

COMMUNICATIONS

CACM.ACM.ORG

OF THE

ACM

01/2012 VOL.55 NO.1

Remembering John McCarthy

Computer Vision
Without Sight


I/O Virtualization

The State of
Phishing Attacks

Interfaces for the
Ordinary User

ACM's FY11
Annual Report

Law and Disorder



ACM Presents Its First-Ever, Free Learning Webinar:

“The Cloud in Your Hands— Marriage of Cloud Computing with Smart Devices”



Presented live on Thursday, January 26, 2 PM EST

Register at <http://learning.acm.org/webinar/current>

Innovations in wireless networking technology are driving our increasingly connected world, with mobile devices gaining acceptance for both professional and private use. The future promises a vast majority of devices relying on cloud services to enhance end user experiences—with services being a natural extension of such devices.

Designed for IT managers and developers, this session will provide an introduction to the exciting new world of cloud-enabled mobile computing. A few complex user scenarios possible with this new paradigm will be discussed, along with a hands-on tutorial for developing such mobile applications on Microsoft's Windows Phone platform.

What you'll learn:

- ↓ Why use cloud services with mobile devices?
- ↓ What user scenarios are possible by marrying cloud services with mobile devices?
- ↓ How do you develop cloud-enabled mobile apps?

Presenters:

- ↓ **Arjmand Samuel**, Senior Research Program Manager, Microsoft Research
- ↓ **Danny Dalal**, Senior Development Lead, Microsoft Research

Moderator:

- ↓ **David B. Johnson**, Rice University, Past President of ACM SIGMOBILE

*Space for this **free** webinar is limited. Register Today!*
Contact timanovsky@hq.acm.org for more information.



Association for
Computing Machinery

Advancing Computing as a Science & Profession



Frederick P.
BROOKS, JR.



Charles P. (Chuck)
THACKER



Charles W.
BACHMAN



Barbara
LISKOV



Richard M.
KARP



Leslie G.
VALIANT



Marvin
MINSKY



Adi
SHAMIR



Richard E.
STEARNS



John
HOPCROFT



Robert E.
KAHN



Vinton G.
CERF



Robert
TARJAN



Ivan
SUTHERLAND



Butler
LAMPSON



Juris
HARTMANIS



Andrew C.
YAO



Donald E.
KNUTH



Dana S.
SCOTT



Raj
REDDY



Fernando J.
CORBATÓ



Edward A.
FEIGENBAUM



Alan C.
KAY



William (Velvet)
KAHAN



Frances E.
ALLEN



E. Allen
EMERSON



Niklaus
WIRTH



Ken
THOMPSON



Leonard
ADLEMAN



Ronald
RIVEST



Edmund
CLARKE



Joseph
SIFAKIS

WHAT DO YOU GET WHEN YOU PUT

32 TURING AWARD WINNERS TOGETHER?

**THE ACM TURING
CENTENARY CELEBRATION**
A ONCE-IN-A-LIFETIME
EVENT THAT YOU'LL
NEVER FORGET...
AND YOU'RE INVITED!

CHAIR:

Vint Cerf, '04 Turing Award Winner,
Chief Internet Evangelist, Google

PROGRAM CO-CHAIRS:

Mike Schroeder, Microsoft Research

John Thomas, IBM Research

Moshe Vardi, Rice University

MODERATOR:

Paul Saffo, Managing Director of Foresight, Discern
Analytics and Consulting Associate Professor, Stanford

WHEN:

June 15th & June 16th 2012

WHERE:

Palace Hotel, San Francisco CA

HOW:

Register for the event at turing100.acm.org

REGISTRATION IS FREE-OF-CHARGE!



THE A.M. TURING AWARD, often referred to as the "Nobel Prize" in computing, was named in honor of Alan Mathison Turing (1912-1954), a British mathematician and computer scientist. He made fundamental advances in computer architecture, algorithms, formalization of computing, and artificial intelligence. Turing was also instrumental in British code-breaking work during World War II.

Departments

- 5 **Editor's Letter**
**Artificial Intelligence:
Past and Future**
By Moshe Y. Vardi
-
- 6 **Letters To The Editor**
Software Engineering Is Engineering
-
- 8 **ACM President's Letter**
By Alain Chesnais
-
- 9 **ACM's FY11 Annual Report**
-
- 14 **BLOG@CACM**
**Understanding CS1 Students;
Defective Software**
Mark Guzdial writes about why teachers must grasp introductory CS students' theories about computing. Bertrand Meyer argues for the necessity of analyzing large-scale software disasters and publishing a detailed technical study.
-
- 16 **CACM Online**
**eBooks Will Abound in the
ACM Digital Library**
-
- 43 **Calendar**
-
- 125 **Careers**

Last Byte

- 136 **Future Tense**
The Near Cloud
Wish I never pulled the plug...
By Daniel H. Wilson

News

- 17 **Better Medicine Through
Machine Learning**
Computers that tease out patterns from clinical data could improve patient diagnosis and care.
By Neil Savage
-
- 20 **Revamping Storage Performance**
Great strides are being made in finding fast alternatives to the slow disks that dominate storage systems, but fast media are not nearly enough.
By Gary Anthes
-
- 23 **Law and Disorder**
International law has always been a murky and Byzantine area. However, the Internet and digital technology have raised the stakes, the risks, and the challenges.
By Samuel Greengard
-
- 26 **Celebration Time**
The centennial celebrations of Alan Turing's birth might help turn a quiet British genius into an iconic global hero.
By Sarah Underwood
-
- 27 **Analyzing Apple Products**
Researchers untangle the complex web of Apple's global supply chain—and offer lessons for managers and policymakers trying to chart the future course of U.S. industry.
By Alex Wright
-
- 28 **John McCarthy, 1927–2011**
Winner of the 1971 A.M. Turing Award, John McCarthy was a founder of artificial intelligence and inventor of the Lisp programming language.
By Paul Hyman

Viewpoints

- 30 **Law and Technology**
**The Yin and Yang of Copyright
and Technology**
Examining the recurring conflicts between copyright and technology from piano rolls to domain-name filtering.
By Randal C. Picker
-
- 33 **The Business of Software**
The Difference Engine
Observations on cognitive diversity and team performance.
By Phillip G. Armour
-
- 35 **Historical Reflections**
**The IBM PC: From Beige Box
to Industry Standard**
Looking back at three decades of PC platform evolution.
By Thomas Haigh
-
- 38 **Viewpoint**
**Interfaces for the Ordinary User:
Can We Hide Too Much?**
Increasing the visibility and access to underlying file structure on consumer devices can vastly improve the user experience.
By Kai A. Olsen and Alessio Malizia
-
- 41 **Interview**
An Interview with Stephen A. Cook
Stephen A. Cook, winner of the 1982 A.M. Turing Award, reflects on his career.
By Philip L. Frana




Practice



- 48 **Creating Languages in Racket**
Sometimes you just have to make a better mousetrap.
By Matthew Flatt
-
- 57 **Bufferbloat: Dark Buffers in the Internet**
Networks without effective AQM may again be vulnerable to congestion collapse.
By Jim Gettys and Kathleen Nichols
-
- 66 **I/O Virtualization**
Decoupling a logical device from its physical implementation offers many compelling advantages.
By Carl Waldspurger and Mendel Rosenblum




Articles' development led by [acmqueue queue.acm.org](http://queue.acm.org)

**About the Cover:**


John McCarthy—creator of Lisp, pioneer of modern artificial intelligence research, and recipient of the 1971 ACM A.M. Turing Award among many other honors—is remembered by friends and colleagues on p. 28. Illustration by Josie Jammet.

Contributed Articles



- 74 **The State of Phishing Attacks**
Looking past the systems people use, they target the people using the systems.
By Jason Hong
-
- 82 **Flexible Experimentation in Wireless Sensor Networks**
Virtual testbeds model them by seamlessly integrating physical, simulated, and emulated sensor nodes and radios in real time.
By Geoff Coulson, Barry Porter, Ioannis Chatzigiannakis, Christos Koninis, Stefan Fischer, Dennis Pfisterer, Daniel Bimschas, Torsten Braun, Philipp Hurni, Markus Anwander, Gerald Wagenknecht, Sándor P. Fekete, Alexander Kröller, and Tobias Baumgartner
-
- 91 **Long-Term Confidentiality of PKI**
How to guarantee files encrypted and transmitted today stay confidential for years to come.
By Chi-Sung Lai, Shang-Ming Jen, and Chia-Yu Lu

Review Articles



- 96 **(Computer) Vision Without Sight**
Computer vision holds the key for the blind or visually impaired to explore the visual world.
By Roberto Manduchi and James Coughlan

Research Highlights

- 106 **Technical Perspective**
Where Do People Draw Lines?
By Frédo Durand
-
- 107 **Where Do People Draw Lines?**
By Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz
-
- 116 **Technical Perspective**
Content-Centric Networking
By Jim Kurose
-
- 117 **Networking Named Content**
By Van Jacobson, Diana K. Smetters, James D. Thornton, Michael Plass, Nick Briggs, and Rebecca Braynard



ACM, the world's largest educational and scientific computing society, delivers resources that advance computing as a science and profession. ACM provides the computing field's premier Digital Library and serves its members and the computing profession with leading-edge publications, conferences, and career resources.

Executive Director and CEO

John White

Deputy Executive Director and COO

Patricia Ryan

Director, Office of Information Systems

Wayne Graves

Director, Office of Financial Services

Russell Harris

Director, Office of SIG Services

Donna Cappel

Director, Office of Publications

Bernard Rous

Director, Office of Group Publishing

Scott E. Delman

ACM COUNCIL

President

Alain Chesnais

Vice-President

Barbara G. Ryder

Secretary/Treasurer

Alexander L. Wolf

Past President

Wendy Hall

Chair, SGB Board

Vicki Hanson

Co-Chairs, Publications Board

Ronald Boisvert and Jack Davidson

Members-at-Large

Vinton G. Cerf; Carlo Ghezzi;

Anthony Joseph; Mathai Joseph;

Kelly Lyons; Mary Lou Soffa; Salil Vadhan

SGB Council Representatives

G. Scott Owens; Andrew Sears;

Douglas Terry

BOARD CHAIRS

Education Board

Andrew McGettrick

Practitioners Board

Stephen Bourne

REGIONAL COUNCIL CHAIRS

ACM Europe Council

Fabrizio Gagliardi

ACM India Council

Anand S. Deshpande, PJ Narayanan

ACM China Council

Jianguang Sun

PUBLICATIONS BOARD

Co-Chairs

Ronald F. Boisvert; Jack Davidson

Board Members

Marie-Paule Cani; Nikil Dutt; Carol Hutchins;

Joseph A. Konstan; Ee-Peng Lim;

Catherine McGeoch; M. Tamer Ozsu;

Vincent Shen; Mary Lou Soffa

ACM U.S. Public Policy Office

Cameron Wilson, Director

1828 L Street, N.W., Suite 800

Washington, DC 20036 USA

T (202) 659-9711; F (202) 667-1066

Computer Science Teachers Association

Chris Stephenson,

Executive Director

COMMUNICATIONS OF THE ACM

Trusted insights for computing's leading professionals.

Communications of the ACM is the leading monthly print and online magazine for the computing and information technology fields. *Communications* is recognized as the most trusted and knowledgeable source of industry information for today's computing professional. *Communications* brings its readership in-depth coverage of emerging areas of computer science, new trends in information technology, and practical applications. Industry leaders use *Communications* as a platform to present and debate various technology implications, public policies, engineering challenges, and market trends. The prestige and unmatched reputation that *Communications of the ACM* enjoys today is built upon a 50-year commitment to high-quality editorial content and a steadfast dedication to advancing the arts, sciences, and applications of information technology.

STAFF

DIRECTOR OF GROUP PUBLISHING

Scott E. Delman

publisher@cacm.acm.org

Executive Editor

Diane Crawford

Managing Editor

Thomas E. Lambert

Senior Editor

Andrew Rosenbloom

Senior Editor/News

Jack Rosenberger

Web Editor

David Roman

Editorial Assistant

Zarina Strakhan

Rights and Permissions

Deborah Cotton

Art Director

Andrij Borys

Associate Art Director

Alicia Kubista

Assistant Art Directors

Mia Angelica Balaquiot

Brian Greenberg

Production Manager

Lynn D'Addesio

Director of Media Sales

Jennifer Ruzicka

Public Relations Coordinator

Virginia Gold

Publications Assistant

Emily Williams

Columnists

Alok Aggarwal; Phillip G. Armour;

Martin Campbell-Kelly;

Michael Cusumano; Peter J. Denning;

Shane Greenstein; Mark Guzdial;

Peter Harsha; Leah Hoffmann;

Mari Sako; Pamela Samuelson;

Gene Spafford; Cameron Wilson

CONTACT POINTS

Copyright permission

permissions@cacm.acm.org

Calendar items

calendar@cacm.acm.org

Change of address

acmhlp@cacm.acm.org

Letters to the Editor

letters@cacm.acm.org

WEB SITE

http://cacm.acm.org

AUTHOR GUIDELINES

http://cacm.acm.org/guidelines

ACM ADVERTISING DEPARTMENT

2 Penn Plaza, Suite 701, New York, NY

10121-0701

T (212) 869-7440

F (212) 869-0481

Director of Media Sales

Jennifer Ruzicka

jen.ruzicka@hq.acm.org

Media Kit acmm mediasales@acm.org

Association for Computing Machinery (ACM)

2 Penn Plaza, Suite 701

New York, NY 10121-0701 USA

T (212) 869-7440; F (212) 869-0481

EDITORIAL BOARD

EDITOR-IN-CHIEF

Moshe Y. Vardi

elic@cacm.acm.org

NEWS

Co-chairs

Marc Najork and Prabhakar Raghavan

Board Members

Hsiao-Wuen Hon; Mei Kobayashi;

William Pulleyblank; Rajeev Rastogi;

Jeannette Wing

VIEWPOINTS

Co-chairs

Susanne E. Hambrusch; John Leslie King;

J Strother Moore

Board Members

P. Anandan; William Aspray; Stefan Bechtold;

Judith Bishop; Stuart I. Feldman;

Peter Freeman; Seymour Goodman;

Shane Greenstein; Mark Guzdial;

Richard Heeks; Rachelle Hollander;

Richard Ladner; Susan Landau;

Carlos Jose Pereira de Lucena;

Beng Chin Ooi; Loren Terveen



PRACTICE

Chair

Stephen Bourne

Board Members

Eric Allman; Charles Beeler; David J. Brown;

Bryan Cantrill; Terry Coatta; Stuart Feldman;

Benjamin Fried; Pat Hanrahan; Marshall Kirk

McKusick; Erik Meijer; George Neville-Nel;

Theo Schlossnagle; Jim Waldo

The Practice section of the CACM

Editorial Board also serves as

the Editorial Board of *acmqueue*.

CONTRIBUTED ARTICLES

Co-chairs

Al Aho and Georg Gottlob

Board Members

Robert Austin; Yannis Bakos; Elisa Bertino;

Gilles Brassard; Kim Bruce; Alan Bundy;

Peter Buneman; Andrew Chien;

Peter Druschel; Blake Ives; James Larus;

Igor Markov; Gail C. Murphy; Shree Nayar;

Bernhard Nebel; Lionel M. Ni;

Sriram Rajamani; Marie-Christine Rousset;

Avi Rubin; Krishan Sabnani;

Fred B. Schneider; Abigail Sellen;

Ron Shamir; Marc Snir; Larry Snyder;

Veda Storey; Manuela Veloso; Michael Vitale;

Wolfgang Wahlster; Hannes Werthner;

Andy Chi-Chih Yao

RESEARCH HIGHLIGHTS

Co-chairs

Stuart J. Russell and Gregory Morrisett

Board Members

Martin Abadi; Stuart K. Card; Jon Crowcroft;

Shafi Goldwasser; Monika Henzinger;

Maurice Herlihy; Dan Huttenlocher;

Norm Jouppi; Andrew B. Kahng;

Daphne Koller; Michael Reiter;

Mendel Rosenblum; Ronitt Rubinfeld;

David Salelin; Lawrence K. Saul;

Guy Steele, Jr.; Madhu Sudan;

Gerhard Weikum; Alexander L. Wolf;

Margaret H. Wright

WEB

Co-chairs

James Landay and Greg Linden

Board Members

Gene Golovchinsky; Marti Hearst;

Jason I. Hong; Jeff Johnson; Wendy E. MacKay



ACM Copyright Notice

Copyright © 2012 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 permissions@acm.org or fax

(212) 869-0481.

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; www.copyright.com.

Subscriptions

An annual subscription cost is included

in ACM member dues of \$99 (\$40 of

which is allocated to a subscription to

Communications); for students, cost

is included in \$42 dues (\$20 of which

is allocated to a *Communications*

subscription). A nonmember annual

subscription is \$100.

ACM Media Advertising Policy

Communications of the ACM and other

ACM Media publications accept advertising

in both print and electronic formats. All

advertising in ACM Media publications is

at the discretion of ACM and is intended

to provide financial support for the various

activities and services for ACM members.

Current Advertising Rates can be found

by visiting <http://www.acm-media.org> or

by contacting ACM Media Sales at

(212) 626-0686.

Single Copies

Single copies of *Communications of the*

ACM are available for purchase. Please

contact acmhlp@acm.org.

COMMUNICATIONS OF THE ACM

(ISSN 0001-0782) is published monthly

by ACM Media, 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



Association for
Computing Machinery





DOI:10.1145/2063176.2063177

Moshe Y. Vardi

Artificial Intelligence: Past and Future

Chess fans remember many dramatic chess matches in the 20th century. I recall being transfixed by the 1972 interminable match between challenger Bobby Fischer and

defending champion Boris Spassky for the World Chess Championship. The most dramatic chess match of the 20th century was, in my opinion, the May 1997 rematch between the IBM super-computer Deep Blue and world champion Garry Kasparov, which Deep Blue won $3\frac{1}{2} - 2\frac{1}{2}$.

I was invited by IBM to attend the rematch. I flew to New York City to watch the first game, which Kasparov won. I was swayed by Kasparov's confidence and decided to go back to Houston, missing the dramatic second game, in which Kasparov lost—both the game and his confidence.

While this victory of machine over man was considered by many a triumph for artificial intelligence (AI), John McCarthy (Sept. 4, 1927–Oct. 24, 2011), who not only was one of the founding pioneers of AI but also coined the very name of the field, was rather dismissive of this accomplishment. “The fixation of most computer chess work on success in tournament play has come at scientific cost,” he argued. McCarthy was disappointed by the fact that the key to Deep Blue's success was its sheer compute power rather than a deep understanding, exhibited by expert chess players, of the game itself.

AI's next major milestone occurred last February with IBM's Watson program winning a “Jeopardy!” match against Brad Rutter, the biggest all-time money winner, and Ken Jennings, the record holder for the longest championship streak. This achievement was also

dismissed by some. “Watson doesn't know it won on “Jeopardy!,” argued the philosopher John Searle, asserting that “IBM invented an ingenious program, not a computer that can think.”

In fact, AI has been controversial from its early days. Many of its early pioneers overpromised. “Machines will be capable, within 20 years, of doing any work a man can do,” wrote Herbert Simon in 1965. At the same time, AI's accomplishments tended to be underappreciated. “As soon as it works, no one calls it AI anymore,” complained McCarthy. Yet it is recent worries about AI that indicate, I believe, how far AI has come.

In April 2000, Bill Joy, the technologists' technologist, wrote a “heretic” article entitled “Why the Future Doesn't Need Us” for *Wired* magazine, “Our most powerful 21st-century technologies—robotics, genetic engineering, and nanotech—are threatening to make humans an endangered species,” he wrote. Joy's article was mostly ignored, but in August 2011 Jaron Lanier, another widely respected technologist, wrote about the impact of AI on the job market. In the not-too-far future, he predicted, it would just be inconceivable to put a person behind the wheel of a truck or a cab. “What do all those people do?” he asked.

Slate magazine ran a series of articles in September 2011 titled “Will Robots Steal Your Job?” According to writer Farhad Manjoo, who detailed the many jobs we can expect to see taken over by computers and robots in the

coming years, “You're highly educated. You make a lot of money. You should still be afraid.”

In fact, worries about the impact of technology on the job market are not only about the far, but also the not too far future. In a recent book, *Race Against The Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, by Erik Brynjolfsson and Andrew McAfee, the authors argue that “technological progress is accelerating innovation even as it leaves many types of workers behind.” Indeed, over the past 30 years, as we saw the personal computer morph into tablets, smartphones, and cloud computing, we also saw income inequality grow worldwide. While the loss of millions of jobs over the past few years has been attributed to the Great Recession, whose end is not yet in sight, it now seems that technology-driven productivity growth is at least a major factor.

The fundamental question, I believe, is whether Herbert Simon was right, even if his timing was off, when he said “Machines will be capable ... of doing any work a man can do.” While AI has been proven to be much more difficult than early pioneers believed, its inexorable progress over the past 50 years suggests that Simon may have been right. Bill Joy's question, therefore, deserves not to be ignored. *Does the future need us?*

Moshe Y. Vardi, EDITOR-IN-CHIEF

Software Engineering *Is* Engineering

I WAS CONFOUNDED by the conclusion of Michael Davis's Viewpoint "Will Software Engineering Ever Be Engineering?" (Nov. 2011), mainly because anything I can do in code I can also do in digital hardware, analog hardware, fluidics, even gears and motors. One could say "Maintenance is a nightmare," but there is ample proof that software maintenance can also be a nightmare with even modest code. By drawing attention to "[m]oderate the interests of the software engineer, the employer, the client, and the users for the public good," Davis exposed an unfortunate turn of phrase that is easily misconstrued. However, the following from the ACM's own "Software Engineering Code of Ethics and Professional Practice" (version 5.2 <http://www.acm.org/about/se-code>) states the principle clearly: "Software engineers shall act consistent with the public interest," which also maps well to the National Society of Professional Engineers' "Code of Ethics for Engineers" (<http://www.nspe.org/Ethics/CodeofEthics/index.html>): "Hold paramount the safety, health, and welfare of the public."

Software engineering is not computer programming. Programming is a constructive activity, like soldering, riveting, and pouring concrete, where excellence requires special knowledge, skill, and ability. The fact that many engineers strive to be competent technicians should not confuse the issue; engineering work and technical work are simply different types of work, often done by the same person.

The overarching engineering paradigm—systems engineering—recruits specialty engineering disciplines, including hardware (such as electrical and mechanical), software (such as applications and embedded), and human factors (such as micro-ergonomics and macro-ergonomics). All are used, knowingly or unknowingly, in the formulation and translation of engineering requirements (design inputs) into engineering specifications (design

outputs), enabling construction of systems with utilitarian and/or aesthetic value. Ideally, systems engineering occurs in a structured, systematic manner with incremental verifications and the physical realization validated against the original requirements. Haphazard design and construction may be "engineering" but is indeed unprofessional and provides no rational means for ensuring the protection of the public's safety, health, and welfare.

It may be reasonable to say developing banking and retail software systems do not require training in advanced mathematics, physics, chemistry, biology, or even social science; the same cannot be said of software systems for nuclear weapons, transportation systems, or medical devices. Designing and programming banking and retail software is, perhaps, more like bookkeeping and accounting practices and thus justifies being taught in business school. Software may be designed by engineers or by accountants, but only the former should be deemed software engineering.

The software development I have been involved in for 50 years has always required my engineering training and experience. My basic training was electrical engineering, not computer science, so, perhaps, my projects are somewhat skewed in this direction.

GM Samaras, Pueblo, CO

Author's Response:

Samaras responded to my quotation of Rule 1.2 of the "Software Engineering Code of Ethics" ("moderate") with a quotation from Principle 1 from the same code ("consistent with the public interest"). He then claimed Principle 1 "maps well" to the paramount clause now common in engineering codes of ethics. Not so. "Hold paramount" means put above other considerations, a somewhat higher standard than "consistent with." He also ignored most of my argument, especially on differences in curriculum between engineering proper and software engineering. His focus is solely on function.

Michael Davis, Chicago, IL

LINQ Limits vs. Computational Reality

As a developer of information systems for more than 15 years, I can say the Language-integrated Query, or LINQ, framework component championed by Erik Meijer in "The World According to LINQ" (Oct. 2011) produces impressive abstractions but also that these abstractions hide many real problems and that projects that rely heavily on LINQ risk hitting a performance wall. The mathematical theory of monads, or computational structures in functional programming, is well established and useful, but today's information systems are not mathematical constructs but real physical systems with real physical and fiscal constraints.

While processor speed and memory capacity are improving, the speed of light will always mean at least 60ms for a signal roundtrip between Tokyo and New York. No matter the location of the data source, it will not be inside a processor and will require significant time to perform selections, projections, and cross-applies. The danger is programmers working with abstractions make it easy to create huge computational loads with simple statements, with systems ultimately paying in terms of performance and operational costs. Needing a more general abstraction or language for extracting and evaluating large, diverse data sources is not the real problem in application programming.

I don't deny that the variety of data processing concepts is indeed a problem but criticize LINQ's (and Meijer's) focus on language. Many languages and formalisms have the power to express how to get the information an application requires. I wish there were fewer of them so they could be used more easily in many different domains, without programmers having to learn and maintain so many.

LINQ is not the solution system architects have been waiting for; it just hides the most challenging problems,

moving them out of sight but not out of existence. A system must still answer a request, and the fewer instructions programmers must provide to the system about what an application needs (and where to find it, along with quantities to accept and cost in terms of response time and resource consumption the application tolerates), the less the system will provide what the sender of the request has asked for.

In real systems, a response to a request is rarely only a collection of information but much more likely a set of desired behavior in time, cost, size, and quality.

Markus Kropf, Thun, Switzerland

Author's Response:

Abstraction means hiding unnecessary details, and most class libraries and frameworks aim to hide specific details to simplify programmers' lives. However, if these details are in fact relevant to an application, then programmers should not use the abstraction or hope the abstraction is leaky (actually a good thing) such that they can dive under the hood in case they need to deal with lower-level concerns.

Erik Meijer, Redmond, WA

Reinstate Jahromi

This letter updates the Letter to the Editor "Justice for Jahromi" (Nov. 2011) where I reported Masaud Jahromi had been imprisoned in Bahrain for months, beginning last April, and not permitted access to a lawyer or formally charged as to his crime and subsequent imprisonment.

The Jahromi family, which is aware of the letter, says it greatly appreciates ACM's support and concern.

The new information I have received is both encouraging and of grave concern. For one thing, Jahromi's exact title has been clarified. He is a professor of telecommunication engineering and chairman of the Telecommunication Engineering Department of Ahlia University in the Kingdom of Bahrain. The encouraging news is he was released on bail September 12, 2011, though it came five months after his arrest. On August 24, he was finally informed of the single charge against him: "participation in an unauthorized rally." The

concern is that, following the arrest, false rumors were circulated that destroyed his social and academic prestige at Ahlia University and led to his suspension from his positions there. His lawyer says there is no official order to suspend Jahromi, and he has the right to be back on duty until the day the court issues its verdict. However, the university has apparently rejected Jahromi's request.

The family is apprehensive about Jahromi's status and understands it depends on the final verdict and that in the interim it is likely to have no income. That status is likewise likely to remain precarious. Since his release, three court sessions have been held, with all subsequently adjourned to later dates. He is now banned from leaving Bahrain, with the next trial session scheduled for December 22. He has been notified that court appearance will be the final pleading session and the judge will appoint a later date to announce the verdict.

Jahromi's arrest for exercising his right of free speech and subsequent lack of due process, including denial of access to a lawyer, is in violation of the International Covenant on Civil and Political Rights, to which Bahrain has acceded.

Jahromi appears to have committed no crime and should be presumed innocent until proven otherwise. We all have a stake in maintaining academic freedoms, including of expres-

sion, at institutions of higher learning worldwide. I urge us all to write to the president of Ahlia University to reconsider the decision to suspend Jahromi from his positions and reinstate him at once, pending the court's final decision:

Professor Abdullah Y. Al-Hawaj
President Ahlia University
P.O. Box 10878
First Floor, Gosi Complex,
Exhibitions Road
Manama, Kingdom of Bahrain
email: aalhawaj@ahliauniversity.edu.bh
fax: +973 17 290083

Jack Minker, College Park, MD

Correction: Bouyer Figures Found

Two in-text figures on pages 80 and 81 were inadvertently dropped from the print version of Patricia Bouyer et al.'s Review Article "Quantitative Analysis of Real-Time Systems Using Priced Timed Automata" (Sept. 2011). The complete version is now available through the ACM Digital Library (<http://dl.acm.org/citation.cfm?id=1995396&CFID=67421339&CFTOKEN=67074515>) and *Communications* Web site (<http://www.cacm.acm.org>). —Ed.

Communications welcomes your opinion. To submit a Letter to the Editor, please limit yourself to 500 words or less, and send to letters@cacm.acm.org.

© 2012 ACM 0001-0782/12/01 \$10.00

Coming Next Month in COMMUNICATIONS

Text-Mining the Voices of the People

Globus Online: Radical Simplification of Data Movement via SaaS

Programming by Optimization: Toward a New Paradigm for Developing High-Performance Software

Progress and Challenges in Intelligent Vehicle-Area Networks

What Have We Learned from Software Engineering?

Incentive Auctions

Emotion and Security

Peer Instruction: A Teaching Method to Foster Deep Understanding

Yet Another Technology Cusp: Confusion, Vendor Wars, and Opportunities

Plus the latest news on botnets, the science of science, and cooperative robots.



DOI:10.1145/2063176.2063206

Alain Chesnais

ACM President's Letter

Early in 2011, IEEE Computer Society (IEEE-CS) President Sorel Reisman and I began discussing how IEEE-CS and ACM could work together more cooperatively. We shared that news with

our members in a "Letter from the Presidents" published in the August issues of both *Communications of the ACM* and *Computer*. We solicited suggestions from our members (<http://cooperation.acm.org>) and received some ideas that are worth pursuing for the benefit of the community, which I discuss here.

On September 22, 2011, ACM Chief Executive Officer John White, Chief Operating Officer Pat Ryan, and I met with IEEE-CS President Sorel Reisman and Executive Director Angela Burgess at ACM headquarters in New York City. The goal of this meeting was to build on the many areas in which ACM and IEEE-CS cooperate. Based on input from our members, we looked at: working together to support the efforts of the Computing in the Core Coalition; improving access to our respective digital libraries; reducing joint member fees; sharing speaker programs; and the possibility of merging the two organizations. In reading through the agreements reached for increased cooperation, it is important to remember that IEEE-CS is one (albeit the largest) of 38 societies of the IEEE. As a result, while some proposals for increasing cooperation can be dealt with directly by IEEE-CS, others require IEEE-CS engaging successfully with "parent" IEEE.

Computing in the Core is a non-partisan advocacy coalition established by ACM to elevate the national profile of computer science education in K-12 within the U.S. and work toward ensuring that computer science

is one of the core academic subjects in K-12 education. Several associations, corporations, scientific societies, and other non-profits who share the goal of seeing real computer science exist and count—particularly at the high school level—have joined the coalition. IEEE-CS has agreed to formally join Computing in the Core and work with ACM and the other members of the coalition on these important issues. At this meeting we signed a memo of understanding whereby IEEE-CS formally joined the Computing in the Core Coalition.

Improved Access to Digital Libraries. We agreed that an important goal would be to work toward seamless access to content for individuals with access (individual subscription or institutional subscription) to both the IEEE-CS's Digital Library (CSDL) and the ACM DL. Both ACM and IEEE-CS are eager to implement this feature; however, because of the relationship between CSDL and Xplore, it may be some time before we can start real discussions with IEEE and the Xplore staff.

Reduced Joint Member Fees. We agreed that we would institute a 20% discount for individuals who were both Members of ACM and Affiliate Members of IEEE-CS. The goal is that an individual should be able to go to one place to join/renew his/her membership at a rate of \$160 (based on current dues levels). ACM is ready to move on this agreement; IEEE-CS needs to hold discussions with IEEE in order to put this reduced joint membership rate into effect.

Shared Speakers Program. We agreed that we would share lists and promote each other's lectureship programs to appropriate subunits. I am pleased to say this is currently under way and that ACM Chapters will be able to request IEEE-CS Distinguished Lecturers and IEEE-CS Chapters will be able to request ACM Distinguished Speakers.

Highlighting Cooperation

It is important to remember that while ACM and IEEE-CS initiated these exploratory conversations last year, in fact, there are many areas where cooperation between the two organizations already exists. Indeed, the level of cooperation between ACM and IEEE-CS is higher than with any other organizations. That said, we felt it would be appropriate to highlight the full extent to which we are working together. We agreed to develop a joint Web page to keep members apprised of what the two societies are doing together (<http://www.acm.org/acm-ieeeecs-coop/>).

The initiatives noted here represent a good start on increasing collaboration and cooperation with IEEE-CS, and the ACM Council fully supports these efforts. In fact, the ACM Council feels that working on increased collaboration between the two societies is the right focus and the right next step to take. To help ensure continued progress, Sorel and I will establish an ACM/IEEE-CS Task Force on Cooperation to oversee the work we have started and to look for additional areas of potential cooperation.

Alain Chesnais, PRESIDENT



ACM membership reached another all-time high for the ninth consecutive year and much of that growth can be traced to the success of ACM's international initiatives.



DOI:10.1145/2063176.2063179

Alain Chesnais

ACM's Annual Report

FY11 was a defining year for ACM as the largest educational and scientific computing society in the world. Many of the initiatives we have set forth over the last few years

have taken root and we now see tangible evidence of their success.

It was a year that solidified ACM's global reach with established hubs in Europe, India, and China finishing their first year of operation with significant growth in membership, chapters, and conferences. Indeed, ACM membership reached another all-time high for the ninth consecutive year and much of that growth can be traced to the success of ACM's international initiatives.

It was also a year that marked great strides in ACM's commitment to improving the image and health of the computing discipline. Computer Science Education Week proved a sizeable success in promoting the importance of computer science education to a mass audience. ACM-W launched a presence (and conference) in India encouraging women from the region to explore the limitless opportunities of a career in computer science. ACM's Educational Policy Committee and Computer Science Teachers Association issued a landmark report on the vital need for the U.S. to create a strong K-12 computer science education program. And the Dot Diva campaign—the NSF-funded grant to ACM and WGBH-Boston to develop a new image for computing—launched a dynamic new Web site advocating computer science as a desirable career path.

And it was a year that introduced major enhancements to the “crown jewel” of Association—the ACM Digital Library—that now boasts in-

creased functionality, sophisticated improvements to the citation pages, and new browser capabilities.

I am particularly pleased to report that it was a year ACM and the IEEE-Computer Society began to explore myriad ways the two organizations might work together to capitalize on the healthy competition that has served our profession well over the decades (see my ACM President's Letter on page 8). And I look forward this year to further examining new ways for the Association to utilize social networking to add value to the membership community.

The following pages summarize some of the highlights of a truly eventful year in the life of ACM. As always, we are indebted to our devoted volunteers, members, and industry partnerships for constantly propelling us to do better.

Alain Chesnais, ACM PRESIDENT

ACM's Annual Report for FY11

ACM, the Association for Computing Machinery, is an international scientific and educational organization dedicated to advancing the arts, sciences, and applications of information technology.

Publications

The centerpiece of ACM Publications is the ACM Digital Library (DL) serving as the primary distribution mechanism for all the Association's publications as well as host to another 12 periodicals and a set of conference proceedings. The DL, now available at 2,650 institutions in 64 countries, boasts an estimated 1.5 million users worldwide. The result of this wide availability led to more than 16 million full-text downloads during FY11.

ACM is committed to increasing the scope of material available via the DL. Last year, over 31,000 full-text articles were added, bringing total DL holdings to 307,000 articles. ACM's *Guide to Computing Literature* is also integrated within the DL. More than 207,000 works were added to the bibliographic database in 2011, bringing the total *Guide* coverage to more than 1.67 million works.

A major reorganization and enhancement of the citation pages in the DL took place this year. Conference proceeding metadata has been enhanced, and there is a new browse capability that creates a view of conference events and summarized views across the entire archive of publications for a particular event.

As of year-end FY11 ACM is the publisher of 78 periodicals, including 40 journals and transactions, eight magazines, and 30 newsletters. During FY11, ACM added 504 conference and related workshop proceedings to its portfolio.

Two new journals were launched last year: *Transactions on Intelligent Systems* and *Technology and Transactions on Management Information Systems*. The ACM Publications Board also approved a new journal: *Transactions on Economics and Computation* to debut in FY12.

SIGAPP resumed publishing *Applied Computing Review* (ACR) in an electronic version only. The publication, which focuses on current trends in the field,

contains invited papers from world-renowned researchers and selected papers presented by prominent researchers attending the SIG's Symposium on Applied Computing.

Education

Running on Empty: The Failure to Teach K-12 Computer Science in the Digital Age, a report released by ACM's Education Policy Committee (EPC) and Computer Science Teachers Association (CSTA), found that approximately two-thirds of the U.S. have few computer science education standards for secondary school education. In addition, most states treat high school computer science courses as simply an elective and not part of a student's core education. This landmark report addressed the vital need to build a strong K-12 computer science education program in schools throughout the U.S. that today largely does not exist.

CSTA is a key partner in ACM's efforts to see real computer science count at the high school level. In FY11, CSTA membership hit an all-time high of 10,000. In addition, CSTA completed work on a major revision to the ACM/CSTA K-12 Model Curriculum.

A new coalition—Computing in the Core (CinC)—was established to engage the federal government to ensure computer science is included in STEM discussions as a distinct discipline. The coalition comprises ACM, CSTA, National Center for Women & Information Technology, ABI, Computing Research Association, Microsoft, Google, and SAS. Through its efforts a major piece of legislation—the Computer Science Education Act—was introduced in the U.S. House of Representatives that would provide federal funding to states to develop learning standards and teacher certification standards for computer science as well as offer funding to states to implement these plans.

The ACM-led CinC once again inspired the U.S. House of Representatives to declare the week of December 5–11, 2010 National Computer Science Education Week. Armed with an enhanced Web site, a larger set of supporters, and

a new model of engagement, the second CSEdWeek succeeded in reaching a vast audience, providing information and activities for students, educators, parents, and corporations advocating computer science education.

ACM and the National Science Foundation (NSF) released a report last April detailing the findings from a joint Strategic Summit on Computing Education Challenges at Community Colleges. The meeting was designed to determine why community colleges are not able to apply successfully for federal funding for computing programs. The recommendations reported in *Digitally Enhancing America's Community Colleges* are designed to spur new grant proposals, inventive initiatives, and dynamic partnership activities to infuse innovation into community college computing courses and programs.

ACM's Education Board and ACM India are joining forces to spearhead a vital initiative to determine how best to improve the quality of undergraduate computing education in India. This cooperative effort also includes key educators in India and is designed to support and strengthen computing education throughout the country.

Professional Development

Two new products for practitioners and managers were introduced in recent months: Tech Packs and Learning Paths were designed to provide a significant learning resource for emerging areas of computing not directly covered by an ACM SIG, conference, or publication. A Tech Pack comprises a set of fundamentally important articles on a particular topic that provide a context and comprehensive perspective on the subject. The first two Tech Packs focused on cloud computing and parallel computing; with 'Packs' on mobility, software as a service, security, and gaming in the works. Learning Paths aim to provide practical educational training tools focused on mastering essential programming languages and technologies. They were conceived as practical entry points into various technologies for software

developers, engineers, architects, and technical managers. The first ACM Learning Path covers Ruby. LPs slated for FY12 include Python and Java.

ACM Queue, the online practitioner's magazine spirited by the Professions Board, continues to grow a loyal following. The Web site received over one million page views and almost 500,000 readers in FY11, up from the previous year.

ACM's Committee on Professional Ethics (COPE) works to nurture an ethical stance by practitioners and those who employ practitioners. In the last year, COPE completed translations of the ACM Code of Ethics into major European languages as well as Arabic, with Hindi almost complete.

SIGUCCS' series of professional development Webinars continued last year with five new online offerings. The topic of staff management is particularly popular among the community and has been a major focus of these new Webinars.

Public Policy

ACM's U.S. Public Policy Council (USACM) educates policymakers in many areas of potential legislation, including bills on Internet monitoring, patent reform, e-voting, privacy and security, and mandatory breach notification. Last year USACM filed comments on numerous regulatory and legislative proposals. In addition, members testified before subcommittees of the U.S. House of Representatives as well as held several meetings with policymakers and staff from the White House, regulatory agencies, and Capitol Hill on a variety of technology- and privacy-related issues.

ACM's Committee on Computers and Public Policy assists ACM in a variety of internationally relevant issues pertaining to computers and public policy and helps make the Association more visible worldwide. Most notably, CCCP's ACM Forum on Risks to the Public in Computers and Related Systems has successfully achieved its intended goal to share and discuss the potential and serious computer-related risks with a global audience.

ACM's commitment to raising the visibility and addressing the issues of women in computing moved to the international arena in FY11, with more

direct links between ACM-W and the new regional councils in India and Europe. The result was a significant ACM-W effort in India last year where several ACM chapters have now been formed and a major ACM-W conference was held in conjunction with Amrita University.

Students

The 35th Annual ACM International Collegiate Programming Contest (ICPC) took place in Orlando, FL, with 105 teams competing in the final round. Earlier rounds of the competition included nearly 25,000 contestants representing 2,070 universities from 88 countries. Financial and systems support for ACM-ICPC is provided by IBM. The top four teams won gold medals as well as employment or internship offers from IBM.

The ACM Student Research Competition (SRC), sponsored by Microsoft Research, continues to offer a unique forum for undergraduate and graduate students to present their original research at well-known ACM-sponsored and co-sponsored conferences before a panel of judges and attendees. This year's SRC saw graduate and undergraduate winners compete against more than 50 participants in contests held at 13 ACM conferences. Their research covered a range of computing innovations that have applications for high-performance computer systems designs, image retrieval systems for astrophysicists, improvements in massively parallel graphics processors, biochips, and assistive technologies for speech-impaired people. SRC Grand Finals winners were recognized at ACM's Awards Banquet in San Jose, CA, last June.

ACM's *XRDS*, the student magazine formerly known as *Crossroads*, was relaunched at the beginning of FY11 with a new look and feel to the quarterly print edition and a contemporary new Web site (<http://xrds.acm.org/>) that provides information and opportunities aimed directly to students, including internships, fellowships, grants, and job postings.

At the SIGDA CADathlon, student teams competed in a set of design automation-related programming problems. Modeled after ICPC, the contest attracted over 20 teams.

ACM Council

PRESIDENT

Alain Chesnais

VICE PRESIDENT

Barbara G. Ryder

SECRETARY/TREASURER

Alexander L. Wolf

PAST PRESIDENT

Wendy Hall

SIG GOVERNING BOARD CHAIR

Vicki Hanson

PUBLICATIONS BOARD CO-CHAIRS

Ronald Boisvert, Jack Davidson

MEMBERS-AT-LARGE

Vinton G. Cerf

Carlo Ghezzi

Anthony Joseph

Kelly Lyons

Mary Lou Soffa

Salil Vadhan

SGB COUNCIL REPRESENTATIVES

Joseph A. Konstan

G. Scott Owens

Douglas Terry

ACM Headquarters

EXECUTIVE DIRECTOR/CEO

John R. White

DEPUTY EXECUTIVE DIRECTOR/COO

Patricia M. Ryan

2010 ACM Award Recipients

A.M. TURING AWARD

Leslie Valiant

ACM-INFOSYS FOUNDATION AWARD IN THE COMPUTING SCIENCES

Frans Kaashoek

ACM/AAAI ALLEN NEWELL AWARD

Takeo Kanade

THE 2011-2012 ACM-W

ATHENA LECTURER AWARD

Judith S. Olson

GRACE MURRAY HOPPER AWARD

Craig Gentry

ACM-IEEE CS 2011

ECKERT-MAUCHLY AWARD

Gurindar S. Sohi

Karl V. Karlstrom

OUTSTANDING EDUCATOR AWARD

Barbara Ericson and Mark Guzdial

OUTSTANDING CONTRIBUTION TO ACM AWARD

Joseph S. DeBlasi

DISTINGUISHED SERVICE AWARD

Reinhard Wilhelm

PARIS KANELAKIS THEORY

AND PRACTICE AWARD

Kurt Mehlhorn

SOFTWARE SYSTEM AWARD

GroupLens Collaborative Filtering

Recommender Systems

Peter Bergstrom

Lee R. Gordon

Jonathan L. Herlocker

Neophytos Iacovou

Joseph A. Konstan

Shyong (Tony) K. Lam

David Maltz

Sean McNee

Bradley N. Miller

Paul J. Resnick

John T. Riedl

Mitesh Suchak

ACM-IEEE CS KEN KENNEDY AWARD

David J. Kuck

DOCTORAL DISSERTATION AWARD

Bryan Jeffrey Parno

HONORABLE MENTION

Benjamin Snyder

Internationalization

FY11 marked the first full year of operation for ACM Europe and ACM India. Both councils were devoted to creating initiatives to further growth in chapters, membership, conferences, and other high-quality ACM activities. By year-end, both councils realized significant advances in all areas. In Europe, the number of professional conferences increased 75%, the number of chapters increased by 19%, and the number of European members increased by 2.3%. In India, the number of members increased 18% and chapters and conferences also saw significant jumps.

ACM China was launched in June 2010. Since that time, the number of members in China increased from slightly over 1,700 at the end of FY10 to almost 12,390 by FY11. Moreover, the number of conferences and chapters increased dramatically during its first year of operation.

Over 500 women attended the First Grace Hopper Conference in India in December 2010. This landmark event, in partnership with the Anita Borg Institute for Women and Technology, was held in Bangalore and featured leaders from industrial, academic, and government communities. Special sessions focused on the role of women in today's technology fields in India, including computer science, IT and research, and engineering.

ACM's Education Board continues to work toward reversing declining enrollments in computing disciplines worldwide. The Board's efforts have made serious inroads with their collaborative work with ACM Europe and ACM India to raise awareness of the professional opportunities afforded with a degree in computer science. The Board also creates initiatives to help foster a positive image of computing among young people.

One of the Association's major membership initiatives during FY11 involved increasing ACM's visibility in South America. The Association reached out to members of the Brazilian Computer Society (SBC) offering a six-month complimentary electronic-only ACM membership. Over 1,000 SBC members eagerly took advantage of this offer and ACM is now beginning the conversion process

Balance Sheet: June 30, 2011 (in Thousands)

ASSETS

Cash and cash equivalents	\$30,273
Investments	59,734
Accounts receivable and other assets	8,195
Deferred conference expenses	5,330
Fixed assets, net of accumulated depreciation and amortization	992
Total Assets	\$104,524

LIABILITIES AND NET ASSETS

Liabilities:	
Accounts payable, accrued expenses, and other liabilities	\$11,628
Unearned conference, membership, and subscription revenue	23,536
Total liabilities	\$35,164
Net assets:	
Unrestricted	63,158
Temporarily restricted	6,202
Total net assets	69,360
Total liabilities and net assets	\$104,524
Education Board accreditation	\$95
USACM Committee	20
Total expenses	\$115

to turn these trial members into full members.

Electronic Community

A new ACM Learning Center, with more user-friendly navigation, went live in early FY11. The customizable center is flexible in accommodating current and future educational and professional development tools and resources. The Learning Center's new "virtual campus" offers books from Safari Books Online and Books 24x7 as well as thousands of courses, virtual labs, and Brainbench exams from Element K for Professional and Student Members. Podcasts, videos, and other resources enhance the lifelong learning experience ACM strives to provide. The Association also built strategic partnerships with major academic institutions to bring exclusive discounts on courses and programs to ACM members, including courses from Stevens Institute of Technology and NYU-ePoly. ACM's new Tech Packs are also part of the Center; offering a wealth of

articles from the ACM DL and conference proceedings that cover a particular subject from all angles.

In the third and final year of the NSF grant to WGBH-Boston and ACM to develop an image for computing, the Dot Diva initiative unveiled a new Web site (<http://www.dotdiva.org/>) designed to expand its efforts in encouraging high school and college-aged girls to explore the vast potential for career and creative opportunities in the computer science field. The initiative was designed to build a positive image for high school girls, and to attract young women as part of a broad-based work force equipped with skills that are fundamental to ensuring future U.S. economic competitiveness and prosperity. The new site features stories, blogs, videos and resources designed to attract girls into the computing field at a critical time in their lives. These resources are tailored to appeal to girls as well as educators, computer science professionals, and parents.

Statement of Activities: Year ended June 30, 2011 (in Thousands)

REVENUE	Unrestricted	Temporarily Restricted	Total
Membership dues	\$8,938		\$8,938
Publications	18,275		18,275
Conferences and other meetings	26,117		26,117
Interests and dividends	1,650		1,650
Net appreciation of investments	5,791		5,791
Contributions and grants	3,634	\$1,351	4,985
Other revenue	232		232
Net assets released from restrictions	870	(870)	0
Total Revenue	65,507	481	65,988
EXPENSES			
Program:			
Membership processing and services	\$1,045		\$1,045
Publications	11,750		11,750
Conferences and other meetings	23,651		23,651
Program support and other	7,724		7,724
Total	44,170		44,170
Supporting services:			
General administration	9,569		9,569
Marketing	939		939
Total	10,508		10,508
Total expenses	54,678		54,678
Increase (decrease) in net assets	10,829	481	11,310
Net assets at the beginning of the year	52,329	5,721	58,050
Net assets at the end of the year	\$63,158*	\$6,202	\$69,360*

* Includes SIG Fund balance of \$30,814K

USACM transitioned to a new Web site with enhanced organization and presentation of information as well as established Twitter and Facebook pages for additional avenues for public policy communication.

The ACM History Committee commissioned a world-class Web site for the Turing Award that will include biographies and lectures from many of the past winners. Plans call for the site to be completed in time for ACM's Celebration of Turing's 100th birthday in June 2012.

The Computer Science Teacher's Association runs a comprehensive Web site for K–12 computing education that provides a wealth of resources and information, including teaching strategies, lesson plans, and curricu-

lum materials. This year, CSTA added podcasts, videos, and curriculum links to the site. Traffic statistics indicate this site to be of vital use for CS educators.

Many ACM SIGs introduced new and/or improved Web sites making significant use of social media to reach out to members. SIGITE unveiled a new site (<http://sigite.org>) that integrates the SIG's entry into social media including Twitter, Facebook, LinkedIn, and a blog. SIGMOBILE maintains an extensive Web site (<http://www.sigmobile.org>) that contains archived copies of most of the SIG's conference Web sites, including all previous years of MobiCom, MobiHoc, and MobiSys.

Conferences

The Grace Hopper Celebration, the world's largest gathering of women in computing, reported record-breaking attendance in Atlanta, GA. The annual event, designed to bring together the research and career interests of women in computing and highlight their accomplishments across industry, academia, and government, drew 2,147 attendees from 29 countries, including 940 undergraduate and graduate students.

The third annual SIGGRAPH Asia conference once again captured a wide spectrum of digital innovations. Over 7,600 professionals from 45 countries attended the conference and exhibition in Seoul, South Korea.

The first ACM International Conference on Bioinformatics and Computational Biology (ACM-BCB 2010) took place in Niagara Falls, NY. The successful event included workshops on protein interaction data, biological network analysis, gene networks, and immune informatics.

KDD-2010 attracted a record number of attendees, bested only by KDD-2000.

Recognition

There were 100 new chapters chartered in FY11. Of the 25 new professional chapters, 23 were internationally based; of the 75 new student chapters, 34 were international.

The ACM Fellows Program recognized 41 members for their contributions to computing and computer science in FY11. The new inductees brought the total number of ACM Fellows to 726.

ACM also named 47 of its members as Distinguished Members in recognition of their individual contributions to practical and theoretical aspects of computing that drive innovation and sustain economic competitiveness. Nearly 40% of the 2010 honorees are from non-North American countries, including France, Germany, Italy, China, Taiwan, Singapore, Australia, India, and South Africa.

The Distinguished Speakers Committee (DSP) expanded its reach in FY11; not only did lectures increase by 37%, many were held worldwide, including India, China, Pakistan, Taiwan, Israel, France, Germany, and the United Arab Emirates. 

The *Communications* Web site, <http://cacm.acm.org>, features more than a dozen bloggers in the BLOG@CACM community. In each issue of *Communications*, we'll publish selected posts or excerpts.



Follow us on Twitter at <http://twitter.com/blogCACM>

DOI:10.1145/2063176.2063180

<http://cacm.acm.org/blogs/blog-cacm>

Understanding CS1 Students; Defective Software

Mark Guzdial writes about why teachers must grasp introductory CS students' theories about computing. Bertrand Meyer argues for the necessity of analyzing large-scale software disasters and publishing a detailed technical study.



Mark Guzdial
"We're too Late for 'First' in CS1"

<http://cacm.acm.org/blogs/blog-cacm/102624>
December 7, 2010

We in computer science education have long argued about how to start the first course. "If they just see X first, they will understand everything in terms of X" where we might replace X with objects, functions, or recursion. We express concern about what will happen if they don't see the right stuff first. You may recall the Edsger W. Dijkstra quote, "It is practically impossible to teach good programming to students that have had a prior exposure to BASIC: As potential programmers, they are mentally mutilated beyond hope of regeneration."

I don't know if any of those beliefs about "what students need to know first" were once true, but I'm quite confident that none of them are true today. *There is no first.* There is no objects-

first, functions-first, hardware-first, or non-BASIC-first. No, it's not that I'm suggesting that there already is computer science in high school—there are way too few teachers and classes, on average, for that to be true. Rather, in a world where students live with Facebook, Wi-Fi, and email, they most certainly have theories about how computing works. There's no way to get to them before they make up their theories. By the time they decide to study computing, they have them.

Leigh Ann Sudol-DeLyser had a nice paper, "Mental Models of Data," at the 2009 Koli Calling conference, where she asked students to talk about the computing in their lives, then she tried to figure out which data structures they were already thinking about. For example, students already realize that a Facebook newsfeed has the newest information on top, and the oldest disappears off the bottom—sounds like these students already recognize a *queue*, even if they don't know the term yet.

I thought of this looking at Carsten Schulte's ITICSE 2010 paper, "Comparison of OOP First and OOP Later: First Results Regarding the Role of Comfort Level," on studying student perceptions of difficulty in an objects-first vs. an objects-later class. This is a follow-up to his earlier ICER paper, where he reported no learning differences between an objects-first and an objects-later class. While I've heard complaints about Schulte's analysis methods, I found his experimental set-up to be as careful as one could possibly be. The *only* differences between his two classes are the order of topics. Objects-first vs. objects-later doesn't matter, but neither does any other "first." His results are really not surprising. We already know that the sequence of topics in a curriculum rarely makes much difference in the learning. Students are really quite adept at getting by with less knowledge, and filling in the gaps as new information comes in.

It's an important open research question: How do students understand the computing around them? What theories do they have? They might not have any—until an error occurs. How they respond to that error suggests what kind of computational model they have. For example, watch a student do a Google or Bing search and then revise it to get better results. How did she revise it? What did she add to get a better result, and why did she think that *that* would work?

We know something about how novices develop naive theories of computing. John Pane's research shows us that

novices tend to develop declarative explanations of software that are driven by events, and don't seem to develop notions of objects at all. The "Common-sense Computing" group has shown us that novices can create algorithms for a variety of problems, although that doesn't really tell us how they think software and software development works in the world around them.

We are now in the same position as educators in physics (or biology, chemistry, or other sciences). Students have theories about how Wii controllers, voicemail menu systems driven by spoken voice commands, touch screens, and Google and Bing search work. *If these novice theories "mutilate" their minds, then it's done, it's happened to everyone, and we'd best just get on dealing with it.* There is no chance to place a theory in their minds before they learn anything else. We have to start from where the students are, and help them develop better theories that are more consistent and more correct. There is no *first*, but we can influence *next*.



Bertrand Meyer
"Again: The One Sure Way to Advance Software Engineering"
<http://cacm.acm.org/blogs/blog-cacm/101891>

January 13, 2011

Once again, bad software has struck. From 7:30 A.M. to late afternoon on November 10, 2010, Internet access and email were unavailable to most customers of Swisscom, the main mobile services provider in Switzerland. Given how wired our lives have become, such outages can have devastating consequences. As an example, customers of some of the largest banks in Switzerland cannot access their accounts online unless they type in an access code, one-time-pad style, sent to their cell-phone when they log in.

That is all the news we will see: Something really bad happened, and it was due to a software bug. A headline for a day or two, then nothing. What we will miss in this case as with almost all software disasters—most recently, the Great Pre-Christmas Skype Outage of 2010—is the *analysis*: what went wrong, why it went wrong, and what is being done to ensure it does not go

BERTRAND MEYER

"In Rahm Emanuel's immortal words, 'You never want a serious crisis to go to waste.'"

wrong again. Systematically applying such analysis is the most realistic technique available today for breakthrough improvements in software quality. The IT industry is stubbornly ignoring it. It is our responsibility as software engineering professionals to change that self-defeating attitude.

I have harped on this theme before^{1,2,3} and will continue to do so until the attitude changes. Quoting from the first reference:

Airplanes today are incomparably safer than 20, 30, 50 years ago: 0.05 deaths per billion kilometers. That's not by accident.

Rather, it's by accidents.

What has turned air travel from a game of chance into one of the safest modes of traveling is the relentless study of crashes and other mishaps. In the U.S. the National Transportation Safety Board has investigated more than 110,000 accidents since it began its operations in 1967. Any accident must, by law, be analyzed thoroughly; airplanes themselves carry the famous "black boxes" whose only purpose is to provide evidence in the case of a catastrophe. It is through this systematic and obligatory process of dissecting unsafe flights that the industry has made almost all flights safe.

Now consider software. No week passes without the announcement of some debacle due to "computers"—meaning, in most cases, bad software. The indispensable Risks Digest Forum⁴ and many pages around the Web collect software errors; several books have been devoted to the topic. A few accidents have been investigated thoroughly; two examples are Nancy Leveson's milestone study of the Therac-25 patient-killing medical device², and Gilles Kahn's analysis of the Ariane 5 crash, which Jean-Marc Jézéquel and

I used as a basis for our 1997 article⁶. Both studies improved our understanding of software engineering, but these are exceptions. Most of what we have elsewhere is made of hearsay and partial information, and plain urban legends—like the endlessly repeated story about the Venus probe that supposedly failed because a period was typed instead of a comma, most likely a canard.

Part of the solution is to use the legal system. For any large-scale software failure in which public money is involved, a law should require the convocation of an expert committee and the publication of a detailed technical analysis. The software engineering community should lobby for the passage of such a law and should not rest until it is enacted.

For private businesses the legal approach may be harder to pursue as some might view it as undue government interference, but it may still be pushed given the obvious public interest in software that works. The scenario would be for the industry to adopt, as a voluntary standard, the principle that every large-scale mishap must automatically lead to an exhaustive and public post-mortem analysis; in Rahm Emanuel's immortal words, "You never want a serious crisis to go to waste."

Until that happens, software will remain brittle. Think of the last time you stepped into a plane, and how different you would have felt if aircraft manufacturers had been allowed, disaster after disaster in the past 70 years, to keep the embarrassing details to themselves and continue business as usual. G

References

1. Bertrand Meyer, The one sure way to advance software engineering, <http://bertrandmeyer.com/2009/08/21/the-one-sure-way-to-advance-software-engineering/>, August 21, 2009.
2. Bertrand Meyer, Dwelling on the point, <http://bertrandmeyer.com/2009/11/29/dwelling-on-the-point/>, November 29, 2009.
3. Bertrand Meyer, Analyzing a software failure, <http://bertrandmeyer.com/2010/05/24/analyzing-a-software-failure/>, May 24, 2010.
4. Peter G. Neumann (moderator), The Risks Digest Forum on Risks to the Public in Computers and Related Systems, <http://catless.ncl.ac.uk/Risks/>.
5. Nancy Leveson, *Safeware: System Safety and Computers*, Addison-Wesley, 1995.
6. Jean-Marc Jézéquel and Bertrand Meyer: Design by contract: the lessons of Ariane, *Computer* 30, 1, January 1997.

