

ACM ANNUAL REPORT FOR FY07

# COMMUNICATIONS

FEBRUARY 2008 VOLUME 51, NUMBER 2

of the ACM

## ALTERNATE REALITY GAMING

WOMEN AND  
MEN IN IT

THE MYTHS  
AND TRUTHS  
OF WIRELESS  
SECURITY

PEER TO PATENT  
PROJECT NEEDS  
EXPERTS

MUNICIPAL  
WIRELESS  
NETWORKS



Association for  
Computing Machinery



Formerly Idea Group Inc.

# NEW YEAR, NEW JOURNALS!

IGI Global, provider of superior journals in computer science and information technology management, is expanding its journal program with three new titles available now, 30 more coming by the end of this year, and a total of 100 by the end of 2009. This rapidly growing collection will provide researchers with unprecedented depth and breadth in coverage of cutting-edge technology topics.



## International Journal of Information Systems and Supply Chain Management

John Wang, Montclair State University, USA

ISSN: 1935-5726, EISSN: 1935-5734  
Individual (Print): US \$125.00  
Institution (Print and Free Online): US \$495.00  
Institution (Online Only): US \$445.00

Examines current, state-of-the-art research in the areas of SCM, and the interactions, linkages, applications, and support of SCM using information systems.



## International Journal of Information Technologies and Systems Approach

David Paradise, Florida State University, USA, and Manuel Mora, Universidad Autónoma de Aguascalientes, México

ISSN: 1935-570X, EISSN: 1935-5718  
Individual (Print): US \$125.00  
Institution (Print and Free Online): US \$495.00  
Institution (Online Only): US \$445.00

Emphasizes the systemic world-view of managerial, organizational and engineering interaction facets provided by modern complex information systems and information technologies.



## Journal of Information Technology Research

Mehdi Khosrow-Pour, Information Resources Management Association, USA

ISSN: 1938-7857, EISSN: 1938-7865  
Individual (Print): US \$115.00  
Institution (Print and Free Online): US \$395.00  
Institution (Online Only): US \$345.00

Presents comprehensive interdisciplinary and refereed research on the most emerging and breakthrough areas of information science and technology.

Recommend these outstanding peer-reviewed journals to your library today. And talk to your librarian about InfoSci-Journals, which provides unlimited, full-text access to all IGI Global journals at a cost equivalent to just a few subscriptions. See [www.infosci-journals.com](http://www.infosci-journals.com) for more information and free trials!

For more information or to order IGI Global's new 2008 journals or InfoSci-Journals contact customer service at [cust@igi-global.com](mailto:cust@igi-global.com) or visit IGI Global's website at [www.igi-global.com](http://www.igi-global.com).

IGI Global • 701 E. Chocolate Ave., Suite 200 • Hershey, PA 17033-1240 USA • 1-866-342-6657  
[cust@igi-global.com](mailto:cust@igi-global.com) • [www.igi-global.com](http://www.igi-global.com)



# CSCW08

November 8-12, 2008  
Hilton San Diego Resort  
San Diego, California, USA

# CALL FOR PAPERS

The 2008 ACM Conference  
on Computer Supported  
Cooperative Work

## Get your CSCW Submissions Ready!

CSCW is the premiere conference on the technical and social aspects of communication, coordination, competition and collaboration. The conference's topics continue to expand as we increasingly use technologies to live, work and play with others. For complete details about CSCW's multiple participation venues, please review the Call for Participation on our website [www.cscw2008.org](http://www.cscw2008.org).

### Venues of Participation

The conference offers many opportunities, some old and some new, for you to participate and collaborate.

<b>Venues of Participation</b>	<b>Submission Deadline</b>
Papers & Notes	18 April 2008
Workshops	2 May 2008
Panels	2 May 2008
Student Volunteers	31 May 2008
Doctoral Colloquium	18 July 2008
Interactive Posters	18 July 2008
Demonstrations	18 July 2008
Videos	18 July 2008

### Information

Questions and requests for information should be sent to [publicity@cscw2008.org](mailto:publicity@cscw2008.org)

Also, see our website [www.cscw2008.org](http://www.cscw2008.org)



## Call For Papers Sept. 23-26, 2008 Beijing, China



### IEEE International Conference on Services Computing (SCC 2008)

July 8 - 11, 2008, Hawaii, USA  
<http://conferences.computer.org/scc/2008>

Services Computing has become a foundational discipline that covers the science and technology of services innovation research that leverages IT and computing technology to model, create, and manage business solutions, scientific applications, as well as modernized services. SCC 2008 continues to bridge the gap between Services Computing and business models with an emerging suite of ground-breaking technology that includes Service Oriented Architecture (SOA), business process integration and management, services engineering and utility/grid computing. The theme of SCC 2008 is "Services: Business, Technology, and Application."

SCC 2008 seeks original, unpublished research and industry/application papers in the following major tracks:

- Foundations of Services Computing
- Business Models and Operations in Services Industry
- Business Process Management and Integration
- SOA Tools, Solutions, and Services

#### SCC 2008 and SERVICES 2008 (Part A) Important Dates

Paper Submission Due : 1/28/2008  
Decision Notification : 3/25/2008  
Camera-Ready Due : 4/18/2008

Sponsored Journals for ICWS & SCC 2008  
Int. Journal of Web Services Research (JWSR) (SCI)  
Int. J. of Business Process Integration & Management (IJBPIM)  
Int. Journal of Grid & Utility Computing (IJGUC)  
IEEE IT Professional Magazine  
IEEE Transactions on Services Computing (TSC)

### IEEE Congress on Services (SERVICES 2008)

<http://conferences.computer.org/services/2008>  
Part A: Jul. 8-11, 2008, Hawaii, USA  
Part B: Sep. 23-26, 2008, Beijing, China

The two flagship conferences, ICWS 2008 & SCC 2008, are strategically part of SERVICES 2008. While ICWS & SCC will continue to focus on high-quality research and industry paper sessions, the theme of SERVICES is to support them. SERVICES 2008 will include the following key events:

Services University, 2008 IEEE Services Computing Contest, 2008 IEEE SOA Industry Summit, 2008 IEEE Services Computing Workshops (SCW), 2008 IEEE Ph.D. Student Symposium on Services Computing, 2008 IEEE Symposium on SOA Standards, Demos and Exhibits, Job Fair, IEEE Body of Knowledge (BoK) on Services Computing.

All papers reviewed and accepted by SERVICES 2008 events in Hawaii and Beijing will be published in the Proceedings of 2008 IEEE Congress on Services (SERVICES 2008).

### IEEE International Conference on Web Services (ICWS 2008)

September 23 - 26, 2008, Beijing, China  
<http://www.icws.org>

ICWS has been a prime international forum for both researchers and industry practitioners to exchange the latest fundamental advances in the state of the art and practice of Web services. ICWS also aims to identify emerging research topics and define the future of Web services. Over the past six years ICWS has grown steadily attracting over 250 participants on a regular basis.

ICWS 2008 will continue to feature research and industry/application papers with a wide range of topics, focusing on the infrastructure of SOA and various aspects of IT services.

## Call For Papers



ICWS 2008 seeks original, unpublished papers reporting new work in all aspects of Web Services (WS). Topics of interest include but are not limited to:

- Extended WS Specifications
- WS Publishing & Discovery
- WS Composition & Integration
- WS Testing
- Security Challenges for WS
- QoS in WS
- Web 2.0 and Web X.o Concepts in WS
- Semantic WS
- WS-based Mobile Computing
- WS Standards and Implementation
- WS Modeling
- Relationship Binding, and Tooling

#### ICWS 2008 and SERVICES 2008 (Part B) Important Dates

Paper Submission Due : 4/7/2008  
Decision Notification : 5/26/2008  
Camera-Ready Due : 6/27/2008

ICWS 2008, SCC 2008, and SERVICES 2008 Conference Proceedings will be published by IEEE Computer Society Press. They will be included in EI Compendex (ICWS 2003-2007, SCC 2004-2007, SERVICES 2007) & SCI (in Plan)

Based on the best practice at SERVICES 2007, all papers/presentations (To be captured in RSS/PodCast/Video formats by ICWS/SCC/SERVICES 2008) will be contributed to the public IEEE Body of Knowledge (BoK) on Services Computing Portal <http://www.servicescomputing.tv>



Sponsored by The Technical Committee on Services Computing  
<http://tab.computer.org/tesc>  
Contact: ZHANGLIJ@IEEE.ORG



## Call For Papers July 8-11, 2008 Hawaii, USA



# Table of Contents


February 2008

## Articles

- 36 **Alternate Reality Gaming** *Jeffrey Y. Kim, Jonathan P. Allen, and Elan Lee*
- 43 **IT Diffusion in Developing Countries** *Eric Shih, Kenneth L. Kraemer, and Jason Dedrick*
- 49 **Are People Biased in Their Use of Search Engines?** *Mark T. Keane, Maeve O'Brien, and Barry Smyth*
- 53 **Alternative Scenarios to the "Banner" Years** *Nolan J. Taylor, Eleanor T. Loiacono, and Richard T. Watson*
- 60 **Sharing Knowledge** *Peter Marks, Peter Polak, Scott McCoy, and Dennis Galletta*
- 66 **The Myths and Truths of Wireless Security** *Alfred Loo*
- 72 **Municipal Broadband Wireless Networks** *Munir Mandviwalla, Abhijit Jain, Julie Fesenmaier, Jeff Smith, Paul Weinberg, and Greg Meyers*
- 81 **Women and Men in the IT Profession** *Vicki R. McKinney, Darryl D. Wilson, Nita Brooks, Anne O'Leary-Kelly, and Bill Hardgrave*
- 85 **Managing Large Collections of Data Mining Models** *Bing Liu and Alexander Tuzhilin*



36



60

## Columns

- 15 **Technology Strategy and Management** *Managing Software Development in Globally Distributed Teams* *Michael A. Cusumano*
- 19 **Viewpoint** *Peer to Patent Needs Your Expertise* *Andy Oram*
- 21 **Legally Speaking** *How to Achieve (Some) Balance in Anti-Circumvention Laws* *Pamela Samuelson, Jerome H. Reichman, and Graeme Dimwoodie*
- 27 **ACM Annual Report for FY07**
- 91 **Technical Opinion** *Record Matching in Digital Library Metadata* *Min-Yen Kan and Yee Fan Tan*
- 104 **Inside Risks** *Software Transparency and Purity* *Pascal Meunier*

## Departments

- 7 **News Track**
- 9 **Forum**
- 14 **SIG Election Candidates**
- 95 **Career Opportunities**



# CALL FOR PAPERS

## 2008 ACM International Conference on Computing Frontiers



May 5 - 7, 2008, Ischia, Italy



Sponsored by ACM - SIGMICRO

### GENERAL CHAIR

Alex Ramirez, UPC, ES

### PROGRAM CO-CHAIRS

Gianfranco Bilardi, Università di Padova, IT  
Michael Gschwind, IBM TJ Watson, US

### STEERING COMMITTEE

Monica Alderighi, INAF, IT  
Utpal Banerjee, Intel, US  
Steven Beaty, Metro State College Denver, US  
Sergio D'Angelo, INAF, IT  
Michel Dubois, U. of Southern California, US  
Kemal Ebcioglu, Global Supercomputing, US  
Gearold Johnson, Colorado State U., US  
Sally A. McKee, Cornell University, US  
Cecilia Metra, University of Bologna, IT  
José Moreira, IBM TJ Watson, US  
Valentina Salapura, IBM TJ Watson, US  
Giacomo Sechi, INAF, IT  
Per Stenström, Chalmers U. of Technology, SE

### FINANCE/REGISTRATION CHAIR

Carsten Trinitis, TU München, DE

### SPECIAL SESSIONS CHAIR

Osman Unsal, BSC, ES

### LOCAL ARRANGEMENTS CHAIR

Claudia Di Napoli, CNR, IT

### PUBLICITY CHAIR

Julita Corbalan, UPC, ES

### LIAISON CHAIR FOR ASIA

Hitoshi Oi, University of Aizu, JP

### PUBLICATION CHAIR

Sergio D'Angelo, INAF, IT

### WEB CHAIR

Greg Bronevetsky, LLNL, US

The increasing needs of present and future computation-intensive applications have stimulated research in new and innovative approaches to the design and implementation of high-performance computing systems. These boundaries between state of the art and revolutionary innovation constitute the computing frontiers that must be pushed forward to provide the computational support required for the advancement of all science domains and applications. This conference focuses on a wide spectrum of advanced technologies and radically new solutions; it is designed to foster communication among many scientific and technological disciplines.

Authors are invited to submit papers on all areas of innovative computing systems that extend the current frontiers of computer science and engineering and that will provide advanced systems for current and future applications.

Papers are sought on theory, methodologies, technologies, and implementations concerned with innovations in computing paradigms, computational models, architectural paradigms, computer architectures, development environments, compilers, and operating environments. Papers are solicited in, but not limited to, the following areas:

- Non-conventional computing
- Next-generation high performance computing, esp. novel high-performance systems (including Cell, GPGPU and custom accelerators)
- Applications, programming models and performance analysis of parallel architectures and novel high-performance systems
- Virtualization and virtual machines
- Grid computing
- Compilers and operating systems
- Workload characterization of emerging applications
- Service oriented architecture (SOA) and system impact
- Supercomputing
- SOC architectures, embedded systems and special-purpose architectures
- Temperature, energy, and variability-aware designs; fault tolerance and reliability
- System management and security
- Quantum and nanoscale computing
- Computational biology
- Reconfigurable computing and architecture prototyping
- Autonomic and organic computing
- Computation intelligence frontiers: theory and industrial applications

All papers will be published in the conference proceedings and will be made available in the ACM Digital Library. In addition, selected papers will be invited to appear in a special issue of the HiPEAC journal.


Submitted manuscripts should not exceed 20 double-spaced, single-column pages, including figures, tables, and references. Submission implies that at least one author will register for the conference and present the paper, if accepted. Submissions must be made electronically as Adobe PDF files through the conference web-site, and must not be simultaneously submitted to any other publication outlet.

### Important Dates

Paper submission deadline:  
**December 7, 2007**

Author notification:  
**January 18, 2008**

Final papers due:  
**February 22, 2008**



# Editorial Pointers



ALTERNATE REALITY GAMING (ARG) IS A RELATIVELY new genre of collaborative entertainment that combines online information and puzzles with real-world events. Known for their elaborate scavenger hunts, ARGs attract players into interactive plots that often blur the distinction between cyberspace and public space; online cues lead players to real-world clues with the goal of finding the ending to the story—an ending they ultimately create themselves.

One recent example of ARGs that generated media interest was the cross-media game “I Love Bees,” or ILB. Its intricate and evolving storyline prompted a diverse group of players to follow cyber clues that led to ringing pay phones (and clues) across the U.S. and the U.K.

The ILB model posed some unique design challenges for the creative team. For this month’s cover story, Jeffrey Kim, Jonathan P. Allen, and Elan Lee retell the plotline for ILB as it unfolded, detailing the design decisions and the ways players ultimately constructed and directed the game. While ILB succeeded in adding new layers and dimensions to the ARG experience and supporting technologies, the authors contend that in the end, it’s all about the story.

ALSO THIS MONTH, MANDVIWALLA ET AL. EXAMINE MUNICIPAL wireless networks and how they represent a new option for broadband access to the Internet that promises anytime, anywhere connectivity. Alfred Loo says the strength of a computer system’s security is always measured by its weakest component, and in most systems that component is the end user. Further, wireless security will never be realized without a coordinated effort among corporations, manufacturers, employers, and end users.

Shih et al. explore global IT diffusion by analyzing data from 44 countries over a 15-year period and found markedly different results for developing and developed countries. McKinney et al. investigate the professional IT experience for men and women and find significant similarities that shine new light on reasons why women are still under-represented in the field. Marks et al. study the ways managers influence how knowledge is shared within an organization. Keane, O’Brien, and Smyth track the role of user bias in deciding which links to follow in rank-ordered results lists. And Taylor, Loiacono, and Watson discuss alternative formats for Web banner ads.

In this month’s “Viewpoint,” Andy Oram calls for technologists to support the Peer to Patent project, an initiative created to address the proliferation of computer-related patents that sometimes hinder innovation rather than promote it. Peer to Patent channels input directly from the public to the examiners at the U.S. Patent and Trademark Office who ultimately decide what patents to grant.

*Diane Crawford*  
EDITOR

## COMMUNICATIONS

of the ACM

A monthly publication of  
the ACM Publications Office

### ACM

2 Penn Plaza, Suite 701  
New York, NY 10121-0701 USA  
(212) 869-7440 FAX: (212) 869-0481

**Group Publisher:** Scott Delman

**Editor:** Diane Crawford

**Managing Editor:** Thomas E. Lambert

**Senior Editor:** Andrew Rosenbloom

**Editorial Assistant:** Zarina Strakhan

**Copyright:** Deborah Cotton

### Contributing Editors

Phillip G. Armour; Hal Berghel;  
Michael A. Cusumano; Peter J. Denning;  
Robert L. Glass; Seymour Goodman;  
Rebecca Mercuri; Peter G. Neumann;  
Pamela Samuelson; Meg McGinity Shannon

**Art Director:** Caren Rosenblatt

**Production Manager:** Lynn D’Addesio

### Advertising

ACM Advertising Department  
2 Penn Plaza, Suite 701, New York, NY 10121-0701  
(212) 869-7440; Fax: (212) 869-0481

### Account Executive:

William R. Kooney: [acm-advertising@acm.org](mailto:acm-advertising@acm.org)

For the latest media kit—including rates—contact:  
Graciela Jacome: [jacome@acm.org](mailto:jacome@acm.org)

### Contact Points

**CACM editorial:** [crawford\\_d@acm.org](mailto:crawford_d@acm.org)

**Copyright permission:** [permissions@acm.org](mailto:permissions@acm.org)

**Calendar items:** [calendar@acm.org](mailto:calendar@acm.org)

**Change of address:** [acmcoa@acm.org](mailto:acmcoa@acm.org)

### Communications of the ACM

(ISSN 0001-0782) is published monthly by the ACM, 2 Penn Plaza, Suite 701, New York, NY 10121-0701. Periodicals postage paid at New York, NY 10001, and other mailing offices.

POSTMASTER: Please send address changes to Communications of the ACM, 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA

Printed in the U.S.A.



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*

DOI: 10.1145/1314215.1314216

# ACM *The Association for Computing Machinery*

ACM (founded 1947) is an international scientific and educational organization dedicated to advancing the art, science, engineering, and application of information technology, serving both professional and public interests by fostering the open interchange of information and by promoting the highest professional and ethical standards.

**Executive Director and CEO:** John White  
**Director, ACM U.S. Public Policy Office:** Cameron Wilson

**Deputy Executive Director and COO:** Patricia Ryan  
**Director, Office of Information Systems:** Wayne Graves  
**Director, Office of Financial Services:** Russell Harris  
**Financial Operations Planning:** Darren Ramdin

**Director, Office of Membership:** Lillian Israel

**Director, Office of Publications:** Mark Mandelbaum  
Bernard Rous

**Deputy Director:** Diane Crawford  
**Deputy Director, Magazine Development:** Jono Hardjowirogo  
**Publisher, ACM Books and Journals:**

**Director, Office of SIG Services:** Donna Cappo

## ACM Council

**President:** Stuart I. Feldman  
**Vice-President:** Wendy Hall  
**Secretary/Treasurer:** Alain Chesnais  
**Past President:** David A. Patterson  
**Chair, SGB Board:** Joseph A. Konstan  
**Co-Chairs, Publications Board:** Ronald Boisvert, Holly Rushmeier

**Members-at-Large:** Michel Beaudouin-Lafon (2000–2008); Chuang Lin (2007–2008); Bruce Maggs (2006–2010); Jennifer Rexford (2007–2008); Barbara Ryder (2000–2008); David S. Wise (2004–2008).

**SGB Council Representatives:** Norman Jouppi (2006–2007); Robert A. Walker (2006–2008); Alexander Wolf (2005–2007)

## Board Chairs and Standing Committees

**Education Board:** Andrew McGettrick  
**SGB Board:** Joseph A. Konstan; **Membership Services Board:** Terry Coatta; **Publications Board:** Ronald Boisvert, Holly Rushmeier  
**Professions Board:** Stephen R. Bourne **USACM Committee:** Eugene Spafford

## SIG Chairs

**SIGACCESS:** Vicki Hanson **SIGACT:** Richard Ladner;  
**SIGAda:** John McCormick; **SIGAPL:** Guy R. Larocque;  
**SIGAPP:** Barrett Bryant; **SIGARCH:** Douglas C. Burger;  
**SIGART:** Maria Gini; **SIGBED:** Lothar Theile;  
**SIGCAS:** Florence Appel; **SIGCHI:** Julie Jacko;  
**SIGCOMM:** Mark Crovella; **SIGCSE:** Barbara Boucher Owens;  
**SIGDA:** Diana Marculescu; **SIGDOC:** Brad Mehlenbacher;  
**SIGecom:** David Pennock **SIGEVO:** Darrell Whitley;  
**SIGGRAPH:** G. Scott Owen; **SIGIR:** Elizabeth Liddy;  
**SIGITE:** Han Reichgelt; **SIGKDD:** Gregory Piatetsky-Shapiro;  
**SIGMETRICS:** Carey Williamson; **SIGMICRO:** Erik Altman;  
**SIGMIS:** Janice Sipior; **SIGMOBILE:** David B. Johnson;  
**SIGMOD:** Raghu Ramakrishnan; **SIGMULTIMEDIA:** Klara Nahrstedt; **SIGOPS:** Doug Terry; **SIGPLAN:** Kathleen Fisher;  
**SIGSAC:** Virgil D. Gligor; **SIGSAM:** Mark W. Geisbrecht;  
**SIGSIM:** Simon J.E. Taylor; **SIGSOFT:** William G. Griswold;  
**SIGUCCS:** Leila Lyons; **SIGWEB:** Ethan Munson

For information from Headquarters: (212) 869-7440

## ACM U.S. Public Policy Office:

Cameron Wilson, Director  
1100 Seventeenth St., NW  
Suite 507  
Washington, DC 20036 USA  
+1-202-659-9711—office  
+1-202-667-1066—fax  
wilson\_c@acm.org

# COMMUNICATIONS

OF THE ACM • *A monthly publication of the ACM Publications Office*

ACM, 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA (212) 869-7440 FAX: (212) 869-0481

## Editorial Advisory Board

Gordon Bell; Hal Berghel; Grady Booch;  
Nathaniel Borenstein; Vinton G. Cerf;  
Kilnam Chon; Jacques Cohen; Larry L. Constantine;  
Jon Crowcroft; Peter J. Denning; Mohamed E. Fayad;  
Usama Fayyad; Christopher Fox; Ravi Ganesan;  
Don Hardaway; Karen Holtzblatt; Pattie Maes;  
Eli Noam; Cherri Pancake; Yakov Rekhter;  
Douglas Riecken; Ted Selker; Dennis Tschritzis;  
Ronald Vetter

## Publications Board

Co-Chairs: Ronald F. Boisvert and Holly Rushmeier  
Board Members: Gul Agha; Michel Beaudouin-Lafon; Carol Hutchins; Mary Jane Irwin; Ee-ping Lim; Keith Marzullo; M. Tamer Ozsu; Vincent Shen; Mary Lou Soffa; Ricardo Baeza-Yates

## ACM Copyright Notice

Copyright © 2008 by Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page.

Copyright for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or fee. Request permission to publish from: Publications Dept. ACM, Inc. Fax +1 (212) 869-0481 or email <permissions@acm.org>

For other copying of articles that carry a code at the bottom of the first or last page or screen display, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, 508-750-8500, 508-750-4470 (fax).

## Subscriptions

Annual subscription cost is included in the society member dues of \$99.00 (for students, cost is included in \$40.00 dues); the nonmember annual subscription is \$189.00. See top line of mailing label for subscription expiration date coded in four digits: the first two are year, last two, month of expiration. Microfilm and microfiche are available from University Microfilms International, 300 North Zeeb Road, Dept. PR, Ann Arbor, MI 48106; (800) 521-0600.

**Single Copies** are \$8 to members and \$17 to nonmembers. Please send orders prepaid plus \$7 for shipping and handling to ACM Order Dept., P.O. Box 11414, New York, NY 10286-1414 or call (212) 626-0500. For credit card orders call (800) 342-6626. Order personnel on duty 8:30–4:30 EST. After hours, please leave message and order personnel will return your call.

## Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you were previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG newsletter at any time, and you do not want this work to appear in the ACM Digital Library, please inform permissions@acm.org, stating the title of the work, the author(s), and where and when published.



# News Track

## PUBCAST OR PERISH?

SciVee.com, funded by the National Science Foundation, is the latest in a growing number of video-sharing startups designed to let scientists broadcast themselves working in the lab or delivering lectures. The media outlet encourages scholars with new research papers to make short videos, or pubcasts, highlighting their key points. The Associated Press reports fans of the niche sites say they help the public—and students—understand the scientific process, while allowing researchers to duplicate one another's results and discourage fraud. Journal editors are quick to caution the leap to video will not displace the practice of reporting scientific results in peer-reviewed journals or at scientific conferences anytime soon. Most journals with online editions are taking a wait-and-see approach about YouTube-type videos, though many routinely add podcasts and other media to accompany papers. Says Stewart Wills, online editor of *Science*: "This is an area we're extremely interested in, but we're still in the embryonic stage."

## TOP 10 IT DISASTERS


ZDNet.co.uk recently posted the Top 10 IT-related disasters and failures of all time. While the list is subjective, and purposely omits incidents that resulted in loss of life, "it does illustrate where faulty hardware and software have cost organizations dearly, both financially and in terms of reputation." With number one being the worst, the list notes:

1. Faulty Soviet early warning system nearly causes WWII (1983)
2. AT&T network collapse (1990)
3. The explosion of the Ariane 5 (1996)
4. Airbus A380 suffers from incompatible software issues (2006)
5. Mars Climate Observer metric problem (1998)

6. EDS and the Child Support Agency (2004)
7. The Y2K problem (1999/2000)
8. When laptops exploded (2006)
9. Siemens and the passport system (1999)
10. Los Angeles Airport flights grounded (2007)

## DECODING THE EYE

A new chip densely packed with electrodes is the first step toward creating artificial retinas that approximate normal vision to the degree they enable people to read. While retinal prostheses have been used in human clinical trials, this new chip, developed by scientists at the University of California, Santa Cruz, promises to add focus and clarity to what is now blurs of motion. *Technology Review* reports the new chip can stimulate and record from individual cells in retinal samples. The technology will provide insight into how the retina codes information and how to mimic that coding—lessons crucial to developing the next generation of retinal implants. In the future, some version of the technology might be used to send visual information down the optic nerve. The 512-



electrode array (the gold circle in the picture) was modeled after the silicon microchip detectors used to capture particles in high-energy physics. Now that scientists have created a technology with such precise control, they are using it to study the language of the retina—one they hope implants will ultimately be able to understand.

## MONA LISA'S SMILE

The *Mona Lisa*, a perennial topic of lore, has been subjected to many experiments in an effort to unlock its mysteries. The latest comes from Paris where engineer Pascal Cotte spent much of the last two years working with cutting-edge digital scanning

# News Track



## OUT OF THE BOX

Intel is looking to join a growing number of high-tech firms forsaking cubicle culture for new workplaces designed to save space and money

tools to reveal the true colors of the renowned painting. His results indicate Leonardo da Vinci's original (finished around 1505) included eyebrows and eyelashes. Moreover, the painter not only repositioned the two fingers on her left hand but originally drew a more expressive smile. BBC News reports that Cotte scanned the painting in the Louvre's laboratory in 2004 using infrared and ultraviolet sensors during the process. He has now spent more than 3,000 hours analyzing data from those scans and his resulting 240-megapixel scan reveals traces of facial hair obliterated by restoration, as well as that Mona was holding a blanket that has all but faded from view. "With just one photo you go deeper into the construction of the painting and understand that Leonardo was a genius," says Cotte.

## HOW SCIENTISTS VIEW NANO

The unknown human health and environmental impact of nanotechnology is a bigger worry to scientists than to the general public, according to a report published in *Nature Nanotechnology*. Based on a survey of U.S. households and leading nanotechnology scientists and engineers, it appears that those with the most insight into the technology, despite its enormous potential, are also the most wary of the health and environmental problems it might pose. These findings are in stark contrast to controversies sparked by the advent of past technologies such as nuclear power and genetically modified foods, which scientists perceived as having lower risks than the public. In fact, the public is more concerned with the loss of privacy from new tiny surveillance devices and the potential loss of U.S. jobs.

while encouraging collaboration among co-workers. The *San Jose Mercury News* reports Intel has been testing three experimental work sites—in Arizona, California, and Oregon—where open areas, comfy armchairs, and extra conference rooms have replaced the cube-ism the chip giant helped popularize. Employees are logging onto a corporate network each morning using wireless connections; their phone numbers follow them wherever they go. Workers use whiteboards outfitted with electronics so drawings and business plans can be



transferred to laptops and emailed to colleagues. Open seating and laptops reign; desks and filing cabinets are old school.

Results of this three-month trial will determine if Intel rolls out the new design to the rest of the company. Cisco Systems's San Jose site (pictured here) concluded its first cube-free year. Employees contend working without assigned seating takes getting used to, but the "team" approach to projects has certainly improved. Sun Microsystems, Google, Intuit, and Network Appliance have also adopted forms of the wall-less work environment, noting significant improvement in collaborative efforts. **G**

Bottom: Gary Reyes/San Jose Mercury News

Send items of interest to [cacm@acm.org](mailto:cacm@acm.org)



# Stop Monitoring Legal Internet Traffic

I strongly disagree with the implication in Steven M. Bellovin et al.'s "Inside Risks" column "Internal Surveillance, External Risks" (Dec. 2007) that because the U.S. is a transit point for so much transcontinental international Internet traffic, U.S. government agencies should be able to intercept international communications transiting the U.S. Internet infrastructure. They said that if the source or destination of the traffic is in the U.S., then the agencies should discard the data before any "further processing is done." This seems to mean that international traffic is fair game. Perhaps a better solution would be if when the source and destination of the traffic are both in the U.S. or both are outside the U.S., then the agencies should discard the traffic before any "further processing is done."

The government has no right to monitor international traffic that legally passes through the U.S. The transit points are by nature as international as the United Nations or the World Bank, and their sovereignty must be respected.

The U.S. should thus set a standard for integrity by not

touching data with an international origin and destination that legally passes through U.S. transit points, unless it has a legal right to do so.

VIR V. PHOHA  
*Ruston, LA*

## *Authors' Response:*

While we may agree with Phoha's sentiment, it is settled law in the U.S. that the Constitution does not grant legal rights to noncitizens outside the U.S. Publishing in *Communications*, an international, technical publication, we wrote as technologists addressing a technical problem. The issue Phoha raises should certainly be addressed by the U.S. Congress, perhaps even by the legal system as well.

STEVEN M. BELLOVIN  
MATT BLAZE  
WHITFIELD DIFFIE  
SUSAN LANDAU  
JENNIFER REXFORD  
PETER G. NEUMANN

## **ACCOUNT FOR UNKNOWN, AS WELL AS KNOWN**

I would like to compliment *Communications* on the quality of two columns: "The Profession of IT" "The Choice Uncertainty Princi-

ple" by Peter J. Denning and the "Viewpoint" "Stop the Numbers Game" by David Lorge Parnas (both Nov. 2007).

The former reminds me of some of the fundamental issues and factors that sometimes stay hidden from most of us and thus are rarely if ever taken into consideration. I'm reminded of something along these lines that happened to me some years ago. An analog-to-digital convertor appeared to be malfunctioning, producing wrong numbers. When we looked it over, it appeared at every stage of the conversion process to be working perfectly, and in fact it was. The problem was very high (for the time) frequency noise on the conversion-completed-ready-to-read line, leading to the register with the number being read prematurely before stabilizing on the correct value.

The cause of the problem was detected only by accident because the time base had been set incorrectly on the oscilloscope. We never thought a premature read-out could be the cause of the apparent malfunction.

Parnas's "Viewpoint" interests me because my work has long

involved science and technology statistics. Some of the most heavily used (and unfortunately frequently abused) indicators are publications and patent counts and citations. They are easy to measure but only poor proxies for what one actually wants to know. The holy grail is an accurate measure of scientific and technological output. Publication counts, citation counts, and patent statistics are most certainly not the holy grail.

In his splendid “Viewpoint,” Parnas pointed out many of the reasons we should not simply be counting publications and/or citations in computer science (or in any other subject). He delivered one of the best critiques I’ve ever read of bibliometrics and the likely effects of its misuse. It would be so good if this message would get across to all those inclined to use (and abuse) such data.

IAN PERRY  
*Brussels, Belgium*

## **DILIGENCE-BASED SECURITY PREFERRED**

The article “Necessary Measures” about security risk assessment by Wade H. Baker et al. (Oct. 2007) took an excellent scenario-based approach to risk assessment rather than trying to assess a single limited type of attack to justify a particular countermeasure. However, the authors fell into the trap of describing security risk assessment through an incomplete example that ultimately was not really risk assessment. However, I do support their conclusion that valid security risk assessment is difficult and may not succeed due to a lack of valid data and excessive complexity.

The example was supposed to

involve a risk assessment for an organization afflicted by known frequency and impact of malware attacks through email, using a 2004 survey of the loss experience of 40 organizations. But the results (total cost \$32.7 billion without countermeasures, \$44,000 with countermeasures) are estimates of actual costs, not risks, as labeled in the article.

The considerable cost of risk assessment and the initial costs of countermeasures, along with various losses, seemed to have been omitted in the example calculations. Lost sales, reputation damage, and regulatory and legal fees were omitted for “pedagogical reasons.” Any of these losses, along with lost talent, project delays, and other business adversity, could exceed all other losses in a single incident.

Security risk deals with adversity that might or might not happen but was not the case in the example. The countermeasures against ongoing adversity can be tested on a trial basis (without risk assessment) and their effects measured to reach conclusions, justify further action, and calculate trade-offs. Meanwhile, the authors included another factor—countermeasure effectiveness—adding yet another dimension of uncertainty and complexity. They could have tried to measure residual risk against countermeasure effectiveness. Unfortunately, they identified only the effectiveness of deterrence, protection, detection, and recovery. Other important factors (such as avoidance, mitigation, audit, transference, insurance, motivation, and correction) should also be considered to

achieve the full value of the countermeasures.

Also not accounted for was the possibility that the identified countermeasures might be effective for many other types of attacks, thus potentially adding significant value to them. Moreover, reducing the risk of one type of attack could increase or reduce the risk of other types of attacks, depending on the strategies and tactics of unknown perpetrators that increase or decrease the organization’s overall security risk.

Another problem is the unknown complexity of the ways risks, frequencies, vulnerabilities, countermeasures, victims, and perpetrators affect one another; for example, a malware attack could result in an extortion attempt.

The main reason security risk is so difficult to measure is that no one can know how unknown perpetrators at unknown times in unknown circumstances might affect vulnerabilities, whether known or unknown. Though good security inhibits the sharing of confidential and potentially damaging loss and security experience information with outsiders, it also severely limits the collection of data.

Intangible risk-based security ought to be replaced with diligence-based security through known countermeasures, standards, laws, regulations, and competitive business practices.

DONN B. PARKER  
*Los Altos, CA*

## **A GOOD SERIAL PROGRAMMER**

In “Parallel Computing on Any Desktop” (Sept. 2007), Ami Marowka explored the principles

involved in Gant charts, CPM, and PERT. It's good to see these things being discussed, but I wonder about all this concern over how to implement multithreads in a multiprocessor environment. It's been done since the 1970s.

This is not to suggest that all parallel computing issues have been resolved. Where once multithreading was almost exclusively invoked by system software and primarily concerned efficient use of system resources, we are now moving that capability into applications to speed task throughput by having multiple parts of the task run simultaneously.

As Marowka framed it, a central issue is when and how to multithread and determine the throughput improvements that might be achieved. Invocation and completion of a thread yields processing overhead, so a poorly defined thread might be at best an exercise in futility or at worst hurt throughput. This introduces the case of the application that simply does not lend itself to multithreading.

I pretty much agree with Marowka, except that the "complexity" issue seemed to be overstated. My take is that a good serial programmer with, say, five years of experience already has the skills needed to take on multithreading.

**DON PURDY**  
*Rutherford, NJ*

---

Please address all Forum correspondence to the Editor, *Communications of the ACM*, 2 Penn Plaza, Suite 701, New York, NY 10121-0701; email: [crawfordd@acm.org](mailto:crawfordd@acm.org).

---

DOI: 10.1145/1314215.1314217

Coming next in

# COMMUNICATIONS

of the ACM

## March 2008

Cover Story: *Urban Sensing: Out of the Woods*

In the world of personal communication, pervasive computing has entered every backpack, purse, and coat pocket in the form of a mobile phone. This contextual swing is characterized as "urban sensing," and it promises a fundamental shift from science and engineering into the realms of politics, aesthetics, interpretation, and motivation. A move, industry observers contend, that warrants careful scrutiny.

## April 2008

Cover Story: *The Psychology of Security*

The most elegant and intuitively designed interface does not improve security if users ignore warnings, choose poor settings, or unintentionally subvert corporate policies. The fundamental user problem in security systems involves how people think of risk and how those thoughts guide their behavior. Risk and uncertainty are extremely difficult concepts for people to evaluate. For designers of security systems, it is important to examine how users evaluate security risks and understand why good users make bad decisions.

## May 2008

Special Section: *Organic User Interfaces*

There's a quiet revolution in display manufacturing technology taking place. New technologies have led to displays that are so thin and flexible they resemble paper. Indeed, we are on the brink of a new paradigm in computer user interface design: one in which computers can have any organic form or shape. One where any object, no matter how complex, dynamic, or flexible its structure, may display information.



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*

# ACM, Uniting the World's Computing Professionals, Researchers, Educators, and Students



Dear Colleague,

At a time when computing is at the center of the growing demand for technology jobs worldwide, ACM is continuing its work on initiatives to help computing professionals stay competitive in the global community. ACM delivers resources that advance computing as a science and profession.

As an ACM member you have access to leading-edge resources including publications and conferences covering the entire computing field. Whether you are pursuing scholarly research, building systems and applications, or managing computing projects, ACM offers opportunities to advance your interests.

## MEMBER BENEFITS INCLUDE:

- Access to ACM's **new Career & Job Center** offering a host of exclusive career-enhancing benefits
- **Free e-mentoring services** provided by MentorNet®
- **Free and unlimited access to 2,500 online courses** from SkillSoft®
- **Free and unlimited access to 600 online books** from Safari® Books Online, featuring leading publishers, including O'Reilly (Professional Members only)
- **Free and unlimited access to 500 online books** from Books24x7®
- A subscription to ACM's flagship monthly magazine, **Communications of the ACM**
- The option to subscribe to the full **ACM Digital Library**
- The **Guide to Computing Literature**, with over one million searchable bibliographic citations
- The option to connect with the **best thinkers in computing** by joining **34 Special Interest Groups** or **hundreds of local chapters**
- **ACM's 40+ journals and magazines** at special member-only rates
- **TechNews**, ACM's tri-weekly email digest delivering stories on the latest IT news
- **CareerNews**, ACM's bi-monthly email digest providing career-related topics
- **MemberNet**, ACM's e-newsletter, covering ACM people and activities
- **Email forwarding service & filtering service**, providing members with a free acm.org email address and high-quality **Postini** spam filtering
- And much, much more!

ACM's worldwide network of more than 83,000 members range from students to seasoned professionals and includes many of the leaders in the field. ACM members get access to this network and the advantages that come from their expertise to keep you at the forefront of the technology world.

Please take a moment to consider the value of an ACM membership for your career and your future in the dynamic computing profession.

Sincerely,

A handwritten signature in blue ink that reads "Stuart I. Feldman". The signature is written in a cursive style and is positioned above the printed name and title.

Stuart I. Feldman  
President  
Association for Computing Machinery



Association for  
Computing Machinery

*Advancing Computing as a Science & Profession*



Association for  
Computing Machinery

Advancing Computing as a Science & Profession

# membership application & digital library order form

Priority Code: ACACM28

## You can join ACM in several easy ways:

### Online

<http://www.acm.org/join>

### Phone

+1-800-342-6626 (US & Canada)  
+1-212-626-0500 (Global)

### Fax

+1-212-944-1318

Or, complete this application and return with payment via postal mail

### Special rates for residents of developing countries:

<http://www.acm.org/membership/L2-3/>

### Special rates for members of sister societies:

<http://www.acm.org/membership/dues.html>

Please print clearly

Name \_\_\_\_\_

Address \_\_\_\_\_

City \_\_\_\_\_ State/Province \_\_\_\_\_ Postal code/Zip \_\_\_\_\_

Country \_\_\_\_\_ E-mail address \_\_\_\_\_

Area code & Daytime phone \_\_\_\_\_ Fax \_\_\_\_\_ Member number, if applicable \_\_\_\_\_

### Purposes of ACM

ACM is dedicated to:

- 1) advancing the art, science, engineering, and application of information technology
- 2) fostering the open interchange of information to serve both professionals and the public
- 3) promoting the highest professional and ethics standards

I agree with the Purposes of ACM:

Signature \_\_\_\_\_

ACM Code of Ethics:

<http://www.acm.org/serving/ethics.html>

## choose one membership option:

### PROFESSIONAL MEMBERSHIP:

- ACM Professional Membership: \$99 USD
- ACM Professional Membership plus the ACM Digital Library: \$198 USD (\$99 dues + \$99 DL)
- ACM Digital Library: \$99 USD (must be an ACM member)

### STUDENT MEMBERSHIP:

- ACM Student Membership: \$19 USD
- ACM Student Membership plus the ACM Digital Library: \$42 USD
- ACM Student Membership PLUS Print CACM Magazine: \$42 USD
- ACM Student Membership w/Digital Library PLUS Print CACM Magazine: \$62 USD

All new ACM members will receive an  
ACM membership card.  
For more information, please visit us at [www.acm.org](http://www.acm.org)

Professional membership dues include \$40 toward a subscription to *Communications of the ACM*. Member dues, subscriptions, and optional contributions are tax-deductible under certain circumstances. Please consult with your tax advisor.

