guelph

Learning to program requires much hard work and dedication. Most students in introductory programming courses struggle to grasp programming concepts in general. While some students keep at it until the key concepts gel, many simply give up instead. Universities in the US, Canada, and elsewhere have reported withdrawal, failure, and D-grade rates approaching 50 percent in introductory computer programming courses.

To address these problems, we recently designed and successfully offered a revolutionary one-semester course that integrates the use of HTML, JavaScript, and Java. This approach departs markedly from using a single general-purpose programming language such as Java or C++.

An introductory course's use of JavaScript provides an environment in which students can excel.

COMMON PROBLEM

Today, most introductory programming courses are based on learning a general-purpose programming language such as Java. Yet many students with no previous computing background find it hard to grasp the concept of programming in general, and object-oriented programming in particular. Our experience shows that many students have no problems with the language itself, but the tools they are using—such as the development environment—make the learning process difficult. For example, the compiler often produces cryptic error messages.

Further, in introductory programming courses assignments usually consist of command-line input and output. Students become frustrated when all they see are error messages and no output. Novice programming students struggle to interpret such cryptic error messages. This problem cannot be solved by teaching GUI programming in Java using, for example, the AWT/Swing packages, because programming in Swing requires a basic knowledge of Java syntax and its object-oriented programming mechanisms.

UNIQUE SOLUTION

The University of Guelph-Humber (www.guelphhumber.ca), an initiative between the University of Guelph and Humber College Institute of Technology and Advanced Learning, offers the Distributed Computing and Communications Systems Technology program. The two institutions combine to give students a fully recognized university honors degree in applied computing and a college diploma in communications systems technology. This four-year cooperative education program is the only one of its kind in Canada.

Students here can study the convergence of computing and telecommunications in a single program, which gives them the applied computing expertise they need to provide working solutions to real-life problems. The program emphasizes systematic approaches to the design and development of secure distributed applications, hands-on experiments in leading-edge wired and wireless technologies, and teamwork projects. In this lab-intensive program, students learn by doing.

To help the program succeed, an innovative curriculum encourages further learning and increases student retention. Our first computer programming course provides a step-by-step approach to programming that requires no previous experience with computers in general or programming in particular.

We designed the course to encourage students to get started with programming, even at a slow pace, rather than risk frustrating and eventually losing them. This course focuses on using computers to build applications for fun: static homepages, interactive homepages using JavaScript, and simple stand-alone Java applications. We accomplished this through *click programming*, a teaching methodology that integrates more than one programming language. By introducing HTML, JavaScript, and Java successively, each topic prepares the students

Continued on page 106

Continued from page 108

for the next, harder topic. The course lasts 13 weeks, with a three-hour lecture each week plus a total of 10 three-hour labs. The course covers HTML for two weeks, JavaScript for four weeks, and Java for seven weeks.

The course's first section introduces file management and basic Internet concepts, the World Wide Web, and HTML. The second section introduces the basic elements of programming in JavaScript, including input/output, decisions and control flow, instantiating objects, invoking methods on objects, defining functions, and handling events. The third section introduces basic Java concepts and the elements of programming in Java, including using primitive data types, instantiating and using objects, invoking methods, dealing with input and output but not file I/O, and designing classes.

We use pair programming (www.pairprogramming.com) for both labs and assignments. This approach lets students share knowledge and learn about tips and tricks that their classmates might be familiar with. In the pairing process, we avoided letting students pair themselves, fearing that friends and possibly strong programmers would pair together. To ensure a fair pairing process, we paired students randomly and rotated pairs for all labs and assignments.

OLD VERSUS NEW APPROACHES

Traditional introductory courses, which focus on teaching a general-purpose language, place several obstacles between students and successful learning—and increase the instructor's workload as well.

General complexity

Several general-purpose programming languages, such as Pascal, Modula-2, C++, and Java, have been widely used in introductory programming courses. Teaching students how to program in any of these languages is difficult, especially if they have no previous computing background. For example, consider the following segment of Java code for the classical Hello World program:

```
public class Hello {
 public static void
main(String argv[]) {
   System.out.println
      ("Hello World!");
 }
}
```

Even in this simple program, many things must be explained in the first lecture. For example, what is the meaning of `public`, `static`, `void`, `main`, `String`, and `[]`? These unknowns could frustrate students immediately in the first week of the course.