Mark Guzdial is a professor at the Georgia Institute of Technology. Bertrand Meyer is a professor at ETH Zurich and ITMO (St. Petersburg) and chief architect of Eiffel Software.

© 2012 ACM 0001-0782/12/01 \$10.00



DOI:10.1145/2063176.2063181

Scott Delman

eBooks Will Abound in the ACM Digital Library

ACM is preparing to offer a new and exciting member benefit for 2012; *eBooks accessible via the ACM Digital Library!* As many ACM members are already aware, the Association has had agreements in place for several years with both Safari Books Online and Books 24x7 that enable current members to download selected computing books from those eBook platforms (approximately 1,200 titles in aggregate between the two separate collections). To date, tens of thousands of ACM members have taken advantage of this free benefit and this program is generally considered an important service to the ACM community.

In an effort to expand on the success of the existing eBooks Collections in the ACM Learning Center (see <http://learning.acm.org/books/index.cfm>), ACM will start to make available the full-text versions of approximately 150 computing eBook titles published by Morgan Kaufmann and Syngress directly via the ACM Digital Library platform. In addition to having full-text access to a growing and amazing collection of academic- and practitioner-oriented eBook titles, ACM members (students and professionals alike) will be able to login with their ACM Web Account credentials and have seamless access to this content directly from the ACM Digital Library. Those familiar with the Digital Library know the platform already indexes and abstracts many books from the computing field. Now, the full-text versions of many of these titles from MK and Syngress will be loaded in the DL for individual use.

Look for this new benefit in the first quarter of 2012 and happy reading!



ACM Member News

LEADING BY EXAMPLE



Inspiration does not often stem from a completely original idea, especially on the Web where

communities build upon existing programs, designs, and applications and find new or innovative ways to use them.

For several years, Scott Klemmer, associate professor of computer science and co-director of the Human-Computer Interaction Group at Stanford University, and colleagues have created tools that better leverage the availability of such online examples to support people's creative work.

"People learn by trying to understand the work of experts," says Klemmer, who was recently awarded the Katayanagi Emerging Leadership Prize, which honors a researcher who demonstrates leadership promise. "We're building software to allow people to gain greater access to these expert examples, and to more readily adapt them to new concepts. If you had all the examples at your fingertips, you'd find better ways to create layouts, forms, and new interaction techniques."

The rapid expansion of computer technology on a global level will only increase the demand for example tools. Applications and devices will need to quickly adapt to the wealth of diverse backgrounds and interests in the international user community, says Klemmer.

It is critical, he stresses, to give credit to the originators of the examples. "There are three reasons to do this," Klemmer says. "First, someone put in the hard work to make it happen, and deserves to be recognized for that. Also, the reader of the program should know where it came from, to understand why it was built that way. And if the example is updated, we can notify the borrowers to tell them that a new version is available."

—Dennis McCafferty


Better Medicine Through Machine Learning

Computers that tease out patterns from clinical data could improve patient diagnosis and care.

MEDICINE CAN BE AS much art as science, a detective story in which doctors rely not only on lab tests and x-rays, but on their own experience and clues from a patient's history to develop diagnoses or predict future health problems. But all of those lab tests, blood pressure readings, magnetic resonance imaging (MRI) scans, electrocardiograms, and billing codes add up to reams of data, which before too long will be joined by individual gene sequencing. Computer scientists are increasingly applying machine learning techniques to all that data, searching for patterns that can aid diagnosis and improve clinical care.

"Machine learning plays, I think, an essential role in medical image analysis nowadays," says Kenji Suzuki, assistant professor of radiology and medical physics at the University of Chicago's Comprehensive Cancer Center. Suzuki has been working on automating the detection of cancerous lesions in images from x-rays or computed tomography scans. Considering that radiologists may miss 12%–30% of lung cancers in such scans, a machine learning tool offers great potential.

Since the mid-1980s, computer sci-



Kenji Suzuki and colleagues' comparison of their rib-suppressed temporal-subtraction (TS) images with conventional TS images: (a) previous chest radiographs, (b) current chest radiographs of the same patient, (c) rib-suppressed TS images with fewer rib artifacts, and (d) conventional TS images.

entists have tried to improve on that performance using feature-based machine learning in which the computer would pick out morphological features, texture differences, and more to identify abnormal tissue. But such cataloging of features still misses some cancers that doctors are able to spot with their own eyes. Sometimes the features the computer is seeking can be subtle or overlap with normal anatomical structures such as bone, making them more difficult to spot. So Suzuki asks the machine to instead focus on the values, such as intensity, of individual pixels. “Because the computing power has increased dramatically in recent years, we can process the pixel values directly,” he says. The resulting system is highly sensitive, achieving up to 97% accuracy.

But one concern is making the program so sensitive that it starts finding nonexistent lesions. In Suzuki’s lung cancer tests, the feature-based algorithm falsely identified five lesions per patient, while the pixel-based method produced less than one false positive per patient and no false negatives. Suzuki says each method produces a different type of false positive, so combining the two approaches leads to the most accurate outcome. Suzuki is now working on expanding the technique to other types of cancer and other imaging methods, such as MRI and PET. “You just need to train the machine learning

As the use of electronic medical records gains acceptance, machine learning is likely to play an even larger role in clinical medicine.

technique with new images,” he says.

A Guide for Diagnosis

At IBM Almaden Research Center, the Advanced Analytics for Information Management (AALIM) project applies machine learning to a wide variety of data—readings of vital signs, tests such as echocardiograms, and demographic information—to chart the medical histories of patients over several years. By comparing the history of hundreds or thousands of people, the system can identify previous patients who are similar to a current patient, then apply collaborative filtering to suggest the best diagnosis and treatment options for the new patient.

“With a large number of pre-diagnosed patient datasets available in electronic health records, physicians can now benefit from the opinion of their peers on cases similar to their patients,” says Tanveer Syeda-Mahmood, head of IBM’s Multimodal Mining for Healthcare project. The hope is that by helping doctors base their decisions on quantitative information, the number of diagnostic errors can be reduced. AALIM, Syeda-Mahmood says, provides “a holistic view of the patient’s condition,” producing one-page summaries, long-term profiles of various measurements of health, and detailed comparisons showing diagnosis, treatment, and outcomes for similar patients.

The computer might, for instance, help a relatively new doctor decide that she needs to consult a specialist. Approximately 5% of cases, says Syeda-Mahmood, are ambiguous enough that even senior clinicians ask other doctors for their opinions, which AALIM easily provides. In emergency rooms, it can cut the time doctors spend flipping through charts by half. Although researchers have tested the system on patient data, it would likely require approval from the Food and Drug Administration before it could be used in a hospital.

Finding ways to use all this medical data often requires new developments in machine learning. For instance, Syeda-Mahmood wanted to give the

Education

Tech-Rich Learning Environments

What’s a solution to classrooms plagued by students’ low retention rates, high withdrawals, and failing grades? Toss out traditional lectures and create a technology-rich environment, say researchers at Rochester Institute of Technology (RIT).

Data from six years of research involving 500-plus undergraduates taking three different engineering courses revealed that 90% of the students said they learned and retained information better when a combination of tablet PCs, collaborative software, and multiple projection screens were used as teaching tools.

The study targeted three foundation courses for

engineering-technology degree programs—pneumatics and hydraulics, applied dynamics, and applied fluid mechanics—all of which suffered from higher-than-average student withdrawals and low retention rates. In one class, for example, 27% of the students had previously received low or failing grades and had to repeat the course or withdraw from it.

“The students were just not getting the material,” says Robert Garrick, associate professor at RIT’s College of Applied Science and Technology. “They weren’t understanding the intended learning objectives.”

In a traditional laboratory setting, students work at

different workstations while the instructor walks around, answers questions, and reviews circuit operations.

For the study, RIT researchers redesigned the courses to include a tablet PC and DyKnow collaborative software for each student. Classes took place in an interactive technology classroom that featured a multiscreen display connected to faculty and student tablet PCs. Class information and any notations were captured, recorded, and archived.

The students told the researchers they preferred the tablets for note taking, in-class work, test preparation, and classroom layout.

The combination of the technology resources improved the visual connection to the material, student-faculty interaction, and it enhanced the modeling of engineering problems, three areas seen as critical to retaining the technical information, explains RIT assistant professor Larry Villasmil.

The next step for the researchers is to study how tech-rich learning environments can work for underrepresented groups, such as deaf and hard-of-hearing students, who make up about 10% of the students in engineering and technology courses.

—Paul Hyman

computer a doctor's ability to recognize some types of heart disease by the characteristic shape of waves produced by an electrocardiogram. She developed a new function, called a constrained non-rigid translation transform, which could identify the similarity between shapes in different ECG readouts.

The Power of Prediction

Although diagnosis and treatment are key aspects of medical care, prediction is also important, especially when it can lead to early interventions. We all know, for instance, that factors such as weight and blood pressure can give an idea of a person's risk of heart disease. But those sorts of risk scores are based on population-wide models, says Shyam Visweswaran, assistant professor of biomedical informatics at the University of Pittsburgh. "If you build a model from a group of people who are kind of similar to the current patient, you might do better," he says. Visweswaran has developed an algorithm that lets a computer use clinical data to learn a model tailored to one specific patient and predict outcomes for that person.

The computer takes all the data it has on the patient, such as age, blood pressure, and lab results, and then picks one variable and builds a model of all the patients in its database who share that variable. It could, for instance, compare everyone in the 50–55 age group.

"With a large number of pre-diagnosed patient datasets available in electronic health records, physicians can now benefit from the opinion of their peers on cases similar to their patients," says Tanveer Syeda-Mahmood.

It builds a model for each variable it can find, looks at which ones best fit the patient at hand, and then averages the best models to make a personalized prediction of that patient's outcome. Whereas a population model only uses a handful of variables considered to be the best—it could be a simple checklist of several risk factors, for instance—this approach can potentially use any of hundreds of variables. One additional advantage is the machine might identify some factor that is predictive, but that medical science was not previously aware of, opening up new areas for research, Visweswaran says.

As with Suzuki's pixel-based processing, this is another machine learning method that has benefited from the growth of processing power. As recently as five years ago it might have taken a half-hour to build all these models. Today it takes less than a minute, so the machine can guide diagnoses in real time during patient visits.

This approach could help predict an intensive-care patient's risk of an infection spreading to other organs, which is a notoriously difficult task, and lead to earlier or more aggressive treatment. It might help doctors decide, for instance, which pneumonia patients need to be admitted to the hospital and which patients could be sent home with antibiotics.

In the informatics program at Children's Hospital Boston, assistant professor Ben Reis and his colleagues are working on predicting a patient's future diagnoses years in advance. They have developed Bayesian models that they call Intelligent Histories, which comb through the standard diagnostic codes used for billing, to find patterns in a patient's history that predict risk. In their first application of the work, they discovered they could identify patients at risk of domestic abuse as much as two years before the doctors seeing those patients first discovered the problem.

Doctors are supposed to screen patients for domestic abuse, but often miss it until the problem becomes acute, says Reis. Not only can the computer aid in screening for known signs of abuse, it also picked up other diagnostic codes in the test that had not been thought of as predictive, such as infections, which might teach doctors something about domestic abuse.

The Children's Hospital Boston team is working to expand their modeling to other types of diagnoses. At the same time, they are refining the machine learning itself. For instance, "we're trying to quantify how the quality of the data that goes into the model affects the results that come out," says Reis.

As the world moves to a greater use of electronic medical records, machine learning is likely to play an even larger role in clinical medicine, researchers predict. Visweswaran says genetic data, in particular, is going to require complicated computational models if it is going to be of value. Soon, experts expect, the cost of gene sequencing will drop to the point that individual genomes will become part of people's medical records, and will be available to the same data mining and pattern recognition approaches being applied to other data.

Genetic data will be too complicated and too voluminous to be handled with old-fashioned charting systems. "You have to have computational tools to query this data," Visweswaran says. "There's no way it can be done on paper." **G**

Further Reading

Gruhl, D., et al.

Aalim: A cardiac clinical decision support system powered by advanced multi-modal analytics, International Medical Informatics Conference, Cape Town, South Africa, Sept. 12–15, 2010.

Reis, B.Y., Kohane, I.S., and Mandl, K.D.

Longitudinal histories as predictors of future diagnoses of domestic abuse: modeling study, *British Medical Journal*, 339, Sept. 2009.

Syeda-Mahood, T., Beymer, D., and Wang, F.

Shape-based matching of ECG recordings, 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, August 23–27, 2007, Lyon, France.

Suzuki K., Zhang J., and Xu, J.

Massive-training artificial neural network coupled with Laplacian-eigen function-based dimensionality reduction for computer-aided detection of polyps in CT colonography, *IEEE Transactions on Medical Imaging* 29, 11, Nov. 2010.

Visweswaran, S., Angus, D.C., Hsieh, M., Weissfeld, L., Yealy, D., and Cooper, G.F.

Learning patient-specific predictive models from clinical data, *Journal of Biomedical Informatics* 43, 5, Oct. 2010.

Neil Savage is a science and technology writer based in Lowell, MA.

© 2012 ACM 0001-0782/12/01 \$10.00

Revamping Storage Performance

Great strides are being made in finding fast alternatives to the slow disks that dominate storage systems, but fast media are not nearly enough.


THE GROWTH in the power of computers, driven for decades by Moore's Law and more recently by increased core counts, coupled with huge improvements in network bandwidth, disk densities, and other metrics, is nothing short of astonishing. Yet practitioners at the high end say we are losing the battle against the data deluge.

"Our society is literally drowning in data," notes Allan Snively, associate director of the San Diego Supercomputer Center at the University of California at San Diego (UCSD). "Data acquisition devices ranging from space telescopes to genomic sequencing machines, and the Internet itself, are producing data almost faster than it can be written to disk."

Indeed, disks are a major part of the problem. Disk-based storage systems are 10 to 100 times slower than a network and thousands of times slower than main memory in delivering data to an application, in part because the data comes from a relatively slow electromechanical device. In recent years, fast flash memories have begun replacing disks in some applications. And manufacturers have started showing promising prototypes of more exotic non-volatile storage devices such as phase-change memory (PCM).

While substituting fast media for slow disks can help, it is by no means the whole solution. In fact, computer scientists at UCSD argue that new technologies such as PCM will hardly be worth developing for storage systems unless the hidden bottlenecks and faulty optimizations inherent in storage systems are eliminated.

A team at UCSD led by Steven Swanson, assistant professor of computer science and engineering, is doing just that. A recent prototype, called



A view of the internals of the Moneta storage array with phase change memory modules installed.

Moneta, bypasses a number of functions in the operating system (OS) that typically slow the flow of data to and from storage. These functions were

A phase-change memory prototype, called Moneta, bypasses a number of functions in the operating system that typically slow the flow of data to and from storage.

developed years ago to organize data on disk and manage input and output (I/O). The overhead introduced by them was so overshadowed by the inherent latency in a rotating disk that they seemed not to matter much. But with new technologies such as PCM, which are expected to approach dynamic random-access memory (DRAM) in speed, the delays stand in the way of the technologies' reaching their full potential. Linux, for example, takes 20,000 instructions to perform a simple I/O request.

Introducing Moneta

Moneta is a prototype high-performance storage array. It uses field-programmable gate arrays to implement a scheduler and a distributed set of memory controllers attached to conventional DRAM emulating PCM. (A similar prototype from the same researchers, called Onyx, uses PCM.)

The scheduler orchestrates Moneta's operations and is able to extract parallelism from some I/O requests. By redesigning the Linux I/O stack and by optimizing the hardware/software interface, researchers were able to reduce storage latency by 60% and increase bandwidth as much as 18 times.

The I/O scheduler in Linux performs various functions, such as assuring fair access to resources. Moneta bypasses the scheduler entirely, reducing overhead. Further gains come from removing all locks from the low-level driver, which block parallelism, by substituting more efficient mechanisms that do not.

"The OS gives you parallelism as an illusion," says Rajesh Gupta, a professor of computer science and engineering at UCSD. "It was designed to allow multiple users to support concurrency while actually doing shared access. But shared access and parallelism are two different things."

A further reduction in latency comes from bypassing an interrupt needed to wake up a thread that sleeps while waiting for completion of an I/O request. Instead, the thread does not sleep but spins in a busy loop.

In the hardware, I/O bandwidth is increased by providing separate queues for reads and for writes, and I/O is balanced by not processing big requests all at once while smaller ones wait.

As a result of these optimizations to software and hardware, Moneta performs I/O benchmarks 9.5 times faster than a RAID array of conventional disks, 2.8 times faster than a RAID array of flash-based solid-state drives (SSDs), and 2.2 times faster than Fusion-io's flash-based SSD, a high-end flash technology.

Actual speedups on large, data-intensive jobs could be even more dramatic, says Snavey. "For the right kind of random data access—such as examining a social network graph—flash is between 10 times and 100 times faster than spinning disk, and Moneta is 10 times faster than flash," he says. "That means some data-mining calculations that might take 24 hours today would only take 15 minutes or even less."

But, cautions Swanson, getting that kind of improvement will require re-writing or significantly reengineering the application software.

Moneta performs I/O benchmarks 9.5 times faster than a RAID array of conventional disks, 2.8 times faster than a RAID array of flash-based solid-state drives (SSDs), and 2.2 times faster than Fusion-io's high-end, flash-based SSD.

Sharing Principles

While Moneta is optimized for PCM, Swanson says many of Moneta's principles could be used with other fast non-volatile memories, such as spin-transfer torque random access memory. "The key takeaway is that when

these new memories appear, we will shift to a place where the software really plays a critical role in the overall performance of the system," Swanson says. "In Moneta, we focused strongly on software for minimizing latency and maximizing concurrency."

Swanson's team is now moving parts of the Moneta software out of the OS and into the storage array itself, in a special library that can be accessed by applications. Additional enhancements to reduce latency will be possible, Swanson says. For example, a database typically uses the standard I/O calls, with their inherent overhead, which are provided by the OS. But software in the new library bypasses that overhead by taking over those calls and talking to the storage hardware directly. "I can now address non-volatile storage directly from my application, just like DRAM," Gupta says. "That's the broader vision—a future in which the memory system and the storage system are integrated into one."

Swanson says taking advantage of these capabilities will require extensive revamping of application software. "We have talked with companies like Oracle and Teradata, and they realize they will have to change a lot of

Overhauling Systems Operations

Onur Mutlu, assistant professor of electrical and computer engineering at Carnegie Mellon University, is considering the use of phase-change memory (PCM) not for storage systems, but for main memory, where conventional DRAM is nearing the limits of its scaling abilities. Like Steven Swanson, assistant professor of computer science and engineering at University of California, San Diego, Mutlu is not looking to improve PCM itself, but the systems operations that surround it.

PCM is thought to be more scalable and ultimately less costly than DRAM, but has access latencies several times slower than DRAM, requires very high write energies, and is much less durable. Mutlu and his collaborators have reorganized PCM buffers in a way that mitigates high-energy writes and hides latency. Their design performs partial writes so that only modified parts of cache are written, reducing wear. And it distributes wear more evenly by shifting the contents of a write along a row of memory cells and by periodically swapping memory segments of high and low write accesses. These techniques can extend the life of PCM, used as memory, from 171 days of continuous use to as much as 20 years, he says.

An advantage of PCM is that, unlike DRAM, it scales favorably. Mutlu says PCM is expected to scale down to 9nm, and it can also gain density by storing multiple bits in a cell. He says no one knows exactly how much more DRAM can scale, but that scaling it beyond 25nm will be extremely challenging.

"PCM definitely has a future in storage," Mutlu says. "Main memory is a lot harder because of very high bandwidth and low latency requirements, so the research we are doing is 10 years out or so." He and his colleagues are not trying to replace DRAM, but augment it in hybrid memory systems. "So you have a limited amount of DRAM in the system and a very large amount of PCM," he explains. "Hardware maintains counters about which pages should go to DRAM vs. PCM, with high-frequency pages going to DRAM."

the way their systems work,” he says. A great deal of the complexity in database management systems lies in the buffer management and query optimization to minimize I/O, and much of that might be eliminated.

Also, storage systems will greatly benefit from changes in the way they access data over a network. Now, when a storage system accesses remote data on disk, it must navigate through the network stack, through file services software on both the local and remote machine, then do all that again coming back. “This change in storage performance is going to force us to look at all these different aspects of computer system design,” Swanson says. “The reach of this is going to be surprisingly broad.”

Swanson and colleagues at other universities are attempting to “catalyze these changes” by forming a consortium—not yet named—of storage system researchers. They include experts from the low levels of the OS, through the application layers, and on up to the data center and network architectures, he says. “The idea is to attack all these layers at once,” says Swanson, “and hopefully demonstrate that it’s worth industry’s time to make these changes to commercial systems.”

Swanson’s group is looking out five to eight years, he says. “The end point of this is you’ll have non-volatile solid-state storage that’s about as fast

“I can now address non-volatile storage directly from my application, just like DRAM,” Rajesh Gupta says. “That’s the broader vision—a future in which the memory system and the storage system are integrated into one.”

as DRAM,” he notes. “That would be an increase of about 2,500 times improvement in latency and bandwidth. This is much faster than Moore’s Law increases. I think it’s the largest increase in any aspect of system performance, in the shortest time, ever. Fully exploiting these memories is going to require making changes throughout the system, but it’s going to be very exciting time.”

Further Reading

Caulfield, A., De, A., Coburn, J., Mollov, T., Gupta, R., and Swanson, S.

Moneta: A high-performance storage array architecture for next-generation, non-volatile memories, *Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture*, Atlanta, GA, Dec. 4–8, 2010.

Akel, A., Caulfield, A., Mollov, T., Gupta, R., and Swanson, S.

Onyx: A prototype phase-change memory storage array, *Proceedings of the 3rd USENIX Conference on Hot Topics in Storage and File Systems*, Portland, OR, June 14, 2011.

Mollov, T., et al.

Understanding the impact of emerging non-volatile memories on high-performance, IO-intensive computing, *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, New Orleans, LA, Nov. 13–19, 2010.

Lee, B., et al.

Phase-change technology and the future of main memory, *IEEE Micro* 3, 1, Jan.–Feb. 2010.

Qureshi, K., Srinivasan, V., and Rivers, J.A.

Scalable high performance main memory system using phase-change memory technology, *ISCA ’09 Proceedings of the 36th Annual International Symposium on Computer Architecture*, Austin, TX, June 20–24, 2009.

Gary Anthes is a technology writer and editor based in Arlington, VA.

© 2012 ACM 0001-0782/12/01 \$10.00

Technology

Undergrads Seek Utility in Tech

A recent survey of technology preferences among college undergraduates offers results both predictable and surprising. As for the former, it should not raise eyebrows that the college computer lab is steadily fading in relevance on campuses, especially in this age of ubiquitous mobile devices.

As for the surprises? It defies convention that spending on physical textbooks, as opposed to digital versions, remains strong. They are the second-most commonly purchased item online (clothing is first). “The adoption of e-textbooks has progressed slowly,” says

Eric Weil, managing partner of Student Monitor LLC, which recently published the survey research in which 1,200 full-time students at 100 campuses participated. “They vastly prefer to buy used textbooks. Why is that? For the same reason they’ve bought used textbooks for decades: The highlighting is already done for them.”

Other notable survey findings include:

- ▶ **Computer labs are not dead.** More than six of 10 students still work on school-issued computers at least once a week. However, only 31% do so at least once a day.
- ▶ **Old-school tech has not**

disappeared. Microsoft Word remains the most popular computer program, used by 72% of students compared to the second-most used, Skype, at 46%. Google Docs ranks far behind Word, accessed by just 18% of undergraduates.

▶ **Smaller is better.** Students like used physical textbooks because they are easy to tote around. And that factor is contributing to the increase of tablet computers. Among undergraduates who plan a computer purchase in the next year, 18% plan to buy a tablet, compared to 10% who did last year.

▶ **Social media keeps surging.** Nearly all undergraduates (98%) have joined Facebook and three

of four use Twitter. Facebook commands nearly four hours of their time every week, and Twitter accounts for 101 minutes weekly. It is not all fun and games either. Even a site like LinkedIn is frequented by 65% of college students, who typically spend more than an hour a week there.

“It isn’t just the seniors going on there looking for a job,” Weil says. “It’s also younger students who are seeking internships or good summer employment. It’s a grown-up site. But it’s a very useful learning tool for undergraduates who want to get serious about their careers.”

—Dennis McCafferty

Law and Disorder

International law has always been a murky and Byzantine area. However, the Internet and digital technology have raised the stakes, the risks, and the challenges.


GETTING PEOPLE AND COUNTRIES to agree on things has never been the simplest of matters. However, in the age of the Internet and digital technology, there is no question that the stakes are greater than ever. “The ability to build a legal framework across nations is an increasingly difficult task,” states Michael Geist, research chair in Internet and e-commerce law at the University of Ottawa.

Cyberlaw—essentially real-world law extended to the virtual world—is at the center of an increasingly contentious battle over rights, responsibilities, and resources. Unfortunately, there is no international court and no global legislature. Antagonists may be individuals, cybergangs, or nation-states. And new forms of crime—resulting from the rise in Internet commerce, as well as digital goods—has created nettlesome challenges that touch diverse areas: theft, defamation, copyright infringement, intellectual property theft, child pornography, espionage and terrorism, to name a few.

“The legal system is struggling to keep up with today’s technology,” states Jonathan Bick, an adjunct professor of Internet law at Rutgers University Law School. Unfortunately, these days, there are more questions than answers. How is technology changing the way countries approach matters as diverse as international crime and content ownership? How is it altering business? And what are governments doing to bring order to cyberspace?

Legal Grievs

Since the emergence of the public Internet in the mid-1990s, people and communications have become intertwined in ways that would have once been unimaginable. Nearly 2.1 billion people—about 30% of the world’s population—now use the Internet. Globally



An official gives a press briefing about cyberattacks at the National Police Agency in Seoul, South Korea on July 8, 2009. South Korean intelligence officials believe North Korea or pro-Pyongyang forces in South Korea are responsible for the disruptive cyberattacks.

connected commerce, supply chains, and workplaces have become the norm. In fact, about U.S. \$10 trillion in global online transactions currently take place and the figure could rise to U.S. \$24 trillion by 2020, according to the Council of Europe (CoE). By contrast, the current gross world product is about \$63 trillion, according to the World Bank.

Of course, where there is money there are thieves and scofflaws. Yet keeping up with a fast-changing digital environment has proven overwhelming. One of the biggest challenges is the simple fact that “what’s illegal in one country may not be illegal in another,” says Pauline C. Reich, director of the Asia-Pacific Cyberlaw, Cybercrime and Internet Security Institute and co-author of *Law, Policy and Technology: Cyberterrorism, Information Warfare and Internet Immobilization*.

But the issues and challenges don’t stop there. “Individual countries can pass whatever laws they like. If you can’t arrest people and enforce the law

it’s not very useful,” Bick notes. The root problem is that there is no such thing as international law. “It’s nothing more than a series of bilateral treaties, conventions, and governments claiming jurisdiction over certain laws,” says Bick. “It’s up to individual countries to decide whether they want to comply with another country’s laws.”

Many crimes, such as identify theft and child pornography, already have clearly established laws and police protocols that span international borders. The prickliest issues revolve around areas such as print, broadcast, and telecommunications, where industrial-age business models and digital-age sharing and stealing redefine usage boundaries. The challenge, says Jan Kleissen, director of standards for CoE, is balancing risk with rights. “The Internet is critical for the exercise of people’s rights and freedoms, as well as for their everyday activities,” he explains.

There is certainly no shortage of legal disputes. For example, the United

Arab Emirates, citing security concerns, suspended data access and text messaging for 500,000 BlackBerry devices in 2010 before it struck a deal with Research In Motion (RIM) that may have placed a server inside the country, with the government able to access it. A similar clash took place in India in 2008. Security agencies there threatened to shut down BlackBerry service unless RIM allowed the government to intercept messages. The company refused, and after a series of discussions, the Indian government relented.

Bick argues that a great deal of the legal jostling taking place in the international arena revolves around economics rather than overarching legal reasons. "In some cases, governments and companies use their own courts, which they control, to promote a competitive advantage." That is nothing new, Geist says. "More powerful countries, such as the U.S., often attempt to tilt international standards in their favor."

Yet it is not only information and money that are in the legal crosshairs. This past June, Japan's parliament enacted legislation criminalizing the creation or distribution of computer viruses. The goal of cracking down on cybercrime seems noble enough, but critics stated that the new law could infringe on constitutionally guaranteed personal liberties. The Japanese Parliament later revised the law to allow legitimate uses. Of course, Japan does not exist in a vacuum. Barring an international treaty, the law is unlikely to

"Individual countries can pass whatever laws they like," notes Jonathan Bick. "If you can't arrest people and enforce the law, it's not very useful."

cut down on malware in any significant way; malware writers can always find a safe haven.

The issues grow even more complex when hackers breach systems in other countries and cybercriminals plant botnets on computers. In July 2009, South Korea accused North Korea of launching cyberattacks against government, news media, and financial Web sites—and spreading botnets. The U.S. was also the target of these attacks. North Korea denied the charges and, as with other similar cases, there is essentially no legal recourse.

Facing the Music

Yet another controversial area is publishing and copyright law. It is no secret that the Internet makes it ridiculously easy to share music, movies, books,

and software. Hackers routinely break encryption codes and once something is posted on a peer-to-peer network it often goes viral within hours. There's essentially no way to put the toothpaste back in the tube.

"Modern computer and communications technologies have pushed the law far beyond what it was intended to address," states Andrew Adams, professor of information ethics at Meiji University. "We're trying to make laws designed to apply to industrial middlemen and a 20th century business model relevant for a digital society. This tension is playing out in the legal and political arena with growing frequency."

The music industry is an example of how industrial- and digital-age models conflict. In the early 2000s, record companies resisted selling digital tracks and instead tried to stamp out file-sharing applications such as Napster. When the industry—essentially the Recording Industry Association of America in the U.S.—began serving lawsuits to those suspected of illegally sharing music (some of whom were grandmothers and children) a backlash emerged. Also, "the industry realized that if you sue 100 people and 100,000 people engage in the act, you're never going to catch up to the problem," Adams says.

To be sure, the dynamics of the business, largely as a result of the Internet, had changed radically. This ultimately forced the music industry to adapt to digital distribution of

Milestones

Computer Science Awards

The U.S. National Science Foundation, ACM, IEEE, the Institution of Engineering and Technology, and the University of Exeter recently honored leading computer scientists.

U.S. PRESIDENTIAL AWARD FOR SCIENCE MENTORING
USACM member Juan E. Gilbert, chairman of the College of Engineering and Science's Human-Centered Computing Division at Clemson University, was among nine individuals and eight organizations named as recipients of the U.S. Presidential

Award for Excellence in Science, Mathematics and Engineering Mentoring.

ACM GORDON BELL PRIZE
A research group from RIKEN, the University of Tsukuba, the University of Tokyo, and Fujitsu Limited were awarded the ACM Gordon Bell Prize in the peak performance category for execution for their research results obtained using the K computer. The award-winning results, presented at SC11, calculated the electron states of silicon nanowires, which have attracted

attention as a core material for next-generation semiconductors.

SEYMOUR CRAY COMPUTER ENGINEERING AWARD

IEEE Computer Society presented Charles Seitz, an architect and designer of innovative computing and communication systems, with the 2011 Seymour Cray Computer Engineering Award in recognition of "innovations in high-performance message-passing architectures and networks."

FARADAY MEDAL

The Institution of Engineering

and Technology awarded the Faraday Medal, its most prestigious award, to Donald E. Knuth, professor emeritus at Stanford University's computer science department, for his contributions to computer science.

LOEBNER PRIZE

AI programmer Bruce Wilcox and his new chatbot, Rosette, won the Bronze Annual Medal in the 21st Loebner Prize Competition held at the University of Exeter. It was the second consecutive year that one of Wilcox's chatbots won the Loebner Prize.

content and eventually drop digital rights management protection. The latter, according to Adams, created an additional problem of antagonizing the vast majority of honest content licensees who desire access to their music across their own digital devices. However, while these events were unfolding record companies watched revenues implode. Forrester Research reports that music industry revenues declined from \$14.6 billion in 1999 to \$6.3 billion in 2009.

Content providers have not given up. Instead they have intensified lobbying efforts. In France, the HADOPI law, enacted in 2009, created a three-strikes procedure that can cut off Internet access for an IP address after repeated copyright violations. In the U.K., the Digital Economy Act 2010 created regulatory code that requires Internet service providers to track and report violations, which could result in penalties, including termination of services.

But the battle lives on. The United Nations recently declared laws with graduated three-strikes provisions a violation of international law. "The level of surveillance needed for applying these rules on shutting down people's Internet connections may be incompatible with EU data-protection rights," says Adams.

Courting Consensus

Amid the chaos, there is a growing effort to create intentional cooperation in the battle against cybercrime. The most visible initiative is the Council of Europe's Convention on Cybercrime, which attempts to join "numerous stakeholders from around the world,"

"Modern computer and communications technologies have pushed the law far beyond what it was intended to address," says Andrew Adams.

including nation-states, nongovernmental organizations that manage Internet resources, business organizations, computer scientists, and Internet users, Kleissen explains.

Thirty-one nation-states, including the U.S., are parties to the convention, while 16 states have signed the agreement. It provides minimum standards for violations revolving around infringement of copyright, computer-related fraud, intellectual property theft, hate crimes, and violations of network security. Unfortunately, Russia and China haven't signed the agreement. And some criticize the convention for falling behind current threats, including botnets, spam, identity theft, and terrorist use of the Internet.

In the end, perhaps only one thing is clear: The years ahead will present enormous challenges. An increasingly globalized and interconnected world translates into growing concerns over online crime. "Jurisdiction may sound like a technical and dusty issue," Kleissen says, "but without properly functioning rules on jurisdiction, the Internet cannot fully develop to its potential. People must have a reasonable expectation of security and privacy." ■

Further Reading

Jaeger, P.T., Lin, J., Grimes, J.M., and Simmons, S.N.

Where is the cloud? Geography, economics, environment, and jurisdiction in cloud computing, *First Monday: Peer-reviewed Journal on the Internet* 14, 5, May 4, 2009.

Jurich, J.P.

Cyberwar and customary international law: The potential of a bottom-up approach to an international law of information operations, *Chicago Journal of International Law* 9, 1, July 2008.

Mehra, S.K.

Law and cybercrime in the United States today, *American Journal of Comparative Law* 58, Supplement 1, June 2010.

Marion, N.E.

The council of Europe's cyber crime treaty: An exercise in symbolic legislation, *International Journal of Cyber Criminology* 4, Issues 1 and 2, Jan.–July 2010 and July–Dec. 2010.

Moore, T. and Clayton, R.

The economics of online crime, *The Journal of Economic Perspectives* 23, 3, Summer 2009.

Samuel Greengard is an author and journalist based in West Linn, OR.

© 2012 ACM 0001-0782/12/01 \$10.00

Internet

Americans and Social Media

Why do Americans use social media? Mainly to stay in contact with the people who matter most to them, according to a new report from the Pew Research Center's Internet & American Life Project.

Two-thirds of the 2,277 adults surveyed by phone in April and May 2011 use social media platforms such as Facebook, Twitter, MySpace, and LinkedIn, and about the same portion of those social media users say staying in touch with current friends and family members is a major reason they use these sites.

Half of the social media users say connecting with old friends that they have lost touch with is also a major reason. Other factors play a much smaller role. For example, 14% of users say connecting with people around a shared hobby or interest is a major reason they use social media, and 9% say making new friends is equally important.

The ability to read comments by public figures such as politicians, celebrities, and athletes does not come into play as a major factor. Three-quarters of users say this plays no role in their decision to use these sites.

Compared with older adults, social media users under the age of 50 are more likely to say the tools help them keep up with existing friends and reconnect with old ones. About 70% of users under 50 say staying in touch with current friends is a major reason they use online social platforms.

"The most significant finding is that for most users, social media is seen primarily as a tool for maintaining existing key ties," says Aaron Smith, senior research specialist at the Pew Research Center and author of the report. "Activities such as meeting potential dating partners or interacting with public figures are much less relevant than deepening bonds with those who are already important."

—Bob Violino

Celebration Time

The centennial celebrations of Alan Turing's birth might help turn a quiet British genius into an iconic global hero.

CELEBRATIONS OF THE life and work of Alan Turing, who was born in London 100 years ago on June 23, 1912, will take place around the world this year, including major conferences, stage adaptations of *Breaking the Code*, and academic meetings and workshops about his research and ideas.

The Turing centennial might provide a catalyst for the mathematics and computing communities to revisit Turing's work, including some of the theories he posed but left unfinished. The centennial is also designed to raise public awareness of a man many call the founding father of computer science, and to place Turing on the pedestal containing science's most eminent figures.

"Many of the events have an academic focus, but we are also trying to stretch the public's knowledge of Alan Turing's history and work," says Barry Cooper, professor of mathematics at the University of Leeds and one of the coordinators of the 2012 Alan Turing Year celebration. Noting the U.K. government's limited support of computer science, Cooper hopes the Turing centennial will bring together a critical mass of people that will demonstrate computer science's importance and improve its image among young people.

Turing Year celebrations in 2012 will be a global affair. Some of the countries planning events include Brazil, Canada, China, Czech Republic, Germany, Israel, Italy, Netherlands, New Zealand, Norway, the Philippines, Portugal, Spain, Switzerland, and the U.S. ACM is planning a two-day celebration (June 15–16) in San Francisco that will bring together 32 A.M. Turing Award winners (see page 1). In the U.K., two major conferences are being held in Cambridge and Manchester, both places where Turing produced influential work. The University of Cambridge conference, "Computability in Europe 2012: How the World Computes," has the ideas of



Turing's computability proposal as its central theme. The Manchester conference, "Celebrating Turing—Mind, Mechanism and Mathematics," will feature lectures by 16 renowned scientists.

"The celebration is interesting because it is not just one special occasion, but a wide scope of events on an international scale," says John Graham-

"Many of the events have an academic focus," says Barry Cooper, "but we are also trying to stretch the public's knowledge of Alan Turing's history and work."

Cumming, vice president of engineering at Causata and a member of the Turing Centenary Advisory Committee. A television film about Turing was shown on U.K. television in November 2011, and several stage adaptations of Hugh Whitmore's *Breaking the Code*, which covers Turing's work at Bletchley Park during World War II, are planned.

In London, the Science Museum, which has never previously featured an exhibition on Turing, will host "An Olympian Mind: Alan Turing and the Dawn of Digital Computing, 1936–1954." "I hope people will realize Turing was an important founder of computer science," says Graham-Cumming, "and think of him as one of the great Britons."

"If the centenary year increases awareness that there is science in computing, that will be a fine thing," says Samson Abramsky, professor of computing at the University of Oxford and a member of the Turing Centenary Advisory Committee. "Turing has a special status for two reasons. Scientifically, he is one of the major founding fathers of the science of computation. His contributions spanning his work in Princeton and Cambridge in the 1930s, at Bletchley Park during the war, and at the National Physical Laboratory and Manchester University in the 1940s and 1950s, have been a major influence on the theory and practice of computation today. Still more, the human aspects of his life, such as his work at Bletchley Park, his prosecution for homosexuality, his presumed suicide and early death, have brought him into the cultural sphere and made him an iconic figure."

Iconic status, at least on a geeky level, might not be too far away for Alan Turing. A Hollywood feature film on his life is in the works. 

Sarah Underwood is a technology writer based in Teddington, U.K.

© 2012 ACM 0001-0782/12/01 \$10.00

Analyzing Apple Products

Researchers untangle the complex web of Apple's global supply chain—and offer lessons for managers and policymakers trying to chart the future course of U.S. industry.

APPLE'S ICONIC PRODUCT boxes boast of their contents' U.S. pedigree with a "Designed in Cupertino" tagline. While most buyers probably suspect the products are made outside the U.S., almost no one outside the famously tight-lipped company fully understands the economics of Apple's manufacturing supply chain.


Enter University of California, Irvine business professor Ken Kraemer who has led a series of far-reaching investigations into how Apple's products are made. The team's findings, as detailed in "Capturing Value in Global Networks: Apple's iPad and iPhone," not only shed light on the intricacies of the modern high-tech value chain, but they may also hold important lessons for managers and policymakers trying to chart the future course of U.S. industry.

The team's analysis reveals that while most Apple components are indeed manufactured abroad—predominantly in Korea, Japan, and Taiwan—the economic benefit accrues overwhelmingly to Apple. So, while each iPad tablet computer sold contributes a theoretical \$275 to the U.S. trade deficit, that transaction comes with enormous benefits to Apple's employees, shareholders, distribution partners, and, by extension, the rest of the U.S. economy.

While the devices may be assembled in China, that country accrues surprisingly little economic benefit. No iPad or iPhone components are actually manufactured in China, and as little as \$10 from each sale goes to Chinese firms, primarily by way of wages paid to workers on the assembly line.

The team's analysis revealed that, accounting for both consumers' and phone carriers' share of the cost, Apple probably retains about 58% of the sales price of a typical iPhone 4 smartphone. That margin drops considerably for the lowest-end iPad, where Apple holds on

Distribution of value for the iPhone based on Ken Kraemer's research.



to roughly 30% of the \$499 sales price. Those margins increase, of course, when Apple sells the products via its Web site or stores.

How did the team unearth so much detail about the inner workings of such a notoriously secretive organization?

"We just dug," says Kraemer. The team, which also included Greg Linden of University of California, Berkeley and Jason Dedrick of Syracuse University, relied heavily on tear-down reports from iSuppli, in which engineers disassemble products and perform painstaking analysis of each component. (The iPad, for instance, contains more than 400 individual parts.) The team then worked with industry analysts, former Apple employees, and various experts to unravel the entire supply chain surrounding the Apple products, making informed estimates about the involved financial transactions. The team is now applying its methodology to other industries, such as analyzing, for example, the U.S. wind power industry.

What lessons can managers derive from this analysis? First and foremost, the study calls into question the received political wisdom that reviving manufacturing will pave the way back to eco-

nomic prosperity. For all the Sturm und Drang around the U.S.'s eroding manufacturing base, the real economic value of high-tech manufacturing and assembly jobs may be quite limited.

Kraemer believes the lesson here is that U.S. companies can learn a lot from the Apple model by focusing their efforts further up the value chain. "We need to think more about the kind of jobs that it makes sense to keep here," he says. In his view, that means creating high-value, high-margin jobs—like design, marketing, and product management—and keeping them close to home.

If other companies can embrace the Apple model, taking advantage of worldwide manufacturing resources while creating well-paying jobs at home, there may be hope for the currently beleaguered U.S. economy.

"Those who decry the decline of U.S. manufacturing too often point at the offshoring of assembly for electronics goods like the iPhone," the authors conclude. "There is simply little value in electronics assembly." □

Alex Wright is a writer and information architect based in Brooklyn, NY.

© 2012 ACM 0001-0782/12/01 \$10.00



John McCarthy, 1927–2011

Winner of the 1971 A.M. Turing Award, John McCarthy was a founder of artificial intelligence and inventor of the Lisp programming language.

THE FIELD OF artificial intelligence (AI) was founded at a conference at Dartmouth College in 1956, with John McCarthy as one of its influential attendees. McCarthy subsequently expanded on the notion of logical AI, writing what appears to be the first paper on the topic, “Programs With Common Sense,” in 1958.

In the paper, he laid out for the first time the specifics of how a program might “know” things and reason with them, according to Stuart Russell, professor of computer science at the University of California, Berkeley. “McCarthy developed what he considered to be the right way to build intelligent systems,” says Russell.

McCarthy passed away on Oct. 23 in Stanford, CA, at 84, and had served on the Stanford faculty for four decades.

He received the A.M. Turing Award in 1971 for his work in AI, the Kyoto Prize in 1988, and the U.S. National Medal of Science in 1991.

While he was most celebrated for his AI work (including coining the term “artificial intelligence”) and for creating the Lisp programming language, he was also one of the first to investigate how to rigorously prove properties of programs. He invented abstract syntax; created the nonmonotonic logic technique called circumscription; and invented the garbage collector. He also made original contributions to ideas for space travel, including what he called the “space fountain,” “a kind of space elevator,” says Pat Hayes, once McCarthy’s research assistant and now senior research scientist at the Institute for Human and Machine Cognition in Pensacola, FL.

“His students included pioneers in logic programming, AI planning, several branches of theoretical computer science, robotics, machine learning, formal ontologies, knowledge representation, cognitive science, formal



In 1966, John McCarthy hosted a series of four simultaneous computer chess matches via telegraph against rivals in the Soviet Union.

philosophy, and much more,” says Hayes. “John was Descartes reincarnated in the 21st century, but with a better sense of humor.”

Indeed, in his 2001 short story, “The Robot and the Baby,” McCarthy cleverly explores the question of whether robots should have emotions.

McCarthy was “a very, very clear thinker regardless of the topic—politics, sociology, the water supply in the San Andreas Basin, anything,” recalls Hayes. “When you brought up a subject, you knew right away he’d already done the math and worked out the numbers. And if he found someone who was more knowledgeable than he on a subject, he’d take them aside, buy them a drink, and start grilling them. He was like a human vacuum cleaner for information from experts.”

McCarthy was in the habit of posting many of his thoughts on the Web.

“Essentially, they were a long list of his thesis ideas that were free for the taking,” says Russell. “He also created a Q&A page with a collection of questions he’d been asked over and over again over

the years...along with the correct answers so that he didn’t have to be bothered answering them one more time.”

Another of McCarthy’s important contributions to computing involved time sharing. Although not the first to conceive the idea, his was the first project to implement a time-sharing system at the Massachusetts Institute of Technology (MIT) in 1957 on a modified IBM 704 computer.

“For John, it was just a necessary solution to a nuisance,” explains Russell. “MIT only had a limited number of computers and many people who wanted to use them. He believed time sharing would be a good solution and so he made it happen. And you can read about it in his ‘Reminiscences on the History of Time Sharing.’ I’m sure he thought he’d made just a minor contribution, but for computer science as a whole, it was one of the major ideas for computer systems.”

Paul Hyman is a science and technology writer based in Great Neck, NY.

© 2012 ACM 0001-0782/12/01 \$10.00

Law and Technology

The Yin and Yang of Copyright and Technology

Examining the recurring conflicts between copyright and technology from piano rolls to domain-name filtering.

THE EMERGENCE OF the Internet has put enormous pressure on the rights model of U.S. copyright law. That model is premised on the notion that copyright holders are entitled to control the making of copies of their works, but technology has made that control somewhere between fragile and nonexistent. Content creators have struggled to restore the control assumed by copyright law. Two recent developments, one pending federal legislation and the second an industry-wide agreement between Internet service providers and content distributors, provide new looks at this ongoing issue.

Technology and copyright have a complex relationship. New waves of technology have created novel expressive opportunities and dramatic improvements in the ability to distribute copyrighted works. But new technology rarely asks permission, and with each technical advance, we have seen new opportunities and new clashes. Perforated rolls for player pianos in the early 1900s came from sheet music and roll producers were not eager

The rights model of the law has not changed—authors are entitled to control copying—but the practical ability to enforce that right has shrunk.

to write checks to copyright holders. Radio saw recorded music as a way to fill the airways even though disks came with a legend stating that the music was not licensed for radio broadcast. And the VCR introduced a new vocabulary—time shifting—and the chance to watch TV on your schedule, not broadcasters' schedules. It did so without offering any compensation to broadcasters or show producers and even created the risk that the financing model for

free broadcast TV would be put at risk by viewers with nimble fingers who fast-forwarded through commercials.

Since at least the advent of Napster, the music industry has struggled to find a strategy to control illegal downloads of music. Technology made it very easy to rip CDs and share the results with the world. The music industry responded with lawsuits, first against Napster, Aimster, and Grokster, and then against individual consumers, leading to prominent examples such as the ongoing saga of Jammie Thomas-Rasset. The suits have been on the whole quite successful, at least as measured by the standards that lawyers use. Grokster lost 9-0 on the question of whether it might be liable for inducing copyright infringement (there was much more division on the question of how the U.S. Supreme Court's prior *Sony* case should apply to this situation). Thomas-Rasset has faced juries multiple times and each time jurors have come back with damage awards—the first time \$1.92 million and second time \$1.5 million—that judges found too high.

Notwithstanding all of that, the

music industry sees these as paper victories, as file sharing has continued largely unabated. In some basic sense, law has failed the music industry. Technology has changed the integrity associated with distributing copies of copyrighted works by making copying easy and worldwide distribution instantaneous. To distribute a copy of the work is to put the means of production into the hands of consumers.

The rights model of the law has not changed—authors are entitled to control copying—but the practical ability to enforce that right has shrunk. The music industry started by chasing firms that were facilitating peer-to-peer file swapping. But this was like chasing quicksilver: even if you got your hands on one version, another would quickly reappear and the hive-mind of the P2P networks would reorganize around the new version.

The litigation clock is wildly out of sync with the speed of P2P organization. Ordinary law enforcement scales poorly and it is easy to see why the content industry would like a scalable way to enforce the key right to control whether copies are made of copyrighted works.

The Rojadirecta Case

In late 2008, Congress passed the Prioritizing Resources and Organization for Intellectual Property Act of 2008. The sole virtue of such a clumsy name is that it shortens to the PRO IP Act. The new act made it possible for the federal government to seize domain names associated with Web sites where allegedly infringing behavior was taking place or being facilitated. And seize it has. On June 30, 2010, the U.S. Immigration and Customs Enforcement bureau launched Operation in Our Sites and seized nine domain names and physical assets connected to commercial movie and television piracy Web sites. The program has expanded with additional domain name seizures for 77 sites in November, 2010 and with three additional sets of seizures of domain names through mid-2011.

Take a closer look at one of these cases. On February 1, 2011, the U.S. government seized the rojadirecta.com and rojadirecta.org domain names. Before the seizures, rojadirecta.com and rojadirecta.org offered up a guide to Internet TV focusing on sports (a lot of what the



Home pages for rojadirecta.com (top) and rojadirecta.me (below).

U.S. calls soccer but the rest of the world calls football). Like the early Napster, Rojadirecta offers links, not direct hosting, to facilitate what it calls P2P TV.

But if you go to those domain names today, when you type the .com or .org sites into a browser after the seizure, you are not offered links to the beautiful game. The URL bar for your browser

will indicate you have indeed reached your intended destination, but when you look at your screen you see three U.S. governmental enforcement seals. The rest of the page briefly sets out the basis for the seizure—a search warrant under two federal statutes—and states that copyright infringement can be a federal crime.

Search rojadirecta.org on the whois database. The current registrant for the domain name is the U.S. Department of Homeland Security with corresponding physical addresses, email addresses, and phone numbers. This is what a domain name seizure looks like. Prior to the seizure, the registrar Go Daddy dealt with Rojadirecta as its customer, but now the federal government has been substituted as its customer and the feds exercise control over the domain name.

But there is more to the story, of course. You do not have to type in a domain name to reach a Web site. That is just a convenience for us humans. You can type in an IP address directly and the seizure of rojadirecta.org means nothing for direct connection to the Rojadirecta Web site through the IP address. And, if that is too clumsy, Rojadirecta solved that problem by setting up new domain names offshore, including relocating to Spain at rojadirecta.es. Rojadirecta immediately announced its new domain names through its accounts on Facebook and Twitter and was back up and running. That is not to say that Rojadirecta does not want its original domain names back—it does and is in litigation over that—but the move to offshore domain names makes clear why the Pro IP Act is not the be-all and end-all for protecting copyrights.

The litigation over the Rojadirecta domain names is at a very early stage. In August, 2011, a federal court rejected Rojadirecta's preliminary efforts to get back its domain names. Part of this analysis turned on the injury that Rojadirecta was suffering from not having access to the old domain names and the court found that that injury was minimal. Why? Because Rojadirecta had been able to set up new domain names outside of the reach of U.S. officials and Rojadirecta could easily inform its users about its new locations. So, to be a tad cynical, the complete ineffectiveness of the seizure meant it did not matter in the short run that Rojadirecta could not use its original domain names. Rojadirecta sought to press a First Amendment claim, but the court left that for subsequent litigation, again on the view that the speech was not blocked but was instead just displaced to a new location.

The Rojadirecta saga should make clear why the content industry is looking for new enforcement tools.

New Enforcement Tools

The Rojadirecta saga should make clear why the content industry is looking for new enforcement tools. New legislation passed in 2008 in the form of the Pro IP Act, domain names seized and yet the activity continues elsewhere outside of the country. What might a solution look like? A technological approach by companies with the market position and financial stakes to make something work, companies with something to lose if they fail to comply with their obligations. Digital rights management was a technical response, but one that embedded the technical protection in the digital object itself, and not in the Internet's infrastructure. Something based in the U.S., so the firms can't just exit overseas. The natural target might be big firms with bottlenecks. This sounds like Internet service providers, search engines, and the like. It sounds like, in fact, S.968, the draft act on Preventing Real Online Threats to Economic Creativity and Theft of Intellectual Property Act of 2011—the Protect IP Act for those of you quick with abbreviations—and the industry-wide memorandum of understanding implemented in July, 2011. (And the House of Representatives has put its own dog in the fight with H.R. 3261, the Stop Online Piracy Act.)

Start with the latter and call it the ISP Memorandum. This is an agreement between key players in the content industry—the Recording Industry Association of America, the Motion Picture Association of America, and a number of the corporate members of the RIAA and the MPPA—and the leading Internet service providers, including corporate entities for Verizon, Comcast, and Time Warner Cable (but without consumers at the bargaining table).

The core of the agreement is a graduated six-step protocol for Internet service providers to respond to customers thought to be engaging in IP infringement. The protocol calls for education of offending consumers through steps such as temporary landing pages before consumers are able to access the Internet generally. Education is backed by a variety of mitigation measures directed at degrading the quality of the service delivered by the Internet service provider such as reduction in download and upload speeds.

The Protect IP Act would impose obligations on intermediaries and infrastructure providers to make it more difficult to find and get to sites that are, in the language of the bill, “dedicated to infringing activities.” The idea behind this is simple: if technology has created the problem of easy file sharing, technology also should provide the solution. This would include, after action by the federal government in court, blocking domain names from resolving to the matching IP address—DNS filtering—and not serving up links to infringing Web sites. And there is no shortage of criticism of these provisions: many law professors are up in arms about the First Amendment, while a white paper by leading technologists suggests the act would accomplish little while interfering with efforts to roll out the new DNS Security Extensions.

Conclusion

The ISP Memorandum and the draft acts represent the current bleeding edge in the ongoing struggle between copyright and technology. We have moved from piano rolls to DNSSEC, but the conflicts recur. The legal regime gives copyright holders the right to control the making of copies, but no one told that to technology. Technological engineering is frequently easier to do than institutional engineering and yet these systems need to coevolve and to do that we need to talk across the disciplines in a coherent way. If we fail to do that, we will produce a sloppy result that will not accomplish anything for law or for technology. 

Randal C. Picker (r-picker@uchicago.edu) is Paul and Theo Leffmann Professor of Commercial Law at The University of Chicago Law School and Senior Fellow at The Computation Institute of the University of Chicago and Argonne National Laboratory.

Copyright held by author.



The Business of Software

The Difference Engine


Observations on cognitive diversity and team performance.

MIKE RAN HIS fingers through his slightly unruly hair and straightened his tie one last time. At the secretary's nod he took a deep breath, pushed the door open, and stepped into his first interview since graduating from college.

Mike's Interview

There were four people at a long table. "Come in Mr. McGinley, please sit down" one of them said, gesturing to the chair across the table. Mike settled in and tried to present a confident air. "Well, Mr. McGinley," said the man directly across from him touching his fingertips together and leaning back "...could you start by telling us why you want to become a computer systems analyst?" Mike was nonplussed. "Uh, what's a 'computer systems analyst'?" he asked after a moment. His questioner's eyebrows raised and he glanced to the left and right at his fellow interviewers. "Why, it's the job you have applied for. It's the job you think your degree in applied mathematics qualifies you for," he said. "No, I applied for a job in Marketing," said Mike, "...and my degree is in anthropology."

There was a brief flurry of activity on the other side of the table as the panel quickly reread Mike's application. "Well, it seems you have been sent to the wrong interview," one of them said, "...please excuse us for a moment." The four of them conferred for a few minutes in tones too quiet for Mike to hear. The lead interviewer cleared his throat and said, "Well, would you *like* to become a computer systems analyst?"



No Computer Qualifications

This scene occurred in the North of England in the early 1970s. Britain's national steel company had initiated a rapid hiring program to staff up its computer departments. They had

The best problem solvers tend to operate in similar ways, so there is little benefit in having more than one of them.

seen how computers were going to change the industry and needed people with the skills to program them. The trouble was universities were not turning out computer scientists and engineers. There were almost no courses or degrees in computers at the time, certainly not enough to satisfy the burgeoning demand. So the company was hiring people with chemistry and math backgrounds. The company was hiring mechanical and metallurgical engineers. They were hiring all kinds of people and this was not necessarily a bad thing.

Difference

In his book *The Difference*, Scott Page showed how heterogeneous groups—teams of people who think differently—can outperform teams of like-

minded people.² In fact he showed in some of his experiments that groups composed of highly (but similarly) skilled agents often perform worse than teams of less skilled, but more diverse, thinking people. As Page provocatively asserts: *diversity trumps ability*.

The reason is quite straightforward: the best problem solvers tend to operate in similar ways, so there is little benefit in having more than one of them. Diverse thinkers, however, may see the same problem from different perspectives and can create more optimal hybrid solutions out of these viewpoints.

Other researchers have noted this. One person on a team might have better spatial processing while another processes information from a linguistic, or a logical-mathematical direction—these three approaches being examples of the different kinds of thinking outlined in Howard Gardner's theory of "*Multiple Intelligences*".¹ Apart from the modality of understanding, described by Gardner, each of us might approach a problem from a different direction. One person might look at a problem and mostly see all the challenges and difficulties that could arise (an all-too-common engineering disease unfortunately). Someone else looking at the same problem may see the possibilities and rewards. Yet another may rely on intuitive gut feel whereas his colleague may simply absorb and process the information as is, without feeling the need to express her judgment at all.^a

An Experiment

Some years ago my wife (a psychologist) and I conducted an experiment on 75 engineering executives and managers at a large telecommunications company. We first gave each person a set of deductive reasoning problems to solve and collected the results. Then we divided the group into 15 teams of five people and conducted some team development activities with them. Following this, we gave each team the same tasks they had completed as individuals

Getting the ideas is a good thing but actually using those viewpoints to create solutions is a better thing.

but this time asked them to solve the problems as a team. Once the teams had finished, but before analyzing the individual and team results, we polled each team on a number of attributes related to their performance as a team—we essentially asked them to rate themselves as to how effective they thought they were. Then we compiled the results and presented our findings. They were quite interesting:

- Most teams not only outperformed the average of their individual scores, they outperformed the best individual score on that team. This is a common finding in team development programs.

- However, a few teams actually scored worse as a team than they had as individuals. Two teams scored much lower, lower even than their lowest individual score.

- When comparing their self-report of team effectiveness against their actual performance we found that those teams that reported lower effectiveness did, in fact, have lower effectiveness. Of the teams that self-reported high effectiveness, however, some were high performance and some were decidedly not.