### RETURN COMPLETED APPLICATION TO:

Association for Computing Machinery, Inc.  
General Post Office  
P.O. Box 30777  
New York, NY 10087-0777

Questions? E-mail us at [acmhelp@acm.org](mailto:acmhelp@acm.org)  
Or call +1-800-342-6626 to speak to a live representative

**Satisfaction Guaranteed!**

### payment:

Payment must accompany application. If paying by check or money order, make payable to ACM, Inc. in US dollars or foreign currency at current exchange rate.

Visa/MasterCard     American Express     Check/money order

Professional Member Dues (\$99 or \$198)    \$ \_\_\_\_\_

ACM Digital Library (\$99)    \$ \_\_\_\_\_

Student Member Dues (\$19, \$42, or \$62)    \$ \_\_\_\_\_

**Total Amount Due**    \$ \_\_\_\_\_

Card # \_\_\_\_\_ Expiration date \_\_\_\_\_

Signature \_\_\_\_\_

# SIGs Announce Candidates for Election

In accordance with ACM Bylaw 6, the following SIGs will hold elections in 2008: SIGCAS; SIGGRAPH; SIGSIM, and SIGUCCS.

ACM Policy and Procedures require that those SIGs holding elections notify their membership of candidates for elected offices. Here is a list of SIGs that have submitted their slate of candidates. Additional information will appear in the March issue of *Communications*.

---

## SIGCAS

### EXECUTIVE COMMITTEE

Florence Appel  
Diana Burley  
Lani Fraizer  
Benjamin Gerber  
Joseph D. Oldham  
Mark Perry  
David G. Robinson  
Steve Sawyer  
Bruce F. Webster

## SIGGRAPH

### Chair

Anthony (Tony) Baylis  
G. Scott Owen

### Vice-Chair

Paul Debevec  
James Foley

## Director at Large

Marie-Paule Cani  
James F. O'Brien

## SIGSIM

### Chair

Osman Balci  
George F. Riley

### Vice-Chair

David W. Bauer  
John A. "Drew" Hamilton, Jr.

### Secretary/Treasurer

Margaret Loper  
Levent Yilmaz

## SIGUCCS

### Chair

Teresa Lockard  
Bob Haring-Smith

### Board of Directors

Nancy J. Bauer  
Timothy Foley  
Jim Bostick  
Diane Jung  
Pat Kohrman  
Patti Mitch  
Alex Nagorski  
Gail Farally-Semerad  
Steven Timmins  
Kelly Wainwright

---

In accordance with the SIG Bylaws, additional candidates may be placed on the ballot by petition. All candidates must be Professional members of ACM, as well as members of the SIG. Anyone interested in petitioning must inform ACM Headquarters (Pat Ryan, ACM, 2 Penn Plaza, Suite 701, NY, NY 10121) and the SIG Secretary of their intent to petition by March 15th.



# Managing Software Development in Globally Distributed Teams

Surveying the basic elements needed for globally distributed teams to function smoothly and efficiently.

Many companies in the U.S., Europe, and Japan are now building or contracting for new software products, components, and custom-designed systems that will be created in globally distributed teams, such as in India or China, and increasingly in Russia, Eastern Europe, Southeast Asia, and Latin America. While taking advantage of lower costs and broader talent pools around the world, the problem with this trend is that the world is not exactly “flat,” as Thomas Friedman of the *New York Times* has asserted. Many companies send programming, engineering, and technical back-office work overseas and get excellent results. But software development can be tricky in many circumstances, which is why globally distributed teams require special skills in system design and project management.<sup>1</sup>

One of the biggest challenges is to apply an iterative (or more popularly called “agile”) style of development when teams are large and not co-located. This is important because a lot of the research over the past decade on software engineering and product development in general has concluded that itera-

the daily builds and milestone approaches of firms like Microsoft to the more nontraditional practices of extreme programming (see my October 2007 column). We also know that, as long as projects do things like show early prototypes to customers, run regression and integration tests on each of



ative approaches have lots of benefits, in contrast to traditional waterfall or sequential methods of developing new products. Iterative techniques make it possible for projects to respond more easily and, while they are in progress, to changing demands from customers or unforeseen challenges from competitors.<sup>2</sup> Iterative software engineering practices range from

their builds, and do frequent design reviews, they can accommodate a lot of change while still keeping quality and budgets under control [1].

But these and other iterative techniques become more difficult

<sup>1</sup>For an excellent treatment of this subject, see E. Carmel, *Global Software Teams: Collaborating Across Borders and Time Zones*. Prentice-Hall, Englewood, NJ, 1998. Full disclosure: I also have had exposure to these processes as a director of a public Indian software company, Patni—NYSE symbol PTI—for the past several years.

<sup>2</sup>See, for example, B. Boehm, “A View of 20th and 21st Century Software Engineering.” In *Proceedings of the 28th International Conference on Software Engineering* (Shanghai, China, 2006).

to use consistently when teams are geographically distributed in different countries and cultures around the world. On the technical side, we often see unclear or too rapidly changing requirements, delays, and mistakes in shifting from design to implementation, unforeseen technical interdependencies that thwart integration of components, and bug-ridden code caused by these and other problems. On the organizational side, there are varying communications, as well as development styles in different countries, as well as differing degrees of facility with the English language. In response, many companies resort to sequential waterfall-like planning and project management when they distribute software or other engineering work globally under the assumption this is a simpler way to deal with communication and coordination problems. This may be true sometimes, such as in embedded software development projects where specifications need to be tightly fixed. But if companies want to maximize their ability to respond to customer needs and changing market conditions, they still need to figure out how to incorporate iterative development techniques.

Managing global teams is common today but is not a new problem. IBM, for example, developed OS/360 in the 1960s and successor operating systems in multiple sites around the U.S. and Europe. Today, though, it seems that Indian software firms, led by Infosys, Tata, Wipro, Satyam, and a

few others, have taken the lead in refining a “global delivery model.”<sup>3</sup>

The best practices I am familiar with begin with an iterative process, adopt waterfall-like practices in the middle implementation phase, and then move back to iterative to complete the project. They work roughly as follows: The software provider (such as an Indian IT services vendor or perhaps Accenture or IBM Global Services) sets up a project office at or near the customer site. Experienced requirements engineers and project leads work closely with the customer to generate a scoping document, preliminary requirements, and even prototypes of part of the product or system functionality.

The next step is to produce more detailed requirements for critical parts of the system, as well as a preliminary development plan and budget. Later, the team based in India (or China or Russia) will do detailed design, coding, unit testing, integration, and system testing—not usually in one phase but in a series of milestones or subprojects representing major parts of the system or components being built. Eventually, the team will bring the completed system back to the customer for final acceptance testing and start iterating again to get the final details right. They should know

<sup>3</sup>For a description of Infosys’s global delivery model, see [www.infosys.com/gdm/default.asp](http://www.infosys.com/gdm/default.asp). For a comparison of Indian and other software projects, see M. Cusumano et al., “Software Development Worldwide: The State of the Practice,” *IEEE Software* 20, 6 (Nov./Dec. 2003), 28–34.

where they are in the project if the development team has been working with the customer closely all along.

But before a software company can make these handoffs and handbacks work relatively smoothly, the globally distributed team must have several basic elements in place.

First, and perhaps the most difficult, is to create an iterative contract with the customer—if possible. Customers often want a fixed price and schedule, especially when they are dealing with low-cost suppliers competing for their business. But, for the software provider, too many things can still go wrong in any large project, and it is easy to lose money on fixed-price contracts if the work spins out of control or if the initial estimating was poorly conducted. One strategy is to persuade the customer to agree to an initial scoping and prototyping engagement lasting from a few days to a few weeks, depending on the size or complexity of the task. After that effort, the software provider should be able to offer a more precise estimate of the cost and time required for the project. Most contracts will then require the customer to pay for the work incrementally as the software is built. But exactly how much to pay could be determined as the work is planned in more detail and as the customers learn more about exactly what they want.

It is essential to do frequent updates with the customer on the state of the project—show the cus-

tomer each week what the team is doing, how far they have progressed versus the schedule milestone, and what the evolving product or system actually looks like. It is critical to keep iterating back and forth between the customer and the developers offshore rather than simply attempt a one-time handoff. No one gets it right the first time. But it is equally important to put some time limits on the changes and move forward in discrete milestones or subprojects so the iterations are not endless.

For these steps to work well, it is essential to break up the system into clearly defined subsystems, components, or increments that separate teams working in parallel can design, build, and test as independent modules, at least to some degree. Each increment also should take no more than a couple of weeks to complete and then integrate with other parts of the system under development; otherwise misunderstandings between the requirements engineers and the programmers and testers will persist and grow. This kind of decomposition or architectural design is easier said than done. It requires considerable thought and planning up front in a project, as well as a deep domain understanding of the type of system or product the customer wants developed. But, with adequate modularity, a distributed project team can have several small groups work independently; they can accommodate frequent changes and later bring their work together—more or less successfully.

In order to keep teams synchronized, the project managers must do what all good iterative projects do, whether they are located in one site or in a dozen—employ short design, build, and test cycles of a couple of days or weeks at the most. Then the team members need to review their designs and code frequently, as well as test their builds with other components from other sites, as often as possible. Some distributed teams prefer to work with what are really mini-waterfall cycles of, say, a week or two and then do their reviews, integration, stabilization (bug fixing), and system testing work intensively for each cycle. Whatever the scheduling, the key is to keep the teams synchronized and the components as de-coupled or as modular as possible to reduce interdependencies and integration “gridlock,” also known as “infinite defect loops.”

Globally distributed projects must involve participants with strong organizational and process skills. It is generally very useful to have a strong project manager leading the entire effort, to have a commonly accepted development process (such as an iterative version of CMM), and to have occasions for the team members to meet each other and communicate directly to solve major issues that might occur.

It is useful to have tools that support distributed development and large teams. Today, there are many in-house and commercial tools available to facilitate daily builds, continuous regression and

integration testing, and version control, as well as configuration management. Distributed teams also rely on email and videoconferencing to communicate. With high-speed data connections, the tools that work for a co-located project also can work for a globally distributed project.

We have advanced far beyond the days when IBM engineers would travel the world carrying reels of magnetic tape whenever they wanted to do a system build for pieces of software made in different locations. So, yes, with the Internet, we can distribute work globally in a way never before possible or economical. But our thinking and management skills still need to catch up with the technology and global realities. The world is flatter than it once was, but when it comes to software development, there are still many potential bumps in the road. Fortunately, we generally know where the bumps are located. **C**

#### REFERENCE

1. MacCormack, A. et al. Trade-offs between productivity and quality in selecting software development practices. *IEEE Software* 20, 5 (Sept./Oct. 2003), 78–85.

---

**MICHAEL CUSUMANO** (cusumano@mit.edu) is a professor at the MIT Sloan School of Management and author of *The Business of Software*, Free Press/Simon and Schuster, 2004.

---

© 2008 ACM 0001-0782/08/0200 \$5.00

---

DOI: 10.1145/1314215.1314218

# DIS 2008

CAPE TOWN · SOUTH AFRICA

*designing interactive systems*

25-27 February 2008

[www.sigchi.org/dis2008](http://www.sigchi.org/dis2008)



SIGCHI



UNIVERSITY OF CAPE TOWN



Association for  
Computing Machinery

**NOKIA**  
Connecting People

**SAP**

**Microsoft**

DESIGN  
INDABA

# Peer to Patent Needs Your Expertise

You can help ensure that computer-related patents are properly reviewed and stop the worthless ones from being granted.



**T**echnologists rarely have an opportunity to influence social and legal policy directly. Peer to Patent ([www.peertopatent.org](http://www.peertopatent.org)) is one such opportunity. It addresses the proliferation of computer-related patents that

sometimes hinder innovation rather than promote it—while offering technologists a chance to guide the system toward healthy equilibrium in a way that makes exemplary use of their expertise. Peer to Patent is a project that every *Communications* reader should at least visit; many will find participation important and rewarding, spending as much or as little time as they'd like.

ACM's mission calls on its members to lend their expertise to help policymakers and institutions determine the proper role of technology. The U.S. patent system is a particularly apt context for exercising this mandate, because inventors in the computer field live under the constant fear of lawsuits based on patents that should never have been granted in the first place [1, 2].

One example of the alarming disruptions that patents can wreak on the spread of computer technology was a 2006 legal clash in which patents exerted by NTP, a patent-holding company, nearly shut down the popular email service on BlackBerry mobile devices. Most of these patents were eventually overturned due to the existence of prior art (evidence that the ideas being patented were in use earlier by

other inventors). Yet these rulings were too late for BlackBerry developer Research in Motion, which eventually paid more than \$600 million to NTP to resolve the lawsuit, a burden compounded by untold legal fees and possible lost business due to public concern over the product's continuing viability.

What if programmers and computer scientists who knew how commonplace the NTP inventions were could have submitted prior art to the patent office, nipping the applications in the bud and averting wasteful lawsuits? That's exactly what Peer to Patent promises for the patent applications available for review on its Web site.

Launched last year in cooperation with the U.S. Patent and Trademark Office, Peer to Patent sets up a well-structured forum at [peertopatent.org](http://peertopatent.org), with training and group discussion, that channels input from the public directly to the USPTO's examiners who decide what patents to grant. As a one-year pilot, Peer to Patent began July 15, 2007 at [www.uspto.gov](http://www.uspto.gov). Its counterparts in Britain and the European Union are now gearing up for similar experiments testing the public's ability and willingness to contribute to the patent-examination process and thereby justify making it more open in the future. Corporate sponsors include General Electric, Hewlett-Packard, IBM, Microsoft, and Red Hat.

Recognizing that computer professionals are not professional patent examiners, the [peertopatent.org](http://peertopatent.org) Web site is attractively designed for efficient participation. You log in, view the full text and figures of patents, offer comments, and submit pointers to

prior art that would help an examiner determine whether an invention merits a patent. Discussion in the site's forums is constructive and focused. Rating prior-art submissions allows the strongest and most relevant ones to rise to the top. Moderators remind participants of the site's goals and culture. In addition to finding and voting on prior art, these forums are places where potential employers and employees can meet one another; so too can potential collaborators looking to staff research projects.

It's important to understand that a patent examiner—an impartial government employee—still makes the ultimate decision as to whether to award a patent. The system is designed to resist gaming. If, say, a company's competitors want to jump on its patents, they can still go right ahead. Prior art is prior art no matter who finds it.

Go to [peertopatent.org](http://peertopatent.org) and click the "View an Application" link for a patent on any familiar-looking technology. Reading patent applications can be intimidating at first; fortunately, resources at [peertopatent.org](http://peertopatent.org) will get you started. I suggest simply jumping in and reading a particular patent's description. Keep a second tab or window open on the browser so you can switch back and forth between text and figures. Believe it or not, a patent is supposed to teach you how to re-create its invention. True, the lawyers who write the applications often seem not to take this responsibility seriously. But most descriptions are decipherable, and after you understand the description, you can deal with the claims that define the ideas to which the applicant wants exclusive rights.

Some patent applications will probably prompt you to think "I'm no expert in this area, but I know colleagues who are." Contacting them is one of the best contributions a participant can make to the system. You can also encourage companies to submit patents for review, earning the reward of expedited handling in the patent office.

Peer to Patent does not represent a total solution to the world's patent problems. Jason Schultz, a staff attorney at the Electronic Frontier Foundation, describes himself as "a big fan" but says, "I don't think the Peer to Patent project alone will fix

what's wrong with the system." Patent lawyers I interviewed are supportive and cautiously optimistic. The computer science researchers I interviewed, while generally viewing the project as valuable, aired some skepticism. For instance, Jon Bentley of Avaya Labs, called the project "an exciting proposal" but warned that people typically volunteer for projects "unencumbered by cash flow," whereas "the essence of patents (well, one essence) is exactly money."

The sheer volume of patent applications (which has grown to several thousand computer-related applications a year in the U.S. alone) could strain the capacity of Peer to Patent volunteers if the system expands from a pilot to become common practice. In addition, some opponents of software patents believe that the act of reviewing them also legitimizes them. If software is indeed a poor fit for the kinds of research and rewards propelled by the patent system, patents can be a hindrance, even if they pass the prior-art test.

There are counterarguments to these objections, but Peer to Patent should be appealing largely because it's an immediate intervention with measurable outcomes. Patent applications will not receive the reviews they need until researchers see them as part of their personal responsibility to society, in the same way they spend time refereeing journal papers, sitting on the boards of nonprofit organizations, and performing university-related duties unrelated to their own central research tasks.

Get a feel for what this experience would be like by giving Peer to Patent a test run. The promising start-up you save from a lawsuit may be your own.

## REFERENCES

1. Abril, P. and Plant, R. The patent holder's dilemma: Buy, sell, or troll. *Commun. ACM*, 50, 1 (Jan. 2007), 36–44.
2. Thatcher, M. and Pingry, D. Software patents: The good, the bad, and the messy. *Commun. ACM* 50, 10 (Oct. 2007), 47–52.

---

**ANDY ORAM** ([andy@oreilly.com](mailto:andy@oreilly.com), [www.praxagora.com/andy](http://www.praxagora.com/andy)) is an editor at O'Reilly Media and a member of Computer Professionals for Social Responsibility.

---

© 2008 ACM 0001-0782/08/0200 \$5.00

---

DOI: 10.1145/1314215.1314219

# How to Achieve (Some) Balance in Anti-Circumvention Laws

Should users of technically protected content be able to give notice and require content owners to take down the technical measure to enable fair uses?

The World Intellectual Property Organization Copyright Treaty (WCT), concluded in 1996, recognizes “the need to maintain a balance between the rights of authors and the larger public interest, particularly education, research, and access to information” in updating international copyright norms to respond to challenges arising from advances in information and communications technologies, including global digital networks.

The treaty calls upon nations to “provide adequate legal protection and effective legal remedies against the circumvention of effective technological measures that are used by authors in connection with the exercise of their rights,” although such rules should not impede acts that are “permitted by law” or otherwise beyond the authority of copyright owners. Unfortunately, the treaty gives no guidance about how nations might implement an anti-circumvention norm that can enable public interest uses of copyrighted works.

## IMBALANCED NATIONAL IMPLEMENTATIONS


While the WCT embodies a negotiated balance between copyright owners and users of digital works, the translation of this balance into the domestic laws of the U.S. and European Union (EU) has been incomplete. When enacting the Digital Millennium Copyright Act (DMCA) of 1998 as the U.S. implementation of the WCT, Congress achieved a reasonable balance of competing interests in its creation of safe harbors from copyright liability for Internet service providers and other intermediaries for the infringing acts of others.

Congress failed to achieve a similar balance of interests, however, when establishing new rules

forbidding circumvention of technical protection measures (TPMs) used by copyright owners to control access to their works and in regulating the manufacture and distribution of technologies primarily designed or produced to enable circumvention of copyright-protective TPMs. Much the same can be said of the European implementation of the WCT.

The resulting imbalance in U.S. and European anti-circumvention rules harms legitimate interests of the public in making fair uses, privileged uses, and other non-

infringing uses of copyrighted works (collectively designated as “public interest uses”). We believe that practical judicial and administrative measures can and should be devised to implement the balanced spirit of the WCT without



reopening the contentious debates leading up to enactment of the DMCA and the EU Directive.

### THE REVERSE NOTICE AND TAKEDOWN CONCEPT

To this end, we propose a “reverse notice and takedown” procedure to achieve some of the balance in anti-circumvention rules the WCT endorsed, but which implementing legislation has thus far failed to deliver. Under this regime, users would be able to give copyright owners notice of their desire to make public interest uses of technically protected copyrighted works, and rights holders would have the responsibility to take down the TPMs or otherwise enable these lawful uses.

We call this a reverse notice and takedown procedure because, in an inversion of the notice and takedown procedure first developed through common law adjudication to determine ISP liability for wrongful acts of users, it is the user who will be giving notice and the content owner who will have a responsibility to take something down. It would effectuate the nascent, but not fully realized, legislative intent to permit public interest uses of technically protected digital content, while at the same time protecting copyright owners against circumvention of TPMs likely to facilitate massive infringements.

A reverse notice and takedown mechanism is a principle capable of numerous implementations. In the U.S., this goal can be achieved through judicial interpretation of

the anti-circumvention rules on a case-by-case basis. It was, after all, the judicial branch that introduced the fair use doctrine into U.S. law and also pioneered the notice and takedown rules to govern ISP liability.

In the heated political climate in which the DMCA was enacted, the measured analysis developed in *Religious Technology Center v. Netcom* was invaluable in shaping ISP liability rules. Unfortunately, no similarly careful judicial assessment was available in the late 1990s to guide Congress about how to achieve an appropriate balance in the anti-circumvention rules.

We believe courts in the U.S. can and should be enlisted in bringing about a balanced approach for dual-use circumvention technologies akin to that developed for the dual-use technologies and services of ISPs. Recent decisions, moreover, provide a basis for this case law evolution. EU member states could implement the reverse notice and takedown regime in their national implementations of the EU Directive.

### CIRCUMVENTION FOR FAIR USES?

The case law in the U.S. is mixed on the question whether it is lawful to circumvent a TPM for fair use purposes. The first judicial opinion to consider this question, *Universal City Studios, Inc. v. Reimerdes*, concluded that Congress had considered, and decided against, allowing circumventions

for fair use or other privileged purposes. It further opined that anti-circumvention rules were consistent with the First Amendment guarantee of free speech, even though the U.S. Supreme Court has said fair use and other limitations are important for the compatibility of copyright and freedom of expression.

Of course, neither Reimerdes nor his co-defendant Eric Corley had attempted to bypass the Content Scramble System (CSS) to make fair uses of DVD movies; rather they had posted DeCCS, software designed to bypass CSS, on their Web sites as part of a protest against the anti-circumvention rules. Judge Kaplan ruled that posting and linking to DeCCS violated the DMCA anti-circumvention rules. Judge Kaplan regarded the anti-circumvention laws to have “fundamentally altered the landscape of copyright” as to technology provider liability.

*Reimerdes* set forth a framework for analyzing anti-circumvention claims that, if followed in subsequent cases, would exclude consideration of virtually all public interest concerns. Under Judge Kaplan’s interpretation, anti-circumvention liability would arise: if a copyright owner has adopted a TPM to control access to its copyrighted works (even if they are persistent access controls such as CSS); and if an unauthorized person has developed a technology that bypassed this TPM (relying, if necessary, on an inference that if the defendant’s technology



bypasses the TPM, it must have been primarily designed or produced to do so).

Under this decision, it is irrelevant whether copyright infringement has occurred—or is even possible—from the availability of a challenged technology. Nor does it matter whether the tool might enable consumers to tinker with a copyrighted work they have purchased. Harm to the copyright owner's interests was presumed.

#### **CIRCUMVENTING LOCK-OUT MEASURES**

It did not take long after the *Reimerdes* decision for canny technology developers to realize that the anti-circumvention rules, as interpreted in that case, could be used to attack competitors in the market for uncopyrightable products and services.

Lexmark, a manufacturer of printers and toner cartridges, claimed that the authentication protocol (or digital handshake) component of copyrighted computer programs installed on chips in its printers and toner cartridges was an access control. Because Static Controls made chips designed and produced to bypass this access control for other printer cartridge manufacturers, Lexmark charged it with violating the DMCA. The trial court, relying heavily on *Reimerdes*, issued a preliminary injunction against Static Controls' manufacture of these chips.

The Sixth Circuit eventually reversed, seemingly on the ground the DMCA does not apply to digital fences limiting access to func-

tional aspects of the printers. The court's reasoning on the anti-circumvention claim is, unfortunately, neither very coherent nor persuasive.

A concurring judge would more forthrightly have invoked the misuse doctrine, so as to "make clear that in the future companies like Lexmark cannot use the DMCA in conjunction with copyright law to create monopolies of manufactured goods for themselves just by tweaking the facts of [a] case."

Chamberlain, the maker of electronic garage-door opening (GDO) devices, made an analogous attempt to use the anti-circumvention rules to foreclose competition in the market for GDOs. It sued Skylink for DMCA violations because it made a universal GDO that bypassed the digitized "lock-out" (access control) components of programs Chamberlain had installed in its GDOs and transmitters. Chamberlain argued that the "plain language" of the DMCA and precedents such as *Reimerdes* and the lower court decision in *Lexmark* provided compelling support for its claim against Skylink. The Federal Circuit strongly disagreed.

#### **FINDING BALANCE IN THE DMCA**

A fundamental premise of the Federal Circuit's interpretation was its perception that Congress had intended the DMCA anti-circumvention rules to be balanced:

The most significant and consistent theme running throughout

the entire legislative history of the anti-circumvention and anti-trafficking provisions of the DMCA...is that Congress attempted to balance competing interests, and "endeavored to specify, with as much clarity as possible, how the right against anti-circumvention would be qualified to maintain balance between the interests of content creators and information users." [citation omitted] The Report of the House Commerce Committee concluded that § 1201 "fully respects and extends into the digital environment the bedrock principle of 'balance' in American intellectual property law for the benefit of both copyright owners and users."

The court observed:

Statutory structure and legislative history both make clear that § 1201 applies only to circumventions reasonably related to [copyright] protected rights. Defendants who traffic in devices that circumvent access controls in ways that facilitate infringement may be subject to liability under § 1201(a)(2)...[D]efendants whose circumvention devices do not facilitate infringement are not subject to § 1201 liability.

Without proof of a nexus between the availability of an allegedly unlawful circumvention tool and the existence, or grave threat, of copyright infringement, § 1201 liability should not be imposed. Thus, it was relevant that Chamberlain had not alleged that Skylink's GDO infringed its copyrights or contributed to any

other infringement.

Under Chamberlain's interpretation of the DMCA, "the owners of a work protected *both* by copyright *and* a technological measure that effectively controls access to that work ... would possess *unlimited* rights to hold circumventors liable ... *merely for accessing that work* even if that access enabled only rights that the Copyright Act grants to the public." The Federal Circuit found this construction of the DMCA "problematic for a number of reasons."

For one thing, Chamberlain's construction of the DMCA "borders on the irrational." Construing the DMCA as though it was concerned only about control over access, and not about rights protected by copyright law, would, moreover, be "both absurd and disastrous," for it would "allow any manufacturer of any product to add a single copyrighted sentence or software fragment to its product, wrap the copyrighted material in a trivial 'encryption' scheme, and thereby gain the right to restrict consumers' rights to use its products in conjunction with competing products." This would "allow virtually any company to attempt to leverage its sales into aftermarket monopolies," even though this would be unlawful under the antitrust laws and the copyright misuse doctrine.

### **PRESERVING FAIR USES IN THE DMCA**

At least as problematic were the implications of Chamberlain's

interpretation of the DMCA for the rights of consumers to make fair uses:

Chamberlain's proposed construction would allow copyright owners to prohibit exclusively fair uses even in the absence of any feared foul use. It would therefore allow any copyright owners through a combination of contractual terms and technological measures, to repeal the fair use doctrine with respect to an individual copyrighted work—or even selected copies of that copyrighted work ... Consumers who purchase a product have the inherent legal right to use that copy of the software. What the law authorizes, Chamberlain cannot revoke.

**C**ontrary to Chamberlain's contention, "the DMCA emphatically did not 'fundamentally alter' the legal landscape governing the reasonable expectations of consumers or competitors; did not 'fundamentally alter' the ways that courts analyze industry practices; and did not render the pre-DMCA history of the GDO industry irrelevant." The Federal Circuit consequently rejected Chamberlain's interpretation of § 1201 "in its entirety."

The Federal Circuit deserves considerable praise for recognizing that balance is a bedrock principle of intellectual property law and for developing a framework for interpreting § 1201 that enables courts to develop a balanced approach to interpretation of the DMCA's anti-circumvention rules

insofar as copyright owners try to use them to block fair and other non-infringing uses of technically protected copyrighted works. Courts interpreting § 1201 should reject *Reimerdes'* unbalanced and overly broad interpretation of § 1201 in favor of the framework set forth in *Chamberlain* and *StorageTek*, which we believe is far more consistent with the letter and spirit of the WCT and with Congressional intent in enacting the anti-circumvention rules.

### **FROM CHAMBERLAIN TO REVERSE NOTICE AND TAKEDOWN**

Building on the insights of *Chamberlain*, courts faced with public interest challenges to the DMCA anti-circumvention rules should develop a notice and takedown approach to facilitate a balancing of the interests of copyright owners and the public in order to enable privileged uses of technically protected works.

Under our proposed approach, a user wanting to make non-infringing uses of technically protected material could demand a right to bypassing of TPMs for legitimate purposes, such as extracting a fair use clip of a movie in order to complete a specified non-infringing project. Copyright owners could be given 14 days either to deny the limited circumvention proposal or to allow it by silence, without prejudice. In case of denial, the user would be entitled to seek a declaratory judgment to vindicate its claim to an entitlement to circumvent a TPM for the

purpose of engaging in specified non-infringing uses.

To become fully operational, this proposal would benefit from standardized procedures concerning the form in which notice should be given to copyright owners in a “reverse notice and take down” demand. It would also require courts to allow those providing needed decryption skills and technology to benefit from the same privileged use exception that a demandant had ultimately vindicated either in court or by silent acquiescence of the copyright owner. Above all, such a regime would particularly benefit from the kind of expeditious, low-cost administrative tribunals proposed in other contexts.

These long-term considerations should not, however, obscure the feasibility or desirability of immediately instituting ad hoc case-by-case judicially devised reverse notice and take down procedures to promote the formation of a jurisprudence of permissible non-infringing uses of technically protected content to complement and supplement the jurisprudence of infringing uses discussed here. Reverse notice and takedown procedures could attenuate the tension between the DMCA anti-circumvention rules and the public interest uses that Congress meant to preserve.

#### **BENEFITS OF REVERSE NOTICE AND TAKEDOWN**

A reverse notice and takedown procedure would contribute to making the DMCA into an instrument that adequately pro-

protects copyrighted works against circumvention that would lead to infringement without creating barriers to entry that thwart new technologies for sharing non-protected matter. It could facilitate licensing to non-profit entities on reasonable terms and conditions, and it could help to frustrate growing tendencies to put public domain matter off limits by encasing it in impenetrable electronic fences.

It could also attenuate the systematic use of digitized, electronic prior restraints on speech, which are likely to eventually provoke constitutional challenges. Indeed, an extension of the reverse notice and take down model could present would-be users of public domain material with a workable choice between sustaining the costs of securing and implementing judicially approved circumvention or purchasing the public domain matter from the vendor at reasonable prices for the sake of convenience.

Once the courts develop normative baselines for dealing with reverse notice and takedown requests an administrative procedure could evolve over time to apply and refine this normative framework. This development could also induce copyright owners to engage in private initiatives consistent with this framework, such as designating circumvention services to which putative public interest users might go to achieve circumvention for non-infringing purposes.

We believe that courts will be able to discern putative public

interest users who are not acting in good faith in making reverse notice and takedown requests and to put in place safeguards to ensure that the reverse notice and takedown regime does not bring about the increased infringements that the DMCA was enacted to avoid.

#### **CONCLUSION**

Although it is not possible in either the U.S. or the EU to write anti-circumvention rules on a completely blank slate, there is flexibility in the legal cultures of both entities to implement a reverse notice and takedown procedure to achieve needed balance in anti-circumvention regulations. Nations that have yet to implement the WCT may find our proposed reverse notice and takedown regime provides a far more balanced way to comply with the treaty than the approach being promoted by U.S. trade negotiators. **C**

---

**PAMELA SAMUELSON** (pam@ischool.berkeley.edu) is the Richard M. Sherman Distinguished Professor of Law and Information at the University of California, Berkeley.

**JEROME H. REICHMAN** (reichman@law.duke.edu) is the Bunyan Womble Professor of Law at Duke Law School.

**GRAEME DINWOODIE** (GDinwoodie@kentlaw.edu) is a Professor of Law at Chicago-Kent College of Law.

---


This column is derived from the article “A Reverse Notice and Takedown Regime To Enable Public Interest Uses of Technically Protected Copyrighted Works.” *Berkeley Technology Law Journal* 22 (2007), 981–1060.

---

© 2008 ACM 0001-0782/08/0200 \$5.00

---

DOI: 10.1145/1314215.1314220



**CONNECT WITH OUR  
COMMUNITY OF EXPERTS.**

**[www.reviews.com](http://www.reviews.com)**



Association for  
Computing Machinery

**Reviews.com**

They'll help you find the best new books  
and articles in computing.

Computing Reviews is a collaboration between the ACM and [Reviews.com](http://Reviews.com).

# ACM Annual Report for FY07

In recent months we have celebrated ACM's 60th Anniversary and the 50th Anniversary of its flagship publication, *Communications of the ACM*. Both events signify the enduring role ACM has played as the conduit for the world's educators, researchers, and professionals to share their common computing interests, inspire new innovation, and reveal their latest research. Indeed, ACM's steadfast commitment to advancing computing as a science and profession worldwide can be traced from the first days of the ENIAC, to improving today's quality of life, to encouraging and educating future generations to join a field of countless opportunities.

In the past year, major corporations have endorsed ACM's global influence in recognizing technical excellence by sponsoring or increasing the cash value of a number of the Association's prestigious awards. The financial value of the A.M. Turing Award, certainly the computing field's greatest and most visible honor, will increase from \$100,000 to \$250,000, thanks to the generosity of co-sponsors Intel and Google. IBM increased ACM's Software System Award to \$35,000. The Doctoral Dissertation Award, now sponsored by Google, is \$20,000 for first place, and \$10,000 for second place. And Google raised its sponsorship of the Grace Murray Hopper Award to \$35,000.

Moreover, ACM recently announced the creation of a new honor, the ACM-Infosys Foundation Award, recognizing

younger scientists whose innovations are having a dramatic impact on the computing field. The award will carry a prize of \$150,000, and was endowed by the Infosys Foundation.

ACM's global outreach has never been stronger, as witnessed by the growth in the number of technical conferences, professional chapters, and Special Interest Group initiatives taking root overseas this year. The Association's devotion to forging professional relationships worldwide was solidified by meetings to establish an ACM presence in the major technology hubs of China and India. ACM's China Task Force held its first two meetings in FY07. An agreement was reached with Tsinghua University to establish an ACM office in Beijing as a focal point for the Association's presence in China. ACM's India Task Force met in Bangalore recently to explore how best to serve this burgeoning audience.

In the coming year, ACM will continue to foster these critical global relationships, as well as build new bridges to connect the computing interests of its international membership, and create new services and products to meet their professional needs. We are working on expanding services for practitioners and a radically improved *Communications*.

As always, we are indebted to ACM's dedicated volunteers, members, staff, and industry partners for their ongoing support in turning promises into reality.

Stuart I. Feldman  
*ACM President*

**ACM**, the Association for Computing Machinery, is an international scientific and educational organization dedicated to advancing the art, science, engineering, and application of information technology.

#### PUBLICATIONS

ACM's Digital Library and *Guide to Computing Literature* continue to flourish in content, usage, and prestige, and remain in great demand worldwide. The number of full-text articles in the DL as of June 2007 was 209,000 and the Guide's bibliographic database increased to over one million entries.

Two new journals were launched in FY07: Transactions on Knowledge Discovery in Data (TKDD) and Transactions on the Web (TWeb). In addition, the ACM Publications Board approved two new publications: *ACM Transactions on Reconfigurable Systems* (TRES), and *ACM Transactions on Computation Theory* (ToCT).

The Publications Board also assembled a group of digital library experts to examine ways to enrich the functionality and services available from ACM's DL product. The discussions, which focused on areas such as search and retrieval technology, open-access, and user-interface improvements, culminated in a list of recommendations for ACM to consider.

Efforts to revitalize the editorial model of *Communications of the ACM* remained in full swing, with focus groups and volunteers coming together throughout the year to exchange ideas for enhancing ACM's premier publication. Details of these discussions and the subsequent plan were published in the January 2008 issue of *Communications*, page 44.

ACM *Queue's* "Queuecasts" had a very successful inaugural year. These podcasts, available

via the magazine's Web site ([www.acmqueue.org/](http://www.acmqueue.org/)), feature industry experts in 15-minute conversations on timely topics.

#### EDUCATION

A new Education Policy Committee (EPC) was formed to influence education policy on behalf of the computing community. The primary focus of the committee is education and advocacy in K–12 science, technology, engineering, and mathematics (STEM) education. The EPC's initial efforts will address the quality and availability of computer science education at the high school level.

The Computer Science Teachers Association (CSTA) continues to support and promote opportunities for K–12 teachers and students to better understand the computing disciplines. CSTA issued a second edition of *A Model Curriculum for K–12 Computer Science*, designed to set the context for computer science within current K–12 education and provide a framework for preparing young students for personal and professional opportunities in the 21st century.

A new brochure to counter the many misperceptions about computing circulating among high school students, their parents, and teachers was prepared jointly by ACM, IEEE Computer Society, and the Association for Information Systems. "Computing Degrees and Careers," provides more accurate information about computing and the career opportunities it offers.

### **PROFESSIONAL DEVELOPMENT**

The diverse offerings available from the Professional Development Centre were enhanced dramatically last year with the expansion of ACM's Online Books and Courses programs. Members now have access to over 1,100 books from Safari and Books24X7—a valuable resource providing reference works on a broad range of subjects in technology, project management, and business.

Web-based professional development courses available for ACM members grew to over 2,200 in FY07. These courses, from leading provider Skill Soft, cover a wide selection of topics from information technology and business opportunities.

ACM's Career and Job Center, powered by JobTarget, offers members the opportunity to view and apply for a variety of highly targeted technology jobs, including hundreds of corporate job postings often not seen on commercial sites. Members can post resumes, create and receive job alerts, and receive live career advice.

The Professions Board is hoping to launch its Best Practices community Web site in the near future. Steady progress was made this year in creating a site to provide answers to practical questions about technologies, approaches, tools, and processes by creating a community where IT professionals can discuss qualified solutions to problems.

### **PUBLIC POLICY**

ACM's U.S. Public Policy Committee

(USACM) continued its efforts to advance the computing discipline by finding a balance between innovation, intellectual property protection, and privacy issues; and for secure and reliable computing systems. Several committee members testified before Congress on the need to reform e-voting policy, the use of large-scale IT systems for identity and employment verification, and the privacy risks of using Social Security numbers in databases. Committee members and Policy Office staff briefed Congressional members on the threat of botnets, the technical issues in a proposed electronic employment verification system, and e-voting. In addition, USACM submitted detailed comments on the security risks inherent in the REAL ID program for new driver's licenses introduced by the Department of Homeland Security. For more information about USACM's activities, visit [www.acm.org/usacm/annualreports/index.html](http://www.acm.org/usacm/annualreports/index.html).

The Committee on Computers and Public Policy (CCPP) serves ACM with respect to a variety of internationally relevant issues pertaining to computing. The committee works to create and maintain global online forums to share and discuss public policy concerns.

### **STUDENTS**

The 31st Annual ACM International Collegiate Programming Contest World Finals was sponsored by IBM and hosted by ACM's Japan Chapter and IBM Tokyo Research Lab.

The preliminaries drew 6,099 teams from 1,756 universities representing 82 countries. Some 88 teams made it to the finals and the 2006 winners were from Warsaw University.

The ACM Student Research Competition is an internationally recognized venue enabling undergraduates and graduate students to share research results and exchange ideas with other students, judges, and conference attendees and understand the practical applications of their research. Rounds of competition are held at a variety of ACM conferences throughout the year.

ACM's Committee on Women in Computing (ACM-W) will provide financial support for undergraduate or graduate women students in CS programs who are interested in attending research conferences. Initially, up to 12 scholarships of about \$500 each will be awarded annually.

SIGARCH made a one-time grant of \$100K toward the endowment of a fellowship for Ph.D. students in high-performance computing.

#### **CONFERENCES**

The Federated Computing Research Conference, held every four years as an umbrella meeting ground for a wide spectrum of affiliated research conferences and workshops, drew a record number of attendees. In fact, most of the conferences within FCRC '07 exceeded attendance expectations.