Pair programming lets students share knowledge with their classmates.

In addition, user input and output must be explained in the first few weeks because students will need to write programs that depend on user input. Reading input from the user, however, is difficult. Consider the following segment of Java code that prompts the user to enter a number:

```
System.out.print("Enter
      a number: ");
BufferedReader br = new
      BufferedReader (
   new InputStreamReader
      (System.in));
String str = br.readLine
      ();
int number = Integer.
      parseInt (str);
```

While streams represent a powerful mechanism in Java, several objects must be instantiated to construct an input stream that reads from the keyboard. Further, the input will be read as a string and must be parsed and converted to an integer. The preceding code does not even include the exception-handling constructs `try { }` and `catch { }`.

Scripted elegance

Compare the preceding segment of Java code with the following two-line JavaScript code that prompts the user for input and parses it:

```
var str = prompt("Enter
   a number: ", "");
var number =
parseInt(str);
```

When this JavaScript script runs in a browser, it will present the user with an accessible and attractive dialog box. Students certainly like using such powerful features because they immediately feel productive—an important benefit for first-year students that encourages them to pursue further learning. To use JavaScript, however, students must learn about HTML because JavaScript is embedded in HTML documents.

Axing the compiler

When a programmer writes in a conventional programming language such as Java, that code must be compiled and run. If any compilation errors occur, the programmer must edit the program, recompile, then run the program again. Making errors in a conventional programming language thus produces a fatal error that prevents the program from executing. Making errors in HTML is *not* fatal: The browser will make its best effort to render the page, but will probably not display it as intended. We did, however, encourage students to write valid HTML documents and not rely on the forgiveness of the browser. It is worth noting that students developed their HTML documents using a text editor of their choice—most used Windows Notepad and didn't rely on an HTML editor.

Once students become comfortable with HTML, the course introduces them to some programming concepts using JavaScript. Students can embed their JavaScript code in HTML documents, a powerful approach that lets them learn programming at a slow pace without worrying about cryptic compiler-error messages.

Next, the course covers JavaScript. During this time, students learn about objects and method invocation. After this, they're well prepared to move on to Java. Given that decisions and control-flow statements have already been covered in JavaScript, when the students move to Java they can start producing interesting Java programs almost immediately.

On the other hand, Java and JavaScript differ significantly with regard to their object models, execution environments, and typing strictness. These differences, especially typing, caused some confusion during the first week of our Java section. This confusion has never become a major problem, however, and students gain a good understanding of typed languages as well as implicit and explicit casting.

We strongly believe that teaching computer programming in the context of simple client-side Web applications provides a motivating framework for students and encourages them to excel. We urge computing educators to try this revolutionary programming approach.

Based on student retention rates, we believe our course has been successful. Overall, performance in programming assignments and exam results have far surpassed those in other introductory programming courses we have taught at universities with students of similar quality. In our first offering during the autumn of 2002, all 34 students passed the course. Student marks spanned a high of 96 percent and a low of 60 percent, averaging 73 percent. Students learned how to create simple and dynamic Web sites and acquired the fundamentals of programming.

Further, feedback collected from students suggests they enjoyed the experience. We believe this approach is applicable not only to novice computing professionals, but also to those who wish to gain more control over what their computer does for them. The course could also enhance the teaching of computing studies in schools. In an effort to increase enrollment in the Distributed Computing and Communications Systems Technology program, we plan to let students from other disciplines take the course. ■

Qusay H. Mahmoud *is an assistant professor in the Department of Computing and Information Science, University of Guelph, and associate chair of the Distributed Computing and Communications Systems Technology program, University of Guelph-Humber. Contact him at qmahmoud@cis.uoguelph.ca.*

Wlodek Dobosiewicz *is a professor and chair of the Department of Computing and Information Science, University of Guelph. Contact him at dobo@cis.uoguelph.ca.*

David Swayne *is a professor in the Department of Computing and Information Science, University of Guelph, and program coordinator of the Distributed Computing and Communications Systems Technology program, University of Guelph-Humber. Contact him at dswayne@cis.uoguelph.ca.*

Editor: Neville Holmes, School of Computing, University of Tasmania, Locked Bag 1-359, Launceston 7250; neville.holmes@utas.edu.au