- One team achieved an almost perfect score. When we looked at the composition of this team, we found it was the only one that did not consist entirely of engineering managers. In fact, it included the general manager, his administrative assistant, and one of the non-technical support staff. And when we looked at the strategy this team used to solve the problem it transpired that it was the administrative assistant who suggested the approach that allowed the team an almost perfect score. To the team's credit, they embraced and built on the ideas even though they did

not come from an executive, an engineer, or a manager.

Team Engine

So what did we learn? Clearly, teams usually perform better than individuals—but we already knew that. Secondly, sufficiently bad teams can make things even worse than the worst solution from one person. Thirdly, if a team reports it is not working well, you can trust them—they really are not working well. But if a team says they are great they might be so bad they don't even know how bad they are. In this case you will have to look at the results they are getting (or not getting). Lastly, getting ideas from other viewpoints is a really good thing, even if the ideas come from people whose skills lie in different areas and whose "status" might be lower.

Getting the ideas is a good thing but actually using those viewpoints to create solutions is a better thing. This does, however, require acceptance and understanding of different perspectives.

Scott Page shows that we might be getting less than optimal results in the business of software if we focus only on technical skills and engineering or IT viewpoints and that it would be helpful to intentionally bring in different skills and perspectives. And learn how to use them, of course.

Mike's Interview Redux

The remainder of Mike McGinley's interview for the job of computer systems analyst consisted mostly of the interview panel carefully but enthusiastically explaining to him what the role would involve and trying to convince Mike that he should consider it. They did a good job. Mike was quite intrigued and graciously accepted the position, even though his degree was not on the list of prerequisites. And so an anthropologist became a computer systems analyst.

And a jolly good one he was too. **C**

References

1. Gardner, H. *Frames of Mind: The Theory of Multiple Intelligences*. Basic Books, New York, NY, 1983.
2. Page, S.E. *The Difference*. Princeton University Press, Princeton, NJ, 2007.

Phillip G. Armour (armour@corvusintl.com) is a senior consultant at Corvus International Inc., Deer Park, IL, and a consultant at QSM Inc., McLean, VA.

Copyright held by author.

^a These are some of the ways of approaching problems described by Edward DeBono in *Six Thinking Hats: An Essential Approach to Business Management*. Little, Brown and Co., New York, NY, 1985.

Historical Reflections

The IBM PC: From Beige Box to Industry Standard

Looking back at three decades of PC platform evolution.

THE IBM PERSONAL Computer was 30 years old last year. The IT world is more interested in the future than the past, so industry pundits used the anniversary primarily to ponder the end of the PC era: the future, they tell us, belongs to phones, tablets, and clouds rather than beige desktop boxes. Yet even if the 400 million or so PCs sold in 2011 were the last ever made it is still clear that no other computer architecture has ever been so important for so long. The PC evolved from a single machine to an industry standard, not just for desktop computers but for notebooks, workstations, and servers. Whether you run Windows, Linux, or even (since 2006) Mac OS you are probably running it on this platform.

But what really do we mean by a “PC” anyway? The Lenovo laptop I used to write this column does not look or act very much like the IBM system I could have received for my ninth birthday, had my parents been able to afford more than the Sinclair ZX81 that actually launched my computing career. And the genuine IBM Portable PC tucked behind my filing cabinet is rarely the machine I reach for first when leaving on a research trip.

No detail of the original PC remains unchanged in its modern descendents, except perhaps the vestigial row of pins still found on some motherboards to keep an internal speaker. This recalls the apocryphal story of a label in the



The IBM PC/XT circa 1983.

Tower of London: “Axe, 12th century (shaft 14th century, head 15th century).” Its buses, sockets, and interfaces have changed, though some capabilities are functionally preserved for backward compatibility though emulation, virtualization, or legacy modes. The results are impressive—I once successfully booted a 1999 PC with a PC-DOS 1.1 diskette from 1982.

Components of the PC Story

The story of the PC platform has five main chapters, and in each one it meant something different. For a year after its very successful 1981 launch the PC platform was a single proprietary computer model, albeit one built largely from standard parts to control costs and speed its introduction. In this it resembled earlier personal computers, such as the Apple II launched

four years earlier. The creation of the PC has been told and retold many times by journalists, so that IBM’s fateful alliance with Microsoft has taken on the character of an origin myth for the modern computer industry.

In the second chapter the PC platform broadened. Buyers could choose between an IBM PC, the more powerful PC/XT (1983) and PC/AT (1984) or a “clone” from another company. Each IBM machine set a new de facto standard, which was quickly replicated by clones claiming to be “100% compatible” with the real thing, even replicating its foibles. The shift was captured by *PC Magazine*, when in 1986 it changed its subtitle from “the independent guide to IBM personal computers” to “the independent guide to IBM-standard personal computing.” CP/M, the previous standard platform for business microcomputers, had run on a diverse range of hardware. In contrast, attempts by companies such as Apricot and DEC to improve their MS-DOS computers at the expense of compatibility were ultimately rejected by the market. The new platform was the combination of IBM hardware and DOS, rather than either in isolation. This “standard” was defined operationally by trying to run Lotus 1-2-3 and Microsoft Flight Simulator, both of which gave the quirks of the IBM hardware a thorough workout.

The third chapter began in 1987 when IBM replaced all its existing personal computer models with a new Personal System/2 range. Reporting

the launch, the *New York Times* quoted Steve Ballmer's opinion that this was "the most important introduction in the short history of personal computers. It's the computer architecture for the next decade."⁵ IBM expected clone makers to start paying license fees to copy the new and heavily patented Micro Channel Architecture bus and other system features. This made sense. The clone industry had worked well to faithfully copy IBM's models and drive down costs, but its biggest early innovation had been putting a handle on the box and squeezing in a tiny screen to create the 28-pound Compaq Portable.

Few expected the PC/AT architecture to survive for long once IBM abandoned it. A Gartner Group analyst was quoted saying "if the IBM 'clone' companies hope to keep their share of the corporate market, they'll have to match IBM's new personal computer architecture."² Yet it was IBM that found itself isolated, and by 1994 it had lost its number-one position in the personal computer market. The long-defunct PC/AT evolved into the basis of a constantly evolving informal standard, under the control of no single company. Clone companies eventually duplicated many of the features IBM introduced with its PS/2 machines, including 3.5-inch disk drives, and new connectors for keyboards and mice, but never adopted the overall PS/2 architecture. The clone firms tried to push "Industry Standard Architecture" as a new term for their platform, but the machines were still generally called "IBM Compatible" even though it was no longer clear what this meant. IBM's own machines were now less compatible than their competitors.

No longer could any single company redefine the PC platform. Yet this was the chapter in which it advanced most dramatically and crushed spirited competition from Apple, Atari, Commodore, and other proprietary alternatives. By the end of this third chapter in 1996 the processor power of a high-end PC (by then a Pentium-based machine) had risen by a factor of around 100 compared to the original PC/AT. Large screens driven by high-performance graphics cards on 32-bit buses had replaced slow and ugly EGA displays. Hard disk capacities had risen from 30 megabytes to several gigabytes. RAM capability had

Tiny choices made by the original PC/AT designers imposed fundamental constraints 10 years later.

grown from 512KB to 32MB, with high-quality audio and CD-ROM drives giving new multimedia power. The operating system advanced from DOS 3.3 to Windows 95 (or the more robust NT 4.0) and the PC finally overcame legacy limitations in its use of memory that had frustrated earlier users.

This innovation took place incrementally and with no central point of control. As a historian of technology and business I find this the most interesting and understudied part of the story. PC companies of the era had little in common with traditional high-technology manufacturers. PC/AT clones of the late 1980s were assembled from a handful of standard parts, each typically built by a different specialist firm. One was the motherboard, into which a processor and memory chips were inserted. This was screwed into a case, along with a power supply. Several key functions were performed by expansion cards plugged into 16-bit slots. A typical configuration filled three slots with a display adapter, a combined parallel and serial card to drive a printer and modem, and a disk controller card. Hard and floppy disks filled two drive bays of standard dimensions. All of these parts, including the motherboard, were available in dozens of variants from different hardware suppliers. (This also facilitated the exodus of component manufacturing to Asia, as suppliers could focus on low-cost niches without needing to engineer or market a whole computer).

Moderately knowledgeable computer users saved money by building their own computers, guided by books or evening classes. This could be done in an hour using no tool more exotic than a Phillips head screwdriver. The more

commercially minded built to order for their friends and eventually set up computer businesses in neighborhood storefronts or, as Michael Dell famously did, in his dorm room. A PC clone "manufacturer" needed to procure only one custom part: a small badge bearing its logo (to be stuck in the standard depression on the top left corner of the case). Everything else could be ordered from a catalogue.

The physical structure of the PC/AT came to define the market structure of the personal computer industry. A PC company, even a big one like Dell, did little original development work. It enhanced its models continually by procuring improved components. Improved parts would pop right in, but only if they precisely fitted the constraints and interfaces evolved from the original PC/AT. Within this framework some kinds of innovation were easy, others hard, and some impossible.

The easiest improvements involved substitution of a single component with no changes to interfaces or the structure of the overall system: faster processors, bigger hard drives, or new components like sound cards. One of the most dramatic component innovations came from Compaq, the largest of the clone makers, when in 1986 it stopped waiting for IBM to make something new to copy and launched the Deskpro 386. This shoehorned Intel's new processor into the existing architecture of the PC/AT. At the time this looked like a stopgap solution. *Newsweek* called it a "calculated risk" and noted that "customers may choose to wait for IBM, fearing that any other computer will be incompatible." However Compaq's engineering was soon duplicated by specialist motherboard producers as a way of keeping the clone ecosystem refreshed with new technology.

Other changes required producers of several components to work together. This was more difficult, but not impossible. For example, the original PC/AT used a complex and expensive hard disk controller. This set a de facto standard (known as ST-506) and rival firms toiled to produce compatible drives and controllers that were better and cheaper. But by the end of the 1980s PC clones were moving over to the new IDE standard, which shifted most of the

control electronics onto the drive itself to lower costs, improve performance, and supported much larger drive sizes. Drive manufacturers cooperated with motherboard producers (and the producers of chipset and BIOS components who supplied them), computer assemblers, and of course Microsoft to incorporate the necessary changes. Another big change was the introduction of several rival higher speed bus standards (EISA, PCI, and VLB), all of which required motherboard and expansion card producers to adopt new technology but preserved the physical dimensions of the card itself.

The most difficult thing to change was what would, in a more conventional kind of high-technology product, have been the easiest: the case design. The standard power supply occupied the right rear corner, with an inconveniently located switch. The case and the motherboard would invariably be produced by different companies. This was not a problem, as long as the motherboard was not too large for the case and had its mounting holes in the right places. It could shrink a little, but not too much as it had to reach the mounting holes and extend far enough to align its expansion slots and keyboard connector correctly with the cutouts in the back of the case. This determined the minimum width of the case and the height was set by the height of the expansion cards, so PCs remained bulky even as workstations shrank into fashionable “pizza boxes.” Any use of custom components raised costs and limited flexibility, though cases did eventually start to feature power switches on the front and special rear mountings for mouse and printer sockets. Even the popular “mini-tower” format just took the traditional desktop layout, shortened it, and turned it sideways.

Tiny choices made by the original PC/AT designers imposed fundamental constraints 10 years later. For example, graphics, network, and sound controllers were rarely integrated onto the motherboard. Why? Not because of any technological limit, but simply because there were no holes on the back of the standard case through which they could protrude to the outside world. On the other hand, hard and floppy disk controllers were widely integrated by

this point as the standard connector could simply be routed inside the case to the motherboard.

The fourth chapter opens with the appearance of machines based on Intel's new ATX motherboard format in 1996. Intel's dominance had been growing from processors into other key motherboard components, so for the first time since IBM had abandoned the PC/AT architecture almost a decade earlier there was a company with the power to introduce a successful new standard format. This revised case design, power supply connections, and system board layout. It included a large and flexible case opening so that peripherals could be connected directly to motherboard sockets. The final major constraint from the original IBM designs had been removed. More and more hardware was merged onto a handful of motherboard chips from Intel or one of a handful of rivals, typically including graphics, sound, and networking functions. People stopped talking about “IBM PC Compatible” computers and spoke instead just of PCs or, more revealingly, of Wintel computers after their key features of Intel x86-based hardware and Microsoft Windows operating systems.

We live in the fifth and perhaps final chapter, which began in the mid-2000s when laptop PCs (and later netbooks) began to outsell desktops. Desktop PCs will soon be akin to V8 engines or tube amplifiers—used only by hobbyists and those obsessed with performance over practicality. But the PC's fundamental architecture lives on in new packages.

Conclusion

The history of technology includes many landmark products and many successful standards. However, the IBM PC is perhaps unique in evolving seamlessly from a single proprietary product to an open standard and the foundation of an entire global industry. Some ideas developed by historians are useful in understanding this—for example, the idea of backward compatibility long predates the computer, as historian Thomas Hughes showed when he created the term “technological momentum” to describe the power of established electrical power systems.¹ The special characteristics of computer

systems, in particular the rapid growth in power, the layering of technologies, and the use of emulation to retain compatibility have given the PC a new kind of evolutionary flexibility.

At the same time, this story reminds us of the importance of paying attention to the mundane details of technology. Who would have guessed that the case layout would be the last main feature of IBM's design to be retained unchanged, and a major constraint on the development of the platform? Yet of course a hole, unlike a processor addressing mode or API, is something that cannot be retained virtually in a special compatibility mode. It is either there or not. As sociologists of technology such as Donald MacKenzie and Trevor Pinch have recently reminded us, no technology can truly transcend the joys and trials of this material world.^{3,4} 

Further Reading

The creation of the original PC is told from an IBM perspective in Chposky, J. and Leonsis, T. *Blue Magic: The People, Power, and Politics Behind the IBM Personal Computer*. Facts on File, NY, 1988.

For a broader and more entertaining look on the early years of the personal computer industry, including insightful analysis of the rise of the clone industry, see Cringely, R.X. *Accidental Empires: How the Boys of Silicon Valley Make their Millions, Battle Foreign Competition, and Still Can't Get a Date*. Addison-Wesley, Reading, MA, 1992.

Historians have started to think about the rise of the PC as a standard in Sumner, J. “Standard and Compatibility: The Rise of the PC Computing Platform.” In “By whose standards? Standardization, stability and uniformity in the history of information and electrical technologies.” Volume 28 of *History of Technology*, J. Sumner and G.J.N. Gooday, Eds., Continuum, London, 2008, 101–127.

References

1. Hughes, T. *Networks of Power: Electrification in Western Society, 1880–1930*. Johns Hopkins University Press, Baltimore, MD, 1983.
2. Knobelsdorff, K.E. IBM's four-month-old PS/2 has put computer world on hold. *Christian Science Monitor* (Aug. 19, 1987).
3. MacKenzie, D. *Material Markets: How Economic Agents are Constructed*. Oxford University Press, NY, 2009.
4. Pinch, T. and Swedberg, R. Eds., *Living in a Material World*. MIT Press, Cambridge, MA, 2008.
5. Sanger, D.E. IBM offers a blitz of new PC's. *New York Times* (Apr. 3, 1987).

Thomas Haigh (thaigh@computer.org) is an associate professor of information studies at the University of Wisconsin, Milwaukee and chair of the SIGCIS group for historians of computing.

Copyright held by author.

Viewpoint

Interfaces for the Ordinary User: Can We Hide Too Much?

Increasing the visibility and access to underlying file structure on consumer devices can vastly improve the user experience.

THERE IS A tendency today, in the operating systems for PCs, tablets, and cellphones, to hide the underlying file structure from users. The idea is to simplify by hiding “technical details.” We challenge this view, show there are situations where the user needs to work on the file level, and offer arguments that controlling the folder structure will simplify many operations.

Modern computer-based devices, including PCs, phones, and tablets, are designed for wide markets. Since these devices are acquired to make our lives simpler and perhaps more interesting, we do not expect to spend a lot of time learning how to use them. Interfaces should therefore be intuitive. The techniques applied to attain this goal go in two opposite directions: hiding and visualizing. By hiding technical details we can reduce the number of objects and processes an ordinary user needs to know. With visualization we can represent the more important objects in such a way that these can be handled by the user. A car is a good example. In earlier generations of automobiles, a driver would have to open the hood to pump gas, check the oil, and perhaps clean the gas filter and carburetor. Today, with automation, drivers do not need to see the engine at all. Technical details are hidden under the hood, while lamps on the dashboard give a warning if there is a problem. That is,

The complexity, extendibility, and flexibility of modern electronic devices make it impossible to hide files and folders.

the panel visualizes what we need to know in order to drive the car.

Manufacturers of computer devices, operating systems, and application software have used these principles of hiding and visualizing for decades. Low-level technical parts, ASCII codes, disk track and sector numbers, program code, and system files are hidden. At the same time care has been taken to visualize more high-level objects such as programs, files, and folders. Based on ideas from Xerox PARC in the late 1970s, the user is allowed to start a program by a click on an icon or to copy a file by “dragging” the file icon from one folder to the other, that is, to operate directly on the visual representations of the underlying objects. The aim is not only to make intuitive interfaces but also to empower the user. With control of the data users can apply their devices to all sorts of tasks.

Now an explorer program—a file manager—becomes crucial. With this the user is offered an overview of all data resources on the PC, often with folders visualized as a tree structure. However, on modern devices, from cellphones to tablets, we see a tendency to hide files and folders. The idea is perhaps to make things simpler by taking hiding a step further, using application-dependent file locations. For example, when the user takes a picture the operating system of the cellphone will store this in a default location. Most users will not be aware of where. When the user wants to view pictures this is done through a photo viewer that shows the contents of the hidden folders. Some devices even come without an explorer program. Where folder and file names are created by the system, an explorer view will be of little or no use anyway.

We argue that this development is counterproductive—that the complexity, extendibility, and flexibility of modern electronic devices make it impossible to hide files and folders. As we shall see, there are many situations where an explorer program is necessary to offer the user full control over the device.

When an Overview Is Needed

The user-oriented interfaces we build on top of the underlying hardware and software break down whenever a device error occurs. It may be an unreadable disk, a corrupted USB device, a change in the underlying file structure, or run-

ning out of disk space. That is, the command, “open document” may result in a “file not found” message. Since we cannot build hardware and software that never fails, it is better to require a basic understanding of important concepts up front than confusing the user whenever an error occurs. Knowledge of the basic technology is also required to comprehend why a backup is needed. Even when backup is automated, for example over a wireless network, the user would still need to understand that no backup will be performed if the network is not available.

A user may have several devices, such as a phone, tablet, camera, and a USB flash storage device. At some time in the future she may also buy a new type of device X. All can be connected to a PC. The traditional strategy has been to present these devices as generic disks, with similar folder structures as for internal devices. The user can then view this structure through the explorer program, and may use standard desktop operations to move or copy files between the devices. This will also work for device type X. We see that the advantage of letting a user work on low-level objects such as files and folders is that on this level one may use a large set of “generic” functions, such as copying, moving, and deleting. Copying an image is of course similar to copying a document or a spreadsheet or any other data type. By viewing external devices as generic disks instead of cameras, phones, or music players we can perform similar operations on these, for example moving files, without learning new specialized operations. Of course, any new type of device, a GPS navigator, an eBook reader, or the aforementioned device X, can be handled in the same way.

In order to simplify and hide details the user's aims must be predictable.



The alternative, that many modern operating systems and drivers apply, is to let an application pertinent to the device type take care of the connection. For example, when a user connects a digital camera to a PC running Windows, the following alternatives are presented as defaults: view, edit, and print. But most users would want to copy the pictures to a folder on the PC at this time. On the iPad, Apple uses synchronization as the default mechanism. When an iPad is connected to a PC the software will automatically synchronize data such as photos or documents between the two devices. But synchronization is a powerful mechanism that can have dramatic effects. The user may not want to transfer all the images from the PC to the tablet; there may not even be storage space

on the tablet for everything. Synchronization may therefore easily result in an “out of memory” error, leaving the average user in a situation that she cannot recover from. If the user has more than one PC, for example an office and a home computer, she may get into real trouble. At home all photos from the home computer will be synchronized with the iPad. But if she later connects the iPad to her office computer, for example to synchronize documents, all pictures on the iPad that are not on the office PC will be deleted.

In order to simplify and hide details the user's aims must be predictable. This will be the case for a car, but not for a flexible computer system where users can download new applications, modify software, add apps, and perform other tasks. Then

ACM's *interactions* magazine explores critical relationships between experiences, people, and technology, showcasing emerging innovations and industry leaders from around the world across important applications of design thinking and the broadening field of the interaction design. Our readers represent a growing community of practice that is of increasing and vital global importance.

interactions
<http://www.acm.org/subscribe>



it becomes important that the user controls the data objects. For example, the owner of a small bakery that offers home delivery may get orders by SMS. The default sorting of messages that his phone offers, with the most recent message at the top, is not what he needs. Instead he may want to have messages ordered on delivery date, or by address. He could have developed an app to move the data to an order entry system, but this will also require access to the data objects. The manufacturers of the cellphone would probably never have envisioned such a use of their device. However, with an overview of the data—in this case, the SMS messages—our baker can use the device for his special purpose.

Even the ordinary user may get into problems when the data structure is hidden. Most cellphones offer one folder only for storing messages. This works fine with a few messages, but navigation becomes difficult when these run into high numbers. Users will encounter problems when they employ their phone as a camera. With just a few images sequential storing works fine, but when the number of images runs into the hundreds or thousands user-controlled structures are required.

Conclusion

A user can benefit when her PC has noted the calendar event she entered on the cellphone, when the phone can provide her with a weather forecast for the current location at any time, or when she can “walk” down and view a street scene using Google. Many of these application-oriented tasks can be automated and the data hidden with little cost to the average user. In other cases we must take the “magic” away. Simplification is important, but should not be taken to the point where we make users less effective. Jakob Nielsen tells us that “Using the system changes the users, and as they change they will use the system in new ways.”¹ That is, users evolve with the system. In fact we can talk about a coevolution of users and systems. Systems may be used in ways that not even the designers dreamed of. Thus it can be a good investment for a computer user to learn some fundamentals about the technology. Giving the user access to data through an explorer program will

Simplification is important, but should not be taken to the point where we make users less effective.

result in much more powerful user capabilities, while hiding the data may stymie innovation (see Santini²).

We can ask why companies that have excelled in usability try to hide files and folders even when the arguments for visibility are so compelling. The professional answer may be they expect that most customers will only use the devices in the most straightforward ways (as planned, restraining the need for flexibility); that they have limited amounts of data (so that memory or disk overflow do not occur); that tablets and phones are connected to one PC only (to avoid synchronization problems); that backup routines can be performed by automatic systems and that users will contact service personnel whenever an error occurs. But there may be another reason. With access to the underlying data structure users can exploit standard formats for images, documents, and for all other data types. It does not matter which camera was used to capture images or on which computer system the images are stored. Users are free to move the data to their next PC or cellphone. If files are hidden, however, they can be stored in proprietary formats, thus making it more convenient for the user to utilize devices from one manufacturer only. **C**

References

1. Nielsen, J. *Usability Engineering*. Academic Press, San Diego, CA, 1993.
2. Santini, S. Is your phone killing the Internet? *IEEE Computer* (Dec. 2010).

Kai A. Olsen (kai.olsen@himolde.no) is a professor in informatics at Molde College and at the University of Bergen, Norway.

Alessio Malizia (alessio.malizia@gmail.com) is an associate professor in the computer science department at the Universidad Carlos III de Madrid, Spain.

Copyright held by author.


Interview

An Interview with Stephen A. Cook

Stephen A. Cook, winner of the 1982 A.M. Turing Award, reflects on his career.

THE CHARLES BABBAGE INSTITUTE holds one of the world's largest collections of research-grade oral history interviews relating to the history of computers, software, and networking. Most of the 400 interviews have been conducted in the context of specific research projects, which facilitate the interviewer's extensive preparation and often suggest specific lines of questions. Transcripts from these oral histories are a key source in understanding the history of computing, since traditional historical sources are frequently incomplete.

The following is a condensed version of an interview with A.M. Turing Award recipient and ACM Fellow Stephen A. Cook, considered one of the forefathers of computational complexity theory. The original interview was conducted by CBI researcher Philip L. Frana in Toronto, Canada, in 2002 following Cook's lecture at the University of Minnesota as part of the Cray Distinguished Speaker Series. After describing his background, including his influential 1971 presentation on "The Complexity of Theorem Proving Procedures," Cook discusses his move to the University of Toronto in 1970 and the reception of his work on NP-completeness, leading up to his 1982 A.M. Turing Award for "contributions to the theory of computational complexity, including the concept of nondeterministic, polynomial-time completeness. The ensuing exploration of the boundaries and



nature of NP-complete class of problems has been one of the most active and important research activities in computer science."

Cook further discusses the feasibility of solving the P versus NP problem, which has recently received renewed attention given increasingly powerful computational capabilities and the decreasing cost of computing. In a September 2009 *Communications* article, Lance Fortnow wrote that Cook's work on computational complexity theory has motivated a great amount of research into decision problems with "feasible" decision procedures in the areas of workplace efficiency, transportation, logistics, and manufacturing. It

also laid the groundwork for discoveries of NP-intermediate integer factorization algorithms like the high-profile Shor's algorithm for breaking public-key cryptography.

The complete transcript of the original interview is available at <http://conservancy.umn.edu/handle/107226>.

Were your parents mathematicians?

My father was a chemist. He got his Ph.D. at the University of Michigan and worked for many years at a subsidiary of Union Carbide. Later he became adjunct professor at the University of Buffalo. My mother eventually had two master's degrees, one in English and one in history. She was mostly a house-

wife but did some teaching at Erie Community College. She taught English for a number of years.

How did you make the decision to attend Michigan in 1957?

My parents were alumni and they met there. Some other relatives also attended Michigan. My mother grew up in Michigan.

I take it you knew that the mathematics and computer science programs were strong at Michigan at that point?

Not especially. I knew that Michigan was generally good academically. I was not thinking of computer science when I went to Michigan. Computers were pretty new then. This was 1957.

What were you thinking about doing when you went to Michigan?

Engineering, actually. I enrolled in the College of Engineering in Engineering Science, and I had quite an interest in electronics. Clarence, NY, is where we were living at the time, and Clarence had, and still has, a prominent citizen named Wilson Greatbatch. Greatbatch was inducted into the National Inventors Hall of Fame for inventing the first implantable artificial cardiac pacemaker. [Greatbatch, who held more than 325 patents and was a recipient of the Lemelson-MIT Prize, passed away September 27, 2011. –Ed.] He was developing the electronics for it while I was in high school so I learned and worked with him, just helping him solder up circuits. This was the early days of transistors, so I got quite interested in electronics.

He recommended Michigan?

No, it was just my parent's alma ma-

I think there was a feeling that there were certain problems that just seemed to be hard.

ter. My older brother had also gone to Michigan.

I don't know too many people who were at Michigan at that time. Did you study with Bernie Galler?

Yes, absolutely. I took a course from him my very first year. It was a one-hour credit course in programming—that was my introduction to programming.

He said you did very well. I took the liberty of asking him about you. He remembers you. Were you learning SNOBOL?

Not SNOBOL, but it was a Michigan product. I remember the Graham-Arden compiler. I can't remember the programming language. There was a homegrown algebraic programming language.

Was it MAD, the Michigan Algorithmic Decoder?

That sounds right. That might have been it.

Do you remember any of your other mentors from those days?

At Michigan, the math guy was Nicholas Kazarinoff. I was in engineering science and I took a calculus course where I performed well and he noticed me. That was really my best subject all along. He encouraged me to jump into a third-year course the second term and so I took an accelerated mathematics program. Eventually I transferred into the Bachelor of Math and Science degree program after two and a half years, and majored in mathematics.

Was it difficult to make that adjustment?

Yes. It was clear that mathematics was my real area. Of course, I was good in mathematics in high school, but I didn't know any mathematicians. I didn't really know what mathematicians did.

Who made the recommendation that you study at Harvard upon graduation?

That was one of the great mathematics departments. I applied at other places too, like Princeton and Berkeley. I don't remember exactly why I ended up at Harvard.

Was Alan Cobham already at Harvard at this time?

Well, no. Before I got there he was a graduate student in the mathematics department and was close to getting his Ph.D. He wrote a thesis, but he didn't bother to complete the minor thesis requirement. Instead, he just went off to work for IBM Research in Yorktown Heights.

Did you have a major professor in mind when you went to Harvard?

No, and I didn't really know what I wanted to do either. I put down algebra as my area. I got more interested in computers when I took a course with my eventual advisor, Hao Wang. He wasn't in the mathematics department; he was in applied physics.

What was your thesis topic?

We didn't have a real master's thesis. The master's degree was something you picked up, just course work really.

So it was just a stepping-stone?

Yes.

In one of your lectures you talk about Cobham's question, "Is multiplication harder than addition?" as being inspirational to you. Was that a turning point?

That was one thing. Yes, he wrote this interesting paper on the intrinsic computational difficulty of functions, which I read. That was an influential paper for sure. There were other things around too. Michael Rabin was interested in the same kind of problems and he had written articles, and then there were other papers. I think I mention them in my Turing Award article.^a

Right. Were Princeton faculty and students visiting Harvard in those days, or were you going down to Princeton?

No, I never went to Princeton. James Bennett's thesis was quite influential, but I never met him. In fact I think he dropped out of the academic picture as soon as he got his degree. I did meet Bob Ritchie. I think he came up to visit, I think that was the only Princeton connection I can remember.

But you were reading their papers?

^a Cook, S.A. An overview of computational complexity. *Commun. ACM* 26, 6 (June 1983), 401–408.

I guess what I provided was a definition and result—the NP-completeness result crystallized it.

I was reading their papers, but not having personal interaction with them.

At some point you picked a thesis advisor.

Yes, I had taken a couple courses from Hao Wang. We got along. He wasn't especially interested in complexity, but he was a logician and he had an interest in computation and he had done work in automatic theorems before.

You received your Ph.D. in 1966. What was your thesis?

It was on the complexity of multiplication,^b so that was right in line with Cobham's question.

You immediately took a job at UC Berkeley in 1966?

Yes, that's correct.

Did you finish in the middle of the term?

No. I finished in the spring and spent the summer in Europe, and then went off to Berkeley.

What was it like to be a Midwesterner who moves out to California? Many people have found that a very easy transition, but others have found it very difficult.

There wasn't a huge difference as far as the atmosphere on the university campus at Berkeley. Of course that was the 1960s so there was a lot more student foment.

You were there at a time of a great deal of student organizing.

There was indeed. The Free Speech

movement was in full swing and there was tear gas on campus.

You remember some of these incidents?

Yes. There were times when we couldn't get to the campus because there were demonstrators in the way.

Did you have a role?

No, I was just an observer. I wasn't politically very active.

You were hired as an assistant professor at Berkeley, not as a lecturer, is that right?

That's right. My position was in the mathematics department and I was actually cross-appointed with something in the computer center. Initially, I had no connection with the computer science department, which was just starting out then.

I realize that research and discovery is not always an evolutionary process, but this must have been a time that was most critical to you in preparing for your 1971 presentation on "The Complexity of Theorem Proving Procedures" at ACM SIGACT. Was NP-completeness something you'd been thinking about hard at Berkeley?

No. I was thinking about complexity issues, but the specific idea of NP-completeness didn't come to me until immediately before giving the paper. It was gelling from other ideas I had been thinking about.

How long had SIGACT been around in 1971?

It was pretty new. I'm pretty sure this was the third meeting.

Had you attended other meetings?

Yes. I had papers in every conference of STOC,^c as we now call it, for about 15 years.

You taught at Berkeley for four years, from 1966 to 1970.

Yes, that's right.

And after four years you decided to move on?

No. I was denied tenure by the math-

Calendar of Events

January 17–19

ACM-SIAM Symposium on Discrete Algorithms, Kyoto, Japan, Contact: David S. Johnson, Email: dsj@research.att.com, Phone: 908-582-4742

January 18–20

ACM International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems, Taipei, Taiwan, Sponsored: SIGDA, Contact: Charlie Chung Ping Chen, Email: cchen@cc.ee.ntu.edu.tw

January 22–28

The 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Philadelphia, PA, Sponsored: SIGPLAN, Contact: John Field, Email: jfield@google.com

January 28–30

ACM International Health Informatics Symposium, Miami, FL, Sponsored: SIGHIT, Contact: Jiming Liu, Email: jiming@comp.hkbu.edu.hk

January 30–February 3

Fourteenth Australasian Computing Education Conference, Melbourne, Australia, Contact: Michael da Raadt, Email: michael.d@moodle.com

February 6–8

International Conference on Agents and Artificial Intelligence, Vilamoura, Algarve Portugal, Contact: Fred Ana, Email: afred@lx.it.pt

February 7–9

Second ACM Conference on Data and Application Security and Privacy, San Antonio, TX, Sponsored: SIGSAC, Contact: Elisa Bertino, Email: bertino@cerias.purdue.edu

^b Cook's thesis title was "On the Minimum Computation Time of Functions."

^c The ACM SIGACT Symposium on the Theory of Computing.

ematics department. Tenure decisions weren't as open then. I don't know what information was presented to the department. The entire math department took a vote—I know that. But I don't know what kind of evidence was presented or what the basis for the decision was.

Was that a pattern? Did you have friends that suffered the same problem, in your small circle?

My natural colleagues tended to be in computer science departments and I think that made a big difference. Subsequently, I had no trouble getting offers from computer science departments. My field may have been a little too new to be accepted in mathematics.

This was right around the time that Dijkstra's tenure was denied in Amsterdam too.

Was that also a mathematics department decision?

Yes. He then went on to Austin, Texas. I wonder if it wasn't the same problem that he faced.

I guess so, in the sense that the field was not completely respectable, mathematically. I had training and there was a strong group of logicians in Berkeley and I had something to do with them, and I think they had some interest in my work, but apparently not enough.

Toronto hired you as an associate professor in 1970?

Yes, that's right. I think I was hired as an associate professor and then got tenured a year later. I had other offers too. I had an offer at Yale.

Why did you pick Toronto over Yale?

Just because the problem is NP-complete does not mean that you should not try to solve it.

The city, certainly, is much nicer, and the department was better established. Yale's computer science program was just starting up and it wasn't clear how well it was going to go. There were several interesting people here. An excellent departmental chair, Tom Hull, really established things here. He did a lot of early hires and set it up as one of the premier departments in the continent.

When was the department established here?

There was first a graduate department and then an undergraduate department. The undergraduate department started up about 1969 or 1970—about when I came.

You settled not far from home.

Yes. I grew up near Buffalo, and we had gone to Ontario resorts in summers.

Was the department as theoretical an institution as it is today?

I think numerical analysis was a strong feature, as it was in many early departments. There were also a couple of physicists, Kelly Gotlieb and Pat Hume. They are both retired now. Kelly is very active in ACM. He's in his 90s, but is co-chair of the ACM awards committee. He may have been the first chair of the department.

Who were your early colleagues?

Allan Borodin was probably closest. His office was right across the hall. He was a Cornell graduate, hired the year before me.

Did you know him previously?

No, I knew vaguely of his work. I met him during my first recruiting trip, and we got along fine.

What kinds of things were you teaching when you first arrived?

My first appointment was cross-appointed in mathematics. That lasted for a year or two and then I switched completely to computer science. I did teach some math courses. I taught a course in logic, for example. I also taught first-year programming—I did that at Berkeley too. I taught first-year computer programming a few times. I did teach graduate courses in my own area, in computational complexity.

Were there not the rivalries up here between the mathematicians and computer scientists?

I'm not sure that the mathematicians totally appreciated my work. I think there was a feeling on their part that they should—a feeling that this is an up-and-coming subject and they should have something to do with it. Yet I didn't have too much to do with the members of the mathematics department, and that's one reason I decided to switch over entirely to the computer science department.

Within a year you had presented your paper on the complexity of theorem proving procedures at the Symposium on the Theory of Computing. Was there an immediate and positive reaction to your paper? The very next year R.M. Karp shows that 21 problems are NP-complete. Was it something that was on a lot of people's minds?

I think so. I think there was a feeling that there were certain problems that just seemed to be hard. Rabin was also interested in these. I remember he was quite interested in the traveling salesman problem, and was trying to find ways to get lower bounds on the complexity. So I would say yes, there was something in the air for sure. I guess what I provided was a definition and a result—the NP-completeness result crystallized it.

[The traveling salesman problem is a graph theory problem that asks, "Given a map of N cities connected by roads, can a salesperson visit all the cities exactly once within some given number of miles?" The traveling salesman problem is a so-called infeasible problem in computer science, and is therefore unlikely to be computable in polynomial time; the problem is only solved in exponential time. Some readers may quibble with this definition; fuller explanations are found in: Flood, M.M. "The Travelling Salesman Problem," *Operations Research* 4 (1956), 61–75; Hoffman, A.J. and Wolfe, P. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*, E.L. Lawler et al., Eds. (Chichester: John Wiley, 1985), 1–16; Reinelt, G. "The Traveling Salesman: Computational Solutions for TSP Applications," *Lecture Notes in Computer Science* 840 (Springer-Verlag, Berlin, 1994). —Ed.]

In your lecture at the University of Minnesota you noted that you need to remind some of your new students that all problems in NP are not hard.

Yes.

I gather that that's how this mythology sprung up that NP problems are hard?

Some people even think NP stand for "Not P," but it stands for non-deterministic polynomial time. So we have these contrasting classes, P and NP. The simplistic assumption is that the P ones are the easy ones, and the NP are the hard ones. And of course P is the subset of NP. It's the NP-complete ones, the subset of NP problems, which are the hard ones.

Much of this is covered in the chapters of just about any introductory computer science textbook.

Yes.

The idea of time being the most important complexity measure seems rather straightforward to me now because I've heard it and read it several places, but it apparently wasn't.

I think time was an important measure. It was Alan Cobham who was trying to think of some intrinsic measure like "work," but in fact his theorem was about the characterization of polynomial time, so that was the thing he talked about—time. Time seemed to be the most obvious measure of complexity. Certainly space memory was also considered right from the start.

You and Richard Karp were colleagues at Berkeley?

We overlapped at Berkeley. He came to Berkeley from IBM a year or two before I left, so I knew him.

So he returned home from the SIGACT symposium, and started looking at these problems more carefully.

Yes, that's right. I think he toured the states talking to people about them, and coming up with new problems.

Is it fair to say, then, that Karp is your 'popularizer'?

Yes. He did a tremendous thing—there's no question about it. I certainly didn't realize there were so many natural computational problems out there that turned out to be NP-complete.

I've always been interested in things that had a logical flavor to them, like formal correctness of programs.

Between 1970 and 1980 you received several grants from the National Research Council to work on this problem and others. In 1975, you are promoted to full professor at the University of Toronto.

Yes.

And then followed a number of awards: the E.W.R. Steacie Memorial Fellowship to support fundamental research essential to the development of science; the Izaak Walton Killam Memorial Research Fellowship from the Canada Council for the Arts.

Yes.

And in that period too, in 1982, you were awarded the ACM A.M. Turing Award for, among other things, your contributions of complexity theory.

Yes. And of course the trigger was the theory of NP-completeness.

In 1985 you became a university professor. Numerous teaching awards followed.

A couple.

You were awarded the CRM/Fields Prize in 1998. What is the Fields Institute? A mathematics institute?

It's a mathematics institute, right. It's on our campus although I guess it's separate. It's a bit like the Isaac Newton Center at Cambridge. They have a building for mathematics research. They sponsor programs and they have emphasis programs in different areas in mathematics.

The *Mathematical Intelligencer* declared the P versus NP problem one of the three greatest math problems of the next century. Where does this perception come from?

It seems to be really relevant to the real world—probably more than the other problems on the list. It's not clear what impact on the world the other very interesting problems have, though they certainly could impact mathematics. If P equals NP, it could have a dramatic effect on the world. More likely P is not equal to NP. There the impact would still be good. It would lead to the possibility of proving cryptographic protocols are secure, which is something we can't hope to do at present.

One of your audience members said that NP-completeness is sometimes identified on the basis of—as he called it—an "unrealistic" example. I gathered he was trying to argue that there was a disconnect between the theory and human experience on some level.

I think he was referring to the fact that some problems are NP-complete but still seem to be, in practice, solvable. So that's a question of what class of inputs you want to use. Every NP-complete problem is easy to solve for some inputs and maybe in some cases the inputs you really are interested in are the easy ones. So, in that case, saying it's NP-complete is misleading. Even for the original NP-complete problem of satisfiability, the fact that it's NP-complete hasn't stopped this big industry of programs that they're solving the satisfiability problem with—in some cases, very dramatic and useful successes. As I mentioned in the talk, they've been able to verify large chunks of a microprocessor by proving unsatisfiable gaps that it causes with the tens of thousands of variables. And so there are certainly some. Just because the problem is NP-complete does not mean that you should not try to solve it.

I'm wondering if this isn't the same stumbling block, on a very general level, that other disciplines are struggling with. In genomics and bioinformatics they now talk about empirical laws that haven't found their theory yet. They say, "Don't worry about proving these things—they're empirically derived from computations."

In bioinformatics there are lots of computational problems as you say, and I'm sure that many of them are NP-complete or NP-hard in their full generality and that just means you have to

change the problem or somehow get around the complication in tractability. Or perhaps the inputs that you're really interested in may not be hard.

Do you consult with the bioinformatics people in Toronto?

We don't have a strong bioinformatics group. We do have people in the medical sciences interested in the subject. When we have talks I always attend them and I'm interested in the subject.

Can you talk about randomizing algorithms and Boolean circuit complexity being a key to P not equal to NP .

That's a possible approach and there's an intriguing connection between Boolean circuit complexity and the P and NP connection.

How often do you get messages from people who say they have solved the problem?

Not that often. I probably get one a month. They don't say they have solved the problem necessarily. Rather, they ask about it and sometimes they'll send a program or an algorithm or they'll ask if a certain approach works. In one case somebody sent a program for solving the Mine Sweeper problem, which is NP -complete. He didn't know what to do with it and he didn't want to tell me the algorithm because he was afraid I would steal it and take the million-dollar award in the Clay Mathematics Institute Millennium Prize competition.

Can the problem be solved?

It's possible. It's not quite that way, of course. There are two ways the P vs. NP question can be solved: P equals NP , or P does not equal NP . Most of us think it will be solved by showing P not equal to NP . But if it is solved by showing P equals NP , then it would have dramatic implications for mathematics and it might—I can't say for sure—but it might lead to solutions to all the other problems.

Why do complexity theorists think that P is not equal to NP ?

I think there are two main reasons. One is that computer scientists are really good at finding efficient algorithms to solve computational problems. We've been doing this now for 30

The fundamental problem in software is assuring that it's correct.

or more years, probably 40 years. There have been detailed courses on it, and mathematical successes. And as far as the NP -complete problems go, many of them are really useful in industry. Lots of people, not just academic computer scientists, but real people in the field—programmers and engineers—have been trying to solve these problems efficiently. And of course they've all failed to solve any of the NP -complete problems, at least in finding provable polynomial solutions for any of them.

So that's the one side. The other side is we think, assuming P not equal to NP , why haven't we proved it? It just seems to be very difficult. It's much easier to find an algorithm to solve a problem to show it's in P than it is to prove it's not in P because you have to rule out every possible algorithm. We know that's difficult and there is this sequence of inclusions of complexity classes; log-space is a subset of P , which is a subset of NP , which is a subset of P -space. And we know the first one: log-space is a proper subclass of the last one, P -space, by a simple diagonal argument, and therefore one of the adjacent conclusions has to be proper but we can't prove any of them are proper, so that's just good evidence we're not good at establishing separations that are there.

Do you expect a winner anytime soon?

Yes it could happen. I'm guessing it's a feasible problem to solve. Maybe we have to develop more techniques to solve it, but it's going to be solved eventually.

One of the other areas that you've been working on is assertions. I interviewed Tony Hoare in 2002; How did you come into contact with axiomatic semantics and assertions? Is that a relatively recent area of interest to you?

I did that work in the 1980s. I've always been interested in things that had a logical flavor to them, like formal correctness of programs. People in this department were interested in formal correctness and so I was aware of Hoare's work. Hoare developed these rules for proving the so-called "Hoare triples." For each instruction he would have a triple that defined the instruction in some sense. But he didn't prove anything about the whole system that he got. So I was just trying to think in some sense, 'These rules must be complete.' I was trying to figure out the sense in which they were complete.

Microsoft hired him to help introduce assertions into their operating systems.

The fundamental problem in software is assuring that it's correct. There are some fundamental problems writing it, but once you've written it you have to somehow debug it and try to develop your confidence that it's correct. From early on, various people have said we need a way of mathematically certifying the software. Coming up with formal assertions of specifications and formal proofs that it meets the specifications with 'assertions' seems a very natural way to do that.

Hoare said to me that, just before he left Oxford, teaching was moving away from knowledge-acquisition through induction to more collaborative group activities and a dialogue-oriented approach. Does that sound familiar to you? Has the teaching here changed over time?

I haven't really changed my method of teaching courses, which is the traditional one of lectures, consulting, office hours, and answering email questions. That's the way most of us still work here. The previous dean had every department develop a seminar course for first-year students so there'd be more close contact with regular faculty members. I guess that's one pressure. I don't know if it's radically different. These seminars were to be small and perhaps a little more informal. We have tutorial sessions led by graduate students in all our courses. The idea is that classes are supposed to be smaller and more interactive. 

Copyright held by author.



Association for
Computing Machinery

Advancing Computing as a Science & Profession

membership application & digital library order form

Priority Code: AD10

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

Professional membership dues include \$40 toward a subscription to *Communications of the ACM*. Student membership dues include \$15 toward a subscription to *XRDS*. 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
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 _____

Article development led by [acmqueue](http://queue.acm.org)
queue.acm.org

**Sometimes you just have
to make a better mousetrap.**

BY MATTHEW FLATT

Creating Languages in Racket

CHOOSING THE RIGHT tool for a simple job is easy: a screwdriver is usually the best option when you need to change the battery in a toy, and `grep` is the obvious choice to check for a word in a text document. For more complex tasks, the choice of tool is rarely so straightforward—all the more so for a programming task, where programmers have an unparalleled ability to construct their own tools. Programmers frequently solve programming problems by creating new tool programs, such as scripts that generate source code from tables of data.

Since programmers often build task-specific tools, one way to make them more productive is to give them better tool-making tools. When tools take the form of program generators, this idea leads to libraries for creating languages that are directly extensible. Programmers may even be encouraged to think about a problem in terms of a language that would better support the task. This approach is sometimes called *language-oriented programming*.³

Racket is both a programming language and a framework for building programming languages. A Racket program can contain definitions that extend the syntax of the language for use later in the same program, and language extensions can be packaged as modules for use in multiple programs. Racket supports a smooth path from relatively simple language extensions to completely new languages, since a programming tool, like any other piece of software, is likely to start simple and grow as demands on the language increase.

As an example task, consider the implementation of a text-adventure game (also known as interactive fiction), where a player types commands to move around in a virtual world and interact with objects:

```
You're standing in a meadow.  
There is a house to the north.  
> north  
You are standing in front of a house.  
There is a door here.  
> open door  
The door is locked.  
>
```

To make the game interesting, a programmer must populate the virtual world with places and things that have rich behavior. Most any programming language could implement this virtual world, but choosing the right language construct (that is, the right tool) to represent each game element is a crucial step in the development process.

The right constructs allow commands, places, and things to be created easily—avoiding error-prone boilerplate code to set up the world's state and connections—while also allowing a programming language's full power to implement behaviors.

In a general-purpose programming language, no built-in language construct is likely to be a perfect fit. For example, places and things could be objects, while commands could be implemented as methods. The game's players, however, don't



call methods but instead type commands that have to be parsed and dynamically mapped to responses for places and things. Similarly, saving and loading a game requires inspecting and restoring the state of places and things, which is partly a matter of object serialization but also of setting variables to unmarshalled values (or else using an indirection through a dictionary for each reference from one object to another).

Some programming languages include constructs—such as overloading or laziness—that a clever programmer can exploit to encode a domain-specific language. The de-

sign of Racket addresses the problem more directly; it gives programmers tools to explicitly extend the programming language with new syntax. Some tasks require only a small extension to the core language, while others benefit from the creation of an entirely new language. Racket supports both ends of the spectrum, and it does so in a way that allows a smooth progression from one end to the other. As a programmer's needs or ambitions grow for a particular task, the programmer can take advantage of ever more of Racket's unified framework for language extension and construction.

The text-adventure example presented here illustrates the progression from a simple embedding in Racket to a separate domain-specific language (including IDE support for syntax coloring), explaining relevant Racket details along the way; no prior knowledge of Racket is necessary. Readers who prefer a more complete introduction to the language should consult *The Racket Guide*.¹

The example is a “toy” in multiple senses of the world, but it is also a scale model of industry practice. Most every video-game developer uses a custom language, including the Racket-based language that is used to implement content for the *Uncharted* video-game series.² Evidently, when billions of entertainment dollars are on the line, the choice of programming language matters—even to the point of creating new, special-purpose languages.

The World in Plain Racket

Our text adventure game contains a fixed set of *places*, such as a meadow, house, or desert, and a fixed set of *things*, such as a door, key, or flower. The player navigates the world and interacts with things using commands that are parsed as either one or two words: a single verb (that is, an intransitive verb, since it does not have a target object) such as *help* or *look*; or a verb followed by the name of a thing (that is, a transitive verb followed by a noun) such as *open door* or *get key*. Navigation words such as *north* or *in* are treated as verbs. A user can save the game using the *save* and *load* verbs, which work everywhere and prompt the user for a file name.

To implement a text-adventure game in Racket, you would start by declaring structure types for each of the three game elements shown in Figure 1.

Racket is a dialect of Lisp and a descendant of Scheme, so its syntax uses parentheses and a liberal grammar of identifiers (for example, `transitive?` is an identifier). A semicolon introduces a newline-terminated comment. Square brackets are interchangeable with parentheses but are used by convention in certain contexts, such as grouping a field name with modifiers. The `#:mutable` modifier declares a field as mutable, since fields are immutable by default.

Figure 1. Distribution of value for the iPhone based on Ken Kraemer's research.

```
(struct verb
  (aliases      ; list of symbols
   Desc         ; string
   transitive?)) ; Boolean

(struct thing
  (name          ; symbol
   [state #:mutable] ; any value
   actions))     ; list of verb-function pairs

(struct place
  (desc         ; string
   [things #:mutable] ; list of things
   actions))    ; list of verb-function pairs
```

Figure 2. Example place definition.

```
(define meadow (place "You're in a meadow."
  (list flower)
  (list (cons south
    (lambda () desert))))))
```

Figure 3. Registering game element definitions.

```
(define names (make-hash)) ; symbol to place/thing
(define elements (make-hash)) ; place/thing to symbol

(define (record-element! name val)
  (hash-set! names name val)
  (hash-set! elements val name))

(define (name->element name) (hash-ref names name))
(define (element->name obj) (hash-ref elements obj))
```

Consequently, the complete implementation of the meadow is:

```
(define meadow (place ...)) ; as in Figure 2.
(record-element! 'meadow meadow)
```

The first struct form in the code binds `verb` to function taking one argument for each field and creating a verb instance. For example, you can define a `south` verb with alias `s` as

```
(define south (verb (list 'south 's)
  "go south" #false))
```

Lisp and Racket programs tend to use strings for the text that is to be shown to an end user—for example, verb descriptions such as "go south". A *symbol*, written with a leading single quote (for example, 'south), is more typically used for an internal name, such as a verb alias.

Given the definition of `south` and a thing, `flower`, you could define a meadow place where the `south` verb moves the player to a desert place as illustrated in Figure 2.

The `list` function creates a list, while `cons` pairs two values. The `cons` function usually pairs an element with a list to form a new list, but here `cons` is used to pair a verb with a function that implements the verb's response. The `lambda` form creates an anonymous function, which in this case expects zero arguments.

When a verb's response function produces a place, such as `desert` in the example, the game execution engine will move the player to the returned place. The game engine's support for saving and loading game state, meanwhile, requires a mapping between places and their names. (Places can be implemented as objects that can be serialized, but restoring a game requires both deserialization and updating Racket-level variables such as `meadow`.) The `record-element!` function implements mappings between names and places as shown in Figure 3.

Things must be defined and registered in much the same way as places. Verbs must be collected into a list to be used by the game's command parser. Finally, the parsing and execution engine needs a set of verbs that work everywhere, each with its response function. All of those pieces form the interesting part of the game implementation, while the parsing and execution engine is a few dozen lines of static infrastructure. See the accompanying sidebar for links to the com-

plete game implementation. The code needed to construct the virtual world is particularly verbose.

Syntactic Abstraction

Although the data-representation choices discussed previously are typical for a Racket program, a Racket programmer is unlikely to write the repetitive code that directly defines and registers places, since it includes so many boilerplate lists, conses, and lambdas. Instead, a Racket programmer would write

```
(define-place meadow
  "You're in a meadow."
  [flower]
  ([south desert]))
```

and would add a `define-place` form to Racket using a pattern-based macro. The simplest form of such a macro uses `define-syntax-rule` as depicted as Figure 4.

The form immediately after `define-syntax-rule` is a *pattern*, and the form after the pattern is a *template*. A use of a macro that matches its pattern is replaced by the macro's template, modulo substitutions of *pattern variables* for their matches. The `id`, `desc`, `thng`, `vrbl`, and `expr` identifiers in this pattern are pattern variables.

Note the `define-place` form cannot be a function. The `desert` expression after `south` is, in general, an expression whose evaluation must be delayed until the `south` command is

Figure 4. The `define-place` macro.

```
(define-syntax-rule (define-place id desc [thng] ([vrbl expr]))
  (begin
    (define id (place desc
                      (list thng )
                      (list (cons vrbl (lambda () expr )))))
    (record-element! 'id id )))
```

Figure 5. The generalized `define-place` and `define-thing` macros.

```
(define-syntax-rule (define-place id desc
  [thng ...]
  ([vrbl expr ] ...))
  (begin
    (define id (place desc
                      (list thng ...)
                      (list (cons vrbl (lambda () expr ))
                          ...)))
    (record-element! 'id id )))
```

Ellipses work in the obvious way, and with this generalized `define-place`, you can put both a cactus and a key initially in the desert and respond to direction verbs other than `north` by staying in the desert:

```
(define-place desert
  "You're in a desert."
  [cactus key]
  ([north meadow]
   [south desert]
   [east desert]
   [west desert]))
```

The macro for a thing is similarly straightforward:

```
(define-syntax-rule (define-thing id
  [vrbl expr ] ...)
  (begin
    (define id
      (thing 'id #false (list (cons vrbl (lambda () expr )) ...)))
    (record-thing! 'id id )))
```

entered. More significantly, the form should bind the variable meadow so that Racket expressions for commands can refer to the place directly. In addition, the variable's source name (as opposed to its value) is used to register the place in the table of elements.

The define-place macro so far matches exactly one thing in a place and exactly one verb and response expression. To generalize to any number of things, verbs, and expressions, you

add ellipses to the pattern as shown in Figure 5.

Verbs are slightly trickier, because you want to make simple verbs especially compact to specify, and you need one kind of pattern for intransitive verbs and another for transitive verbs. The following example illustrates the target syntax:

```
(define-verbs all-verbs
  [quit])
```

```
[north (= n) "go north"]
[knock _]
[get _ (= grab take) "take"])
```

This example defines four verbs: quit as an intransitive verb with no aliases; north as an intransitive verb with alias n and a preferred description "go north"; knock as a transitive verb (as indicated by the underscore) with no aliases; and get as a transitive verb with aliases grab and take and preferred description "take". Finally, all of these verbs are collected into a list that is bound to all-verbs for use by the game's command parser.

Implementing the define-verbs form requires a more general kind of pattern matching to support different shapes of verb specifications and to match = and _ as literals. An implementation of define-verbs can defer the work of handling an individual verb to a define-one-verb macro, which uses define-syntax and syntax-rules as shown in Figure 6.

The define-place, define-thing, and define-verb macros are examples of *syntactic abstraction*. They abstract over repeated patterns of syntax, so that a programmer can avoid boilerplate code and concentrate on the creation of interesting verbs, places, and things.

The revised game implementation, which has a compact and readable implementation of the virtual world, is available online (see the accompanying sidebar for links).

Syntactic Extension

A Racket programmer who is interested in writing a single text-adventure game would likely stop extending the language at this point. If the text-adventure engine should be reusable for multiple worlds, however, a Racket programmer is likely to take a step beyond syntactic abstraction to *syntactic extension*.

The difference between abstraction and extension is partly in the eye of the beholder, but extension suggests that functions such as place and record-element! can be kept private, while define-place is exported for use in the world-defining module with implementation-independent semantics. In the world-defining module, macros

Figure 6. The define-one-verb and define-everywhere macros.

```
(define-syntax define-one-verb
  (syntax-rules (= _)
    [(one-verb id (= alias ...) desc)
     (define id (verb (list 'id 'alias ...) desc #false))]
    [(one-verb id _ (= alias ...) desc)
     (define id (verb (list 'id 'alias ...) desc #true))]
    [(one-verb id)
     (define id (verb (list 'id) (symbol->string 'id) #false))]
    [(one-verb id _)
     (define id (verb (list 'id) (symbol->string 'id) #true))]))
```

The = and _ in parentheses after syntax-rules indicate that = and _ are literals, rather than pattern variables, in the patterns that follow. Each pattern afterward has a corresponding template. Thus, in

```
(define-verbs all-verbs
  ....
  [get _ (= grab take) "take"])
```

the define-verbs expansion turns the last clause into a define-one-verb use:

```
(define-one-verb get _ (= grab take) "take")
```

This matches the first pattern of one-verb and expands into:

```
(define get (verb (list 'get 'grab 'take) "take" #true))
```

Finally, a define-everywhere form is created for defining verb responses that work throughout the world, as needed for verbs such as save and load:

```
(define-syntax-rule (define-everywhere id ([vrb expr] ...))
  (define id (list (cons vrb (lambda () expr)) ...)))

(define-everywhere everywhere-actions
  ([quit (begin (printf "Bye!\n") (exit))]
   [save (save-game)]
   [load (load-game)]
   ...))
```

Figure 7. The typed compile-time structure.

```
(begin-for-syntax
  (struct typed
    (id ; an identifier
     type) ; a string
    #:property prop:procedure (lambda (self stx) (typed-id self))))
```

An identifier is written as a symbol, but with a # prefix, so that

```
(typed #'gen-desert "place")
```

such as `define-place` have the same status as built-in forms such as `define` and `lambda`.

To make this shift, you can put the `define-verbs`, `define-place`, `define-thing`, and `define-everywhere` definitions in their own module, called `world.rkt`.

```
#lang racket
(require "txtadv.rkt")
(define-verbs ....)
(define-everywhere ....)
(define-thing ....) ...
(define-place ....) ...
```

This module imports `txtadv.rkt`, which exports `define-verbs` as well as functions used in verb responses such as `save-game` and `load-game`. Meanwhile, `txtadv.rkt` keeps private the structures and other functions that implement the world data types.

```
#lang racket
(provide define-verbs define-thing
         define-place define-every-
           where
         save-game
         load-game
         ....)
(struct verb ....)
....
(define-syntax-rule (define-verbs
  ....) ....)
....
```

The `#lang racket` line that starts each module indicates that the module is implemented in the racket language. In `world.rkt`, `require` additionally imports both the syntactic extensions and functions that are exported by the `txtadv.rkt` module.