SIGGRAPH 2006, held in Boston, also

exceeded attendance expectations, registering almost 20,000 artists, researcher scientists, developers, filmmakers, and academics from 80 countries. More than 230 companies exhibited, including 76 first-time exhibitors.

ASSETS 2006 featured novel computer applications for attendees with disabilities. Assistive technology research from the University of Washington was available for people with hand and hearing limitations, including the vocal joystick and a sign language cell phone that employs video compression.

#### **LOCAL ACTIVITIES**

The Membership Services Board chartered 41 new chapters in FY07, six of which were international professional chapters. Of the 35 new student chapters, 11 were internationally based.

#### **INTERNATIONAL**

Great inroads were made in ACM's initiatives to establish roots in the major technology centers of China and India. A China Task Force and India Task Force were established this year and initial meetings were held to build professional relationships and secure the best methods for sharing ACM's publications, services, and resources with these influential countries.

CSTA has been working on outreach programs with Mexico and several European associations. Moreover, CSTA hosted a workshop in Great Britain that examined issues



related to K–12 education, and plans for a joint symposium in Israel are under way for later this year.

The Education Board has been working to broaden European participation in computing education activities. The Board worked in an international advisory capacity with the European Union-funded Euro-Inf project, which was established to devise criteria for the accreditation of degrees in informatics across Europe at the undergraduate and masters level.

SIGGRAPH approved the concept of a new conference in Asia. The first SIGGRAPH Asia will be held at the end of the year in Singapore.

SIGMOD is making a determined effort to expand its international initiatives. Their initial steps forged close relationships with societies throughout Europe, China, and Japan.

#### **ELECTRONIC COMMUNITY**

The ACM History Committee revamped its Web site ([www.acm.org/history.acm.org](http://www.acm.org/history.acm.org)) and has made the first set of oral histories of former ACM Presidents and Turing Award recipients available in the Digital Library. Each interview is conducted by a professional historian and provides a sense of the ACM's role in computing history, as the interviews parallel the development of the field.

The Education Board and CSTA joined forces to create a dynamic Web site offering information on educational and career options in the computing field. The online material is

designed for high school students and educators and is available in English and Spanish.

SIGWEB has been building new member-only services into its Web site and expanding the information found within. New features of its site include a collection of relevant Ph.D. dissertations and a series of interviews with active researchers from the SIGWEB community.

#### **RECOGNITION**

The ACM Fellows Program, established in 1993 to honor outstanding ACM members for their achievements in computer science and IT, inducted 41 new fellows in FY07, bringing the total count to 594.

The first recipients of ACM's distinguished member grades were announced last spring. The Association recognized 77 Senior Members who have at least 10 years of professional experience and who have demonstrated performance and accomplishment that set them apart from their peers. ACM also recognized 17 Distinguished Engineers, 24 Distinguished Scientists, and eight Distinguished Members—all having the distinction of at least 15 years of professional experience and who have demonstrated significant accomplishments or made a significant impact on the computing field.

The first recipients of the Frank Anger Memorial SIGBED/SIGSOFT Student Award were announced in FY07. The award promotes interdisciplinary research between embedded systems and software engineering.

## BALANCE SHEET

JUNE 30, 2007 (IN THOUSANDS)

---

### ASSETS:

CASH AND CASH EQUIVALENTS	\$21,328
INVESTMENTS	47,406
ACCOUNTS RECEIVABLE AND OTHER CURRENT ASSETS	4,008
CONTRIBUTION RECEIVABLE	1,250
DEFERRED CONFERENCE EXPENSES AND OTHER ASSETS	5,192
FIXED ASSETS, NET OF ACCUMULATED DEPRECIATION AND AMORTIZATION	2,022
<b>TOTAL ASSETS</b>	<b>\$81,206</b>

### LIABILITIES AND NET ASSETS:

#### LIABILITIES:

ACCOUNTS PAYABLE, ACCRUED EXPENSES, AND OTHER LIABILITIES	\$ 7,876
UNEARNED CONFERENCE, MEMBERSHIP, AND SUBSCRIPTION REVENUE	19,076
<b>TOTAL LIABILITIES</b>	<b>26,952</b>

#### NET ASSETS:

UNRESTRICTED	48,796
TEMPORARILY RESTRICTED	5,458
<b>TOTAL NET ASSETS</b>	<b>54,254</b>
<b>TOTAL LIABILITIES AND NET ASSETS</b>	<b>\$81,206</b>

#### OPTIONAL CONTRIBUTIONS FUND—PROGRAM EXPENSES (\$000)

EDUCATION BOARD ACCREDITATION	\$50
USACM COMMITTEE	20
<b>TOTAL EXPENSES</b>	<b>\$70</b>

## STATEMENT OF ACTIVITIES

YEAR ENDED JUNE 30, 2007 (IN THOUSANDS)

	UNRESTRICTED NET ASSETS	TEMPORARILY RESTRICTED NET ASSETS	TOTAL
<b>REVENUE:</b>			
MEMBERSHIP DUES	\$8,832		\$8,832
PUBLICATIONS	14,131		14,131
CONFERENCES AND OTHER MEETINGS	23,530		23,530
INTEREST AND DIVIDENDS	2,121		2,121
NET APPRECIATION OF INVESTMENTS	3,694		3,694
CONTRIBUTIONS AND GRANTS	3,108	5,089	8,197
OTHER REVENUE	218		218
NET ASSETS RELEASED FROM RESTRICTIONS	665	(665)	----
<b>TOTAL REVENUE</b>	<b>\$56,299</b>	<b>4,424</b>	<b>\$60,723</b>
<b>EXPENSES:</b>			
<b>PROGRAM:</b>			
MEMBERSHIP PROCESSING AND SERVICES	\$889		\$889
PUBLICATIONS	10,683		10,683
CONFERENCES AND OTHER MEETINGS	21,012		21,012
PROGRAM SUPPORT AND OTHER	5,992		5,992
<b>TOTAL PROGRAM EXPENSES</b>	<b>\$38,586</b>		<b>\$38,586</b>
<b>SUPPORTING SERVICES:</b>			
GENERAL ADMINISTRATION	\$8,708		8,708
MARKETING	1,394		1,394
<b>TOTAL EXPENSES</b>	<b>\$48,688</b>		<b>48,688</b>
INCREASE (DECREASE) IN NET ASSETS	7,611	4,424	12,035
NET ASSETS AT THE BEGINNING OF THE YEAR	41,185	1,034	42,219
NET ASSETS AT END OF YEAR	\$48,796*	\$5,458	\$54,254*

\*INCLUDES SIG FUND BALANCE OF \$23,939

## ACM COUNCIL

PRESIDENT	<i>Stuart I. Feldman</i>
VICE PRESIDENT	<i>Wendy Hall</i>
SECRETARY/TREASURER	<i>Alain Chesnais</i>
PAST PRESIDENT	<i>David A. Patterson</i>
SIG GOVERNING BOARD CHAIR	<i>Joseph A. Konstan</i>
PUBLICATIONS BOARD CO-CHAIRS	<i>Ronald Boisvert, Mary Jane Irwin</i>
MEMBERS-AT-LARGE	<i>Michel Beaudouin-Lafon, Bruce Maggs Barbara Ryder, Kevin Scott Jeanette Wing, David S. Wise</i>
SGB COUNCIL REPRESENTATIVES	<i>Norman Jouppi, Robert A. Walker Alexander Wolf</i>
ACM HEADQUARTERS	
Executive Director/CEO	<i>John R. White</i>
Deputy Executive Director/COO	<i>Patricia M. Ryan</i>

## 2006 ACM AWARD RECIPIENTS

A.M. TURING AWARD	<i>Frances E. Allen</i>
ACM/AAAI ALLEN NEWELL AWARD	<i>Karen Spärck Jones</i>
GRACE MURRAY HOPPER AWARD	<i>Daniel Klein</i>
ACM-IEEE CS ECKERT-MAUCHLY AWARD	<i>Mateo Valero</i>
OUTSTANDING CONTRIBUTION TO ACM AWARD	<i>David S. Wise</i>
DISTINGUISHED SERVICE AWARD	<i>Susan L. Graham</i>
PARIS KANELAKIS THEORY AND PRACTICE AWARD	<i>Robert Brayton</i>
ACM PRESIDENT'S AWARD	<i>Eugene H. Spafford</i>
SOFTWARE SYSTEM AWARD	<i>Bertrand Meyer</i>
SIAM/ACM AWARD IN COMPUTATIONAL SCIENCE AND ENGINEERING	<i>Chi-Wang Shu</i>
DOCTORAL DISSERTATION AWARD	<i>Yi-Ren Ng</i>
HONORABLE MENTION	<i>Aseem Agarwala</i>

**CALL FOR PARTICIPATION**

# **CTS 2008**

**Irvine, California, USA**



## **The 2008 International Symposium on Collaborative Technologies and Systems**

**May 19 – 23, 2008**

**The Hyatt Regency Irvine Hotel  
Irvine, California, USA**

### **Important Dates:**

Paper Submission Deadline -----	<b>January 7, 2008</b>
Workshop/Special Session Proposal Deadline -----	<b>December 14, 2007</b>
Tutorial/Demo/Panel Proposal Deadline -----	<b>December 21, 2007</b>
Notification of Acceptance -----	<b>February 8, 2008</b>
Final Papers Due -----	<b>March 3, 2008</b>

### **Conference Co-Chairs:**

**Bill McQuay**, Air Force Research Laboratory, Wright Patterson AFB, USA  
**Waleed W. Smari**, University of Dayton, USA

**For more information, visit the CTS 2008 web site at:**

**<http://cisedu.us/cis/cts/08/>**



**In cooperation with the ACM, IEEE, IFIP**

By Jeffrey Y. Kim, Jonathan P. Allen,  
and Elan Lee

# ALTERNATE REALITY GAMING

*Millions update the state of the game on the way to a common conclusion, in one case to help the Operator regain control of a spaceship and bring her crew back to the future.*

Alternate reality games represent a new genre of digital gaming designed to blur the distinction between a player's experience in the digital world inside the game and the real world outside the game. Combining online information and real-world events, ARGs bring gamers together to collectively solve puzzles and advance a game's storyline. Part of what characterizes an ARG is that the game universe is not explicitly limited to a particular piece of software or set of digital content. A typical ARG would not even acknowledge or promote the fact that it is a game, yet every Web site or discussion group may contain and reveal a potential clue.

---



One of the first successful ARGs was “The Beast,” a murder-mystery problem-solving game launched in 2001 to promote the Steven Spielberg movie “Artificial Intelligence.” Its first clues (known as the “rabbit hole”) were planted in the images in movie advertising posters. People who followed the clues found Web sites suggesting that a character named Evan Chan had been murdered in the A.I. movie universe. In order to identify the murderer, the gamers formed their own discussion group called “Cloud-makers” ([www.cloudmakers.org/](http://www.cloudmakers.org/)). The game attracted significant mainstream media and public interest. At its peak in July 2001, it had an estimated three million unique visitors, putting ARGs on the entertainment map.



Figure 1. Margaret's Honey bottle.

**T**wo features of a successful ARG are a compelling storyline and collaborative game play. In it, a plotline is narrated and delivered through multiple communication channels, including Web pages, email messages, phone calls, and print-based mailings. Gamers use them to track the story's progress. In two of the first ARGs to achieve a critical mass of more than a million participating gamers—"The Beast" and "ILoveBees" (2004)—compelling plot development made it possible for the designers to attract, retain, and increase the number of players. The collaborative nature of game play is another important ARG feature. Game players construct their own means of interacting through Web-based discussion boards, email messages, and real-world gatherings. Through this interaction, players communicate with one another, share their knowledge, offer interpretations of the storyline, and gather the information necessary to progress toward the game's conclusion.

ARGs are quite different from another more widely played digital game genre: massively multi-

player online games. MMOG gamers collectively play in a shared digital world that provides a persistent game universe in which a large number of gamers (often more than 20,000 at a time) play against computer-generated characters or other MMOG players. Whether individually or as a team, or "clan," gamers aim for a relatively well-defined set of goals (such as killing enemies, seizing a castle, or strengthening a gamer's in-game character). In the MMOG world, gamers generally play their chosen roles, engage in combat, or create collective real-time strategies in war or commerce.

ARGs are not designed to create a compelling experience through 3D graphics or simulated battles. Rather, they provide shared scenarios through which gamers interact and collaborate to construct an eventual ending to the story. ARG gamers' enjoyment depends on the kind of shared experience they have with one another. Without it, they will not stay interested in collective puzzle solving and information gathering. The gamers' sustained, active, and voluntary participation is the most important condition of the ARG experience. Without it, the game stops evolving and ultimately ceases to exist.

An ARG sustains player interest by generating new content based on the constantly updated state of the game. If an existing puzzle is solved, the game's storyline is designed to move forward to sustain gamer interest. If players find a clue and share it with the rest of the gaming community, the storyline must be updated in real time. A typical story update cycle takes a week, enough time for gamers to reflect on the story before new action takes place or a new puzzle is released [1]. However, if gamers move more quickly than the designers ("Puppetmasters") anticipate, the designers must react immediately to maintain the game's appeal. Timing is key. A stale game loses both the smaller group of active players and the larger body of participants watching the action unfold. ARGs are both collaborative game and spectator sport.

The time-dependent and interactive nature of



ARGs makes them like theatrical productions, spontaneous and responsive to the audience. The challenging aspects of creating an engaging ARG include compelling plotlines, regular content delivery, collective puzzles of the appropriate difficulty level, player feedback monitoring, and self-organizing player groups—all while maintaining strict secrecy and leaving no opportunity to test before going live.

### LAUNCHING AN ARG (“ILOVEBEES”)

The ARG “ILoveBees” (ILB) game was launched in July 2004 to promote an upcoming release of the Xbox game Halo2. From the beginning, ILB offered a stream of mysterious clues, initially through two different channels: In one, FedEx packages were sent to approximately 20 people active in the digital gaming industry. Inside was a bottle of honey from Margaret’s Honey in San Francisco (see Figure 1). Suspended in the honey were nine letters spelling I-L-O-V-E-B-E-E-S. The editor at an ARG player Web site (argn.com) received the same package and immediately posted his thoughts on the site. In the other channel that same weekend, the *ilovebees.com* address was displayed for a fraction of second as the Halo2 trailer was being shown at movie theaters. Visitors to *www.ilovebees.com* found a Web site that appeared to have been infested by a strange artificial intelligence program most notable for Web images.

ILB’s primary aim was to attract gamer and media interest in the Halo2 release. The ILB design team created a Web site (*www.ilovebees.com*) that seemed to be infected by a mysterious computer virus from the future (as in Halo2). For the next few weeks, ARG gamers and Halo fans did not know what to do with the game. Only after investigating clues left on the Web site did they figure out their main objective. They would have to help an AI program (“the Operator”) that crash-landed on Earth find its way back to the



Figure 2. ILOVEBEES player.

the invading Covenant army.

### ARG DESIGN TEAM

Before taking on development of ILB, its design team already had a successful experience with the ARG “The Beast.” This helped the lead game architects assemble a design team of three storywriters, a community lead, a number of audio-file producers, a technical support group, and several voice actors. The total ILB design team included no more than 30 members, significantly fewer than a typical MMOG product-development team. The ILB design team had three main responsibilities:

**Storyteller.** The lead writer, science-fiction author Sean Stewart, led the story writers. Before the game’s launch, the ILB team interacted with the Halo2 team (Bungee and Microsoft’s Xbox Division) to keep the game story consistent with the Halo2 universe.

**Community lead.** While most sound and image production was outsourced, the design team maintained tight control over the monitoring of players. Because timing and the mystery of the story are key elements of any ARG, the design team did not have the luxury of beta testing or focus-group testing. Rather, it relied on round-the-clock monitoring of the community. The community lead (Jane McGonigal, then a doctoral student at the University of California, Berkeley) monitored player activities and communication throughout the entire game. The community lead checked every known source of information written about the game, including discussion groups (such as *bees.netninja.com* and the “haunted apiary” at *unfiction.com*), as well as mainstream media coverage and blogs. The commu-

nity lead filed a daily report to the design team every morning the game was being played. With the report, the design team could make instant story updates and adjustment as needed. Gamers were locked in a contest to solve new puzzles. And, with so many gamers collaborating, the pace of the game often moved more quickly than the design team had anticipated.

*Technology support and sound effects.* The technology group managed Web page design, site security, and maintenance. The audio-production group created audio files that, when assembled in the correct order, provided stories from the Halo universe. More audio files and special sound effects were added as the game progressed.

### ARG PLAYERS

From the beginning, ILB attracted gamers from other ARGs, as well as from various MMOGs (Halo fans). The design team saw exponential growth as the campaign went on. The initial number of gamers was about 100, then increased to 500, 10,000, 250,000, and eventually three million (over only 12 weeks). During the first week of play, gamers posted 50 new comments every 30 seconds [4]. In addition to more than a million message-board posts, the gamers logged an average of 33,000 lines of chat per day [3].

**I**LB players fell into two categories: active players and bystanders who watched the active players solve the puzzles and advance the game; the 100,000 active players were only a fraction of the three million unique visitors to the ILB site. The ILB players were not limited to the typical game-playing demographic: teenage boys. Though the ILB design team did not survey the player community, it felt that the players' demographic characteristics were similar to their previous experience. For "The Beast," a fan Web site survey with several thousand responses showed 50% of responding gamers were female. This contrasted sharply with the typical gender distribution of digital games: 80% male, 20% female. However, most ILB gamers belonged to the traditionally most active player age group: 16 to 25.

The ILB gamers played together as a group (see Figure 2), often gathering (physically) for game activities (such as receiving phone messages). They cre-



Figure 3.  
Payphone  
crowd.

ated discussion groups to share game information, including dedicated ARG sites ([www.unfiction.com](http://www.unfiction.com)) and general Internet collaboration sites ([wikipedia.org](http://wikipedia.org)). One of the first items posted by experienced ARG gamers was a message asking how they would play the game together. Within two weeks, they began to get a sense of the ILB story. Volunteers posted summaries and possible interpretations of the storyline on various Web sites. Throughout the game, these sites offered a well-organized and coherent storyline as ILB progressed. Contributors also connected plot developments with detailed information from the Halo game universe. Information at these sites was so well organized that even the ILB design team found the summaries helpful. The ILB gamers collectively found clues, solved puzzles, and shared their thoughts on the deeper meanings in the story. Players assumed that the ILB design team knew exactly how the game would unfold and therefore would always be a step ahead of the players. When the game concluded in November 2004, ILB gamers were genuinely surprised to hear the design team say the gamers themselves had control over how the plot unfolded.

### STORY CONSTRUCTION AND DELIVERY

Gamers had to assemble the story of the Operator (the crash-landed AI program) from various story fragments delivered by the ARG designers. To provide story narratives the ILB design team relied on three primary channels:

- Hidden html code, email exchanges, sound files, and images purportedly created by the Operator;
- Voice clips sent to payphones; and
- A blog maintained by an imaginary character in the game.

www.ilovebees.com/404

The ILB design team used them to blur the boundaries between in-game and out-of-game reality and experience; players exchanged email messages with the in-game AI programs and communicated verbally with the Operator and with each other.

The ILB design team tried two new approaches to ARG delivery: more extensive game community monitoring and nondigital, as well as digital, media to deliver the game experience. Throughout the game (three months), the community lead checked gamer Web sites, analyzed wiki updates, and read Internet relay chat activities on a daily basis. This information helped the designers adjust the game to reflect recent player actions. For example, the designers assumed that gamers would sympathize with the Sleeping Princess, another in-game AI program on the *ilovebees.com* server. In an unexpected turn of events, some gamers felt otherwise and informed the “bad” AI program about the Princess’s hiding place. The designers had no choice but to let the Princess be captured by the enemy in the next round of the ILB’s story update.

**T**he ILB design team applied various media (such as email, blogs, and phone calls) to quickly deliver pieces of narrative to as many and as diverse a group of gamers as possible. The most notable technical innovation in ILB was the use of public payphones as a delivery mechanism (throughout the U.S. and the U.K.). The GPS coordinates of selected payphones and narrative transmission times were announced on the ILB Web site. Once the narratives were confirmed to have been heard via payphone, the design team made the audio clips available on the ILB site. The designers tried to ensure that game play was not too difficult. In order to keep entry barriers as low as possible, the design team hired a group of payphone-location-seekers “axon hunters,” who searched for suitable payphones and documented the relevant GPS coordinates. This selection process ensured that the only entry barrier involved the geographical location of the gamers in the U.S. and the U.K.

The gamers relied on email to communicate with the in-game characters. Thousands of gamers traveled to retrieve the messages at payphones, sharing their experience on specialized discussion boards. The gamers also collectively learned the syntax and meaning of the special language embedded in

*ilovebees.com* Web pages. The language itself was not overly complex but was arcane enough to give gamers a sense of camaraderie. These collaborative experiences contributed to a shared sense of playing together. Eventually, in Aug. 2004, the number of unique visitors to the ILB site reached three million. All the voice clips (412 in total, plus 70 other clips embedded in Web images) were heard at payphones and compiled by players into a coherent sequence that revealed the secret reason for the spacecraft’s crash landing. The game ended when the gamers helped the Operator regain control of the spaceship and bring her crew back to the future.

In November 2004, live events in major U.S. cities, including Boston, Chicago, and San Francisco, allowed ILB gamers to meet one another and speak to the game’s designers. In the end, the ILB designers reached and surpassed their goals in terms of numbers of unique visitors for generating and sustaining player interest in the ARG (see Figure 3).


#### **FUTURE ARG**

As ARGing have become increasingly popular, new variations have pushed the boundaries of the genre. New ARGs are experimenting with stories based on real-world and political topics (such as the World Without Oil simulation [5]). Instead of being tied to an existing story or product, ARGs (such as *Perplex City* [2]) are creating their own standalone storylines and game products. Other ARGs tied to specific products or storylines are being created unofficially by fans. In one celebrated case, the producers of the online video series *Lonelygirl15* were so impressed by an unofficial “fan fiction” ARG that they made it the official ARG for the series ([www.lonelygirl15.com](http://www.lonelygirl15.com)). As long as players show interest, the ARG genre will continue to diversify.

Perhaps the most significant development for ARGs is that they are going mainstream. As ARG-like games are used in more movies, TV shows, and consumer-goods marketing, mainstream audiences may demand less secretive and complex gameplay. This will also create more experimentation with different kinds of ARGs and measures of ARG performance.

Bringing ARGs into the mainstream increases the importance of ease of use and low barriers to entry. Because the gameplay itself can be complex (spanning multiple Web sites, phone lines, email systems, and physical locations), ARGs must learn to accommodate both hardcore and casual players. While a hardcore player might be willing to leave home in the middle of a tornado to answer a ringing payphone, the casual player must be able to find a way into the game simply by visiting a Web site or an online forum. The

# ACM TRANSACTIONS ON INTERNET TECHNOLOGY



**Transactions on Internet Technology (TOIT).** This quarterly publication encompasses many disciplines in computing—including computer software engineering, middleware, database management, security, knowledge discovery and data mining, networking and distributed systems, communications, and performance and scalability—all under one roof. TOIT brings a sharper focus on the results and roles of the individual disciplines and the relationship among them. Extensive multi-disciplinary coverage is placed on the new application technologies, social issues, and public policies shaping Internet development. **Subscribe Today!**

## PRODUCT INFORMATION

ISSN: 1533-5399      Price: \$39 Professional Member  
Order Code: (140)      \$34 Student Member  
                                 \$160 Non-Member  
                                 \$14 Air Service (for residents  
                                 outside North America only)

## TO PLACE AN ORDER

### Contact ACM Member Services

Phone: 1.800.342.6626      Email: acmhelp@acm.org  
(U.S. and Canada)      Mail: ACM Member Services  
+1.212.626.0500      PO Box 11414  
(Global)      New York, NY  
Fax: +1.212.944.1318      10286-1414 USA  
(Hours: 8:30am—4:30pm, Eastern Time)

[www.acm.org/pubs/periodicals/toit](http://www.acm.org/pubs/periodicals/toit)



Association for Computing Machinery  
Advancing Computing as a Science & Profession  
[www.acm.org](http://www.acm.org)

AD26

opportunity must always be available for players to be more involved, but as ARGs enter the mainstream, they must offer more casual, compact, and immediately entertaining experiences to attract the largest audience possible.

What will never change is the importance of a compelling story that attracts and keeps people in an ARG universe over the course of months, even years. For both experienced and casual players, an ARG's success depends on the simplicity of its storyline. An ARG must be easy to understand and simple enough to describe to friends, even those who know nothing about technology. The ability to spread a simple intriguing message virally among players will continue to be a key to any ARG's appeal. Also, the community leads' ability to know players and sense their daily reaction to the story will remain central to a successful ARG. Like any form of mass collaboration, the future of ARGs is in the hands of the players and the people watching them play in both the real and the virtual worlds. **C**

## REFERENCES

1. Kim, J. Interview with Sean Stewart (May 2004).
2. Kurs, S. Final clue in the world's biggest treasure hunt. *The Sunday Times* (Jan. 29, 2007); [technology.timesonline.co.uk/tol/news/tech\\_and\\_web/personal\\_tech/article1288060.ece](http://technology.timesonline.co.uk/tol/news/tech_and_web/personal_tech/article1288060.ece).
3. McGonigal, J. Interaction statistics. *42 Entertainment Community Internal Report* (Aug. 27, 2004).
4. McGonigal, J. Alternate reality gaming meets Halo. *42 Entertainment Community Internal Report* (July 23, 2004).
5. Olsen, S. Provocative politics in virtual games. CNet News (Mar. 28, 2007); [www.news.com/Provocative-politics-in-virtual-games/2100-1043\\_36171089.html?tag=item](http://www.news.com/Provocative-politics-in-virtual-games/2100-1043_36171089.html?tag=item).

**JEFFREY Y. KIM** (jykim@u.washington.edu) is deputy director of the Institute for National Security Education and Research in the Information School at the University of Washington, Seattle, and an adjunct assistant professor in the Graduate School of Culture Technology at the Korea Advanced Institute of Science and Technology.

**JONATHAN P. ALLEN** (jpallen@usfca.edu) is an associate professor of information systems in and chair of the Technology, Innovation, and Entrepreneurship Department of the School of Business and Management at the University of San Francisco.

**ELAN LEE** (elan@fourthwallstudios.com) is founder and chief designer of Fourth Wall Studios and was creative director and lead designer of "I Love Bees," "The Beast," "Last Call Poker," and other 42 Entertainment projects.

The authors are grateful to Sean Stewart of Fourth Wall Studios and Bob Fagan of 42 Entertainment for providing valuable insider tales concerning the design of ILB.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2008 ACM 0001-0782/08/200 \$5.00

DOI: 10.1145/1314215.1314222

---

*By* Eric Shih, Kenneth L. Kraemer,  
*and* Jason Dedrick

# IT DIFFUSION IN DEVELOPING COUNTRIES

*Policymakers need to recognize that developing economies have different drivers for IT investment than their wealthier brethren.*

**T**here is widespread belief among international agencies and development specialists in the potential value of information technology (IT) to support economic and human development [11, 12]. Some question whether IT alone can have a major impact on the standard of living in developing countries, but most see it offering access to vital information and services such as weather forecasting, commodity prices, health care, and education. However, a significant digital divide exists between richer and poorer countries in the use of IT and the availability of complementary assets such as telecommunications networks and skilled IT professionals. This gap has led to a public debate about what can be done to promote greater IT use so that developing countries can achieve the types of benefits already being enjoyed in the industrialized world.

The problem currently is diagnosed by some as resulting from a lack of affordability of computer hardware, with various low-cost computers such as a \$100 laptop being offered as solutions. Others view the problem as part of a broader set of issues that include poverty, lack of infrastructure and inadequate education. To inform these policy discussions, it is important to understand the factors that influence IT use at the country level and whether there are differences in these factors between developed and developing countries. Policy efforts based on incorrect assumptions are likely to have little impact on IT diffusion or economic development.

In order to offer empirically based insights into the drivers of IT use, we analyzed data from 44 countries over a 15-year period and found markedly different results for developing and developed countries. These results have implications for government policymakers and others interested in promoting IT investment for economic development.

Many in the development community believe that IT can serve as a catalyst to help poorer countries accelerate development [11, 12]. As later adopters, they have access to lower cost, easier-to-use technology, and can learn from the experience of developed countries in using IT. However, studies that compare developed and developing countries find that IT investment has a positive and significant relationship with productivity growth at the macroeconomic level in developed countries, but not in developing countries [6].

One likely reason is the low level of IT investment in poorer countries. Even in the U.S., the impacts of IT on productivity and growth only became widely recognized by economists in the late 1990s, after decades of cumulative investment [7]. Bell and Pavitt [2] argue that growth in developing countries comes from technological accumulation, which is "... an evolutionary process of continuous innovation and imitation." Given their historically low levels of IT spending, it is likely that most developing countries have not reached a level of accumulated investment needed to achieve measurable productivity gains. Hence the question of what factors influence country-level IT investment is very important, as IT

investments over time will influence when countries might achieve significant economic impacts.

Only a few empirical studies have looked at the factors influencing the level of IT investment. One study of 11 Asia-Pacific countries found that IT investment was associated with diffusion of telecommunications infrastructure, education levels, technical skills, and the percent of the economy in services industries [8]. Another study of 89 countries found that computer hardware imports, an indicator of IT investment, were associated with educational attainment, openness to imports and property rights protection [3].

Another study, based on a flexible accelerator investment model, finds differences between developed and developing countries in terms of factors influencing IT investment [10]. Here, we further examine the issue of developed and developing country factors using a seemingly unrelated regression (SUR) model that focuses on factors that can be addressed by policy choices of both governments and international organizations.

#### FACTORS LINKED TO IT INVESTMENT

Theoretical and empirical studies have noted that the process of technology diffusion occurs unevenly across national boundaries because of differences in national

environments and links to the global economy. Although various studies have shown that IT investment is correlated with the level of national wealth, other factors have been shown to be significant as well: resources for technology investments, structure of the economy, complementary assets, and openness to external influences (see Figure 1).

*Resources for technology investments.* Investing in new technologies requires the availability of capital, either from external sources such as foreign direct investment and foreign aid or from internal sources such as equity markets and domestic loans and credits. For both developing and developed countries, access to loans and credits is likely to be important. For developed countries, we would expect that access to capital via the equity markets would also be important, as those countries have more well-developed stock markets that reward companies for making productivity-enhancing investments in IT. For most

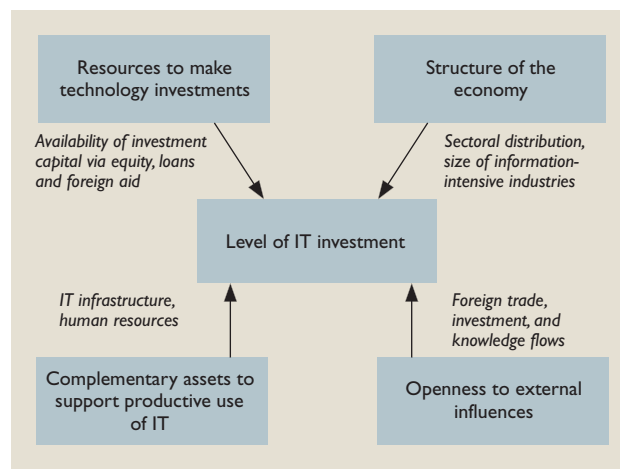


Figure 1. Conceptual framework.

developing countries, we would not expect equity markets to be mature enough to be a significant source of capital. However, we would expect foreign aid to play a significant role, as it can serve as a substitute for scarce domestic capital and also stimulate private investment.

*Structure of the economy.* When we look across countries, we find a wide disparity in the extent to which national economies have evolved from agriculture to manufacturing and finally to information-intensive services. Unlike agriculture, manufacturing, or even retail services, which involve production or movement of physical goods, sectors such as financial services primarily involve manipulation and transmission of information. In financial services, the use of IT is much more pervasive, as IT can greatly increase the efficiency and reduce the cost of manipulating and transmitting information. We thus would expect countries with larger financial services sectors to have higher rates of investment in IT. Earlier research has found a significant positive association between the size of a country's services sector and IT investment [3, 8, 9].

We further expect the positive impact of the financial services sector on IT investment to be more pronounced in developed countries than developing countries. For developed countries, there is likely to be some point at which the increasing scale of the finance sector requires accelerating investments in IT to handle the volume and complexity of transactions through increased automation.

*Complementary assets.* IT use requires the presence of complementary assets such as a telecommunications infrastructure and skilled human resources. The importance of a telecommunications infrastructure is supported by empirical studies showing a positive association between teledensity indicators and IT investment [3, 8, 9].

Also vital is the presence of human resources with appropriate skills and access to sources of information on how to use the technology. Educated workers more readily adjust to the implementation of new technologies and an educated workforce reduces opposition to social changes associated with adoption of new technologies [9]. Also, human resources such as IT professionals, engineers, and management specialists

are likely to be especially important, not only for their technical knowledge, but for their access to information through personal networks. Strong empirical support has been found between education levels and IT use at the country level [3, 8, 9].

However, it is possible that the impact of complementary assets is less significant in developed countries, which have widespread diffusion of telecommunications and high education levels. Once

Type of variable	Factor	Measure	Definition
Independent variables	Resources to make technology investments	Log (AIDGDP)	Foreign aid as a percent of GDP
		Log (CDTGDP)	Loans and credits made to the private sector as a percent of GDP
		Log (MARGDP)	Total market capitalization as percent of GDP
	Structure of the economy	Log (FIRE)	Finance, insurance and real estate as percent of total country employment
		Complementary assets	Log (TELDEN)
	Log (EDTERP)		Tertiary school enrollment as percent of relevant age group
	Openness to foreign influences	Log (TRDGDP)	Trade as percent of GDP
		Log (FDIGDP)	Inward foreign direct investment as percent of GDP
	Control variable	Country wealth	Log (GDP/CAP)
Dependent variable	IT investment (Diffusion)	Log (IT/GDP)	IT spending on hardware as percent of GDP

Figure 2. Measures and definitions for diffusion factors.

the level of complementary assets reaches a certain level, the marginal impact of an additional phone line, or of an extra percentage point in tertiary education, may be diminished. In contrast, we would expect that the impact of complementary assets would be greater in developing countries, which are still in the process of creating adequate levels of such assets [6].

*Openness to external influences.* Effective use of IT requires a broad range of knowledge, both technical and managerial, much of which can be found beyond the borders of any country. Foreign trade facilitates the diffusion of such knowledge across borders as it “provides channels of communication that stimulate cross-border learning of production methods, product design, organizational methods, and market conditions” [4]. Additionally, foreign direct investment (FDI) has a positive impact on technical progress in the host country [1].

We would expect that greater external openness should lead to more rapid diffusion of technologies into a country. Multinational corporations tend to bring with them business practices that rely more heavily on IT, and thus are more likely to invest in IT themselves and require that their suppliers make similar investments. They also bring knowledge of how to use IT productively. Openness to trade and FDI also may expose a national economy to greater interna-

tional competition, driving IT investment by local firms as a tool for survival. Finally, doing business internationally may force firms to adopt IT in order to meet the requirements of foreign suppliers or customers.

Thus, we would expect countries with higher levels of trade and FDI, relative to GDP, would invest more in IT. We expect that the impacts would be more significant for developing countries, which are likely to be farther behind the global frontier in adopting IT-enabled business practices and thus should benefit more from external sources of knowledge.

## METHODS, DATA, AND MODELS

We model IT spending against the foregoing factors to identify which are determinants of IT investment in the entire sample of countries. We then divide the sample into developed and developing countries, in order to test the hypothesis that these groups will have different characteristics in terms of IT investment drivers.

We used multiple measures to capture the factors hypothesized to influence the level of investment in IT as shown in Figure 2. And because IT investments are likely to correlate significantly with the wealth or level of economic development in a country, we control for wealth, and measure it in terms of gross domestic product (GDP) per capita.

*Data.* We use spending on computer hardware as a percent of GDP as a measure of IT investment. International Data Corporation (IDC) provided data on IT investments from 1985 to 1999 for 44 countries. The series captures the value of shipments, which is the revenue paid to vendors for hardware and systems. Countries included in our analysis and their respective beginning and ending year of available data are presented in Table 1. Data on the structure of national economies comes from the International Labour Organization's *2000 Yearbook of Labour Statistics* and World Bank's *2004 World Development Index*. We standardized monetary findings in current year U.S. dollars.

Developing Countries			Developing Countries		
	Years		Years		Years
1	Australia	1985-99	1	Argentina	1991-99
2	Austria	1985-99	2	Brazil	1985-99
3	Belgium	1988-99	3	Bulgaria	1991-99
4	Canada	1985-99	4	Chile	1985-99
5	Denmark	1985-99	5	China	1987-99
6	Finland	1985-99	6	Colombia	1985-99
7	France	1985-94	7	Czech Republic	1993-99
8	Germany	1991-99	8	Egypt	1989-99
9	Ireland	1985-99	9	Greece	1985-99
10	Italy	1985-99	10	Hungary	1992-99
11	Japan	1985-99	11	Indonesia	1985-99
12	Netherlands	1985-99	12	Israel	1985-99
13	New Zealand	1986-99	13	South Korea	1985-99
14	Norway	1985-99	14	Malaysia	1985-99
15	Singapore	1985-99	15	Mexico	1988-99
16	Spain	1985-99	16	Philippines	1985-99
17	Sweden	1985-99	17	Poland	1985-99
18	Switzerland	1991-99	18	Portugal	1985-99
19	United Kingdom	1985-99	19	Romania	1996-99
20	United States	1985-99	20	Russia	1993-99
			21	Slovak Republic	1994-99
			22	Slovenia	1993-99
			23	Turkey	1985-99
			24	Venezuela	1985-99

\* In current year US\$

Table 1. Data availability by country.

countries as all others (n=24). The developed countries have an average GDP per capita of \$18,945 over the 15 years in our sample, while developing countries averaged \$2,749.

We used SUR and estimated two sets of parameters for developed and developing countries.<sup>2</sup> We did this instead of estimating two different equations because we assume that there exists some factor, such as global economic conditions, that impacts both developed and developing countries and is not captured in the model.<sup>3</sup> The SUR model resulted in a higher adjusted R<sup>2</sup> than the initial regression model (0.866) (see Table 3). The SUR results contrast with the initial regression model and provide us with a clearer picture of differences in determinants across the two samples.

## DEVELOPING COUNTRIES HAVE DIFFERENT IT DRIVERS

The results provide empirical support for most of our predictions about the factors shaping IT investments across national economies. The factors explain 80% of the variation in IT investments. These results indicate the usefulness of our conceptual model (see Fig-

<sup>1</sup>For countries that receive no foreign aid, it would have been impossible to take natural log transformations. Therefore, a small constant of 0.0001 was added to level of foreign aid of all countries at all years to avoid censoring.

<sup>2</sup>The fixed effects pooled model is

$$\ln(Y_{i,t}) = \beta_1 \ln(X_{1,i,t}) + \beta_2 \ln(X_{2,i,t}) + \dots + \beta_7 \ln(X_{9,i,t}) + \sum_{i=1}^N \omega_i \alpha_i + \sum_{t=1}^{T-1} \lambda_t \gamma_t + \varepsilon_{it}$$

<sup>3</sup>As a caution against the possibility of reverse causation, we conducted a Granger causality test where reverse causation is absent when  $f(\ln x_{i,t} | \ln x_{i,t-1}, \ln X_c, \ln y_{t-1}) = f(\ln x_{i,t} | \ln x_{i,t-1}, \ln X_c)$  where  $\ln X_c$  is a set of  $\ln x$ 's other than  $\ln x_t$ . That is, all other independent variables are used as control variables in the test of causation. A simple linear model was used in estimation. Test results indicate that all  $F_{(1,538)}$ 's  $> 1.198$ ,  $p > 0.10$ , which suggests that  $\ln y_t$  does not cause  $\ln x_t$  and therefore reverse causation is not an issue.

*Model estimation.* Our model combined data in a cross-section and over time in one single model. To correct for skewed data, natural log transformation was applied to the variables.<sup>1</sup> The regression estimates are reported in Table 2.

We hypothesized that factors driving IT investments may differ between developed and developing countries for reasons presented in the conceptual framework. We define developed countries as those the World Bank defines as high-income countries (n=20), and developing



ure 1). More importantly, the SUR model provides strong empirical support for our hypothesis that the factors shaping investment in developing countries differ from those in developed ones. Here, we discuss our major findings, organized around the four categories in the conceptual framework (see Figure 1).

*Resources for technology investments.* We tested three different types of financial resources that can be channeled into IT investments and found that for the full sample, the level of foreign aid a country receives (AIDGDP), capitalization of equity markets (MARGDP), and level of credits and loans to the private sector (CDTGDP) all influence IT investments. However, when examining the differences between developed and developing countries, we found differences in the type of resources each utilizes. In the developed countries the only significant variable was market capitalization, while in developing countries both foreign aid and credits and loans to the private sector were significant. This makes sense in that foreign aid can be a significant source of capital for developing countries while developed countries rarely receive such aid. On the other hand, developed countries generally have more well-developed equity markets from which to raise capital for IT investments.

*Structure of the economy.* For the full set of countries, IT investment is positively and significantly related to the share of the economy in information-intensive industries (FIRE). The split analysis of developed and developing countries revealed, however, that the effect is significant only in the developed countries. This finding seems reasonable given developing countries generally have a much smaller share of the economy in information-intensive sectors, so that the size of these sectors has not reached critical mass as significant drivers of IT investment. On the other hand, developed countries have likely surpassed that threshold.

Independent Variables	Coefficient	SE
Log (AIDGDP)	0.035**	0.007
Log (CDTGDP)	0.005**	0.001
Log (MARGDP)	0.021**	0.009
Log (FIRE)	0.124**	0.018
Log (TELDEN)	0.003**	0.001
Log (EDTERP)	-0.013	0.029
Log (TRDGDP)	0.139	0.087
Log (FDIGDP)	0.020**	0.008
Log (GDPCAP)	0.558**	0.087
R <sup>2</sup> = 0.816 Adj. R <sup>2</sup> = 0.792 SE = 0.125 N = 521 * p < 0.10 ** p < 0.05		

**Table 2.** Factors shaping IT investments across all countries.

Independent Variables	Developing Countries		Developing Countries	
	Coefficient	SE	Coefficient	SE
Log (AIDGDP)	-0.009	0.010	0.018*	0.010
Log (CDTGDP)	0.002	0.001	0.004**	0.001
Log (MARGDP)	0.044*	0.025	-0.015	0.009
Log (FIRE)	0.101**	0.023	0.032	0.027
Log (TELDEN)	0.001	0.001	0.007**	0.001
Log (EDTERP)	-0.014	0.013	0.019**	0.004
Log (TRDGDP)	-0.141	0.172	0.632	0.104
Log (FDIGDP)	0.016*	0.009	0.047**	0.012
Log (GDPCAP)	-0.057	0.148	0.934**	0.107
R <sup>2</sup> = 0.866 Adj. R <sup>2</sup> = 0.846 SE = 0.108 N = 521 * p < 0.10 ** p < 0.05				

**Table 3.** Factors shaping IT investments in developing and developed countries.

*Complementary assets.* We did not find an effect for education (EDTERP) or density of telephone lines (TELDEN) to significantly impact IT investments for the full sample. However, the SUR model revealed that a significant effect does exist for developing countries but not developed countries. This supports our hypothesis that telecommunications infrastructure has reached a point of saturation in developed countries and therefore has an insignificant marginal impact on IT investment. On the other hand, most developing countries lag behind in telecommunication infrastructure; thus increases in teledensity can still foster additional investment in computers. Similar arguments can be made for human resources. Developed countries already have the necessary human capital in place while developing countries are building their human capital; therefore, we only observe an effect of human capital in developing countries.

*Openness to external influences.* There was no effect for level of international trade (TRDGDP) on IT investment either overall or for the split country analyses. Based on this result, we would conclude that trade in itself does not create competitive pressure or carry with it IT-enabled business practices sufficient to have an impact on IT investment.

However, level of foreign investment (FDIGDP) was significantly related to IT investments as hypothesized. Although positive and significant for both groups, the effect of foreign investment was stronger in developing countries. This is consistent with the argument that developing countries have more to gain from inflows of knowledge associated with foreign investment.

## CONCLUSION

The analysis leads to two general conclusions that have important practical implications. First, our study shows that wealth is the single most important factor influencing IT investment but that other factors are significant as well. This raises the question of how important GDP per capita is relative to other

factors for developing countries, given that relative wealth is difficult to directly influence through policy changes, at least in the short run. To estimate relative effects, we look at the standardized betas in the model and find the highest standardized beta goes to GDPCAP (0.879), followed by FDIGDP (0.559), and EDTERP (0.236). This suggests that the next two factors, which are both directly addressable by policy choices, have roughly as much influence as national wealth on IT investment.

Second, the factors driving diffusion are different for developing economies than for developed ones. The availability of investment resources (loans and foreign aid), the level of complementary assets, and openness to foreign investment all play a role in driving IT investment in developing countries. Again, these factors can be influenced by national development policies as well as by financial aid from international development agencies.

#### POLICY IMPLICATIONS

For developing countries to realize the potential benefits of IT, policymakers should look for ways to promote IT investment as well as developing investment resources, complementary assets, and openness to external influences. The empirical findings suggest several policy recommendations.

- *Resources for technology investments.* The analysis suggests that availability of loans and credit is crucial for developing countries, which means that the maturity and dynamics of the financial system is a key variable for those countries. Case studies have shown that banks in some developing countries perceive IT investment as risky, and they can be hesitant to extend credit for such investments. In such cases, government policy could increase the availability of credit through the banking system by providing loan guarantees or other incentives. The connection between foreign aid and IT investment suggests a role for developed country governments and international organizations such as the World Bank in providing capital to developing countries specifically for IT investments.
- *Complementary assets.* Increasing investment in telecommunications infrastructure, which is usually best accomplished by policies that introduce competition into that sector, will promote IT diffusion. Over the longer term, increasing tertiary education levels will also be beneficial in developing skills needed for IT use. While our data only measured the impact of tertiary education in general, there is evidence from country case studies

that focused efforts to train people in IT skills is important as well.

- *Openness to external influences.* Encouraging foreign investment by removing restrictions and improving the environment for foreign capital is likely to have a major impact. In cases such as Mexico and Brazil, economic liberalization that led to investment by foreign multinationals stimulated IT use. These multinationals required suppliers to adopt IT, and created competitive pressure for domestic firms to invest in IT [5]. Although the benefits of FDI may be greater for firms engaged in the global economy, there may be spillovers for purely local firms through local competition and knowledge transfer. **C**

#### REFERENCES

1. Barrell, R., and Pain, N. Foreign direct investment, technological change, and economic growth within Europe. *Economic Journal* 107, 445 (1997) 1770–1786.
2. Bell, M. and Pavitt, K. Technological accumulation and industrial growth contrasts between developed and developing countries. *Industrial and Corporate Change* 2, 2 (1993), 157–210.
3. Caselli, F., and Coleman II, W.J. Cross-country technology diffusion: The case of computers. *The American Economic Review*, 91, 2 (2001).
4. Coe, D.T., Helpman, E., and Hoffmaister, A.W. North-south R&D spillovers. *The Economic Journal* 107, 440 (1997), 134–149.
5. Dedrick, J., Kraemer, K.L., Palacios, J.J., Tigre, P.B. and Botelho, A.J.J., Economic liberalization and the computer industry: comparing outcomes in Brazil and Mexico. *World Development* 29, 7 (July).
6. Dewan, S., and Kraemer, K.L. Information technology and productivity: Evidence from country-level data. *Management Science* 46, 4 (2000), 548–562.
7. Jorgenson, D.W. Information technology and the U.S. economy. *American Economic Review* 91, 1 (2001), 1–32.
8. Kraemer, K. L., and Dedrick, J. Payoffs from investment in information technology: Lessons from Asia-Pacific region. *World Development* 22, 12 (1994), 1921–1931.
9. Robison, K.K., and Crenshaw, E.M. Post-industrial transformations and cyber-space: A cross-national analysis of Internet development. *Social Science Research* 31 (2002), 334–363.
10. Shih, E, Kraemer, K.L. and Dedrick, J. Research note: Determinants of country-level investment in information technology. *Management Science* 53, 3 (March), 521–528.
11. United Nations Development Program (UNDP). *Human Development Report 2001: Making New Technologies Work for Human Development*. Oxford University Press, New York, 2001.
12. World Bank Group. *Information and Communication Technologies: A World Bank Group Strategy*. Washington D.C., 2002.

---

This research has been supported by grants from the U.S. National Science Foundation and the IBM Corporation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors only.

---

**ERIC SHIH** (e\_shih@yahoo.com) is an associate professor of marketing at the Graduate School of Business, Sungkyunkwan University, Korea.

**KENNETH L. KRAEMER** (kkraemer@uci.edu) is a research professor in the Paul Merage School of Business and Co-Director of the Personal Computing Industry Center (PCIC) at the University of California, Irvine.

**JASON DEDRICK** (jdedrick@uci.edu) is Senior Fellow and Co-Director of the Personal Computing Industry Center (PCIC) at the University of California, Irvine.

---

© 2008 ACM 0001-0782/08/0200 \$5.00

---

DOI: 10.1145/1314215.1314223

By Mark T. Keane, Maeve O'Brien,  
and Barry Smyth

# ARE PEOPLE BIASED *in Their Use of Search Engines?*

**S**earch engines are among the most-used resources on the Internet. Google [2], for example, now hosts over eight billion items and returns answers to queries in a fraction of a second—thus realizing some of the more far-fetched predictions envisioned by the pioneers of the Web [1]. Here, we assess whether people are biased in their use of a search engine; specifically, whether people tend to click on those items that are presented as being the most relevant in the search engine's results list (those items listed at the top of the results list). To test this bias hypothesis, we simulated the Google environment systematically reversing Google's normal relevance ordering of the items presented to users. Our results show people do manifest some bias, favoring items at the top of results lists, though they also seek out high-relevance items listed further down a list. Whether this bias arises from people's implicit trust in search engines such as Google, or some other effect, is a subject for further discussion later in this article.

*Assessing user search behavior when deciding which links to follow in rank-ordered results lists.*

---

The Web provides access to an unparalleled volume of information at time costs that are orders of magnitude lower than those required for traditional media. The critical jumping-off point for this vast repository is typically provided by the results returned by search engines to user queries: Google [2], for example, returns results to over 100 million queries daily, queries that are typically two words long [12]. Like many search engines, Google uses the collective intelligence of the Web to rank-order pages of relevance to a particular query. Each page in this ordered list is typically summarized by a clickable title, some snippets of the page's content (with highlighted matching content words) and a Web address link.

A rational searcher might be expected to assess each of these page summaries against their information need and click on the one that appears as the most relevant. People may not search in such a way, however. Instead, they may manifest biases—they may simply click on top-listed results without checking the results against what is on offer, for example. Such biases, if they exist, could be due to users coming to implicitly trust a search engine. That is, over time, as a search engine consistently delivers relevant pages toward the top of its results lists, users might come to assume the top results are indeed the best. Alternatively, such biases may arise because of a user's tendency to "satisfice," or stop at the first item that contains the sort of content sought, rather than looking for the "best" result among hundreds of relevant results.

The key manipulation in this study was to compare users' responses when they received results lists in their normal ordering versus a systematically reversed order. If people are biased in their search then they will not notice that the relevance rankings have been reversed. That is, they should respond identically to the normal and reversed lists clicking on results placed first at the top of the lists. If they are not biased then they should respond differently to the normal and reversed lists; specifically, they should hunt down the reversed list to find highly relevant items. To presage our results, the truth seems to lie somewhere between these two extremes. There is definite evidence of bias in people's Google searches; they tend to click on first-listed items, though people also sometimes seek out highly relevant results lower down results lists.

## METHOD

Thirty science undergraduates at University College in Dublin, Ireland were paid to participate in the

study. Participants were asked to answer 16 questions on computer science (for example, "Who invented Java?") by running as many queries as they liked on the simulated Google environment. The interface was designed to have the look and feel of Google, and all participants reported they thought it was Google they were using. The simulated system was built using comprehensive search logs from a live user trial in which a separate group was asked to answer the same 16 questions [4]. All of the queries used by this group were stored, as were all of the results lists returned to these queries. We then created a database linking specific queries to their appropriate results list. This database sat behind our simulated Google interface and was used to systematically counterbalance the presentation of results lists in either a normal or reversed ordering when a user entered a query.

The results lists returned to a given query were presented in either their original relevance ordering or a reversed ordering in a counterbalanced way across trials of the experiment. The order in which questions were presented was randomized for each participant to counteract any learning effects that might occur in the course of a trial. Participants were instructed to limit the number of query terms to two or less, to parallel typical user behavior (the average Web user enters between one and three query terms [12]). For each question we recorded the number and names of the queries entered and the search results clicked on by users. The timing of each transaction was also recorded. Participants were asked to complete a form detailing their answers to the questions and sessions averaged 1.5 hours.

We also carried out a ranking post test to see whether people agree with Google's relevance ordering of results. This post test was carried out on a sample of the results lists using a new group of 14 students. These participants were asked to manually rank the presented results lists from the search experiment (on a 1–10 scale from "most likely to select" to "least likely to select"). A sample of 16 results lists from the experiment were used, based on those results lists returned to the most frequently used query for each of the 16 questions. This sample thus covered those results lists that contributed most to any effects found in the experiment. Each participant received the results lists in a randomized order and the results in each list were also randomized for every participant. This procedure was adopted to ensure an accurate assessment of people's relevance ranks, independent of any possible bias effect. People were given one hour to complete this ranking task during which participants only completely ranked a subset of the presented results sets.

## RESULTS AND DISCUSSION

The dependent measure was the users first click, that is the first chosen link by a given user in a returned results list to a given query. The data was analyzed in a 2 (condition; normal versus reversed) x 10 (relevance rank; 1–10) design treating queries as the random factor. That is, for each query we recorded the proportion of people that chose a particular ranked result, noting whether this occurred in a list that was normal or reversed. The two-way analysis of variance (ANOVA) with condition and relevance rank revealed a main effect of relevance rank [ $F(9,319) = 102.14$ ,  $p < 0.01$ ,  $MSe = 0.89$ ], and a reliable interaction between the condition and relevance rank [ $F(9,319) = 11.31$ ,  $p < 0.01$ ,  $MSe = 0.10$ ]. Tukey's post-hoc comparisons of the interaction showed there were reliable differences between the first-click frequencies for the 1st, 9th, and 10th relevance ranks (see the figure here).

These results clearly indicate people's first clicks in the normal and reversed conditions are not identical, providing evidence that people are partially biased in their search activities. Items with the highest-relevance ranks (items ordered first by Google) are chosen 70% of the time in the normal condition, but this rate drops to 10% in the reversed condition. In contrast, the 9th and 10th relevance-ranked items are chosen more often (13% and 41%, respectively) in the reversed condition than in the normal one (2% and 2%, respectively). Intermediately ranked items are much the same across both conditions.

The significance of what is happening is readily apparent if one considers the data by position in the results lists. The accompanying figure shows that when lower-relevance-ranked items are positioned first and second in the results list (as they are in the reversed condition) they are chosen more often by users, despite their limited relevance. In contrast, when the highest-relevance items are positioned last in the results list (in the reversed condition) they are being chosen by users considerably less often. In short, users are, in part, misled by the presented order of the items. However, sometimes people deliberately hunt out the highly relevant items even when they are located at the very bottom of the returned list.

The post test showed there is close agreement between people's rank of returned results and those rankings posed by Google. People's mean rankings of the sampled results lists correlate highly with the search engine rankings ( $r^2 = 0.9124$   $t = 9.13$ ;  $df = 8$ ;  $p < 0.0001$ ). This result shows us that the items Google presents as the best are considered by people to be the best too. It is interesting this finding occurs even when people have been given the results lists in a randomly re-ordered form, suggesting highly relevant items in each results list were easily identifiable. This post test also sheds some light on another issue regarding the relevance topology of the

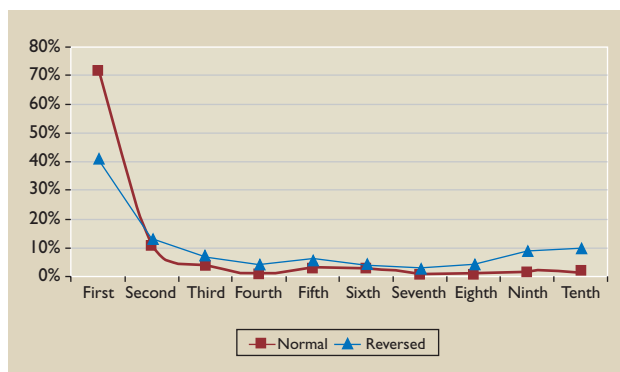
results lists. One concern about the evidence is that the first 10 results in each list are approximately equal in relevance and the real relevance differences only begin around the 100th or 200th ranked items. If this were the case then the search behavior observed would only apply to results lists with

flat, relevance topologies. This concern is partly answered by the correlation reported here, but not fully.

To get a better idea of the actual relevance topology we analyzed the rankings produced by people in the post test in a different way. For each of the 16 results lists sampled, we noted the mean rating given by people to each result in the list. If the relevance topology is flat for these lists then these mean ratings should all be approximately equal (recall, order effects are controlled for this data by randomization). However, this is not what we found. There is a huge variety of different topologies for the results in each list; a few have a single highly relevant item (with a mean rank of 1 or 2), others have several results given high mean ranks, while others have a linearly increasing relevance topology. This finding suggests our random selection of questions for the experiment generated a random selection of different relevance topologies; these are presumably representative of the topologies generated by Google. Furthermore, they are not all flat but hugely varied.

## GENERAL DISCUSSION

Our study results clearly show people are partially biased in their search behavior. While it is known that people have a fondness for items at the beginning of written lists, the novelty of our study is that it demonstrates such effects within a search engine



The normalized percentage of first clicks by position in the list in the normal and reversed conditions.

context (through our systematically controlled forward-reversed paradigm). So, given that we have evidence of such bias, the difficult question to answer is “Why?”

Recently, Joachims et al. [5] carried out an independently conducted similar experiment to the present one using an eye-tracking paradigm that interprets its findings as being due to people’s development of an implicit trust in search engines. However, other recent work does not conclusively support the “trust hypothesis,” for instance, O’Brien and Keane [7] have found this bias exists even when search results are presented as simple text lists.

**A**n alternative possibility is that the bias is a function of “satisficing” search heuristics on behalf of users, where users seek satisfactory as opposed to optimal results. Our findings seem closer to this type of search behavior in that we only find a partial bias; that is, people do sometimes search to the bottom of the list to find the highly relevant items. Also, O’Brien and Keane [7] have observed different click patterns across different result distributions where a highly relevant result coming before many irrelevant results stands a greater chance of being chosen over the same highly relevant result preceded by other relatively relevant results, for instance.

O’Brien and Keane [8] have modeled users interacting with search results adopting a satisficing strategy, accommodating both the findings presented in the current study and eye-tracking evidence that suggests users tend to evaluate results sequentially, deciding immediately whether or not to click [6]. O’Brien and Keane’s model [8] predicts users will, in general, tend to click on top results over results lower down the list, though this tendency should not be as strong when the relevance of the top results is weakened. They [8] demonstrate how the model, across a number of trials, approximates the aggregate search behaviors of large numbers of users searching the Web.

On the whole, our results suggest search engines could misleadingly overpromote an initially popular page because, having placed it at the top of the results list, it is clicked on preferentially by users, in turn increasing the likelihood of it being placed first, being clicked on, and so on (see also [3, 7, 8, 10]). This problem obviously applies to search engines that rely on histories of previous user choices (for example, [11]), but it could also apply to those linkage-based algorithms such as Google’s PageRank [9] because the top-of-the-list pages are more likely to end up as the chosen link on people’s Web pages. Search engine designers may

need to design systems to overcome such effects (for a solution see [10]), and it is clear that future information delivery systems have much to learn from such detailed analyses of user search behavior. **C**

## REFERENCES

1. Berners-Lee, T., Cailliau, R., Groff, J., and Pollermann, B. World Wide Web: The information universe. *Electronic Networking: Research, Applications, and Policy* 2, 1 (1992), 52–58.
2. Brin, S. and Page, L. The anatomy of a large-scale hypertextual Web search engine. In *Proceedings of the 7th International World Wide Web Conference (WWW '98)* (1998), 107–117.
3. Cho, J. and Roy, S. Impact of search engines on page popularity. In *Proceedings of the 13th International World Wide Web Conference (WWW '04)* (2004).
4. Freyne, J. and Smyth, B. Collaborative search: A live user trial. In *Proceedings of the 26th European Conference on Information Retrieval Research (ECIR '04)* (2004).
5. Joachims, T. et al. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the Conference on Research and Development in Information Retrieval (SIGIR '05)* (2005).
6. Klöckner, K., Wirschum, N., and Jameson, A. Depth- and breadth-first processing of search result lists. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)* (2004).
7. O’Brien, M. and Keane, M.T. Modeling result-list searching in the World Wide Web: The role of relevance topologies and trust bias. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (2006).
8. O’Brien, M. and Keane, M.T. Modelling user behaviour using a search engine. In *Proceedings of the 12th International Conference on Intelligent User Interfaces (IUI '07)* (2007).
9. Page, L., Brin, S., Motwani, R., Winograd, T. *The PageRank Citation Ranking: Bringing Order to the Web* (1998).
10. Pandey, S., Roy, S., Olston, C., Cho, J., and Chakrabarti, S. Shuffling a stacked deck: The case for partially randomized ranking of search engine results. In *Proceedings of the 31st international Conference on Very Large Data Bases (VLDB '05)* (2005).
11. Smyth, B., Freyne, J., Coyle, M., Briggs, P., and Balfe, E. I-SPY—Anonymous, community-based personalization by collaborative meta-search. In *Proceedings of the 23rd SGAI International Conference on Innovative Techniques Applications of Artificial Intelligence* (2003).
12. Spink, A., Wolfram, D., Jansen, B.J., and Saracevic, T. Searching the Web: The public and their queries. *Journal of the American Society of Information Science and Technology* 53, 2 (2001), 226–234.

---

**MARK T. KEANE** (mark.keane@ucd.ie) is the chair of the School of Computer Science and Informatics and Vice President for Innovation at University College Dublin, Ireland.

**MAEVE O'BRIEN** (maeve.m.obrien@ucd.ie) is a Ph.D. candidate in the School of Computer Science and Informatics at University College Dublin, Ireland.

**BARRY SMYTH** (barry.smyth@ucs.ie) is the digital chair of the School of Computer Science and Informatics at University College Dublin, Ireland, and the founder and chief technical officer of Changing Worlds Limited, Dublin, Ireland.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2008 ACM 0001-0782/08/0200 \$5.00

DOI: 10.1145/1314215.1314224

By Nolan J. Taylor, Eleanor T. Loiacono, and  
Richard T. Watson

# *Alternative Scenarios to* THE “BANNER” YEARS

Online advertising spending continues to grow, increasing by 35% from 2005 to 2006 to reach \$16.9 billion. This trend is expected to persist with a projected 2011 figure of \$36.5 billion [8]. Despite this trend consumers have shown increasing annoyance with online advertising, resorting to specialized software and restrictive Web browser settings in an effort to reduce the number of pop-ups and other forms of advertising [1]. Although the banner format is still the most prominent form of Web advertising, corporations are starting to give alternative formats serious consideration.

One potential alternative to the banner is the use of background graphics, which present information in a less obtrusive manner. Web pages typically employ background images for aesthetic reasons, but rarely to convey specific information. Both the foreground and background of Web pages provide a means of exposing consumers to brands and logos. Unlike traditional

*A test of alternative  
formats to Web banner  
ads.*

advertising, the intent of these ads is to increase product awareness rather than direct sales generation through click-through. The potential use of ads for increasing brand salience is not unreasonable, given that most banner ads are merely seen and not acted upon. Ironically, while most Web pages use some form of background imagery, little is known about the effects of background image exposure on consumer perceptions in an online advertising context.

Background images possess several features that make them an attractive alternative to traditional advertising formats, such as banner ads. First, background images can be larger in size because increasing the size of these images does not affect the text layout. Second, the background images can increase a visitor's exposure time because they are visible even if the user scrolls to the bottom of the page. Finally, being an integral part of a page, background exposures are not targeted by the assortment of ad blocking software currently on the market. This ensures the visitor is exposed to the ad even if other traditional ad formats are blocked. This study compares the recall and recognition of the typical (468 x 60) banner ads to two versions of background wallpaper ads: background images that use the entire page and sidebar images that use the left margin of the page.

#### WEB ADVERTISEMENT FORMATS

The influence of advertising can be both covert and overt. Exposure to information may influence a person's opinions, attitudes, or choices without his or her explicit recollection of the causal event [11]. Viewers are able to recall ad banners even after only a single exposure and these exposures may result in significant attitude changes [3]. These attitude changes, in turn, can ultimately influence behavior. On the Web, behavioral intentions include a wide range of activities, including actions to add a Web page to a user's revisit consideration set (for example, creating a bookmark or desktop shortcut), a return visit to a particular Web page, or an actual purchase. For attitude and behavioral changes, remembering a brand, consciously or unconsciously, is a key first step.

Even as advertisers have standardized on a relatively small number of advertising formats, changes in Web standards such as Dynamic HTML and scripting languages extensions (including Macromedia's Flash, Microsoft's ActiveX, and Apple Computer's Quicktime) offer new opportunities for conveying messages to Web users. Banner ads, the de facto format for Web advertising, are colorful, rectangular images that are hyperlinked to more detailed brand information. These images are prominently displayed and are usually one of the first objects to appear as a



Web page is loaded.

Figure 1. Foreground (banner) format.

A seemingly less-intrusive alternative solution to banner ads is to place the image (or a modified version) in the background. The HTML standard provides for the inclusion of an image that appears superimposed over Web page text. Prior research has examined a number of aspects of these backgrounds, including how background complexity influences recall [5] and purchase intention [12], as well as how background "priming" affects product preferences [9].

In general, HTML background images can take one of two forms. In the first, the image is superimposed over the text (that is, background). Because these images are in the background, the brand is actually intertwined with the main content of the Web page. In the second, the image is separate from the main text of the page. This is achieved by widening the text margins to expose part of the background image. Thus, like banners, these images are prominently displayed yet still enjoy the advantage of being HTML background images.

Background images have a number of advantages over the traditional banner format. One such advantage is that scrolling down the page does not result in the banner image being placed outside of the viewing area. This scroll-off effect associated with the banner format places an upper limit on a viewer's attention time. A solution to scroll-off and its effects on exposure time is to increase the size of the image relative to the page. Increasing the size of images in the foreground, however, carries with it the penalty of displacing non-advertising information and slower loading times. These penalties have been recognized by a growing number of advertisers who have resorted to separate non-scrolling areas for banner advertising, either within the page (a frame) or as a separate window (a pop-up). Using HTML background images allows the ad information to be visible to the viewers as long as they remain on the page. In general, a background image on a Web page will have a longer expo-





Figure 2. “Background” format.

sure time than a foreground image because of the latter’s susceptibility to the scroll-off effect. Longer exposure time increases the opportunity for a viewer to notice the stimuli and, as a result, the opportunity to form a mental representation of the stimuli [7].

A second beneficial effect is that HTML background images result in a repetition effect because images in the background are repeated to fill the full area of a page. The advantage is not in the repetition itself, but rather because the background ad image has a higher probability of being one of the last items seen (and thus remembered) on a page by a visitor. Previous research has suggested that “recency effects” (from a user having just seen an image or piece of text) have a positive influence on information retention [10]. Thus, while foreground images ordinarily have a memory advantage due to their prominence [6], exposure time and recency effects should serve to a minimum to equalize the memory disadvantages of the less prominently placed HTML background images. Given these advantages, one would expect users would recall and recognize ads placed in the background or sideground (termed sidebar ads) of a site with greater ease and frequency than those placed in the foreground. Sidebar ads, with both prominence and repetition advantages, are posited to perform the same or better than background ads.

To test these alternative ad formats an online experiment was conducted using ads of equal size and layout to avoid confounding. Subjects were 155 undergraduate students enrolled in an introductory psychology class at a large state university; 61% were male.

**Research Method.** A traditional banner ad was tested against each of the background format and sidebar format, which in turn were both tested using two levels of repetition (scrolling with the text stationary). In the scrolling condition, the ad image and the text moved in synchronization, while in the stationary (watermark) condition, the image did not move even when the text was scrolled. The result was

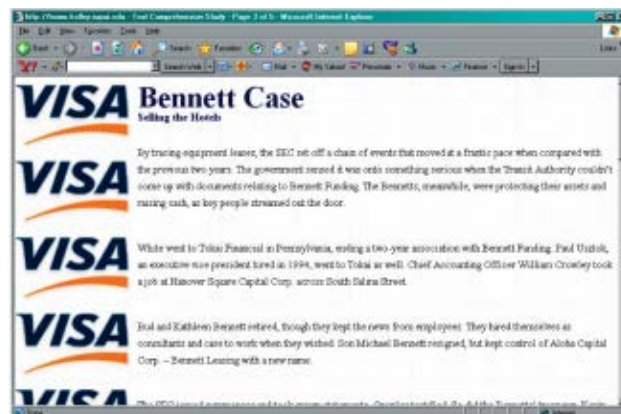


Figure 3. “Sidebar” format.

a two- (ad format) by-two (repetition) between-subjects design with the banner ad serving as the control condition for comparison purposes.

**Stimuli and Procedures.** Besides the traditional ad banner (see Figure 1), two variants of the HTML background ads were created and tested (see Figure 2 and Figure 3). To control for possible effects due to ad size, the banner image was used to generate both the sidebar and background images while altering neither the dimensions of the non-white space portion of the ad nor the content of the image. The background format was faded slightly to improve the readability of the text on which it would be superimposed. Credit cards, a high familiarity product category, were chosen as the focus of this research. Furthermore, the tested brand (Visa) is a high familiarity brand within this group.

In order to avoid undue focus on the ads, subjects were told the purpose of the research was to examine reading comprehension. A business case (Fall from Grace: Anatomy of the Bennett Funding Collapse), discussing the ethical questions regarding a “reputable” business that used an information system glitch to its advantage, was used as a filler task so that subjects would not spend undue attention on the advertisements. The case was divided into five approximately equal-length Web pages. The ad appeared on the third page in all treatments. In order to more closely simulate normal Web conditions, the other four case pages contained non-brand distracter images so that subjects would not give undue attention to the treatment page. Thus, all five pages were relatively uniform in appearance. (In the case of the banner page, a faint background containing no text or brand information was included so that all five case pages had a background). Scripts were used to randomly assign treatments, record responses, and measure the time spent on the treatment page. At the end of the case, subjects answered questions related to the advertisements, the content of the story, and a number of demographic items.

To ensure the validity of our results, we held certain factors constant. Prior exposures to the Web site and specific ads used were controlled through the development of unique Web pages and ads for this particular study. All subjects read the same business case. There was no mention of the brands chosen for this study (or their competitors) within the case. Subjects were not previously exposed to any similar tests.

**Measures.** Independent measures included the manipulated factor (ad format and ad repetition) as well as gender, Web familiarity, and exposure time as controls. Web familiarity was measured using the item “How would you rate your experience level with using the Web?” with endpoints “not familiar” (1) and “familiar” (7). Exposure time to the ad was also recorded. A script was used to measure the time between the initial display of the page containing the ad and the initial display of the following page. For simplicity the view times were converted into three categories: short (less than one minute), medium (one to three minutes), and long (greater than three minutes). It should be noted that each measured view time necessarily overestimates the actual view time for the banner format because the length of the case text precluded reading the entire page without the banner being scrolled out of view.

Two dependent measures also were included, those being brand recall and recognition. Brand recall was measured by asking subjects to list as many brand names as they could remember from the case they just viewed. Ad recognition was measured by presenting a list of brands that included the pictured brand (Visa) and six other brands that did not appear in the case pages. Subjects were asked to indicate whether or not they recognized the brand from the exercise.

**Results.** Results were analyzed via separate logistic regressions for recall and recognition. The average recall and recognition rates are reported in Table 1. Analysis using the banner format (foreground) as the reference group indicated that banner image recall performance was similar to that of sidebar (43% vs. 49%;  $p=0.70$ ) and background format performance (43% vs. 36%;  $p=0.091$ ). Banners were remembered similarly to sidebar images (91% vs. 90%;  $p=0.39$ ) but more than background images (91% vs. 66%;  $p=0.002$ ). Analysis comparing background images to sidebar images indicate sidebar images are no better in

terms of recall (49% vs. 36%;  $p=0.081$ ), but that sidebar ads are superior to background ads in terms of ad recognition (90% vs. 66%;  $p=0.002$ ). In addition, longer exposure times and watermark images (as opposed to scrolling images) resulted in better recall, while differences in recognition were attributable to ad format alone.

One important question is whether alternative formats would be viewed negatively with respect to current standards. Negative effects for the alternative formats could, in turn, negatively impact the attitude toward the brand as well. Attitude toward each ad format was measured via three seven-point Likert scales ranging from strongly disagree (1) to strongly agree (7).

These include, “I enjoyed the advertisements” and “These advertisements were annoying” (reverse scaled). The measures displayed a high level of reliability (Cronbach’s alpha = 0.77). An analysis of variance (ANOVA) was conducted comparing average

attitude ratings of the three formats using a second group of 157 respondents. Attitudes toward the traditional ad format ( $M=3.30$ ,  $SD=0.99$ ) did not differ significantly for either background ( $M=3.38$ ,  $SD=0.98$ ) or sidebar ( $M=3.72$ ,  $SD=1.07$ ) treatments ( $F(2,150)=2.59$ ;  $p=0.78$ ). Thus, the alternative formats were just as acceptable as the current banner standard.

## MANAGERIAL IMPLICATIONS

The results of this study suggest alternative ad placements can act as a viable alternative to banner ads when increasing brand salience is the goal. Sidebar ads (non-scrolling) proved to be the most effective format. The performance of background ads was slightly worse than the performance of banners, but this may have been due in part to the fading process necessary to ensure text readability and the position of the text that was superimposed over the image. Given these findings, it is clear that sidebar advertisements are a sound alternative to banner ads, though further research is warranted. Consistent with previous banner ad research [3], these findings show the significant impact alternative ads can have on advertised brands even with a single exposure (see Table 2).

What is revealed by this research is that, like previous studies of “traditional” media outlets (print, television, and movies), ads that are less conspicuous to viewers (such as sidebar and background ads) still affect their brand recollection and recognition. Fur-

Format	Repetition	Recall	Recognition
Foreground		42.9%	91.4%
Sidebar	Watermark	58.1%	93.5%
	Scrolling	40.0%	86.7%
	<b>Total</b>	<b>49.2%</b>	<b>90.2%</b>
Background	Watermark	42.4%	75.8%
	Scrolling	27.6%	55.2%
	<b>Total</b>	<b>35.5%</b>	<b>66.1%</b>

Table 1. Brand recall and recognition means.

thermore, because such ads cover a larger area of the screen than banners, they are less susceptible to the ad avoidance (scroll-off) behavior typically given to traditional banners. In the current study, we restricted the alternative formats to banner-sized images but, in actual applications, larger ads could be used to potentially achieve even higher levels of recall. Thus, companies seeking to expand their brand awareness through Web advertising have additional options not previously considered. Given that companies will undoubtedly want to retain the ability to receive direct response from ads (click-through) in some situations, we see two scenarios for alternative ad formats: first, as enhancers to clickable formats, and second, as primers to future brand information, including clickable and traditional ads, ads in other media, and the brand itself (see Table 3).

The findings also suggest marketing managers would be wise to redirect some funds toward these alternative ad formats given the increasing popularity of ad blocking software. Banner ad blocking software firms, such as AdSubtract, tout the benefits of their software (most notably faster download times) through customer feedback. Through the use of the HTML background option used in this study, sidebar and background placements allow adequate download time while still allowing the brand ad to load.

### LIMITATIONS AND FUTURE RESEARCH

As with any research, there are limits to this study that must be understood. First, the sample, drawn from a university population, is relatively young (the median age is 19 years old). Although younger users have historically made up the majority of Web users, the Web user population is becoming increasingly diverse. Further research on Web users of varying ages is required in order to generalize across age groups.

Second, this study employed a business case as its focus that required subjects to search within rather than between pages. Because subjects were not searching for any particular facts within each page, they proceeded from top to bottom with little or no deviation. In cases where subjects search for particular facts within a page, search patterns are likely to be less linear and result in different ad exposure patterns and exposure times. Thus, it is necessary to further explore if the results obtained here hold true across

Expected Results	Supported?	
	Recall	Recognition
Ads on the Web placed in the background (across entire page) are remembered equal to or more than those placed in the foreground (banner) ads.	YES (B=F)	NO (B<F)
Ads on the Web placed in the sidebar (left margin of page) are remembered equal to or more than those placed in the foreground (banner) ads.	YES (S=F)	YES (S=F)
Ads on the Web placed in the sidebar (left margins of page) are remembered equal to or more than background (across entire page) ads.	YES (S=B)	YES (S>B)

F= Foreground (banner), B=Background, S=Sidebar

Table 2. Summary of results.

different task types and viewing patterns.

Third, this study focused on a highly familiar brand, Visa, with readily recognizable symbol, name, and brand colors. Although this undoubtedly enhanced recognition and recall across all three treatments, treatment differences may be less (or more) pronounced with less familiar brands. Future research should therefore examine less familiar brands in order to give advertisers information on how to market the two types of brands (high and low recognition) most effectively online.

Finally, the effects of adding animation or interactivity to banner ads was not included in this study for control purposes. Although previous research has





Alternative Ads as Enhancers		As a <b>memory enhancer</b> – In conjunction with clickable ads, alternative formats may be used to increase the overall exposure time, therefore leading to higher recall/recognition levels.
		As an <b>affect enhancer</b> – Alternative formats of established brands can be used to improve the attitude toward newer brands and brand extensions through the power of association.
Alternative Ads as Primers		As a <b>brand primer</b> – Displayed in multiple locations throughout the site, the alternative formats can be used to increase the receptiveness to the brand or advertisement for the brand at a point in the future. Similar methods are frequently employed in film and television (brand placements).
		As an <b>advertisement primer</b> – Alternative formats can be used to convey low-detail symbols or logos as a prelude to the full detail ads of an ad campaign. Previously, similar methods have been employed for new product launches (for example, in 1989 Nissan "pre-announced" its Infiniti series with an ad spot that never showed the product).

Table 3. Potential uses for alternative Web advertisement formats.

indicated that animation does not enhance a consumer's memory (recall and recognition) of online banner ads [2, 4], sidebar animation may have a greater impact and is worthy of a follow-up investigation. It is also possible that the alternative formats can be combined (for example, sidebar and background together) or used in conjunction with traditional banner ads and subsequent research should include both competitive and complementary aspects of these formats. **C**

## REFERENCES

1. Armitt, C. One-third of Internet users opposed to intrusive ads. *New Media Age* (2004), 12.
2. Bayles, M.E. Designing online banner advertisements: Should we animate? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (Minneapolis, MN, 2002), ACM Press, 363–368.
3. Briggs, R. and Hollis, M. Advertising on the Web: Is there response before click-through? *Journal of Advertising Research* 37, 2 (1997), 33–45.
4. Burke, M., Gorman, N., Nilsen, E. and Hornof, A. Banner advertisements hinder visual search and are ignored. *Extended Abstracts of ACM CHI 2004: Conference on Human Factors in Computing Systems*, (Vienna, Austria, 2004), 1139–1142.
5. Danaher, P.J. and Mullarkey, G.W. Factors affecting online advertising recall: A study of students. *Journal of Advertising Research* 43, 2 (2003), 252–267.
6. DeLorme, D.E. and Reid, L.N. Moviegoers' experiences and interpretations of brands in films revisited. *Journal of Advertising Research* 28, 2 (1999), 71–94.
7. Gibson, L.D. What can one TV exposure do? *Journal of Advertising Research* 36, 2 (1996), 9–18.
8. Internet Advertising Bureau and PriceWaterhouseCoopers. Internet advertising revenues grow 35% in '06, hitting a record close to \$17 billion. May 23, 2007; [www.iab.net/NEWS/pr\\_2007\\_05\\_23.asp](http://www.iab.net/NEWS/pr_2007_05_23.asp).
9. Mandel, N. and Johnson, E.J. When Web pages influence choice: Effects of visual primes on experts and novices. *Journal of Consumer Research* 29, 2 (2002), 235–245.
10. Miller, N. and Campbell, D. Recency and primacy in persuasion as a function of the timing of speeches and measurements. *Journal of Abnormal and Social Psychology* 59, 1 (1959), 1–9.
11. Smith, R.E. and Swinyard, W.R. Information response models: An integrated approach. *Journal of Marketing* 46, 1 (1982), 81–93.
12. Stevenson, J.S., Bruner II, G.C. and Kumar, A. Webpage background and viewer attitudes. *Journal of Advertising Research* 40, 1/2 (2000), 29–34.

**NOLAN J. TAYLOR** ([notaylor@iupui.edu](mailto:notaylor@iupui.edu)) is an assistant professor in the Kelley School of Business at Indiana University, Indianapolis.  
**ELEANOR T. LOIACONO** ([eloiacon@wpi.edu](mailto:eloiacon@wpi.edu)) is an assistant professor at Worcester Polytechnic Institute, Worcester, MA.  
**RICHARD T. WATSON** ([rwatson@terry.uga.edu](mailto:rwatson@terry.uga.edu)) is the J. Rex Fuqua Distinguished Chair for Internet Strategy in the Terry College of Business at the University of Georgia, Athens.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2008 ACM 0001-0782/08/0200 \$5.00

DOI: 10.1145/1314215.1314225



**University of Puerto Rico at Mayaguez**  
**College of Engineering**  
**Department of Electrical and Computer Engineering Doctoral Program in CISE**

The Department of Electrical and Computer Engineering (ECE) of the University of Puerto Rico at Mayaguez (UPRM) invites applications for tenure-track positions in Computer Sciences and Engineering. The Department is interested in attracting faculty in the area of software engineering, databases or networking. Candidates are expected to make important contributions to research activities being conducted under the doctoral program in Computing and Information Sciences and Engineering (CISE), either by enhancing and strengthening current projects or starting new research tracks. Applicants must possess a PhD degree in computer science or computer engineering or closely related field, and demonstrate strong potential for excellence in research. If applicant's native language is not Spanish, applicant should be able to communicate fluently in English, and by the end of the tenure-track appointment be able to communicate effectively in Spanish.