Since macro binding is part of the Racket language, as opposed to being implemented as a separate preprocessor, macro bindings can work with module imports and exports the same as variable bindings. In particular, the definition of the `define-verbs` macro can see the verb constructor function because of the rules of lexical scope, while code in the `world.rkt` module cannot access `verb` directly because of the same scoping rules. Since a use of `define-verbs` in `world.rkt` expands to a use of `verb`, considerable language machinery is required for Racket to maintain lexical scope in the

Where to Find the Code Online

COMPLETE GAME IMPLEMENTATION:

<http://queue.acm.org/downloads/2011/racket/0-longhand/txtadv+world.rkt>
<http://queue.acm.org/downloads/2011/racket/0-longhand/README.txt>

REVISED GAME IMPLEMENTATION:

<http://queue.acm.org/downloads/2011/racket/1-monolith/txtadv+world.rkt>
<http://queue.acm.org/downloads/2011/racket/1-monolith/README.txt>

MODULAR GAME IMPLEMENTATION:

<http://queue.acm.org/downloads/2011/racket/2-modules/txtadv.rkt>
<http://queue.acm.org/downloads/2011/racket/2-modules/world.rkt>
<http://queue.acm.org/downloads/2011/racket/2-modules/README.txt>

GAME WITH A TXTADV.RKT LANGUAGE:

<http://queue.acm.org/downloads/2011/racket/3-module-lang/txtadv.rkt>
<http://queue.acm.org/downloads/2011/racket/3-module-lang/world.rkt>
<http://queue.acm.org/downloads/2011/racket/3-module-lang/README.txt>

CODE FOR THE GAME WITH STATIC CHECKS:

<http://queue.acm.org/downloads/2011/racket/4-type/txtadv.rkt>
<http://queue.acm.org/downloads/2011/racket/4-type/world.rkt>
<http://queue.acm.org/downloads/2011/racket/4-type/README.txt>

GAME CODE WITH NONPARENTHESES SYNTAX:

<http://queue.acm.org/downloads/2011/racket/5-lang/txtadv-reader.rkt>
<http://queue.acm.org/downloads/2011/racket/5-lang/txtadv.rkt>
<http://queue.acm.org/downloads/2011/racket/5-lang/world.rkt> <http://queue.acm.org/downloads/2011/racket/5-lang/README.txt>

CODE FOR GAME WITH A DRRACKET PLUG-IN:

<http://queue.acm.org/downloads/2011/racket/6-color/txtadv.rkt>
<http://queue.acm.org/downloads/2011/racket/6-color/world.rkt>
<http://queue.acm.org/downloads/2011/racket/6-color/README.txt>
<http://queue.acm.org/downloads/2011/racket/6-color/lang/color.rkt>
<http://queue.acm.org/downloads/2011/racket/6-color/lang/reader.rkt>

presence of macro expansion, but the result is that syntactic extension is easy for programmers.

The modular game implementation is available online (see the accompanying sidebar for links).

Module Languages

Although the `world.rkt` module cannot directly access constructor functions such as `verb`, the module still has access to all of the Racket language and, via `require`, any other module's exports. More constraints on `world.rkt` may be appropriate to ensure that assumptions of `txtadv.rkt` are satisfied.

To exert further control, you can convert `txtadv.rkt` from a module that exports a language extension to one that exports a language. Then, instead starting with `#lang racket`, `world.rkt` starts with

```
#lang s-exp "txtadv.rkt"
```

For now, `s-exp` indicates that the language of `world.rkt` uses S-expression notation (that is, parentheses), while `txtadv.rkt` defines syntactic forms. Later, the S-expression and syntactic-form specifications are combined into a single name, analogous to `#lang racket`.

Along with changing `world.rkt`, you can change `txtadv.rkt` to export everything from racket:

```
#lang racket
(provide define-verbs ....
         (all-from-out racket))
....
```

Instead of `(all-from-out racket)`, you could use `(except-out (all-from-out racket) require)` to withhold the `require` form from `world`.

rkt. Alternatively, instead of using `all-from-out` and then naming bindings to withhold, you could explicitly export only certain pieces from racket.

The exports of `txtadv.rkt` completely determine the bindings that are available in `world.rkt`—not only the functions, but also syntactic forms such as `require` or `lambda`. For example, `txtadv.rkt` could supply a `lambda` binding to `world.rkt` that implements a different kind of function than the usual `lambda`, such as functions with lazy evaluation.

More commonly, a module language can replace the `module-begin` form that implicitly wraps the body of a module. Specifically, `txtadv.rkt` can provide an alternate `module-body` that forces `world.rkt` to have a single `define-verbs` form, a single `define-everywhere` form, a sequence of `define-thing` declarations, and a sequence of `define-place` declarations; if `world.rkt` has any other form, it can be rejected as a syntax error. Such constraints can enforce restrictions to limit the power of the `txtadv.rkt` language, but they can also be used to provide domain-specific checking and error messages.

The game implemented with a `txtadv.rkt` language is available online (see the accompanying sidebar for url

information). The `module-begin` replacement in the implementation requires `define-verbs` followed by `define-everywhere`, then allows any number of other declarations. The module must end with a `place` expression, which is used as the starting location for the game.

Static Checks

The `define-verb`, `define-place`, and `define-thing` forms bind names in the same way as any other Racket definition, and each reference to a verb, place, or thing is a Racket-level reference to the defined name. This approach makes it easy for verb-response expressions, which are implemented in Racket, to refer to other things and places in the virtual world. It also means, however, that misusing a reference as a thing can lead to a runtime error. For example, the incorrect reference to `desert` as a thing in

```
(define-place room
  "You're in the house."
  [trophy desert]
  ([out house-front]))
```

triggers a failure only when the player enters room, and the game engine fails when trying to print the things within the place.

Many languages provide type checking or other static types to ensure the absence of certain runtime errors. Racket macros can implement languages with static checks, and macros can even implement language extensions that perform static checks within a base language that defers similar checks to runtime. Specifically, you can adjust `define-verb`, `define-place`, and `define-thing` to check certain references, such as requiring that the list of initial things in a place contain only names that are defined as things. Similarly, names used as verbs with responses can be checked to ensure they are declared as verbs, suitably transitive or intransitive.

Implementing static checks typically requires macros that are more expressive than pattern-matching macros. In Racket, arbitrary compile-time code can perform the role of expander for a syntactic form, because the most general form of a macro definition is

```
(define-syntax id transformer-expr )
```

where *transformer-expr* is a compile-time expression that produces a function. The function must accept one argument, which is a representation of a use of the `id` syntactic form, and the function must produce a representation of the use's expansion. In the same way that `define-syntax-rule` is shorthand for `define-syntax` plus `syntax-rules` and a single pattern, `syntax-rules` is shorthand for a function of one argument that pulls apart expressions of a certain shape (matching a pattern) and constructs a new expression for the result (based on a template).

The compile-time language that is used for *transformer-expr* can be different from the surrounding runtime language, but `#lang racket` seeds the language of compile-time expressions with essentially the same language as for runtime expressions. New bindings can be introduced to the compile-time phase with `(require (for-syntax ...))` instead of just `require`, and local bindings can be added to the compile-time phase through definitions wrapped with `begin-for-syntax`.

Figure 8. The revised `define-place` macro with type declaration.

```
(define-syntax-rule (define-place id ....)
  (begin
    (define gen-id (place ....)) ; as before
    (define-syntax id (typed #'gen-id "place"))
    (record-element! 'id id)))
```

Figure 9. The revised `define-place` macro with type checking.

```
(define-syntax-rule (define-place id
                                desc
                                [thng ...]
                                ([vrb expr ] ...))
  (begin
    (define gen-id
      (place desc
        (list (check-type thng "thing") ...)
        (list (cons (check-type vrb "intransitive verb")
                    (lambda () expr))
              ...)))
    (define-syntax id (typed #'gen-id "place"))
    (record-element! 'id id)))
```

For example, to check for verbs, things, and places statically, `begin-for-syntax` can define a new typed structure as illustrated in Figure 7 to associate the binding `gen-desert` to the type "place". The `#:property prop:procedure` clause in the declaration of `typed` makes a typed instance act as a function (for reasons explained later). The function takes one argument in addition to the implicit `self` argument, but it ignores the argument and returns the typed instance's `id`.

You can use `typed` by changing the `define-place` form to bind a place name `id` to a compile-time typed record. At the same time, `define-place` binds a generated name `gen-id` to the runtime place record shown in Figure 8.

Since a typed record acts as a function, a use of `id` expands to `gen-id`, so `id` still can be used as a direct reference to the place. At the same time, other macros can look at the `id` binding and determine that its expansion will have the type "place".

Other macros inspect types by using a `check-type` macro. The implementation of `check-type` is in the complete code online, but its essential feature is that it uses a compile-time function `syntax-local-value` to obtain the compile-time value of an identifier; the `check-type` macro then uses `typed?` to check whether the compile-time value is a type declaration, in which case it uses `typed-type` to check whether the declared type is the expected one. As long as the type check passes, `check-type` expands to its first argument.

The `define-place` macro uses `check-typed` to check whether the list of things at the place contains only names that are defined as things. The `define-place` macro also uses `check-typed` to check whether verbs that have responses in the place are defined as intransitive verbs (see Figure 9).

The `define-one-verb` macro must change to similarly declare each verb as either type "transitive verb" or "intransitive verb". The `define-thing` macro changes to declare its binding as a "thing", and it checks that each handled verb is defined as a "transitive verb".



Racket macros can implement languages with static checks, and macros can even implement language extensions that perform static checks within a base language that defers similar checks to runtime.



See the sidebar for the online availability of the code for the game with static checks.

The implementation of `check-form` uses `syntax-case`, which provides the pattern-matching functionality of `syntax-rules`, but pairs each pattern with an expression rather than a fixed template.

New Syntax

A Racket programmer who defines a custom text-adventure language for other Racket programmers is especially likely to stop at this point. If the text-adventure language is to be used by others who are less familiar with Racket, however, a different notation may be appropriate. For example, others may prefer a notation such as the following from `world.rkt`:

```
#lang reader "txtadv-reader.rkt"

===VERBS===
north, n
  "go north"
get _, grab _, take _
  "take"
...
===EVERYWHERE===
save
  (save-game)
load
  (load-game)
...
===THINGS===
---cactus---
  get
    "Ouch!"
...
===PLACES===
---desert---
  "You're in a desert."
    [cactus, key]
north start
south desert
...
```

In this notation, instead of forms such as `define-verbs` and `define-everywhere`, sections of the program are introduced by tags such as `===VERBS===` and `===EVERYWHERE===`. Names in the `===VERBS===` section implicitly define verbs, listing aliases afterward through a comma-separated sequence followed by an optional description of the verb. Similarly, each name in the

===EVERYWHERE=== section implicitly defines the response to a verb; the responses are still written as Racket expressions, but they could be in any alternate notation, if desired. Each thing and place is defined by its own subsection, such as `---cactus---`, with per-object verb responses in the same way as in `===EVERYWHERE===`.

Non-S-expression syntax is enabled in `world.rkt` by starting with `#lang reader "txtadv-reader.rkt"` instead of `#lang s-exp "txtadv.rkt"`. The reader language constructor, unlike the `s-exp` language constructor, defers parsing of the program's text to an arbitrary parsing function that is exported by the named module, which in this case is `txtadv-reader.rkt`. The parser from `txtadv-reader.rkt` is responsible for processing the rest of the text and converting it into S-expression notation, including the introduction of `txtadv.rkt` as the module language for the parsed `world.rkt` module.

More precisely, a reader function parses input into a syntax object, which is like an S-expression that is enriched with lexical-context and source-location information. It also acts as the representation of code for macro-transformer arguments and results. The syntax-object abstraction provides a clean separation of character-level parsing and tree-structured macro transformations. The source-location part of a syntax object automatically connects the result of macro expansion back to the original source; if a runtime error occurs in the code generated from `world.rkt`, then the error can point back to the relevant source.

The game code with nonparentheses syntax is available online (see the sidebar for urls).

The parser in `txtadv-reader.rkt` is implemented in an especially primitive way with regular expressions. The Racket distribution includes better parsing tools such as Lex- and Yacc-style parser generators.

IDE Support

One of the benefits of S-expression notation is that a programming environment's functionality adapts easily to syntactic extension, since syntax coloring and parentheses matching

can be independent of macro expansion. Some of those benefits are intact with the new syntax for describing a world, since the parser keeps source locations with identifiers and since the code ultimately expands to Racket-level binding forms. For example, the Check Syntax button in DrRacket can automatically draw arrows from the binding instance of `cactus` to each bound use of `cactus`.

DrRacket needs more help from the language implementer for IDE features, such as syntax coloring, that depend on the character-level syntax of the language. Filling in this piece of the sample text-adventure language takes two steps:

- Install the language's reader as a `txtadv` library collection instead of relying on a relative path such as `txtadv-reader.rkt`. Moving to the namespace of library collections allows DrRacket and the program to agree on which language is being used (without requiring project-style configuration of the IDE).

- Add a function to the `txtadv` reader module that identifies additional support for the language, such as a module that implements on-the-fly syntax coloring. Again, since DrRacket and the module use the same specification of the module's language, the syntax color can be precisely tailored to the module's language and content.

The code for the game with a DrRacket plug-in for syntax coloring is available online. You will find links in the accompanying sidebar.

This plug-in colors the program according to the game language's syntax instead of Racket's default rules, highlighting lexical syntax errors in red.

More Languages

The source code of the Racket distribution includes dozens of unique `#lang` lines. The most common is `#lang racket/base`, a stripped-down variant of `#lang racket`. Other common lines include `#lang scribble/manual` for documentation sources, `#lang racket/unit` for externally linkable components, `#lang scheme` for legacy modules, and `#lang setup/infotab` for library metadata. Most Racket languages use S-expression notation, but `scribble/manual` is a notable

exception; even parentheses-loving Racketeers concede that an S-expression is a poor notation for documentation prose.

Different languages in the Racket distribution exist for different reasons, and they use Racket's language-creation facilities to different degrees. Racket developers do not create new languages lightly, but the benefits of a new language sometimes outweigh the cost of learning a language variant. These benefits are as readily available to Racket users as to the core Racket developers.

Racket's support for S-expression languages and language extensions is particularly rich, and the examples in this article only scratch the surface of that toolbox. Racket's toolbox for non-S-expression syntax is still evolving, especially with respect to composable parsers and language-triggered IDE plug-ins. Fortunately, Racket's `#lang` protocol moves most of the remaining work out of the core system and into libraries. This means that Racket users are as empowered as core Racket developers to develop improved syntax tools. 

Related articles on queue.acm.org

DSL for the Uninitiated

Debasish Ghosh

<http://queue.acm.org/detail.cfm?id=1989750>

The World According to LINQ

Erik Meijer

<http://queue.acm.org/detail.cfm?id=2024658>

OCaml for the Masses

Yaron Minsky

<http://queue.acm.org/detail.cfm?id=2038036>

References

1. Flatt, M., Fidler, R.B. PLT. 2011. The Racket Guide; <http://docs.racket-lang.org/guide>.
2. Liebgold, D. Functional mzScheme DSLs in game development. Presented at Commercial Users of Functional Programming (2011).
3. Ward, M. Language-oriented programming. *Software—Concepts and Tools* 15, 4 (1994), 147–161.

Matthew Flatt is an associate professor in the School of Computing at the University of Utah, where he works on extensible programming languages, runtime systems, and applications of functional programming. He is one of the developers of the Racket programming language and a coauthor of the introductory programming textbook, *How to Design Programs* (MIT Press, 2001).



Networks without effective AQM may again be vulnerable to congestion collapse.

BY JIM GETTYS AND KATHLEEN NICHOLS

Bufferbloat: Dark Buffers in the Internet

TODAY'S NETWORKS ARE suffering from unnecessary latency and poor system performance. The culprit is bufferbloat, the existence of excessively large and frequently full buffers inside the network. Large buffers have been inserted all over the Internet without sufficient thought or testing. They damage or

defeat the fundamental congestion-avoidance algorithms of the Internet's most common transport protocol. Long delays from bufferbloat are frequently misattributed to insufficient bandwidth and this misinterpretation of the problem leads to the wrong solutions being proposed.

Congestion is an old problem on the Internet, appearing in various forms with different symptoms and causing major problems. Buffers are essential to the proper functioning of packet networks, but overly large, unmanaged, and uncoordinated buffers create excessive delays that frustrate and baffle end users. Many of the issues that create delay are not new, but their collective impact has not been widely understood. Thus, buffering problems have been accumulating for more than a decade. We strive to present these problems with their impacts so the

community can understand and act upon the problem and, we hope, learn to prevent future problems.

This article does not claim to be the first to identify the problems of excessive buffering, but it is instead intended to create a wider understanding of the pervasive problem and to give a call to action.

Internet Buffers and Congestion

The latency a packet experiences in a network is made up of transmission delay (the time it takes to send it across communications links), processing delay (the time each network element spends handling the packet), and queuing delay (the time spent waiting to be processed or transmitted). Paths between communicating endpoints in the Internet are typically made up of many hops with links of different rates or bandwidths; the smallest

bandwidth along the path is referred to as the bottleneck bandwidth. Packets cannot arrive at the destination any faster than the time it takes to send a packet at the bottleneck rate; without effective use of the network the delay can be much worse.

Latency along the path—the time from the start of a packet’s transmission by the sender until the packet is received at the destination—can be much longer than the time it takes to simply transmit the packet at the bottleneck rate. To maintain a steady flow of packets at the maximum rate, the “packets in flight” must be sufficient to fill the “pipe” of latency between sender and destination. Buffers are placed in front of a communications link in case packets arrive while the link is in use, thus requiring storage while the previous arrivals are serviced. The important location for buffers is at the path bottleneck, but the critical fast-to-slow transition can be different for different paths, different in the reverse path, and with dynamic bandwidths can change along the same path.

Figure 1 shows the relationship between throughput and delay for a packet network. System throughput is the fastest rate at which the count of packets transmitted to the destination by the network is equal to the number of packets sent into the network. As the number of packets in flight is increased, the throughput increases un-

til packets are being sent and received at the bottleneck rate. After this, more packets in flight will not increase the received rate. If a network has large buffers along the path, they can fill with these extra packets and increase delay.

A network with no buffers has no place for packets to wait for transmission, thus, extra packets are dropped. If the sending rate is increased, the loss rate increases correspondingly. Although any received packets will have a short delay, the usefulness of the received data stream would likely be impacted. To operate without buffers, arrivals must be completely predictable and smooth; thus, global synchronized timing is critical to avoiding loss. Such networks are complex, expensive, and restrictive (that is, they lack the flexibility of the Internet). A well-known example of a bufferless network is the analog telephone network before packet switching took over. Adding buffers to networks and packetizing data into variable-size packets was part of the fundamental advance in communications that led to the Internet. The history of Internet congestion and its solution is the story of trying to find the optimal way to deploy and use buffers in a network. That story is still being written, but some of the lessons of the past are being ignored.

The fundamental transport protocol of the Internet is TCP/IP. TCP’s per-


sistence is testimony both to the robust and flexible design of the original algorithm and to the excellent efforts of the many researchers and engineers who have tuned it over the decades. TCP made use of the idea of pipesize and the knowledge there was reasonable but not excessive buffering along the data path to send a *window* of packets at a time—originally sending the entire window into the network and waiting for its acknowledgment before sending more data.

The early Internet was plagued by insufficient buffering. Even under moderate loads, the packets in flight of one or more connections could arrive at a bottleneck link in a burst and be dropped because of insufficient bandwidth. This led to heavy losses and the plummeting throughput associated with congestion collapse. Internet researchers and engineers had to advocate for sufficiently large buffers to avoid this poor network utilization. Congestion collapse hit a large part of the Internet in 1986. The network became clogged with retransmitted packets while *goodput* slowed to a trickle. As part of the solution, slow-start and congestion-avoidance algorithms were added to TCP and rapidly deployed throughout the Internet. They enabled the early Internet to recover and set the stage for the rapid growth in the 1990s with the adoption of World Wide Web applications.

These TCP additions attempt to keep the network operating near the inflection point where throughput is maximized, delay is minimized, and little loss occurs. A sender-destination pair’s TCP tries to determine the pipe-size between them and to keep exactly that number of packets in flight throughout the data transfer. Since networks are shared and conditions change along the path, the algorithms continually probe the network and adapt the number of packets in flight. The slow-start algorithm (*slow* relative to the algorithm it replaced) attempts to make a first guess as to how fast TCP may operate by an initial exponential-growth phase in transmission rate. When the first packet loss is detected, TCP reduces its sending rate and enters the congestion-avoidance phase.

At the advent of congestion control in TCP, the recommendation for buffer

Figure 1. Throughput and delay.



sizing was to have a bandwidth-delay product (BDP's) worth of buffer, where bandwidth is the bottleneck link and delay is the round-trip time (RTT) between sender and destination. The rationale is that such a buffer can hold an entire "flight" of packets should they all arrive at the bottleneck link in a burst. To apply this rule, the bandwidth used in link buffer sizing was that of the immediately outgoing link, since the location of the actual bottleneck is unknown. Similarly, a canonical value was suggested for the RTT: 100ms, a continental delay for the U.S. and Europe.

Once adequate buffers became routine, another problem could occur: the buffers were now part of the pipe that TCP is so good at filling. Filling these buffers would cause delay to increase, and persistently full buffers lack the space to absorb the routine *burstiness* of a packet network. John Nagle's seminal work in 1985¹⁰ first drew attention to the consequences of large buffering. While working on TCP congestion-avoidance algorithms, Van Jacobson recognized the "persistently full buffer" issue in 1989, culminating in the development of Random Early Detection (RED) with Sally Floyd in 1993.⁵

A number of implementations, variations, imitations, and reports on RED's use are available in the literature.¹⁴ These are generically termed active queue management (AQM), which attempts to keep the queues at the

bottleneck from growing too large by monitoring the growth of the packet queue and signaling the sender's TCP to slow down by dropping (or marking) packets in a timely fashion. Different approaches have been taken to monitoring the packet queue and making the drop (or mark) decision. The Internet Research Task Force (IRTF) urged the deployment of active queue management in the Internet, publishing an RFC in 1998, popularly known as "the RED manifesto."²

Note that packet loss for TCP is not in itself a problem but is *essential* for its functioning in the face of congestion. The excessive and consecutive packet losses that come from persistently full buffers do present a problem, which is what the "warning" drops of AQM prevent (in addition to long delays).

The truth is, AQM is not widely or consistently configured and enabled in routers and it is completely unavailable in many devices. Furthermore, the existence of cheap memory and the misguided desire to avoid packet loss have led to larger and larger buffers being deployed in the hosts, routers, and switches that make up the Internet. It turns out this is a recipe for bufferbloat. Evidence of bufferbloat has been accumulating over the past decade, but its existence has not yet become a widespread cause for concern. The next section outlines Jim's personal journey of discovery.

Jim's Journey: (Re)Discovering Latency

"The Internet is slow today, Daddy." This had become a frequent refrain in the Gettys household. When I would attempt to debug the problem, like a will-o'-the-wisp, it would usually vanish. On several occasions symptoms occurred long enough for me to waste significant amounts of time on my ISP's support line before they vanished. I attributed the recurring problem to the doubtful quality of the cable to my house or equipment damage from a lightning strike. Since my job is research in immersive teleconferencing, I knew I had to dig into the intermittent poor network problem, if only for myself.

An Enlightening Lunch. Suspecting features of my cable provider to be part of the problem, I met with Comcast's Rich Woundy, who provided a number of new issues to consider:


- The "big-buffers" problem, which David Clark (Internet network architect, currently senior research scientist at MIT) had warned about several years earlier.

- Broadband measurement studies have been indicating overly large edge buffers.

- AQM is MIA—many ISPs are running without any AQM even in circumstances where they really should.

- A group at UC Berkeley's ICSI (International Computer Science Institute) had developed a very fine tool

Figure 2. Smokeping from Jim's home to MIT.



for network diagnosis called Netalyzr (<http://netalyzr.icsi.berkeley.edu>).

The next day, I recorded my “smoking gun,” a smokeping (<http://oss.oetiker.ch/smokeping/>) plot, while moving 20GB of data from my house to MIT (see Figure 2). The uncongested RTT of this path is less than 10 milliseconds, yet I was experiencing more than 1.2 seconds of latency, significant packet loss, and painful Web browsing. I suspended the rsync a few times to read my email, and, as can be seen on the plot, this would almost

instantly “fix” my home network. This was the “Daddy, the Internet is slow today” phenomenon; the fact that the delays went away when I suspended my large data transfer explained why the problem disappeared when I went looking for it; in order to debug the network, I was stopping the work that was inducing the problem.

Packet Captures at Home and “Abroad.” I took a capture of a large file transfer over the same path. Scrolling through this capture with Wireshark (<http://www.wireshark.org>) showed pe-

culiar behavior: obvious bursts of terrible behavior containing hundreds of duplicate ACKs, multiple retransmits, out of order packets, among others, on about 10-second periods followed by long periods of what looked like normal behavior. A plot of the data revealed 500KB in flight over a 10ms path. My uplink bandwidth was 2Mbps, so the true BDP was 2.5KB. I could expect use of a 100ms RTT for buffer sizing to result in 25KB, but 500KB was an order of magnitude larger. The one-second RTT is consistent with emptying a 500KB buffer at 2Mbps. To remove uncertainty, I repeated the experiment directly plugged into the cable modem and saw the same results.

I repeated my tests over my in-laws’ fiber broadband service in New Jersey. Again, the results showed much more data in flight and much larger RTT times than expected: 250KB outstanding on a 20ms path and 200ms latency with almost the same shape as on my cable. Over subsequent weeks, I added to my data sets by visiting local libraries and other targets of opportunity; the pattern was the same wherever I went. Finally, I had collected enough disturbing data to be consistent with the big-buffers problem, and I suspected that the problem was endemic among all technologies and providers.

Calling the Experts. I posted the packet traces to a group of TCP experts: Dave Clark, Dave Reed, Scott Bradner, Greg Chesson, Van Jacobson, and Vint Cerf. Their feedback revealed that extreme buffering created an artificially large pipe size and that packet discards occur according to *tail drop*—that is, when a packet arrives to a full buffer, it is dropped. The packet’s destination is unaware of the dropped packet until the entire bloated buffer has been transmitted, which can take many times the uncongested RTT. TCP expects timely notification of packet loss for correct operation. With such large buffers, TCP’s slow-start algorithm doesn’t see any drops and thus greatly overestimates the correct pipe size and requires multiple packet drops before TCP can enter its congestion-avoidance phase.

Jacobson provided plots of the data, reproduced in figures 3 and 4. The shape of the window-size evolution is characteristic of the CUBIC implementation of TCP,¹³ which is the

Figure 3a. Packet RTT and window size: Over five-minute trace.





Figure 3b. Packet RTT and window size: Detail over 70 seconds.



default in Linux. CUBIC's initial window growth is like Jacobson's original algorithm, but it "flattens out" at an apparently stable pipe-size estimate for a while.⁸ If it does not detect a congestion event (that is, a packet drop), then it ramps the window up quickly. In the trace in Figure 3, about five seconds pass without a drop, whereupon TCP ramps up the window in an attempt to find a new operating point; after 10 seconds the buffer is full, and packet loss occurs. (Because of the offload engine, the RTTs are from the last packet of the jumbogram it receives.) Notice the RTT ramps up quickly to about 1.2 seconds and mostly stays there. The three-second RTT spikes show where massive dropping took place when the buffer became full. These are followed by a drop in window size and RTT. This, in turn, causes TCP to shut down the window size. The window-size curve illustrates that sometimes the algorithm gets a drop before it goes into CUBIC's second probing stage.

Figure 4 shows the goodput (determined from the ACKs) versus time. The initial brief period of nearly 10Mbps is a result of Comcast's PowerBoost feature and is followed by a steady 2Mbps, showing that I was getting all the bandwidth I expected. Figure 4b shows the RTT seen at each window size. The lower data set clearly results from the PowerBoost phase and the upper data set is the subsequent 2Mbps phase. These points show exactly the situation shown abstractly in Figure 1: the delay is about 10ms initially; then the window size increases, the buffer fills up, and the delay (as measured by the RTT) increases linearly.

TCP should approximately share a bottleneck link between competing flows. The impact of bufferbloat on TCP's behavior is subtle and profound in two ways:

- For TCP congestion avoidance to be useful to people using that link, a TCP connection causing congestion must react quickly to changes in demand at the bottleneck link, but TCP's reaction time is quadratic to the amount of overbuffering. A link that is 10 times overbuffered not only imposes 10 times the latency, but also takes 100 times as long to react to the congestion. Your short, interactive TCP connection

loses completely to any long-lived flow that has saturated your link.

- The long-lived flow's inability to respond to congestion can cause complete starvation on competing transfers (yours or anyone who shares the link). A local service may be overbuffered by another factor of 10 times compared to a remote service.

Closing the Case on Broadband Bufferbloat. Evidence of excessive buffering accumulated over the course of the past decade is finally sufficient to motivate systematic study.

A 2007 study of nearly 2,000 hosts connected through cable and DSL companies in Europe and North America focused on measuring the residential

"last mile." The results showed that upstream queues for DSL were frequently in excess of 600ms, and those for cable generally exceeded one second.⁴

Netalyzr, a measurement tool for the last mile or access link, has been key to the exposure of bufferbloat. A 2010 study of 130,000 measurement sessions revealed widespread, severe overbuffering in the broadband edge.⁹ (The results are used here with permission of the authors.) Figure 5 is a scatterplot of bandwidth plotted against inferred buffer capacity, each point representing a single Netalyzr test session. The solid diagonal lines indicate the latency, in seconds, exposed by Netalyzr's buffer test. The tests show excessive latencies

Figure 4a. Uplink bandwidth during trace.





Figure 4b. Scatterplot of RTT experienced vs. window size during trace.



in both downlinks and uplinks in all broadband technologies. Since Netalyzr tops out at 20Mbps and bounds the test length at five seconds, the situation is clearly worse than shown.

Focusing only on cable customers, the same study showed the equipment had two dominant buffer sizes: 128KB and 256KB (for reference, a 3Mbps uplink would take 340ms to empty a 128KB buffer; and a 1Mbps uplink would take about one second). The Netalyzr authors note the difficulty of sizing buffers for the wide range of operational access rates, both from different service levels and from dynamically varying rates. Case closed.

Where There's Smoke, There's Usually Fire. Observation of 8-second latency at my home router sparked installation of OpenWrt (www.openwrt.org) for further investigation. I set the router-transmit queue to zero but saw no effect on latency. The WiFi link from my laptop was of poor quality (resulting in a bandwidth of around 1Mbps), so the bottleneck link was my WiFi link—and since my test was an upload, the bottleneck was in my laptop rather than in my router! I finally realized that AQM is not just for routers; outbound bottlenecks could easily be at the host's queue, and WiFi is now frequently the bottleneck link.

Manipulating the Linux transmit queue on my laptop reduced latency

about 80%; clearly, additional buffering was occurring somewhere. “Smart” network interface chips today usually support large (on the order of 256 packets) ring buffers that have been adjusted to maximize throughput over high-bandwidth paths on all operating systems. At the lowest Wi-Fi rate of 1Mbps, this can add three seconds of delay. Device-driver ring buffers need careful management, as do all other buffers in operating systems. A single packet of 1,500 bytes is 12ms of latency at 1Mbps; you can see the amount of buffering must adjust dynamically very quickly over two orders of magnitude so as not to sacrifice bandwidth or latency.

Compounding this problem, modern operating systems adjust socket-buffer sizes in response to observed delay; so operating system and driver bufferbloat can cause a cascade of excessive buffering higher in the network stack, resulting in still higher latencies in applications.

Bufferbloat is not just in broadband. In 2009, Dave Reed (Internet network architect, now with SAP Labs) reported problems in 3G networks: he saw high RTTs without packet loss and correctly diagnosed the cause.¹² Very high latencies were observed to the point where packets may be delivered but so late that they are seldom useful; people time out before packets.¹¹

Broadband and wireless bufferbloat are also the root causes of most of the poor Internet performance seen at many hotels and conferences.


Though the edge is more easily measured, there are some reports of congestion in the core. The RED manifesto has usually been ignored, so there are “dark” buffers hidden all over the Internet.

The Road to Hell is Paved with Good Intentions

In the past decade, not only was AQM not deployed, but new factors, unknown at the time of the RED manifesto, also exacerbated the problems of full buffers. The early Internet had slow links and a very small number of simultaneous data transfers sharing these links. Wireless did not exist. The first residential Internet systems connected personal computers through low-bandwidth links into an Internet of relatively high-speed links. The Internet has evolved to a very bandwidth-rich core. Today residential and small business Internet connections increasingly connect customers' high-bandwidth stub networks through smaller bandwidth links into this core. Bottlenecks at the Internet's edge can easily move between the wireless access (when its bandwidth is low) and the provider's uplink, both of which can have highly variable bandwidths.

Memory also became cheap; you cannot buy RAM chips small enough for the buffering in edge devices, and these devices have no mechanisms for self-limitation. Commodity network devices now span many downward-compatible generations: Ethernet has gone from 10Mbps to 10Gbps; wireless operates from 1Mbps to 100 or more Mbps; and cable from 10Mbps to, soon, several hundred Mbps. The result is a single buffer statically sized for larger bandwidths but much too large for lower-bandwidth links. For example, the 256 packets of buffering found in many of today's 802.11 device drivers alone translates to more than three seconds at 1Mbps, which is all the bandwidth you may have on some wireless networks. Complicating this is that recommendations about the amount of buffering have been influenced by early Internet

Figure 5. Plot reproduced from ICSI's Netalyzr studies.



problems of insufficient buffers, thus erring on the side of larger buffers, perhaps unaware that AQM is rarely used or unavailable.

Wireless links and networks are increasingly part of the edge access and are even more variable than broadband bandwidth: moving a device a few centimeters can change rates by one to two orders of magnitude; and because wireless is a shared medium, this also affects rates. Since bandwidth can vary by a factor of 100 at short time scales, static buffering is *never* appropriate.

A number of approaches to speeding up Web access contribute to transient access-link bloat by dumping large numbers of packets onto these links simultaneously.

Revisiting the Bandwidth-Delay Product

The efficacy of BDP-sized buffers is in question. As pointed out in a presentation at ACM SIGCOMM in 2004, BDP is not appropriate for highly multiplexed core links.¹ The rationale for maintaining a BDP buffer still applies at the network edge where a single flow can congest a link. The problem is in determining that BDP. Bandwidth variations of two or more orders of magnitude clearly play havoc with the bandwidth. At the same time, the 100ms delay assumption has been weakened by the advent of content delivery networks (CDNs) and other services engineered to bring common RTTs down to 10ms–30ms. Thus, even if an access link is a constant bandwidth and its buffer is sized to 100ms, it may still be 3–10 times too large.

For more than a decade, TCP tuning has been focused on improvements needed for high-BDP environments where large packet windows are required to achieve good throughput. These new algorithms are not in themselves nefarious, focusing on efficiently filling the pipe, but the researchers have unconsciously worked with a model of high bandwidth and AQM-enabled buffers. When the large pipe size comes from buffers rather than bandwidth, the algorithms efficiently fill those buffers, resulting in large delays. Controlling buffers makes it possible for one TCP to work well everywhere, a solution that is preferable to



Operating systems and hardware have an amazing number of buffer hiding places. As software and hardware is updated, more sources of bloat can be uncovered.



attempting to create a version of TCP specifically for access links.

Clearly there cannot be a “correct” static amount of buffering in the global Internet: *A self-adaptive AQM is the only viable long-term solution in any device (including your computer or smart phone) with a network buffer.*

AQM for the Modern World

In early 1998, Kathleen Nichols discovered flaws in RED and started to work with Jacobson to make improvements. At that time, the main concern was finding an algorithm that could be configured for any link by setting a single-rate parameter, as well as developing a viable approach to tracking persistent queue while ignoring short-term bursts.⁷ Subsequent research tried to fix some of the flaws but failed to create an AQM that could keep persistent buffers short without overdropping. Network operators faced only with algorithms requiring expert manual configuration that might hurt them have understandably been unwilling to enable and configure AQM.

In the ensuing decade wireless has been widely deployed, bringing wildly varying bandwidth to many edge links, cable Internet access has become common, and a device’s access bandwidth can easily vary by two orders of magnitude. It is now obvious that any AQM algorithm that does not take as an input the rate at which data leaves the buffers cannot work in today’s highly variable bandwidth environment. Clearly without such an algorithm, bufferbloat will be difficult to defeat.

Surprising to most, AQM is essential for broadband service, home routers, and even operating systems: It is not just for big Internet routers.

When does overbuffering hurt? Overbuffering hurts anytime you saturate a link; for example:

- ▶ Copying a file over the Internet.
- ▶ Running old versions of BitTorrent or other file-sharing application.
- ▶ Sending/receiving email messages to Grandma with pictures attached.
- ▶ Uploading video to YouTube.
- ▶ Web browsing, which can hurt you or others momentarily.

The saturated link can be anywhere, in either or both directions in the path: easiest and most common to see is the

operating system, wireless link, and broadband service.

Why is overbuffering a problem? Oversized buffers fill and cause delay, destroying many uses of the Internet:

- Stock traders and gamers do not want their competition to have even a 1ms advantage over them.

- To play music, jitter (variation in delay) and latency must be controlled and kept below 100ms.

- For something to “feel” attached to your hand (perfect rubber banding), latencies need to be below the 30ms range; for keyboard echoing to be imperceptible, 50ms.

- Speed-of-light latency dominates voice over IP (VoIP) over long-haul networks, making access latencies critical for keeping end-to-end latency below 150ms (the longtime telephony metric).

- Excessive packet loss induced by bufferbloat may cause DNS lookup failure.

- Essential network protocols such as ARP, DHCP, RA, and ND all presume timely response and can fail without it.

- Web browsing becomes painful as delays go from hundreds of milliseconds to multiple seconds.

Many service providers would like to be able to provide low-latency services in their networks to customers, whether remote gaming, hosted desktop systems, or backup. Solving the bufferbloat problem is necessary for

their successful deployment.

The Tip of the Iceberg

Operating systems and hardware have an amazing number of buffer hiding places. As software and hardware is updated, more sources of bloat can be uncovered. In particular, as older TCPs are replaced with modern ones, users not currently bloating their access buffers may suddenly experience much longer delays, (for example, Windows XP does not enable TCP window scaling, so it never has more than 64KB in flight at once).

Current commonly used network performance tests fail to test latency simultaneously with bandwidth: a link must become saturated for queuing delays to become obvious. It can take 10 seconds to fill the buffers of a broadband device, home router or operating system, and most consumer broadband tests do not test for that long—thus missing bufferbloat.

Excessive access delays tend to be written off as network congestion. Employing larger backbone pipes and rationing bandwidth use cannot improve performance for the users congesting access uplinks or viewing downloads through bloated buffers at the provider edge.

Mitigations

There are glimmers of hope. DOCSIS (Data over Cable Service Interface Specification) was modified in spring 2011

allowing cable operators to reduce buffering in cable modems. This mitigation will not take effect until 2012 at best and will require cable-modem firmware upgrades or (most likely) modem replacement, as well as motivated and knowledgeable operators.


Proper solutions for Web browsers can improve access-link behavior. These include HTTP/1.1 pipelining and Google’s SPDY (<http://dev.chromium.org/spdy>), both of which can achieve better performance, reduce the total number of packets and bytes transferred, and enable TCP to function better.

Some mitigations are simple and direct for the knowledgeable. A home router or your laptop, for example, almost never operates in the high-bandwidth environment for which the operating system has likely been tuned. Adjusting buffering in the operating system and/or device drivers can make a major improvement over the defaults. Unfortunately, while these adjustments may be accessible in your laptop, they may not be accessible in your home router or handheld devices.

Bandwidth shaping can be used to prevent bottleneck buffers from filling, but at a cost in bandwidth. Contrast the smokeping result in Figure 6 with that in Figure 2; almost two orders of magnitude improvement is not bad.

Most importantly, static, unmanaged buffers are inappropriate for modern network elements. Network architects and designers must adjust.

Figure 6. Smokeping of file transfer from Jim’s house to MIT after mitigations.



To address this, Nichols and Jacobson have resumed work on a robust, adaptive AQM.

Be Part of the Solution

The situation may worsen before it improves, and immediate action is necessary. Potential solutions must be subjected to rigorous testing and analysis before being widely deployed; otherwise, existing problems can be made worse. Unfortunately, today there is a distinct lack of funding for the kinds of performance monitoring, tuning, and improvement that characterized the early Internet.

The first step is to make the problem apparent. Consumer tests are important (for example, Speedtest.net, SamKnows, M-Labs (<http://www.measurementlab.net/>), Netalyzr), but better tests that point to the correct offender, usable by everyone, are badly needed. Consumer tests often perpetuate the mythology that more bandwidth means higher “speed,” and better marketing metrics are essential. Stuart Cheshire’s famous “It’s the Latency, Stupid” rant should be taken to heart.³

An open source project, CeroWrt, is under way at bufferbloat.net using OpenWrt to explore potential solutions, including AQM. Please help. A wide range of testing is needed for confidence in any algorithm. Since our operating systems are commodities and are used in today’s home routers, home-router bufferbloat is a direct result of host bufferbloat. Solve one, and you solve the other.

Unfortunately, since bufferbloat misleads TCP’s congestion-avoidance algorithm with respect to the effective pipesize, modern networks without effective AQM may again be vulnerable to congestion collapse from saturated edge buffers creating packet delays measured in seconds. Congestion collapse has been reported in a large-scale network, requiring complete shutdown and careful restart of the entire network to regain (temporary?) stability.

We are flying on an Internet airplane in which we are constantly swapping the wings, the engines, and the fuselage, with most of the cockpit instruments removed but only a few new instruments reinstalled. It crashed before; will it crash again?



The important location for buffers is at the path bottleneck, but the critical fast-to-slow transition can be different for different paths, different in the reverse path, and with dynamic bandwidths can change along the same path.



Acknowledgments

We would like to thank Dave Clark, Dave Reed, Vint Cerf, Van Jacobson, Vern Paxson, Nick Weaver, Scott Bradner, Rich Woundy, Greg Chesson, Dave Täht, and a cast of hundreds. 

Related articles on queue.acm.org

Whither Sockets?

George V. Neville-Neil

<http://queue.acm.org/detail.cfm?id=1538949>

Network Virtualization: Breaking the Performance Barrier

Scot Rixner

<http://queue.acm.org/detail.cfm?id=1348592>

You Don't Know Jack about Network Performance

Kevin Fall, Steve McCanne

<http://queue.acm.org/detail.cfm?id=1066069>

References

1. Appenzeller, G., Keslassy, I. and McKeown, N. Sizing router buffers. In *Proceedings of ACM SIGCOMM* (Portland, Aug. 2004).
2. Braden, R. et al. Recommendations on queue management and congestion avoidance in the Internet, RFC2309 (April 1998).
3. Cheshire, S. It's the latency, stupid (1996); <http://rescomp.stanford.edu/~cheshire/rants/Latency.html>.
4. Dischinger, M. et al. Characterizing residential broadband networks. In *Proceedings of Internet Measurement Conference* (San Diego, CA, Oct. 24–27, 2007).
5. Floyd, S. and Jacobson, V. Random Early Detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking* (Aug. 1993).
6. Jacobson, V. Notes on using RED for queue management and congestion avoidance. Talk at North American Network Operators' Group (1998) 13; <ftp://ftp.ee.lbl.gov/talks/vj-nanog-red.pdf>.
7. Jacobson, V. Notes on Using RED for Queue Management and Congestion Avoidance; <ftp://ftp.ee.lbl.gov/talks/vj-nanog-red.pdf>, see also <http://www.nanog.org/mtg-9806/agen0698.html>.
8. Jacobson, V. and Karels, M. Congestion avoidance and control. In *Proceedings of SIGCOMM '88*, (Aug. 1988).
9. Kreibich, C. et al. Netalyzr: Illuminating the edge network. In *Proceedings of the Internet Measurement Conference* (Melbourne, Australia, Nov. 1–3, 2010).
10. Nagle, J. On packet switches with infinite storage. In *Proceedings of Network Working Group RFC 9790* (Dec. 1985); www.ietf.org/rfc/rfc970.txt.
11. Reed, D.P. Congestion collapse definition (2009); <http://mailman.postel.org/pipermail/end2end-interest/2009-September/007769.html>.
12. Reed, D.P. What's wrong with this picture (2009); <http://mailman.postel.org/pipermail/end2end-interest/2009-September/007742.html>.
13. Rhee, I. and Xu, L. CUBIC: A new TCP-friendly high-speed TCP variant. *ACM SIGOPS* 42, 5 (2008).
14. Villamizar, C. and Song, C. High-performance TCP in ANSNET. *Computer Communications Review* 24, 5 (1994), 45–60.

Jim Gettys is at Alcatel-Lucent Bell Labs, USA, where he works on bufferbloat, as a properly working low-latency Internet is required for immersive teleconferencing. He was vice president of software at the One Laptop per Child project, editor of the HTTP/1.1 standard, and is one of the original designers of the X Window System.

Kathleen Nichols is the founder and CTO of Pollere Inc., a consulting company working in both government and commercial networking. She has 30 years of experience in networking, including a number of Silicon Valley companies and as a cofounder of Packet Design.

Decoupling a logical device from its physical implementation offers many compelling advantages.

BY CARL WALDSPURGER AND MENDEL ROSENBLUM

I/O Virtualization

THE TERM *VIRTUAL* is heavily overloaded, evoking everything from virtual machines running in the cloud to avatars running across virtual worlds. Even within the narrower context of computer I/O, virtualization has a long, diverse history, exemplified by logical devices that are deliberately separate from their physical instantiations.

For example, in computer storage, a logical unit number (LUN) represents a logical disk that may be backed by anything from a partition on a local physical drive to a multidisk RAID volume exported by a networked storage array. In computer networking, a virtual private network (VPN) represents a logically isolated private network, where the isolation is provided using cryptographic methods to secure data that may in fact traverse the public Internet. In computer architecture, an IOMMU (I/O memory-management unit) translates I/O-virtual memory addresses to corresponding physical memory addresses, making direct memory access by devices safe and efficient. Other examples of virtualization


include virtual LAN (VLAN), N_Port ID virtualization (NPIV), Intel Virtualization Technology for Directed I/O (VT-d), and Multiroot I/O Virtualization (MR-IOV).

The common theme is decoupling the logical from the physical, introducing a level of indirection between the abstract and the concrete. Such indirection has proven to be remarkably powerful and versatile. Modern virtualization platforms exploit indirection and abstraction in numerous ways.

A virtual machine (VM) is a software abstraction that behaves as a complete hardware computer, including virtualized CPUs, RAM, and I/O devices. A virtualization software layer, known as a *hypervisor*, provides the level of indirection that decouples an operating system and its applications from physical hardware. The term *guest* is commonly used to distinguish the layer of software running within a VM; a guest operating system manages applications and virtual hardware, while a hypervisor manages VMs and physical *host* hardware.

Although IBM invented and commercialized mainframe VMs many decades ago, VMs didn't make the leap to commodity hardware until the late 1990s, when VMware pioneered efficient virtualization on x86 platforms. Since then, virtualization has experienced a resurgence of interest in both industry and academia. Today, VMs are commonplace in many computing environments and nearly ubiquitous in enterprise data centers and cloud-computing infrastructures.

Since virtualization is a broad topic, and the universe of I/O devices is large and diverse, this article focuses on some representative I/O systems issues in VM-based systems, primarily in the context of a single physical host. After highlighting key benefits and challenges, we explore various implementation approaches and techniques that have been leveraged to enable flexible, high-performance I/O virtualization.



Benefits

Many of the benefits of virtualized systems depend on the decoupling of a VM's logical I/O devices from its physical implementation. Examples range from the ability to multiplex many VMs on the same hardware to advanced virtualization features such as live migration and enhanced security.

At the most basic level, decoupling enables time- and space-multiplexing of I/O devices, allowing multiple logical devices to be implemented by a smaller number of physical devices. Applications of virtualization such as server consolidation or running het-

erogeneous operating-system environments on the same machine rely on this feature. As inexorable trends create ever more powerful hardware, its not surprising that much of it remains seriously underutilized. The ability to multiplex logical I/O devices onto physical ones allows both administrators and automated systems to drive I/O devices at higher utilization and achieve better hardware efficiency. Much of virtualization's rapid adoption over the past decade can be attributed to the significant cost savings resulting from such basic partitioning and server consolidation.

Decoupling provides for flexible mappings between logical and physical devices, facilitating seamless portability. By supporting mappings of logical I/O devices to physical devices with different yet semantically compatible interfaces, virtualization makes VMs portable, even across heterogeneous systems. The same VM image can be run on computers with different I/O devices and configurations, with the I/O virtualization layer providing the necessary conversion.

Decoupling also enables popular VM features such as the ability to suspend and resume a VM and the ability

to move a running VM between physical machines, known as live migration. In both of these applications, active logical devices must be decoupled from physical devices and recoupled when the VM resumes after being saved or moved.

This virtualization layer may also change mappings to physical devices, even when the VM itself does not move. For example, by changing mappings while copying storage contents, a VM's virtual disks can be migrated transparently between network storage units, even while remaining in active use by the VM. The same capability can be used to improve availability or balance load across different I/O channels. For example, in a storage system with multiple paths between the machines and storage, the virtualization layer can rebind mappings to mask failures or avoid delays that might occur because of contention on paths.

I/O virtualization provides a foothold for many innovative and beneficial enhancements of the logical I/O devices. The ability to interpose on the I/O stream in and out of a VM has been widely exploited in both research papers and commercial virtualization systems.

One useful capability enabled by I/O virtualization is device aggregation, where multiple physical devices can be combined into a single more capable logical device that is exported to the VM. Examples include combining mul-

multiple disk storage devices exported as a single larger disk, and network channel bonding where multiple network interfaces can be combined to appear as a single faster network interface.

New features can be added to existing systems by interposing and transforming virtual I/O requests, transparently enhancing unmodified software with new capabilities. For example, a disk write can be transformed into replicated writes to multiple disks, so that the system can tolerate disk-device failures. Similarly, by logging and tracking the changes made to a virtual disk, the virtualization layer can offer a time-travel feature, making it possible to move a VM's file system backward to an earlier point in time. This functionality is a key ingredient of the snapshot and undo features found in many desktop virtualization systems.

Many I/O virtualization enhancements are designed to improve system security. A simple example is running an encryption function over the I/O to and from a disk to implement transparent disk encryption. Interposing on network traffic allows virtualization layers to implement advanced networking security, such as firewalls and intrusion-detection systems employing deep packet inspection.

Challenges

While virtualization offers many benefits, it also introduces significant challenges. One is achieving good I/O

performance despite the potential overhead associated with flexible indirection and interposition. Complex resource-management issues such as scheduling and prioritization are introduced by multiplexing physical devices across multiple VMs, further impacting performance. Another challenge is defining appropriate semantics for virtual devices and interfaces, especially when faced with complex physical I/O devices or system-level optimizations.


In many systems, a nontrivial performance penalty is associated with indirection. The same can be true for virtualized I/O, since I/O operations must conceptually traverse two separate I/O stacks: one in the guest managing the virtual hardware; and one in the hypervisor managing physical hardware. The longer I/O path affects both latency and throughput, and imposes additional CPU load.

Indeed, I/O-intensive workloads on some early virtualization systems suffered a virtualization penalty larger than a factor of two. Since then, further research, optimizations, and hardware acceleration have reduced this penalty into the noise for an impressive set of demanding production workloads. Somewhat counterintuitively, virtualized systems have even outperformed native systems on the same physical hardware, overcoming native scaling limitations by instead running several smaller VM instances in a scale-out configuration.

Figure 1 depicts the flow of an I/O request in a virtualized system. When an application running within a VM issues an I/O request, typically by making a system call, it is initially processed by the I/O stack in the guest operating system, which is also running within the VM. A device driver in the guest issues the request to a virtual I/O device, which the hypervisor then intercepts. The hypervisor schedules requests from multiple VMs onto an underlying physical I/O device, usually via another device driver managed by the hypervisor or a privileged VM with direct access to physical hardware.

When a physical device finishes processing an I/O request, the two I/O stacks must be traversed again, but in the reverse order. The actual device posts a physical completion interrupt, which is handled by the hypervisor. The

Figure 1. An I/O request issued by an application is processed first by the guest operating system I/O stack running within the VM, and then by the hypervisor I/O stack managing physical hardware.



hypervisor determines which VM is associated with the completion and notifies it by posting a virtual interrupt for the virtual device managed by the guest operating system. To reduce overhead, some hypervisors perform virtual interrupt coalescing in software, similar to the hardware batching optimizations found in physical cards, which delay interrupt delivery with the goal of posting only a single interrupt for multiple incoming events.

Interposition can incur additional overhead by manipulating I/O requests such as inspecting network packets to perform security checks or encrypting disk writes transparently. In some cases, the interposition costs are negligible, especially compared with high-latency operations such as I/O to traditional rotating media. In other cases, even making an extra in-memory copy of I/O data may be prohibitively expensive—for example, for fast networks with extremely high packet rates. To improve performance, some hypervisors parallelize portions of this processing, offloading work to additional processor cores. Of course, when there is contention for CPUs, this leaves fewer cores available for running VMs.

Managing resources in virtualized systems presents additional challenges. Although each VM is presented with the illusion of having its own dedicated virtual hardware, in reality the hypervisor must multiplex limited physical hardware across multiple VMs of varying importance, mapping their virtual resources onto available physical resources. At the most basic level, contention for a physical device will result in scheduling delays for some VMs. At a minimum, the hypervisor must prevent VMs from monopolizing resources and denying service to others.

More generally, the hypervisor should provide some measure of performance isolation or quality-of-service controls to reflect the relative importance or absolute requirements of diverse VM workloads. The ability to express resource-management policies is especially important when physical resources are shared across multiple users or organizations, as is common in multitenant cloud-computing environments. Several hypervisors support a relative weight control, where a VM's allocation is directly pro-



I/O virtualization provides a foothold for many innovative and beneficial enhancements of the logical I/O devices.



portional to its weight. Some also provide absolute reservation and limit settings, which bound a VM's minimum and maximum allocation, regardless of system load.

For I/O devices that can be accessed concurrently by VMs on different hosts, such as networked storage arrays, resource management requires distributed algorithms to schedule requests fairly and efficiently. Virtualization platforms have only recently started offering sophisticated solutions capable of providing end-to-end quality of service for VM I/O bandwidth and latency.

Scheduling may also impact VM performance in subtler ways. For example, contention for CPU resources can cause problems for TCP networking performance. TCP connections rely on accurate round-trip time (RTT) estimates in order to perform flow control and adjust window sizes appropriately. A VM, however, may be descheduled for tens or even hundreds of milliseconds while a packet is pending. As a result, CPU time-multiplexing can distort a VM's RTT values, causing its congestion windows to grow too slowly, which degrades throughput significantly. To solve this problem, some researchers have proposed offloading more TCP functionality to the hypervisor. Another option would be to present VMs with virtual network interface controller (NIC) hardware that supports optional TCP Offload Engine (TOE) functionality, as found in some physical NICs.

The very idea of adding TCP offload capabilities to a virtual NIC highlights the difficulty of choosing appropriate semantics for virtual hardware. This is especially true for devices with even more complex interfaces, such as modern graphics cards. At one extreme, virtual hardware can have an interface identical to a physical device. This approach has the compelling advantage of compatibility with all software that already supports (or will support) the physical device. Unfortunately, such transparency usually comes at the cost of emulating a fairly complex virtual device interface that was not designed to support virtualization efficiently. At the other extreme, virtual hardware can have a completely new hypervisor-specific interface, designed explicitly to be simple and efficient.

A related challenge is ensuring that virtualization faithfully preserves the semantics that software expects of physical devices. For example, some virtualization systems boost I/O performance by leveraging a hypervisor-level buffer cache between the VM and physical storage. While caching reads does not introduce any problems, caching writes would violate the durability semantics relied upon by guest file systems, databases, and other software. In a native, unvirtualized system, an I/O completion indicates that a write has been committed. Hypervisors must use a write-through cache to preserve this property, although some provide an explicit option to trade off safety for performance by relaxing this constraint.

Direct memory access (DMA) illustrates additional safety and performance issues. It enables an I/O device to read and write host RAM directly without involving the CPU, which is critical for achieving high-performance I/O rates. Unfortunately, giving devices the ability to use DMA to reach arbitrary physical memory locations is risky, especially since the majority of operating-system bugs result from misbehaving device drivers. As discussed in the next section, virtualization systems can ensure strict isolation between VMs by employing various approaches, such as leveraging hardware IOMMU functionality at both the guest and hypervisor levels.

Approaches

The classic way of implementing I/O virtualization is to structure the software in two parts: an emulated virtual device that is exported to the VM and a back-end implementation that is used by the virtual-device emulation code to provide the semantics of the device. Modern hypervisors support an I/O virtualization architecture with a split implementation, as shown in Figure 2, where a virtual machine can select among different virtual device interface emulation front-ends as well as multiple different backend implementations of the device. For example, the virtual machine can be configured with an IDE, a SCSI, or a paravirtualized disk device that is implemented using a file, a local disk, or a storage area network (SAN). Here, we describe the way these are implemented in modern virtualization systems and discuss some of the



The classic way of implementing I/O virtualization is to structure the software in two parts: an emulated virtual device exported to the VM and a back-end implementation used by the virtual-device emulation code to provide the semantics of the device.



available optimization options.

As a concrete example, consider a legacy PC I/O device such as an IDE disk. The operating system in the VM would call a device driver to launch disk read or write requests. The device driver would include OUT instructions that program the operation type, device number, disk-sector number, length, and buffer-memory location for the operation. The driver assumes that the device will use DMA to transfer the contents between memory and the disk device and then raise an interrupt when it is done. For the emulated device to work correctly, the emulation software must catch and interpret the OUT instructions to determine the correct operation and its arguments; perform an emulation of the operation by storing or fetching the requested block of storage, using the emulation of the architecture's DMA capability to read or write memory; and finally raise the proper interrupt signal on the VM to notify the driver that the request is finished.

While the device emulation code is specific to the particular device being emulated (for example, an IDE disk), the semantics of the operations being performed are general and frequently constructed so that the same device emulation can access multiple different back-end implementations. For example, with virtual disks, the back-end implementation could be as simple as forwarding the request untransformed to a native physical IDE controller or as complex as implementing the storage for the virtual disk as a file in a host operating-system file system, as in many desktop virtualization products. In the latter case, the back end must generate host operating-system file-system read and write operations for the file containing the virtual disk contents in order to perform emulated virtual disk-sector read and write operations.

A pluggable structure for back-end implementations makes it easy to generate new capabilities for virtual devices. A disk-storage back end can implement an emulated CD-ROM device for a VM simply by accessing a file containing its ISO (International Organization for Standardization) image. Similarly, the snapshot and undo capabilities in modern virtualization systems can be

implemented by logging rather than overwriting a virtual-disk file, enabling the virtualization layer to control which version of the disk is visible to the VM.

Along with flexibility and innovative features come potential performance penalties, which can vary greatly. A back end optimized for a server machine that efficiently shuttles emulated reads and writes to a portion of a local disk may have very low virtualization overheads. In contrast, the penalty may be relatively large for a desktop-virtualization back end where the data is stored in an encrypted file on a network file system used by the host operating system. Not only would every read and write request traverse the file system and networking code of the host operating system and the network file system, but additional encryption and decryption overhead would also be incurred for each sector.

Much current research in I/O virtualization is focused on either new interposition functionality that solves some problem or optimizations for reducing the overheads associated with the virtualization. Optimization has been an especially important goal in server virtualization and consolidation, where overheads directly impact metrics such as the number of VMs supported per server. Beyond back-end optimizations, reducing I/O-virtualization overhead requires decreasing virtual-device emulation costs. Several recent optimizations have tried exactly this using either software or modifications to the I/O device hardware.