The ECE Department offers a Baccalaureate and Masters in Computer Engineering and the PhD in CISE jointly with the UPRM Department of Mathematics. For further information please visit [www.ece.uprm.edu](http://www.ece.uprm.edu) and [www.phd.cise.uprm.edu](http://www.phd.cise.uprm.edu) or call 1-787-833-3338.

Applications with a curriculum vitae, three reference letters, and MS and PhD original transcripts, should be sent no later than February 29, 2008, to CISE Search Committee University of Puerto Rico at Mayaguez P.O. Box 5028 Mayagüez, PR 00681-5028

Or in pdf format to [ciseph@ece.uprm.edu](mailto:ciseph@ece.uprm.edu) Appointments will be made by July 1st, 2008.

UPRM is an equal opportunity affirmative action employer.

## University of Zurich

### Department of Informatics

The Department of Informatics at the University of Zürich is inviting applications for a tenured faculty position in the area of Human Computer Interaction.

We seek applications for an associate or full professorship with a strong research focus on concepts, tools and empirical studies in human-computer interaction. Excellent candidates in the early stages of their academic careers, with a strong research record but that do not readily qualify for a full professorship, as well as women, are strongly encouraged to apply.

The University of Zürich is the largest university of Switzerland and regularly ranked among the top international research universities. Applications should include a detailed resume, research and teaching statements, and the names and addresses of at least three references.

Primary consideration will be given to applications received by March 3, 2008.

Additional information on the department as well as the position can be found at <http://www.ifi.uzh.ch/profhci>

Applicants can also contact the chair of the search committee Prof. Abraham Bernstein, Ph.D. <http://www.ifi.uzh.ch/~bernstein>

essays, {craft, art, science} of software, python, eclipse, agile development, onward!, {generative, functional} programming, .net, open source, concurrency, smalltalk, aspects, second life, ruby, service-orientation, objects, embedded, ultra large scale systems, objective-c, {model, test}-driven development, c#, passion, fun!, agents, design patterns, domain-specific languages, wiki, use cases, movies, product-lines, java, lightning talks, refactoring, plop



## DEFINE THE FUTURE OF SOFTWARE

[www.oopsla.org/submit](http://www.oopsla.org/submit)

### CONFERENCE CHAIR

**GAIL E. HARRIS**

Instantiated Software Inc.  
chair@oopsla.org

### PROGRAM CHAIR

**GREGOR KICZALES**

University of British Columbia  
papers@oopsla.org

### ONWARD! CHAIR

**DIRK RIEHLE**

SAP Research  
onward@oopsla.org

### CALL FOR PAPERS

March 19, 2008

Due date for Research Program, Onward!, Development Program, Educators' Symposium, Essays and proposals for Tutorials, Panels, Workshops and DesignFest

July 2, 2008

Due date for Development Program Briefs, Doctoral Symposium and Student Volunteers



Association for  
Computing Machinery

NASHVILLE CONVENTION CENTER, NASHVILLE, TN  
October 19 - 23, 2008

BY PETER MARKS,  
PETER POLAK, SCOTT MCCOY, AND  
DENNIS GALLETTA

*How managerial prompting, group  
identification, and social value orientation  
affect knowledge-sharing behavior.*

## SHARING KNOWLEDGE


Knowledge management systems (KMS) have taken on a prominent role in both information systems research and practice. Such systems promise the potential for contributing to an organization's strategic advantage by unlocking knowledge heretofore housed only in the minds of certain organizational members. This potential has led to a rapid increase in the volume and types of systems for sharing knowledge. Recently, many issues and strategies relevant to KMS have been identified [7].

Sharing knowledge throughout an organization has intuitive appeal. If organizational members share valuable information freely with other members, the organization's responsiveness and effectiveness can be greatly augmented by preventing those members from having to repeatedly solve the same problems. In an environment of organizational sharing, a KMS can readily save time and money for both providers and users of knowledge. For example, more than a decade ago Microsoft created its "Knowledge Base" (KB), which was the first place employees checked when a software problem surfaced. Microsoft's decision to

---

ILLUSTRATION BY ROBERT NEUBECKER

Two overlapping speech bubbles, one blue and one red, containing illegible cursive handwriting.



move the KB to a customer-focused extranet is ample testimony to its usefulness; now customers query the system without having to contact a fee-based support center, saving customers the fees and/or call, and saving Microsoft from having to build, staff, and manage a response center.

Because a KMS requires individuals to share their knowledge, success is not guaranteed; several factors can interact to hinder knowledge sharing [2]. One significant problem is that individuals often gain considerable power when they hold unique knowledge [5], especially when it is perceived to be of high value. They are sometimes hesitant to give up their competitive advantage by sharing it. Therefore, KMS adoption and use might require substantial encouragement by management (prompting), intense loyalty to the firm (group identification), or perhaps even simply personal values that support sharing (social value orientation).

#### PRIOR RESEARCH

Researchers have studied how sharing behavior relates to managerial prompting, group identification, and social value orientation [4]. Here, we describe the previous research on these three facilitating factors, the design of our experiment, our findings, and managerial implications.

**Knowledge sharing.** Knowledge in a KMS can be considered to be a public good, an asset available to all members of a community or organization regardless of whether the members contributed to its constitution [9]. Examples of public goods include public television, public parks, and clean air. Unfortunately, the fundamental issue with public goods is the “free rider” problem, whereby certain individuals will utilize a public good without contributing their fair share. Thorn and Connolly [11] asserted that information in a shared database will generally be undersupplied, as is the case with other public goods. Individuals are not likely, therefore, to share their valuable, personally held information through KMS.

Positive motivators have been identified that entice individuals to contribute even when personal costs of doing so are high. Factors that are likely to affect knowledge sharing include repeated direction (prompting) by managers, identification with a group engaged in a common task, and the person’s fundamental predisposition to sharing. Each factor is described in more detail here.

**Managerial prompting.** To influence employees’ attitudes, behaviors, and activities, organizations sometimes exert control [1] to achieve particular goals. Prompts are cues provided by managers to remind workers what is expected of them and can be informally described in highly descriptive yet colloquial language as

“nagging.” We do not attach value judgments about such prompts; they are simply an attempt by management to influence the behavior of individuals within an organization. As described here, prompting workers with meaningful reminders can have a strong desired effect on behavior [4]. Therefore, we expect that in situations where management encourages knowledge-sharing behavior by prompts or reminders, employees will contribute more knowledge to a KM system.

**Group identification.** Identification with a group will create allegiances that motivate individual behavior to benefit the group. Group identification has been examined in public goods research to determine why some individuals share more willingly than others in resource dilemmas [8].

Group identification accounts for the contributions from people to a common good [4]. Although usually considered an individual difference variable, it was found in De Cremer and Van Vugt [4] that, when certain managerial prompts reminded individuals about their group membership, “selfish” individuals increased their cooperation. Therefore, we expect to find that people who strongly identify with their team will share more knowledge for the benefit of the whole group. Individuals who do not feel they belong to the group will withhold valuable knowledge and contribute less to a KMS.

**Social value orientation.** Individuals can be described by three different profiles: as collectivists, competitors, or individualists [10]. Collectivists try to ensure they are sharing any returns in an equal manner with others. Competitors try to maximize their return by producing the maximum amount of separation between themselves and others. They are not as concerned with maximizing their own utility in the traditional sense of making a maximum amount of revenue. Rather, they obtain their utility by maximizing the distance between themselves and other competitors. So, competitors seek to increase the distance between themselves and others. In comparison, individualists try to maximize their own return regardless of what others are doing.

In many studies focusing on individuals’ contributions, the three categories are often combined into two groups: individualists and competitors are labeled proself and collectivists are labeled prosocial. Proself individuals are most concerned that the outcome in any social dilemma is based on their ability to make a gain. Prosocial individuals are satisfied when group collective goals are met in social dilemmas. It is proposed here that a person’s orientation (prosocial or proself) has an impact on the decision to share in a public goods dilemma. Therefore, we expect that prosocial individuals will contribute more knowledge to a KMS than proself individuals.



**Other factors.** Managerial prompting, group identification, and social value orientation are not the only factors that could be considered when studying elements that impact knowledge sharing. However, they are related through the work of De Cremer and Van Vugt [4] because they are common elements in most organizational environments. Therefore, they are pertinent factors in the study of organizational dynamics in modern business environments.

**Interactions.** It has been shown in other contexts that the three principal variables described here affect the amount of sharing by individuals when examined in isolation. It would be more interesting to investigate how these variables interact among themselves and how they together as a whole affect the amount of knowledge sharing. It is expected that proself individuals will make fewer contributions for the good of the group, but the institution of managerial prompts can lead them to contribute. On the other hand, it is plausible to expect that prosocial individuals do not find the prompts for sharing necessary, as they already understand and appreciate the need to maximize the good of the group. Therefore, we propose that managerial prompts will be more effective in increasing sharing with proself individuals than with prosocial persons.

Similarly, proself individuals are expected to be affected more by group identification cues than prosocial individuals. Prosocial individuals are likely to already be thinking about the group when faced with group identification cues, while proself individuals might be led to consider indirect rewards from being on a winning team with such cues. This effect was shown in another context in De Cremer and Van Vugt [4]. Therefore, group identification cues will be more important to proself individuals than prosocial ones when knowledge sharing is to be maximized.

For prosocial individuals, who already are expected to contribute, managerial prompting and group identification cues are expected to be largely unnecessary, even ineffective. It is therefore expected that the differences in knowledge sharing will be even greater when proself individuals will be given group identification and managerial reminders at the same time.

**Method.** A laboratory experiment was conducted utilizing computers at two large U.S. universities to test the previously described predictions. Seventy-six undergraduate business students participated in the experiment in

exchange for extra credit toward their course grades. Students were seated in a computer lab and told the number of points to be awarded would be determined by their performance, composed of their own reward plus a share of their group's reward. The sessions were carried out in class under the researchers' supervision.

The study was set up as an investment game in which participants were to make a series of profit-maximizing decisions. The decisions focused on subjects' need to balance contributions of information for the common good against retaining that information for their own good, as we will describe in more detail here. Although

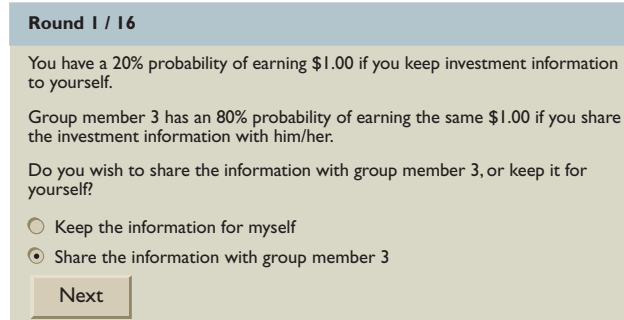
individuals were told they formed groups of six members each, all decisions were made individually with no possible interaction with other participants.

The total profit of a player consisted of the player's own gains plus a bonus awarded if the player's team achieved a predetermined level of contributions (40% of the total). Whether to keep information for oneself for probable direct rewards or to contribute it to the team for probable indirect rewards was the dilemma at the heart of the investment decision. On one hand, players were enticed to keep information to themselves and increase their individual gains; on the other hand, they were drawn toward contributing the information to the team and securing the group bonus.

The relationship between profit and information sharing was implemented by a series of decision-making situations (rounds) offered to every player. In each round we displayed investment information and two probabilities the player was facing: a probability of individual gain should the player keep the information and a probability of a team gain should the player decide to share the information with another team member.

After two practice decisions, the players were presented with 16 consecutive information-sharing decisions, illustrated in the figure here. The rounds were identical except for the varying probabilities assigned to individual and team profit successes. The complete table of probabilities is shown in Table 1. Upon completion of the decision rounds, subjects were asked to provide demographic and other information.

We simulated *Managerial Prompts* by delivering three strategically timed pop-up messages while the task was under way for subjects in the "prompted" treatment. The first message appeared before round 1 and stated that the system would monitor the overall number of contributions to the team by its members, and



An experimental round.

should these contributions be insufficient to reach the bonus, each member would be notified. Regardless of the individual subject's level of sharing, the second message appeared after round 6 and warned the player that contributions to the team were too low and needed to be increased in order to obtain the bonus. The third message appeared after round 12 and displayed the same warning as the second message.

The manipulation of *Group Identification* was achieved by the indication to half of the subjects by the researcher that the game was being played in a competition with a named rival university. It was hoped that subjects who received such indication would identify with their group of classmates as an "us against them" situation. A similar manipulation was used successfully in Van Vugt and De Cremer [12].

*Social Value Orientation* reflected an individual preference and therefore was measured as part of the demographic variables. The Decomposed Games instrument adapted from Messick and McClintock [10] was used to classify the respondents into proself (individualistic and competitive) and prosocial (cooperative) categories.

## RESULTS

It appears that managerial prompts indeed resulted in more information sharing. Both the total sharing and three individual segments of sharing opportunities were examined, because the prompts were offered twice. Interestingly, the first segment did not stimulate differences, which would indicate that subjects did not pay much attention to the first managerial prompt. However, the other segments and the total did elicit significant differences. Overall, subjects who were provided with managerial prompts shared on average 7.34 times while subjects without such prompts shared on average 5.93 times (out of 16).

Group identification subjects (those exposed to a message about beating a rival school) did not seem to share any more than non-group identification subjects (6.84 vs. 6.32, respectively). While the means are in the opposite direction, they are not significantly different, and no conclusions can be drawn about group identification as a main effect.

Round Number	Payoff probability to player (if information not shared)	Payoff probability to team (if information shared)
1	0.2	0.8
2	0.8	0.2
3	0.2	0.6
4	0.6	0.2
5	0.4	0.8
6	0.8	0.4
7	0.2	0.4
8	0.4	0.2
9	0.4	0.6
10	0.6	0.4
11	0.6	0.8
12	0.8	0.6
13	0.3	0.5
14	0.5	0.3
15	0.5	0.7
16	0.7	0.5

Table 1. Payoff probabilities.

Unexpectedly, 47 of the 76 subjects responded as proself individuals, while only 10 responded as prosocial individuals. The other 19 were indeterminate. Consistent with our expectations, the small number of prosocial subjects did contribute significantly more information (8.10 times) than the proself majority (6.26 times). When turning to interactions among the factors, we predicted that proself individuals would

be affected more by managerial prompts than would prosocial individuals. Results were consistent with that expectation: the sharing scores for prosocial individuals receiving managerial cues were not significantly different (8.4 with prompts vs. 7.8 without prompts;  $p=0.362$ ), and in contrast, the sharing scores were significantly different for proself individuals receiving managerial prompts (7.15 for managerial prompt subjects and 5.59 for subjects without prompts).

We asserted that proself individuals would be more sensitive to group identity cues than prosocial individuals. We were unable to test this assertion because of distribution of subjects in the sample. There were only two prosocial subjects in the no-group-identity cell.

Finally, a three-way interaction among the three experimental factors was tested. The small number of prosocial subjects in the non-group identification cell leads us to interpret any results from this test with caution. However, even with the sample difficulties, the three-way interaction was indeed significant.  $R^2$  was 0.301 and adjusted  $R^2$  was 0.210, indicating that the ANOVA explained a substantial amount of variance for the small sample; the means are shown in Table 2.

## CONCLUSION

This experimental study explored the direct and interaction effects of managerial prompts, group identification, and social value orientation on knowledge sharing. The results indicate that knowledge sharing in organizations can be encouraged by management's reminders of the importance of the goal, as well as reminders about rivals. It is also important for organizations to hire personnel with prosocial traits if they want to depend on them to share. This builds on previous literature that identifies important organizational cultural variables as important elements in knowledge sharing [3].

Organizational culture might also interact with national culture in potentially interesting ways, and additional research might investigate if the phenomena are related. For instance, a future study might find that in societies that are more collectivist in orientation (such as several countries in Asia), different tactics could be used to foster prosocial attitudes than those in more individualistic societies (such as the U.S.) [6].

Although group identification in this sample did not seem to make a difference, the results should be taken only as a directive for closer study. Many organizations use tactics to foster such an image, such as posting mottos on a wall, having retreats and annual meetings with motivational speakers, or communicating regularly the importance of the organization's strong standing in the market. Such tactics to make employees desire to beat the competition could provide productivity benefits, while employees seek improved job security, larger bonuses, or even satisfaction that comes with success. Group identification, therefore, remains an intuitive and promising goal in spite of the nonsignificant findings in this study.

This research also has some limitations. The research method was experimental. The method was selected in order to adequately manipulate the variables and environment. However, this method does sometimes suffer from the ability to generalize results across a population when there are multiple other factors that can also influence behavior. Students' expectations of reward might have played a role in their behavior during the study, and incentives in a real situation need to be examined to see if they are consistent with the behavior in the study.

While managerial prompting might indeed enhance contributions in actual organizational situations as it did in this study, its effectiveness in an employment setting might need further study. This experiment was meant to discover if prompting had any promise, and future studies must be conducted to determine its efficacy over time. In a situation where a pop-up screen is shown multiple times per hour, such a tactic might be completely ignored or ineffective after several days. Personal prompting might be most useful in a variety of situations, but there could very well be systems that intermittently monitor performance and provide occasional advice through prompts that are used sparingly. Users would quickly learn the system is keeping track of their individual progress in a formal way and might adjust their behavior

Treatment	No Prompts		Prompts	
	Mean	Standard Deviation	Mean	Standard Deviation
<b>No Group</b>				
Proself	<b>5.61</b>	3.15	<b>7.75</b>	2.73
Prosocial	<b>16.00</b>	.	<b>10.00</b>	.
<b>Group</b>				
Proself	<b>5.56</b>	1.42	<b>6.25</b>	2.96
Prosocial	<b>5.75</b>	2.06	<b>8.00</b>	1.41

Table 2. Total sharing.

according to appropriate incentive systems.

Future research is needed to discover other important antecedents of knowledge sharing, both separately and in combination with those explored here. The full promise of KMS will follow a firmer understanding of how to foster the sharing of knowledge by organizational members. If individuals learn how to overcome natural tendencies to retain personal competitive advantage based on proprietary information, we might realize such large organizational gains that the individuals will be even better off than they would be if everyone jealously guarded their valuable knowledge. **C**

## REFERENCES

1. Birnberg, J. Some reflections on the evolution of organizational control. *Behavioral Research in Accounting* 10 (1998), 27–46.
2. Bock, G., Zmud, R., and Kim, Y. Behavioral intention formation in knowledge sharing: Examining the roles of extrinsic motivators, social psychological forces, and organizational climate. *MIS Quarterly* 29, 1 (2005), 87–111.
3. Davenport, T., DeLong, D., and Beers, M. Successful knowledge management projects. *Sloan Management Review* (1998), 43–57.
4. De Cremer, D. and Van Vugt, M. Social identification effects in social dilemmas: A transformation of motives. *European Journal of Social Psychology* 29 (1999), 871–893.
5. Goodman, P. and Darr, E. Computer-aided systems and communities: Mechanisms for organizational learning in distributed environments. *MIS Quarterly* 22, 4 (1999), 417–440.
6. Hofstede, G. *Culture's Consequences: International Differences in Work-Related Values*. Sage Publications, Beverly Hills, CA, 1980.
7. King, W., Marks, P., and McCoy, S. The most important issues in knowledge management. *Commun. ACM* 45, 9 (Sept. 2002), 93–97.
8. Kramer, R.M. and Goldman, L. Helping the group or helping yourself? Social motives and group identity in resource dilemmas. In D.A. Schroder, Ed., *Social Dilemmas: Perspectives on Individuals and Groups*. Praeger, Westport, CT, 1995, 49–67.
9. Marwell, G. and Oliver, P. *The Critical Mass in Collective Action: A Micro-Social Theory*. Cambridge University Press, New York, 1993.
10. Messick, D. and McClintock, C. Motivational basis of choice in experimental games. *Journal of Experimental and Social Psychology* 4 (1968), 1–25.
11. Thorn, B. and Connolly, T. Discretionary data bases: A theory and some experimental findings. *Communication Research* 14, 5 (1987), 512–528.
12. Van Vugt, M. and De Cremer, D. Leadership in social dilemmas: The effects of group identification on collective actions to provide public goods. *Journal of Personality and Social Psychology* 76, 4 (1999), 587–599.

**PETER MARKS** (pete-marks@us.army.mil) is the Chief Information Officer at Brooke Army Medical Center in San Antonio, TX.

**PETER POLAK** (ppolak@miami.edu) is an assistant professor at the School of Business Administration, University of Miami, Coral Gables, FL.

**SCOTT MCCOY** (scott.mccoy@mason.wm.edu) is an assistant professor in the Mason School of Business at the College of William and Mary, Williamsburg, VA.

**DENNIS GALLETTA** (galletta@katz.pitt.edu) is AIS President, an AIS Fellow, and a professor of business administration at the Katz Graduate School of Business, University of Pittsburgh, Pittsburgh, PA.

© 2008 ACM 0001-0782/08/0200 \$5.00

DOI: 10.1145/1314215.1314226

By Alfred Loo

# THE MYTHS AND TRUTHS OF WIRELESS SECURITY

*Technology will never cure all wireless security ills.  
It will take a coordinated effort involving corporations, manufacturers,  
employers, and end users to fight the fight.*

The recent security breach at a U.S. credit-card processing center put over 40 million accounts at risk of fraud [11]. This event was quite shocking to many of us—the day-to-day users of credit cards. We expect a financial institution of this stature to be aware of security threats and their serious consequences as it processes very sensitive and important data. An organization of such a size should be able to employ a team of security experts. Indeed, researchers in IT security are very active [5, 9], and research papers in this area are abundant. We tend to believe that security experts in consumer finance

---



# THE STRENGTH OF A COMPUTER SYSTEM'S SECURITY IS ALWAYS MEASURED BY ITS WEAKEST COMPONENT.

In most systems, the weakest components are the end users, particularly when they are accessing the corporation's databases with wireless facilities at home.

organizations have adopted all possible measures and such breaches are highly unlikely. However, this event once again proved Murphy's Law—whatever can go wrong, will go wrong.

Organizations with fewer resources are even more vulnerable than the above-mentioned processing center. The incident is a wake-up call for every company to review its security measures. The strength of a computer system's security is always measured by its weakest component. In most systems, the weakest components are the end users, particularly when they are accessing the corporation's databases with wireless facilities at home. The reasons are simple.

- Wireless facilities for computers are fairly new for end users and there is a great deal of misunderstanding over their security features.
- Although computer facilities within the organization are protected by computer experts, connections from employees' homes are not.
- Research papers and books on wireless security are written for technical people. Most users, including some computer professionals, are not able to understand these materials.

Wireless routers have been available for several years, are extremely easy to install, and are inexpensive—most cost less than \$100. A single router provides users with the following conveniences:

- It is quite common for families to have several computers, usually situated in different rooms. To share broadband communication, the installation of network cables would be a time-consuming and expensive task, depending on the distances between rooms. Wireless routers solve this problem easily.
- Users can move their computers easily around the home, as long as they remain within the coverage of the wireless signal. The computer is still able to access the Internet via the wireless facility and save hours of cable reconnection.

## THE MYTHS

A seasoned application programmer with a master's degree who uses a wireless router at home recently expressed his belief that his connection was secure. It is not. If an application programmer puts his faith in myths, we cannot expect anyone with less computer knowledge to be aware of the security problems.

In theory, nothing can be 100% safe. Indeed, even if an owner tries to protect his/her house from burglary by installing good locks, determined thieves can still break in provided they have the right equipment and enough time. If they cannot break the lock, they can break the door, or a window, or drill a hole in the wall, dig a tunnel, use explosives, and so on. However, all of these methods will take more time and make more noise. Carrying more equipment will increase the chance the thieves will be detected or even arrested before or after the task. Rational thieves will compare the costs of these methods with the estimated valuables in the house. They might decide it is not worthwhile to break into the house provided the owner has the right degree of security.

Wireless communication is very similar to this scenario. However, while many users believe their "doors," at least, are locked, the truth of wireless communication is the doors are not even closed and are actually wide open [3, 10]. Intruders can simply walk in. A wireless router user might believe the following:

- Wireless routers have been in the consumer market for several years. Many people are using them so they must be safe. There are enough security features in routers, so communications through them are automatically protected.
- It is time-consuming to hack a computer, so attackers go after corporate computers. There is nothing valuable enough in home communications to attract hackers. Users have little to lose. There will not be any serious consequences even if their transmissions are intercepted by hackers.

## THE TRUTHS

Most wireless routers have some security features.

However, these features are optional and usually turned off. The installation of a router is quite easy, and the user can complete it within minutes. During the initialization process, a small installation program will ask them to type in the account name and password they use for their ISP. Users might believe they are protected by the password afterward; indeed, the account name and password are stored in the router. When the user turns on their computer, the router logs into the ISP automatically. The computer then communicates with the router directly and accesses the Internet through it. The router does not ask the computer to supply the account name and password again. As a result a hacker's computer can connect to the router without any knowledge of the account.

Hackers do not require specific hacking tools, as operating systems can find nearby routers and connect to them almost automatically. Hackers can then perform illegal operations such as hacking other computers, spreading viruses, organizing terrorist activities, and so on. They do not need to get into users' houses or plug in any cables. They might be sitting inside a car [1, 2, 6] parked on the road near a user's house. A hacker might also live next door with nothing more than a notebook computer and wireless LAN card. It is extremely difficult for users to detect this hacking process, as hackers do not need account names, passwords, IP addresses, or any identification to the ISP. It is impossible to trace them afterward.

Hacking can take place automatically if proper software is installed on the computer. For example, a hacker could leave a computer on in a van while masquerading as technicians from telephone or electricity companies. They do not need to stay in the van during the hacking processes. Antennas and amplifiers can extend the area of wireless coverage. A study [7] shows it is possible to attack a target from 20 miles away. This kind of hacking is more difficult to detect.

As the router is actually sending/receiving messages for the computers, it is impossible for the ISP to detect the number of users. This will create headaches for the users and governments. When the police come to the doors of the legitimate users, they will have a difficult time explaining what is going on. Governments will also have problems in identifying hackers from ISP records.

There is another illusion in the initiation process. Many routers have a default administrator account "admin" (with the same default password "admin"). Users can change the password or create a new account. They might believe this is another level of protection. Actually, this account is used only when users want to view or revise the settings in the router.

Hackers do not need this piece of information to access the Internet through the router. This illusion only gives users a false sense of security.

As the security features are usually turned off, communications between computers and routers are transmitted in plain text. Messages can be intercepted easily. The hacker only needs an ordinary computer with a wireless LAN card to intercept the signal.

### THE CONSEQUENCES

Many companies have developed applications that use Internet technologies as they offer many advantages. Employees can access the databases of the company with an Internet browser. Naturally, employees can also access the systems with their home computers so they can continue to work at home.

Such advantages are accompanied by additional risks. Many systems rely on only account names and passwords to authenticate users. If a hacker can obtain this information while a user is accessing corporate databases, every security feature installed by the corporation will fail and the unauthorized intrusion will go undetected.

Some corporations protect their log-in pages with SSL or similar techniques. Such traditional methods for wired communication can also be effective for wireless networks. The transmissions are encrypted if the corporation is using such techniques. Users might believe that further encryption is not necessary. However, they will also access Web sites that do not require high levels of security, and a patient hacker can obtain hints from these non-secured communications. For example, a user might use their favorite car model as a password. This word could be found in their communications as they frequently search for information about this model from the Internet.

It is common for users to have multiple accounts. Many Web sites providing free information also ask their readers to create new accounts. As these Web pages do not involve any money or confidential information, encryption techniques might not be used.

In theory, users should use different passwords (and account names) for different accounts. In practice, many people use the same password for multiple purposes as it is extremely difficult to remember all accounts and passwords. If account names and passwords are intercepted in unsecured transmissions, hackers can spoof users and break into company databases.

Even if users do use different passwords, it is still dangerous not to protect passwords for non-secured

## SECURITY WILL NEVER BE PERFECT as hackers can always find new methods to crack systems.

Web sites. By analyzing user communications, hackers can detect patterns in the composition of passwords. It is difficult for people to remember a combination of characters and numbers that have no meaning. A common practice is to take words and numbers that are familiar and combine them.

For instance, a user might choose the first four characters of his son's name (for example, "davi" from "David"), combined with the year of his son's birthday (say, 1992). To make it a little more difficult to crack, the user reverses the order so the final password is 2991ivad. This password is not bad because it is easy to remember and it does not use a complete word from a dictionary. However, an assiduous hacker will find out the pattern if the password can be intercepted. The hacker can then try the user's corporate account with his daughter's name and year of birth (by adding or subtracting a small number from the 1992 in the example). If this does not succeed, the hacker can try the name of the user's wife, or something similar. Hackers can obtain these pieces of personal information by monitoring the communications of the user and his family members. In other words, such information could help to minimize the number of trials needed to break into a computer system.

By eavesdropping on user communications, hackers can find out the users' favorite or frequently used words and Web sites. This kind of information can also speed up the hacking process. Thus, it is important to encrypt all communications even when there is no confidential information in Web pages.

Finding the executives of a target corporation is also easy. The names of executives can be found on their company Web sites. Sometimes there are even photos and brief biographies on the same page. Hackers can follow the executives' cars and find their residences easily. Armed with this knowledge, hackers are ready to carry out their jobs.

### CORPORATE RESPONSIBILITIES

Parents do not leave their babies in the jungle and

expect them to protect themselves. Nevertheless, many organizations do not take care of the home Internet connections of their employees. The widespread use of wireless routers by employees is a threat to many computer systems.

There are many ways to attack wireless communication, so it is almost impossible to make it absolutely safe. There are also basic weaknesses in the technologies that are currently being used in most routers (for example, WEP [4, 5, 8]). However, we should force the hacker to invest more time and install more hardware/software for each attack. This will increase the hacker's risk of being caught. Corporations should urgently review their security, with the following measures easy to implement.

- Corporations should educate their employees to the risks and appropriate countermeasures. From time to time, they could offer brief seminars with materials explained in laymen's terms.
- Corporations should remind users to be vigilant and report suspected intruders to the authorities as soon as possible.
- Users should be taught to turn on the easy and common security features of their routers. These features can be deployed in a few minutes provided that the router manufacturer builds software with a good human-computer interface. It is not possible to discuss all features in this article, but the following examples should suffice.
  - > Every LAN card has a MAC address. Users can store the MAC address in the router and ask it only to accept computers with the right MAC address.
  - > Users should turn on the encryption feature in the router.
  - > Wireless routers broadcast their Service Set Identifier (SSID) to surrounding computers. Computers need this name for proper connection to the router. The user should change this name often and disable router broadcasting. This measure will mean that hacking is more time-consuming.



- > Every router comes with an administrator account that has a default password. Some users do not bother to change this password. Once hackers log into routers, they can modify the settings and/or turn off all security mechanisms. Users should change the password (and account name) to protect the settings.
- Most router installations are easy. However, some routers come with poor human-computer interfaces in the parts that activate the security features. It is difficult for users to select a good router on their own. Corporations should evaluate routers and recommend only products with good security features and human-computer interfaces. If enough corporations are offering these kinds of recommendations, it will force manufacturers to pay more attention to developing these areas.
- Digital certificates and VPN should be adopted as they provide higher levels of security in the long run.

#### MANUFACTURER RESPONSIBILITIES

There are flaws [5] in the current security technologies of wireless routers. Security experts are proposing new standards [4, 5] to overcome these problems, but manufacturers can improve security in the meantime. Security will never be perfect as hackers can always find new methods to crack systems. Nevertheless, the following measures could be implemented quickly and easily.

- Router manufacturers should make their products safer by turning on security features as default settings. Users should be forced to turn them off if they do not need the facilities, and they should be warned of the risks.
- Manufacturers should incorporate small resident programs that allow the checking of the number of users at a particular time. Log files should be maintained so that users can check whether there are any intruders. These programs could be executed in the user's computer or the router (if the router has enough memory).
- Once a new hacking method is discovered, it should be announced on the manufacturer's Web site or through email. Updated software and/or firmware (if any) that can combat the new attack should be available for downloading from the manufacturer's Web site.

#### USER RESPONSIBILITIES

- Users must follow the security procedures of their employers. It is their duty to safeguard the system

by using computers properly.

- Users should invest their time in learning how to use wireless communication safely, such as by attending seminars offered by their employers.
- Users should be vigilant at all times. For example, many computer systems display the last log-in time when users log in. These records should be checked to detect intruders.

#### CONCLUSION

Technology alone will never be able to solve all security problems. Enhancement of the coordination between employers, end users, and wireless facilities manufacturers is constantly required. Users should understand it is their obligation to protect their employers' computer systems by understanding the risks and appropriate countermeasures, and that it would be worthwhile investing their time in updating their knowledge. Employers and manufacturers should make this process as easy as possible. Indeed, common sense, constant vigilance, and up-to-date knowledge are the best weapons in the fight against hacking. **□**

#### REFERENCES

1. Associate Press. Florida man charged with stealing Wi-Fi: Practice is common, but arrests are unusual. MSN, July 6, 2005; [www.msnbc.msn.com/id/8489534](http://www.msnbc.msn.com/id/8489534).
2. Bangeman, G. Illinois WiFi freeloader fined US\$250. *Arc Technica*, Mar. 2006; <http://arstechnica.com/news.ars/post/20060323-6447.html>.
3. Berghel, H. Wireless infidelity I: War driving. *Commun. ACM* 47, 9 (Sept. 2004), 21–26.
4. Chen, J., Jiang, M. and Liu, Y. Wireless LAN security and IEEE802.11i. *IEEE Wireless Commun.* 12, 1 (2005), 27–36.
5. Edney, J. and Arbaugh, W. *Real 802.11 Security*. Addison Wesley, Reading, PA, 2004.
6. Leary, A. Wi-Fi cloaks a new breed of intruder. *St. Petersburg Times*, July 4, 2005.
7. Maxim, M. and Pllio, D. *Wireless Security*. McGraw Hill, 2002.
8. McCullough, J. *Caution! Wireless Networking*. Wiley, 2004.
9. Pfleeger, C. *Security in Computing*. Prentice Hall, Englewood, NJ, 2006.
10. Reuters News. *Wireless Networks Easy to Hack*. Aug. 8, 2005; [www.itweb.co.za/sections/internet/2005/0508081002.asp?S=Reuters&A=REU&O=FPW](http://www.itweb.co.za/sections/internet/2005/0508081002.asp?S=Reuters&A=REU&O=FPW).
11. Schneier, B. CardSystems exposes 40 million identities. *Schneier on Security* (June 2005); [www.schneier.com/blog/archives/2005/06/cardsystems\\_exp.html](http://www.schneier.com/blog/archives/2005/06/cardsystems_exp.html).

---

**ALFRED LOO** ([alfred@ln.edu.hk](mailto:alfred@ln.edu.hk)) is an associate professor in the Department of Computing and Decision Sciences at Lingnan University, Hong Kong.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

---

© 2008 ACM 0001-0782/08/0200 \$5.00

---

DOI: 10.1145/1314215.1314227

# MUNICIPAL BROADBAND WIRELESS NETWORKS

*“...people lack many things: jobs, shelter, food, health care and drinkable water. Today, being cut off from basic telecommunications services is a hardship almost as acute as these other deprivations, and may indeed reduce the chances of finding remedies to them.”*

UN Secretary General, Kofi Annan,  
in a keynote address to the International Telecommunication Union,  
Oct. 9, 1999.

**I**n the global information society, the availability of Internet access has become critical for the world's citizens and organizations. Nevertheless, certain geographical areas and populations lag behind others in terms of Internet access, and a number of communities and governments around the world are actively seeking solutions to speed up the process of universal broadband Internet availability [4]. Municipal wireless networks (MWNs)<sup>1</sup> represent a new option for providing broadband access to the Internet.

MWNs are needed because private sector Internet service providers tend to focus their services toward more financially attractive markets and consumers,

---

<sup>1</sup>A MWN is defined as a wireless Internet access network that is created with active local leadership and involvement. MWNs are typically based on a point to multipoint link between base station and subscriber equipment. A base station is an outdoor antenna connected to a wired Internet backbone that sends data wirelessly to subscriber equipment such as laptops. Typical usage of a MWN involves using the built-in wireless card included in most laptops to connect to the Internet.

*By Munir Mandviwalla, Abhijit Jain, Julie Fesenmaier,  
Jeff Smith, Paul Weinberg, and Greg Meyers*

---

and thus tend to neglect certain constituencies [10]. According to economic theory, it is appropriate for government to intervene in the provision of a service if the service utilizes public goods and has the potential to result in significant positive externalities [1]. Arguably, wireless Internet service fulfills these requirements because it utilizes the publicly owned radio spectrum, and can increase the potential for new innovations and societal benefits as more citizens get connected.

MWNs are also more feasible today because recent developments in standards and technology have made broadband wireless more convenient, effective, and a viable alternative to cable, DSL, and other wired technologies [11]. There is also an expectancy that municipalities will find it relatively easy to set up wireless networks because they have ready and free access to physically dispersed assets such as lamp posts and traffic lights that are ideal for installing base stations [11]. Declines in prices have made the necessary network infrastructure increasingly affordable. In the U.S. estimates for deployment range from \$40,000 to \$74,000 per square mile [2], this is a small fraction of the costs of other public undertakings such as road maintenance and law enforcement.

MWNs also provide visionaries with a means to

realize dramatic changes in society. New applications based on ubiquitous and affordable broadband Internet access can have a significant positive impact on the socioeconomic health and status of entire regions. Concepts such as e-government and e-healthcare could set in motion radical ideas such as changes in the political process, delivery of health care, or more specific changes such as control of automobile traffic patterns through wireless mesh networks [9]. Proponents are also driven by the idea of digital convergence where traditional telephony and broadcast entertainment systems are superseded by a wireless Internet pipeline that delivers all network and broadcast services at a lower cost to the consumer [8].

It is therefore not surprising that since 2003, more than 200 communities around the world are actively seeking to build MWNs. These range from small communities such as Grand Haven, MI to large metropolitan areas such as Paris, France. In summary, MWNs are interesting because they represent a feasible alternative to Internet access, there is already a worldwide grassroots effort to create MWNs, and it is likely that new emergent applications will follow widespread deployment.

There is little research-based guidance available on the core issues and challenges that must be addressed

MWNs are interesting because **THEY REPRESENT A FEASIBLE ALTERNATIVE TO INTERNET ACCESS**, there is already a worldwide grassroots effort to create MWNs, and it is likely that new emergent applications will follow widespread deployment.

by a community during the process of implementing a MWN [7]. In this article, we offer recommendations based on our experience as participants in an expert group<sup>2</sup> that, under assignment from the city's Mayor's office, developed a proposal for a MWN for one of the largest cities in the U.S.—Philadelphia [6]. The recommendations are presented in the form of a stage-based implementation framework (see Figure 1). The framework can guide communities that seek to implement new MWN projects and assist researchers in identifying emerging issues. The investigation included a diverse array of projects and therefore our findings are likely relevant globally. Nevertheless, non-U.S. readers should also consider findings that are regionally proximate.<sup>3</sup>

Our investigation was based on:

- An analysis of the expectations and requirements of more than 110 stakeholders and future consumers of the proposed MWN,
- A comparative analysis of 25 MWN projects, primarily in the U.S.,
- A review of the related academic and professional literature, and
- An evaluation of pilot implementations specific to the Philadelphia MWN.

The process of MWN implementation can be understood in terms of three distinct stages. During

the first stage, a community and its local government need to simultaneously consider the needs of diverse stakeholders, identify manifold goals and address challenging policy issues. The second stage involves simultaneously thinking through potential applications, selecting the appropriate wireless technology, and planning the management and funding of the MWN. The last stage involves actual MWN implementation. Each stage will require a process of evaluation and reflection to rationalize potentially competing visions and imperatives. Here, we further elaborate each stage.

### STAGE 1: IDENTIFYING MWN GOALS

In this stage, a community identifies and sorts out competing goals, considers the needs of diverse stakeholders, and addresses key policy issues.

*Goals.* The 25 different MWN projects analyzed for this study address a range of goals.<sup>4</sup> These typically take the form of a primary goal and several accompanying secondary goals. Primarily, MWNs are envisioned as a solution to address the “digital divides” in Internet access. Because digital divides will not disappear simply from providing access to the Internet, MWN projects often also incorporate education and workforce retraining programs. For example, a social service agency in Philadelphia called the People's Emergency Center has set up wireless networks in homeless shelters, and provides training to enable residents to learn about welfare-to-work programs and acquire technology-related life skills.

Among secondary goals, local governments want to

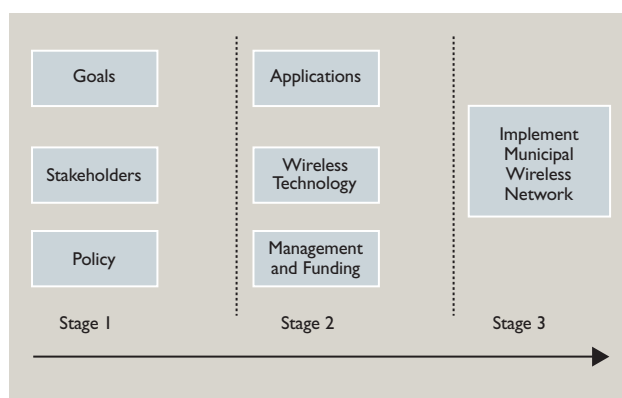


Figure 1. MWN framework.

<sup>2</sup>The expert group was formed in the spirit of a public-private partnership, and was composed of members from the government, the non-profit sector, the corporate sector, and the academic community.

<sup>3</sup>There is thriving MWN activity in Europe, for example, the Legible City Initiative in Bristol, U.K. and the Intelligent Bus Project in Paris, France [5].

<sup>4</sup>Due to space constraints, a detailed analysis and listing of all these projects has been omitted from this paper. The sidebar “Grand Rapids, Michigan” presents one illustrative example. See [6] for details on the Philadelphia project.

save on Internet access costs by eschewing Internet service providers (ISPs) in favor of the MWN. By ensuring everyone has Internet access, municipal governments aim to eventually improve service and save costs through electronic delivery of municipal services to citizens. MWNs are expected to boost the attractiveness of communities for prospective visitors because tourists and convention attendees can benefit from on-demand access to local information. A long-term objective is that citizens will be empowered to participate more fully in the political process by organizing, debating political issues, and acquiring information via the Internet. MWNs are also considered a tool for urban renewal, for if they deliver the expected economic and social benefits it should lead to an improvement in the economic health and social well-being of served communities.

In addition to the social objectives listed here, MWN projects can earmark quotas for local entrepreneurs to help with building the MWN or to provide services that utilize the MWN. In this way MWNs may substantially reduce the barriers to entry in such markets that are now dominated by a few large firms. For businesses, MWNs are a way to extend or differentiate their services without having to worry about building network infrastructure. For example, many hotels already provide wireless access for customers, and a MWN will allow them to take advantage of economies of scale and offer their customers a far wider range of services without having to build underlying infrastructure.

There is no clear evidence that the digital divide can be successfully addressed by MWNs. Further research is needed to assess whether the stated benefits can be realized and whether the secondary goals are sufficient to justify the investment.

*Stakeholders.* It is important to recognize that different stakeholders may have diverse and sometimes contradictory goals they wish to achieve via a MWN.

Stakeholder	Interest
State or City Government	Build the image of the area and address the needs of political constituents such as voters who live in low-income neighborhoods. Use the network to save on the cost of internet access for government employees.
Municipal Services	Reduce communication costs and deliver improved services such as real estate assessors using hand-held wireless computers to submit reports.
Underserved/Disadvantaged Individuals	Free or low-cost Internet access and related computing and training programs associated with a MWN may provide lifestyle, education, and job-related improvements.
Community Residents	Provide a low-cost alternative to ISPs. Ubiquitous access may improve lifestyles, enable grassroots projects for urban renewal, and enable electronic interaction between and among citizens and government. Residents are also concerned about the potential tax burden, and may have ideological reservations about the advisability of a municipality offering high-tech services when there might be problems with basic services such as trash collection.
Local Tourist Industry	Tourists and business visitors such as convention attendees would benefit from on-demand broadband access and may also use the network to access tourist data about the local community. A MWN could thus help communities to market themselves as a desirable destination.
Small and startup businesses	Provide a lower-cost alternative to for-fee services. In addition, entrepreneurs could market new services and applications based on the MWN. For example, given an extant wireless infrastructure it is much more feasible to implement ideas that depend on ubiquitous Internet access such as providing content on PDAs for museums or zoos.
Large and mid-sized corporations	Will likely use private networks because of security and reliability concerns but see a MWN as a tool to enable telecommuting. They may support MWN efforts to promote community goodwill and potential improvement in community human capital.
Telecoms and ISPs	A MWN is a potential threat to existing network service providers, though their expertise is likely to be needed for the construction of any MWN. Some firms are seeking legal protection against MWN.
Non-profit and community groups	Serve as a catalyst for delivery of services and information and can be used for training and work force development programs.
Utility / Transportation / Health care	Such organizations often have staff out in field locations and could design new applications, such as for reading utility meters or for providing access to patient records to nurses providing hospice care.
Higher Education	Colleges and universities often have their own wireless networks but could further improve student services or provide new curricula by taking advantage of low-cost broadband off-campus access.
Public Schools	Could realize significant cost savings by avoiding investment in telecommunications infrastructure. Schools could also start experimenting with parent-driven and/or at-home student learning initiatives.
Outside the municipality	Adjacent areas may jump on board to extend the "zone" and/or may feel competitive pressure to provide comparable services.

Table 1. MWN stakeholders.

As part of the Philadelphia project we organized 13 focus groups with a total of 110 participants that represented a diverse set of interests. Table 1 presents the results in a framework that categorizes and analyzes interests of different stakeholders. The framework is generic and should prove useful in identifying potential conflicts and contradictory goals in advance.

The stakeholders in Table 1 range from beneficiaries of a MWN to those needed to build, sustain, or fund the MWN. The MWN thus potentially offers an environment that can enable us to better understand how new forms of "grassroots" or "community-based" technology initiatives are formed and sustained. Implementers will need to carefully manage the communication with such diverse partners and incorporate their issues and needs in the initial planning.

*Policy.* There is a policy question about whether—and under what conditions—government entities should get involved in the creation of MWNs. Internet service providers have started mounting legal challenges. For example, in the state of Pennsylvania,

telecommunication providers successfully lobbied for the establishment of a law that forbids municipalities in most of the state from creating a MWN. However, after considerable debate, the city government of Philadelphia was able to receive an exemption from this law. So far, telecoms (such as the Baby Bells) and ISPs have successfully lobbied 14 states in the U.S. to pass laws prohibiting municipalities from building wireless networks because they are considered unfair competition [3]. The threat to entrenched interests goes beyond telecoms and ISPs because wireless Internet is seen as a disruptive technology that can displace traditional wired telephony, cellular service, and broadcast entertainment by delivering equivalent services at a lower cost to the consumer [8]. The telecoms and ISPs may have a point because municipalities in general have an advantage over them as they have ready and free access to public assets that can be used to mount network infrastructure equipment.

**A** major policy challenge for a municipality is that MWNs come with an assortment of legal implications, involving matters such as who bears the responsibility for data security, related liabilities, operational performance, and customer service management. Finally, since MWNs operate over the public radio spectrum, they may face governmental regulatory constraints and licensing requirements.

There remain fundamental unresolved policy questions such as:

1. Are MWNs an example of the government competing with private business?
2. Does free access to public facilities provide municipalities with an unfair advantage in competing with ISPs?
3. Should public funds be used to support MWNs when the stakeholders are often entities such as state or city governments and other community groups that do not have experience in providing Internet service?
4. Is there a need for new uniform legislation to handle the often-competing goals among MWN stakeholders? Clearly, implementers should first seek to understand and assess the

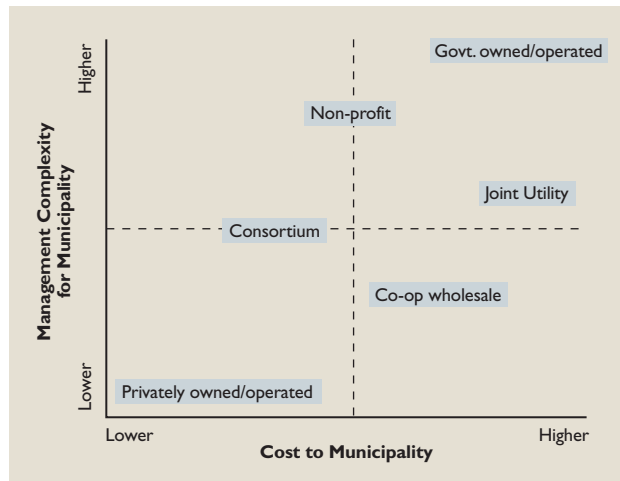


Figure 2. Comparing MWN management and funding models.

legislative landscape in their community before embarking on the creation of a MWN.

## STAGE 2: FLESHING OUT MWN INFRASTRUCTURE AND FUNDING

Once the goals, stakeholders, and policy issues surrounding a MWN are resolved in Stage 1, managers must identify the kinds of applications a MWN will support. Managers must also select appropriate wireless technologies, obtain sites for installing the receivers, and plan the management and funding. Moreover, ongoing evaluation and reflection is needed to rationalize competing visions, maintain public interest, and to stay abreast of new technologies.

*Applications.* The most common and compelling application of a MWN has been to create free or low-cost Internet access. There are also several secondary types of applications. For example, public safety institutions such as police, fire, and ambulance departments could be provided with a back-up communication network or with mission-critical information. Police in San Mateo, CA, for instance, can access law enforcement databases through MWN-enabled car laptops. Voice communications and value-added applications are especially relevant in underdeveloped areas where landline telephone service is poor, and in countries where mobile operators charge high rates for voice and data communications. For example, voice-over-Wi-Fi telephone service is now available in Rio Rancho, NM, through a MWN developer. In Houston, the city government is working on a system that will allow citizens to read, monitor, and update payments in parking meters via the Internet.

Given the many different applications of a MWN, researchers must assess how end users will adopt this new technology, what forces determine its usefulness, and what the MWN is ultimately most useful for. MWNs also offer visionaries a new “sandbox” to envision applications that are based on ubiquitous Internet access. Even if MWNs fail as Internet access mechanisms, the pilot projects may produce new innovative social service or entrepreneurial applications because this will be the first time in history that

Attributes / Model	Sample Key Issues	Government owned and operated	Government owned and privately operated	Joint power authority or public utility	Non-for-profit owned and operated	Cooperative wholesale model	Enterprise/Private Consortium owned and operated	Privately owned and operated
<b>Overall Description</b>		Municipality owns and operates the infrastructure	Municipality owns the infrastructure; subcontracts operation to a private party	Similar to a utility which operates on a for-profit mission	Organization takes advantage of tax regulations governing non-profits	Municipality owned, and capacity is sold at wholesale to private providers	Owned and operated by private enterprise, fees shared with government	Owned and managed by private enterprise
<b>Governance</b>								
<b>Ownership</b>	Who owns the infrastructure (antenna)	Municipality	Municipality	Utility	Not for profit organization	Municipality	Private organizations/ consortium	Privately owned
<b>Policies</b>	Who has access? What is appropriate usage?	Commission sets and implements policy	Commission sets policy that is implemented by private party	Regulatory guidelines from municipality, operational policies set by utility managers	Regulatory guidelines from municipality, operational policies set by managers	Commission sets and implements policy	Commission sets policy that is implemented by private parties	Commission sets policy that is implemented by private enterprise
<b>Management</b>	Who is responsible for connection problems?	Municipality	Private parties, with oversight from municipality	Utility managers with oversight from municipality	Organization managers with oversight from municipality	Municipality	Private parties with oversight from municipality	Private enterprise with no or minimal oversight from municipality
<b>Finances</b>								
<b>Startup funding</b>	Who will pay for the startup costs?	Grant, donation, government allocation, loan/bond	Grant, donation, government allocation, loan/bond, private funding	Grant, donation, government allocation, load/bond	Grant, donation	Grant, donation, government allocation, loan/bond	Grant, donation, government allocation, loan/bond, private funding	Loan, private and/or equity funding
<b>Continuing "Profit / Costs"</b>	How will the profits/deficits be handled?	Invested in municipality and upgrades	Divided between profit taking and investment in infrastructure	Invested in upgrades to the system	Invested in upgrades to the system.	Invested in upgrades to the system.	Profits shared between operators and upgrades	Operated no different from any private business
<b>Upgrade</b>	How will upgrades be implemented and funded?	Funded through selectively applied fees	Funded through profits	Funded through profits	Funded through profits	Funded through sale of excess capacity	Funded through profits	Funded through profits

**Table 2. MWN management and funding models.**

so many different people will have ubiquitous, untethered Internet access.

*Wireless technology.* Applications require supporting infrastructure, and there are several wireless technologies available for creating a MWN including Wi-Fi, Mesh, and WiMAX. The sidebar “Wireless Technology Tutorial” provides a glimpse of the available technologies. The development of standards is both pushing and in some ways hindering MWNs. The low-cost widespread availability of 802.11b demonstrated the potential and stimulated demand, but the slower spread of newer, more robust standards may stifle further growth. Standards have often been the catalyst of widespread adoption of

new Internet technologies; in the wireless network space, the real innovators and drivers may be home-grown “standards” being developed by the grassroots efforts of early adopters. Implementers in the short-term should keep a careful watch on both current and future standards to avoid premature obsolescence.

*Management and funding.* The rapid evolution of technology, and the fact that the management of such projects is not a core competency of municipalities, creates significant risk and challenges for the management and funding of MWNs. Complicating factors include dealing with arbitrary and inflexible budgetary, regulatory, social, ideological, and political constraints such as “the project must be completed at zero cost to the taxpayers” or “the project must

Given the many different applications of a MWN, researchers must assess how end users will adopt this new technology, what forces determine its usefulness, and **WHAT THE MWN IS ULTIMATELY MOST USEFUL FOR.** MWNs also offer visionaries a new “sandbox” to envision applications that are based on ubiquitous Internet access.

## WIRELESS TECHNOLOGY TUTORIAL

**Wi-Fi.** Includes products based on the IEEE 802.11 specifications. The 802.11b standard is widely available today in laptops and other wireless devices. Wi-Fi hotspots are spreading rapidly and Wi-Fi will continue to dominate the last few meters of delivery infrastructure.

**Mesh.** Consists of multiple interconnected nodes where at least one node is connected to the Internet. Data is transmitted across multiple short hops for high bandwidth and can be rerouted through alternate nodes if that node fails. The upcoming IEEE 802.11s (estimated 2008) will provide standardization for Mesh.

**WiMAX.** Worldwide Interoperability for Microwave Access products are based on the 802.16 standard. The Fixed WiMAX standard can provide very high throughput point-to-point connections wirelessly as far as 30 miles for several hundred users at different quality of service levels. The new mobile WiMAX (IEEE 802.16e) standard (estimated 2007) will likely address concerns of standardization while improving connectivity.

include computer training for the third district.” Moreover, there are very few successful projects available to emulate.

Table 2 presents a comparison of different management and funding models currently being considered in the projects analyzed for this study [4]. Figure 2 provides a map that compares the models based on cost and management complexity.

The key advantage of the *government owned and operated* model is that the municipality can directly ensure provision of services in digital divide areas and preserve other “social consciousness” aspects. However, the project could become a drain on tax funds and will not benefit from the discipline created by the private industry focus on the bottom line and equity markets. Municipalities are also typically not best practice leaders on project management and technology implementation, and wireless standards are evolving so rapidly that a poor or slow decision could cripple a project. Conversely, a completely municipality-led project will likely be simpler to manage and implement and may be ideally suited for small communities.

The *government owned and privately operated* model leverages the capabilities of current private providers while still retaining final control over direction and service. Internet service providers who are not made partners may be disadvantaged and may initiate both political and legal challenges against the idea of using public funds to sustain Internet services. Moreover, a single source (monopolistic) arrangement may not provide enough incentive to maintain customer service levels and upgrades. Clearly, the *private ownership and operation* model will have the discipline of the bottom line and will free the municipality from the burden of embarking on a risky new venture. The disadvantage is that the social consciousness aspect of the initiative may get lost and the project could be shuttered for purely financial reasons, and that may not be acceptable. The *joint power authority or public utility* model can address the funding challenges created from municipality ownership and can also provide the discipline of a profit motive. Such an organization may be well suited to handle the complex logistics associated with large infrastructure projects, but may, over time, become inefficient and resistant to change. A failure may still place the burden back on the municipality to bail out the initiative.

The strategies of private providers and municipalities continue to evolve. Deals where the government-side assumed little risk are disappearing and private providers are looking for greater balance in risk exposure. Current trends also point toward a splintering of municipal strategies. Whereas larger cities are more likely to outsource development, small communities are developing MWNs on their own. Several small communities have been successful in obtaining community development grants from the state and federal government to jumpstart MWN projects. Such projects typically include social improvement and digital divide goals, which are less attractive to private providers. Some cities are also encouraging projects that cater to different constituencies. For instance, in Marquette, MI, the city government supported a university effort to cater to off-campus students in the college town, and also supported a purely commercial project.

MWNs may eventually disappear and become absorbed into the larger telecommunications landscape. Before that happens though, **SUCH NETWORKS WILL ALTER OUR NOTION OF UNIVERSAL BROADBAND INTERNET ACCESS.**



Management researchers and policymakers face challenging questions: How should the MWN be best managed and organized given diverse goals, legal constraints, potential for innovation, and rapidly evolving technology? Which model will lead to social improvement, regional development, and/or profit?

### STAGE 3: IMPLEMENTING THE MWN

In the final implementation stage, the network is built and made available to subscribers and the results are evaluated against the goals. A “build it and they will come” strategy could fail if potential users choose to ignore a municipal-led effort, or lack the resources needed to take advantage of it. MWNs are outside the core competency of most municipalities. Can a municipal-led implementation deliver acceptable quality of service and customer service levels? To ensure success, a MWN implementation will require a parallel marketing and education campaign to counter widely held negative perceptions about the competency of municipalities. Municipalities should also consider outsourcing arrangements for handling management and operational aspects of MWNs.

### CONCLUSION

The convergence of need and availability has created the driving force behind the municipal wireless network. Social advocates see a low-cost, highly feasible solution that can have communitywide impact; technologists see the potential for new applications based on ubiquitous Internet access; civic leaders and politicians see an opportunity to dramatically impact the image of their constituent regions; and entrepreneurs see the potential for participating in the lucrative telecommunication industry. Parts of Philadelphia’s network are now operational, but the long-term impacts remain to be studied. For researchers, there are compelling questions related to assessing the role of government intervention, understanding adoption issues, and gauging economic and social impacts.

City	Grand Haven, Michigan				
Population	Approximately 12,000				
Physical Size	Approximately six square miles				
MWN built in year	2004				
MWN built by	Azulstar Networks (formerly named Ottawa Wireless)				
Objectives	Improve quality of life, augment safety, attract businesses, boost tourism, assist education				
Cost	Approximately \$40,000/sq. mile				
Business Model	Azulstar Networks has five-year non-exclusive license to install, operate and maintain network. Board of Light and Power granted permission to use its utility poles for installing radios. City gets five percent of revenues every three months.				
Pricing Model	Available service plans: the first four are for individuals, the last two are for businesses:				
	Monthly Plans	Base Monthly Price	Download Speed (bits/sec)	Upload Speed (bits/sec)	Daily Usage
	Free	Free	256k	60k	1 hour
	Entry	\$14.95	512k	200k	Unlimited
	Extreme	\$29.95	1.5M	300k	Unlimited
	Pro	\$49.95	3M	400k	Unlimited
	Sym 1.5M	\$249.95	1.5M	1.5M	Unlimited
	Sym 3M	\$479.95	3M	3M	Unlimited
Detailed Tech	A single fiber supplies Internet traffic to several hundred WiFi radios (802.11a, b, g) located around the city. Each of these repeaters contains a five-foot-high antenna, and has a range of two to three blocks. Customers mount an antenna outside their windows to pick up the signals.				
Speed	Average download rate of 1.5Mbps				
Network Reliability	Claimed to be greater than 99.7%—even under poor conditions (rain, fog, snow, high winds)				
Other Benefits	Point-to-point dedicated VPN connections, access available 20 miles into Lake Michigan (custom-built ship-to-shore marine Wi-Fi radio), video surveillance				
Sources	<a href="http://www.azulstar.com">www.azulstar.com</a> <a href="http://www.muniwireless.com/article/articleview/4263/1/23/">www.muniwireless.com/article/articleview/4263/1/23/</a> <a href="http://www.walkersands.com/Grand-Haven-First-Citywide-WiFi.htm">www.walkersands.com/Grand-Haven-First-Citywide-WiFi.htm</a>				

MWNs may eventually disappear and become absorbed into the larger telecommunications landscape. Before that happens though, such networks will alter our notion of universal broadband Internet access. We believe that every municipality worldwide should explore MWNs because many social benefits are attainable and the costs are relatively low and can be externally funded. Implementation of MWNs has significant potential to create new economic opportunities and spur technological innovation. However, there is a danger that municipalities will react to vocal interest groups and change their direction, create new “state monopolies,” and stifle new innovation to preserve investments. We recommend that implementers should identify the key stakeholder constituency (such as underserved populations), and create organizational structures to avoid a loss of focus, create organizational and financial structures that will help prevent monopolistic tendencies, and allow identification and adoption of new technological innovations (such as new higher bandwidth wireless standards) and, form creative partnerships with diverse entities to underwrite costs and bring in relevant expertise.

The stage-based framework presented in this article

represents an experienced-based first-step in understanding issues related to MWNs, provides guidelines for practitioners to act on, and presents issues for future research. ■

## REFERENCES

1. Cornes, R. and Sandler, R. *The Theory of Externalities, Public Goods, and Club Goods*, Cambridge University Press, Cambridge, 1986.
2. Earthlink Wins Phila. WiFi Contract. *The Atlanta Business Chronicle*, (Oct. 4, 2005).
3. Fisher, S. Legislative Barriers Against Municipal Wireless (WiMax): And You Thought This Was About Technology?! [www.informit.com](http://www.informit.com), (Oct. 21, 2005); [www.informit.com/articles/article.asp?p=418865&rl=1](http://www.informit.com/articles/article.asp?p=418865&rl=1).
4. Gillett, S.E., Lehr, W.H. and Osorio, C. Local government broadband initiatives. *Telecommunications Policy* 28, 7/8 (Aug./Sept. 2004), 537–558.
5. Hallett, T. Public sector rewarded for wireless innovation. *Silicon.com*, (Apr. 21, 2005); [networks.silicon.com/lans/0,39024663,39129740,00.htm](http://networks.silicon.com/lans/0,39024663,39129740,00.htm).
6. Jain, A., Mandviwalla, M., and Banker, R. Government as catalyst: Can it work again with wireless internet access? *Public Administration Review*, Nov./Dec., 2007.
7. Jain, A., and Mandviwalla, M. Preprototype perceived usefulness of a municipal wireless network: sources of variation among prospective users. In *Proceedings of the Twelfth Americas Conference on Information Systems (AMCIS-12)*, Aug. 4–6, 2006, Acapulco, Mexico.
8. Kessler, A. When brotherly love goes wireless. *The Wall Street Journal*, (Sept. 1, 2005).
9. Mackie, K. PacketHop debuts mobile mesh solution. *Broadband Wireless Online*, (Nov. 10, 2003).
10. Rao, B. and Parikh, M.A. Wireless broadband networks: The U.S. experience. *International Journal of Electronic Commerce* 8, 1 (2003), 37–53.
11. Sirbu, M.A., Lehr, W.H., and Gillett, S.E. Evolving wireless access

- technologies for municipal broadband. *Government Information Quarterly* 23, 3–4 (2006), 480–502.
12. Wireless Philadelphia Business Plan. The Wireless Philadelphia Executive Committee, Feb. 9, 2005.

---

**MUNIR MANDVIWALLA** ([mandviwa@temple.edu](mailto:mandviwa@temple.edu)) is an associate professor and Executive Director in the Fox School of Business, Temple University, Philadelphia, PA.

**ABHIJIT JAIN** ([ajain@nmu.edu](mailto:ajain@nmu.edu)) is an assistant professor in the Cisler College of Business, Northern Michigan University, Marquette, MI.

**JULIE FESENMAIER** is the Associate Research Director for the Cochran Research Center in Fox School of Business at Temple University, Philadelphia, PA.

**JEFFREY W. SMITH** is a consultant at Merck and Co, Inc, West Point, PA.

**PAUL WEINBERG** is Senior Fellow and Professor in the Fox School of Business, Temple University, Philadelphia, PA.

**GREGORY S. MEYERS** is Head of IT, Americas, for Novartis Consumer Health.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2008 ACM 0001-0782/08/0200 \$5.00

DOI: 10.1145/1314215.1314228



UC SANTA CRUZ

## University of California, Santa Cruz Computer Science Department

The UCSC Computer Science Department seeks qualified applicants for the following tenured/tenure-track faculty positions within the Baskin School of Engineering:

### Assistant or Associate Professor

Interests: Outstanding applicants with research excellence in computational aspects of videogame design such as artificial intelligence, real-time animation/graphics and HCI.

### Assistant Professor

Interests: Outstanding applicants with research excellence in computational digital media in the context of videogames, including game design, game studies and game art.

Position: 810-08

The department has strong graduate and undergraduate programs. Research and instruction are supported by excellent computing facilities and state-of-the-art laboratories in the new Engineering 2 building. UCSC is close to Silicon Valley and has strong ties with many of the high technology companies in the area. Faculty salaries are competitive and opportunities for consulting are extensive.

A detailed job description and application instructions are available at [www.soe.ucsc.edu/jobs/810-081pgflyer.pdf](http://www.soe.ucsc.edu/jobs/810-081pgflyer.pdf) For full consideration, materials should reference position number 810-08 and arrive by January 4, 2008 (if Digital Media) or February 20, 2008 (if Computer & Video Game Design). UCSC is an EEO/AA/IRCA Employer.

By Vicki R. McKinney, Darryl D. Wilson,  
Nita Brooks, Anne O'Leary-Kelly, and Bill Hardgrave

# Women and Men in the IT Profession

*Fewer women entering IT drives the  
underrepresentation problem.*

Fewer women work in or achieve the same job levels as men in corporate America today. In the IT profession, recent estimates (2006) suggest that women make up only 26% of IT professionals in the U.S. where they're outnumbered by men six to one in leadership positions [3, 6]. What explains these lower levels of professional achievement for women in IT? Unfortunately, few facts are available to inform the answer, though numerous suppositions have been offered, including that women have less innate ability or interest in the "hard" sciences, that their educational experience dissuades them from careers in IT, and that they are simply less com-

fortable working in what is, and has always been, a predominately male environment. However, the validity of such assumptions has not been tested extensively in the IT arena. What is needed is a large-scale study able to distinguish fact from supposition. Here, we report on such a study we conducted in 2003 that examined the attitudes and experience of 815 male and female IT professionals working in multiple organizations across the U.S. (see the sidebar "Study Respondents and Scales").

We examined several key questions regarding gender differences in IT: Do women share the same motivations or reasons for entering the IT profession

as men?; Are women as well socialized as men into the profession?; Do women have the same types of experience as men in the IT work force?; and Do men and women develop similar attitudes regarding the desirability of the profession? They helped pinpoint potential causes for the lower numbers of women in the IT profession. We used a Web-based survey to collect data from a variety of IT professionals from U.S. organizations.

### MOTIVATIONS, SOCIALIZATION, EXPERIENCE

A better understanding of why women are so underrepresented in the IT profession begins with the reasons for a person entering the profession in the first place. For example, common wisdom suggests that men may be more attracted than women to technology-related careers [6]. If true, then men seem more likely to weather inevitable career bumps because they feel a strong association with the work itself. To assess this, we asked our participating IT professionals about the influence a variety of motivators had on their decisions to go into IT, including: love of technology/computers; using state-of-the-art equipment; opportunity for task variety; opportunity for gratifying work; opportunity for advancement; opportunity for job autonomy; job security; level of income; ease of entry into the profession; flexible work hours; and professional prestige.

Our results suggest that men and women share some but not all motivations for entering IT (see Table 1). Both groups cited opportunity for job autonomy, advancement, task variety, professional prestige, income, using state-of-the-art equipment, and gratifying work. There were also notable differences; for example, men were significantly more likely than women to identify “love of technology/computers” as a key motivator. Women, on the other hand, more often indicated that “job security,” “ease of entry,” and “flexible work hours” were primary reasons for entering the profession. This pattern of results suggests that factors in the work itself are more

Item	Males	Females	Different
Opportunity for job autonomy	5.11 (1.34)	5.17 (1.42)	no
Opportunity for advancement	5.34 (1.36)	5.46 (1.16)	no
Opportunity for task variety	5.37 (1.32)	5.54 (1.26)	no
Prestige of the profession	4.08 (1.68)	4.05 (1.64)	no
Income	5.77 (1.51)	5.66 (1.17)	no
Using state-of-the-art technology	4.71 (1.55)	4.54 (1.56)	no
Opportunity for gratifying work	5.46 (1.24)	5.56 (1.33)	no
Love of technology/computers	5.40 (1.56)	4.86 (1.63)	yes (p<.001)
Job security	5.47 (1.38)	5.74 (1.25)	yes (p<.05)
Ease of entry into the profession	4.00 (1.71)	4.28 (1.54)	yes (p<.05)
Flexible working hours	4.43 (1.76)	4.96 (1.74)	yes (p<.001)

Notes:  
 • Seven-point Likert scales used, with 7=Important and 1=Unimportant;  
 • Cell value = mean (standard deviation)

Table 1. Motivation for entering the IT profession.

important in the career decision making of male IT professionals, while factors around the job (such as flexible hours) are more important in the decision making of female IT professionals.

Profession-related socialization refers to individuals’ degrees of learning about their chosen profession [2]. For example, in IT, professionals must acquire information about the technical and social aspects of their work in order to be effective. This information is often passed along via socialization agents (such as role models willing to share career and job-related

information). Women with less access to role models or lower levels of learning around the technical or social aspects of the profession may be less likely to remain or succeed in the profession, providing a plausible explanation for their underrepresentation. We asked our sample of IT professionals to report on their socialization.

One of the more common aspects of socialization is the degree of comfort an individual might share with other professionals. Socialization is often used to reference an individual’s interaction with others in social settings. Role modeling is another important aspect of socialization into a profession, reflecting the assistance received from other professionals in learning the norms of the profession. The nature of IT suggests that confidence in technical skills and comfort with the profession’s language are also important to an individual’s socialization in the profession.

Men and women in our survey both generally reported a similar level of experience with role models. Women, even in a predominately male work environment did not report a significant difference from men in the influence role models had on their careers in IT. This surprising finding does not support previous assumptions that the lack of females in IT means a lack of role models for women, which was assumed to be a disadvantage for women [1].

Similarly, men and women in the survey reported comparable levels of learning and comfort around the social aspects of the profession despite stereotypes that suggest women are drawn more to social interaction

[7]. However, our surveyed male IT professionals also reported stronger socialization with regard to the technical aspects of the profession, including familiarity with its language and confidence concerning their own skills (see Table 2).

Another plausible explanation for women's underrepresentation in IT could be that women have fewer positive experiences once in the profession; for example, female IT professionals may encounter more difficulty balancing work and family [4], possibly encouraging them to leave IT and pursue other careers.

Our findings uncovered only one significant gender difference across a variety of work-related experiences. Female and male IT professionals alike reported similar levels of experience regarding the work-family conflict, feelings of burnout, perceptions of work load, perceptions of fair treatment in job scheduling, assignment of job responsibilities, pay and other rewards, and perceptions of supervisor support related to family issues. They differed in regard to their perceptions of supervisor support related to their careers. This finding indicates that women perceive greater support in meeting career goals, recognizing opportunities, and improving their job performance.

The degree of "affective connection" can also influence how long a person stays in an occupation [9]. Women who feel less positive about the IT profession may be more likely to leave. A lack of attachment to the IT profession may explain why fewer women are in IT careers. To test this lack of attachment we examined three attitudes that reflect an individual's attachment to the IT profession: career satisfaction, professional identification, and intention to leave.

Career satisfaction reflects individuals' general contentment with their chosen careers. We would expect

Item	Males	Females	Different
<b>Socialization</b>			
Role Models	4.38 (1.30)	4.15 (1.36)	no
Comfort with social aspects of profession	5.27 (1.05)	5.20 (1.03)	no
Comfort with technical language	5.33 (1.21)	4.95 (1.20)	yes ( $p < .001$ )
Confidence in technical skills	5.19 (1.09)	4.73 (1.07)	yes ( $p < .001$ )
<b>Experience</b>			
Work/family conflict	3.39 (1.62)	3.40 (1.71)	no
Burnout	3.48 (1.60)	3.42 (1.64)	no
Perceived workload	3.99 (1.49)	3.82 (1.55)	no
Fair treatment	4.94 (1.26)	5.02 (1.16)	no
Supervisor support: family-related	5.07 (1.36)	5.24 (1.33)	no
Supervisor support: career-related	3.95 (1.46)	4.19 (1.51)	yes ( $p < .05$ )
<b>Attitudes</b>			
Career satisfaction	4.72 (1.26)	4.84 (1.29)	no
Professional identification	3.74 (1.30)	3.77 (1.16)	no
Turnover intention	3.24 (1.47)	3.20 (1.66)	no

Notes:  
 • Seven-point Likert scales used, with 7=Important and 1=Unimportant;  
 • Cell value = mean (standard deviation)

Table 2. Survey respondents' socialization, experience, and attitudes.

that individuals with high levels of career satisfaction would want to remain in their current professions. Professional identification reflects the degree to which people define themselves in terms of their chosen professions. Despite feelings of strong professional identification, leaving the IT profession would be tantamount to professionals leaving part of themselves behind. Moreover, intention to turnover indicates individual judgments about the likelihood that the individuals will leave the IT profession in the near future.

We found no significant gender differences for these measures of attitude. Male and female IT professionals in the study reported similar levels of satisfaction with their IT careers. They also reported similar (strong) levels of professional identification with the profession. Finally, and perhaps most important to the question at hand, we found no significant gender differences in intention to leave the profession.

## IMPLICATIONS

Our motivation for undertaking this research was to more clearly understand why women are underrepresented in IT. Two general reasons are possible: an "input problem," wherein women do not enter IT, and a "throughput problem," wherein women enter IT but then leave. Our study emphasized the latter. That is, we focused primarily on experience in and attachment to the profession, finding surprisingly little evidence of a difference between the genders.

How should these results be interpreted? While we found few differences between men and women in IT, we do not conclude that there are no gender differences in the treatment and/or experience of all IT workers. The study examined some (not all) work-related experience in a sample of IT professionals. It is possible that if we asked about other types of career experience (such as promotional opportunities) or included a different sample of IT workers, we might have found evidence of differential treatment or experience between male and female professionals.

## STUDY RESPONDENTS AND SCALES

We collected data from several large U.S. companies and targeted IT associations. The 815 respondents held a variety of IT positions, including computer programmer, computer analyst, computer scientist, and computer engineer, as well as others, including manager, administrator, project manager, and help-desk specialist. They came from a variety of industries, including health care (19%); IT services/software (25%); transportation (29%); government (10%); and other (17%). The majority of the sample was male (61%) and Caucasian (80%); African-American (11.5%); Hispanic (3.6%); and other (4.9%). Of the sample, 61% had an undergraduate degree and 15% had a master's degree.

We used previously tested and published instruments, except for motivation (still not available). Based on an extensive literature review, a motivation scale was created by the research team. A total of 56 people responded to our Web-based pilot survey used to test the instrument. We made minor changes to the instrument based on our analysis. The Cronbach's alphas for the main survey's scales were all acceptable at 0.71 or higher. **C**

Similarly, research examining experience at different career stages would inform and supplement our findings. However, the results we report here are based on a large sample of IT professionals, suggesting that the lack of gender differences in our findings should not be dismissed.

A possible interpretation of our results is the IT industry faces more of an "input problem" than a "throughput problem." While men and women do not differ in how they identify with the profession, there is evidence that women enter technology-related professions in lower numbers than men [5]. This is consistent with our findings here that more male IT professionals say they entered the profession because they love technology and computers, and more male IT workers report high levels of comfort with their own technology-related skills. These factors are specific to the IT profession. Female respondents indicated that they entered the profession because it provides benefits like job security, flexible work hours, and ease of entry, factors not specific to the IT profession.

Furthermore, the fact that we found few gender differences in the work experience among IT professionals implies it may be the lower numbers of women entering IT, not the larger number of women leaving IT, that drives the underrepresentation problem. Women in our sample were not significantly different from men in their identification with IT. This suggests that even greater attention is needed on "sup-

ply-side" issues (such as young girls' knowledge of computer careers, including computer science, early female identification with IT, and attention to learning styles [8] in computer education).

Meanwhile, the IT workplace appears to be a positive place for female professionals. This is at least the case for the organizations participating in our study, as indicated by the lack of significant difference in role modeling and career-related supervisor support. The fact that women in our sample reported comparable experience and attitudes as those of their male counterparts is good news for a profession that has been and remains male dominated. In our study, women and men were more alike than different in their experience in and attachment to the IT profession. **C**

## REFERENCES

1. Ahuja, M. Women in the information technology profession: A literature review, synthesis, and research agenda. *European Journal of Information Systems* 11, 1 (Mar. 2002), 20–34.
2. Chao, G., O'Leary-Kelly, A., Wolf, S., Klein, H., and Gardner, P. Organizational socialization: Its content and consequences. *Journal of Applied Psychology* 79, 5 (1994), 730–743.
3. D'Agostino, D. Where are all the women IT leaders? *eWeek* (Oct. 1, 2003); [www.eweek.com/article2/0,1895,1309599,00.asp](http://www.eweek.com/article2/0,1895,1309599,00.asp).
4. Igbaria, M. and Parasuraman, S. Status report on women and men in the IT workplace. *Information Systems Management* 14, 3 (1997), 44–53.
5. Information Technology Association of America. *Untapped Talent: Diversity, Competition, and America's High-Tech Future: Executive Summary* (June 21, 2005); [findarticles.com/p/articles/mi\\_m0EIN/is\\_2003\\_May\\_5/ai\\_101190621](http://findarticles.com/p/articles/mi_m0EIN/is_2003_May_5/ai_101190621).
6. National Center for Woman & Information Technology. *NCWIT Scorecard 2007: A Report on the Status of Women in Information Technology*; [www.ncwit.org/pdf/2007\\_Scorecard\\_Web.pdf](http://www.ncwit.org/pdf/2007_Scorecard_Web.pdf).
7. Rasmussen, B. and Hapnes, T. Excluding women from the technologies of the future? *Futures* (Dec. 1991), 1107–1119.
8. Sax, L. *Why Gender Matters*. Doubleday Publishing, New York, 2005.
9. Scott, C., Connaughton, S., Diaz-Saenz, H., Maguire, K., Ramirez, R., Richardson, B., Shaw, S., and Morgan, D. The impacts of communication and multiple identifications on intent to leave. *Management Communication Quarterly* 12, 3 (1999), 400–435.

---

**VICKI R. MCKINNEY** ([mckinney.vicki@gmail.com](mailto:mckinney.vicki@gmail.com)) is an organizational consultant.

**DARRYL D. WILSON** ([ddwilson@odu.edu](mailto:ddwilson@odu.edu)) is an associate professor in the Department of Information Technology and Decision Sciences at Old Dominion University, Norfolk, VA.

**NITA BROOKS** ([nita.brooks@quinnipiac.edu](mailto:nita.brooks@quinnipiac.edu)) is an assistant professor in the Information Systems Management Department in the School of Business at Quinnipiac University, Hamden, CT.

**ANNE O'LEARY-KELLY** ([aokelly@walton.uark.edu](mailto:aokelly@walton.uark.edu)) is a professor in the Management Department in the Sam M. Walton College of Business at the University of Arkansas, Fayetteville, AR.

**BILL HARDGRAVE** ([bhardgrave@walton.uark.edu](mailto:bhardgrave@walton.uark.edu)) is an associate professor in the Information Systems Department of the Sam M. Walton College of Business at the University of Arkansas, Fayetteville, AR.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

By Bing Liu and Alexander Tuzhilin

# MANAGING LARGE COLLECTIONS OF DATA MINING MODELS

*Data analysts and naive users alike in information-intensive organizations need automated ways to build, analyze, and maintain very large collections of data mining models.*

**Model building** is a key objective of data mining and data analysis applications. In the past, such applications required only a few models built by a single data analyst. As more and more data has been collected and real-world problems have become more complex, it has become increasingly difficult for that data analyst to build all the required models and manage them manually. Building a system to help data analysts construct and manage large collections of models is a pressing issue.

Consider a credit-card marketing application. The credit-card-issuing company wishes to build models describing the behavior of small segments of customers, or microsegments. Examples are middle-age customers with children in college living in zip

code 10012 and graduate engineering students at university XYZ. A large credit-card company might have to deal with tens of thousands of such microsegments, each involving dozens of different models. Therefore, it may need to build and support hundreds of thousands of models. Similar problems also occur in personalization applications and e-commerce.