One of the software techniques for lowering emulation costs is to reduce the number of trap-and-emulate operations the hypervisor is required to do to perform an I/O operation. For example, the legacy PC IDE interface uses eight-bit OUT instructions to communicate with the disk controller. Communicating the sector number, buffer address, and length requires multiple such instructions, causing repeated traps into the hypervisor to run the device emulation code. Using the driver for an alternative disk device such as a SCSI disk can achieve the same functionality with far fewer traps, greatly reducing the emulation overheads.

Even further reduction can be achieved by optimizing the communication between the VM software and

the device emulation. In modern operating-system environments such as Windows and Linux, it is possible to install device drivers that communicate the request's arguments to the hypervisor's device emulation code directly, with minimal overhead. This approach of using virtual hardware optimized for the virtualization layer rather than matching any particular real device is referred to as *paravirtualization*. In practice, most modern virtualization platforms support an emulated legacy device for compatibility, as well as providing an optional paravirtual device for higher performance.

As an example, a paravirtualized disk interface could have the device emulation code accept commands via a memory segment shared between the driver and the emulation, allowing communication of commands with practically zero overhead. The emulation code simply passes the command to the optimized back-end implementation. Examples include Xen's virtual block-device front-end driver and VMware's PVSCSI guest disk driver.

For an I/O device used by only a single VM, where the back-end implementation is mostly passing through the driver commands from the VM to the device, it is tempting to pass through the device directly, assigning it to the VM exclusively. Consider the example of a high-performance NIC that is used by only one of the VMs running on a computer. It is relatively easy to configure the CPU virtualization so the x86 instructions that talk to the device can be connected directly to the device and incur zero I/O virtualization overheads. This *pass-through mode* can eliminate both the device emulation and back-


end implementation overheads.

Although pass-through mode can remove I/O virtualization overheads, it introduces several limitations and implementation challenges that have slowed its deployment. Aside from the obvious limitation that each pass-through device can be used by only a single VM, pass-through forms a coupling between the hardware and the VM. As a result, many of the portability benefits of virtualization are lost, along with key benefits such as live migration and features that depend on the ability to interpose on I/O.

One of the biggest challenges of pass-through mode affects devices that use DMA. The fundamental problem is that the driver in the VM will program device DMA using the guest's notion of memory addresses, which differ from the real memory addresses in which the VM's memory resides. This is not only incorrect, but also a large safety and security problem since the device could read and write memory potentially belonging to the hypervisor or some other VM. To make this work, the VM's driver must translate memory addresses to use the correct real memory before programming them into the device. This exposes the driver to the details of hypervisor memory virtualization and still has safety problems since bugs in the driver can result in incorrect translations.

To eliminate both the limitations and the challenges of pass-through, device builders have modified their hardware to be aware of the virtualization layer. To handle the limitation of exclusive pass-through-only devices, such virtualization-aware hardware exports

Figure 2. Modern hypervisors split device virtualization into a front-end emulation and a backend implementation. For example, virtual machines can be configured with IDE, SCSI, or paravirtualized disk devices that are implemented either as a file, a local disk, or a SAN storage area network. The emulated and physical devices, such as the SCSI front-end and local SCSI back-end, may differ in size and other attributes.



multiple interfaces, each of which can be attached to a different VM. As a result, each VM is given its own directly accessible pass-through copy of the device. For example, a virtualization-aware NIC could have many personalities that look and act as if many separate NICs had been directly mapped into different VMs.

Additional hardware support is needed to address the challenges of performing DMA operations directly involving VM memory. A memory management unit is employed to map the DMA operation's memory addresses into the correct locations in the VM's memory. This mapping hardware—the IOMMU—is programmed for each VM attached to a device with the mappings of where the VM resides in memory. Each DMA request is run through the IOMMU, which routes the request to or from the correct location in the real machine memory or generates an error if the request is not valid. The IOMMU allows the driver in the VM to program device DMA using its virtualized notion of memory addresses, while still allowing the hypervisor to decide where VM memory is actually located in physical machine memory. The IOMMU also provides a level of safety, ensuring that even buggy driver software in the guest cannot generate DMA accesses to locations outside the VM.

Although IOMMUs can safely and efficiently allow virtualization-aware I/O devices to access the memory of a virtual machine directly, there are implications for some of the more sophisticated memory virtualization operations in modern hypervisors that rely upon dynamic page remapping. Consider features such as overcommitting memory, where the hypervisor can reclaim RAM via techniques such as demand-paging VM memory to secondary storage; memory compression; or transparent memory sharing where identical pages can be de-duplicated by sharing them read-only between multiple virtual machines. These features require that certain accesses to VM memory take faults and invoke hypervisor actions before they are allowed to proceed. DMA devices need to respect this, and hence the devices need to support something similar to page faults on DMA operations where the hypervisor is invoked before the

DMA is permitted to finish. The ability to tolerate arbitrary delays on DMA operations can have a much deeper impact on the changes needed to make an I/O device virtualization-aware.

While production hypervisors employ hardware IOMMUs and other hardware-enforced memory-mapping techniques to guarantee isolation between VMs, they have not yet included IOMMUs in the virtual hardware presented to VMs. A virtual IOMMU (vIOMMU) would allow a guest operating system to defend against its own buggy device drivers, as in a native system.

Recently, researchers have developed new IOMMU emulation techniques for efficiently providing vIOMMUs to guests. Even more significantly, the same approach facilitates a more flexible form of device pass-through, where a VM is allowed to interact with a directly-assigned I/O device without hypervisor intervention. Since the guest exposes which regions of its memory are currently involved in DMA operations to the vIOMMU, the hypervisor is able to modify mappings for other memory regions safely. By interposing only on vIOMMU operations, it is possible to achieve near-native I/O performance while preserving the hypervisor's ability to manage, re-map, and overcommit memory.

Conclusion

Decoupling a logical device from its physical implementation offers many compelling advantages. A single physical device can be multiplexed, allowing it to act as many virtual devices, improving hardware utilization. Abstracting away details about specific hardware and physical location makes seamless migration possible. The level of indirection between virtual and physical also provides a convenient hook for interposing on I/O operations transparently, enabling new capabilities such as replication, load balancing, encryption, and security checks.

A key challenge in I/O virtualization is achieving these benefits with minimal overhead. A number of clever software and hardware approaches have been devised to achieve high-performance indirection and interposition, including paravirtualization and virtualization-aware devices. Resource-management issues, such as scheduling and prioritization, become important when

device multiplexing is used to consolidate different workloads onto the same physical hardware. Defining clean interfaces and appropriate semantics for virtual devices is also challenging.

I/O virtualization remains an active area of research and development in both academia and industry. Although we focused here on systems issues from the perspective of an individual physical machine, the broader context of I/O virtualization includes an enormous range of work on distributed systems and the fabrics that connect their virtual and physical components. The increasing prevalence and commercial success of systems based on virtual machines is certain to fuel demand for new virtualization optimizations and I/O capabilities. C

Related articles on queue.acm.org

The Virtualization Reality

Simon Crosby, David Brown

<http://queue.acm.org/detail.cfm?id=1189289>

Network Virtualization: Breaking the Performance Barrier

Scot Rixner

<http://queue.acm.org/detail.cfm?id=1348592>

The Cost of Virtualization

Ulrich Drepper

<http://queue.acm.org/detail.cfm?id=1348591>

Suggested Reading

1. Ahmad, I., Gulati, A. and Mashtizadeh, A. vIC: Interrupt coalescing for virtual machine storage device IO. In *Proceedings of the 2011 Usenix Annual Technical Conference* (June 2011).
2. Amit, N., Ben-Yehuda, M., Tsafir, D. and Schuster, A. IOMMU: Efficient IOMMU emulation. In *Proceedings of the 2011 Usenix Annual Technical Conference* (June 2011).
3. Gamage, S., Kangarlou, A., Kompella, R.R. and Xu, D. Opportunistic flooding to improve TCP transmit performance in virtualized clouds. In *Proceedings of the Second ACM Symposium on Cloud Computing* (Oct. 2011).
4. Gulati, A., Ahmad, I. and Waldspurger, C. PARDA: Proportional allocation of resources for storage access. In *Proceedings of the Seventh Conference on File and Storage Technologies* (Feb. 2009).
5. Sugerman, J., Venkitachalam, G. and Lim, B.-H. Virtualizing I/O devices on VMware workstation's hosted virtual machine monitor. In *Proceedings of the 2001 Usenix Annual Technical Conference* (June 2001).

Carl Waldspurger is currently working with stealth-mode startups. For most of the previous decade, he was a principal engineer at VMware, where he was responsible for core resource management and virtualization technologies, including scheduling, memory management, and distributed systems.

Mendel Rosenblum is an associate professor in the computer science and electrical engineering departments at Stanford University. His research interests include system software, distributed systems, and computer architecture. He was cofounder of VMware Inc. and active in the R&D of VMware's virtualization platform.

© 2012 ACM 0001-0782/12/01 \$10.00

ADVANCE YOUR CAREER WITH ACM TECH PACKS...

For Serious Computing Professionals.



Searching through technology books, magazines, and websites to find the authoritative information you need takes time.

That's why ACM created "Tech Packs."

- Compiled by **subject matter experts** to help serious computing practitioners, managers, academics, and students understand today's vital computing topics.
- Comprehensive **annotated bibliographies**: from ACM Digital Library articles, conference proceedings, and videos to ACM Learning Center Online Books and Courses to non-ACM books, tutorials, websites, and blogs.
- **Peer-reviewed** for relevance and usability by computing professionals, and **updated** periodically to ensure currency.

Current topics include **Cloud Computing**, **Parallel Computing** and **Enterprise Architecture**. In development are Mobility, Business Intelligence/Data Management, Security, and more.

Suggestions for future Tech Packs? Contact:

Yan Timanovsky
ACM Education Manager
timanovsky@hq.acm.org

Access to ACM resources in Tech Packs is available to ACM members at <http://techpack.acm.org> or <http://learning.acm.org>.



Association for
Computing Machinery

Advancing Computing as a Science & Profession

DOI:10.1145/2063176.2063197

**Looking past the systems people use,
they target the people using the systems.**

BY JASON HONG

The State of Phishing Attacks

PHISHING IS A kind of social-engineering attack in which criminals use spoofed email messages to trick people into sharing sensitive information or installing malware on their computers. Victims perceive these messages as being associated with a trusted brand, while in reality they are only the work of con artists. Rather than directly target the systems people use, phishing attacks target the people using the systems. Phishing cleverly circumvents the vast majority of

an organization's or individual's security measures. It doesn't matter how many firewalls, encryption software, certificates, or two-factor authentication mechanisms an organization has if the person behind the keyboard falls for a phish.

On the surface, phishing attacks may seem to be a variant of spam. However, such attacks can lead to damaging losses in terms of identity theft,^{14,25} sensitive intellectual property and customer information, and national-security secrets.

Phishing attacks are also increasingly pervasive and sophisticated. Phishing has spread beyond email to include VOIP, SMS, instant messaging, social networking sites, and even massively multiplayer games.^{4,6,35} Criminals have also shifted from sending mass-email messages, hoping to trick anyone, to more selective "spear-phishing" attacks that use relevant contextual information to trick specific victims.

Academic research and commercial work in phishing is a dynamic area combining social psychology, economics, distributed systems, machine learning, human-computer interaction, and public policy. In 2006, Jakobsson and Myers²⁰ published an overview of how phishing works and what countermeasures were available at the time. This article serves as an introduction, as well as overview, of the

» key insights

- **Phishing attacks initially targeted general consumers, aiming to steal identity and credit-card information, but evolved to also include high-profile targets, aiming to steal intellectual property, corporate secrets, and sensitive information concerning national security.**
- **Developers must go beyond blaming users if they expect to deploy effective countermeasures against phishing attacks.**
- **The three general strategies for protecting end users from phishing scams: make things invisible, develop better user interfaces, and provide effective training.**



current state of phishing. It starts with how phishing attacks work, why people fall for them, the debate over the actual damage they cause, and finally a survey of countermeasures against phishing.

Anatomy of an Attack

Phishing attacks involve three major phases: The first is potential victims receiving a phish; the second is the victim taking the suggested action in the message, usually to go to a fake Web site but can also include installing malware or replying with sensitive information; and the third is the criminal monetizing stolen information.

Fake phishing email. Most phishing email messages use social techniques rather than technical tricks to fool end users. Conveying urgency is a well-known method used by criminals to misdirect people's attention;³⁴ an example is pretending to be a system administrator warning people about a new attack, urging them to install the attached patch. Another is notifying people there have been multiple failed logins for their account and they must verify their account now or risk dire consequences.

Appealing to people's sense of greed is an ancient technique now adapted to

the digital world. One phish the author of this article almost fell for about five years ago was filling out a survey for a bank in return for a small amount of money. The survey seemed innocuous until it asked for a bank-account number for depositing funds. So-called Nigerian 419 scams, offering "free" money in exchange for helping the sender move large amounts of money, also fall into this category. However, such obvious get-rich-quick scams are morphing to appeal to other emotions. Phishers today might pose as a relief agency asking for help with a recent natural disaster or as a random person appealing to

prurient interests, as in, say, “see Britney Spears naked.”

More sophisticated spear-phishing attacks use specific knowledge of individuals and their organizations; for example, an attack on military personnel might contain an invitation to a general’s retirement party, asking recipients to click on a link to confirm they will attend. People who wouldn’t normally fall for phish might in this case, due to the context. Jagatic et al.¹⁹ experimented in 2007 with how to exploit social-network information, showing that people were 4.5 times more likely to fall for phish sent from an existing contact over standard phishing attacks. Criminals indeed heavily target online social-networking sites partly for this reason.

Spear-phishing is also being used against high-level targets, in a type of attack called “whaling”; for example, in 2008, several CEOs in the U.S. were sent a fake subpoena along with an attachment that would install malware when viewed.²⁶ A *Communications* blog entry¹⁶ outlined several successful spear-phishing attacks in late 2010 and early 2011, with victims including the Australian Prime Minister’s office, the Canadian government, the Epsilon mailing list service, HBGary Federal, Oak Ridge National Laboratory, and RSA SecurID.

Setting up fake Web sites. Most phishing attacks try to convince people to go to a fake site where personal information is collected. To host a fake site, scammers use free Web space and a compromised machine or register a new domain.²⁷

When registering new domains, criminals look for names similar to the site they want to impersonate; for example, impersonating eBay, scammers might register ebay-login.com. Criminals also commonly use homograph attacks that exploit the visual similarity of characters; for example, bankofthevest.com⁸ uses two v’s to look like a w. Internationalized domain names facilitate this kind of attack, since characters in different language sets may appear identical.

However, in practice, criminals have opted for even simpler approaches. One is to put the domain name in plain sight, as in, say, paypal.com.phishsite.com. Surprisingly, many attacks make no attempt to disguise the destination site,



It doesn’t matter how many firewalls, encryption software, certificates, or two-factor authentication mechanisms an organization has if the person behind the keyboard falls for a phish.



relying on people’s lack of understanding of URLs. Unfortunately, even these simple tactics still fool many people.

When phishing attacks were just starting, scammers would create Web pages by hand, so they tended to be of poor quality, often including misspellings and hotlinks to images on the original site. The majority of phishing sites today are created with toolkits that might let a phisher specify what legitimate page to copy and where to direct stolen data, then generate all needed content. In 2008, Cova et al.⁷ identified more than 500 working kits. One surprising finding was that over one-third of these toolkits would send phished information to a location different from the one specified by the phisher, targeting inexperienced criminals who would do the work (and bear the legal risk) of breaking into sites.

When phishing attacks began, law enforcement, industry, and academic researchers were not organized in preventing and responding to attacks. However, as countermeasures (such as blacklisting and takedowns) were deployed (discussed later in the section on invisible attacks), criminals began introducing new techniques, thus starting an arms race that continues today. The most innovative approach so far is called “fast flux,” using a large pool of proxies and domain names to hide the true location of a phish. Fast flux makes it more difficult to blacklist sites since many URLs must be checked manually. Finding and taking down offending sites is also difficult since more work is needed to find the actual server. While an average phishing site lasts an average of 62 hours before being taken down, sites using fast flux tend to last an average of 196 hours.²⁷

Monetizing stolen information. The final phase of phishing is the monetization of stolen information. In some cases, the path is direct (such as when stealing banking credentials). In other cases, the path is convoluted (such as when stealing credentials for online games and social networking sites). Criminals have shown ingenuity here; for online games, they might transfer all of a victim’s virtual gold to an accomplice, then sell it to other players for real money. These attacks are common enough that Blizzard Entertain-

ment, creator of the popular online game *World of Warcraft*, sells special authenticators and offers in-game gifts for using them.⁵

Phishing on social networks is also somewhat indirect in terms of monetization. One attack involves notifying the victim's friends that the person is in trouble and needs money fast. Another involves using compromised accounts to spread malware; for example, the Koobface worm sends messages to a victim's friends urging them to go to a site that contains malware. Another is to steal victims' passwords and break into their email and bank accounts, working all too well since many people reuse passwords and existing password-reset mechanisms send responses directly to an individual's email address.

Marketplaces have also developed for acquiring and trading legitimate credentials. In the early days of phishing, phishers might use stolen credentials directly; today, many sell such credentials through underground networks to other criminals. These purchasers in turn might recruit unsuspecting people as "mules" to launder money and goods, reduce the risk the criminals face, and circumvent existing countermeasures; for example, some so-called "work at home" jobs involve receiving money transfers into the mule's bank account, with the funds actually coming from a hacked bank account. The mule then wires the money to a different account in another country, keeping a small commission. Such activities are illegal, and many perpetrators have been indicted around the world.²²

This evolution in how stolen credentials are monetized is due to specialization and perceived risk. A person good at creating phishing sites might not necessarily be good at stealing money from the accounts, especially given increasing vigilance by banks and law enforcement. Thus, rather than risk being traced, a phisher could opt to sell stolen information to others who are less risk averse.

Many researchers have examined how criminals trade stolen information on open Internet Relay Chat (IRC) channels. Herley and Florencio¹⁵ found that criminals often sell credentials for pennies on the dollar, explaining the

situation as a classic case of a marketplace for lemons. Given the anonymity of IRC, sellers find it easy to swindle purchasers by offering fake credentials or selling the same ones multiple times. Likewise, it is also easy for law enforcement and banks to offer honeypot credentials; as such, it is difficult for buyers to assess the quality of stolen data before buying. This asymmetric information about sellers and their goods leads buyers to dramatically lower what they are willing to pay.

Why We Fall for Attacks

An unfortunate response by the technically savvy is to dismiss end users as stupid and gullible, but it overlooks the fact that phishers deliberately exploit the poor usability of many interfaces that provide few cues for assessing the legitimacy of email messages and Web sites. Moreover, a deeper understanding of end-user motivations, beliefs, and mental models is essential for the security community to build effective countermeasures.

Dhamija et al.⁸ conducted one of the earliest studies (2006) investigating why people fall for phishing scams, asking participants to identify various Web sites as legitimate or fake. They found that good phishing sites fooled 90% of their participants and that most browser cues were opaque. Many par-

ticipants incorrectly judged sites based on their content and how professional they appeared, not realizing that Web pages are easily copied. Dhamija et al. also found that even experienced participants in the study had trouble with picture-in-picture attacks showing screenshots of a Web browser at a given site (see Figure 1). Picture-in-picture attacks point to an even greater challenge—that many people cannot differentiate between the browser "chrome," or buttons and URL area, that can mostly be trusted, and the browser content area, where attackers can show whatever they want.

Also in 2006, Downs et al.⁹ conducted a complementary study examining phishing email messages. As in Dhamija et al.,⁸ Downs et al. found their participants used basic, often incorrect heuristics in deciding how to respond to email messages; for example, some participants reasoned that since the business already had their information, it would be safe to give it again.

Sheng et al.³⁰ conducted a follow-up study involving a large-scale survey examining demographics and phishing susceptibility. Surprisingly, they found women were more vulnerable to phishing than men, primarily due to women having less exposure to technical knowledge. They also found that younger participants (ages 18 to 25) per-




Figure 1. Picture-in-picture phishing attacks show an image of a Web browser in the content area of a Web browser and fool even experienced users; adapted from Jackson et al.¹⁸

formed worse than all other age groups, possibly due to fewer years of experience online, less exposure to training materials, and less aversion to risk.

How Bad?

The Anti-Phishing Working Group (<http://www.antiphishing.org>) is an international consortium of law enforcement, industry, and academic researchers devoted to combating Internet scams and online fraud. For phishing attacks against consumers, the peak was in the third quarter of 2010 when APWG identified more than 115,000 unique phishing email messages worldwide, along with more than 150,000 unique phishing sites worldwide.³ Subsequent APWG reports have shown a slow downward trend.

However, estimates of damage caused by phishing varies widely, ranging from \$61 million per year¹⁴ to \$3 billion per year²⁵ of direct losses to victims in the U.S. The main problem is a lack of data from banks and other institutions that suffer losses; as such, these estimates are heavily dependent on the methods used and assumptions made by the organizations compiling and reporting the statistics.

While there is not yet full agreement among security analysts regarding how to calculate direct damages, there is increasing agreement that the indirect costs of phishing are substantial. One bank the author has spoken to said it cost it about \$1 million per attack in terms of call-center costs, recovery costs, and actual money (relatively small) that could not be recovered. A more difficult metric to measure is the damage to personal and corporate reputations. In presentations on the economics of computer security, Cormac Herley of Microsoft Research captured the problem succinctly, asking: "What is the first thing you think of when you hear the words 'Nigerian businessman'?"

Estimates of direct costs to the public also fail to capture the damage from specialized spear-phishing attacks. A number of successful high-profile attacks were reported in 2011, victimizing agencies in the Australian and Canadian governments, the Epsilon mailing list service, Gmail, Lockheed-Martin, Oak Ridge National Laboratory, and RSA.¹⁶

In 2009, the Operation Aurora attacks used spear-phishing and malware to target a number of organizations, most notably Adobe, Google, Symantec, and Yahoo!. In many cases, the attackers stole source code and other intellectual property. However, there are no good estimates as to the damage caused by spear-phishing, due to victims' unwillingness to share information and the basic difficulty in assessing damages.

Countermeasures

Given the risks of phishing, what can individuals and organizations do to protect themselves? Taking an end-user perspective, three general strategies are: make things invisible, so users do not have to do anything different; provide better user interfaces that either make the situation more obvious to users or offer additional protection; and train end users to proactively recognize and avoid phishing attacks. All three are needed to provide the strongest possible protection against phishing attacks.

Make it invisible. The first line of defense is to prevent phishing attacks from reaching end users in the first place. Solutions include filtering phishing email messages, blocking fake sites, and taking down fake sites.

Filtering phishing email. A large body of research covers spam detection. However, research on detecting phishing email messages is sparse, in part because phishing is a relatively new phenomenon and because phishing email messages sent by technically savvy criminals are so convincing. Fette et al.¹¹ developed the first email phishing filter (in 2007) identifying several features that are highly indicative of phishing, including, for example, URLs that use different domain names. Researchers and developers alike have since then explored additional features and machine-learning techniques.

An alternative to heuristics is authentication and verification technologies. For example, the Sender Policy Framework (SPF) uses the Simple Mail Transfer Protocol (SMTP) to reject forged email addresses. Another is DomainKeys Identified Mail (DKIM), verifying the DNS domain of a sender and message integrity. However, all have proved difficult to deploy on a large scale and do not provide protection

against several kinds of phishing attacks;¹³ for example, while focusing on preventing email spoofing, attackers easily create alternative fake addresses.

Blocking phishing sites. There are two general ways of detecting phishing Web sites: heuristics that examine the URL, HTML, and server characteristics to classify sites; and manually verified blacklists.

For heuristics, researchers have investigated a number of ideas using machine learning; examples include looking for patterns in URLs,¹² words in Web pages,¹ and search engines.³⁹ Researchers have also looked at linguistic characteristics of Web pages, identifying the brand name a Web page claims to be.³⁷ The effectiveness of all such techniques is reasonable, with true positive rates (correctly identifying a phishing site) of 90% or better and false positive rates (incorrectly labeling a legitimate site as phish) approaching 1% or less.

The best-known anti-phishing blacklists are operated by Google, Microsoft, and PhishTank, each containing URLs manually verified as phish. Google's is integrated with Firefox and Chrome, so no special action is required of end users to protect themselves. Microsoft's is integrated with Internet Explorer. And PhishTank's uses a wisdom-of-crowds approach to identify phish, letting people submit potential phish. Once enough other people vote that a submission is indeed a phish, it is added to their blacklist. Since October 2006, PhishTank has received close to four million votes from volunteers, labeling more than half-a-million phishing sites.²⁸

Several commercial browser add-ons are designed to block phish. Since they can be installed in Web browsers, their effectiveness can be evaluated empirically. For example, in 2009, Sheng et al.³³ examined major blacklists and browser tools, showing that zero-hour protection offered by blacklists had a false positive rate of 0% but a true positive of less than 20%. Even after 12 hours, the best blacklist identified only 83% of phish. They also found that deployed heuristics were somewhat effective in identifying phish but were used only to warn people in the Web browser rather than block likely phishing sites.

Sheng et al.'s research identified a gap between research and industry in terms of true positives. Academic research has generally focused on heuristics and machine-learning techniques with very good true positives though somewhat high false positives. These heuristics are good at identifying phishing sites not seen before. On the other hand, industry relies primarily on blacklists, which have middling true positives but no false positives. However, blacklists do not generalize well to future unseen cases, can be slow to respond to zero-hour attacks, and are easily overwhelmed by automatically generated URLs, a tactic phishers have already adopted.

In follow-up work, Sheng et al.³¹ probed the issue of heuristics vs. blacklists by interviewing people in industry, law enforcement, and academia, finding that concern over liability for false positives is the major barrier to deploying more aggressive heuristics. However, the first few hours of an attack are critical for blocking it, as a substantial fraction of users will have read their email by the time blacklists are updated. Jagatic et al.¹⁹ found that during regular work hours, most users who fell for a phishing attack did so in the eight hours following the start of the attack.

Sheng et al.³¹ identified several ways to ameliorate the situation; foremost is to clarify the legal issues surrounding false positives. Another is to have a central clearinghouse for phish, rather than piecemeal efforts that take longer to identify phish due to duplicated effort. A third is for researchers to develop better heuristics that minimize false positives. An early example of such heuristics was developed by Xiang et al.³⁸ observing that many phish are near or exact duplicates because they are generated by toolkits. Once a phish is on a blacklist, other copies of it can be identified quickly and blocked with virtually no risk of false positives. Using probabilistic-matching methods, the obvious countermeasure of adding noise can also be mitigated.

Taking down phishing sites. Several organizations identify and take down phishing sites, and private mailing lists help share information about fake sites, as well as find contact information for specific ISPs and Web sites.

When phishing sites are taken down, end users who click on a phish are typically shown a “page not found” error. One innovation developed by APWG and Carnegie Mellon University is to have ISPs and take-down providers replace the phishing page with a training message, teaching people who click on phishing email messages about such attacks. The APWG landing page,² in use since September 2008, is available in several languages. As of April 2010, it has been displayed in place of 1,285 phishing pages and viewed almost 200,000 times.¹⁷ While measuring the effect of the landing page is difficult, it is a step in the right direction, offering multiple ways of protecting people worldwide.

Better interfaces. The second major strategy for protecting people is to provide better interfaces. The following paragraphs cover innovations in warnings, support for properly identifying Web sites, and authentication.

A general problem with security warnings is that users often close them the instant they appear, a perfectly rational behavior, as many warnings are so obtuse people don't understand what the problem is or what they should do. Other warnings annoyingly

interrupt what people are trying to accomplish. Warning notifications can also be too subtle, with people not even seeing them.

A “passive indicator” warns of potential dangers without interrupting the user's task. In contrast, “active indicators” force users to notice the warnings by interrupting them. Studies by Wu et al.³⁶ and Egelman et al.¹⁰ found passive warnings ineffective in protecting people from phishing scams, as they are easily missed.

Egelman et al. also examined the effectiveness of active anti-phishing warnings in Firefox and Internet Explorer 7; see Figure 2 for an example of Firefox's active warning. Using simulated phishing attacks, they found no participants fell for phishing attacks when seeing Firefox's warning, but, surprisingly, half of the participants using IE did. Egelman et al. analyzed the results using a model from the warning sciences describing how people see, understand, believe, and act on warnings in the regular physical world. Using this framework, they found that most people simply did not “see” the warning in IE, since it looked like a standard “page not found” warning. Several participants also did not believe the warn-




Figure 2. The active warnings used by Mozilla Firefox when blocking phishing pages are more effective than passive warnings.




Figure 3. Extended Validation certificates in Microsoft's Internet Explorer and Mozilla Firefox.

ing, thinking Microsoft would not put them at risk, and went on to give sensitive personal information. In response to this work, Microsoft redesigned its anti-phishing warnings in IE8.

Techniques have been developed to help people identify the sites they are on. However, it is unclear how much they help in practice; for example, Extended Validation (EV) certificates are a special kind of certificate with specific guidelines for verifying that an organization purchasing the certificate is legitimate. When a site with an EV certificate is loaded, the browser's URL bar changes automatically to show the site's brand name (see Figure 3). However, a 2007 study by Jackson et al.¹⁸ found EV certificates ineffective in protecting people from phishing attacks.

Using SiteKey, a technique used by many financial organizations worldwide, users first select a secret picture. When logging in, they see if the picture is displayed to verify they are on the right site. However, Schecter et al.²⁹ found SiteKey suffers from the same problem as passive indicators, in that the absence (or presence) of an indicator is easily missed or even rationalized away by the typical user.

An alternative to indicators is to improve the way users sign into sites. Two-factor authentication (2FA) strength-

ens authentication by requiring two separate ways to prove one's identity. One of the most common forms is key fobs with a periodically changing number synchronized with a remote server. Users login by using both their password and this number. While 2FA increases the cost of conducting phishing attacks, phishers have also developed workarounds, including switching to real-time man-in-the-middle attacks using malware (such as the Zeus Trojan horse).

Train users. Training is the third way of protecting people from phishing scams. Though an essential aspect of computer security, it is also arguably the least-popular approach, given the inherent challenges of motivating people to be secure, as well as the fact that training does not guarantee complete protection (though neither does any other solution today).

Many Web sites offer advice on how to identify phishing sites. A 2010 study by Kumaraguru et al.²⁴ found this kind of information useful in helping people identify fake Web sites but only if they actually read the material. In an earlier study, Kumaraguru et al.²³ found that simply emailing anti-phishing material was ineffective, because people were habituated to receiving such warnings and thought they already knew how to protect themselves.

Two lines of research have sought to address these problems: The first is micro games designed to teach people about phish. Micro games are a popular format for games played for short periods of time. Sheng et al.³² developed a micro game for computer security called Anti-Phishing Phil (see Figure 4) that teaches about browser address bars, domain names, and phishing pages, then tests users on what they learned. Phil incorporates many ideas from learning science, a body of empirical research that seeks to understand the best methods for learning and retaining knowledge. An example principle is "conceptual-procedural," holding that high-level concepts should be interleaved with concrete procedures on how to achieve given tasks. An evaluation of Phil with more than 4,500 people demonstrated it improved novices' ability to identify phish by 61% while also dramatically lowering false positives.

The second approach is "embedded training," teaching people in the specific context of use in which they would normally be attacked. Embedded training is in contrast to other forms of security training that might take place in a classroom and give people few opportunities to test what they've learned. Kumaraguru et al.²⁴ developed an embedded training system called PhishGuru that sends simulated phishing email messages to people. If they fall for one, they see an intervention that teaches them about phishing and how to protect themselves. In a study with more than 500 participants, Kumaraguru et al. found this approach led to a 45% reduction in falling for phish even a month after being trained, helping lead to creation of the APWG landing page² described earlier.

Conclusion

This article has emphasized criminals' tenacity and creativity, a trend that will only continue. We will also likely see increased spear-phishing and whaling attacks, as phishers look for vulnerable targets with valuable information.

Blurring traditional security perimeters, phishing also causes new problems for organizations. Even their lawyers and accountants can be attacked to surreptitiously gain access to




Figure 4. Anti-Phishing Phil is a micro game that teaches people how to identify phishing scams.

documents. Facebook and other social media provide more contextual details that can be used for spear-phishing attacks. Individual employees falling for phish in one context can cause headaches for their organizations over reused passwords. Finally, instant messaging, VOIP, SMS, and other relatively new ways of communicating offer criminals new vector options for delivering attacks.

On the positive side, law enforcement, industry, and academic researchers are getting better organized in terms of reporting phishing attacks, sharing information, analyzing data to identify trends, and focusing resources. More organizations are devoted to combating online fraud, including APWG, the Internet Crime Complaint Center (IC3) (<http://www.ic3.gov/>), and National Cyber-Forensics and Training Alliance (NCFTA) (<http://www.ncfta.net/>). There are also initiatives dedicated to educating people about phishing scams, including StaySafeOnline.com. Law enforcement has stepped up efforts in gathering evidence and cooperating with international partners in shutting phishing sites and phishing gangs. Legislators in the U.S. have also passed laws to explicitly spell out what phishing is and the related penalties, including California's Anti-Phishing Act of 2005,²¹ though such laws face many of the same challenges as anti-spam laws in terms of attackers being outside a particular jurisdiction, the sheer number of attacks, and limited resources available to law enforcement.

Phishing will continue to be an arms race. Since any communication medium can be used for phishing, it is also a problem that can never be solved. The best we can hope for is to blunt the worst aspects of phishing and continue to work on better ways to prevent, detect, and respond to this new form of a very old crime.

Acknowledgments

Special thanks to the Supporting Trust Decisions group (<http://cups.cs.cmu.edu/trust/>) at Carnegie Mellon University and Wombat Security Technologies (<http://www.wombatsecurity.com/>) for their contributions and comments. 

References

1. Abu-Nimeh, S., Nappa, D., Wang, X., and Nair, S. A comparison of machine learning techniques for phishing detection. In *Proceedings of The Anti-Phishing Working Group's Second Annual eCrime Researchers Summit* (Pittsburgh, PA, Oct. 4–5, 2007), 60–69.
2. Anti-Phishing Working Group. APWG & Carnegie Mellon University's phishing education landing page; <http://education.apwg.org/r/en/>
3. Anti-Phishing Working Group. *Phishing Activity Trends Report: Third Quarter Report*, Jan. 2010; http://apwg.org/reports/apwg_report_Q3_2009.pdf
4. Arthur, C. Facebook hit by phishing attack. *The Guardian* (Apr. 30, 2009); <http://www.guardian.co.uk/technology/2009/apr/30/facebook-phishing-scam>
5. Blizzard Entertainment. *Battle.net Authenticator FAQ*; http://us.blizzard.com/support/article.xml?locale=en_US&articleId=24660
6. Cavalli, E. World of Warcraft phishing attempts on the rise. *Wired* (Apr. 29, 2009); <http://www.wired.com/gamelife/2009/04/world-of-warcraft-phishing-attempts-on-the-rise/>
7. Cova, M., Kruegel, C., and Vigna, G. There is no free phish: An analysis of 'free' and live phishing kits. In *Proceedings of the Second USENIX Workshop on Offensive Technologies* (San Jose, CA, July 28, 2008). Usenix; <http://portal.acm.org/citation.cfm?id=1496706>
8. Dhamija, R., Tygar, J.D., and Hearst, M.A. Why phishing works. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Quebec, Apr. 24–27). ACM Press, New York, 2006, 581–590; <http://portal.acm.org/citation.cfm?id=1124861>
9. Downs, J.S., Holbrook, M.B., and Cranor, L.F. Decision strategies and susceptibility to phishing. In *Proceedings of the SOUPS Symposium on Usable Privacy and Security* (Pittsburgh, July 12–14). ACM Press, New York, 2006.
10. Egelman, S., Cranor, L.F., and Hong, J.I. You've been warned: An empirical study of the effectiveness of Web browser phishing warnings. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Florence, Italy, Apr. 5–10). ACM Press, New York, 2008, 1065–1074.
11. Fette, I., Sadeh, N., and Tomasic, A. Learning to detect phishing emails. In *Proceedings of the 16th International World Wide Web Conference* (Banff, Canada, May 8–12, 2007), 649–656.
12. Garera, S., Provos, N., Chew, M., and Rubin, A.D. A framework for detection and measurement of phishing attacks. In *Proceedings of the WORM Workshop on Rapid Malcode* (Alexandria, VA, Nov. 2). ACM Press, New York, 2007; <http://portal.acm.org/citation.cfm?id=1314391>
13. Görling, S. An overview of the Sender Policy Framework as an anti-phishing mechanism. *Internet Research* 17, 2 (2007), 169–179.
14. Herley, C. and Florencio, D. A Profitless endeavor: Phishing as a tragedy of the commons. In *Proceedings of the New Security Paradigms Workshop* (Lake Tahoe, CA, Sept. 22–25, 2008).
15. Herley, C. and Florencio, D. Nobody sells gold for the price of silver: Dishonesty, uncertainty, and the underground economy. In *Proceedings of Workshop on the Economics of Information Security* (London, June 24–25, 2009).
16. Hong, J. Why have there been so many security breaches recently? Blog@CACM (Apr. 27, 2011); <http://cacm.acm.org/blogs/blog-cacm/107800-why-have-there-been-so-many-security-breaches-recently/fulltext>
17. Hong, J.I. Statistical analysis of phished email users intercepted by the APWG/CMU phishing education landing page. In *Proceedings of the Anti-Phishing Working Group Counter eCrime Operations Summit IV* (Sao Paulo, Brazil, May 11–13, 2010); http://www.antiphishing.org/events/2010_opSummit.html
18. Jackson, C., Simon, D.R., Tan, D.S., and Barth, A. An evaluation of extended validation and picture-in-picture phishing attacks. In *Proceedings of the 11th International Conference on Financial Cryptography* (Trinidad/Tobago, Feb. 12–15, 2007), 281–293.
19. Jagatic, T.N., Johnson, N.A., Jakobsson, M., and Menczer, F. Social phishing. *Commun. ACM* 50, 10 (Oct. 2007), 94–100.
20. Jakobsson, M. and Myers, S. *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft*. Wiley-Interscience, 2006.
21. Keizer, G. California enacts tough anti-phishing law. *InformationWeek* (Oct. 3, 2005); <http://informationweek.com/news/171202672>
22. Krastev, N. U.S. indicts dozens from Eastern Europe in Internet theft scheme. *Radio Free Europe* (Oct. 1, 2010); http://www.rferl.org/content/US_Indicts_Dozens_From_Eastern_Europe_In_Internet_Theft_Scheme/2173545.html
23. Kumaraguru, P., Rhee, Y., Sheng, S. et al. Getting users to pay attention to anti-phishing education: Evaluation of retention and transfer. In *Proceedings of the Anti-Phishing Working Group's Second Annual eCrime Researchers Summit* (Pittsburgh, Oct. 3–5, 2007), 70–81.
24. Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L.F., and Hong, J.I. Teaching Johnny not to fall for phish. *ACM Transactions on Internet Technology* 10, 2 (2010), 1–31.
25. Litan, A. Phishing attack victims likely targets for identity theft. Gartner Group, May 2004.
26. Markoff, J. Larger prey are targets of phishing. *New York Times* (Apr. 16, 2008); <http://www.nytimes.com/2008/04/16/technology/16whale.html>
27. Moore, T. and Clayton, R. Examining the impact of Website take-down on phishing. In *Proceedings of the Anti-Phishing Working Group's Second Annual eCrime Researchers Summit* (Pittsburgh, Oct. 3–5, 2007), 1–13.
28. PhishTank. *PhishTank Stats, 2011*; <http://www.phishtank.com/stats.php>
29. Schechter, S.E., Dhamija, R., Ozment, A., and Fischer, I. The emperor's new security indicators: An evaluation of Website authentication and the effect of role playing on usability studies. In *Proceedings of the IEEE Symposium on Security and Privacy* (Washington, D.C., 2007), 51–65.
30. Sheng, S., Holbrook, M.B., Kumaraguru, P., Cranor, L.F., and Downs, J.S. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Atlanta, Apr. 10–15). ACM Press, New York, 2010, 373–382.
31. Sheng, S., Kumaraguru, P., Acquisti, A., Cranor, L.F., and Hong, J.I. Improving phishing countermeasures: An analysis of expert interviews. In *Proceedings of the Fourth Anti-Phishing Working Group eCrime Researchers Summit* (Tacoma, WA, Oct. 20–21, 2009).
32. Sheng, S., Magnien, B., Kumaraguru, P. et al. Anti-Phishing Phil: The design and evaluation of a game that teaches people not to fall for phish. In *Proceedings of the Third Symposium on Usable Privacy and Security* (Pittsburgh, July 18–20). ACM Press, New York, 2007, 88–99.
33. Sheng, S., Wardman, B., Warner, G., Cranor, L.F., Hong, J.I., and Zhang, C. An empirical analysis of phishing blacklists. In *Proceedings of the Sixth Conference on Email and Anti-Spam* (Mountain View, CA, July 16–17, 2009).
34. Stajano, F. and Wilson, P. Understanding scam victims: Seven principles for systems security. *Commun. ACM* 54, 3 (Mar. 2011), 70–75.
35. Verisign. *Fraud Alert: Phishing: The Latest Tactics and Potential Business Impact*. White Paper, 2009; <http://www.verisign.com/static/phishing-tactics.pdf>
36. Wu, M., Miller, R.C., and Garfinkel, S. Do security toolbars actually prevent phishing attacks? In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Montréal, Apr. 24–27). ACM Press, New York, 2006, 601–610.
37. Xiang, G. and Hong, J.I. A hybrid phish detection approach by identity discovery and keywords retrieval. In *Proceedings of the International World Wide Web Conference* (Madrid, Apr. 20–24, 2009), 571–580.
38. Xiang, G., Rose, C., Hong, J.I., and Pendleton, B. A hierarchical adaptive probabilistic approach for zero-hour phish detection. In *Proceedings of the ESORICS 15th European Symposium on Research in Computer Security* (Athens, 2010), 571–589.
39. Zhang, Y., Hong, J.I., and Cranor, L.F. Cantina: A content-based approach to detecting phishing Web sites. In *Proceedings of the 16th International World Wide Web Conference* (Banff, Canada, May 8–12, 2007), 639–648.

Jason I. Hong (jasonh@cs.cmu.edu) is an associate professor in the School of Computer Science and Human Computer Interaction Institute at Carnegie Mellon University, Pittsburgh, PA.

© 2012 ACM 0001-0782/12/01 \$10.00

DOI:10.1145/2063176.2063198

Virtual testbeds model them by seamlessly integrating physical, simulated, and emulated sensor nodes and radios in real time.

BY GEOFF COULSON, BARRY PORTER, IOANNIS CHATZIGIANNAKIS, CHRISTOS KONINIS, STEFAN FISCHER, DENNIS PFISTERER, DANIEL BIMSCHAS, TORSTEN BRAUN, PHILIPP HURNI, MARKUS ANWANDER, GERALD WAGENKNECHT, SÁNDOR P. FEKETE, ALEXANDER KRÖLLER, AND TOBIAS BAUMGARTNER

Flexible Experimentation in Wireless Sensor Networks

WIRELESS SENSOR NETWORKS (WSNs) play a key role in the emerging “real-world Internet,” with several large-scale WSNs being deployed; see, for example, Bernat² and Dudek et al.⁹ However, WSN development is inherently complex, involving hardware design, embedded and distributed programming, heterogeneity, scale, and unpredictable environmental changes. Addressing this complexity, testbed-based experimentation (recommended by Weiser³²) is increasingly the norm for developing and optimizing WSN systems in a controllable environment prior to deployment.

The WSN research community has historically relied on three main approaches to testbed-based experimentation: physical, simulation, and emulation. However, researchers appreciate that each involves

significant drawbacks when used in isolation (see the sidebar “Physical Testbeds vs. Simulation vs. Emulation”). Therefore, they seek to combine all three to enable a more complete evaluation of the system being developed. Unfortunately, each approach requires different coding styles and tools, forcing researchers to expend significant effort reimplementing their systems for different tools/platforms/approaches. As a remedy, techniques have been developed to reduce the transitioning effort among the three approaches, but further work is needed to address the emerging requirement for more flexible experimental facilities.

Our work abstracts the concept of testbeds to yield virtual testbeds (VTBs) programmed similarly regardless of whether their underlying realization is physical, simulated, or emulated. VTBs are private, custom-designed, per-experiment, virtualized testbed instances that enable developers to seamlessly combine and/or interchange physical elements, including sensor nodes and radios, with simulations and emulations of these elements.

We are developing a reference implementation of the VTB abstraction on top of a large-scale federated physical testbed infrastructure (see Figure 1), augmenting the inherent flexibility of the VTB abstraction in terms of scal-

» key insights

- **Physical, emulated, and simulated elements of wireless sensor networks can be seamlessly mixed to gain massive-scale “virtual testbeds” with desired trade-offs involving fidelity, repeatability, and network size.**
- **The relative speed difference between Internet links and low-power wireless radios allows sensor nodes at physically distant testbed sites to interconnect (virtually) over the Internet.**
- **The WISEBED approach of recursively and hierarchically applying a common Web Services API to physical, simulated, and federated or hybrid testbeds facilitates transparent composition of virtual testbeds that can be accessed and controlled by common tools.**



ability (VTBs combine physical nodes from different federated sites) and heterogeneity (VTBs incorporate a variety of equipment types).

Research Challenges

In this article, rather than survey the many purely physical WSN testbeds (such as Chun et al.⁴), we focus on work that seeks to “virtualize” WSN testbeds to some extent or federate physical testbeds; as far as we know, no previous work has sought to do both. Note that the term “virtualization” is commonly used in a variety of contexts, as in “slicing” in systems like PlanetLab and the simultaneous running of multiple operating systems on a single processor, as in VMware. We reserve the term exclusively

for systems that mix physical WSN testbeds with elements of simulation and emulation.

In terms of virtualization, several research efforts have sought to bridge the gap between physical execution and simulation. For example, Li et al.²³ and TOSSIM²² used TinyOS component graphs to generate discrete-event simulations that can be augmented with a small number of attached physical nodes (such as three reported in Li et al.²³). Node software in Girod et al.¹⁵ and Park et al.²⁷ was executed inside a simulator, but the simulated nodes communicated through physical radios. And in Österlind et al.²⁶ and Wen et al.³³ node software was executed iteratively on physical devices and in simulation, with vari-

ous arbitration and timing schedules applied. All these projects involved a useful mix of physical and simulated elements, but that mix was fixed; for example, TOSSIM cannot support simulated nodes with real radios, and Wen’s cannot mix physical nodes with simulated nodes. In contrast, we strive for generality, so physical, simulated, and emulated elements can all be combined as desired.


Another feature of these efforts is that programmers must significantly adapt their code to work with the given mix of physical and simulated elements. In contrast, we strive for transparency so the VTB looks as much as possible like a purely physical testbed, with the goal being to minimize the overhead of transferring an experimen-

Physical Testbeds vs. Simulation vs. Emulation

Here we outline the three main approaches to providing WSN testbeds. Our notion of VTBs allows users to arbitrarily combine elements of all three (see the figure here).

Physical testbeds (such as those described in Chun et al.,⁴ Dutta et al.,¹¹ Ertin et al.,¹³ Handziski et al.,¹⁸ Tutornet,¹² and Werner-Allen et al.³⁴) excel at high-fidelity evaluation of mature WSN designs, as well as detailed planning for real-world deployments. However, physical testbeds for WSN systems tend to be small in scale, expensive to maintain, and time-consuming to set up. They also typically lack flexibility, often offering only a single, fixed, connectivity topology and limited heterogeneity (such as only a single type of sensor node, radio, operating system, or programming language). They also tend to be limited in their programmability at lower levels of the system; for example, many use fixed operating systems and networking stacks. They are also often unsuited to experimentation scenarios requiring repeatability of experiments, as many relevant operating parameters are beyond user control (such as local radio interference due to infrastructure and other experiments).

Three-cornered testbed design space for WSN experimentation.



Simulation (such as described in Fekete et al.¹⁴ and Levis et al.²²) is useful for quickly trying out new ideas and for investigating the behavior of new protocols and mechanisms in varied topologies at large scale and in a repeatable manner. The most notable drawback is a lack of fidelity, often making it unrealistic to simulate fully at the instruction-execution level and with high-fidelity radio or power-consumption characteristics. While such limitations are not necessarily problematic in traditional network environments, where simulators (such as NS-2²¹) are prominent, they represent significant drawbacks in WSN environments where resource scarcity and incidental physical characteristics are of the essence. Simulation alone is therefore of limited use in planning for real-world WSN systems and deployments.

Emulation (such as described in Girod et al.¹⁶ and Wu et al.³⁶) is situated between physical reality and simulation. Whereas simulation abstractly models target systems, emulation duplicates the functionality of one system in terms of another system and is therefore capable of much greater fidelity than simulation while potentially offering greater flexibility than a purely physical testbed. Emulation is a much less exploited approach in the WSN testbed context despite much potential; for example, emulation in the form of network overlay technology could be used to support different inter-node connectivity patterns in a physical testbed. Alternatively, a battery-based power supply on a physical node could be emulated by interposing a mains electricity-powered hardware module degrading power over time.

tal system to real-world deployment. Finally, all these projects were small in scale and therefore of limited application in the development of large-scale WSN deployments.

Turning now to work that applies the federation concept to physical WSN testbeds,^a Ohio State University¹⁷ combined the infrastructure and software of the Kansei testbed with the GENI facility to provide a unified solution; the Senslab project³⁰ aimed to unify four discrete heterogeneous testbeds into a single testbed of 1,000 nodes; and Cooperating Objects Network Of

Excellence (CONET)⁶ employed a REST API to provide uniform Web-based access to different testbeds. These efforts highlight the promise of testbed federation in pursuit of scalability, but none support integration of simulated and emulated elements or deliver the transparency of federation offered by our VTB abstraction.

VTB Abstraction

VTBs provide the abstraction of a user-designed private WSN testbed in which some testbed elements are physical, some are simulated, and others are emulated. Users design their VTB in such a way that its mix of physicality, simulation, and emulation is appropriate to their goals. They then instantiate their VTB, deploy their

software onto it, and observe the outputs and behavior of their experimental systems as if they were running on a dedicated physical testbed.

VTB elements can include:

Sensor modality. Examples are temperature sensors and pollutant sensors;

Sensor data. What the sensors observe (such as the current temperature and pollutant levels);

Nodes. CPU+memory+radio devices to which sensor modalities are attached;

Node power. Power-supply characteristics (such as battery or solar) and remaining energy of nodes;

Node connectivity. Nodes in broadcast range of each node and the volatility of their relationships; and

Node mobility. The movement patterns nodes are subjected to if, for example, they are mounted on public-transport vehicles.

Users augment VTB physical nodes by attaching emulated sensor modalities not supported by hardware. Alternatively, they can foster experiment repeatability by feeding emulated (scripted) sensor data to their nodes (physical, simulated, or emulated). They can also increase the scale of their VTB by selecting physical nodes from different sites in the underlying federation, combining simulated nodes with a core set of physical nodes, and even multiplexing multiple emulated nodes onto a single physical node. They can also employ emulated node power to facilitate repeatable battery-life investigations.

The emulation of node connectivity is the biggest single source of VTB flexibility, realized through virtual links offering emulated real-time connectivity between pairs of nodes according to the characteristics of different types of radio hardware. Using virtual links, testbed users are able to explore different connectivity patterns atop a set of physical nodes with fixed physical connectivity. For example, they might add emulated connectivity between pairs of nodes with no physical connectivity, including across sites and between any combination of physical, emulated, and simulated nodes. They can also remove unwanted physical connectivity to, say, build an emulated ring topology on top of an underlying physical mesh. Taking this to extremes, they can even

^a The concept of a federated testbed is not new; projects (such as PlanetLab and Panlab) have long applied similar ideas, though not in a WSN context.




Figure 1. WISEBED physical testbed environment. In this federation of physical WSN testbed sites (nine today), each differs in its choice of hardware, software, and physical layout.

emulate node mobility (of physical, simulated, and emulated nodes) by dynamically changing the emulated connectivity between nodes according to a time-based script.³

From the user’s point of view, a VTB would appear as an instance of a Web services interface called iWSN supporting a comprehensive set of testbed-related operations, including loading experimental code and collecting results. Users typically employ a GUI-based front-end to mediate access to their VTB rather than interact with their iWSN instance directly.

The user’s view of the entire VTB-based WSN experimentation process is captured in a series of (potentially iterated) steps (see Figure 2):

Experimental software development. Users employ our software development kit (SDK, <http://www.wisebed.eu>) to develop and synthesize software that can be deployed in a VTB;

VTB specification. They specify a custom VTB that meets the requirements of their experiment in terms of its physical, simulated, and emulated elements;

VTB reservation. They contact our reservation system to request their custom-specified VTB be instantiated at a particular time for a particular duration. The reservation system re-

serves a set of underlying resources able to support the VTB and returns a “reservation key” that uniquely identifies the reservation and serves as a promise that a VTB (as specified) will be made available at the requested time;


VTB instantiation. When the pre-specified reservation time arrives, users ask the instantiation system to redeem their reservation keys. In response, the instantiation system builds the user-specified VTB using the previously reserved underlying resources, loads the user’s software onto

the nodes of the VTB (along with any SDK runtime services), and returns a dedicated iWSN instance; and

VTB operation. Users run their experimental software on their newly instantiated VTB under the control of a GUI that mediates access to the VTB in real time.

To show the applicability of this approach, consider an experimental scenario in which experimenters want to select a broadcast protocol for use in a large-scale WSN deployment.² This experimentation requires a testbed with three properties: scale and topology

Figure 2. User view of VTB-based WSN system development.



close to that of the target deployment environment; significant use of physical nodes and radios (important, as the forwarding decisions made by broadcast protocols typically rely on physical features like RSSI thresholds difficult to simulate with high confidence); and repeatability. This combination would be extremely challenging for most existing testbed environments.

Using it, experimenters might begin by designing a simple VTB underpinned by simulated nodes and connectivity. Such a VTB could be executed on a desktop PC backed by a suitable simulator engine (such as Shawn¹⁴). Running their experimental code would give an initial feel for the general behavior of the candidate protocols and likely critical areas. They might then create a second, more sophisticated VTB underpinned by federated physical testbeds augmented with virtual links to yield a configuration with scale and topology close to that of the target deployment environment. The same experimental code can then be executed on this higher-fidelity VTB, using emulated sensor input to drive the experiments in a repeatable manner. It could be argued that using virtual links here might compromise fidel-

ity somewhat, but the flexibility of the approach makes running a range of what-if experiments straightforward with different virtual-link configurations to build confidence in the stability and fidelity of the results.

Software Infrastructure

Here, we expand on key aspects of the underlying functionality required to support the VTB abstraction in our federated environment. We omit consideration of VTB reservation (less central to the structure and performance of VTBs and implemented like a number of existing purely physical testbeds), as well as discussion of our security provision (to ensure users are properly authenticated and authorized to use specific resources at specific times); for more, see the WISEBED Project.³⁵

The SDK includes tools for design, implementation, synthesis, and deployment of the experimental software that will run atop a user's VTB. It is based on a lightweight software component model⁷ offering a simple-to-use modular development process. The component model is runtime reconfigurable, so different components can be deployed as the experiment proceeds, facilitating exploration of what-

if scenarios.

The SDK is also the primary means by which we address the key research challenge of transparency, whereby users are protected from having to re-implement their code when moving from physical to simulated to emulated and mixed testbed environments. This transparency is achieved through an abstraction layer that selectively exposes low-level APIs (such as drivers) so both application-level software and systems-oriented software can run unchanged atop physical, simulated, or emulated nodes; our Lorien operating system²⁸ is a good example of systems-oriented software that runs over our low-level APIs.

Besides its abstraction layer, the SDK provides a set of runtime services selectively configured into the user's software build by the instantiation system, comprising these components:

Generic node management. Supports implementation of the generic node-management operations supported by the iWSN interface, including, for example, functions that support the pinging of the node or respond to requests for current battery status;

Sensor emulator. "Pretends" to be sensors attached to the node and supply a stream of emulated sensor data (generated internally or driven by an external source); and

Radio stacking framework. Inserts pseudo-network-device-driver components in front of the bottom-level radio device driver, a key element of our "virtual links" implementation.

The SDK currently operates across a range of platforms, including Conti-ki,¹⁰ iSense,⁵ Lorien,²⁸ TinyOS,¹⁹ and the Shawn WSN simulator.

VTBs are specified through an XML schema called WiseML,³⁵ a multi-purpose format also used to encapsulate experimental output data. For VTB specification, WiseML supports specification of the following system elements:

Node-related information. Included is information on node type, whether the node should be physical, simulated, or emulated, the coordinates of the node in a global 3D space, and the sensor modalities supported; and

Connectivity information. Helps define potential connectivity between pairs of nodes in terms of virtual links; specifying a virtual link between two




Figure 3. The iWSN interface, a Web services interface providing uniform management access to all testbed types, whether VTBs, the underlying physical testbeds of the federated physical environment, or purely simulation-based testbed instances (simulator engines).

nodes in the VTB enables one-hop uni-directional communication between these nodes.^b

Users typically generate WiseML-based VTB specifications through a GUI-based tool. Such tools are independent of the rest of the software infrastructure, and it is possible to imagine a range of them with varying degrees of sophistication. A typical tool might display graphical representations of the available physical testbeds laid out in space, with each testbed represented as a graph and one-hop reachability by directed edges. With these graphs, users could select a subset of the available physical nodes to be allocated to their VTB. They could also add simulated nodes from a drop-down menu and selectively add and delete edges between arbitrary nodes to modify the topology of their VTB.

The VTB instantiation system gets involved when users redeem a reservation key issued by the reservation system. The system gathers the physical resources derived from the reservation stage, configuring/connecting them according to their original VTB specification. The instantiation system (and operation step) relies on the iWSN interface providing uniform access to all testbed realizations and supporting the federation and hierarchical composition of testbeds (see Figure 3).

The instantiation process involves creating an empty VTB, then using its iWSN interface to populate it with the required nodes, connectivity, and other parameters. For example, where a user wants a VTB to include physical nodes from multiple sites, the instantiation system creates an overarching VTB into which it inserts selected nodes from the various sites; it then sets up virtual links between nodes to create the required federated configuration.

A similar approach to federating underlying testbeds is applied in experiments involving simulated nodes whereby the instantiation system ex-

ecutes one or more simulator engines. These engines are simulator processes instantiated on a suitably located and resourced server machine; we primarily use the Shawn WSN simulator mentioned earlier. As in physical sites, the simulator engine is abstracted as a testbed in its own right, exporting an iWSN instance that can be federated and “virtually linked” with testbed elements in other VTB instances to build up the user’s required configuration.

The software behind the iWSN interface’s virtual link operations uses the SDK’s radio-stacking framework discussed earlier to deploy “pseudo” radio drivers that appear to software on the node as “real” radio drivers. To disable physical connectivity between two nodes, the pseudo-driver on one node simply drops packets originating from the other. To establish a virtual link where there is no physical connectivity, the pseudo-driver on the sending node transparently diverts (selected) outgoing packets to virtual-link software running on a server, and the pseudo-driver on the receiving node inserts incoming packets arriving from the virtual-link software.