The traditional approach is to aggregate the data into large segments, then use domain knowledge combined with “intelligent” model-building methods to produce a few good models. Intelligent means selecting the right functions and model types based on automatic algorithms and domain-expert knowledge. This approach reduces the number of models. However, it does not eliminate the need for

a large number of models in practice and is not ideal because some important characteristics of smaller models are lost in the aggregated models. The approach thus represents a compromise due to a human analyst's own limited intellectual capacity. With today's computing power, building a large number of models is not an issue. The bottleneck is the human analyst.

An initial approach to developing such tools [7] involved an application in Motorola, Inc. The system (called "Opportunity Map"), which has been in regular use since 2006, includes a set of customer-designed rule-manipulation and visualization operations. However, the approach lacks a formal foundation, and new features are constantly being added in an ad hoc manner. Model-management research should give the approach a theoretical foundation and extend it to other types of models rather than only to rules.

**Our goal** here is twofold. First, we want to make the argument that a generic model-management system is needed to facilitate large-scale model-building applications. And second, since model management is a relatively unexplored problem, we want to identify the main issues and problems and set up a research agenda, rather than present specific solutions, which is a separate research problem for the data mining community.

The idea of model management has been studied in the IS community since the mid-1970s [11]. However, the focus was on organizational and business issues and operations research and management-science-type models (such as mathematical programming, production and distribution, and transportation models) [5]. No industrial-strength systems were built, and little work was done on ways to manage statistical and data mining models. Interest in model management tailed off after an empirical study of 192 firms showed these firms had used about 200 models each on average in 1998 [12]. The database community also studied model management [1], but models there referred to schemas and metadata of database systems, which are quite different from data mining models.

Several data mining researchers proposed data mining query languages (such as DMQL [2] and M-SQL [10]). The main idea was to integrate data mining and database management systems and extend SQL to include operators for constrained rule mining. M-SQL also allows limited querying of mined rules. M-SQL was significantly improved in [9] by the Rule-QL language, which has a formal rule cal-

culus based on the full set of first-order-logic expressions and is able to query multiple rulebases. The query part of the proposed work aims to expand Rule-QL to other models.

The integration of data mining and databases is called "inductive databases" [4], focusing on constrained pattern/rule mining within the database engine. However, predictive models are also important. The number of predictive models can be quite large and thus in need of effective management.

On the industrial front, Microsoft SQL Server 2005 provides model-building tools (such as decision trees, neural networks, and association rules). Model building and testing can be done through a data mining query language like SQL. However, such languages are limited, useful mainly for testing a model to determine model accuracy.

Oracle Data Mining (ODM) supports a range of data mining models and model-assessment and model-inspection capabilities. Model metadata can be queried through SQL (via the PL/SQL API). Data mining techniques are embedded in ODM's database engine; in Microsoft SQL Server, the data mining server and the database server are separate entities.

SAS acknowledged the importance of generating and managing models by adding model-management capabilities to its Enterprise Miner. Users are thus able to "register" their models with the SAS Metadata Server for subsequent retrieval. SAS also provides a Web-based model repository viewer for the user-analyst to browse the deposited models.

The Java Data Mining (JSR-73) standard [3] enables the embedding of model-management functions into applications. For example, it allows the importing and exporting of multiple model representations, as well as the querying of model metadata.

Although these initial approaches toward embedding model-management functions constitute an excellent starting point, much more needs to be done to develop a deeper understanding of model management. We propose to pursue this work further by focusing on automating a significant part of the model building, management, and analysis process for applications in which a large number of models are built.

We also note that little prior work was done on automated analysis of modelbases, yet it remains a major and challenging research topic. When dealing with a large number of models, it is necessary to analyze them using automatic techniques to find useful models, delete useless ones, identify weaknesses of models, and suggest repair actions. Manual analysis based on the experience of analysts alone is not feasible.



## BUILDING A LARGE NUMBER OF MODELS

Model management involves the following major tasks:

*Building.* Semiautomated or automated generation of a large number of models and the organization and storage of the models;

*Analyzing.* Querying and analyzing models in the modelbase; and

*Maintaining.* Keeping models and modelbases up to date as the environment and corresponding data change.

Traditionally, data analysts constructed models manually by choosing dependent and independent variables, along with the type and the structure of the model, then fitting the model against the data. This approach is fine for generating a limited number of models but does not scale to applications dealing with hundreds of thousands of models. Therefore, it is important to develop effective and efficient methods for generating large modelbases.

The first step in the process is to determine the data model of the modelbase and define its schema. Heterogeneous models can be organized in modelbases by grouping them based on either application or model type. In the former, models belonging to the same application are stored in the same table; in the latter, models of the same type are stored in the same table, so, for example, all decision trees are stored in a separate table, and all the logistic regressions are stored in yet another table. Although each method has certain advantages, we focus on the latter and assume that models are grouped together based on type. This approach is more natural because it is more convenient to design a schema for a single type of model than it is for multiple types; different models may require different representations and have different model properties. The schema of the modelbase may contain the following fields for supervised models:

*ModelID.* The key attribute uniquely identifying a model in the modelbase;

*TrainData.* A pointer to the data file used for building the model;

*TestData.* Like TrainData but pointing to the test data to compute model performance; it may be optional since a separate test data set may not always be needed, as in cross validation;

*Model.* The “object” in which the model is stored; and

*Model property attributes.* For defining properties of the model derived by accessing the model attributes; for example, in the case of decision trees,

certain statistics (such as the number of nodes in a tree) can be precomputed and stored as model-property attributes, which can vary depending on model type.

This schema may need to be extended or customized to suit application-specific needs.

Once modelbase schemas are designed, the modelbase must be populated with models. A scalable approach is to generate models and populate the modelbase semiautomatically. That is, the analyst iteratively and interactively formulates requests of the form “for data set X build the models of type Y and of the form Z” where X is the data set defined by the TrainDataID identifier of the modelbase, Y is the type of model to be generated (such as decision tree and SVM), and Z is an expression specifying a template (or constraint) defining the models that should be built and stored. For example, analysts can build all the decision trees having “Purchase\_Decision” as a class attribute and the “Income” variable as the root node. This approach must be complemented with various methods for selecting the sets of most important attributes, performing the right transformations to the data, and assisting analysts in imposing better and more meaningful constraints Z on the parameters of the models to be generated.

## ANALYZING MODELBASES

Analysis aims to improve users’ understanding of the models stored in modelbases. This can be done either manually (or semiautomatically) by a domain expert using analysis tools or automatically by a software system deploying model-inferencing techniques:

*Manual model analysis.* The manual approach includes techniques and tools that allow data analysts to examine and evaluate large collections of heterogeneous models. The most effective techniques include the modelbase query language, model-examination operators, and modelbase usage analysis and reporting tools;

The query language should support multiple modelbases of heterogeneous models (such as logistic regression, decision trees, and rules). It should also be based on first-order logic, supporting a set of model-independent and model-specific functions for extracting certain properties from these models. An example of a model-independent function is the predictive accuracy of the model; corresponding examples of queries utilizing such functions would be “retrieve all models with predictive accuracy greater than 95%” and “retrieve the model with the highest predictive accuracy.” The latter query would require a self-join

and can be expressed in a modelbase calculus as:

$$\{m \mid MB(m) \wedge (\Rightarrow n) (MB(n) \Rightarrow p(m) > p(n))\}$$

where  $MB$  is the modelbase and  $p(m)$  is the function returning the predictive accuracy of model  $m$ .

The modelbase query language should support functions pertaining to particular types of models (such as number of nodes in a decision tree or beta-coefficients in a logistic regression model). A query can then request retrieval of all decision trees with number of nodes greater than  $n$  or logistic regression models with at least one beta-coefficient greater than 1. The query language should also be able to access the data used for training and testing the model, leveraging this access to ask questions about the models and data.

In [8], we showed that although SQL can be used to express a large class of model-management queries, it is not sufficient for expressing all important queries and hence needs to be extended. For example, the query “find minimal association rules” (those whose lefthandside and righthandside do not contain the lefthandside and righthandside of any other rule respectively) could not be expressed with the standard SQL over the basic modelbase schema in which each association rule constitutes a model. However, it is possible to use macros over certain types of schema to express the query in SQL, though the resulting query is extremely slow.

If the underlying modelbase schema is XML-based, it is possible to use an XML-based query language (such as XQuery) to query modelbases. To do this, it is necessary to extend the language with functions pertaining to specific modelbases, as in SQL-based queries. Yet another issue is query optimization. All these issues constitute interesting research topics and are described in [8].

**In addition** to query languages, we expect online-analytical-processing-style browsing and examination operators to be implemented to let end users slice and dice modelbases in search of good and underperforming models. SAS Enterprise Miner already supports some of this basic functionality by providing model-browsing capabilities.

We also expect modelbase usage analysis and reporting tools to be added to manual model-analysis systems. Such tools are the equivalent of report generators in modern database management systems, helping data analysts identify the usage patterns of various models or classes of models in the modelbase.

*Automated model analysis.* When automatically analyzing a modelbase, the analyst looks to identify:

*Underperforming models.* Models whose performance needs improvement;

*Dominated models.* Models (dominated by other models) that can be removed from the modelbase; and

*Missing new models.* Models not in the modelbase but that should be added to it to enhance the modelbase’s overall capability.

*Identification and modification of underperforming models.* Assume that the analyst measures the performance of models  $m$  in modelbase  $M$  using some performance measure  $\mu(m)$  (such as the predictive accuracy of  $m$ ). The analyst may want to identify underperforming models  $m$  that can be improved by being replaced with better-performing models  $m'$ , such that  $\mu(m') > \mu(m)$ . These better-performing models  $m'$  can be obtained from  $m$  through several techniques: One changes parameters of a model based on similar but more successful models. For example, assume that an SVM model  $m$  from modelbase  $M$  predicting online purchases for young customers in Chicago performs poorly. A similar SVM model  $m' \in M$  for young customers in New York performs much better. The analyst can then take the kernel function and parameters from the New York model  $m'$ , transfer them to the Chicago model  $m$ , and test them on  $m$  to see if  $m$ 's performance improves. As another technique, model  $m$  can be replaced by a better model  $m'$  as follows: Assume that SVM is used as model  $m$  for the young customers in Chicago and Naive Bayes as model  $m'$  for young customers in New York and that  $\mu(m') > \mu(m)$ . The analyst can then try to replace SVM with the Naive Bayes model for young customers in Chicago.

Other techniques for identification and modification of underperforming models might also be available; developing them is an interesting research problem.

*Identification of dominated models.* Structural dominance is a type of dominance, whereby one model is better than another model in terms of structure. For example, in applications in which each association rule is treated as a single model, if we have association rules  $X \rightarrow Z(c, s)$  ( $c$  is the confidence and  $s$  is the support) and  $X, Y \rightarrow Z(c', s')$ , and  $c > c'$  and  $s > s'$ , then the first rule dominates the second rule, and only the first rule should be kept [6]. Developing efficient methods that would effectively remove dominated models from the modelbase represents an interesting research problem that would require studies of dominance relationships for different types of models.

*Adding new models to the modelbase.* In order to

identify which models are missing from a modelbase (and therefore should be added to it), we need to perform some form of automated inference on the modelbase similar to the logical inference used in logic, artificial intelligence, and deductive databases. In the most general case, we might define and compute the closure of models, as in the concept of logical implication used in logic and closure of functional dependencies in databases. Addressing model inferencing is difficult in its most general form, though it is possible to focus initially on the most tractable subproblems.

An example of such a subproblem is the development of methods producing new models  $m'$  from existing models  $m' \in M$  such that  $m'$  dominates  $m$  as described earlier. This entails developing efficient methods for discovering such dominating models  $m'$ , given the initial set of models  $M$ .

### MAINTAINING MODELBASES

Model performance changes over time based on the changes to the environment and to the corresponding data. In application or deployment mode, it is crucial to monitor model performance on real-world tasks. One way to do this is for the analyst to consider a model-performance measure  $\mu(m)$  and a threshold value for each model  $m$ . If a model is no longer accurate, that is, if its performance is below the threshold, it should be modified or removed from the modelbase and a new model constructed based on the new data. The monitoring can be done periodically or continuously, depending on the application. In addition to monitoring the performance of individual models, the collective performance of groups of models can also be monitored by tracking each individual model's performance statistics. The traditional life-cycle issue of models is also important. Mechanisms for tracking the construction, evaluation, certification, deployment, and deletion of models should be provided to analysts and constitute interesting research topics.

### CONCLUSION

Although several initial approaches to model management have been proposed by leading computer science and IS researchers, none covers the full scope of the model-management capabilities we have described here. We've studied this problem and identified several research issues that should be addressed. The most challenging is how to develop automated modelbase analysis tools, and researchers will need many years to explore it fully. On the other hand, many problems pertaining to model building and maintenance (such as the automated generation and storage of large modelbases and manual model

analysis tools) are doable and likely to be solved much sooner.

Building and managing very large modelbases is a pressing issue. The development of model-management systems would help data analysts and naive users alike build better models, work with many more models, and make data mining models a common resource in an enterprise. This promise will make model management an even more important research topic in data mining. **□**

### REFERENCES

1. Bernstein, P. Applying model management to classical meta data problems. In *Proceedings of the Conference on Innovative Data Systems Research* (Asilomar, CA, Jan. 5–8, 2003), 209–20.
2. Han, J., Fu, Y., Wang, W., Koperski, K., and Zaiane, O. DMQL: A data mining query language for relational databases. In *Proceedings of the SIGMOD Workshop on Data Mining and Knowledge Discovery* (May 1996), 27–34.
3. Hornick, M., Yoon, H., and Venkayala, S. Java data mining (JSR-73): Status and overview. In *Proceedings of the Second International Workshop on Data Mining Standards, Services, and Platforms* (Seattle, Aug. 22, 2004), 23–29.
4. Imielinski, T. and Mannila, H. A database perspective on knowledge discovery. *Commun. ACM* 39, 11 (Nov. 1996), 58–64.
5. Krishnan, R. and Chari, K. Model management: Survey, future directions, and a bibliography. *Interactive Transactions of OR/MS* 3, 1 (2000); itorms.iris.okstate.edu/doc.html.
6. Liu, B., Hsu, W., and Ma, Y. Pruning and summarizing the discovered associations. In *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Diego). ACM Press, New York, 1999, 125–134.
7. Liu, B., Zhao, K., Benkler, J., and Xiao, W. Rule interestingness analysis using OLAP operations. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Philadelphia, Aug. 20–23). ACM Press, New York, 2006, 297–306.
8. Tuzhilin, A., Liu, B., and Hu, J. *Building and Querying Large Modelbases, Working Paper CeDER-05-16*. Stern School of Business, New York University, New York, 2005.
9. Tuzhilin, A. and Liu, B. Querying multiple sets of discovered rules. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Edmonton, Alberta, Canada). ACM Press, New York, 2002, 52–60.
10. Virmani, A. and Imielinski, T. M-SQL: A query language for database mining. *Journal of Data Mining and Knowledge Discovery* 3, 4 (Dec. 1999), 373–408.
11. Will, H. Model management systems. In *Information Systems and Organization Structure*, E. Grochla and N. Szyperki, Eds. Walter de Gruyter, Berlin, 1975, 468–482.
12. Wright, G., Chaturvedi, A., Mookerjee, R., and Garrod, S. Integrated modeling environments in organizations: An empirical study. *Information Systems Research* 9, 1 (Mar. 1998), 64–84.

---

**BING LIU** (liub@cs.uic.edu) is a professor in the Department of Computer Science at the University of Illinois at Chicago.

**ALEXANDER TUZHILIN** (atuzhili@stern.nyu.edu) is a professor of information systems and NEC Faculty Fellow in the Information, Operations & Management Sciences Department of the Leonard N. Stern School of Business at New York University, New York.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

---

© 2008 ACM 0001-0782/08/200 \$5.00

DOI: 10.1145/1314215.1314230

# ACM Digital Library

[www.acm.org/dl](http://www.acm.org/dl)

## The Ultimate Online INFORMATION TECHNOLOGY Resource!



Powerful and vast in scope, the **ACM Digital Library** is the ultimate online resource offering unlimited access and value!

The **ACM Digital Library** interface includes:

- **The ACM Digital Library** offers over 40 publications including all ACM journals, magazines, and conference proceedings, plus vast archives, representing nearly two million pages of text. The ACM DL includes full-text articles from all ACM publications dating back to the 1950s, as well as third-party content with selected archives. [www.acm.org/dl](http://www.acm.org/dl)



- **The Guide to Computing Literature** offers an enormous bank of more than one million bibliographic citations extending far beyond ACM's proprietary literature, covering all types of works in computing such as journals, proceedings, books, technical reports, and theses! [www.acm.org/guide](http://www.acm.org/guide)

- **The Online Computing Reviews Service** includes reviews by computing experts, providing timely commentary and critiques of the most essential books and articles.



Available only to ACM Members.

Join ACM online at [www.acm.org/joinacm](http://www.acm.org/joinacm)

To subscribe to the ACM Digital Library, contact ACM Member Services:

Phone: 1.800.342.6626 (U.S. and Canada)

+1.212.626.0500 (Global)

Fax: +1.212.944.1318

Hours: 8:30 a.m.-4:30 p.m., Eastern Time

Email: [acmhelp@acm.org](mailto:acmhelp@acm.org)

Mail: ACM Member Services

General Post Office

PO Box 30777

New York, NY 10087-0777 USA

ACM Professional Members can add the ACM Digital Library for only \$99 (USD). Student Portal Package membership includes the Digital Library. Institutional, Corporate, and Consortia Packages are also available.



Association for  
Computing Machinery

Advancing Computing as a Science & Profession

# Record Matching in Digital Library Metadata

Using evidence from external sources to create more accurate matching systems.

**W**hen data stores grow large, data quality, cleaning, and integrity become issues. The commercial sector spends a massive amount of time and energy canonicalizing customer and product records as their lists of products and consumers expand. An Accenture study in 2006 found that a high-tech equipment manufacturer saved \$6 million per year by removing redundant customer records used in customer mailings. In 2000, the U.K. Ministry of Defence embarked on the massive “The Cleansing Project,” solving key problems with its inventory and logistics and saving over \$25 million over four years.

In digital libraries, such problems manifest most urgently not in the customer, product, or item records, but in the metadata that describes the library’s holdings. Several well-known citation lists of

computer science research contain over 50% duplicate citations, although none of these duplicates are exact string matches [2]. Without metadata cleaning, libraries might end up listing multiple records for the same item, causing circulation problems, and skewing the distribution of their holdings. In addition, when different authors share the same name (for example, Wei Wang, J. Brown), author disambiguation must be performed to correctly link authors to their respective monographs and articles, and not to others. Metadata inconsistencies can be due to problems with varying ordering of fields, type of delimiters used, omission of fields, multiple representations of names of people and organizations, and typographical errors.

When libraries import large volumes of metadata from sources that follow a metadata standard, a manually compiled set of rules

called a crosswalk may be used to transform the metadata into the library’s own format. However, such crosswalks are expensive to create manually, and public ones exist only for a few, well-used formats. Crucially, they also do not address how to detect and remove inexact duplicates. As digital libraries mine and incorporate data from a wider variety of sources, especially noisy sources, such as the Web, finding a suitable and scalable matching solution becomes critical.

Here, we examine this problem and its solutions. The de-duplication task takes a list of metadata records as input and returns the list with duplicate records removed. For example, the search results shown in the figure here are identical and should have been combined into a single entry. It should be noted that many disciplines of computer science have instances of similar inexact match-

ing problems, and as such this problem has many names, including de-duplication, data cleaning, disambiguation, record linkage, entity resolution, attribution, and plagiarism detection. While these variant problems differ in specifics, a common key operation is to determine whether two data records match. We explain this problem and generate awareness of the approaches espoused by different communities. For a detailed review, we urge readers to consult the individual papers and a more detailed survey paper [5].

## UNINFORMED STRING MATCHING

In its most basic form, record matching can be simplified as string matching, which decides whether a pair of observed strings refer to the same underlying item. In such cases, we use the similarity between the strings to calculate whether they are coreferential. When such pairwise similarity measures are viewed as kernel-comparison operations, the record-matching problem can be cast as a clustering problem. If two records' similarity exceeds a threshold, they are considered two variants of the same item. String similarity measures can be classified as either set- or sequence-based, depending on whether or not ordering information is used.

Set-based similarity considers the two strings as independent sets of characters  $S$  and  $T$ , such as the Jaccard measure, defined as the ratio of the intersection of the sets over the union (see Equation 1).

$$\frac{|S \cap T|}{|S \cup T|}$$

Equation 1. Jaccard measure.

Cosine similarity, borrowed from information retrieval, views both sets as vectors and calculates the angle between the vectors, where a smaller angle indicates higher similarity. Alternatively, asymmetric measures, such as degree of similarity (see Equation 2) may

$$\frac{|S \cap T|}{|S|}$$

Equation 2. Degree of similarity.

be more appropriate when one string is more important to match than the other.

Sequence-based measures can be generally cast as edit distances. They measure the cost of transforming one ordered string into the other. Typically, the transformation cost is measured by summing the cost of simple incremental operations, such as insertion, deletion, and substitution.

Hybrids of both set- and sequence-based measures are often used. For example, when the string is a series of words, a sequence-based measure may be employed for individual tokens, but the string as a whole may be modeled as a set of tokens [3].

## INFORMED SIMILARITY AND RECORD MATCHING

Library metadata records contain

a wide variety of data—personal names, URLs, controlled subject headers, publication names, and years. Viewing the list as a database table, each of these columns may have its own notions for what is considered acceptable variation (“Liz” = “Elizabeth”; “Comm. of the ACM” = “CACM”; 1996 ≠ 1997). Knowing what type of data exists in a column can inform us of what constitutes similarity and duplication. As such, string similarity measures are usually weighted differently per column.

Certain data types have been studied in depth. In fact, the need to consolidate records of names and addresses in government and industry pioneered research to find reliable rules and weights for record matching. In set-based similarity, tokens may be weighted with respect to their (log) frequency, as is done in information retrieval models. In sequence-based edit operations, a spectrum of weighting schemes has been used to capture regularities in the data, basically by varying the edit cost based on the position and input. For example, in genomic data, sequences often match even when a whole substring is inserted or deleted; the same is true when matching abbreviations to their full forms. In census data, the initial letters of people's names are rarely incorrect.

Such models need to set parameters, such as the cost for each type of edit operation in a principled way. Fortunately, data-driven methods have emerged to learn

optimal weights from training data (see [2, 12]).

### GRAPHICAL FORMALISMS FOR RECORD MATCHING

Graphical formalisms are becoming popular for record matching.

Typically, columns or whole records are viewed as nodes in a graph with edges connecting similar nodes, allowing global information to be incorporated in the

disambiguation process. One may assign similarity values to edges and identify cliques of high weights as matching nodes.

A common manifestation of graphical formalisms in disambiguation tasks is social networks, such as collaboration networks. Social network analysis methods, such as centrality and betweenness, can be applied. For example, in author disambiguation we may be able to attribute two papers to the same “Wei Wang” when the co-author lists do not have common names but share names with a third paper; the two nodes are connected by a path through a third node. Yet another work uses network cuts and random walks in the collaboration network of actors to disambiguate names in the Internet Movie Database [7].

Consolidating records using one column of data can sometimes cascade and benefit matching on other columns of the same data. This incremental approach can resolve duplicates when true



Searching for “computers and intractability” on Google Scholar.

matching records do not exceed a global similarity threshold before individual fields in the records are merged. Graphical formalisms, such as dependency graphs [4] or conditional random fields [11], nicely model incremental record matching, enabling the propagation of contextual similarity.

Graphical formalisms in the form of generative probabilistic models have also been suggested. In the author-disambiguation problem, we can view authors as members of collaborative groups. This model first picks out collaborative groups and then assigns authors within these groups to generate references. We can then run this model in the opposite direction to infer which collaborative group (thus which disambiguated author) is responsible for a particular work [1]. Such graphical models have outperformed methods using pairwise compar-

isons in accuracy but have yet to demonstrate efficiency on large data sets.

### REDUCING COMPLEXITY

Since digital libraries often contain large numbers of records, brute-force pairwise comparison is often infeasible. As of 2005, the estimated number of independent articles and monographs in computer science research alone exceeded 2.2 million [9], an amount

unsuited for  $O(n^2)$  complexity. (Log) linear time matching algorithms are needed to scale to such large metadata sets.

Observations show the ratio of true record matches to non-matches is very low; it is very unlikely two randomly picked records refer to the same item. Thus, a computationally cheap similarity measure is often used to first separate such implausible matches. These blocking (or canopy) methods map records into a set of blocks in linear time. For example, we can construct a block for all records that have the token “J” and another block for all records that have the token “Brown.” Records containing both tokens would be members of both blocks. More computationally expensive similarity measures can then be confined to run only within each block, where records have a non-zero probability of matching.

Constructing an optimal block algorithm requires tuning parameters for the proper number of

blocks, overlap between blocks, and size of the blocks. These parameters can be either rigorously controlled to bound the complexity of the inner comparison [6] or learned from data or sampling [8].

## CONCLUSION

Matching problems matured to become a research issue as early as the 1940s, probably due to the analysis of census data or medical records [6]. Since then, advances have been made both on better theoretical models for weighted matching and proofs for error bounds and optimality.

One promising direction lies with graphical models, which can be tailored to model the underlying structure of the specific record-matching scenario. A difficulty in applying these models is in complexity; modeling the structure more accurately requires a more complex graphical model, which in turn creates complexity in the inference procedure. A way of reducing this complexity further would help propel these models for large-scale data sets.

Bringing more knowledge to bear on the problem may also help. Noisy sources, such as the Web, can be seen as a treasure trove of statistical information for matching—if carefully cleaned and utilized. This is especially fruitful for library metadata, as information about authors, titles, and publishers is readily available on the Web. Motivated by similar approaches in information retrieval research, we have leveraged Web search results when disambiguating author names in citation lists [10]. Our study showed that using evidence from such external sources alone can achieve the same disambiguation performance as using the record data itself. We can also ask

humans for help directly—by distinguishing which parts of the matching process are easier for humans to solve than machines. The classic Fellegi-Sunter model [12] defines a check zone where uncertain matches are given to human experts to manually check. Similar to approaches used in computer vision, active learning based on manual disambiguation can help create more accurate matching systems. Elusive, domain-specific matching knowledge may be easier to capture by having human experts solve example problems rather than asking them to code explicit rules.

It is unclear whether advances in record matching have kept up with the pace at which information is becoming widely available today. In the world of digital libraries, metadata inconsistencies are a significant barrier to locating and collating knowledge that end users and reference librarians have had to adapt to. In some cases, humans resort to using external sources of information to (in)validate a possible match. As more information becomes Web accessible, we expect mining such external sources for knowledge will play an increasingly useful role in matching problems. Incorporating such external yet accessible knowledge gathering as an active component of matching algorithms will be a valuable research direction. **□**

## REFERENCES

1. Bhattacharya, I. and Getoor, L. A latent dirichlet model for unsupervised entity resolution. In *Proceedings of the SIAM International Conference on Data Mining* (Apr. 2006), 47–58.
2. Bilenko, M. and Mooney, R.J. Adaptive duplicate detection using learnable string similarity measures. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Aug. 2003), 39–48.

3. Cohen, W.W., Ravikumar, P., and Fienberg, S.E. A comparison of string distance metrics for name-matching tasks. In *Information Integration on the Web (IIWeb)*, (Aug. 2003), 73–78.
4. Dong, X., Halevy, A., and Madhavan, J. Reference reconciliation in complex information spaces. In *Proceedings of the ACM SIGMOD International Conference on Management of Data* (June 2005), 85–96.
5. Elmagarmid, A.K., Ipeirotis, P.G., and Verykios, V.S. Duplicate record detection: A survey. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 19, 1 (Jan. 2007), 1–16.
6. Hernández, M.A. A generalization of band joins and the merge/purge problem. Ph.D. dissertation, Columbia University, Mar. 1996.
7. Malin, B., Airoidi, E., and Carley, K.M. A network analysis model for disambiguation of names in lists. *Computational and Mathematical Organization Theory* 11, 2 (July 2005), 119–139.
8. Michelson, M. and Knoblock, C.A. Learning blocking schemes for record linkage. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)* (July 2006).
9. Petricek, V. et al. A comparison of on-line computer science citation databases. In *Proceedings of the European Conference on Research and Advanced Technology for Digital Libraries (ECDL)*, (Sept. 2005), 438–449.
10. Tan, Y.F., Kan, M.-Y., and Lee, D. Search engine driven author disambiguation. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL)* (June 2006), 314–315.
11. Wellner, B. et al. An integrated, conditional model of information extraction and coreference with application to citation matching. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)* (July 2004), 593–601.
12. Winkler, W.E. and Thibaudeau, Y. An application of the Fellegi-Sunter Model of record linkage to the 1990 U.S. Decennial Census. Technical Report RR91/09, U.S. Bureau of the Census, 1991.

---

**MIN-YEN KAN** (kanmy@comp.nus.edu.sg) is an assistant professor at the School of Computing, National University of Singapore and leads the Web Information Retrieval/Natural Language Processing Group (WING). **YEE FAN TAN** (tanyee@comp.nus.edu.sg) is a Ph.D. candidate at the School of Computing, National University of Singapore.

---

© 2008 ACM 0001-0782/08/0200 \$5.00

---

DOI: 10.1145/1314215.1314231





### **Allegheny College**

#### **Department of Computer Science**

The Department of Computer Science invites applications for a tenure-track position beginning Fall 2008. Qualifications include a Ph.D. in computer science. Applicants with interests in graphics, media computation, computational science, or data management and security are preferred; however, candidates from other areas of computer science will be considered.

Applicants must provide evidence of ability to teach effectively at the undergraduate level, commitment to liberal arts education, and continuing contribution to the discipline. Responsibilities include teaching and advising undergraduates, guiding students in senior research projects, and contributing to the college-wide freshman/sophomore seminars that emphasize speaking and writing. Salary will be competitive; start-up funds are available. Rank will be commensurate with credentials and experience.

Allegheny College is a highly selective private liberal arts college with a dedicated faculty of teacher-scholars. Information about the Department and its traditional and applied computing majors is available on the Web at <http://cs.allegheny.edu>.

Send letter of application, vita, statement of teaching and research interests, applicable transcripts, and arrange to have three letters sent from references, at least one of whom can comment on teaching, to: Dr. Robert S. Roos, Associate Professor and Chair, Department of Computer Science, Allegheny College, Meadville, PA 16335; email: [robert.roos@allegheny.edu](mailto:robert.roos@allegheny.edu).

Review of applications will begin immediately. Allegheny College is an Equal Opportunity Employer with a strong institutional commitment to develop a diverse faculty and staff. Women and members of other under-represented groups are especially encouraged to apply.

### **Bilkent University**

#### **Computer Engineering and Electrical Engineering Departments**

The Computer Engineering (CS) and Electronics and Electrical Engineering (EE) Departments of Bilkent University (<http://www.bilkent.edu.tr>) are looking for two Postdoctoral Researchers and Ph.D. students to strengthen their 3D graphics and video research. The researchers will be a member of a team that coordinates the European 3DTV and 3DPHONE projects, and develops 3D graphics and 3D video techniques for next-generation mobile devices. The Postdoctoral Researcher in CS will participate in 3D graphics research. Good knowledge of graphics rendering/animation algorithms, GPUs, 3D user interfaces, game engines, is desirable. The Postdoctoral Researcher in EE will participate in multi-view 3D video research. Experience in multi-view video coding and streaming, video codec standards, is desirable. Salary is 2000-2400 Euros/month (AFTER TAXES) for 3 years (Feb 2008-Feb 2011). To Apply: visit <http://www.cs.bilkent.edu.tr/~tcapin/projects/3dphone/postdoc.htm> for details and application. Search will continue until the position is filled.

### **Boston University**

#### **Metropolitan College Department of Computer Science**

Boston University's Metropolitan College seeks to fill a full-time faculty position in the Dept. of Computer Science (<http://csmet.bu.edu/>) starting September 1, 2008. The successful candidate will hold a Ph.D. in CS, engineering or related field, with expertise in Multimedia, Interactive Web Technologies, Computer Graphics, Game Engineering or related fields and have a strong interest in applied research. Teaching experience is required; professional experience in one or

more of these areas is preferred. Duties include teaching six courses (some on line) per academic year, participating in curriculum development, mentoring part-time faculty and students, and administration in the Department, College and University. Applications to: Chairman, Search Committee, Computer Science Metropolitan College, Boston University, 808 Commonwealth Avenue, Boston MA 02215 or e-mail (preferred) to [mattslow@bu.edu](mailto:mattslow@bu.edu). Application review begins December 17 and continues until the position is filled. Boston University is committed to increasing the diversity of our faculty. Members of underrepresented groups are encouraged to apply. Boston University is an affirmative action, equal opportunity employer.

### **Case Western Reserve University**

#### **Department of Electrical Engineering and Computer Science**

Faculty Positions Division of Computer Science Department of Electrical Engineering and Computer Science Case Western Reserve University Cleveland, Ohio 44106-7071 The Department of Electrical Engineering and Computer Science at Case Western Reserve University is seeking Computer Science faculty candidates at all levels. We invite applicants in our thrust areas Biomedical Informatics and Computational Biology, Information Networks, Smart Power Systems/Power Informatics, and Safe, Secure, and Reliable Software. Also, to support these thrust areas, we invite applicants in knowledge and data management, data mining, information retrieval, computer graphics and visualization, artificial intelligence, machine learning, computer networks, distributed systems, software engineering, computer security, human-computer interaction, and algorithms. See <http://www.zmo@eecs.case.edu/www.eecs.case.edu/> for more details. Please submit CV,

#### **ACM POLICY ON NONDISCRIMINATORY ADVERTISING**

◆ ACM accepts recruitment advertising under the basic premise the advertising employer does not discriminate on the basis of age, color, race, religion, gender, sexual preference, or national origin. ACM recognizes, however, that laws on such matters vary from country to country and contain exceptions, inconsistencies, or contradictions. This is true of laws in the United States of America as it is of other countries. ◆ Thus ACM policy requires each advertising employer to state explicitly in the advertisement any employer restrictions that may apply with respect to age, color, race, religion, gender, sexual preference, or national origin. (Observance of the legal retirement age in the employer's country is not considered discriminatory under this policy.) ACM also reserves the right to unilaterally reject any advertising. ◆ ACM provides notices of positions available as a service to the entire membership. ACM recognizes that from time to time there may be some recruitment advertising that may be applicable to a small subset of the membership, and that this advertising may be inherently discriminatory. ACM does not necessarily endorse this advertising, but recognizes the membership has a right to be informed of such career opportunities. ◆

# Career Opportunities

research and teaching interest statements, and three recommendation letters to: Faculty Search Committee Division of Computer Science Department of Electrical Engineering and Computer Science Case Western Reserve University 10900 Euclid Avenue Cleveland, OH 44106-7071. In employment as in education, Case is committed to affirmative action and equal opportunity. Women and minorities are encouraged to apply. Case is a recipient of a National Science Foundation ADVANCE institutional transformation grant to increase the participation of women in science and engineering.

## College of St. Scholastica Computer Science Department

The CS/CIS Department at the College of St. Scholastica seeks a dynamic faculty member to join our program as a tenure-track, Assistant Professor starting Fall 2008. Candidates should have demonstrated commitment to teaching undergraduate courses and mentoring students. A graduate-level degree in a computer related field is required, as is excellent teaching and communication skills. Additional desired qualifications include: ability to effectively teach a wide variety of

introductory and advanced undergraduate classes, experience with online education, and relevant industry experience.

The CS/CIS Department offers a student-centered, undergraduate degree in CS/CIS and an online Masters in IT Leadership. The department embraces creative teaching methods and offers a wide range of classroom and online courses for computer and health information management majors. The Department is one of the oldest in the field, serving students for 23 years, and is now part of the newly formed School of Business and Technology (SB&T). Review of applications will start immediately and continue until the position is filled. AA/EOE

## Columbia University Department of Computer Science

The Department of Computer Science is seeking applicants for a tenure-track position at either the junior or senior level in parallel and distributed computing. Applicants should have a Ph.D. in a relevant field, and have demonstrated excellence in research and the potential for leadership in the field. Senior applicants should also have demonstrated excellence in teaching

and continued strong leadership in research.

Our department of 34 tenure-track faculty and 1 lecturer attracts excellent Ph.D. students, virtually all of whom are fully supported by research grants. The department has close ties to the nearby research laboratories of AT&T, IBM (T.J. Watson), NEC, Siemens, Telcordia Technologies and Verizon, as well as to a number of major companies. Columbia

University is one of the leading research universities in the United States, and New York City is one of the cultural, financial, and communications capitals of the world. Columbia's tree-lined campus is located in Morningside Heights on the Upper West Side.

Applicants should submit a CV, summaries of research and teaching interests, and the names and email addresses of at least three references by filing an online application at <http://www.cs.columbia.edu/facultyrecruit>

Review of applications will begin on January 1, 2008.

Columbia University is an Equal Opportunity/Affirmative Action Employer.

We encourage applications from women and minorities.



The School of Electrical Engineering and Computer Science at Oregon State University invites applications for up to three tenure-track positions in Computer Science. The School of EECS strongly encourages teamwork and collaboration within the School, and with other departments and universities. We are particularly interested in candidates who can contribute richness and depth to our Graphics/Visualization, End-User Software Engineering and Machine Learning groups. The following areas are strong possibilities for collaboration with these groups: Computer Vision; Human Computer Interaction; Natural Language Processing; Parallel and Distributed Computing (including multi-core and data center computing); Programming Languages; Software Engineering; and Theoretical Computer Science (including algorithms and optimization). Applicants should have an earned doctorate in Computer Science/Computer Engineering and demonstrate a strong commitment to high-quality undergraduate and graduate teaching and the development of a vibrant research program. OSU is one of only two American universities to hold the Land Grant, Sea Grant, Sun Grant, and Space Grant designation and is the only Oregon institution recognized for its "very high research activity" (RU/VH) by the Carnegie Foundation for the Advancement of Teaching. With a faculty of 45, the School of EECS enrolls 1300 undergraduate and 300 MS/PhD students. For more information, including instructions for application, visit <http://www.eecs.oregonstate.edu>. OSU is an AA/EOE.



Careers with Mass Appeal

## College of Engineering

### Department of Computer and Information Science

The Department of Computer and Information Science at the University of Massachusetts Dartmouth invites applications for one full-time tenure-track Assistant Professor starting September 1, 2008. Applicants must have a Ph.D. in Computer Science or closely related field and demonstrate a strong commitment to excellence in teaching, scholarly accomplishments, and active research interests in broadly defined areas of computing with special emphasis on intelligent information systems. The successful candidate will be expected to teach and develop computer science courses for both undergraduate and graduate students, to establish a research program leading to archival publications and external funding that supports graduate students, and to be motivated to engage in collaborative and interdisciplinary research.

The Department has an enrollment of approximately 250 computer science majors, offers a rigorous ABET/CAC-accredited undergraduate major in Computer Science (with focus on object-oriented software design and development), Software Engineering Option, and a well established Master's program in Computer Science. Preparations to offer a Ph.D. in Computer Science are also underway. Current faculty are involved in research in areas of software development and software engineering, foundations of computer science, parallel and distributed computing, object-oriented systems, human/computer interaction, database systems, neural networks, pattern recognition, computer vision, and computer security. You may review the entire job description at [www.umassd.edu/hr/jobs.cfm](http://www.umassd.edu/hr/jobs.cfm).

To be considered for this position applicants should send a letter of interest, resume, samples of publications, a brief statement about teaching philosophy and research plans, and the names, mail and e-mail addresses, and telephone numbers of at least three professional references to Faculty Search (CIS), Office of the Dean, University of Massachusetts Dartmouth, College of Engineering, 285 Old Westport Road, N. Dartmouth, MA 02747-2300.

For more information about the College of Engineering please go to <http://www.umassd.edu/engineering/coe/>. Evaluation of application materials will begin immediately and continue until the position is filled. An official transcript will be required for finalists.

*The University of Massachusetts Dartmouth is an EEO/AA employer.*

# University of Waterloo



## **DePaul University in Chicago, School of CTI**

### **Computer Game Development**

The School of CTI at DePaul University in Chicago invites applications for a full-time tenure-track position in Computer Game Development to beginning in September 2008. Our program is highly cross-disciplinary, linked closely with CS, SE, Computer Graphics, Animation, Digital Cinema and Interactive Media. The program has grown to over 150 undergraduate majors since 2004, and benefits from strong ties to the Chicago-area game industry. Ideal candidates will have an active interest in game research and/or development, and be able to teach in multiple areas of game programming and/or computer graphics. Industry experience is a plus. Job responsibilities include teaching, advising, service, supervising student game projects and continued creative activities and/or scholarship in the field. Candidates should have a PhD, MFA or equivalent degree by the date of appointment. Candidates without PhD or MFA may be eligible for adjunct or full-time non-tenure track positions and are welcome to apply. For priority consideration, apply by January 15, 2008 to [game\\_search@cti.depaul.edu](mailto:game_search@cti.depaul.edu). For application details, see job posting at <http://www.cti.depaul.edu/news/jobs.asp>. CTI is interested in recruiting and maintaining a diverse group of faculty. Members of all underrepresented groups, women, veterans, and persons with disabilities are invited and encouraged to apply.

## **East Carolina University College of Technology and Computer Science**

Tenure track assistant professor position starts August 20, 2008. The successful candidate will be committed to teaching at the undergraduate and graduate levels and will be expected to maintain an active research program in software engineering.

**Minimum Qualifications:** Possess, or complete by August 2008, a PhD from an appropriately accredited institution in computer science, software engineering, or a closely related area.

Screening will begin December 15, 2007, and continue until the position is filled. Candidates must complete a candidate profile and submit a vita and a letter of application online at [www.jobs.ecu.edu](http://www.jobs.ecu.edu). Please send official transcripts and three current letters of reference (one or more of which should address teaching) to:

Software Engineering Search Committee  
College of Technology and Computer  
Science  
Suite C-122  
Science and Technology Building  
East Carolina University  
Greenville, NC 27858

Applications are invited for one or more David R. Cheriton Chairs in Software Systems. These are senior positions and include substantial research support and teaching reduction. Candidates with outstanding research records in software systems (very broadly defined) are encouraged to apply. Successful applicants who join the University of Waterloo are expected to be leaders in research, have an active graduate student program and contribute to the overall development of the School.

A Ph.D. in Computer Science, or equivalent, is required, with evidence of excellence in teaching and research. Rank and salary will be commensurate with experience, and appointments are expected to commence during the 2008 calendar year. The Chairs are tenured positions.

With over 70 faculty members, the University of Waterloo's David R. Cheriton School of Computer Science is the largest in Canada. It enjoys an excellent reputation in pure and applied research and houses a diverse research program of international stature. Because of its recognized capabilities, the School attracts exceptionally well-qualified students at both undergraduate and graduate levels. In addition, the University has an enlightened intellectual property policy which vests rights in the inventor: this policy has encouraged the creation of many spin-off companies including iAnywhere Solutions Inc., Maplesoft Inc., Open Text Corp and Research in Motion. Please see our website for more information:

<http://www.cs.uwaterloo.ca/>

Applications should be sent by electronic mail to

[cs-recruiting@cs.uwaterloo.ca](mailto:cs-recruiting@cs.uwaterloo.ca)

or by post to

Chair, Advisory Committee on Appointments  
David R. Cheriton School of Computer Science  
200 University Avenue West  
University of Waterloo  
Waterloo, Ontario  
Canada N2L 3G1

An application should include a curriculum vitae, statements on teaching and research, and the names and contact information for at least three referees. Applicants should ask their referees to forward letters of reference to the address above. Applications will be considered as soon as possible after they are complete, and as long as positions are available.

The University of Waterloo encourages applications from all qualified individuals, including women, members of visible minorities, native peoples, and persons with disabilities. All qualified candidates are encouraged to apply; however, Canadian citizens and permanent residents will be given priority.

Fall 2007

Official transcripts and proper documentation of identity and employability will be required at the time of employment.

Equal Opportunity / Affirmative Action Employer

View complete job posting at  
<https://ecu.peopleadmin.com/applicants/Central?quickFind=56907>.

### Eastern Kentucky University

#### Assistant Professor, Computer Science

Department of Computer Science at Eastern Kentucky University has an opening for a tenure-track position at the Assistant Professor level for Fall 2008. A Ph.D. in Computer Science or related field is preferred, but ABD minimally required. Degree must be from a regionally (SACS, Middle States, North Central, etc.) accredited and recognized international institution. Preference will be given to candidates with specialties in computer security, data mining, or multimedia. We expect a strong commitment to teaching, service, and continued scholarly activity. The successful candidate will be teaching an array of courses including service, undergraduate and graduate level courses in Computer Science.

For detailed job description and application, visit <http://www.cs.eku.edu/facultyposition.htm>. Must apply online at <http://jobs.eku.edu> (search requisition number 0602390). Offers of employment are contingent upon satisfactory background check and educational credential verification. Eastern Kentucky University is an EEO/AA institution that values diversity in its faculty, staff and student body. In keeping with this commitment, the University welcomes applications from diverse candidates and candidates who support diversity.

### Harvard University Division of Engineering and Applied Sciences Center for Research on Computation and Society <http://www.crcs.deas.harvard.edu>

Postdoctoral Fellowships Center for Research on Computation and Society, Harvard University

The Center for Research on Computation and Society solicits applications for its Postdoctoral Fellows Program for the 2007-08 academic year. Fellows are given an annual stipend of

\$55,000 for up to three years to engage in a program of original research, and provided with additional funds for travel and research support.

For the past several years, we have recruited outstanding fellows in the specific focus areas of Privacy and Security. This year, we are broadening our search to additional areas related to Computation and Society including (but not limited to):

- Health and Medical Computing Technology
- Human Computer Interaction
- Automated Decision Making
- Privacy and Security

For the coming year we are especially interested in candidates whose work has connection with the developing uses of computer systems in health care. Given the purview of the Center, we are looking for researchers with broad interdisciplinary interests, desire to work with both computer scientists and colleagues from other disciplines, and demonstrated interest in connecting their research agenda with societal issues. Fellows work with computer science faculty and others from across Harvard University on their own research.

Harvard University is an Affirmative Action/Equal Opportunity Employer. We are particularly interested in attracting women and under represented groups to participate in the CRCS Fellows Program.

Applicants should send a cover letter, CV, research statement, and copies of up to three research papers, and have up to three letters of reference sent to:

- Postdoctoral Fellows Program
- Center for Research on Computation and Society
- Maxwell-Dworkin Laboratory — 247
- Cambridge, MA 02138

The application deadline for full consideration is February 15, 2008. Candidates are urged to visit <http://www.crcs.deas.harvard.edu/> for further information.

### Humboldt-Universitaet zu Berlin

Assistant Professor, Information Systems  
Humboldt's School of Business and Economics is looking for an Assistant Professor in Information Systems (Wirtschaftsinformatik) starting October 2008 or sooner. Applicants should hold a Ph.D. in a related field and have published in the areas of IT strategy, enterprise software (ERP), IT in the supply chain, ubiquitous computing, or privacy and security. Knowledge of German is not required. Qualified women are especially invited to apply. E-mail your application to Dean Oliver Guenther at [bujarek@wiwi.hu-berlin.de](mailto:bujarek@wiwi.hu-berlin.de).

# VCU

Virginia Commonwealth University

## COMPUTER SCIENCE POSITIONS

### School of Engineering

The Computer Science Department at Virginia Commonwealth University (VCU) invites applications for two tenure-track positions. Candidates with demonstrated evidence of active externally funded research at all academic levels will be considered. All specializations will be considered although areas of special interest are cyber-security and biomedical informatics. The successful candidates will have the ability to work with diverse populations in an interdisciplinary environment.

Review of applications will begin on January 15, 2008, and will continue until filled. To apply, submit a CV, the names and e-mail addresses of three references, and a summary of research and teaching interests by e-mail to: [cmcs\\_search@vcu.edu](mailto:cmcs_search@vcu.edu)

The Computer Science Department resides in a new state-of-the-art building of the School of Engineering, designed to support its rapidly expanding research programs and initiatives. VCU, the largest university in Virginia, is an urban, research I extensive institution with a richly diverse university community and commitment to multicultural opportunities. It is located in historic Richmond, Virginia's capital city. Richmond overlooks the James River and is within two hours of Washington, DC, the Atlantic Coast, and the Appalachian Mountains.

*VCU is an Equal Opportunity, Affirmative Action Employer. Women, minorities and persons with disabilities are encouraged to apply.*

# Career Opportunities

## Hendrix College

### Computer Science Department

Hendrix College, a central Arkansas liberal arts college, announces a 2-year, full-time, non-tenure-track position in computer science, starting Aug 2008. An M.S. is required; a Ph.D. is preferred. Candidates from all research specialties are welcome. Responsibilities include teaching five courses a year, distributed across all undergraduate levels, and directing undergraduate research. Applications will be reviewed upon receipt. Full information at [ozark.hendrix.edu](http://ozark.hendrix.edu)

## Merrimack College

### Computer Science Department

The Computer Science Department at Merrimack College invites applications for an assistant professor for a one year position (continuation or conversion to tenure track may be possible) to begin Fall, 2008. A Ph.D. in Computer Science preferred. Excellence in teaching and continued scholarly activity is expected. Candidates are required to send a letter of application, curriculum vita, copies of graduate transcripts, statement of teaching philosophy and research interests, and three letters of reference to Professor Mary G. Noonan, Chair, Department of Computer Science, Merrimack College, 315 Turnpike Street, Box N5, North Andover, MA 01845. Email submissions also accepted ([mary.noonan@merrimack.edu](mailto:mary.noonan@merrimack.edu)) preferred. Permanent US work authorization and/or US Citizenship is required. Review of applicants will begin January 1, 2008

## Middle East Technical University Northern Cyprus Campus Academic Positions (PhD) in Computer Science & Engineering

The applicants are expected to possess proven research and teaching records. They should be able to teach in the CS&E core areas such as Algorithms, Operating Systems, Data Comm., Software Engineering, DBMS, Principles of Programming Languages, Formal Languages and possibly specialized elective courses. For application process please go to <http://www.ncc.meta.edu.tr/academicpositions08-09/> or send mail to [nccapos@metu.edu.tr](mailto:nccapos@metu.edu.tr)

## Mississippi State University Head Department of Computer Science and Engineering

Applications and nominations are being sought for the Head of the Department of Computer Science and Engineering ([www.cse.msstate.edu](http://www.cse.msstate.edu)) at Mississippi State University. This is a 12-month tenure-track position.

Part of the Bagley College of Engineering,

the department has approximately 325 undergraduate majors, 70 graduate students, and 18 tenured and tenure-track faculty. The department offers undergraduate programs in Computer Science and Software Engineering, and jointly administers the undergraduate program in Computer Engineering with the Department of Electrical and Computer Engineering. At the graduate level, we offer M.S. and Ph.D. degrees in Computer Science and faculty also direct graduate students in Computational Engineering and Computer Engineering. Certificates in Software Engineering, Information Assurance, and Computational Biology are also available. Research expenditures total about \$3 million dollars annually and the university as a whole is ranked 48th among 271 U.S. institutions in computer science expenditures. Research areas for the department are high-performance computing, artificial intelligence, graphics and visualization, computer security, and software engineering. Three current faculty members have received NSF CAREER awards. Our computer security area has been designated a National Center of Academic Excellence in

Information Assurance Education (CAEIAE) by the National Security Agency (NSA).

Mississippi State University is a comprehensive land-grant institution with approximately 17,000 students and about 1,000 faculty members. The university is a leader in high performance computing, housing a supercomputer in the top 20 among U.S. universities. The university's main campus is located in Starkville, Mississippi, a vibrant community approximately 2 hours from Jackson MS, Birmingham AL, and Memphis TN.

The successful Head will provide:

- Vision and leadership for nationally recognized computing education and research programs
- Exceptional academic and administrative skills
- A strong commitment to faculty recruitment and development

Applicants must have a Ph.D. in computer science, software engineering, computer engineering, or a closely related field. The successful candidate must have earned national recognition by a



## Windows Kernel Source and Curriculum Materials for Academic Teaching and Research.

The Windows® Academic Program from Microsoft® provides the materials you need to integrate Windows kernel technology into the teaching and research of operating systems.

The program includes:

- **Windows Research Kernel (WRK):** Sources to build and experiment with a fully-functional version of the Windows kernel for x86 and x64 platforms, as well as the original design documents for Windows NT.
- **Curriculum Resource Kit (CRK):** PowerPoint® slides presenting the details of the design and implementation of the Windows kernel, following the ACM/IEEE-CS OS Body of Knowledge, and including labs, exercises, quiz questions, and links to the relevant sources.
- **ProjectOZ:** An OS project environment based on the SPACE kernel-less OS project at UC Santa Barbara, allowing students to develop OS kernel projects in user-mode.

*These materials are available at no cost, but only for non-commercial use by universities.*

For more information, visit [www.microsoft.com/WindowsAcademic](http://www.microsoft.com/WindowsAcademic)  
or e-mail [compsci@microsoft.com](mailto:compsci@microsoft.com).

THE UNIVERSITY OF TENNESSEE  
AT CHATTANOOGA

Dean of the College of  
Engineering and Computer Science

The University of Tennessee at Chattanooga invites applications and nominations for the position of Dean of the College of Engineering and Computer Science. Candidates should have a terminal degree in engineering, computer science, or a closely related discipline; a successful record as a teacher and scholar; and administrative experience. Experience in business or industry will be an added value.

Applicants should submit a current curriculum vita with a letter explaining how their personal and professional strengths fit the qualifications for this position, and contact information for three references. Applicants will be notified before references are contacted. Please submit application materials electronically to Rent Consulting Group, LLC, at [rgent@rentconsultinggroup.com](mailto:rgent@rentconsultinggroup.com). Nominations are encouraged and can be made to the same email address. Review of completed applications will begin on February 25, 2008 and will continue until the position is filled. The preferred starting date is July 1, 2008. Visit the University at [www.utc.edu/administration/academicaffairs/](http://www.utc.edu/administration/academicaffairs/).

The University of Tennessee at Chattanooga is an equal employment opportunity, affirmative action/Title VII/Title IX/Section 504/ADA/ADEA institution.



**Nicholls State University**  
**Department of**

**Mathematics and Computer Science**

The Department of Mathematics and Computer Science at Nicholls State University invites applications for the John A. Brady, Sr. Endowed Chair in Computer Science beginning Fall 2008.

Placement at the associate professor or full professor level is expected. Applicants must be eligible to work in the United States.

Candidates must have a strong record of excellence in teaching, publishing, and funded research. The Brady Chair is a leadership position intended to enhance the reputation of the department and university. Duties will include faculty mentoring, innovative curriculum development, and the recruitment of quality students.

Applicants should submit a letter of application, current vita, copies of transcripts, a creative portfolio outlining how the objectives of the chair are to be met, and at least three letters of recommendation.

These materials should be sent to: Search Committee for Endowed Chair in Computer Science, Department of Mathematics and Computer Science, P.O. Box 2026, Nicholls State University, Thibodaux, LA 70310.

Nicholls State University is an EOE/AA/ADA employer. Women and ethnic minorities are encouraged to apply. Please visit <http://www.nicholls.edu/about> for more information about the university and the community.

distinguished record of accomplishments in computer-science education and research. Demonstrated administrative experience is desired, as is teaching experience at both the undergraduate and graduate levels. The successful candidate must qualify for the rank of professor.

Please provide a letter of application outlining your experience and vision for this position, a curriculum vita, and names and contact information of at least three professional references. Application materials should be submitted online at

<http://www.jobs.msstate.edu/>

Screening of candidates will begin February 22, 2008 and will continue until the position is filled. Mississippi State University is an AA/EOE institution. Qualified minorities, women, and people with disabilities are encouraged to apply.

**Mississippi State University**  
**Faculty Position in Computational**  
**Biology**

The Department of Computer Science and Engineering (<http://www.cse.msstate.edu>) invites applications for a tenure-track faculty position at the Assistant or Associate Professor level in the area of computational biology. Candidates for the position are expected to hold a Ph.D. in computer science, computational biology, or a closely related field (ABDs may be considered). The person filling this position will be affiliated with the MSU Institute for Digital Biology (IDB), a university-level multi-disciplinary research institute that merges MSU's strengths in engineering and biology to solve problems related to health, nutrition, biofuels, food safety, bio-security and agriculture. IDB faculty are currently funded by NSF, USDA, DOD, NIH, and DOE. MSU has a state-of-the-art facility for genomics and proteomics (the Life Science and Biotechnology Institute). Research areas within the department include high performance computing, artificial intelligence, graphics and visualization, computer security, and software engineering.

Mississippi State University is the largest university in the State of Mississippi with approximately 1000 faculty and 17,000 students. The Department of Computer Science and Engineering has 18 tenured and tenure-track faculty and offers academic programs leading to the B.S., M.S. and Ph.D. in Computer Science, and a B.S. in Software Engineering. The department offers a certificate in Computational Biology and jointly administers the B.S. in Computer Engineering. Faculty also direct Ph.D. students in Computational Engineering and Computer Engineering and work with a number of additional on-campus research centers including the High Performance Computing Collaboratory, the Institute for Neurocognitive Science and

Technology, and the Sustainable Energy Center. Seven faculty members in the department have been recognized by NSF CAREER awards. Department research expenditures total around three million dollars per year. Mississippi State University is ranked 48th among 271 U.S. institutions in R&D expenditures in engineering by the National Science Foundation.

Please provide a letter of application, a curriculum vita, research and teaching statements, and names and contact information of at least three professional references. Application materials should be submitted online at <http://www.jobs.msstate.edu/>, or emailed to [office@cse.msstate.edu](mailto:office@cse.msstate.edu). Screening of candidates will begin February 28, 2008 and will continue until the position is filled. Mississippi State University is an AA/EOE institution. Qualified minorities, women, and people with disabilities are encouraged to apply.

**Northern Arizona University**  
**Department of Computer Science**

The Department of Computer Science invites applications for a tenure track position in Computer Science to begin August 2008. Although outstanding candidates in all areas of computer science will be considered, we are particularly interested in applicants with a background in Software Engineering or Computer Security. The position requires a doctorate in Computer Science or a closely related field. Candidates must be particularly committed to undergraduate education, and must be able to teach core Computer Science undergraduate courses as well as upper-division and graduate courses in their specialty area. Candidates must also have a strong commitment to developing a scholarly record. Please see <http://www.cens.nau.edu/> for further information. To apply, send a cover letter identifying the position of interest, resume, and the names of three references to: CS Search Committee Chair, College of Engineering and Natural Sciences, Northern Arizona University, PO Box 15600, Flagstaff, AZ 86011-5600. The search will remain open until the position is filled or closed; however, the screening committee will begin reviewing applications on January 15, 2008. NAU is one of three state universities in Arizona and is an Equal Opportunity/Affirmative Action Employer that is responsive to the needs of dual-career couples. We are committed to a diverse faculty and seek individuals with a commitment to teaching and working in diverse communities. Minorities, women, veterans and persons with disabilities are encouraged to apply.

**Southern Illinois University**  
**Edwardsville**  
**Department of Computer Science**

The Department of Computer Science at

## Career Opportunities

Southern Illinois University Edwardsville invites applications for one tenure-track position at the Assistant Professor level beginning in August 2008. Applicants are required to have a doctorate in Computer Science or a related field.

Located twenty miles from downtown St. Louis, SIUE has an enrollment of over 13,000 students. At SIUE, good teaching is of primary importance and research leading to publication is required for promotion and tenure.

Part of the SIUE School of Engineering, the Department offers an ABET-accredited B.S. degree program, as well as B.A. and M.S. degree programs. The twelve full-time faculty and nearly 200 majors employ labs for software engineering, HCI, robotics, virtual reality, and networking, all maintained by the department.

To apply, send a letter of application, résumé, transcripts, and at least three letters of reference to: Computer Science Search Committee, SIUE Campus Box 1656, Edwardsville, IL 62026-1656. Review of applications will begin on February 25, 2008.

SIUE is an Equal Opportunity/Affirmative Action Employer. Applications from women and minorities are especially encouraged. For more details, see [www.cs.siue.edu/employment](http://www.cs.siue.edu/employment).

### Stanford University Computational Geosciences

Stanford University School of Earth Sciences invites applications for a senior tenure-track faculty appointment at either the Associate or Full Professor level in the area of computational geosciences. We welcome applicants with strong skills in computational theory and practice, as well as a working familiarity with numerical methods for large-scale problems, parallelization paradigms, and modern computer systems. The successful applicant will have research experience or interests in applications to Earth and environmental problems, including but not limited to one or more of the following areas: energy, water, other natural resources, atmospheres, oceans, fluid dynamics, geodynamics, geomechanics, hazards, seismology, electromagnetics, inversion, optimization, and imaging of the earth surface and interior. Experience with large geoscience datasets is desirable. Strong interest in or experience with research collaboration and teaching across earth and environmental science disciplines is highly desirable. We expect the successful applicant to lead development and growth of the Stanford Center for Computational Earth and Environmental Science. The Center, with its shared high productivity computing resources, seeks to expand research and educational opportunities in computational geosciences. This appointment will be with one (or jointly with two) of the four departments in the School of Earth Sciences: Energy Resources Engineering, Geological and Environmental Sciences,

Geophysics, and Environmental Earth System Science (proposed). Further information about the School of Earth Sciences and this search can be found at <http://pangea.stanford.edu/> Stanford University is an equal opportunity employer and is committed to increasing the diversity of its faculty. It welcomes nominations of and applications from women and minority groups, as well as others who would bring additional dimensions to the university's research, teaching and service missions. Please apply online (<http://pangea.stanford.edu/about/jobs.php>) with the following material (.pdf only): cover letter, curriculum vitae, a statement outlining research and teaching experience and interests, and the names and addresses of four referees. Address material to Computational Geosciences Search Committee. Applications received by January 31, 2008 will receive full consideration, though the position will remain open until the appropriate applicant is identified.

### 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 a tenure-track position at the Assistant/Associate Professor level beginning in Fall 2008. Salary and startup packages are competitive. We are especially interested in candidates with specialization in (a) Embedded Systems and Compilers or (b) Ubiquitous Computing/Information Access or (c) Information Security or (d) 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-tech companies such as IBM, Lockheed-Martin, BAE and Universal Instruments provide opportunities for collaboration. Binghamton borders the scenic Finger Lakes region of New York.

Submit a resume and the names of three references to the url address:

<http://binghamton.interviewexchange.com>

First consideration will be given to applications that are received by March 1, 2008. Applications will be considered until the positions are filled.

Binghamton University is an equal opportunity/affirmative action employer.

### Assistant Professor in Computer Forensics & Computer Security

The Department of Computer Engineering and Computer Science invites applications for a tenure-track position at the assistant professor level to begin July 1, 2008. A Ph.D. in computer science / computer engineering, or equivalent is required with research and teaching emphasis in the areas of computer forensics, computer security, information assurance, Internet forensics, and embedded systems. Candidates are expected to have a strong teaching and research background, and will participate in interdisciplinary activities. Screening of application materials will begin immediately. Details and application information can be found at [http://louisville.edu/speed/ccs/new\\_web/index.shtml](http://louisville.edu/speed/ccs/new_web/index.shtml)

*The University of Louisville is an equal opportunity institution and does not discriminate against persons on the basis of race, age, religion, sex, disability, color, sexual orientation, national origin or veteran status.*

### VASSAR COLLEGE Computer Science Department

Vassar College seeks applications for a Visiting Assistant Professor in the Computer Science Department starting in the 2008-2009 academic year. Vassar College is an equal opportunity/affirmative action employer and is actively committed to diversity within its community. Members of historically under-represented groups are especially encouraged to apply.

A commitment to excellence in undergraduate teaching and research is expected. The Ph.D. in computer science is required. Applicants with a background in any area of Computer Science will be considered. Interest and expertise in modeling and simulation, computational science, or databases is a plus. All candidates must be able to cover courses in the core areas of Computer Science.

Vassar College has been successfully building a strong undergraduate program in Computer Science. Introductory courses are taught using Java and C++. The department has two Linux laboratories for introductory and advanced instruction. Faculty are provided with Unix workstations and personal computers. For more information please go to the Computer Science department website at <http://www.cs.vassar.edu/>.

Review of applications will begin January 1, 2008 and continue until the position is filled. Send vita and three letters of reference to Nancy Ide, Chair, Department of Computer Science, Box 732, Vassar College, 124 Raymond Avenue, Poughkeepsie, New York 12604-0732. E-mail: [csdept@cs.vassar.edu](mailto:csdept@cs.vassar.edu).

### Texas A&M University Department of Visualization

Visualization Faculty - Multiple Positions: The Department of Visualization at Texas A&M University seeks tenured/tenure-track faculty. Candidates must demonstrate experience in collaborative efforts that cross-disciplinary lines. A terminal degree or equivalent industry experience is required. Responsibilities include research and/or creative work, teaching and advising at both the graduate and undergraduate levels, and service to the department, university, and the field. The position will be open until filled. For details, e-mail [search@viz.tamu.edu](mailto:search@viz.tamu.edu) or visit <http://www-viz.tamu.edu/announcements/facultysearch07>. Texas A&M University is an Affirmative Action/Equal Opportunity Employer committed to diversity.

### The Hong Kong Polytechnic University Department of Computing

The Department invites applications for Assistant Professors in most areas of Computing, including but not limited to Software Engineering / Biometrics / Digital Entertainment / MIS and Pervasive Computing. Applicants should have a PhD degree in Computing or closely related

fields, a strong commitment to excellence in teaching and research as well as a good research publication record. Initial appointment will be made on a fixed-term gratuity-bearing contract. Re-engagement thereafter is subject to mutual agreement. Remuneration package will be highly competitive. Applicants should state their current and expected salary in the application. Please submit your application via email to [hrstaff@polyu.edu.hk](mailto:hrstaff@polyu.edu.hk). Application forms can be downloaded from <http://www.polyu.edu.hk/hro/job.htm>. Recruitment will continue until the positions are filled. Details of the University's Personal Information Collection Statement for recruitment can be found at <http://www.polyu.edu.hk/hro/jobpics.htm>.

### University College London (UCL)

#### Department of Computer Science

The Department of Computer Science at University College London (UCL) seeks applications for an academic faculty position in the area of computer graphics and virtual environments, at the rank of Lecturer, Senior Lecturer, or Reader (the first equivalent to Assistant Professor and the latter two equivalent to Associate Professor in the US system), commensurate with qualifications and experience. We are looking for internationally outstanding researchers who also have an enthusiasm for teaching and communicating ideas. The Department has broad strength in computer science with a strong orientation to 'experimental computer science'. The Virtual Environments and Computer Graphics (VECG) group works in the broad area of virtual environments, in particular real-time global illumination, image-based rendering, 3D interaction, tracking, real-time systems and augmented reality. The group conducts fundamental graphics algorithms and systems research, and underpins this with user studies to evaluate the impact that technological developments have. Applicants must hold a PhD in Computer Science or a closely related field by the time they begin their appointment. They will be evaluated chiefly on the significance and novelty of their research to date, and their promise for leading a group in a fruitful programme of research. Further details about UCL CS, the post, and how to apply may be found at: <http://www.cs.ucl.ac.uk/vacancies>. All application materials must reach UCL by the 15th of February, 2008.

### University of California Berkeley

#### Engineering and Computer Science Department

The University of California Berkeley invites applications for several approved tenure-track

positions in Electrical Engineering and Computer Sciences at the Assistant Professor level, and one approved position at the Associate or Full Professor level, beginning Fall 2008, subject to budgetary approval. We also consider possible joint appointments with other Berkeley departments.

Applicants should have (or be about to receive) a Ph.D. in Computer Science, Electrical Engineering, Computer Engineering, or a related field, evidence of ability to establish and pursue a program of high quality research, and a strong commitment to graduate/undergraduate teaching. Prioritizing candidates' overall quality and promise over sub-area of specialization, we seek applicants interested in creating innovative and far-reaching solutions to important problems in electrical engineering and computer science. We also welcome applicants working in interdisciplinary areas such as computational biology, nanoelectronics, or the uses of computing in the interests of society.

Applications should include a resume, statements of research and teaching interests, selected publications, and the names of three references who will send recommendations. Review begins November 15, 2007; candidates are urged to apply by that date. The application period closes February 15, 2008, and applications received after that date will not be considered. To apply, go to URL: <http://www.eecs.berkeley.edu/Faculty-Jobs/>

If you do not have Internet access, you may mail your application materials to:

EECS Search Committee  
c/o Jean Richter, 253 Cory  
Hall, UC Berkeley, Berkeley,  
CA 94720-1770.

Online applications are strongly encouraged. Recommenders providing letters should submit them directly via the URL listed above by January 18, 2008. Reference letters are NOT requested directly by the department. Recommenders may view the UC Berkeley Statement of Confidentiality at <http://apo.chance.berkeley.edu/evaltr.html>.

University of California is an Equal Opportunity, Affirmative Action Employer.

### University of Central Oklahoma College of Mathematics and Science

The University of Central Oklahoma, College of Mathematics and Science, Department of Computer Science is seeking a full-time tenure track assistant professor for the Department of Computer Science. The successful candidate should have an earned Ph.D in Computer Science or closely related field. Applications are accepted online <http://jobs.ucok.edu/> Refer to position #A000335. Open until filled.



## Advertising in Career Opportunities

- **How to Submit a Classified Line Ad:** Send an e-mail to [kooney@acm.org](mailto:kooney@acm.org). Please include text, and indicate the issue/or issues where the ad will appear, and a contact name and number.
- **Estimates:** An insertion order will then be e-mailed back to you. The ad will be typeset according to CACM guidelines. NO PROOFS can be sent. Classified line ads are NOT commissionable.
- **Rates:** \$295.00 for six lines of text, 40 characters per line. \$80.00 for each additional three lines. The MINIMUM is six lines.
- **Deadlines:** Five weeks prior to publication date of the issue (which is the first of every month) Latest deadlines: <http://www.acm.org/publications>
- **Career Opportunities Online:** Classified and recruitment display ads receive a free duplicate listing on our website at <http://campus.acm.org/careercenter>
- Ads are listed for a period of 45-days
- **For More Information Contact:**  
WILLIAM KOONEY  
Account Manager  
at 212-626-0687 or  
[kooney@hq.acm.org](mailto:kooney@hq.acm.org)



# Career Opportunities

## **University of Kentucky** **Department of Computer Science** **Assistant Professor**

The University of Kentucky Computer Science Department invites applications for a tenure-track position beginning August 15, 2008 at the assistant professor level. Candidates should have a PhD in Computer Science. Review of credentials will begin on January 15, 2008, and will continue until a suitably qualified candidate is found.

We are interested in candidates with expertise in databases, data mining, knowledge discovery, information security, and bio/medical informatics. The position features opportunities for participation in the Center for Clinical and Translational Services, a newly established interdisciplinary effort involving 15 colleges on the University of Kentucky campus.

The University of Kentucky Department of Computer Science awards B.S., M.S., and Ph.D. degrees. The Department has 23 faculty members committed to excellence in education, research and service. It has about 200 undergraduate and 150 graduate students.

The Department has strong research programs in distributed computing, computer networks, computer vision and graphics, artificial intelligence, scientific computing, cryptography, and information-based complexity.

To apply, a UK Academic Profile must be submitted at [www.uky.edu/ukjobs](http://www.uky.edu/ukjobs) using job #SP520197. Questions should be directed to HR/Employment (phone 1-859-257-9555 press 2 or email [ukjobs@email.uky.edu](mailto:ukjobs@email.uky.edu)), or Diane Mier ([diane@cs.uky.edu](mailto:diane@cs.uky.edu)) in the Computer Science Department.

We are accepting applications now. The application deadline is March 1, 2008, but may be extended as needed. Upon offer of employment, successful applicants must undergo a national background check as required by University of Kentucky Human Resources.

The University of Kentucky is an equal opportunity employer and encourages applications from minorities and women.

## **University of New Mexico** **Department of Computer Science** **Requisition M48143**

The Scalable Systems Laboratory in the University of New Mexico Department of Computer Science seeks a highly-qualified researcher for kernel-level research, programming, and software development, as well as other applied computer science research based on current lab grant and contract needs. A well-qualified candidate will have significant experience with kernel-level systems programming in C, C++, and/or assembly language, network protocol design and implementation, and/or system software development for high-performance or cluster computing systems.

Significant technical writing experience and experience working in an academic research environment are desired but not required.

The UNM Scalable Systems Lab is an active computer science research group with projects in configurable operating systems and virtualization techniques, sensor networks, high-end I/O and monitoring systems, and large-scale communication systems. The laboratory comprises multiple faculty members, postdoctoral researchers, and graduate and undergraduate researchers, and actively collaborates with researchers at Sandia, Los Alamos, and Oak Ridge National Laboratories, Intel Corporation, IBM Research, and Sun Microcomputers. Laboratory facilities include a number of small and medium scale cluster and sensor systems, with access to large-scale systems available through laboratory industrial and government partners. Funding for laboratory research comes from a wide range of sources, including industrial partners, the National Science Foundation, U.S. Department of Energy, U.S. National Nuclear Security Administration, and the Defense Threat Reduction Agency.

Go to: <http://ejobs.unm.edu> to apply.

505-277-6947

EO/AA

## **University of New Mexico** **Department of Electrical and Computer Engineering (ECE)**

The Department of Electrical and Computer Engineering (ECE) at the University of New Mexico invites applications for a full-time faculty appointment in all areas of Computer Engineering at the Assistant Professor level in a status leading to a tenure decision within the department to begin in the Fall 2008. Exceptional candidates at the Associate or Full Professor level may be considered for hire with tenure.

- The position requires the development of a world-class research program, supervising graduate students, and teaching undergraduate and graduate courses in Electrical/Computer Engineering. Candidates should demonstrate the potential to establish a high-profile research program as well as have the necessary skills to be an effective teacher and mentor.

- The ECE Department has 32 tenure-track faculty, 3 endowed chairs, and 24 research faculty, and is currently undergoing a period of aggressive growth. Sponsored research expenditures for the 2006/07 academic year were over \$12M. The Department interacts extensively with the federal laboratories in New Mexico as well as national and local high-tech companies in the Albuquerque area.

- The Albuquerque area as a whole is experiencing tremendous growth, fueled by the state's investment in high-tech industries. The city was selected by Forbes Magazine (May

2006 issue) as one of the best places in the country to start a business.

- For application requirements and further information about the ECE Department, please view the complete position announcement on the department's website at [www.ece.unm.edu](http://www.ece.unm.edu)

- For best consideration, applications must be received by February 15, 2008.

The position will remain open until filled.

- The University of New Mexico, a Carnegie Doctoral Research institution with Very High Research Activity, is an equal opportunity/affirmative action employer and educator.

## **University of North Carolina at Chapel Hill**

### **Department of Computer Science**

Applications are invited for a tenure track faculty member at the level of Assistant Professor without tenure appointment in the Department of Computer Science to begin on or after July 1, 2008. We particularly invite applications from candidates with expertise in bioinformatics, computer security, databases, data mining, physically-based simulation, and robotics. Outstanding candidates who would strengthen or complement other existing research groups in the department will be welcomed.

Applicants are expected to have a doctorate in computer science or computer engineering, or must present an equivalent level of professional accomplishment. Information and details on the application procedure for this search is available at <http://www.cs.unc.edu/Admin/Faculty/Recruiting/>. Minorities and women are encouraged to apply. The University of North Carolina is an equal opportunity, affirmative action employer.

## **York University**


### **Department of Computer Science**

York University, Toronto: The Department of Computer Science and Engineering invites applications in Biomedical Engineering with specialization in medical devices and instrumentation at the Assistant or Associate Professor level in the tenure track stream. Details are available at <http://yorku.ca/acadjobs>. York University is an Affirmative Action Employer.

## **Louisiana State University in Shreveport**

### **Computer Science Department**

LSUS invites applications for 2 tenure track positions—Chair of the Computer Science Dept. and the Abe Sadoff Endowed Chair for Bioinformatics—starting August 2008. Additional information about these positions can be found at <http://www.lsus.edu/jobs> LSUS is an EOE/AA employer.



## Software Transparency and Purity

**M**any software programs contain unadvertised functions that upset users when they discover them. These functions are not bugs, but rather operations intended by their designers to be hidden from end users. The problem is not new—Trojan horses and Easter eggs were among the earliest instances—but it is increasingly common and a source of many risks. I define *software transparency* as a condition that all functions of software are disclosed to users. Transparency is necessary for proper risk management. The term “transparency” should be used instead of “fully disclosed” to avoid confusion with the “full disclosure” of vulnerabilities.

There is a higher standard to be named, because disclosure does not by itself remove objectionable functions. They pose risks while being irrelevant to the software’s stated purpose and utility, and are foreign to its advertised nature. Freedom from such functions is a property that needs a name: loyalty, nonperfidiousness, fidelity, and purity come to mind, but none of them seems exactly right. For the purposes of this column, I shall call it *purity*. “Pure Software” can theoretically exist without disclosure, but disclosure would be a strong incentive, as previously discussed by Garfinkel (see [www.technologyreview.com/Infotech/13556/?a=f](http://www.technologyreview.com/Infotech/13556/?a=f)). Purity does not mean free of errors or unchanged since release. It’s possible for pure software to contain errors or to be corrupted. The following examples illustrate some of the risks from opaque and impure software.

In 2004, the digital video recording (DVR) equipment maker TiVo was able to tell how many people had paused and rewound to watch Janet Jackson’s wardrobe malfunction in the televised Super Bowl. People could opt out of the data collection by making a phone call. The privacy policy, if it was read, did mention some data collection, but did not disclose its full extent and surprising detail. Very few would likely have opted-in to allow this foreign function.

Software purity as a desirable property is highlighted by some of the differences between the GNU Public License (GPL) v2 and v3. The changes can be viewed as intended to protect the capability to remove unwanted functionality from software, including firmware based on GPL code (for example, TiVo).

In 2005, the anti-cheating Warden software that was installed with the World of Warcraft online game was found to snoop inside computers. Some people love knowing it is there, whereas others find it distasteful but are unable to make a convincing argument that it is malicious spyware. Despite being authorized by the End-User License Agreement (EULA), it poses risks that were not made clear, through undisclosed, objectionable behaviors.

Also in 2005, copy prevention software unexpectedly present on Sony BMG CDs was installed surreptitiously when users attempted to play a CD on their computer. It was later recognized as a rootkit. Ironically, it was reused to attack the Warden.

In 2007, people who had paid for Major League Baseball videos from previous years found they were unable to watch them anymore because of a broken Digital Rights Management (DRM) system, because the server providing authorization was decommissioned without warning. Fragile DRM systems, such as those requiring an available server, are undesirable because of the risks they present while being foreign to the advertised features or content.

Also in 2007, Microsoft Live OneCare surreptitiously changed user settings when installed to enable automatic updates and re-enable Windows services that were disabled on purpose; this is documented obscurely. Whereas it was not malicious, it caused many problems to users and system administrators and was vehemently protested. Surreptitious functions pose risks, even if well intentioned.

Software transparency and purity are often valued but not explicitly identified. Beyond the obvious information security risks to users, opaque or impure software also poses business risks in the form of loss of reputation, trust, goodwill, sales, and contracts. It may be that transparency alone is enough for some purposes, and others may also require software purity. An explicit requirement of whichever is appropriate would decrease risks. **□**

---

**PASCAL MEUNIER** ([pmeunier@cerias.purdue.edu](mailto:pmeunier@cerias.purdue.edu)) is a research scientist at Purdue University. His teaching and research include computer security and information assurance.

---

DOI: 10.1145/1314215.1314232



Software Development  
Conference & Expo West

March 3-7, 2008

Santa Clara Convention Center  
Santa Clara, CA

**EARLY BIRD  
DISCOUNT**  
Register by FEB 8  
**SAVE UP  
TO \$300**

**STROUSTRUP &  
SUTTER ON C++**  
JOIN C++ CREATOR  
BJARNE STROUSTRUP  
AND RENOWNED C++  
EXPERT HERB SUTTER  
FOR AN IN-DEPTH,  
TWO-DAY SESSION.



**JOIN THE  
REVOLUTION**

FEATURING  
**OVER  
200**  
EDUCATIONAL  
SESSIONS

**PLUS:**

Keynotes, Expo, Birds-  
of-a-Feathers, Panels,  
Case Studies, Parties  
and other Special  
Events!

**EDUCATIONAL SESSIONS INCLUDE:**

LANGUAGES & PLATFORMS:

- C++
- JAVA PROGRAMMING
- .NET DEVELOPMENT
- RUBY

ESSENTIAL TECHNOLOGY:

- NEXT-GEN WEB
- WEB SERVICES/SOA
- IT FOR DEVELOPERS

SOFTWARE DEVELOPMENT LIFECYCLE:

- MODELING & DESIGN
- PEOPLE, PROCESS & METHODS
- REQUIREMENTS & ANALYSIS
- SECURITY
- TESTING & QUALITY
- BUSINESS OF SOFTWARE

**REGISTER TODAY AT [WWW.SDEXPO.COM](http://WWW.SDEXPO.COM)**

The 1st ACM SIGGRAPH Conference and Exhibition in Asia  
www.siggraph.org/asia2008

## New Horizons

### ACM SIGGRAPH launches the premiere SIGGRAPH Asia in Singapore

Programs include:

- Papers
- Sketches and Posters
- Courses
- Art Gallery
- Computer Animation Festival
- Educators Program
- Emerging Technologies
- Exhibition

***Calling all creative researchers, artists,  
digital innovators and animators!***

This is your opportunity to present your stellar work or attend the 1st ACM SIGGRAPH conference and Exhibition in Asia.

#### ***Queries?***

Contact Conference and Exhibitions  
Management at [asia2008@siggraph.org](mailto:asia2008@siggraph.org) or  
Tel: +65 6500 6700



SIGGRAPHASIA2008

Conference and Exhibition on Computer Graphics and Interactive Techniques  
**Singapore, 10-13 December 2008**

Held in  
**UNIQUELY  
Singapore**  
[visitsingapore.com](http://visitsingapore.com)