Note this virtual link creation delivers only the basic “plumbing.” On top of it, the instantiation system inserts, where required, an additional server process that models a specific radio channel. For this purpose, we are experimenting with the OMNeT++ simulator, finding it fast enough to handle virtual link packets in real time.³¹ For a more detailed discussion of the implementation of virtual links see Baumgartner et al.¹

On completion of the instantiation step, users are given a dedicated iWSN instance through which they interact with their newly instantiated VTB as if it were a private physical testbed; for example, they can then call iWSN services to deploy, monitor execution, and dynamically reconfigure their code. These operations may be carried out programmatically through the iWSN interface directly via Web services calls but are more commonly carried out through a program interacting with the iWSN interface on behalf of the user. We refer to such programs as controllers; like VTB-specification tools, they are distinct architectural elements decoupled from the rest of

^b Two virtual links are needed to represent mutual connectivity between two nodes, allowing users to model situations in which a node A can send to a node B, but B cannot send to A. The WiseML specification of a virtual link includes the unique identifiers of the nodes that will participate in the link, as well as the link’s characteristics in terms of its “link quality indicators,” packet error rate, and other parameters.




Figure 4. TARWIS controller graphical user interface.

the system and can evolve or be replaced independently.

We have implemented two: The first is a scriptable command-line controller that enables batch-mode control of experiments and is particularly useful in experiments that must be executed repeatedly with varied parameters. The second is an interactive GUI-based system called Testbed

Table 2. Time for receiving packet bursts sent from physical to simulated nodes.

Number of pairs	Rx time (seconds)
1	0.109
5	0.359
10	0.452
15	0.710
20	0.881
30	1.001

Management Architecture for Wireless Sensor (TARWIS) networks²⁰ that offers a Web-based user interface; it displays the attached VTB and allows users to interactively issue iWSN commands to individual nodes (such as to reset, flash, reprogram, and reboot nodes and to send commands) and store their output. Individual sensor-node output can be displayed in individual windows, with connectivity/topology information displayed graphically (see Figure 4). Experiment results are stored in a WiseML file for post-experiment archiving and analysis. Note, due to the common use of the iWSN interface, TARWIS, or indeed any controller, can also be used to interact in exactly the same way with physical sites and simulator engines, as well as with VTBs.

The implementation just outlined represents a substantial proof-of-concept demonstration of the viability of

the VTB abstraction, though there is room for further development, especially in terms of VTB-specification tools; other interesting areas (such as node mobility and emulated sensors) have also yet to be explored in detail. More generally, the implementation highlights and validates the modular and pluggable nature of the software infrastructure; alternative implementations of many of the architectural elements (such as specification, reservation, instantiation, and operating systems) can be provided independently. We see this modular design as crucial in testbed environments that aspire to grow and co-evolve with their user base.

Experimental Validation

The central issue for the viability of the VTB concept is the extent to which physical, simulated, and emulated nodes are able to work and communicate seamlessly with one another in real time. Seamless operation implies an experimental validation of VTBs (see Figure 5) should focus on three areas:

Real-time performance of virtual links between physical nodes. Is the connectivity offered by virtual links a suitable basis for emulating the characteristics of real physical radio-based connectivity? Is it possible to support a virtual link capable of being indistinguishable from a real link in terms of, say, message-transmission rates and latencies?

Real-time performance of virtual links between physical and simulated nodes. Can nodes running in a simulator server “keep pace” with physical and emulated nodes elsewhere in the VTB in terms of sourcing and sinking messages at a rate and latency, whereby the system as a whole executes seamlessly?; and

Per-node resource overhead of the SDK’s runtime software. The VTB abstraction imposes this overhead on sensor nodes as an experiment runs on the VTB. The overhead is mainly subsumed within the SDK’s abstraction layer and runtime services, manifesting mainly in terms of memory occupancy.

Performance of virtual links between physical nodes. The evaluation strategy here is to establish that virtual links perform at least as well as real radio links, giving users the potential to

Table 1. Message latencies in virtual links under various scenarios; results are averaged over more than 1,000 messages in each case.

Hardware/OS	Physical Radio	Intra-site	Inter-site
ScatterWeb/TinyOS	75ms	5ms	53ms
TelosB/Contiki	40.2ms	5ms	53ms
iSense/iSense	7ms	5ms	53ms




Figure 5. Overall evaluation strategy.




Figure 6. Node software configuration used in an SDK evaluation.

use any spare capacity to model real radio characteristics (such as using a network simulator like OMNeT++ mentioned earlier).

Specifically, we measure the time required to transmit a message between two physical nodes in three cases:

- ▶ Baseline, via one-hop physical radio link;
- ▶ Intra-site, via a virtual link that traverses from one node to a PC via direct universal asynchronous receiver/transmitter (UART) connection, then to a second PC via Gigabit Ethernet, and finally to the second node via UART again; and
- ▶ Inter-site, via a virtual link, as with intra-site, except it uses the Internet instead of Gigabit Ethernet, interconnecting a site in Lübeck, Germany, with another site in Lancaster, U.K.

We examine these three cases on three different hardware platforms: ScatterWeb, TelosB, and iSense (see Table 1). In both the intra-site and inter-site cases the back-end message transport systems represent the dominant cost, so the results across all platforms are similar. The main observation is that intra-site is significantly faster than using physical hardware radios for all three platforms considered, demonstrating the viability of emulating sensor-node communication in VTBs. Due to the raw speed advantage a generous amount of spare capacity is available for the modeling-link characteristics and radio-contention scenarios prior to packet delivery. Inter-site offers less spare capacity; the relatively slow times are attributable to our current VTB implementation's use of Simple Object Access Protocol (SOAP) encapsulation for messages sent across the Internet. Even so, at least in the case of ScatterWeb/TinyOS, the VTB provides significant spare capacity in which to model radio characteristics. In the future, we expect more efficient transport protocols to provide considerably larger spare capacity; raw ping times between the two sites in question suggest approximately 23ms of one-way latency on average.

Virtual links between physical and simulated nodes. In this experiment we use a physical testbed of 30 iSense nodes connected via Ethernet to an Intel dual-core 2.5Ghz PC with 3GB of

RAM running the Shawn WSN simulator engine. We execute scenarios in which varying numbers of physical nodes connect one-to-one to “partnered” simulated nodes in the Shawn environment. We then measure how well the simulated nodes “keep up” with the physical nodes sending packets at realistic rates. To model a realistically demanding scenario, we employ packet bursts of 10B x 20B packets with 5ms inter-packet gaps. These bursts represent a far higher data rate than would be encountered in a typical WSN-based experimental scenario.

Given this setup, we obtained the results in Table 2 for different numbers (one to 30) of active sender-receiver pairs. The “Rx time” figure, or time to receive all packets, for the single-pair case shows a 10-packet burst is received in 0.109 seconds. Given that most realistic low-power WSN application scenarios employ sample periods on the order of seconds, this figure is perfectly acceptable. Furthermore, as the number of pairs increases, the overhead scales very well, so despite the inherent serialization imposed by the simulator (any event-based simulator would be the same), the receiving rate remains well within the operating range of typical experimental scenarios.

Per-node memory overhead. The per-node memory overhead of the SDK's runtime software is incurred primarily by its common abstraction layer. To evaluate memory overheads, we instrument a typical sensor network application in which nodes sample data every five seconds, sending it to a parent node and also forwarding data received from child nodes. We implement the application twice: once using only native OS facilities, representing the baseline, and once using the SDK along with the chosen operating system.^c As outlined in Figure 6, the SDK-based implementation uses abstraction-layer APIs for radio communication, including the radio-stacking framework, sensing data input, and timing. Both implementations run on TelosB motes with 48KB of program memory and 10KB of RAM.

^c We used Contiki¹⁰ for the measurements reported here but found similar results through other operating systems.

Within these categories of memory the SDK's overhead is measured as follows:

Program memory overhead. From the machine code implementing the common abstraction layer; in our test application, it consumes 5,331B (10.8% of total program memory) compared to 24,750B by Contiki and 950B by the application itself; and

RAM Overhead. The amount associated with the SDK in the example application is 356B (3.6% of total RAM) compared to 1,200B by Contiki and 368B by the application itself.

These figures show that achieving the convenience of a common abstraction layer over all platforms and VTB operation modes—physical, simulated, and emulated—comes at very reasonable cost; the overheads are sufficiently low as to be a minimal impediment on current hardware. Moreover, memory costs are incurred on a pay-for-what-you-use basis depending on which parts of the common abstraction layer are employed in a given experiment, per our component-based approach.

Conclusion

We have proposed and motivated the VTB abstraction, offering comprehensive fidelity/flexibility trade-offs in WSN testbed-based experimentation, as well as the benefits of a federated physical testbed environment.

The VTB implementation is operational and used by an increasing number of experimenters worldwide. We are also developing additional extensions to “harden” it into a more widely available public facility. We are also exploring more experimental scenarios, particularly those involving virtual mobility.³

The underlying federated platform is growing beyond the nine original sites established by the WISEBED project. In particular, a “smart city” site in Santander, Spain, developed by the SmartSantander project (EU FP7 project ICT-257992, www.smartsantander.eu), is being added to the federation and will make available physical sensors already deployed in a real-world environment. In addition, universities in Brazil and Argentina are in the process of becoming WISEBED sites and will enable us to further broaden the

federation approach, push scalability, and increase heterogeneity.

Full documentation of all systems and interfaces discussed here, along with the source code of the implementations, is available at <http://www.wisebed.eu>.

Acknowledgments

This work was partially supported by the European Union under contract IST-2008-224460 (WISEBED). We would like to acknowledge all those who contributed to the design and implementation of the VTB concept in particular and to WISEBED in general. 

References

- Baumgartner, T., Chatzigiannakis, I., Danckwardt, M., Koninis, C., Kröller, A., Mylonas, G., Pfisterer, D., and Porter, B. Virtualising testbeds to support large-scale reconfigurable experimental facilities. In *Proceedings of the Seventh European Conference on Wireless Sensor Networks, Lecture Notes in Computer Sciences*, vol. 5970, J. Silva, B. Krishnamachari, and F. Boavida, Eds. (Coimbra, Portugal, Feb. 17–19). Springer, Berlin/Heidelberg, 2010, 210–223.
- Bernat, J. SmartSantander: The path towards the smart city vision. In *Proceedings of the First ETSI M2M European Telecommunications Standards Institute Machine to Machine Workshop* (Sophia Antipolis, France, Oct. 19–20, 2010).
- Braun, T., Coulson, G., and Staub, T. Towards virtual mobility support in a federated testbed for wireless sensor networks. In *Proceedings of the Sixth Workshop on Wireless and Mobile Ad Hoc Networks* (Kiel, Germany, Mar. 10–11, 2011).
- Chun, B.N., Buonadonna, P., AuYoung, A., Ng, C., Parkes, D.C., Shneidman, J., Snoeren, A.C., and Vahdat, A. Mirage: A microeconomic resource allocation system for sensor network testbeds. In *Proceedings of the Second IEEE Workshop on Embedded Networked Sensors* (Sydney, May 30–31). IEEE Computer Society Press, Washington, D.C., 2005, 19–28.
- Coalesenses GmbH. iSense. Lübeck, Germany, 2006; <http://www.coalesenses.com>
- Cooperating Objects Network of Excellence. *Common Abstractions for Testbed Federation*. FP7-ICT-2007-2-224053, 2011; http://www.cooperating-objects.eu/fileadmin/research/testbed-federation/CONET_D33.pdf
- Coulson, G., Blair, G., Grace, P., Taiani, F., Joolia, A., K. Lee, K., Ueyama, J., and Sivaharan, T. A generic component model for building systems software. *ACM Transactions on Computer Systems* 26, 1 (Mar. 2008), 1–42.
- Crossbow. MicaZ. Milpitas, CA, 2007; <http://www.xbow.com/Products/productdetails.aspx?sid=164>
- Dudek, D., Haas, C., Kuntz, A., Zitterbart, M., Krüger, D., Rothenpieler, P., Pfisterer, D., and Fischer, S. A wireless sensor network for border surveillance (demo). In *Proceedings of the Seventh ACM Conference on Embedded Networked Sensor Systems* (Berkeley, CA, Nov. 4–6). ACM Press, New York, 2009, 303–304.
- Dunkels, A., Gronvall, B., and Voigt, T. Contiki: A lightweight and flexible operating system for tiny networked sensors. In *Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks* (Tampa, FL, Nov. 16–18). IEEE Computer Society Press, Washington, D.C., 2004, 455–462.
- Dutta, P., Hui, J., Jeong, J., Kim, S., Sharp, C., Taneja, J., Tolle, G., Whitehouse, K., and Culler, D. Trio: Enabling sustainable and scalable outdoor wireless sensor network deployments. In *Proceedings of the Fifth International Conference on Information Processing in Sensor Networks* (Nashville, Apr. 19–21). ACM Press, New York, 2006, 407–415.
- Embedded Networks Laboratory. *Tutornet Project: A Tiered Wireless Sensor Network Testbed*. University of Southern California, Los Angeles, 2009; <http://enlusc.edu/projects/tutornet/>
- Erin, E., Arora, A., Ramnath, R., Naik, V., Bapat, S., Kulathumani, V., Sridharan, M., Zhang, H., Cao, H., and Nesterenko, M. Kansei: A testbed for sensing at scale. In *Proceedings of the Fifth International Conference on Information Processing in Sensor Networks* (Nashville, Apr. 19–21). ACM Press, New York, 2006, 399–406.
- Fekete, S. P., Kröller, A., Fischer, S., and Pfisterer, D. Shawn: The fast, highly customizable sensor network simulator. In *Proceedings of the Fourth International Conference on Networked Sensing Systems* (Braunschweig, Germany, June 6–8). IEEE Computer Society Press, Washington, D.C., 2007.
- Girod, L., Ramanathan, N., Elson, J., Stathopoulos, T., Lukac, M., and Estrin, D. Emstar: A software environment for developing and deploying heterogeneous sensor-actuator networks. *ACM Transactions on Sensor Networks* 3, 3 (Aug. 2007).
- Girod, L., Stathopoulos, T., Ramanathan, N., Elson, J., Estrin, D., Osterweil, E., and Schoellhammer, T. A system for simulation, emulation, and deployment of heterogeneous sensor networks. In *Proceedings of the Second International Conference on Embedded Networked Sensor Systems* (Baltimore, Nov. 3–5). ACM Press, New York, 2004, 201–213.
- Global Environment for Network Innovations. *GENI-fying and Federating Autonomous Kansei Wireless Sensor Networks Technical Report*. Ohio State University, Columbus, OH; <http://groups.geni.net/geni/wiki/KanseiSensorNet>
- Handziski, V., Kopke, A., Willig, A., and Wolisz, A. *Twist: A Scalable and Reconfigurable Wireless Sensor Network Testbed for Indoor Deployments*, Technical Report TKN-05-008. Technische Universität Berlin, Berlin, Germany, Nov. 2005.
- Hill, J., Szewczyk, R., Woo, A., Hollar, S., Culler, D., and Pister, K. System architecture directions for networked sensors. *SIGOPS Operating Systems Review* 34, 5 (Nov. 2000), 93–104.
- Hurni, P., Wagenknecht, G., Anwander, M., and Braun, T. A testbed management architecture for wireless sensor network testbeds (TARWIS). In *Proceedings of the Seventh European Conference on Wireless Sensor Networks, Posters, and Demos* (Coimbra, Portugal, Feb. 17–19). Springer-Verlag, Berlin/Heidelberg, 2010.
- Information Sciences Institute. *The Network Simulator NS-2*. 2007; University of Southern California, Los Angeles; <http://www.isi.edu/nsnam/ns/>
- Levis, P., Lee, N., Welsh, M., and Culler, D. Tossim: Accurate and scalable simulation of entire TinyOS applications. In *Proceedings of the First International Conference on Embedded Networked Sensor Systems* (Los Angeles, Nov. 5–7). ACM Press, New York, 2003, 126–137.
- Li, W., Zhang, X., Tan, W., and Zhou, X. H-TOSSIM: Extending TOSSIM with physical nodes. *Wireless Sensor Networks* 1, 4 (Nov. 2009), 324–333.
- Lipphardt, M., Hellbrok, H., Pfisterer, D., Ransom, S., and Fischer, S. Practical experiences on mobile interbody-area-networking. In *Proceedings of the Second International Conference on Body Area Networks* (Florence, Italy, June 11–13). Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, Brussels, 2007.
- Memsic Corporation. *TelosB*. Andover, MA, 2005; <http://www.memsic.com/support/documentation/wireless-sensor-networks/category/7-datasheets.html?download=152%3Atelosb>
- Österlind, F., Dunkels, A., Voigt, T., Tsiftes, N., Eriksson, J., and Finne, N. SensorNet checkpointing: Enabling repeatability in testbeds and realism in simulations. In *Proceedings of the Sixth European Conference on Wireless Sensor Networks* (Cork, Ireland, Feb. 11–13). Springer-Verlag, Berlin/Heidelberg, 2009, 343–357.
- Park, S., Savvides, A., and Srivastava, M.B. Sensorsim: A simulation framework for sensor networks. In *Proceedings of the Third ACM International Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile Systems* (Boston, Aug. 6–11). ACM Press, New York, 2000, 104–111.
- Porter, B., Coulson, G., and Roedig, U. Type-safe updating for modular WSN software. In *Proceedings of the Seventh IEEE Conference on Distributed Computing in Sensor Systems* (Barcelona, June 27–29). IEEE Press, Washington, D.C., 2011, 1–8.
- Project Sun SPOT. Oracle Labs, Redwood Shores, CA, 2006; <http://www.sunspotworld.com>
- SensLAB. *Very Large-Scale Open Wireless Sensor Network Testbed*. Lyon, France, 2010; <http://www.senslab.info/>
- Staub, T., Gantenbein, R., and Braun, T. Virtualmesh: An emulation framework for wireless mesh and ad

hoc networks in omnet++. *Simulation: Transactions of the Society for Modeling and Simulation International* (special issue on software tools, techniques, and architectures for computer simulation) (Jan. 2011), 66–81.

- Weiser, M. Some computer science issues in ubiquitous computing. *Commun. ACM* 36, 7 (July 1993), 75–84.
- Wen, Y., Zhang, W., Wolski, R., and Chohan, N. Simulation-based augmented reality for sensor network development. In *Proceedings of the Fifth International Conference on Embedded Networked Sensor Systems* (Sydney, Nov. 6–9). ACM Press, New York, 2007, 275–288.
- Werner-Allen, G., Swieskowski, P., and Welsh, M. Motelab: A wireless sensor network testbed. In *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks* (Los Angeles, Apr. 25–27). IEEE Press, Piscataway, NJ, 2005.
- WISEBED Project. *Wireless Sensor Network Testbeds*. Lübeck, Germany; <http://www.wisebed.eu/>
- Wu, H., Luo, Q., Zheng, P., He, B., and Ni, L.M. Accurate emulation of wireless sensor networks. In *Proceedings of Network and Parallel Computing* (Wuhan, China, Oct. 18–20). Springer-Verlag, Berlin/Heidelberg, 2004, 576–583.

Geoff Coulson (geoff@comp.lancs.ac.uk) is a professor of distributed computing in the School of Computing and Communications in the Faculty of Applied Sciences at Lancaster University, Lancaster, U.K.

Barry Porter (barry.porter@comp.lancs.ac.uk) is a research associate in middleware and systems research in the School of Computing and Communications in the Faculty of Applied Sciences at Lancaster University, Lancaster, U.K.

Ioannis Chatzigiannakis (ichatz@cti.gr) is director of Research Unit 1 of the Computer Technology Institute & Press and adjunct faculty in the Computer Engineering & Informatics Department of the University of Patras, Rion, Greece.

Christos Koninis (koninis@cti.gr) is a research fellow at Research Unit 1 of the Computer Technology Institute & Press and a Ph.D. candidate in the Computer Engineering & Informatics Department of the University of Patras, Rion, Greece.

Stefan Fischer (fischer@itm.uni-luebeck.de) is a professor of networks and distributed systems in the Institute for Telematics of the University of Lübeck, Lübeck, Germany.

Dennis Pfisterer (pfisterer@itm.uni-luebeck.de) is a senior researcher and tenured lecturer in the Institute of Telematics at the University of Lübeck, Lübeck, Germany.

Daniel Bimschas (bimschas@itm.uni-luebeck.de) is an assistant researcher in the Institute of Telematics at the University of Lübeck, Lübeck, Germany, and lead software engineer of the WISEBED WSN testbed reference implementation and co-author of the WISEBED APIs.

Torsten Braun (braun@iam.unibe.ch) is professor of computer science and head of the research group Computer Networks and Distributed Systems at the University of Bern, Bern, Switzerland.

Philipp Hurni (hurni@iam.unibe.ch), **Markus Anwander** (anwander@iam.unibe.ch), and **Gerald Wagenknecht** (wagen@iam.unibe.ch) were all Ph.D. students in Braun's research group where they developed the TARWIS system.

Sándor P. Fekete (s.fekete@tu-bs.de) is a professor of algorithmics in the Computer Science Department of the Braunschweig Institute of Technology, Braunschweig, Germany.

Alexander Kröller (a.kroeller@tu-bs.de) is an assistant professor of algorithmics in the Computer Science Department of the Braunschweig Institute of Technology, Braunschweig, Germany.

Tobias Baumgartner (t.baumgartner@tu-bs.de) is a research scientist in the Algorithms Group of the Computer Science Department of the Braunschweig Institute of Technology, Braunschweig, Germany.

How to guarantee files encrypted and transmitted today stay confidential for years to come.

BY CHI-SUNG LAIH, SHANG-MING JEN, AND CHIA-YU LU

Long-Term Confidentiality of PKI

The ubiquitous cryptographic Public Key Infrastructure (PKI) faces a multitude of privacy-protection risks. A notable issue is long-term security, which can be deconstructed into long-term authenticity and long-term confidentiality. Authenticity has been widely discussed over the past decade, while confidentiality

has generally been neglected. As the factorization of RSA advances, there is increased urgency to refresh the confidentiality of existing instances of PKI with longer-duration validity. Unfortunately, cryptographers have not come up with a realistic solution to the question of how to guarantee long-term confidentiality, the most challenging unaddressed open problem from previous works. In this article, we formalize the problem by defining

the concept of a Privacy-Free Window (PFW) where the previously protected file is now at risk. By taking advantage of a PKI property called “asymmetric secrecy,” we give a specific solution addressing PFW. This method can be further developed to extend the originally defined security duration of some PKIs and other cryptographic tools. We also furnish an algorithm to verify existing protocols and recommend actions for maintaining security as a PFW occurs.

PKI applications are everywhere in modern information technology, including e-commerce, Secure Sockets Layer/Transport Layer Security, and citizen-to-government and government-to-government applications (such as those involving tax reporting, medical insurance, and passports). However, maintaining PKI security is increasingly complicated since cryptographic attacks are more sophisticated than ever.

» key insights

- **Upgrading all PKI instances is systemically and logistically costly since PKI structures and scope are so variable.**
- **We define PFW and quantify long-term confidentiality to highlight the insecure period encountered by encrypted files.**
- **The asymmetric secrecy property is a practical, software-based low-cost solution requiring negligible changes to existing system hardware.**

Long-term security is an important cryptographic issue, as discussed by Lenstra and Verheul⁷ and can be divided into long-term authenticity and long-term confidentiality. Long-term authenticity focuses on the validation of digital signatures that ensures integrity, non-repudiation, and authentication over a period of time. Long-term confidentiality emphasizes privacy protection and securing a secret over a period of time and was viewed by Buchmann and Vollmer² as the most challenging open issue in PKI security. Most previous work emphasized authenticity in long-term security, as discussed in RFC 3126: Electronic Signature Format for Long-term Electronic Signatures.¹¹ Signatures are usually used to sign plaintext, so privacy is not the main concern and can be repackaged to extend validity. Other research (such as Buchmann and Vollmer²) focused on mitigating risk due to advanced computer performance (such as quantum computing) and modern cryptanalysis technology. Using quantum cryptographic methods (such as Okamoto et al.¹⁰) or applying multiple cryptosystems² is usually recommended against this kind of weakness.

To the best of our knowledge, most of the work we have just outlined focused on long-term authenticity rather than on long-term confidentiality, especially in the realm of PKI. A key point is that plaintext is disclosed with the signature, so authenticity may be addressed after signing. On the other hand, confidentiality is protected through a form of ciphertext; once encrypted data is sent, the original participants no longer have control and should assume that attackers are able to eavesdrop and archive the encrypted data. Once ciphertext is communicated, it cannot be recalled for re-encryption. Its security depends on adequate measures determined (or guessed) before sending. These are some of the reasons achieving long-term confidentiality is more difficult than achieving long-term authenticity.

Here, we use the RSA public-key algorithm¹² to illustrate the risk of unintended exposure, since most PKIs implement RSA to provide security. PKI faces two main threats: quantum



PFW is the insecure period starting when the protected file may be compromised (due to advancing cryptanalysis) and ending with its planned expiration.



computation and advancing factorization. If and when quantum computation is achieved, integer factorization in sub-exponential time will render RSA obsolete. However, a more pressing risk is the factorization progress of RSA. In August 1999, Cavallar et al.⁴ factored a 512b (155 decimal digits) RSA modulus, and in December 2009, Kleinjung et al.⁶ factored a 768b (232 decimal digits) RSA modulus. Though factoring a 1,024b RSA modulus is thousands of times more difficult than factoring a 768b modulus, it is expected the 1,024b RSA modulus will be factored by 2020.^{1,8} The estimated data and recommendations concerning the most effective ways to ensure security of various algorithms can be found in Lenstra and Verheul,⁷ NIST FIPS SP800-57,¹ and other works.

For this article, we set aside the risk from yet unrealized theoretical quantum methods to focus on solving factorization challenges facing PKIs in the near future. An important concept is that the RSA algorithm (based on factorization) is considered secure today. The deeper issue is the length of the keys used in the RSA algorithm. Most current PKI users utilize RSA keys less than or equal to 1,024b, with factoring of this key length expected by 2020. Using longer public keys may be a sound solution, but intrinsic to PKI is infrastructure, and upgrading an existing infrastructure is more difficult than creating a software-only solution. Here, we develop a low-cost framework for addressing long-term confidentiality within PKI.

Privacy-Free Window

Legal and business documents have their own retention requirements and must be kept secret for a specified period. If a file is encrypted by a PKI exchange session (or secret) key, the file's security depends not only on the session key but also on other keys, including those used to protect the session key during the exchange.

Long-term confidentiality of a file is contingent on ensured secrecy within the specified retention period so it remains uncompromised by unauthorized third parties. The scenario we explore here involves two participants who want to transmit a secret file without a shared key. Each pos-

sesses a public-key pair stored on an integrated circuit (IC) card issued by a certificate authority (CA). The participant and the CA do not require a pre-shared key. Since the security of the file relies on the encrypting session key that in turn relies on the security of the related encrypting public keys and session keys (used to protect against disclosure to third parties), the security of the file and all related keys are interdependent.

We define PFW as the insecure period starting when the protected file may be compromised (due to advancing cryptanalysis) and ending with its planned expiration (see Figure 1). PFW can serve as an indicator for quantifying long-term confidentiality. Based on retention period and key strength, not all files will exhibit a PFW. If a PFW exists, then long-term confidentiality has been compromised; the table here lists several parameters to illustrate the risks associated with the PFW.


In Figure 1, the file is encrypted at the current time t_c and must be kept secret until expiration time t_e at the end of retention period. At time t_p , the public keys generated at t_c are factored, and the protected file is no longer safe, since the session keys protected by the public keys can be obtained. Even when public keys are not factored, session keys may be independently compromised at time t_s through advanced cryptanalysis. If either event happens before t_e , the protected file is considered compromised. The PFW for a file under an associated protocol can formally be expressed as

$$\min(t_p, t_s) > t_c \text{ (broken keys are not used for encryption)}$$

$$PFW = [\min(t_p, t_s), t_e]$$

If the signature is in jeopardy, the responsible entities can easily extend the validity of the signature through encapsulation. Standards (such as RFC3126¹¹) are available for encapsulation, but long-term confidentiality is not well researched. Previously proposed approaches would usually apply safer cryptographic methods, using longer encryption keys or shorter retention periods. These measures are

Figure 1. Privacy-free window.



viable in software but often fail to account for hardware limitations.

The concept of session keys in our scenario is a software issue. A sound recommendation for long-term secure communications is proactive selection of algorithms and keys with adequate security levels suitable for the entire retention period. For example, the Advanced Encryption Standard (AES) with a 128b key and 10-round encryption is adequate for achieving confidentiality for at least the next 20 years.¹ Even stronger AES can provide more robust protection. This implies it is trivial to make $t_s > t_p$ and is also a prerequisite for the following discussion. However, algorithms and keys can be limited by non-software circumstances. For example, a PKI may contain millions of IC cards, as in the case of citizen identification. In this instance, card renewal involves not only cost but also logistical concerns arising from updating or renewing cards within a reasonable amount of time. To surmount this problem we must contend with the bottleneck of infrastructure, as well as update existing PKI systems with minimal change. In the following sections, we propose a cost-efficient concept to mitigate the current PFW problem within PKI and maintain security for longer periods than originally envisioned.

Asymmetric Secrecy Property in PKI

PKI users commonly use RSA with 1,024b keys, while CAs commonly use 2,048b to 4,096b keys for safety reasons. For example, the Ministry of the Interior's Certification Authority of Taiwan (MOICA) uses a CA with 2,048b public/private keypair and administers more than 2.48 million users holding National Identification

Parameters of events and their times.

Time	Description
t_c	The time a file is encrypted.
t_p	The public keys used in the PKI at time t_c are factored or compromised.
t_s	The session keys used in the PKI at time t_c are compromised.
t_e	The expiration time of a file.

IC cards with 1,024b public/private keypair.⁹ The asymmetric secrecy property manifests in PKI because the CA usually uses cryptographic parameters with higher bit keys than normal users. This means the user-to-CA link is much safer than the CA-to-user link, since the former is protected by the CA's public key, whereas the latter is protected by the user's public key. By exploiting this property to employ a new protocol on the existing infrastructure, security issues may be mitigated without having to renew the entire infrastructure.

Consider the following example of the PFW utilizing a modified instance of an existing protocol. Kerberos,⁸ which Burrows et al.³ proved was logically secure for authentication, assumes pre-shared session keys between participants (users) and the server (CA). Our Kerberos analogue (see Figure 2) presupposes no shared keys between participants and server, using instead public keys to provide equivalent security (proved by Burrows's method³) at the time of encryption. In this scenario, A and B are principals, and S is the server. In message 1, key exchange is initiated by A , who expresses to S the desire to communicate with B . In message 2, S responds with a message encrypted by A 's public

key PU_A containing a ticket encrypted by B 's public key PU_B , the session key K_{AB} , and other temporary parameters. Principal A then sends the received ticket to share K_{AB} with B . Finally, B replies to A using a message encrypted by K_{AB} to complete the key exchange.

Since all participants are able to generate suitable session keys, as the software solution assumes in our scenario, all session keys and symmetric encryption algorithms are adequately robust and able to achieve long-term confidentiality required during the retention period. We also assume encrypting time $t_c = 2012$ and a retention period for the encrypted file of 15 years, or $t_e = 2027$.

In this example, the steps pertaining to long-term confidentiality consist of message 2 in which S responds to A and message 3 in which A connects to the requested resource, B . In

message 2, K_{AB} is encrypted through 1,024b PU_A . According to Kleinjung et al.,⁶ RSA-1024 may be factored around 2019. If factored as predicted, a PFW will exist from 2019 to 2027. PU_B in Step 3 suffers from the same problem. When it occurs, the scheme will not provide long-term confidentiality.

Figure 3 is our modified protocol taking advantage of the asymmetric secrecy property inherited from Kerberos. To ensure the feasibility of authentication, see the online appendix, where we prove the necessary authenticating behaviors—results (1) to (4) in the appendix—are maintained. They are analogous to the original Kerberos functions, as outlined by Burrows et al.³ Compared to Kerberos, our modified protocol involves two major changes: First, an unnecessary ticket in message 2 is removed, since it was redundant in Burrows³; and sec-

ond, we used the asymmetric secrecy property to exchange new safer session keys for server-to-principal links, originally protected by the principals' public keys. In our protocol, we define two new session keys, K_{AS} and K_{BS} , with exchanges protected by the safe principal-to-server links in messages 1 and 3-2. It then uses K_{AS} and K_{BS} to protect the server-to-principal links in messages 2 and 3-3.

Since authentication functionalities are proved, we now examine confidentiality in our protocol. In message 1 and message 3-2, the public key PU_{KS} is at 2,048b to 4,096b, adequate for keeping the session keys secure until the retention period ends in 2027, according to Barker et al.¹ and Lenstra and Verheul.⁷ Aside from these steps, session keys K_{AS} , K_{BS} , and K_{AB} are assumed secure since they can be regenerated as needed by principals to satisfy their security requirements.

The server in Kerberos can be viewed as a particular kind of CA. Our protocol demonstrates how the asymmetric secrecy property we have described here helps some PKIs extend their lifetimes through a low-cost method. Ours is primarily software-based for handling inherent hardware constraints, allowing use of the stronger cryptographic capabilities of servers to compensate for unplanned lapses in security assurance. If the target PKI is in a hierarchical structure, the inference is that the root CA may hold the longest public/private keypair. The asymmetric secrecy property may be extended upward to include higher-level CAs providing a longer period of security by virtue of longer keys.

Discussion

Our algorithm (in the appendix) provides criteria for determining the existence of the PFW. No matter which cryptographic method is employed (such as Elliptic Curve Cryptography and RSA), the algorithm is always suitable for evaluating the long-term confidentiality of a protocol. As long as the asymmetric secrecy property holds, the protocol can achieve long-term confidentiality. With the proposed algorithm, if the PFW does not exist, the original protocol is used since the modified protocol is slightly more complex and there is no

Figure 2. Kerberos analogue (key exchange and authentication example).





Figure 3. Modified protocol with asymmetric secrecy property.



benefit from using the asymmetric secrecy property.

Messages protected by participants' public keys in Figure 2 that were once secure are rendered insecure by a PFW. The vulnerability of these encrypting public keys is compensated in Figure 3 by supplemental messages to reinstate acceptable security using the asymmetric secrecy property. These messages consist of communications between participants and the server to exchange secure session keys through the secure channels provided by the server's public key.

Along with existing communication between participant and server, some of these supplemental messages are concatenated with existing messages. For example, the first two messages in Figure 2 and Figure 3 perform equivalent actions, though the communication in Figure 3 derives added benefit from the server's longer keys. When inter-participant communications require additional security (such as message 3 in Figure 2, originally protected by PU_B), another two messages (such as 3-2 and 3-3 in Figure 3) must be introduced to achieve the security requirement. Note the functionality of message 3 in Figure 2 is replaced by messages 3-1 to 3-3 in Figure 3.

In our key-exchange protocol, though two extra messages are required to maintain long-term confidentiality, this added load is required only when the participants are constructing new secure communications channels when a PFW exists. The key exchanges are performed only once per session. This additional overhead is marginal and should be within the CA's capacity. Even if that capacity is exhausted, additional servers can scale the service and solve the capacity problem in a centralized manner.

Conclusion

Security of cryptographic algorithms is the most important element in network applications concerning confidentiality and authenticity. All these network activities are based on trust due to cryptography. Modern cryptographic methods provide robust tools for short-term security, but how can they guarantee files encrypted today are also secure until their planned expiration date? Many

measures have been proposed in the literature, but most focus on the signature rather than privacy and lack realistic considerations. In this article, we have defined PFW to highlight the insecure period encountered by an encrypted file and quantify its long-term confidentiality.

PKI is the most common application of cryptography but suffers from a lack of long-term confidentiality. The article's most important contribution is to show how to utilize a very significant property in PKI we call the asymmetric secrecy property. By exploiting it, we provide a practical software-based, low-cost solution requiring little change to existing system hardware. Our secure modified protocol gives credence to the proposed method. An algorithm for evaluating protocols for key and data exchange is described in the appendix, along with two examples requiring different retention periods and related reactions.

Upgrading all instances of PKI is a systemic issue, since various PKIs involve different structures and purposes. However, the concept we have discussed—how to ensure long-term confidentiality—may help mitigate PKI shortcomings. Furthermore, the asymmetric secrecy concept can also be extended to upgrade existing cryptographic tools (such as Kerberos) and other protocols. Addressing PFW through uncomplicated, low-cost solutions is especially appealing today, ahead of the arrival of quantum computing. In this way, long-term security, especially involving privacy, can be ensured for the near future.

Acknowledgment

We wish to express our gratitude to Michael Chang (macst34@gmail.com) for discussions and his review of the English in this article. 

References

1. Barker, E., Barker, W., Burr, W., Polk, W., and Smid, M. *NIST Special Publication 800-57: Recommendation for Key Management*. National Institute of Standards and Technology, Gaithersburg, MD, May 2007; <http://csrc.nist.gov/publications/PubsSPs.html>
2. Buchmann J., May, A. and Vollmer, U. Perspective for cryptographic long-term security. *Commun. ACM* 49, 9 (Sept. 2006), 50–55.
3. Burrows, M., Abadi, M., and Needham, R. A logic of authentication. *ACM Transactions on Computer Systems* 8, 1 (Feb. 1990), 18–36.
4. Cavallar, S., Dodson, B., Lenstra, A.K., Lioen, W., Montgomery, P.L., Murphy, B., Riele, H., Aardal, K., Gilchrist, J., Guillerm, J., Leyland, P., Marchand, J., Morain, F., Muffett, A., Putnam, C., and Zimmermann,

- P. Factorization of a 512-bit RSA modulus. In *Proceedings of the 19th International Conference on Theory and Application of Cryptographic Techniques*, B. Preneel, Ed. Springer-Verlag, Berlin, Heidelberg, 2000, 1–18.
5. ElGamal, T. A public-key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory* 3, 4 (July 1985), 469–472.
6. Kleinjung, T., Aoki, K., Franke, J., Lenstra, A.K., Thomé, E., Bos, J.W., Gaudry, P., Kruppa, A., Montgomery, P.L., Osvik, D.A., Riele, H., Timofeev, A., and Zimmermann, P. Factorization of a 768-bit RSA modulus. In *Proceedings of the 30th Annual Conference on Advances in Cryptology*, T. Rabin, Ed. Springer-Verlag, Berlin, Heidelberg, 2010, 333–350.
7. Lenstra, A.K. and Verheul, E.R. Selecting cryptographic key sizes. In *Proceedings of the Third International Workshop on Practice and Theory in Public Key Cryptography*, H. Imai and Y. Zheng, Eds. Springer-Verlag, London, 2000, 446–465.
8. Miller, S.P., Neuman, B.C., Schiller, J.I., and Saltzer, J.H. Kerberos authentication and authorization system. In *Project Athena Technical Plan*, Section E.2.1. MIT, Cambridge, MA, Oct. 1988; <http://web.mit.edu/Saltzer/www/publications/athenasplan/e.2.1.pdf>
9. MOICA (Certificate Authority of the Ministry of the Interior of Taiwan). Dec. 2010; <http://moica.nat.gov.tw/html/en/index.htm>
10. Okamoto, T., Tanaka, K., and Uchiyama, S. Quantum public-key cryptosystems. In *Proceedings of the 20th Annual International Cryptology Conference on Advances in Cryptology*, M. Bellare, Ed. Springer-Verlag, London, 2000, 147–165.
11. Pinkas, D., Ross, J., and Pope, N. *Electronic Signature Formats for Long-Term Electronic Signatures*. IETF RFC 3126, Sept. 2001; <http://www.rfc-editor.org/rfc/rfc3126.txt>
12. Rivest, R., Shamir, A., and Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21, 2 (Feb. 1978), 120–126.

Chi-Sung Lai (1956–2010) was a much-loved and respected distinguished professor in the Department of Electrical Engineering of National Cheng Kung University, Tainan City, Taiwan.

Shang-Ming Jen (smjen.tw@gmail.com) is a Ph.D. student in the Department of Electrical Engineering of National Cheng Kung University, Tainan City, Taiwan.

Chia-Yu Lu (joylu.tw@gmail.com) is a Ph.D. student in the Department of Electrical Engineering of National Cheng Kung University, Tainan City, Taiwan.

Computer vision holds the key for the blind or visually impaired to explore the visual world.

BY ROBERTO MANDUCHI AND JAMES COUGHLAN

(Computer) Vision Without Sight

MORE THAN 20 million people in the U.S. live with visual impairments ranging from difficulty seeing, even with eyeglasses, to complete blindness. Vision loss affects almost every activity of daily living. Walking, driving, reading and recognizing objects, places, and people become difficult or impossible without vision. Thus, technology that can assist visually impaired (VI) people in at least some of these tasks may have a very relevant social impact.

Research in assistive technology for VI people has resulted in some very useful hardware and software tools in widespread use. The most successful products to date include text magnifiers and screen readers, Braille note takers, and document scanners with optical character recognition (OCR). This article focuses specifically on the use of computer vision systems and algorithms to support VI people in their daily tasks. Computer vision seems like a natural choice for these applications—in a sense, replacing the lost sense of sight with an “artificial eye.” Yet, in spite of the success of computer vision technology

in several other fields (such as robot navigation, surveillance, and user interface), very few computer vision systems and algorithms are currently employed to aid VI people.

In this article, we review current research work in this field, analyze the causes of past failed experiences, and propose promising research directions marrying computer vision and assistive technology for the VI population. Our considerations stem in large part from our own direct experience developing technology for VI people, and from conducting the only specific workshop on Computer Vision Applications for the Visually Impaired, which was held in 2005 (San Diego), 2008 (Marseille), and 2010 (San Francisco).

The VI Population

The VI community is very diverse in terms of degree of vision loss, age, and abilities. It is important to understand the various characteristics of this population if one is to design technology that is well fit to its potential “customers.” Here is some statistical data made available by the American Foundation for the Blind: Of the 25 or more million Americans experiencing significant vision loss, about 1.3 million are legally blind (meaning their visual field in their best eye is 20 degrees or less or their acuity is less than 20/200), and only about 290,000 are totally blind (with at most some light

» key insights

- **Computer vision and mobile computing are powerful tools with great potential to enable a range of assistive technologies for the growing population of blind and visually impaired people.**
- **The actual needs of this population must drive the development of this technology in order for it to be truly useful and likely to be widely adopted. To this end, blind and visually impaired users must be involved at all stages of design, research, and development.**
- **Particular attention must be paid to the development of appropriate user interfaces for this technology that respect the needs of blind and visually impaired users.**



perception). Since the needs of a low-vision person and of a blind person can be very different, it is important not to overgeneralize the nature of visual impairment. Another important factor to be considered is the age of a VI person. Vision impairment is often due to conditions such as diabetic retinopathy, macular degeneration, and glaucoma that are prevalent at later age. Indeed, about one-fourth of those reporting significant vision loss are 65 years of age and older. It is important to note that multiple disabilities in addition to vision loss are also common at later age (such as hearing impairment due to presbycusis or mobility impairment due to arthritis). In regard to the younger U.S. population, about 60,000 individuals 21 years old or younger are legally blind. Of these, fewer than 10% use Braille as their primary reading medium.

Application Areas

Mobility. In the context of assistive technology, mobility takes the meaning of “moving safely, gracefully, and comfortably;”³ it relies in large part on perceiving the properties of the immediate surroundings, and it entails avoiding obstacles, negotiating steps, drop-offs, and apertures such as doors, and maintaining a possibly rectilinear trajectory while walking. Although blind people are more in need

of mobility aids, low-vision individuals may also occasionally trip onto unseen small obstacles or steps, especially in poor lighting conditions.

The most popular mobility tool is the white cane (known in jargon as the long cane), with about 110,000 users in the U.S. The long cane allows one to extend touch and to “preview” the lower portion of the space in front of oneself. Dog guides may also support blind mobility, but have many fewer users (only about 7,000 in the U.S.). A well-trained dog guide helps maintain a direct route, recognizes and avoids obstacles and passageways that are too narrow to go through, and stops at all curbs and at the bottom and top of staircases until told to proceed. Use of a white cane or of a dog guide publicly identifies a pedestrian as blind, and carries legal obligations for nearby drivers, who are required to take special precautions to avoid injury to such a pedestrian.

A relatively large number of devices meant to provide additional support, or possibly to replace the long cane and the dog guide altogether, have been proposed over the past 40 years. Termed Electronic Travel Aids or ETAs,³ these devices typically utilize different types of range sensors (sonars, active triangulation systems, and stereo vision systems). Some ETAs are meant to simply give an indication of

the presence of an obstacle at a certain distance along a given direction (clear path indicators). A number of ETAs are mounted on a long cane, thus freeing one user’s hand (but at the expense of adding weight to the cane and possibly interfering with its operation). For example, the Nurion Laser Cane (no longer in production) and the Laser Long Cane produced by Vistac use three laser beams to detect (via triangulation) obstacles at head-height level, while the UltraCane (formerly BatCane) produced by Sound Foresight uses sonar on a regular cane to detect obstacles up to height level. A different type of ETA—the Sonic Pathfinder, worn as a special spectacle frame, and the Bat K-Sonar, mounted on a cane—use one or more ultrasound transducers to provide the user with something closer to a “mental image” of the scene (such as the distance and direction of an obstacle and possibly some physical characteristics of its surface.)

In recent years, a number of computer vision-based ETAs have been proposed. For example, a device developed by Yuan and Manduchi⁴⁰ utilizes structured light to measure distances to surfaces and to detect the presence of a step or a drop-off at a distance of a few meters. Step and curb detection can also be achieved via stereo vision.²⁵ Range data can be integrated through time using a technique called “simul-




Figure 1. Crosswalk system for providing guidance to VI pedestrians at traffic intersections. (a) Blind user “scans” the crosswalk by panning cellphone camera left and right, and system provides feedback to help user align himself to crosswalk before entering it. (b) Schematic shows that system announces to user when the Walk light is illuminated.

aneous localization and mapping” (SLAM), allowing for the geometric reconstruction of the environment and for self-localization. Vision-based SLAM, which has been used successfully for robotic navigation, has been recently proposed as a means to support blind mobility.^{26,28,37} Range cameras, such as the popular PrimeSense’s Kinect, also represent a promising sensing modality for ETAs.

Although many different types of ETAs have appeared on the market, they have met with little success by the intended users so far. Multiple factors, including cost, usability, and performance, contribute to the lack of adoption of these devices. But the main reason is likely the fact that the long cane is difficult to surpass. The cane is economical, reliable, and long lasting, and never runs out of power. Also, it is not clear whether some of the innovative features of newly proposed ETAs (longer detection range, for example) are really useful for blind mobility. Finally, presenting complex environmental features (such as the direction and distance to multiple obstacles) through auditory or tactile channels can easily overwhelm the user, who is already concentrated on using his or her remaining sensory capacity for mobility and orientation.

Neither the long cane nor the dog guide can protect the user from all types of hazards. One example is given by obstacles that are at head height (such as a propped-open window or a tree branch), and thus are beyond the volume of space surveyed by the cane. In a recent survey of 300 blind and legally blind persons,²¹ 13% of the respondents reported they experience head-level accidents at least once a month. The type of mobility aid (long cane or dog guide) does not seem to have a significant effect on the frequency of such accidents. Another type of hazard is represented by walking in trafficked areas, and in particular crossing a street. This requires awareness of the environment around oneself as well as of the flow of traffic, and good control of one’s walking direction to avoid drifting out of the crosswalk. Technology that increases the pedestrian’s safety in these situations may be valuable, such as a mobile phone system us-

ing computer vision to orient the user to the crosswalk and to provide information about the timing of Walk lights^{12,13} (see Figure 1).

Wayfinding. Orientation (or wayfinding) can be defined as the capacity to know and track one’s position with respect to the environment, and to find a route to a destination. Whereas sighted persons use visual landmarks and signs in order to orient themselves, a blind person moving in an unfamiliar environment faces a number of hurdles:²⁰ accessing spatial information from a distance; obtaining directional cues to distant locations; keeping track of one’s orientation and location; and obtaining positive identification once a location is reached.

According to Loomis et al.,²⁰ there are two main ways in which a blind person can navigate with confidence in a possibly complex environment and find his or her way to a destination: piloting and path integration. Piloting means using sensory information to estimate one’s position at any given time, while path integration is equivalent to the “dead reckoning” technique of incremental position estimation, used for example by pilots and mariners. Although some blind individuals excel at path integration, and can easily retrace a path in a large environment, this is not the case for most blind (as well as sighted) persons.

Path integration using inertial sensors or visual sensors has been used extensively in robotics, and a few attempts at using this technology for blind wayfinding have been reported.^{9,18} However, the bulk of research on wayfinding has focused on piloting, with very promising results and a number of commercial products already available. For outdoor travelers, GPS represents an invaluable technology. Several companies offer GPS-based navigational systems specifically designed for VI people. None of these systems, however, can help the user in tasks such as “Find the entrance door of this building,” due to the low spatial resolution of GPS reading and to the lack of such details in available GIS databases. In addition, GPS is viable only outdoors. Indoor positioning systems (for example, based on multilateration from WiFi beacons) are gaining momentum, and

it is expected they will provide interesting solutions for blind wayfinding.

A different approach to wayfinding, one that does not require a geographical database or map, is based on recognizing (via an appropriate sensor carried by the user) specific landmarks placed at key locations. Landmarks can be active (light, radio, or sound beacons) or passive (reflecting light or radio signals). Thus, rather than absolute positioning, the user is made aware of their own relative position and attitude with respect to the landmark. This may be sufficient for a number of navigational tasks, for example, when the landmark is placed near a location of interest. For guidance to destinations that are beyond the landmark’s “receptive field” (the area within which the landmark can be detected), a route can be built as a set of waypoints that need to be reached in sequence. Contextual information about the environment can also be provided to the VI user using digital map software and synthetic speech.¹⁴

The best-known beaconing system for the blind is Talking Signs, now a commercial product based on technology developed at The Smith-Kettlewell Eye Research Institute.^a Already deployed in several cities, Talking Signs uses a directional beacon of infrared light, modulated by a speech signal. This can be received at a distance of several meters by a specialized handheld device, which also demodulates the speech signal and presents it to the user. RFID technology has also been proposed recently in the context of landmark-based wayfinding for the blind.¹⁶ Passive RFIDs are small, inexpensive, and easy to deploy, and may contain several hundreds of bits of information. The main limitation of RFID systems is their limited reading range and lack of directionality.

A promising research direction is the use of computer vision to detect natural or artificial landmarks, and thus assist in blind wayfinding. A VI person can use their own cellphone, the camera pointing forward, to search for landmarks in view. Natural landmarks are distinctive environmental features that can be detected robustly, and used for guidance either using an existing map¹¹ or by matching against possibly

^a <http://talkingsigns.com>

geotagged image datasets.^{10,19} Detection is usually performed by first identifying specific key points in the image; the brightness or color image profile in the neighborhood of these key points is then represented by compact and robust descriptors. The presence of a landmark is tested by matching the set of descriptors in an image against a dataset formed by exemplar images collected offline. Note that some of this research work (for example, Hile et al.¹¹) was aimed to support navigation in indoor spaces for persons with cognitive impairments. Apart from the display modality, the same technology is applicable for assistance to visually impaired individuals.

Artificial landmarks are meant to facilitate the detection process. For example, the color markers developed by Coughlan and Manduchi^{5,22} (see Figure 2) are designed so as to be highly distinctive (thus minimizing the rate of false alarms) and easily detectable with very moderate computational cost (an important characteristic for mobile platforms such as cellphones with modest computing power). A similar system, designed by researchers in Gordon Legge's group at the University of Minnesota, uses retro-reflective markers that are detected by a "Magic Flashlight," a portable camera paired with an infrared illuminator.³³

Artificial landmarks can be optimized for easy and fast detection by a mobile vision system. This is an advantage with respect to natural landmarks, whose robust detection is more challenging. On the other hand, artificial landmarks (as well as beacons such as Talking Signs) involve an infra-

structure cost—they must be installed and maintained, and represent an additional element to be considered in the overall environment design. This trade-off should be considered carefully when developing wayfinding technology. It may be argued that the additional infrastructure cost could be better justified if other communities of users in addition to the VI population would benefit from the wayfinding system. For example, even sighted individuals who are unfamiliar with a certain location (for example, a shopping mall), and cannot read existing signs (because of a cognitive impairment, or possibly because of a foreign language barrier), may find a guidance system beneficial. Under this perspective, even the signage commonly deployed for sighted travelers can be seen as a form of artificial landmarks. Automatic reading of existing signs and, in general, of printed information via mobile computer vision, is discussed next.

Printed Information Access. A common concern among the VI population is the difficulty of accessing the vast array of printed information that normally sighted people take for granted in daily life. Such information ranges from printed documents such as books, magazines, utility bills, and restaurant menus to informational signs labeling streets, addresses, and businesses in outdoor settings as well as office numbers, exits, and elevators found indoors. In addition, a variety of "non-document" information must also be read, including LED/LCD displays required for operating a host of electronic appliances such as microwave ovens, stoves, and DVD

players, and barcodes or other information labeling the contents of packaged goods such as grocery items and medicine containers.

Great progress has been made in providing solutions to this problem by harnessing OCR, which has become a mature and mainstream technology after decades of development. Early OCR systems for VI users (for example, the Arkenstone Reader and Kurzweil Reading Machine) were bulky machines that required the text to be read be imaged using a flatbed scanner. More recent incarnations of these systems have been implemented in portable platforms such as mobile (cell) phones (for example, the KNFB reader^b) and tablets (for example, the IntelReader^c), which allow the user to point the device's camera toward a document of interest and have it read aloud in a matter of seconds. It is important to note that an important challenge of mobile OCR systems for VI users is the difficulty of aiming the camera accurately enough to capture the desired document area; thus, an important feature of the KNFB user interface is that it provides guidance to the user to help him/her frame the image properly.

However, while OCR is effective for reading printed text that is clearly resolved and fills up most of the image, it is not equipped to find text in images that contain large amounts of unrelated clutter, such as an image of a restaurant sign captured from across the street. The problem of text detec-

b <http://knfbreader.com>

c <http://www.intel.com/about/companyinfo/healthcare/products/reader/index.htm>




Figure 2. Experiments with a blind user searching for a landmark (represented by a color marker placed on the wall) using a cellphone camera.²²

tion and localization is an active area of research^{4,29,35,36} that addresses the challenge of swiftly and reliably sorting through visual patterns to distinguish between text and non-text patterns, despite the huge variability of text fonts and background surfaces on which they are printed (for example, the background surface may be textured and/or curved) and the complications of highly oblique viewing perspectives, limited or poor resolution (due to large distances or motion blur), and low contrast due to poor illumination. A closely related problem is finding and recognizing signs²⁴ characterized by non-standard fonts and layouts that may encode important information using shape (such as stop signs and signs or logos labeling business establishments).

To the best of our knowledge, there are currently no commercially available systems for automatically performing OCR in cluttered scenes for VI users. However, Blindsight Corporation's^d Smart Telescope SBIR project seeks to develop a system to detect text regions in a scene and present them to a partially sighted user via a head-mounted display that zooms into the text to enable him/her to read it. Mobile phone apps such as Word Lens go beyond the functionality offered by systems targeted to VI users, such as KNFB, in that they detect and read text in cluttered scenes, though these newer systems are intended for normally sighted users.

Research is under way to expand the reach of OCR beyond standard printed text to "non-document" text such as LED and LCD displays,³² which provide access to an increasingly wide range of household appliances. Such displays pose formidable challenges that make detection and reading difficult, including contrast that is often too low (LCDs) or too high (LEDs), the prevalence of specular highlights, and the lack of contextual knowledge to disambiguate unclear characters (for example, dictionaries are used in standard OCR to find valid words, whereas LED/LCD displays often contain arbitrary strings of digits).

Another important category of non-document text is the printed informa-



Although many different types of Electronic Travel Aids have appeared on the market, they have met with little success by the intended users so far. Multiple factors, including cost, usability, and performance, contribute to the lack of adoption of these devices.



tion that identifies the contents of packaged goods, which is vital when no other means of identification is available to a VI person (for example, a can of beans and a can of soup may feel identical in terms of tactile cues). UPC barcodes provide product information in a standardized form, and though originally designed for use with laser scanners there has been growing interest in developing computer vision algorithms for reading them from images acquired by digital cameras, especially for mobile cell platforms (for example, the Red Laser app^e). Such algorithms⁸ have to cope with noisy and blurred images and the need to localize the barcode in a cluttered image (for example, taken by a VI user who has little prior knowledge of the barcode's location on a package). Some research in this area^{17,31} has specifically investigated the usability of these algorithms by VI persons, and at least one commercial system (DigitEyes^f) has been designed specifically for the VI population. Finally, an alternative approach to package identification is to treat it as an object recognition problem,³⁸ which has the benefit of not requiring the user to locate the barcode, which comprises a small portion of the entire surface of the package.

Object Recognition. Over the past decade, increasing research efforts within the computer vision community have focused on algorithms for recognizing generic "objects" in images. For example, the PASCAL Visual Object Classes Challenge, which attracts dozens of participants every year, evaluates competing object recognition algorithms from a number of visual object classes in challenging realistic scenes.⁸ Another example is Google Goggles, an online service that can be used for automatic recognition of text, artwork, book covers, and more. Other commercial examples include oMoby, developed by IQ Engines, A9's SnapTell, and Microsoft's Bing Mobile application with visual scanning.

Visual object recognition for assistive technology is still in its infancy,

e <http://redlaser.com>

f <http://digit-eyes.com>

g Other benchmarking efforts include the TREC Video Retrieval Evaluation and the Semantic Robot Vision challenge.

d <http://www.blindsight.com/>

with only a few applications proposed in recent years. For example, Winlock et al.³⁸ have developed a prototype system (named ShelfScanner) for assistance to a blind person while shopping at a supermarket. Images taken by a camera carried by the user are analyzed to recognize shopping items from a known set; the user is then informed about whether any of the items in his or her shopping list is in view. LookTel, a software platform for Android phones developed by IPPLEX LLC,³⁹ performs real-time detection and recognition of different types of objects such as bank notes, packaged goods, and CD covers. The detection of doors (that can be useful for wayfinding applications) has been considered in Yang and Tian.³⁹

A Human in the Loop? The goal of the assistive technology described so far is to create the equivalent of a “sighted companion,” who can assist a VI user and answer questions such as “Where am I?” “What’s near me?” “What is this object?”

Some researchers have begun questioning whether an automatic system is the right choice for this task. Will computer vision ever be powerful enough to produce satisfactory results in any context of usage? What about involving a “real” sighted person in the loop, perhaps through crowdsourcing? For example, the VizWiz system² uses Amazon’s Mechanical Turk to provide a blind person with information about an object (such as the brand of a can of food). The user takes a picture of the object, which is then transmitted to Mechanical Turk’s remote work force for visual analysis, and the results are reported back to the user. The NIH-funded “Sight on Call” project by the Blindsight Corporation addresses a similar application. However, rather than relying on crowdsourcing, it uses specially trained personnel interacting remotely with the visually impaired user, on the basis of video streams and GPS data taken by the user’s cellphone and transmitted to the call center.

Interfaces

Each one of the systems and algorithms described here furnishes some information (for example, the presence of an obstacle, the bearing of a

landmark, or the type and brand of items on a supermarket’s shelf) that must be presented to the VI user. This communication can use any of the user’s remaining sensory channels (tactile or acoustic), but should be carefully tailored so as to provide the necessary information without annoying or tiring the user. The fact that blind persons often rely on aural cues for orientation precludes the use of regular headphones for acoustic feedback, but ear-tube earphones and bonephones³⁴ are promising alternatives. In the case of wayfinding, the most common methods for information display include: synthesized speech; simple audio (for example, spatialized sound, generated so as it appears to come from the direction of the landmark); auditory icons;⁶ haptic point interface,²³ a modality by which the user can establish the direction to a landmark by rotating a handheld device until the sound produced has maximum volume; and tactual displays such as “tappers.”²⁷

One major issue to be considered in the design of an interface is whether a rich description of the scene, or only highly symbolic information, should be provided to the user. An example of the former is the vOICE, developed by Peter Mijer, that converts images taken by a live camera to binaural sound. At the opposite end are computer vision systems that “filter” incoming images to recognize specific features, and provide the user with just-in-time, minimally invasive information about the detected object, landmark, or sign.

Usability

Despite the prospect of increased independence enabled by assistive technology devices and software, very few such systems have gained acceptance by the VI community as yet. Here, we analyze some of the issues that, in our opinion, should be taken into account when developing a research concept in this area. It is important to bear in mind that these usability issues can only be fully evaluated with continual feedback from the target VI population obtained by testing the assistive technology as it is developed.

Cosmetics, Cost, Convenience. No one (except perhaps for a few early adopters) wants to carry around a device that

attracts unwanted attention, is bulky or inconvenient to wear or to hold, or detracts from one’s attire. Often, designers and engineers seem to forget these basic tenets and propose solutions that are either inconvenient (for example, interfering with use of the long cane or requiring a daily change of batteries) or simply unattractive (for example, a helmet with several cameras pointing in different directions). A forward-looking extensive discussion of design for disability can be found in the beautiful book *Design Meets Disability* by G. Pullin.

Cost is also an important factor determining usability. Economics of scale is hardly achievable in assistive technology given the relatively small size of the pool of potential users, and the diversity of such a population. This typically leads to high costs for the devices that do make it to the market, which may make them unaffordable by VI users who in many cases are either retired or on disability wages.

Performance. How well should a system work before it becomes viable? The answer clearly depends on the application type. Consider, for example, an ETA that informs the user about the presence of a head-level obstacle. If the system produces a high rate of false alarms, the user will quickly become annoyed and turn the system off. At the same time, the system must have a very low missed detection rate, lest users hurt themselves against undetected obstacles, possibly resulting in medical (and legal) consequences. Other applications may have less stringent requirements. For example, in the case of a cellphone-based system that helps one find a certain item in the grocery store, no harm will be caused to the user if the item is not found or if the wrong item is selected. Still, poor performance is likely to lead to users abandoning the system. Establishing functional performance metrics and assessing minimum performance requirements for assistive technology systems is still an open and highly needed research topic.

Mobile Vision and Usability. The use of mobile computer vision for assistive technology imposes particular functional constraints. Computer vision requires use of one or more cameras to acquire snapshots or video

streams of the scene. In some cases, the camera may be handheld, for example, when embedded in a cellphone. In other cases, a miniaturized camera may be worn by the user, perhaps attached to one's jacket lapel or embedded in one's eyeglasses frames. The camera's limited field of view is an important factor in the way the user interacts with the system to explore the surrounding environment: if the camera is not pointed toward a feature of interest, this feature is simply not visible. Thus, it is important to study how a visually impaired individual, who cannot use feedback from the camera's viewfinder, can maneuver the camera in order to explore the environment effectively. Of course, the camera's field of view could be expanded, but this typically comes at the cost of a lower angular resolution. Another possibility, explored by Winlock et al.,³⁸ is to build a panoramic image by stitching together several images taken by pointing the camera in different directions.

It should be noted that, depending on the camera's shutter speed (itself determined by the amount of light in the scene), pictures taken by a moving camera may be blurred and difficult or impossible to decipher. Thus, the speed at which the user moves the camera affects recognition. Another important issue is the effective frame rate, that is, the number of frames per second that can be processed by the system. If the effective frame rate is too low, visual features in the environment may be missed if the user moves the camera too fast in the search process. For complex image analysis tasks, images can be sent to a remote server for processing (for example, the LookTel platform³⁰), in which case the speed and latency are determined by the communication channel. Hybrid local/remote processing approaches, with scene or object recognition performed on a remote sever and fast visual tracking of the detected feature performed by the cellphone, may represent an attractive solution for efficient visual exploration.

Thus, a mobile vision system for assistive technology is characterized by the interplay between camera characteristics (field of view, resolution), computational speed (effective achiev-

able frame rate for a given recognition task), and user interaction (including the motion pattern used to explore the scene, possibly guided by acoustic or tactile feedback).

Preliminary research work has explored the usability of such systems for tasks such as wayfinding²² and access to information embedded in bar codes.^{17,31}

Conclusion

Advances in mobile computer vision hold great promise for assistive technology. If we can teach computers to see, they may become a valuable support for those of us whose sight is compromised or lost. However, decades-long experience has shown that creating successful assistive technology is difficult. Far too often, engineers have proposed technology-driven solutions that either do not directly address the actual problems experienced by VI persons, or that are not satisfactory in terms of performance level, ease of use, or convenience. Assistive technology is a prime example of user-centered technology: the needs, characteristics, and expectations of the target population must be understood and taken into account throughout the project, and must drive all of the design choices, lest the final product result in disappointment for the intended user and frustration for the designer. Our hope is that a new generation of computer vision researchers will take on the challenge, arm themselves with enough creativity to produce innovative solutions, and have humbleness to listen to the persons who will use this technology.

In closing, we would like to propose a few novel and intriguing application areas that in our opinion deserve further investigation by the research community.

Independent Wheeled Mobility. One dreaded consequence of progressive vision loss (for example, due to an age-related condition) is the ensuing loss of driving privileges. For many individuals, this is felt as a severe blow to their independence. Alternative means of personal wheeled mobility that do not require a driving license could be very desirable to active individuals who still have some degree of vision

left. For example, some low-vision persons reported good experiences using the two-wheel Segway, driven on bicycle lanes.¹ These vehicles could be equipped with range and vision sensors to improve safety, minimizing the risk of collisions and ensuring that the vehicle remains within a marked lane. With the recent emphasis on sensors and machine intelligence for autonomous cars in urban environments, it is only reasonable that the VI community should soon benefit from these technological advances.

Blind Photography. Many people find it surprising that people with low vision or blindness enjoy photography as a recreational activity. In fact, a growing community of VI photographers takes and shares photos of family and friends, of objects, and of locations they have visited; some have elevated the practice of photography to an art form, transforming what would normally be considered a challenge (the visual impairment) into an opportunity for creativity. There are numerous Web sites (for example, <http://blindwithcameraschool.org>), books, and art exhibitions focused on this subject, which could present an interesting opportunity for computer vision researchers. A variety of computer vision techniques such as face detection, geometric scene analysis and object recognition could help a VI user correctly orient the camera and frame the picture. Such techniques, when coupled with a suitable interface, could provide a VI person with a feedback mechanism similar to the viewfinder used by sighted photographers.

Social Interaction. Blindness may, among other things, affect one's interpersonal communication skills, especially in scenarios with multiple persons interacting (for example, in a meeting). This is because communication in these situations is largely nonverbal, relying on cues such as facial expressions, gaze direction, and other forms of the so-called "body language." Blind individuals cannot access these nonverbal cues, leading to a perceived disadvantage that may result in social isolation. Mobile computer vision technology may be used to capture and interpret visual cues from other persons nearby, thus em-

powering the VI user to participate more actively in the conversation. The same technology may also help a VI person become aware of how he or she is perceived by others. A survey conducted with 25 visually impaired persons and two sighted specialists¹⁵ has highlighted some of the functionalities that would be most desirable in such a system. These include: understanding whether one's personal mannerisms may interfere with social interactions with others; recognizing the facial expressions of other interlocutors; and knowing the names of the people nearby.

Assisted Videoscripting. Due to their overwhelmingly visual content, movies are usually considered inaccessible to blind people. In fact, a VI person may still enjoy a movie from its soundtrack, especially in the company of friends or family. In many cases, though, it is difficult to correctly interpret ongoing activities in the movie (for example, where the action is taking place, which characters are currently in the scene and what they are doing) from the dialogue alone. In addition, many relevant nonverbal cues (such as the facial expression of the actors) are lost. Videodescription (VD) is a technique meant to increase accessibility of existing movies to VI persons by adding a narration of key visual elements, which is presented to the listener during pauses in the dialogue. Although the VD industry is fast growing, due to increasing demand, the VD generation process is still tedious and time consuming. This process, however, could be facilitated by the use of semiautomated visual recognition techniques, which have been developed in different contexts (such as surveillance and video database indexing). An early example is VDManager,⁷ a VD editing software tool, which uses speech recognition as well as key-places and key-faces visual recognition.

Acknowledgments

Roberto Manduchi was supported by the National Science Foundation under Grants IIS-0835645 and CNS-0709472. James Coughlan was supported by the National Institutes of Health under Grants 1 R01 EY018345-01, 1 R01 EY018890-01 and 1 R01 EY01821001A1. 

References

- Ackel, W. A Segway to independence. *Braille Monitor*, 2006.
- Bigham, J.P., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R.C., Miller, R., Tatarowicz, A., White, B., White, S. and Yeh, T. VizWiz: Nearly real-time answers to visual questions. In *Proc. ACM Symposium on User Interface Software and Technology* (2010).
- Blasch, B., Wiener, W. and Welsh, R. *Foundations of Orientation and Mobility*. AFB Press, Second Edition, 1997.
- Chen, X. and Yuille, A. Detecting and reading text in natural scenes. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (2004).
- Coughlan, J. and Manduchi, R. Functional assessment of a camera phone-based wayfinding system operated by blind and visually impaired users. *International Journal on Artificial Intelligence Tool* 18, 3 (2009), 379–397.
- Dingler, T., Lindsay, J. and Walker, B.N. Learnability of sound cues for environmental features: Auditory icons, earcons, spearcons, and speech. In *Proc. International Conference on Auditory Display*, (2008).
- Gagnon, L., Chapdelaine, C., Byrns, D., Foucher, S., Heritier, M. and Gupta, V. A computer-vision-assisted system for Videodescription scripting. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2010).
- Gallo, O. and Manduchi, R. Reading 1-D barcodes with mobile phones using deformable templates barcodes with mobile phones using deformable templates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, in press.
- Hesch, J.A. and Roumeliotis, S.I. Design and analysis of a portable indoor localization aid for the visually impaired. *International Journal on Robotics Research* 29, (Sept. 2010), 1400–1415.
- Hile, H., Liu, A., Borriello, G., Grzeszczuk, R., Vedantham, R. and Kosecka, J. Visual navigation for mobile devices. *IEEE Multimedia* 17, 2 (2010), 16–25.
- Hile, H., Vedantham, R., Cuellar, G., Liu, A., Gelfand, N., Grzeszczuk, R. and Borriello, G. Landmark-based pedestrian navigation from collections of geotagged photos. In *Proc. International Conference on Mobile and Ubiquitous Multimedia* (2008).
- Ivanchenko, V., Coughlan, J. and Shen, H. Crosswatch: A camera phone system for orienting visually impaired pedestrians at traffic intersections. In *Proc. International Conference on Computers Helping People with Special Needs* (2008).
- Ivanchenko, V., Coughlan, J. and Shen, H. Real-time walk light detection with a mobile phone. In *Proc. International Conference on Computers helping people with Special Needs* (2010).
- Kalia, A.A., Legge, G.E., Ogale, A. and Roy, R. Assessment of indoor route-finding technology for people who are visually impaired. *Journal of Visual Impairment & Blindness* 104, 3 (Mar. 2010), 135–147.
- Krishna, S., Colbry, D., Black, J., Balasubramanian, V. and Panchanathan, S. A systematic requirements analysis and development of an assistive device to enhance the social interaction of people who are blind or visually impaired. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2008).
- Kutyukin, V. and Kutianawala, A. Accessible shopping systems for blind and visually impaired individuals: Design requirements and the state of the art. *The Open Rehabilitation Journal* 2 (2010).
- Kutianawala, A. and Kutyukin, V. An eyes-free vision-based UPC and MSI barcode localization and decoding algorithm for mobile phones. In *Proc. Envision Conference*, (2010).
- Ladetto, Q. and Merminod, B. Combining gyroscopes, magnetic compass and GPS for pedestrian navigation. In *Proc. Int. Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation* (2001).
- Liu, J., Phillips, C. and Daniilidis, K. Video-based localization without 3D mapping for the visually impaired. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2010).
- Loomis, J.M., Golledge, R.G., Klatzky, R.L. and Marston, J.R. Assisting wayfinding in visually impaired travelers. *Applied Spatial Cognition: From Research to Cognitive Technology*. Lawrence Erlbaum Assoc., Mahwah, NJ, 2007, pages 179–202.
- Manduchi, R. and Kurniawan, S. Mobility-related accidents experienced by people with visual impairment. *AER Journal: Research and Practice in Visual Impairment and Blindness*, in press.
- Manduchi, R., Kurniawan, S. and Bagherinia, H. Blind guidance using mobile computer vision: A usability study. In *Proc. ACM SIGACCESS Conference on Computers and Accessibility* (2010).
- Marston, J.R., Loomis, J.M., Klatzky, R.L., Golledge, R.G. and Smith, E.L. Evaluation of spatial displays for navigation without sight. *ACM Transactions on Applied Perception* 3, 2 (2006), 110–124.
- Mattar, M.A., Hanson, A.R. and Learned-Miller, E.G. Sign classification using local and meta-features. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2005).
- Pradeep, V., Medioni, G. and Weiland, J. Piecewise planar modeling for step detection using stereo vision. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2008).
- Pradeep, V., Medioni, G. and Weiland, J. Robot vision for the visually impaired. In *Proc. Workshop on Applications of Computer Vision for the Visually Impaired* (2010).
- Ross, D.A. and Blasch, B.B. Wearable interfaces for orientation and wayfinding. In *Proc. ACM SIGACCESS Conference on Computers and Accessibility* (2000).
- Saez, J. and Escolano, F. Stereo-based aerial obstacle detection for the visually impaired. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2008).
- Sanketi, P., Shen, H. and Coughlan, J. Localizing blurry and low-resolution text in natural images. In *Proc. IEEE Workshop on Applications of Computer Vision* (2011).
- Sudol, J., Dialameh, O., Blanchard, C. and Dorsey, T. Looktel: A comprehensive platform for computer-aided visual assistance. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2010).
- Tekin, E. and Coughlan, J. An algorithm enabling blind users to find and read barcodes. In *Proc. IEEE Workshop on Applications of Computer Vision* (2009).
- Tekin, E., Coughlan, J. and Shen, H. Real-time detection and reading of LED/LCD displays for visually impaired persons. In *Proc. IEEE Workshop on Applications of Computer Vision* (2011).
- Tjan, B.S., Beckmann, P.J., Roy, R., Giudice, N. and Legge, G.E. Digital sign system for indoor wayfinding for the visually impaired. In *Proc. Workshop on Computer Vision for the Visually Impaired* (2005).
- Walker, B.N. and Lindsay, J. Navigation performance in a virtual environment with bonephones. In *Proc. International Conference on Auditory Display* (2005).
- Wang, K. and Belongie, S. Word spotting in the wild. In *Proc. European Conference on Computer Vision* (2010).
- Weinman, J.J., Learned-Miller, E. and A. R. Hanson, A.R. Scene text recognition using similarity and a lexicon with sparse belief propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31, (Oct. 2009), 1733–1746.
- Wilson, J., Walker, B.N., Lindsay, J., Cambias, C. and Dellaert, F. SWAN: System for wearable audio navigation. In *Proc. IEEE International Symposium on Wearable Computers* (2007).
- Winlock, T., Christiansen, E. and Belongie, S. Toward real-time grocery detection for the visually impaired. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2010).
- Yang, X. and Tian, Y. Robust door detection in unfamiliar environments by combining edge and corner features. In *Proc. Workshop on Computer Vision Applications for the Visually Impaired* (2010).
- Yuan, D. and Manduchi, R. Dynamic environment exploration using a virtual white cane. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (2005).

Roberto Manduchi (manduchi@soe.ucsc.edu) is an associate professor in the Department of Computer Engineering at University of California, Santa Cruz.

James Coughlan (coughlan@ski.org) is Senior Scientist at The Smith-Kettlewell Eye Research Institute, San Francisco, CA.

research highlights

P. 106

Technical Perspective Where Do People Draw Lines?

By Frédo Durand

P. 107

Where Do People Draw Lines?

By Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher,
Heather Stoddart Barros, Adam Finkelstein,
Thomas Funkhouser, and Szymon Rusinkiewicz

P. 116

Technical Perspective Content-Centric Networking

By Jim Kurose

P. 117

Networking Named Content

By Van Jacobson, Diana K. Smetters, James D. Thornton,
Michael Plass, Nick Briggs, and Rebecca Braynard

Technical Perspective

Where Do People Draw Lines?

By Frédo Durand

ONCE UPON A time, computer graphics focused exclusively on realism and the generation of synthetic images that follow the laws of optics. The field eventually realized that other pictorial styles, such as line drawing, offer compelling depictions that can be more visually economical, better at focusing attention and abstracting out unimportant areas, and be more aesthetically pleasing. After all, a significant amount of the pictures generated by artists and designers are not photorealistic; indeed, a great many of them are line drawings. For example, educational diagrams and user manuals often use such simplified style to better eliminate superfluous detail and focus attention. The breadth of situations in which line drawings are used makes it important to develop algorithms to generate them either automatically or in a user-assisted manner.

Inspired by artists' work, a new sub-field called non-photorealistic rendering has emerged that seeks to imitate traditional media such as pencil drawing or oil painting. Whereas this new area has contributed exciting algorithms and vastly broadened the variety of visual styles that can be created with computers, the establishment of clear problem statements and evaluation metrics has proven challenging.

In contrast to photorealistic computer graphics, which can be formulated as the solution of the light propagation equation, non-photorealistic styles are elusive. Consider line drawing. What is a set of line strokes that can convey a 3D shape from a given viewpoint? This issue touches upon art and perception, but these two fields offer few answers. The effectiveness of line drawing has been demonstrated with perceptual studies, but the question of what lines should be drawn

was ignored and such studies relied on hand-drawn pictures for which the mathematical link between lines and shape cannot be easily investigated.

Computer graphics researchers have had to base line drawing algorithms on hypotheses and intuitions about the relationship between local properties of a shape and the location of appropriate lines. Before computer graphics became interested in the issue, the great mathematician David Hilbert hypothesized that differential geometry holds the answer and that lines should be drawn at so-called parabolic curves. However, once computers enabled the automatic extraction of such curves from 3D models, it became clear they do not lead to compelling drawings. New definitions of lines based on a variety of differential geometry properties were proposed and the quality of computer-generated drawings improved.

We now have many available mathematical feature-line definitions, but the crucial question remains the precise characterization of where skilled human artists place lines to convey a 3D shape. This is the issue tackled by Cole et al. in the following paper. The

The following study is exciting not only for computer graphics, but also for the understanding of human perception and art.

central challenge to relate artists' drawings and shape is registration. One need not only a drawing of a known shape, but also the precise spatial registration of stroke locations onto the shape. Tracing the drawing on top of a photorealistic rendering of the shape could solve registration but would likely bias artists. The authors address this by asking artists to first draw the shape on white paper, and then reproduce their drawing by tracing it on top of a photorealistic rendering of the shape. This method enabled the gathering of a set of precisely registered drawings and shapes.

This study is exciting not only for computer graphics, but also for the understanding of human perception and art. Line drawings are intriguing because they offer stimuli that are very different from real scenes, and that our ancestors have not encountered during their evolution. Understanding their effectiveness at conveying scenes might shed light on the strategy used by our visual system and what locations of a shape carry more visual information.

One might argue that handmade pictures are just one baseline and that the real metric should be the effectiveness at conveying shape, regardless of how similar a drawing is to that of the artist's. The authors have studied this question as well in a companion paper entitled "How Well Do Line Drawings Depict Shape?," which I highly recommend. Together, these two papers dramatically enhance our understanding of line drawing and will lead to more effective algorithms. 

Frédo Durand (fredo@mit.edu) is an associate professor in Electrical Engineering and Computer Science at the Massachusetts Institute of Technology, and a member of the Computer Science and Artificial Intelligence Laboratory.

© 2012 ACM 0001-0782/12/01 \$10.00

Where Do People Draw Lines?

By Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz

Abstract

This paper presents the results of a study in which artists made line drawings intended to convey specific 3D shapes. The study was designed so that drawings could be registered with rendered images of 3D models, supporting an analysis of how well the locations of the artists' lines correlate with other artists', with current computer graphics (CG) line definitions, and with the underlying differential properties of the 3D surface. Lines drawn by artists in this study largely overlapped one another, particularly along the occluding contours of the object. Most lines that do not overlap contours overlap large gradients of the image intensity and correlate strongly with predictions made by recent line-drawing algorithms in CG. A few were not well described by any of the local properties considered in this study. The result of our work is a publicly available data set of aligned drawings, an analysis of where lines appear in that data set based on local properties of 3D models, and algorithms to predict where artists will draw lines for new scenes.

1. INTRODUCTION

The goal of our work is to characterize the mathematical properties of line drawings made by human artists. Specifically, we aim to draw relationships between the locations of lines drawn by artists and properties of the surface geometry, lighting, and viewing conditions at those locations. This type of analysis can both guide the future development of line-drawing algorithms in computer graphics (CG) and provide artists and observers with a precise vocabulary for characterizing and discussing where lines on a model are drawn. This paper describes a study in which art students were asked to make line drawings (Figure 1) that “convey the shape” of 3D models shown to them as rendered images. The study balances the competing concerns of allowing the artists to draw freely and of acquiring useful data. The artists were asked to make two drawings on paper: first a freehand drawing, then a tracing to register their drawing to a faint, photorealistic image of the model. The registered drawings may be used to characterize the differences between specific artists, study correlations between human-drawn and computer-generated lines, and to provide training data for synthesis of new line drawings.

We provide a statistical analysis of the locations where artists drew lines with the geometric, viewpoint, and lighting characteristics of the underlying 3D scene. The analysis supports several conclusions. First, human line drawings, made under our controlled conditions, are quite consistent with each other. Second, most of the areas where artists consistently drew lines can be described by well-known, simple mathematical properties, such as the locations of occluding contours and large gradients of image intensity. Third,

current line-drawing algorithms can help explain many of the lines that do not lie in those areas but cannot explain all the artists' lines. We believe that this paper in no way exhausts the possible investigations that can be performed with this data. We therefore make our drawings and models freely available, in hopes that other researchers continue in this line of inquiry.

1.1. Background and related work

Principles of drawing. Artists have for centuries studied the principles of how to make drawings. Books codifying these principles line the shelves of any major bookstore.^{9, 16, 21} Although these texts tend to emphasize the more artistically salient concerns of composition, motion, passion, and mystery, some also offer advice on using lines to convey texture, bulk, and shading, noting that even sparse line drawings are sufficient for the viewer to identify shape. Some explicitly identify known line types as candidates for drawing (e.g., contours and ridges²³ or specific feature lines on a known shape such as the nose¹⁸). Nevertheless, little more is said about *where* on a figure to place lines in order to best convey shape—this decision making process seems to be learned through trial and error over years of practice by individual artists.

Algorithmic line drawing. Inspired by the effectiveness and aesthetic appeal of human line drawings, scientists have investigated algorithms for generating line drawings. The mathematician Felix Klein reportedly¹¹ asked his students to draw parabolic lines (lines of zero Gaussian curvature) over a bust of Apollo, believing they would expose some aspects of the aesthetics of the sculpture. More recently, researchers have explored smooth silhouettes,¹⁰ suggestive contours and highlights,^{7, 8} geometric ridges and valleys,¹⁷ and apparent ridges,¹³ (Figure 2). Each algorithm has strengths and weaknesses. Silhouette lines (more precisely, occluding contours) are obviously important, but alone they tend to leave a drawing too sparse. Suggestive contours connect to occluding contours and often lie in naturally important areas but do not exist at all on convex shapes. Geometric ridges help define convex shapes but often seem to exaggerate curvature. Apparent ridges use a more sophisticated view-dependent curvature metric but still seem to exaggerate curvature in some cases and tend to be noisy. Finally, while not explicitly designed as a line-drawing method, intensity edge detection in shaded images by algorithms like that of Canny² is so straightforward that it is common in image editing software like Photoshop. These lines are surprisingly effective for many applications but require tuning intensity thresholds and are brittle in the face of image noise.

A previous version of this paper was published in the *ACM Transactions on Graphics* 27, 3; *Proceedings of SIGGRAPH 2008*, (Aug. 2008).

Figure 1. Where people draw lines. Average images composed of 107 drawings show where artists commonly drew lines in our study.





Figure 2. CG line drawings. Recent research has proposed several ways to create a line drawing from a 3D model (a). We examined smooth silhouettes (b), suggestive contours (c), geometric ridges and valleys (d), apparent ridges (e), and image edges (f).



With the myriad line-drawing options now available, it is natural to ask which are appropriate for a specific situation, or which more closely resembles what a human would draw. Recent efforts^{8, 13, 14} include direct comparisons between their results and artists' renderings. However, these comparisons are informal and generally intended to illustrate the inspiration for the work, not to evaluate the results of the algorithm. In contrast, the data set presented here makes possible formal comparisons by directly associating 3D models, lighting, and camera angles with human drawings. **Hand-drawn data sets.** Hand-drawn or hand-annotated data sets can be useful wherever visual comparison between an algorithm and human performance is desired. In computer vision, human annotations have been used as reference data for image segmentation,¹⁵ object recognition,²² and 3D model segmentation.³ Segmentation and line drawing are related but distinct: segmentation boundaries are an obvious source of lines but generally do not provide compelling line drawings on their own.

Two recent studies directly investigated drawings by human artists of 3D models. Isenberg et al.¹² compared viewers' perceptions of hand-drawn versus computer-generated pen-and-ink illustrations. Phillips et al.¹⁹ conducted a study similar to ours, in which artists were asked to draw synthetic, blobby shapes from a range of prompt types. Among other differences from that work, our study includes

a separate tracing and registration step that allows greater accuracy in the analysis of artists' lines.

2. STUDY DESIGN

The study is designed to capture the relationships between the locations where human artists draw lines and the mathematical properties of the model's surface and appearance at those locations. To achieve this goal in a way that supports detailed analysis, several important choices must be made: what drawing style to consider; what models, views, and lighting conditions to use as prompts; how to present these prompts to the artists; what instructions to give the artists; and how to scan and process the drawings. The following sections describe each of our design decisions in detail.

2.1. Artistic style

The first challenge in designing the study is to decide on a style of drawing that is narrow enough that all artists have roughly similar intentions while drawing, yet flexible enough for each artist to exercise individual ingenuity. We balance these goals by focusing on line drawing that includes only feature lines, with no hatching or shading (examples appear in Figure 5). This choice of style was made for two reasons. First, it is a simple style that is familiar to most artists and yet expressive enough to depict shape. Second, it matches the style generated by several NPR rendering algorithms recently proposed in the CG literature^{7, 13}. By asking the artists to draw in the same style as the computer algorithms, we can learn both about the human drawings (by using the vocabulary of the algorithms) and the computer drawings (by using statistical correlations with human tendencies).

We give each artist verbal and written instructions to make drawings with "lines that convey the shape" of an object. We do not provide instructions about whether lines should represent shape features, lighting features, or anything else. However, we specifically ask the artists to refrain from including lines that represent area shading or tone features, such as stippling or hatching.

2.2. Prompt selection

A second design decision is to select 3D models and rendering parameters to use when producing prompts (images depicting a shape for the artists to draw).

Our first concern is to provide images from which the artists can easily infer shape but that are not so familiar that they apply domain-specific (“idiomatic”) knowledge when drawing. This consideration not only rules out overly abstract or complicated 3D surfaces (i.e., shapes unlike anything in common experience) but also rules out objects with strong semantic features (e.g., human faces) and ones commonly drawn in art classes (e.g., fruit). It also suggests that multiple views of the shape be provided as prompts so that ambiguities in one view are resolved by another. Finally, prompt images should be photorealistic to avoid confusing artists that are not familiar with classic CG rendering artifacts such as hard shadows and lack of indirect illumination.

For productive analysis, the set of prompts should include pixels with a wide variety of mathematical properties (e.g., high image gradients, surface critical points, etc.), and these features should be separated spatially to allow clear distinctions to be made. This consideration rules out objects containing only large, planar facets (few interesting surface features), convex objects (no concave surface features), and other surfaces with few inflections. Rather, it suggests blobby objects with many curved surfaces.

Finally, the objects must be relatively simple, without much fine scale detail. Otherwise, the artists may be tempted to abstract or simply omit important features.

Based on these criteria, we select 12 models of 4 object types for our study: (a) four bones, (b) two tablecloths, (c) four mechanical parts, and (d) two synthetic shapes (Figure 3). We synthesize four prompt images for each model, one for each combination of two different viewpoints and two lighting conditions. The two viewpoints are always 30° apart (so that large parts of each model can be seen from both viewpoints) and are carefully chosen to distribute surface features across the image. By providing prompts with different lighting and different viewpoints for the same model, we can analyze image-space properties in isolation from object-space ones.


We generate our images using YafRay,²⁵ a free raytracing package capable of global illumination using Monte Carlo pathtracing. The models are rendered using a fully diffuse, gray material, and thus take on the color of the lighting environment. For lighting, we use the Eucalyptus Grove and Grace Cathedral high-dynamic-range environment maps captured by Debevec.⁶

2.3. Line drawing registration

The final and most difficult part of the study design is to engineer a system that is able to register line drawings made by artists to pixels of a prompt image with great accuracy.

Designing such a system is challenging because there is a trade-off between allowing the artist to draw in a natural manner (e.g., with pencil on a blank sheet of paper) versus including constraints that facilitate accurate registration between prompts and line drawings. On one hand, the drawing process surely must not bias the locations of lines made by the artist, and thus it is not a good idea to have the artist *compose* a drawing directly over the image prompt. On the other hand, the process must provide enough registration accuracy to distinguish between important mathematical properties at nearby pixels in the prompt. This problem is

Figure 3. 3D models. The 12 models from our study, each shown in one of two views and one of two lighting conditions. Groups (a and b) are scans of real objects; (c and d) are synthetic.




particularly difficult since freehand drawings can be geometrically imprecise, and the intended location of every line is only known by the artist.

Our design balances these trade-offs with a simple two-step process. The artist is given two sheets of $8.5'' \times 11''$ paper for each line drawing (Figure 4). The *prompt page* (shown on the left) contains multiple full color views of the prompt shape, one of which is large ($6.5'' \times 4.75''$) and is called the *main view*. The *drawing page* (shown on the right) contains two boxes, each the same size as the main view. The top box is initially blank, while the bottom box contains a faint version of the main view.

The artist is asked to complete the drawing page by first folding the page vertically in half so that only the blank space at the top is visible (Figure 4, left). Using the viewing page for reference, the artist draws the prompt shape in the blank space, just as if they were making a normal sketch. When finished, the artist unfolds the drawing page and copies their freehand drawing onto the faint image on the bottom of the same page. During the copying step, the artist is asked to change the shape of their lines to match the target rendering but not to change the number or relative position of the lines. In effect, the artist is asked to perform a non-linear warp of their original drawing onto the target shape. A typical result is shown on the right side of Figure 4.

We scan the drawing page with a flatbed scanner, locate fiducials included in the corners of the page, and then use the fiducials to register the traced lines with the 3D model rendered from the main viewpoint. An adaptive thresholding method is used to convert the scanned gray-scale image into a binary image so that all the artist’s lines, regardless of strength, are included in the binary image. We then use a thinning operator to narrow the lines in the binary image down to the width of one pixel. The final result is a

Figure 4. Making a drawing. With the drawing page folded in half, the artist makes a freehand drawing while referring to the prompt page (left). The completed drawing page (right) contains a freehand drawing and a registered drawing.



1024 × 768 pixel binary image containing a single pixel wide approximation of the human artist's lines. While this procedure takes up to twice as long as a single drawing (e.g., it requires the artist to draw every line twice), it achieves a nice balance between the design trade-offs: the line drawings are composed in a freehand manner familiar to artists, while the intended locations of every line on the 3D surface can be inferred with great accuracy.

2.4. Data collection


This line drawing and registration procedure was repeated for 29 artists, most of them were enrolled in one of four art classes (two composed of middle and high school students, one of adult evening students, and another of college students). Two of the participants were professional artists. Each artist completed up to 12 prompts. Every participant completed a questionnaire listing his/her gender, age, and number of years of art training. In all, there were 22 females and 7 males. The ages ranged from 10 to 54 years, with an average of 22; and the participants reported an average of 6 years of art training (this number should be taken with a grain of salt, as some participants reported only training at the college level, while others reported all art classes).

Every artist was provided a folder with 1 page of instructions, 12 prompt pages, and 12 corresponding drawing pages (one for each model). The folders were arranged such that no artist could draw the same model more than once, and prompts for models, viewpoints, and lighting conditions were arranged in shuffled order to reduce effects of training on our analysis.

The artists were given brief verbal instructions (“draw lines that convey shape” and “be sure to copy every line from your freehand drawing over the faded image below”) and then told to complete line drawings at their own pace for as long as they had time. Most of the art classes were scheduled for a 2-h block, and each line drawing took 10–15 min, on average (with time split around two-thirds for drawing freehand and one-third for tracing lines over the faded image). Each participant completed an average of 7.5 drawings—only one participant (a professional artist) completed all 12 available in his folder.

In all, 208 line-drawing images were collected. Generally speaking, the artists followed the directions well, produced

Figure 5. Example drawings. Three drawings of the screwdriver model from the same view (a–c), and the average of 14 drawings of the same view (d).



line drawings that convey shape effectively, and were careful when tracing lines over the faded image (some example line drawings are shown in Figure 5). However, in some cases, the artists clearly were not careful in the registration step, failing to follow even the exterior outline of the shape. Since accurate registration of lines to image features is essential for meaningful results, we cull these tracings from our analysis. To do this in an unbiased way, we assume that inclusion of the exterior outline is common to all human line drawings and eliminate from our data set any drawings where less than 90% of the exterior is within 1 mm of a human-drawn line. The remaining 170 line drawings form the basis for our analysis.

3. RESULTS


We can investigate a number of questions by comparing the captured line drawings and the synthetic images provided as prompts to the artists. We ask not only how artists' drawings overlap with one another but also how they overlap with lines generated by CG algorithms and how they can be predicted from local properties of the underlying surface and rendered image.

3.1. How similar are the artists' drawings?

The first and most basic analysis we perform is to measure the similarity between artists' drawings of the same prompts. We can show consistency between artists qualitatively by superposing drawings on top of each other and visualizing how much they overlap. For example, Figure 6a shows each artist's drawing in a separate color. In this example, the artists agree very closely with each other in most areas, especially along obvious features such as boundaries and occluding contours, but differ in exactly where they place lines in the right part of the rockerarm.

In order to quantify consistency, we compute a histogram of pairwise distances between artists' drawings (Figure 6b). For every pixel in every drawing, we record the distance to the closest pixel in every other drawing of the same prompt and then observe how often these distances lie within the tolerance of the tracing procedure (1 mm). Across all prompts, approximately 75% of human drawing

Figure 6. Consistency of artists' lines. (a) Five superimposed drawings by different artists, each in a different color, showing that artists' lines tend to lie near each other. (b) A histogram of pairwise closest distances between pixels for all 48 prompts. Approximately 75% of the distances are less than 1 mm.



pixels are within 1 mm of a drawn pixel in all other drawings for that prompt.


3.2. Do known CG lines describe artists' lines?

A natural question to ask is how well currently known line-drawing algorithms can describe the human artists' lines. In our analysis, we consider the following line-drawing algorithms: image intensity edges,² geometric ridges and valleys (as defined by Ohtake et al.¹⁷), suggestive contours,⁷ and apparent ridges.¹³ For the object space methods (ridges and valleys, suggestive contours, and apparent ridges), we always include the exterior boundary and interior occluding contours in the generated drawing. For Canny edge detection, we always include the exterior but not the interior contours, since they are not necessarily image intensity edges.

Quantifying comparisons between drawings. In order to compare an artist's drawing and a computer-generated drawing quantitatively, we use the standard information retrieval statistics of *precision* and *recall* (PR). Here, precision is defined as the fraction of pixels in the CG drawing that are near any pixel of the human drawing. Recall is defined as the fraction of pixels in the human drawing that are near any line of the CG drawing. We define "near" by choosing a distance threshold—we use 1 mm.

As an example, consider comparing the set of five human drawings shown in Figure 6a with the lines generated by the apparent ridges algorithm (Figure 7). The output of the apparent ridges algorithm is not only a set of lines but also a "strength" value at each line point. In general, we expect stronger lines to be more important and thus more likely to match the artists' lines. We thus generate a series of binary apparent ridges images, each consisting of all points with strength above a given threshold. The PR of each drawing compared with this set of images is shown as a dotted pink line in Figure 7. As the strength

Figure 7. Precision and recall example. Left: apparent ridges are compared with five artist drawings. Solid line (highlighted) is the average PR for the set of drawings. Black dots indicate occluding contours only. Right: an example drawing (black) with apparent ridges overlaid (red, widened by 1 mm on each side). The PR for this example is circled (80% P, 88% R).



threshold is lowered more lines are produced, typically causing recall to increase and precision to go down, yielding a sloping line in the PR graph. For completeness, we allow the PR plot to extend to $P = 1.0, R = 0.0$ (defined as a blank image) and directly downward to $P = 0.0$ from the highest recall obtained by the algorithm. Since each PR curve is defined for $P = [0, 1]$, we can compute an average curve by combining points along lines of fixed precision. The PR values for occluding contours alone are plotted as black dots and are not averaged.


Comparing CG lines in combination. In order to combine line definitions fairly, we use computer-generated drawings with a fixed 80% precision. We then classify each pixel in each human drawing by the nearby CG lines. Pixels that lie near a single-line definition are considered to be explained only by that definition, while pixels that lie near multiple definitions are considered explained by all the nearby definitions.

To visualize the results, we create bar charts that partition the lines into object space definitions (blue), image intensity edges (green), or both (brown). Looking at the results in Figure 8a, we find that the large majority of lines are described by both image intensity edges (Canny edges) and an object space definition. Of the remainder, slightly more lines are explained by the combined object space approaches than by image edges alone.

Lines that are explained only by image edges account for at most 5% of all classified lines at 80% precision. We can also break down the human lines by intuitive categories, such as exterior and interior occluding contours and everything else (Figure 8b). Across all model groups, exterior contours alone account for between 35% and 50% of all classified pixels. Interior occluding contours account for between 10% and 20% of all classified pixels, while all other definitions make up 20%–35%.

Can CG lines characterize artists' tendencies? Given a way of describing an artist's drawing in terms of CG line types, it is possible to investigate whether those descriptions can characterize the similarities and difference between artists' styles or tendencies. For example, it may be possible to characterize whether certain artists tend to draw certain geometric features (e.g., ridges) more than other artists do. In such cases, the CG line definitions provide a vocabulary to discuss

Figure 8. Categorizing artists' lines. (a) Fraction of all lines explained by image based lines only, object based lines only, and both. (b) Fraction of all lines explained by the exterior contours, interior occluding contours, and all other object space lines.



features of human line drawings.


Figure 9 shows a simple example of this type. Two drawings of the same prompt (twoboxcloth with Grace Cathedral lighting) are compared by the composition of CG line types. The colored bars indicate the fraction of the drawing made up by each line type. In this case, however, each set of bars represents a single drawing. One immediate difference between the drawings is that artist A drew more lines besides the contours. Non-contour lines account for 26% of artist A's drawing and only 13% of artist B's drawing. The bulk of the difference between the artists is in the use of ridge-like lines (green, yellow, and pink bars). Artist A drew ridge-like lines along the top of the shape, while artist B did not. This visual difference is evident from the statistics, which show a large fraction of geometric ridges and apparent ridges in artist A's drawing and almost none in artist B's drawing.

3.3. Can local properties explain lines?

While it is interesting to investigate the relationship between artists' lines and the lines commonly used in CG, a more fundamental question is how artists' lines relate to differential properties of images and surfaces. The analysis above addresses this question indirectly, since each CG definition is based on a set of local properties, but it is restricted to the relationships suggested by the known line-drawing algorithms. To address this question, we take a classic data mining approach. For every pixel of every prompt, we compute: (1) a feature vector x of properties derived from the 3D surface and 2D rendered image and (2) an estimated probability that a line will be included at the corresponding location in an artist's line drawing. Our goals are to learn a function $f(x)$ that estimates the probability p of an artist drawing at a point (regression) and to understand which combinations of properties are most useful for building such a function (feature importance).

Choosing local properties. To build the feature vector for each pixel, we compute 15 local properties of three types commonly used in image processing, CG, and differential geometry. First, we consider four image-space properties of the rendered image prompt: luminance, gradient magnitude after $\sigma = 2$ pixels Gaussian blur (ImgGradMag), and minimum and maximum eigenvalues of the image Hessian (corresponding to extrema in second derivative of luminance (ImgMinCurv and ImgMaxCurv, respectively).

Figure 9. Comparison of two drawings by different artists. Two drawings of the same prompt show significant visual differences. These differences are reflected in the statistics, especially in the use of ridge-like lines (green). RV: ridges and valleys, AR: apparent ridges, SC: suggestive contours (see Figure 2).



In general, we expect that lines are more likely near image edges (ImgGradMag is large) and at ridges and valleys of luminance (where ImgMinCurv and ImgMaxCurv are large).

Second, we consider view-independent, differential properties of the visible point on the 3D surface, including the maximum (κ_1), minimum (κ_2), mean ($(\kappa_1 + \kappa_2)/2$), and Gaussian ($\kappa_1 \kappa_2$) curvatures (SurfMaxCurv, SurfMinCurv, SurfMeanCurv, and SurfGaussianCurv, respectively). In most cases, we expect lines to occur in areas where these expressions are large, though it has also been observed that lines are drawn near parabolic lines ($\kappa_1 \kappa_2 = 0$).

Third, we consider view-dependent properties that correspond to specific definitions for computer-generated lines. Corresponding to the definition of ridges and valleys, we take the derivative of the largest principal curvature in the corresponding principal direction (SurfMaxCurvDeriv), which is zero at ridges and valleys. Corresponding to occluding contours, we compute the dot product between normal and view vectors ($N \cdot V$). Corresponding to apparent ridges and valleys, we compute the largest view-dependent principal curvature (ViewDepCurv) and its derivative in the corresponding apparent principal direction (ViewDepCurvDeriv), which are large and zero, respectively, at apparent ridges. Corresponding to suggestive contours, we compute the radial curvature (RadialCurv) and its derivative in the radial direction (RadialCurvDeriv), which are zero and large, respectively, at suggestive contours. Finally, corresponding to principal highlights, we compute radial torsion, which is zero at principal highlights.

Finally, we estimate the probability, p , of an artist drawing at a pixel by averaging the registered drawings of all artists for the same prompt and blurring with a Gaussian filter to account for tracing errors ($\sigma = 0.5$ mm).

Predicting lines by regression. Several of the computed properties clearly can be used to distinguish pixels where artists draw from where they do not (Figure 10). However,

an interesting question is whether combinations of those properties can be used to predict where artists will draw more accurately than any of them alone. To investigate this question, we have experimented with several regression models, including linear regression, radial basis functions, regression trees, and several others. As an example, Figure 11a shows a regression tree built with the M5P package in Weka²⁴ to predict the set of line drawings for one view of the twoboxcloth model shown in Figure 3. Figure 11a shows the prediction of p resulting from this simple tree, while Figure 11b provides a visualization of which pixels sort into which leaves of the tree (pixels in the image are colored to match the text of the leaf).

In this example, several properties are combined by the decision tree to predict p , starting with ImgGradMag at the root. The set of properties chosen is instructive, as it suggests that they provide the highest incremental value in predicting p (at the start of tree building). Of course, many properties are correlated, and the decision tree may be non-optimal, so an alternative tree may have produced similar or better predictions. Nevertheless, it is interesting to see how non-trivial combinations of local properties can be used to make predictions—even though the tree was purposely kept small in this example, it still is able to provide a plausible (albeit coarse) prediction for where artists draw lines (Figure 11a). If we consider deeper trees or other regression models, we are able to predict p from x more accurately.

3.4. Which local properties are most important?

In our data mining framework, it is not only possible to predict where artists will draw but also to examine which local features are most important when building such a regression model. For example, Random Forests¹ estimate the importance of every feature to its model by building a large number of decision trees trained on different subsets of the data.¹ For each feature m of each built tree, the error observed in predictions for the “out of bag” data (the part held out of training) is computed and compared to the error that is observed when values of feature m are permuted. The difference between these errors, averaged and normalized, is reported as the “importance” of feature m . For this analysis, we make the assumption that almost all occluding contours ($N \cdot V = 0$) are drawn by artists (Figure 10), and so exclude any pixel within 1 mm of a contour from the training set.

Table 1 shows the relative feature importance as computed with the Random Forest implementation of Breiman and Cutler in R for the remaining pixels of all drawings in our study.²⁰ The first four columns report the importance of features (rows) estimated when training on models from each of one type (bones, cloth, mechanical, and synthetic), while the rightmost column reports the average over the whole data set.

The results indicate that image-space intensity gradient magnitude is the feature among the tested set that is most useful in predicting the probability that an artist will draw at a particular location in our study (e.g., the average prediction error is largest if values of the image-space gradient magnitude are randomized). While image-space discontinuities often appear at the same place as boundary

Figure 10. Example local surface features. Top: the frequencies of pixels near artists’ lines (blue), and away from artists’ lines (green, dashed), as functions of local surface properties. Bottom: pixels near artists’ lines as a fraction of the total. Pixels where $N \cdot V \approx 0$ or where the Sobel response is high are very likely near a line.





Figure 11. Decision tree for predicting where artists will draw.

(Left) decision tree learned from prompts of bones, (a) predicted probabilities of where artists will draw for this view (black is high probability), (b) a visualization of which pixels fall into which leaves of the tree. Note that this tree was purposely kept small for didactic purposes, yielding coarse prediction.



contours and occluding contours ($N \cdot V = 0$), the locations where those contours appear have been excluded from this study. So, this result suggests that image-space intensity gradients away from the contours are also highly correlated with artist line locations. Of course, this is not surprising, as ridges, valleys, and shadow boundaries are commonly drawn by artists. However, it is a bit surprising how all the simple image-space features (which do not require a 3D model to compute) are so important relative to the other more complex properties that have been the focus of recent research in CG.

3.5. Which CG lines are most important?

We use Random Forests to compute importance of the CG line definitions studied in Section 3.2 for predicting where artists draw lines. For this analysis, we compute a new feature vector for every pixel storing the strength for every CG line definition. Note that strength is only defined at pixels where the algorithm would draw a line (e.g., zeros of maximum curvature derivative for ridges). At all other pixels, strength is always zero. We then recompute the Random Forests with the new feature vectors.

Table 1. Local property “importance.”

	Feature	Bone	Cloth	Mech	Synth	Avg
Image space	ImgGradMag	31.3	36.0	73.8	147.8	72.2
	ImgMaxCurv	38.0	15.8	55.5	64.4	43.4
	ImgMinCurv	15.1	15.3	23.4	56.6	27.6
	ImgLuminance	20.2	19.8	33.9	33.6	26.9
View dependent	N · V	23.6	13.9	31.3	36.9	26.4
	ViewDepCurv	21.5	17.2	49.8	10.1	24.7
	ViewDepCurvDeriv	22.8	14.4	31.9	9.5	19.7
	RadialCurvDeriv	19.2	15.0	29.8	8.0	18.0
	RadialTorsion	14.6	10.3	27.8	7.2	15.0
View independent	RadialCurv	14.8	10.3	26.2	7.2	14.6
	SurfMaxCurvDeriv	16.9	11.0	27.3	8.9	16.0
	SurfMaxCurv	13.9	8.8	25.1	7.6	13.9
	SurfMinCurv	13.9	8.1	27.0	5.1	13.5
	SurfMeanCurv	14.1	8.9	22.5	7.0	13.1
	SurfGaussianCurv	13.1	8.5	25.7	4.9	13.1

The relative importance of local properties in predicting the probability of an artist drawing at a particular location. Columns show results for Random Forests trained on different subsets of the data.

Table 2 shows that strong image-space gradients in illumination (Canny edges) still provide the strongest cues for artists to draw lines, even in relation to other CG line definitions.

4. CONCLUSION

This study allows us to make several quantitative conclusions about how people draw 3D shapes. First, we observe that artists are consistent. In our study, artists drew nearly 75% of their lines at a location that is within 1 mm of *all* other artists drawing from the same prompt. These overlaps appear mainly at occluding contours, which comprise 57% of all lines drawn. Among the other lines, large gradients in image intensity (as measured by image-space gradient magnitude) provide the best single predictor for where artists will draw under the conditions of our study. Lines generated by Canny edge detection on a prompt image cover 76% of artists’ lines with 80% precision. These lines are almost entirely overlapped (95%) by lines predicted by object-space line definitions commonly found in CG. The three object space definitions together cover 81% of the artists’ lines at the same precision. We find that each of the four CG line definitions explains some artists’ lines that the others do not.

The cumulative output of the four line-drawing algorithms considered in this paper cover only 86% of artists’ lines at 80% precision. We believe that some of the remaining lines could be explained by new line definitions based on local properties and that clever combinations of properties could improve precision. Our data-mining analysis provides one way to find new definitions. We made an initial investigation of data-driven line drawing in our earlier article.⁴

We suspect, however, that a more fruitful direction for future research will be to investigate the *global* decisions artists make, rather than the local ones. For example, Figure 12 shows a case where artists chose to draw lines on locally weak ridge and valley features, while omitting lines along locally stronger ridges and valleys. This choice is consistent between several artists. When creating the

Table 2. CG line definition “importance.”

Feature	Bone	Cloth	Mech	Synth	Avg
Canny edges	18.2	37.2	50.9	145.0	53.4
Apparent Ridges	8.7	11.2	21.2	77.9	24.8
Ridges and valleys	6.8	7.4	24.4	77.1	24.5
Suggestive contours	9.8	11.9	17.4	1.6	11.3

Results akin to those in Table 1, but for features derived from CG line definitions.

flange drawing, several artists mentioned that they omitted lines that were “implied.” Since implied features depend on context, they are not describable with local properties alone. Our earlier article and associated data set^a show more examples of these effects.

Other directions. This study suggests several directions for further work, some of which we have already begun to investigate.

One question is the perceptual effectiveness of sparse line drawings. We placed heavy restrictions (feature lines only, no hatching or shading) on the methods the artists were allowed to use, and some artists complained that their resulting drawings did not effectively represent the subject shape. To investigate this issue, we performed a perceptual study of how well sparse line drawings depict 3D shape, using the same set of models and subset of the drawings we collected here.⁵ The study subjects were asked to estimate the surface normal at many points on the drawings. We found that people interpret certain shapes almost as well from a line drawing as from a shaded image and that there was little difference between the best CG line-drawing techniques and the artists’ drawings. We confirmed, however, that some shapes (e.g., the cervical and vertebra models) are difficult to effectively depict with sparse lines, even for a skilled artist.


The results of the perceptual study provide additional motivation for expanding the range of styles we study to include shading or even artistic stylization. It would be particularly interesting to understand the middle ground between a black-and-white line drawing and a fully realistic shaded image. Perhaps there is a slightly enriched visual style that can depict even difficult shapes effectively, without requiring the full, photorealistic representation. The methodology of collecting artists’ work and applying perceptual evaluation should provide the necessary tools to investigate these questions. We believe future results in this area will be of interest not only to CG researchers but also to researchers studying machine and biological vision.

Acknowledgments

First and foremost, we would like to thank the artists that participated in our study, including those from Art Collaborations and Princeton University. Many people provided helpful suggestions, feedback, and encouragement for this project, particularly Doug DeCarlo, Maneesh Agrawala, Fredo Durand, Eitan Grinspun, and the members of the Princeton Graphics Group. We also thank the NSF (CNS-0406415, IIS-0612231, CCF-0702672, CCF-0702672,

^a <http://www.cs.princeton.edu/gfx/proj/LD3D/>

Figure 12. Subtle line selection decisions by artists. The red lines (solid boxes) in this composite (a) are unexplained at 80% precision, but can be characterized as geometric valleys. However, the artists have omitted locally stronger valley features (dotted boxes), as shown by the maximum curvature of the model (b).



IIS-0511965, and CCF-0347427) and Google for funding to support this project and the Aim@Shape, VAKHUN, and Cyberware repositories for the models used in the study. ■

References

1. Breiman, L. Random forests. *Mach. Learn.* 45, 1 (2001), 5–32.
2. Canny, J. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 8, 6 (1986), 679–698.
3. Chen, X., Golovinskiy, A., Funkhouser, T. A benchmark for 3d mesh segmentation. In *ACM SIGGRAPH 2009 Papers* (New York, 2009), 1–12.
4. Cole, F., Golovinskiy, A., Limpaecher, A., Barros, H. S., Finkelstein, A., Funkhouser,

- T., Rusinkiewicz, S. Where do people draw lines? *ACM Trans. Graph. (Proceedings of SIGGRAPH)* 27, 3 (Aug. 2008).
5. Cole, F., Sanik, K., DeCarlo, D., Finkelstein, A., Funkhouser, T., Rusinkiewicz, S., Singh, M. How well do line drawings depict shape? *ACM Trans. Graph. (Proceedings of SIGGRAPH)* 28 (Aug. 2009).
6. Debevec, P. Rendering synthetic objects into real scenes. In *Proceedings of SIGGRAPH 1998* (Orlando, FL, 1998), 189–198.
7. DeCarlo, D., Finkelstein, A., Rusinkiewicz, S., Santella, A. Suggestive contours for conveying shape. *ACM Trans. Graph.* 22, 3 (2003), 848–855.
8. DeCarlo, D., Rusinkiewicz, S. Highlight lines for conveying shape. In *Proceedings of NPAR 2007* (Aug. 2007).
9. Guptill, A.L. *Rendering in Pen and Ink*. Watson-Guptill Publications, New York, 1976.
10. Hertzmann, A., Zorin, D. Illustrating smooth surfaces. In *Proceedings of SIGGRAPH 2000* (New Orleans, LA, 2000), 517–526.
11. Hilbert, D., Cohn-Vossen, S. *Geometry and the Imagination*. American Mathematical Society, 1999.
12. Isenberg, T., Neumann, P., Carpendale, S., Sousa, M.C., Jorge, J.A. Non-photorealistic rendering in context: an observational study. In *Proceedings of NPAR 2006* (2006), 115–126.
13. Judd, T., Durand, F., Adelson, E.H. Apparent ridges for line drawing. *ACM Trans. Graph.* 26, 3 (2007), 19.
14. Lee, Y., Markosian, L., Lee, S., Hughes, J.F. Line drawings via abstracted shading. *ACM Trans. Graph.* 26, 3 (2007), 18.
15. Martin, D., Fowlkes, C., Tal, D., Malik, J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings of the Eighth IEEE International Conference on Computer Vision, 2001. ICCV 2001* (2001), 416–423, Vol. 2.
16. Meyer, S.E., Avilez, M. *How to Draw in Pen and Ink*. Roundtable Press, New York, 1985.
17. Ohtake, Y., Belyaev, A., Seidel, H.-P. Ridge-valley lines on meshes via implicit surface fitting. *ACM Trans. Graph.* 23, 3 (2004).
18. Peck, S.R. *Atlas of Human Anatomy for the Artist*. Oxford University Press, London, U.K., 1982.
19. Phillips, F., Casella, M.W., Gaudino, B.M. What can drawing tell us about our mental representation of shape? *J. Vis.* 9, 5 (Sept. 23, 2005), Article 52.
20. R. R. *A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2005.
21. Ruskin, J. *The Elements of Drawing*. Elipson Classics, 1895.
22. Russell, B., Torralba, A., Murphy, K., Freeman, W. Labelme: a database and web-based tool for image annotation. *Int. J. Comp. Vis.* 77, 1 (2008), 157–173.
23. Smith, S. Looking at line. In *Complete Guide to Drawing and Painting*. R. D. Editors, ed., Readers Digest, 1997, 13.
24. Witten, I.H., Frank, E. *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd edition. 2005.
25. YafRay. *YafRay 0.0.9: Yet Another Free Raytracer*. www.yaf-ray.org, 2008.

Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz, Princeton University, Princeton, NJ.

© 2012 ACM 0001-0782/12/01 \$10.00



Association for
Computing Machinery

Advancing Computing as a Science & Profession



You've come a long way.
Share what you've learned.



ACM has partnered with MentorNet, the award-winning nonprofit e-mentoring network in engineering, science and mathematics. MentorNet's award-winning **One-on-One Mentoring Programs** pair ACM student members with mentors from industry, government, higher education, and other sectors.

- Communicate by email about career goals, course work, and many other topics.
- Spend just **20 minutes a week** - and make a huge difference in a student's life.
- Take part in a lively online community of professionals and students all over the world.



Make a difference to a student in your field.
Sign up today at: www.mentornet.net
Find out more at: www.acm.org/mentornet

MentorNet's sponsors include 3M Foundation, ACM, Alcoa Foundation, Agilent Technologies, Amylin Pharmaceuticals, Bechtel Group Foundation, Cisco Systems, Hewlett-Packard Company, IBM Corporation, Intel Foundation, Lockheed Martin Space Systems, National Science Foundation, Naval Research Laboratory, NVIDIA, Sandia National Laboratories, Schlumberger, S.D. Bechtel, Jr. Foundation, Texas Instruments, and The Henry Luce Foundation.

Technical Perspective

Content-Centric Networking

By Jim Kurose

MUCH HAS CHANGED in the 50 years since the invention of packet switching and the early network designs and deployments that would evolve into today's Internet. Network size has grown from a number of nodes that could be counted on our hands to a number estimated to be upward of 800 million network-attached devices and an equally large number of intermittently connected mobile devices. We have also witnessed the transition to commercial Internet use in the 1990s; the creation of the World Wide Web, P2P networks, Voice-over-IP, and streaming media services; and the increasing importance of security and privacy.

And yet, amid this change, the Internet's architecture and a number of the key protocols that form its technical foundation were conceived and subsequently standardized more than 30 years ago. Certainly, there have been changes and improvements in the intervening years, but it's remarkable indeed that the designs that began with Cerf and Kahn's original 1974 landmark paper¹ have proven so durable in the face of so much change—a tribute to the prescience and wisdom of the Internet's original designers.

While early computer networks focused on “delivering data between computers or between computers and terminals,”¹ today's Internet is arguably more concerned with connecting people with content and information. Cisco's Visual Networking Index noted that P2P and video traffic together account for nearly 80% of consumer Internet traffic in 2011. In such a content-centric worldview, *what* a person wants, rather than *where* it is located, is what matters most; content, rather than the server on which content resides, becomes the starting point. A number of the current networked applications, including content distribution networks (CDNs), publication-subscribe (pub/sub) systems, and P2P file-sharing applications, are already content-focused. But they do so as applications running on top of the Inter-

net and (importantly) there are many separate CDNs, pub/sub systems, and P2P file-sharing services.

With an increased focus on content and with nearly 50 years of networking research to draw on, the time may be right for a new network architecture—one in which content distribution and retrieval, rather than host-to-host packet delivery, is the core function supported in each and every network node. This new architecture might literally be built upon the existing Internet, in much the same way the original Internet layered itself on top of (that is, used the services of) the existing network at that time—the telephone network. At the same time, more recent considerations initially missing from early Internet designs,² such as mobility (of nodes and of content location) and security (including content protection and trust), can be more holistically incorporated.

The following paper by Jacobson et al. presents a bold vision for such a content-centric network (CCN), presenting a design philosophy, architecture, and number of specific mechanisms. CCN requests for content indicate the name of content of interest (but not a location); CCN nodes are then collectively responsible for finding that named content and returning it to the requestor. One can discern lessons well-learned from the Internet in CCN's design, including the adoption of a best effort service model—allowing a content request/response to fail occasionally (perhaps due to node failure or mobility-related topology changes) and performing recovery at the end systems—and using hierarchical aggregatable names and longest prefix matching to scalably route requests to named content.

One can also discern the influence of more recent architectural innovations in the networking research community, including an emphasis on control by the receiver (who is, after all, the ultimate consumer of the content and the one therefore best positioned

to make a decision) rather than sender control, and dynamic multipath transfers. And, of course, one sees innovations resulting from named content items being the unit of data transfer—in-network caching, breadcrumbs that plumb a reverse data path, inherent balancing of local node-to-node request/response flows, and more. We also see content security (including integrity, access control, trust, and provenance), network security (denial-of-service) and mobility all addressed in this seminal paper.

The architectures of large-scale networks (whether for communication, electrical, or transportation) seem to evolve in a style of punctuated equilibrium—slow evolutionary changes between short periods of profound change. The telephone network evolved from human control to wired logic control to stored-program control of electronic switches to today's IP-based control networks. Given the long lifetime of today's Internet architecture, the intense interest in future Internet architecture in the U.S., the E.U. and Asia (including various proposals for content- and information-centric networking), we may well be approaching a point of rapid change in communication network architecture. The following paper offers a perspective on the past as well as glimpse of what the future may hold. 

References

1. Cerf, V. and Kahn, R. A protocol for packet network interconnection. *IEEE Transactions on Communications Technology* 22, 5, 627–641.
2. Clark, D. The design philosophy of the DARPA Internet protocols. In *Proceedings of 1988 ACM SIGCOMM* (Stanford, CA, Aug. 1988), 102–111.

Jim Kurose (kurose@cs.umass.edu) is a Distinguished University Professor (and past chairman) in the Department of Computer Science and Executive Associate Dean of the College of Natural Sciences at the University of Massachusetts, Amherst, MA.

Networking Named Content

By Van Jacobson, Diana K. Smetters, James D. Thornton, Michael Plass, Nick Briggs, and Rebecca Braynard

Current network use is dominated by content distribution and retrieval yet current networking protocols are designed for conversations between hosts. Accessing content and services requires mapping from the what that users care about to the network's where. We present Content-Centric Networking (CCN) which uses content chunks as a primitive—decoupling location from identity, security and access, and retrieving chunks of content by name. Using new approaches to routing named content, derived from IP, CCN simultaneously achieves scalability, security, and performance. We describe our implementation of the architecture's basic features and demonstrate its performance and resilience with secure file downloads and VoIP calls.

1. INTRODUCTION

The engineering principles and architecture of today's Internet were created in the 1960s and 1970s. The problem networking aimed to solve was *resource sharing*—remotely using scarce and expensive devices like card readers or high-speed tape drives or even supercomputers. The communication model that resulted is a conversation between exactly two machines, one wishing to use the resource and one providing access to it. Thus IP packets contain two identifiers (addresses), one for the source and one for the destination host, and almost all the traffic on the Internet consists of TCP conversations between pairs of hosts.

In the 50 years since the creation of packet networking, computers and their attachments have become cheap, ubiquitous commodities. The connectivity offered by the Internet and low storage costs enable access to a staggering amount of new content—500EB were created in 2008 alone.⁷ People value the Internet for *what* content it contains, but communication is still in terms of *where*.

We see a number of issues that affect users arising from this incompatibility between models.

- **Availability:** Fast, reliable content access requires awkward, pre-planned, application-specific mechanisms like CDNs and P2P networks, and/or imposes excessive bandwidth costs.
- **Security:** Trust in content is easily misplaced, relying on untrustworthy location and connection information.
- **Location-dependence:** Mapping content to host locations complicates configuration as well as implementation of network services.


The direct, unified way to solve these problems is to replace *where* with *what*. Host-to-host conversations are a networking abstraction chosen to fit the problems of the 1960s. We argue that *named data* is a better abstraction for today's communication problems than *named hosts*.

We introduce *Content-Centric Networking* (CCN), a communications architecture built on named data. CCN has no notion of host at its lowest level—a packet “address” names content, not location. However, we preserve the design decisions that make TCP/IP simple, robust, and scalable.

Figure 1 compares the IP and CCN protocol stacks. Most layers of the stack reflect bilateral agreements; e.g., a layer 2 framing protocol is an agreement between the two ends of a physical link and a layer 4 transport protocol is an agreement between some producer and consumer. The only layer that requires universal agreement is layer 3, the network layer. Much of IP's success is due to the simplicity of its network layer (the IP packet—the thin “waist” of the stack) and the weak demands it makes on layer 2, namely: stateless, unreliable, unordered, best-effort delivery. CCN's network layer (Section 3) is similar to IP's and makes fewer demands on layer 2, giving it many of the same attractive properties. Additionally, CCN can be layered over anything, including IP itself.

CCN departs from IP in a number of critical ways. Two of these, *strategy* and *security*, are shown as new layers in its protocol stack. CCN can take maximum advantage of multiple simultaneous connectivities (e.g., ethernet and 3G and bluetooth and 802.11) due to its simpler relationship with layer 2. The *strategy* layer (Section 3.3) makes the fine-grained, dynamic optimization choices needed to best exploit multiple connectivities under changing conditions. CCN secures content itself (Section 4), rather than the connections over which it travels, thereby avoiding many of the host-based vulnerabilities that plague IP networking.

Figure 1. CCN moves the universal component of the network stack from IP packets to named content chunks.



A previous version of this paper was published in *Proceedings of ACM's CoNEXT Conference 2009* (Rome, Italy, Dec. 1–4, 2009), ACM, NY.

2. CCN NODE MODEL

CCN communication is driven by the consumers of data. There are two CCN packet types, Interest and Data (Figure 2). A consumer asks for content by broadcasting its Interest over all available interfaces. Any node hearing the Interest and having data that satisfies it can respond with a Data packet (content chunk). Data is transmitted only in response to an Interest and consumes that Interest.^a Since both Interest and Data identify the content chunks being exchanged by name, multiple nodes interested in the same content can share transmissions over a broadcast medium using standard multicast suppression techniques.³

Data “satisfies” an Interest if the ContentName (CCN name) in the Interest packet is a prefix of the ContentName in the Data packet. CCN names are opaque, binary objects composed of a (explicitly identified) sequence of components (see Figure 3). Names are typically hierarchical, so this prefix match is equivalent to saying that the Data packet is in the name subtree specified by the Interest packet (see Section 3.2). IP uses this convention to resolve the $\langle net, subnet, host \rangle$ hierarchical structure of IP addresses and experience has shown it allows for efficient, distributed hierarchical aggregation of routing and forwarding state while allowing for fast lookups.^b One implication of this matching is that

Interests may be received for content that does not yet exist—allowing a publisher to generate that content on the fly in response to a particular query. Such *active names* allow CCN to transparently support a mix of statically cached and dynamically generated content, as is common in today’s Web. Name prefixes may also be context-dependent such as */ThisRoom/projector* to exchange information with the display projector in the current room or */Local/Friends* to exchange information with any friends in the local (e.g., broadcast) environment.^c

The basic operation of a CCN node is very similar to an IP node: A packet arrives on a *face*, a longest-match look-up is done on its name, and then an action is performed based on the result of that lookup.^d The core CCN packet forwarding engine has three main data structures: the FIB (Forwarding Information Base), ContentStore (buffer memory), and PIT (Pending Interest Table).

The FIB is used to forward Interest packets toward potential source(s) of matching Data. It is almost identical to an IP FIB except it allows for a list of outgoing faces rather than a single one. This reflects the fact that CCN is not restricted to forwarding on a spanning tree. It allows multiple sources for data and can query them all in parallel.

The ContentStore is the same as the buffer memory of an IP router but has a different replacement policy. Since each IP packet belongs to a single point-to-point conversation, it has no further value after being forwarded downstream. Thus IP “forgets” about a packet and recycles its buffer immediately after forwarding (MRU replacement). CCN packets are idempotent, self-identifying, and self-authenticating, so each packet is potentially useful to many consumers (e.g., many hosts reading the same newspaper or watching the same YouTube video). To maximize the probability of sharing, which minimizes upstream bandwidth demand and downstream latency, CCN remembers arriving Data packets as long as possible (LRU or LFU replacement).

The PIT tracks Interests forwarded upstream toward content source(s), so returned Data can be sent downstream to its requester(s). In CCN, only Interest packets are routed, and as they propagate upstream toward potential Data sources, they leave a trail of “bread crumbs” for a matching Data packet to follow back to the original requester(s). Each PIT entry is a “bread crumb.” PIT entries are erased immediately after being used to forward a matching Data packet (Data “consumes” an Interest). PIT entries for Interests that never find a matching Data are eventually timed out (a “soft state” model—consumers are responsible for re-expressing an Interest if they still want the Data).

Figure 2. CCN packet types.





Figure 3. Example data name.



^a Interest and Data packets are thus one-for-one and maintain a strict flow balance. A similar flow balance between data and ack packets is what gives TCP its scalability and adaptability⁸ but, unlike TCP, CCN’s model works for many-to-many multipoint delivery (see Section 3.1).

^b While CCN names are variable length and usually longer than IP addresses, they can be looked up as efficiently. The structure of an IP address is not explicit but instead implicitly specified by the contents of a node’s forwarding table. Thus it is very difficult to apply modern $O(1)$ hashing techniques to IP lookups. Instead, $\log(n)$ radix tree search (software) or parallel but expensive TCAMs (high end hardware) are typically used. Since the CCN name structure is explicit, ContentNames can easily be hashed for efficient lookup.

^c This last example would use the explicit identity information created by CCN signing to allow friends to rendezvous via a fixed name rather than via complex enumeration or probing strategies, i.e., the name says what they want to communicate and the signature says who they are in the context of the name, e.g., “a friend in the local environment.”

^d We use the term *face* rather than *interface* because packets are not only forwarded over hardware network interfaces but also exchanged directly with application processes within a machine, as described in Section 5.

When Interest packets arrive on a face, a longest-match lookup is done on the ContentName. (Full details for Interest and Data packet processing are described in Jacobson et al.,¹¹ we provide a short overview of Interest processing here.) When an Interest arrives on a face, if there is a matching ContentObject in the ContentStore, it will be returned on the same face the Interest arrived on. If there is no ContentObject available, the PIT is checked for an existing pending Interest. If there is already a matching entry, the arrival face for the new Interest is added to the list of faces in the PIT entry used to send ContentObjects to requesters when they arrive. If there is not already an existing PIT entry, the FIB table is checked for forwarding information. If there is a corresponding entry, the Interest is forwarded and added to the PIT.

Data packet processing is relatively simple since Data is not routed but simply follows the chain of PIT entries back to the original requester(s). A longest-match lookup of a Data packet's ContentName is done upon arrival. A ContentStore match means the Data is a duplicate, so it is discarded. A FIB match means there are no matching PIT entries, so the Data is unsolicited and it is discarded.^e A PIT match (there may be more than one) means the Data was solicited by Interest(s) sent by this node. The Data is (optionally) validated (see Section 4.1) then added to the ContentStore. A new list of faces is created with the union of the matching PIT entry face lists minus the Data packet arrival face. The Data packet is then sent out each face on this new list.

The multipoint nature of Interest-based data retrieval provides flexibility to maintain communication in highly dynamic environments. Any node with access to multiple networks can serve as a content router between them. Using its cache, a mobile node may serve as the network medium between disconnected areas, or provide delayed connectivity over intermittent links. The Interest/Data exchange also functions whenever there is local connectivity. For example, two colleagues with laptops and ad hoc wireless could share corporate documents normally in an isolated location with no connectivity to the Internet or their organization.

3. TRANSPORT AND ROUTING

CCN transport is designed to operate on top of unreliable packet delivery services, including the highly dynamic connectivity of mobile and ubiquitous computing. Thus Interests, Data, or both might be lost or damaged in transit, or requested data might be temporarily unavailable. To provide reliable, resilient delivery, CCN Interests not satisfied in some reasonable period of time must be retransmitted. Unlike TCP, CCN senders are stateless and the final consumer is responsible for reexpressing unsatisfied

Interests if it still wants the data. A receiver's strategy layer is responsible for retransmission on a particular face (since it knows the timeout for the upstream node(s) on the face) as well as selecting which and how many of the available communication interfaces to use for sending interests, how many unsatisfied interests should be allowed, the relative priority of different interests, etc.

3.1. Reliability and flow control

One Interest retrieves at most one Data packet. This basic rule ensures that *flow balance* is maintained in the network and allows efficient communication between varied machines over networks of widely different speeds. Just as in TCP, however, it is possible to overlap data and requests. Multiple Interests may be issued at once, before Data arrives to consume the first. The Interests serve the role of window advertisements in TCP. A recipient can dynamically vary the window size by changing the number of Interests it issues. We show the effect of such pipelining later in Section 5.2. Since CCN packets are independently named, the pipeline does not stall on a loss—the equivalent of TCP SACK is intrinsic.

In a large network, the end-to-end nature of TCP conversations means there are many points between sender and receiver where congestion can occur from conversation aggregation even though each conversation operates in flow balance. The effect of this congestion is delay and packet loss. The TCP solution is for endpoints to dynamically adjust their window sizes to keep the aggregate traffic volume below the level where congestion occurs.⁸ The need for this congestion control is a result of TCP's flow balance being end-to-end. In CCN, by contrast, all communication is local, so there are no points between sender and receiver that are not involved in their balance. Since CCN flow balance is maintained at each hop, there is no need for additional congestion control techniques in the middle of a path. This is not the same as hop-by-hop flow control, where backpressure between adjacent nodes is used to adjust resource sharing among continuous flows. CCN does not have link FIFO queues but rather LRU memory (cache) decoupling hop-by-hop feedback control loops and damping oscillations.

3.2. Sequencing

TCP conversations between hosts identify data by simple sequence numbers. CCN needs something more sophisticated because consumers are requesting individual content chunks from large collections of data and many recipients may share the same Data packets. Locating and sharing data is facilitated by using hierarchical, aggregatable names that are at least partly meaningful to humans and reflect some organizational structure of their origin, rather than just the sequence in an ephemeral conversation. Despite this extra richness in CCN names, their transport function in Interests is exactly the same as that of sequence numbers in TCP ACKs: specifying the next Data the recipient requires.

Before explaining how the next Data is identified, we first describe the names in more detail. As mentioned,

^e "Unsolicited" Data can arise from malicious behavior, Data arriving from multiple sources, or multiple paths from a single source. In the latter cases the first copy of the Data that arrives consumes the Interest, so duplicate(s) will not find a PIT entry. In all cases the Data should be discarded since that preserves flow balance and helps guarantee stable operation under arbitrary load.

names are hierarchically structured so that an individual name is composed of a number of *components*. Each component is composed of a number of arbitrary octets—variable-length binary values that have no meaning to CCN transport. Names must be meaningful to some higher layer(s) in the stack to be useful, but the transport imposes no restrictions except the component structure. Binary encodings of integers or other complex values may be used directly without conversion to text for transmission. Name components may even be encrypted for privacy. For notational convenience we present names like URIs with '/' characters separating components, as in Figure 3, but these delimiters are not part of the names and are not included in the packet encodings. This example illustrates the application-level conventions currently used to capture temporal evolution of the content (a version marker, *_v*, encoded as FD followed by an integer version number) and its segmentation (a segment marker, *_s*, encoded as 00 followed by an integer value which might be a block or byte number or the frame number of the first video frame in the packet). The final component of every Data packet name implicitly includes a SHA256 digest of the packet.^f

An Interest can specify precisely what content chunk is required but in most cases the full name of the next Data is not known, so the consumer specifies it *relative* to something whose name is known. This is possible because the CCN name tree can be totally ordered (siblings are arranged in lexicographic order) and thus relations like *next* and *previous* can be unambiguously interpreted by the CCN transport without any knowledge of name semantics.

3.3. Rich connectivity, mobility, and strategy

Machines today typically have multiple network interfaces and are increasingly mobile. Since IP is restricted to forwarding on spanning trees, it is difficult for IP to take advantage of more than one interface or adapt to the changes produced by rapid mobility. CCN packets cannot loop, so CCN can take full advantage of multiple interfaces. CCN talks *about* data, not *to* nodes, so it does not need to obtain or bind a layer 3 identity (e.g., IP address) to a layer 2 identity such as a MAC address. Even when connectivity is rapidly changing, CCN exchanges packets as soon as it is physically possible to do so. Furthermore, since CCN Interests and Data are paired, each node gets fine grained, per-prefix, per-face performance information for adaptively choosing the “best” face for forwarding Interests matching some prefix (see Section 5.3).^g

3.4. Routing

Routing has recently experienced a resurgence of research activity. Today there are a variety of interesting and effective candidate solutions for most routing problems. Any routing scheme that works well for IP should also work well for CCN,

^f The digest component is not transmitted since it is derivable. It exists so that an Interest or a ContentName can unambiguously and exactly name any content chunk.

^g In IP, route asymmetry generally makes it impossible for an interior node to learn if an interface or route is actually functioning since it only sees one side of a conversation.

because CCN's forwarding model is a strict superset of the IP model with fewer restrictions (no restriction on multi-source, multi-destination needed to avoid looping) and the same semantics relevant to routing (hierarchical name aggregation with longest-match lookup). CCN provides an excellent vehicle to implement a routing protocol's transport: at the heart of most routing transport protocols is something very similar to CCN's information-oriented guided-diffusion flooding model, since they have to function in the pre-topology phase where peer identities and locations are not known. CCN's robust information security model enables almost automatic routing infrastructure protection.

4. CONTENT-BASED SECURITY

CCN is built on the notion of *content-based security*: protection and trust are properties of the content itself, not of the connections over which it travels. Current IP networks trust content based on the server id and the connection type used for retrieval. This intertwines content and network infrastructure security, thereby forcing clients to retrieve content from the original source to trust it. In CCN all content is digitally signed and private content is encrypted. Embodying security in content, not hosts, reduces the trust burden on network intermediaries, opening the network to wide participation.^h This section gives an overview of CCN's core security design. Additional background and motivation are described in Smetters and Jacobson.¹⁹

4.1. Content validation

Rich and robust content-based security models can be constructed if a consumer can assess just three things about any content object^{2, 5, 18}: *integrity* (is the object intact and complete), *pertinence* (what question does it answer), and *provenance* (who claims this is an answer). TCP/IP's a 16-bit checksum provides explicit, though weak, integrity, but pertinence and provenance are, at best, implicit in the source and destination addresses. CCN's pertinence is explicit in the ContentName and provenance explicit in the key used to sign the packet (“signed info” in Figure 2). Thus the packet's signature both checks integrity and authenticates the binding between name, key, and content. Standard public key signatures are used, so anyone, not just the communication endpoints, can validate them.

Consumers generally validate content in the context of an application and user dependent trust model (Section 4.2). Anything can validate that a Data was signed by the key it claims was used to generate its signature, even without attaching any real-world meaning to that key. This minimal verification is useful to detect packet corruption and to defend against network attacks. For example, consumers can request content by publisher as well as by name in the face of spurious or malicious data. To prevent work-factor attacks, no node is obligated to validate every Data.

^h CCN signing and encryption may require consumers to retrieve keys. For TCP/IP this requires side information and special purpose application protocols but in CCN keys are simply ContentObjects like everything else and, in general, the mechanism and network state required to get a content chunk are sufficient to also get all of its keys.


In particular, CCN content routers may choose to check all, some or none, as their resources allow, and dynamically adapt in response to detected attack or consumer advice.

4.2. Managing trust

Although CCN moves data in a peer-to-peer fashion, it provides *end-to-end* security between content publisher and consumer. Consumers determine whether received content is acceptable or trustworthy. CCN's trust is *contextual*, narrowly determined in the context of particular content and the purpose for which it will be used. Different consumers can even use different trust models for the same content. For example, a bank might require that a deed be signed by someone authorized by the courts while a backup server might only care that its signature validates. This model is more flexible and scalable than mandating a one-size-fits-all approach such as trusting any publisher who has paid an annual fee to some root certifying authority.

The basic primitive of content-based security—authenticated bindings from names to content—can be used to implement mechanisms for establishing higher-level trust. CCN's signed bindings between names and content act in essence to *certify* that Data. For example, when a name refers to an individual or organization and the content is a public key, the result is essentially a digital certificate. This allows CCN to support traditional mechanisms such as hierarchical PKI for establishing trust in keys. CCN also supports more general, non-hierarchical models such as SDSI^{2,5,18} where keys are mapped to identities via locally controlled *namespaces*. For example, the members of an organization might be recognized because their keys are certified by the organization itself, not because they are validated by a third-party like Verisign. If we know and trust one of `parc.com`'s employees, we get their key any time we receive content from them and by following its key locators (certification chain) can easily get the key of `parc.com` (Figure 4). Thus, starting from a small number of public keys authenticated using a variety of user-friendly mechanisms (e.g., personal contact, organizational membership, public experience^{16,20}), one can use SDSI's model to infer trust in a large number of publishers.

Figure 4. CCN trust establishment can associate content namespaces with publisher keys.



Evidence-based security. The CCN content model creates many secured relationships: Each piece of data is bound to its name and publisher's key; names are implicitly bound to other names by the naming hierarchy and application-level naming conventions; publishers are bound to other publishers by the key certification graph. As a consumer accumulates content, this web of trusted relationships can be used to automatically infer trust in entire collections of content. For example, once trust has been established for the block of video shown in Figure 4, that decision can be applied to any block of any version of `/parc.com/george/videos/WidgetA.mpg` signed by the same key. Similarly, a block of the video signed by a different key is automatically suspicious.

This notion can be extended beyond collections defined by naming hierarchy. As Section 3.2 notes, the full name of a ContentObject contains a cryptographic digest of its contents (a *self-certifying name*^{4,6,14,15}) and the identity (key) of its publisher. Embedding a link using the full name of the target ContentObject creates a *secure reference*^{4,12,13,17} that can resolve to only one object. Such references can be used to express *delegation*, saying that the publisher *P* of a link named *N* with target (*N'*, *P'*) intends the name *N* to refer to whatever publisher *P'* refers to by target name *N'*. They can also be used to build up a network of trust in content where signed Data certify the other content chunks they securely reference. For example, having decided to trust the content of web page *A*, browsers may automatically trust the images, ads, etc., that *A* securely links to. This trust is fine-grained since the linked content is only considered valid within the context of *A*.

4.3. Content protection and access control

The primary means of controlling access to CCN content is encryption. CCN does not require trusted servers or directories to enforce access control policies: No matter who stumbles across private content, only authorized users are able to decrypt it.

Encryption of content, or even names or name components, is completely transparent to the network—to CCN, it is all just named binary data (though efficient routing may require that some name components remain in the clear). Decryption keys are distributed as ContentObjects. Name conventions, encapsulated in programmer-friendly libraries, make it easy to find and decrypt the key needed by an authorized user to decrypt a given piece of content. CCN does not mandate any particular encryption or key distribution scheme. Arbitrary, application-appropriate access control models can be implemented simply by choosing how to encode and distribute decryption keys for particular content.

4.4. Network security and denial-of-service

CCN's design protects it from many classes of network attack. Authenticating all content, including routing and policy information, prevents data from being spoofed or tampered with. It is difficult to target malicious packets at host vulnerabilities since CCN Interests address content, not hosts. Effective attacks against CCN must use denial of service to either hide legitimate content or drown it in a sea of spurious packets.

TCP/IP content originates at some server and an attacker can easily hide it by placing a filter anywhere on the single

shortest path between the targeted set of clients and that server. CCN content can be supplied by anything that has a copy and every CCN node can use any and all of its interfaces simultaneously to locate and retrieve a copy. Thus hiding content is exponentially more difficult for the attacker. It must establish a filtering perimeter around its targets that covers all paths to all possible copies of the content. Any copy that makes it through the perimeter immediately becomes a new source that will virally propagate the content to all interested clients.

Drowning (DDoS) attacks can be mounted against sources of CCN content but are substantially more difficult than they are with TCP/IP. The flow balance between Interest and Data prevents any sort of Data flooding, so attackers must attack via Interest packets. Say an attacker marshals a horde of zombies to simultaneously generate interests in some ContentName. If they all use the same name, the Interests will be aggregated (at most one pending Interest in a name is ever forwarded over any link) and no flood will result. So they must all use different names under the targeted prefix. If the different names refer to actual ContentObjects, those objects will be cached everywhere along the paths from the content source(s) to the zombies; thus the flood near the source will quickly clear as Interests are satisfied by downstream cached copies of the Data.ⁱ If the zombie's names are randomly generated then their Interests will never be satisfied by a matching Data and will time out. Thus every intermediate node learns that many bogus Interests are being generated for the targeted prefix. Nodes can decide to temporarily rate limit such Interests (similar to the push-back strategy used against TCP/IP DDoS) or simply prioritize them lower than Interests that are resulting in Data responses.^j In either case the effect of the attack on legitimate traffic is minimized.

4.5. Policy controls

CCN also provides tools that allow an organization to exercise control over where their content will travel. Routers belonging to an organization or service provider can enforce *policy-based routing*, where content forwarding policy is associated with content name and signer. For example, PARC might have a “content firewall” that only allows Interests from the Internet to be satisfied if they are requesting content under the /parc.com/public namespace. An organization could also publish its policies about what keys can sign content under a particular name prefix (e.g., PARC could require that all content in the /parc.com namespace be signed by a key certified by a /parc.com root key), and have their content routers automatically drop content that does not meet those requirements, without asking those routers to understand the semantics of the names or organizations involved. Finally, Interests

could in certain cases be digitally signed, enabling policy routing to limit what namespaces or how often particular signers may query.

5. EVALUATION

In this section we describe and evaluate the performance of our prototype CCN implementation. Our current implementation encodes packets in the ccnb compact binary XML representation using dictionary-based tag compression. Our CCN forwarder, ccnd, is implemented in C as a userspace daemon. Interest and Data packets are encapsulated in UDP for forwarding over existing networks via broadcast, multicast, or unicast.

Most of the mechanics of using CCN (ccnd communication, key management, signing, basic encryption and trust management) are embodied in a CCN library. This library, implemented in Java and C, encapsulates common conventions for names and data such as encoding segmentation and versioning in names and representing information about keys for encryption and trust management. These conventions are organized into *profiles* representing application-specific protocols layered over basic CCN Interest-Data.

This architecture has two implications. First, the security perimeter around sensitive data is pushed into the application; content is decrypted only inside an application that has rights to it and never inside the OS networking stack or on disk. Second, much of the work of using CCN in an application consists of specifying the naming and data conventions to be agreed upon between publishers and consumers.

All components run on Linux, Mac OS X[™], Solaris[™], FreeBSD, NetBSD, and Microsoft Windows[™]. Cryptographic operations are provided by OpenSSL and Java.

5.1. Data transfer efficiency

TCP is good at moving data. For bulk data transfer over terrestrial paths it routinely delivers app-to-app data throughput near the theoretical maximum (the bottleneck link bandwidth). TCP “fills the pipe” because its variable window size allows for enough data in transit to fill the bandwidth $\times$ delay product of the path plus all of the intermediate store-and-forward buffer stages.⁹ CCN's ability to have multiple Interests outstanding gives it the same capability (see Section 3.1) and we expect its data transfer performance to be similar to TCP's.


To test this we measured the time needed to transfer a 6MB file as a function of the window size (TCP) and number of outstanding Interests (CCN). The tests were run between two Linux hosts connected by 100 Mbps links to our campus ethernet. For the TCP tests the file was transferred using the test tool *ttcp*. For the CCN tests the file was pre-staged into the memory of the source's ccnd by requesting it locally.^k This resulted in 6278 individually named, signed CCN Data packets (ContentObjects) each with 1KB of data (the resulting object sizes were around 1350 bytes).

ⁱ This does result in a distributed cache poisoning attack that must be addressed in the CCN node's cache replacement policy.

^j A CCN content router has a limit on the number of pending Interests it will allow on any link (generally related to the bandwidth $\times$ delay product of that link). The router can choose if it wants to hold or discard arriving Interests over the limit and how it selects Interests up to the limit.

^k This was done so the measurement would reflect just communication costs and not the signing cost of CCN content production.

^l Since CCN transacts in packet-sized content chunks, the TCP window size was divided by the amount of user data per packet to convert it to packets.

Figure 5. Bulk data transfer performance.

Results can be seen in Figure 5.¹ CCN requires five times the pipelining of TCP, 20 packets vs. 4, to reach its throughput asymptote. This is an artifact of the additional store-and-forward stages introduced by our prototype's totally unoptimized user-level implementation vs. Linux TCP's highly optimized in-kernel implementation. TCP throughput asymptotes to 90% of the link bandwidth, reflecting its header overhead (payload to packet size ratio). CCN asymptotes to 68% of the link bandwidth. Since this test encapsulates CCN in IP/UDP, it has all the overhead of the TCP test plus an additional 22% for its own headers. Thus for this example the bulk data transfer efficiency of CCN is comparable to TCP but lower due to its larger header overhead.^m

5.2. Content distribution efficiency


The preceding sections compared CCN vs. TCP performance when CCN is used as a drop-in replacement for TCP, i.e., for point-to-point conversations with no data sharing. However, a major strength of CCN is that it offers automatic, transparent sharing of all data, essentially giving the performance of an optimally situated web proxy for all content but requiring no pre-arrangement or configuration.

To measure sharing performance we compared the total time taken to simultaneously retrieve multiple copies of a large data file over a network bottleneck using TCP and CCN. The test configuration is shown in the inset of Figure 6 and consisted of a source node connected over a 10 Mbps shared link to a cluster of six sink nodes all interconnected via 1 Gbps links.ⁿ The machines were of various architectures (Intel, AMD, PowerPC G5) and operating systems (Mac OS X 10.5.8, FreeBSD 7.2, NetBSD 5.0.1, Linux 2.6.27).

The sinks simultaneously pulled a 6MB data file from the source. For the TCP tests this file was made available via an http server on the source and retrieved by the sinks using curl. For the CCN tests this file was pre-staged as described in Section 5.1. For each test, the contents of the entire file were retrieved and we recorded the elapsed time for the last

^m Most of the CCN header size increase vs. TCP is due to its security annotation (signature, witness, and key locator).

ⁿ We used a 10Mbps bottleneck link to clearly show saturation behavior, even with only a small number of nodes.

Figure 6. Total transfer time vs. the number of sinks.

node to complete the task. Multiple trials were run for each test configuration, varying the particular machines which participated as sinks.

Test results are shown in Figure 6. With a single sink, TCP's better header efficiency allows it to complete faster than CCN. But as the number of sinks increases, TCP's completion time increases linearly while the CCN performance stays constant. Note that since the performance penalty of using CCN vs. TCP is around 20% while the performance gain from sharing is integer multiples, there is a net performance win from using CCN even when sharing ratios (hit rates) are low. The win is actually much larger than it appears from this test because it applies, independently, at every link in the network and completely alleviates the traffic concentrations we now see at popular content hubs and major peering points. For example, today a popular YouTube video will traverse the link between youtube.com and its ISP millions of times. If the video were distributed via CCN it would cross that link once. With the current architecture, peak traffic loads at aggregation points scale like the total consumption rate of popular content. With CCN they scale like the popular content creation rate, a number that, today, is exponentially lower.

5.3. Voice-over-CCN and the strategy layer

To demonstrate how CCN can support arbitrary point-to-point protocols we have implemented Voice-over-IP (VoIP) on top of CCN (VoCCN). Complete details and performance measurements are given in Jacobson et al.¹⁰ In this section we describe a test that uses a VoCCN call to demonstrate the behavior and advantages of CCN's strategy layer.

As described in Section 3.3, when the FIB contains multiple faces for a content prefix, the strategy layer dynamically chooses the best. It can do this because CCN can send the same Interest out multiple faces (since there is no danger of looping) and because a CCN node is guaranteed to see the Data sent in response to its Interest (unlike IP where the request and response paths may be almost entirely disjoint). These two properties allow the strategy layer to run experiments where an Interest is occasionally sent out all faces associated with the prefix. If a face responds faster than the current best, it will become the new best and be used


exclusively for the prefix's Interests, until it is time for the next experiment (e.g., after 200 packets, when there is a change in carrier or SSID, or when an Interest does not get a response and times out).

To test this mechanism we ran our *linphone*-based VoCCN client between two Linux 2.6.27 machines (a 3.4 GHz Intel P4 and a 2.66 GHz Intel Core2 Duo) each connected to two isolated wired 1 Gbps ethernet networks. The *linphone* default is to packetize audio into 20 ms frames, so audio activity resulted in a constant 50pps source of RTP packets. As measured by voice quality, the performance of our secure VoCCN prototype was equivalent to that of stock *linphone*. No packets were lost by either client, however a few VoCCN packets (<0.1%) were dropped for arriving too late.

We conducted failover tests by manually disconnecting and reconnecting network cables. Figure 7 shows the traffic on both links during one of these tests. The strategy layer initially picks link B, but at 15 s into the call it switches to link A in reaction to some small variance in measured response time, then switches back to link B at 40 s. At 45 s we unplugged link A but this had no impact since link B was being used at the time. At 60 s link A was reconnected. At 82 s link B was disconnected. The strategy layer switches to link A, and there is a small excursion above 50pps in the link A traffic rate as the packets produced during the failure detection time are retrieved from the upstream CCND. At 95 s link B is reconnected; traffic remains on link A. At 120 s link A is disconnected, and the CCN strategy layer switches back to link B, but the failure detection took longer this time as shown by the large traffic burst on B immediately following the switch. There is one more spontaneous switch at 160 s then the call terminates at 165 s.

The failover behavior is not coded into our client but arises entirely from the CCN core transport. The small delay for the final failover reflects the preliminary state of our current implementation (it does not listen for “carrier lost” notifications from the ethernet driver, so failure detection is timeout rather than event driven). It is interesting to note that after failing over, the client retrieves the missing data from CCN: a few packets were delayed but none were lost.

Figure 7. CCN automatic failover.



6. CONCLUSION

Today's network *use* centers around moving content, but today's *networks* still work in terms of host-to-host conversations. CCN is a networking architecture built on IP's engineering principles, but using named content rather than host identifiers as its central abstraction. The result retains the simplicity and scalability of IP but offers much better security, delivery efficiency, and disruption tolerance. CCN is designed to replace IP, but can be incrementally deployed as an overlay—making its functional advantages available to applications without requiring universal adoption.

We implemented a prototype CCN network stack, and demonstrated its usefulness for both content distribution and point-to-point network protocols. We released this implementation as open source and it is available from Project CCNx™.¹

Acknowledgments

We thank Paul Stewart, Paul Rasmussen, and Simon Barber for substantial help with implementation. We also thank Ignacio Solis, Marc Mosko, Eric Osterweil, and Dirk Balfanz for fruitful discussion and advice.

References

1. Project CCNx™. <http://www.ccnx.org>, Sep. 2009.
2. Abadi, M. On SDSI's linked local name spaces. *J. Comput. Secur.* 6, 1–2 (Oct. 1998), 3–21.
3. Adamson, B., Bormann, C., Handley, M., Macker, J. *Multicast Negative-Acknowledgement (NACK) Building Blocks*. IETF, Nov. 2008. RFC 5401.
4. Clarke, I., Sandberg, O., Wiley, B., Hong, T.W. Freenet: A distributed anonymous information storage and retrieval system. In *Lecture Notes in Computer Science 2009* (2001), 46.
5. Ellison, C.M., Frantz, B., Lampson, B., Rivest, R., Thomas, B.M., Ylonen, T. *SPKI Certificate Theory*, Sep. 1999. RFC2693.
6. Fu, K., Kaashoek, M.F., Mazières, D. Fast and secure distributed read-only file system. *ACM Trans. Comput. Syst.* 20, 1 (2002), 1–24.
7. Gantz, J.F. et al. IDC—The expanding digital universe: A forecast of worldwide information growth through 2010. Technical report, Mar. 2007.
8. Jacobson, V. Congestion avoidance and control. In *SIGCOMM*, 1988.
9. Jacobson, V., Braden, R., Borman, D. *TCP Extensions for High Performance*. IETF—Network Working Group, The Internet Society, May 1992. RFC 1323.
10. Jacobson, V., Smetters, D.K., Briggs, N., Plass, M., Stewart, P., Thornton, J.D., Braynard, R. VoCCN: Voice-over content-centric networks. In *ReArch*, 2009.
11. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M., Briggs, N., Braynard, R. Networking named content. In *CoNext*, 2009.
12. Koponen, T., Chawla, M., Chun, B.-G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I. A data-oriented (and beyond) network architecture. In *SIGCOMM*, 2007.
13. Kubiawicz, J. et al. OceanStore: An architecture for global-scale persistent storage. *SIGPLAN Not.* 35, 11 (2000), 190–201.
14. Mazières, D., Kaminsky, M., Kaashoek, M.F., Witchel, E. Separating key management from file system security. In *SOSP*, 1999.
15. Moskowitz, R., Nikander, P. *Host Identity Protocol Architecture*. IETF—Network Working Group, May 2006. RFC 4423.
16. Osterweil, E., Massey, D., Tseng, B., Zhang, B., Zhang, L. Security through publicity. In *HOTSEC '06*, 2006.
17. Popescu, B.C., van Steen, M., Crispo, B., Tanenbaum, A.S., Sacha, J., Kuz, I. Securely replicated web documents. In *IPDPS*, 2005.
18. Rivest, R.L., Lampson, B. SDSI—A simple distributed security infrastructure. Technical report, MIT, 1996.
19. Smetters, D.K., Jacobson, V. Securing network content. PARC Technical Report, Oct. 2009.
20. Wendlandt, D., Andersen, D., Perrig, A. Perspectives: Improving SSH-style host authentication with multi-path probing. In *USENIX*, 2008.

Van Jacobson (van@parc.com),
PARC, Palo Alto, CA.

James D. Thornton (jthornton@parc.com),
PARC, Palo Alto, CA.

Michael Plass (plass@parc.com),
PARC, Palo Alto, CA.

Nick Briggs (briggs@parc.com),
PARC, Palo Alto, CA.

Rebecca Braynard (rbraynar@parc.com),
PARC, Palo Alto, CA.

Diana K. Smetters (smettters@alum.mit.edu),
Now at Google. Work was done while at PARC.

Azusa Pacific University Full Time Faculty Position with Azusa Pacific University

The Department of Computer Science at Azusa Pacific University invites applications for full-time faculty positions in our cutting-edge B.S. and M.S. programs emphasizing practical, information technology. A Ph.D. and previous experience in teaching and software development are preferred. Applicants are sought in selected program emphases (programming languages, software engineering, and information technology). All candidates must have a religious commitment compatible with the mission of this evangelical Christian university. APU is a Christian, coeducational, liberal arts university. For more details and information on the application procedures, visit www.apu.edu/provost/employment/positions/. For more information about the department, visit www.apu.edu/clas/computerscience/. Position subject to final funding. Azusa Pacific University does not discriminate on the basis of race, color, national origin, gender, age, disability, or status as a veteran in any of its policies, practices, or procedures. Women and minorities are encouraged to apply.

Boise State University Assistant Professor

The Department of Computer Science at **Boise State University** is seeking one to three tenure-track faculty member(s) at the **assistant professor** level. Candidates specializing in the areas of databases, software engineering, and cyber-physical systems are especially encouraged to apply. Applicants in other areas will be considered. An earned Ph.D. in Computer Science or a related field is required by the date of hire. Successful candidates will be expected to make balanced contributions in both teaching and research and to coordinate with faculty and local industry to develop and sustain funded research programs.

For more information, please visit us online at <http://coen.boisestate.edu/cs/opportunities>.

Qualified applicants should send a cover letter, CV, a statement of research interests, a statement of teaching interests, and contact information for at least three professional references to search@cs.boisestate.edu.

Review of applications begins on January 20, 2012 and will continue until position(s) are filled. Boise State University is strongly committed to achieving excellence through cultural diversity. The University actively encourages applications and nominations of women, persons of color, and members of other underrepresented groups. EEO/AA Institution, Veterans preference may be applicable.

California State University, Chico Assistant Professor

California State University, Chico Dept. of Computer Science has two full time tenure track Asst. Prof positions, starting 8/2012. EOE Employer. Please see the full Announcement at: <http://csci.ecst.csuchico.edu/jobs>

Carnegie Mellon University Faculty Position in Societal-Scale Data Analysis Assistant Professor

The H.J. Heinz III College of Carnegie Mellon University is seeking tenure-track candidates at the assistant professor level with strong training in statistical machine learning and data mining and a demonstrated commitment to addressing challenging applications in important societal or management contexts. Please see the full position announcement at www.heinz.cmu.edu/ssda-faculty-search. Apply Email: faculty-recruiting@andrew.cmu.edu. Carnegie Mellon University is an AA/EO employer.

Carnegie Mellon University, **Silicon Valley** Faculty Position Assistant Professor

We are seeking to fill a faculty position (Assistant Professor) for our growing software engineering program. The faculty member will play key roles in expanding our software engineering research program, in teaching, and in recruiting. The ideal candidate for this position will have a strong interest, along with relevant experience, in applying software engineering to system development.

Our MS programs in software engineering and software management currently enroll more than 150 students, both full-time and part-time, many of whom are currently employed at top Silicon Valley companies. These MS programs are project-based and follow a learning-by-doing approach, with no formal exams and infrequent lectures. Both daytime and evening classes are typically small. Faculty members work directly as advisors to student project teams on their deliverables, teaching knowledge and skills on a just-in-time basis. Carnegie Mellon University Silicon Valley (CMU-SV) has been steadily increasing its research activities, emphasizing mobility, networking, and security. We are looking to build our research capabilities in areas such as agile methods, mobile and embedded system development, and cloud computing. CMU-SV also offers PhD degrees through our affiliation with the Elec-

trical and Computer Engineering Department at Carnegie Mellon's Pittsburgh campus.

Please provide us with your curriculum vitae, a statement about your previous experience, and letters from five references. Starting date is August, 2012, or sooner.

We are unable to provide international relocation for this position. More information on CMUSV may be found at <http://sv.cmu.edu>.

Send applications and queries to SEFacultySearch@sv.cmu.edu

Carnegie Mellon University does not discriminate in admission, employment, or administration of its programs or activities on the basis of race, color, national origin, sex, handicap or disability, age, sexual orientation, gender identity, religion, creed, ancestry, belief, veteran status, or genetic information.

College of William and Mary Assistant Professor

We invite applications for a tenure-track assistant professor position in Computer Science for Fall 2012. We are interested in individuals with research expertise in software engineering, operating systems, or data mining and machine learning, but exceptional applicants from other areas of computer science will be considered. Applicants must hold a Ph.D. in computer science or a related field at the time of appointment and must have a strong research record and an interest in teaching.

The College of William & Mary, consistently ranked in the elite group of the Best National Universities-Doctoral by U.S. News and World Report, has committed to a multi-year effort to strengthen its computer science research program and, as a consequence, the department has been the home of multiple NSF and DOE Career Awards. Teaching loads are competitive with top research computer science departments in support of the department's research expectations for faculty. The department currently consists of fifteen faculty members who support B.S., M.S., and Ph.D. programs. More information about the department and university can be obtained at <http://www.cs.wm.edu>.

Applicants should submit a current resume, research and teaching statements, and the names of at least three references. Applicants must apply using William & Mary's online recruitment system at: <http://jobs.wm.edu>.

Columbia University Department of Computer Science Open-rank Faculty Position in Computer Science

Columbia Engineering's Department of Computer Science at Columbia University in New York

City invites applications for tenured or tenure-track faculty positions. One or more appointments at the assistant professor, associate professor and full professor, will be considered.

Columbia Engineering's strategic theme areas are Health, Information and Sustainability, and the successful candidate should contribute to the advancement of the department in these areas by developing an externally funded research program, being a thought leader in the profession, contributing to the undergraduate and graduate educational mission of the Department and providing active service to professional societies. The successful candidate is expected to establish multidisciplinary research and educational collaborations with academic departments and units across Columbia University. The Department is especially interested in qualified candidates who can contribute, through their research, teaching, and/or service, to the diversity and excellence of the academic community.

Applications are specifically sought in any of the areas that fall under the umbrella of Computer Systems, Software, Artificial Intelligence, Theory, and Computational Biology, with particular emphasis on, but not limited to: Computer Graphics, Human-Computer Interaction, Simulation, and Animation. Applicants whose work contributes to CyberBioPhysical Systems, where the biological, physical, and digital worlds fuse, are also encouraged to apply.

Candidates must have a Ph.D. or its professional equivalent by the starting date of the appointment. Applicants for this position at the Assistant Professor and Associate Professors

without tenure must have the potential to do pioneering research and to teach effectively. Applicants for this position at the tenured level (Associate or Full Professor) must have a demonstrated record of outstanding research accomplishments, excellent teaching credentials and established leadership in the field.

Candidates should apply online at:

academicjobs.columbia.edu/applicants/Central?quickFind=55402

and should submit electronically the following: curriculum-vitae including a publication list, a description of research accomplishments, a statement of research/teaching interests and plans, contact information for three people who can provide letters of recommendation, and up to three pre/reprints of scholarly work. The position will close no sooner than 12/31/2011, and will remain open until filled.

Applicants can consult www.cs.columbia.edu for more information about the department.

Columbia is an affirmative action/equal opportunity employer with a strong commitment to the quality of faculty life.

Dartmouth College

Department of Computer Science

Roth Family Distinguished Professorship

The Department of Computer Science at Dartmouth College invites applications for the inaugural Roth Family Distinguished Professorship. We seek candidates with a strong academic

or industry track record in the general area of Digital Arts (including, but not limited to, Computer Graphics, Computer Vision, Visualization, Human-Computer Interaction, Design & Media Arts). Candidates at the level of full professor or senior-level associate professor will be considered.

Dartmouth is home to a growing program in the Digital Arts with affiliated faculty and students in Computer Science, English, Film & Media Studies, Mathematics, Music, Psychology, Studio Art, and Theater. In the coming years we expect continued investment in the Digital Arts. The Roth Family Distinguished Professor will help shape what is expected to be a leading undergraduate and graduate program in Digital Arts.

The Computer Science department (www.cs.dartmouth.edu) is home to 17 tenured and tenure-track faculty members whose research encompasses the areas of digital arts, graphics, vision, algorithms, theory, systems, security, robotics, and computational biology. The Computer Science department is in the School of Arts and Sciences, and it has strong Ph.D. and M.S. programs and outstanding undergraduate majors. The department is affiliated with Dartmouth's M.D.-Ph.D. program and has strong collaborations with the Tuck School of Business and Dartmouth Medical School.

Dartmouth College is located in Hanover, New Hampshire. Dartmouth has a beautiful, historic campus, located in a scenic area on the Connecticut River. Recreational opportunities abound in all four seasons. Dartmouth hosts an annual film



香港城市大學
City University
of Hong Kong

Worldwide Search for Talent

City University of Hong Kong is a dynamic, fast-growing university that is pursuing excellence in research and professional education. As a publicly-funded institution, the University is committed to nurturing and developing students' talent and creating applicable knowledge to support social and economic advancement. Currently, the University has six Colleges/Schools. Within the next two years, the University aims to recruit **100 more scholars** from all over the world in various disciplines, including **science, engineering, business, social sciences, humanities, law, creative media, energy, environment**, and other strategic growth areas.

Applications and nominations are invited for :

Professor/Associate Professor/ Assistant Professor

Department of Computer Science [Ref. A/657/49]

Duties : Teach undergraduate and postgraduate courses, supervise research students, conduct research and perform duties as assigned by the Head of Department.

Requirements : A PhD in Computer Science or related disciplines. Candidates should have an excellent proven record of research and teaching in Computer Science or related areas including Multimedia, Software Engineering, Computer Networks, Databases, Internet Technologies, Information Security, Computer Systems, Algorithms, and Computational Biology.

Salary and Conditions of Service

Remuneration package will be driven by market competitiveness and individual performance. Excellent fringe benefits include gratuity, leave, medical and dental schemes, and relocation assistance (where applicable). Initial appointment will be made on a fixed-term contract.

Information and Application

Further information on the posts and the University is available at <http://www.cityu.edu.hk>, or from the Human Resources Office, City University of Hong Kong, Tat Chee Avenue, Kowloon, Hong Kong [Fax : (852) 2788 1154 or (852) 3442 0311/ email : hrojob@cityu.edu.hk].

Please send the nomination or application in the form of an application letter enclosing a current CV to the Human Resources Office. **Applications will receive full consideration until the positions are filled.** Please quote the reference of the post in the application and on the envelope. The University reserves the right not to fill the positions. Personal data provided by applicants will be used strictly in accordance with the University's personal data policies, a copy of which will be provided upon request.

The University also offers a number of visiting positions through its "CityU International Transition Team" for current graduate students and for early-stage and established scholars, as described at http://www.cityu.edu.hk/provost/cityu_international_transition.htm.

City University of Hong Kong is an equal opportunity employer and we are committed to the principle of diversity. We encourage applications from all qualified candidates, especially those who will enhance the diversity of our staff.

festival, as well as renowned musical and theatrical performers. Convenient public transportation to Boston and New York is available. Airports with commercial service are located 15 minutes away (Lebanon, New Hampshire) and 75 minutes away (Manchester-Boston Regional Airport).

Applicants are invited to send their CV, research statement, teaching statement, and names of at least four references. All material and inquiries should be sent to roth2012@cs.dartmouth.edu. Application review will start on January 15, 2011.

Applicants should arrange to have letters of recommendation sent directly by the recommender, either by email or physical mail, to:

Roth Family Professorship Search
Department of Computer Science
Dartmouth College
6211 Sudikoff Laboratory
Hanover, NH 03755
roth2012@cs.dartmouth.edu

Dartmouth College is an equal opportunity/affirmative action employer and encourages applications from women and members of minority groups.

Dickinson College

Two Year Visiting Computer Science Position

Dickinson College invites applications for a two-year visiting position in Computer Science, beginning Fall 2012. Applications being accepted at <https://jobs.dickinson.edu>

Applicants should have an MS or higher in CS or a related field and a commitment to innovative teaching and student/faculty research in liberal arts setting.

Florida International University (FIU)

Multiple tenure-track and tenured faculty positions at all levels

FIU is a multi-campus public research university located in Miami, a vibrant, international city. FIU offers more than 180 baccalaureate, masters, professional and doctoral degree programs to over 42,000 students. As one of South Florida's anchor institutions, FIU is worlds ahead in its local and global engagement and is committed to finding solutions to the most challenging problems of our times.

The School of Computing and Information Sciences seeks exceptionally qualified candidates for multiple tenure-track and tenured faculty positions at all levels. Outstanding candidates are sought in areas of bio/medical/health informatics, computer architecture, computer graphics, large-scale data management, search, and visualization, human-computer interaction (HCI), networking, programming languages, robotics and game theory, and telecommunication. Exceptional candidates in other areas will be considered as well. Preference will be given to candidates who will enhance or complement our existing research.

Ideal candidates for junior positions should have a record of exceptional research in their early careers. Candidates for senior positions must have an active and proven record of excellence in funded research, publications, and professional



THE CHINESE UNIVERSITY OF HONG KONG



Applications are invited for:-

Faculty of Engineering

The Faculty of Engineering is vigorously seeking several faculty posts at Professor / Associate Professor / Assistant Professor levels with prospect for substantiation in the following interdisciplinary areas to pursue new strategic research initiatives, to fill up vacancies to complement current/planned strengths in its Departments and to teach in the new curriculum to be launched in the academic year 2012-2013. Successful candidates will be expected to develop a significant independent research programme with external funding, supervise postgraduate students, and teach both undergraduate and postgraduate courses.

Interdisciplinary Areas/Programmes

(1) Security (Ref. 1112/076(255)/2)

Security is identified as one of the Faculty's strategic research areas, to be developed in the Department of Computer Science and Engineering and Department of Information Engineering. Talented candidates are sought to complement existing efforts and create new synergies. Candidates in cryptography and/or network/system security are welcome to apply. This area is supported by the University's Focused Investment in Information Sciences.

(2) Biomedical Engineering Programme (Ref. 1112/077(255)/2)

Biomedical Engineering is a new academic programme launched by the Faculty of Engineering in 2010-2011 in close collaboration with the Faculty of Medicine, and is targeted to become full-fledged by 2012-13. The programme is administratively hosted in the Department of Electronic Engineering with academic support from various disciplines. It is looking for excellent candidates for junior faculty openings. Joint appointments across relevant Departments will be considered with reference to the specific disciplines of the successful candidates.

Other Openings at the Departments in the Faculty

(3) Department of Computer Science and Engineering (Ref. 1112/078(255)/2)

Computer Engineering

The Department is looking for a leader and a couple of young and aspiring professors for added momentum to its Computer Engineering programme, and in particular, talents in the areas of multi-core architecture and embedded systems.

(4) Department of Electronic Engineering (Ref. 1112/079(255)/2)

Nanotechnology

The Department is looking for excellent candidates in the area of nanotechnology. The Department aims to recruit a faculty member with expertise particularly in low-dimensional materials and devices. Specific research areas may be related to carbon based materials and devices (i.e. carbon nanotubes, graphene nanoelectronics/optoelectronics, molecular/polymeric electronics and photonics), nanomaterials for advanced energy technologies, etc.

(5) Department of Information Engineering (Ref. 1112/080(255)/2)

Internet Services and Applications

The Department is looking for strong candidates in the area of Internet services and applications, e.g. online social networks and Web analytics.

(6) Department of Mechanical and Automation Engineering (Ref. 1112/081(255)/2)

(i) Energy; (ii) MEMS; (iii) Nano System; (iv) Material Technology; (v) Advanced Manufacturing

The Department is going to launch a new undergraduate programme of Energy Engineering in 2012-13 and is looking for candidates with excellent track records in the areas of renewable energy technologies, such as wind power, kinetic energy harvesting, smart metering and smart grid. Candidates in the areas of MEMS, nano system, materials technology and advanced manufacturing are also welcome to apply.

(7) Department of Systems Engineering and Engineering Management (Ref. 1112/082(255)/2)

(i) Financial Engineering; (ii) Operation Research/Supply Chain Management/Service Engineering

The Department welcomes candidates with expertise in financial engineering, operations research/supply chain management/service engineering.

For posts (1) to (7): Applicants should have a PhD degree and a good scholarly record demonstrating potential for teaching and research excellence. Appointments will normally be made on contract basis for two to three years initially commencing August 2012, which, subject to performance and mutual agreement, may lead to longer-term appointment or substantiation later. Applications will be accepted until the posts are filled.

Salary and Fringe Benefits

Salary will be highly competitive, commensurate with qualifications and experience. The University offers a comprehensive fringe benefit package including medical care, a contract-end gratuity for appointments of two years or longer, and housing benefits for eligible appointees. Further information about the University and the general terms of service for appointments is available at <http://www.cuhk.edu.hk/personnel>. The terms mentioned herein are for reference only and are subject to revision by the University.

Application Procedure

Please send full resume, copies of academic credentials, publication list with abstracts of selected published papers, details of courses taught and evaluation results (if any), a research plan, a teaching statement, together with names, addresses and fax numbers/e-mail addresses of three to five referees to whom the applicants' consent has been given for their providing references (unless otherwise specified), to the Dean, Faculty of Engineering by e-mail to recruit@erg.cuhk.edu.hk. Applicants are asked to indicate the Department and area(s) in which they are interested. The Personal Information Collection Statement will be provided upon request. Please quote the reference number and mark 'Application - Confidential' on cover.

service, as well as a demonstrated ability to develop and lead collaborative research projects. In addition to developing or expanding a high-quality research program, all successful applicants must be committed to excellence in teaching at both graduate and undergraduate levels. An earned Ph.D. in Computer Science or related disciplines is required.

Florida International University (FIU), the state university of Florida in Miami, is ranked by the Carnegie Foundation as a comprehensive doctoral research university with high research activity. The School of Computing and Information Sciences (SCIS) is a rapidly growing program of excellence at the University, with 31 faculty members and 1,400 students, including 65 Ph.D. students. SCIS offers B.S., M.S., and Ph.D. degrees in Computer Science, an M.S. degree in Telecommunications and Networking, and B.S., B.A., and M.S. degrees in Information Technology. SCIS has received approximately \$12.6M in the last three years in external research funding, has six research centers/clusters with first-class computing infrastructure and support, and enjoys broad and dynamic industry and international partnerships.

HOW TO APPLY:

Applications, including a letter of interest, contact information, curriculum vitae, and the names of at least three references, should be submitted directly to the FIU J.O.B.S Link website at <https://www.fiujobs.org>; refer to Position # 33334. The application review process will begin on January 16, 2012, and will continue until the position is

filled. Further information can be obtained from the School website <http://www.cis.fiu.edu>, or by e-mail to recruit@cis.fiu.edu.

FIU is a member of the State University System of Florida and is an Equal Opportunity, Equal Access Affirmative Action Employer.

George Mason University **Department of Computer Science** **Term Assistant Professor-Computer Science**

The George Mason University, Department of Computer Science within the Volgenau School of Engineering invites applications for a renewable, non tenure-track Assistant Professor position in computer science beginning Fall 2012.

Responsibilities include teaching undergraduate computer science courses as well as service duties associated with the department's undergraduate degree programs. For information on the department, visit our Web site: <http://cs.gmu.edu/>.

Minimum qualifications include a graduate degree, preferably a Ph.D. in computer science, software engineering or related field, and a strong commitment and demonstrated excellence in teaching.

For full consideration please submit a faculty application and application materials online at <http://jobs.gmu.edu> for position number F9550Z. To apply, you will need a statement of professional goals including your perspective on teaching, a complete C.V. with publications and teaching history, and the names of three references.

The review of applications will begin on February 13, 2012, and will continue until the position is filled.

GMU is an equal opportunity/affirmative action employer. Women and minorities strongly encouraged.

Georgia Institute of Technology **Tenure-Track Faculty**

The School of Computational Science and Engineering of the College of Computing at the Georgia Institute of Technology invites applications for tenure-track faculty positions. Applications at all levels of service will be considered. Applicants must have an outstanding record of research, a sincere commitment to teaching, and interest in engaging in substantive interdisciplinary research with collaborators in other disciplines.

We encourage applications from all areas of computational science and engineering. Applicants with expertise in areas related to high-performance computing, scientific computing, large-scale data analytics and systems, visualization, computational systems biology, modeling & simulation, or complex systems are especially encouraged to apply.

Applications will be considered until open positions are filled, however, to receive full consideration, applications should be submitted online through <https://recruiting.cc.gatech.edu/cse> by December 16, 2011. The application material should include a full academic CV, teaching and research statements, a list of at least three references and up to three publications.

Applicants are encouraged to clearly identify in their cover letter the area(s) that best describe their research interests. Georgia Tech is an Affirmative Action/Equal Opportunity Employer. Applications from women and under-represented minorities are strongly encouraged.

Harvard University **Assistant Professor**

The Harvard School of Engineering and Applied Sciences (SEAS) seeks applicants for positions at the level of tenure-track assistant professor in the fields of Computer Science and Applied Math/Computer Science, with an expected start date of July 1, 2012.

Position 1: Computer Science. We welcome outstanding applicants in all areas of computer science. We are particularly interested in systems, broadly defined, including compilers, programming languages, distributed systems, databases, networking, and operating systems. Applicants will apply on-line at <http://academicpositions.harvard.edu/postings/3825>

Position 2: Applied Math/Computer Science. We welcome outstanding applicants in all areas of applied mathematics or theoretical computer science. We are particularly interested in topics at the boundary or intersection of these fields, including optimization, applied probability, scientific computing, combinatorics and graph theory, approximation algorithms, and numerical analysis. Applicants will apply on-line at <http://academicpositions.harvard.edu/postings/3824>.



The Hong Kong Polytechnic University is the largest government-funded tertiary institution in Hong Kong in terms of student number. It offers programmes at Doctorate, Master's, Bachelor's degrees and Higher Diploma levels. It has a full-time academic staff strength of around 1,200. The total consolidated expenditure budget of the University is in excess of HK\$4 billion per year.

DEPARTMENT OF COMPUTING

Founded in 1974, the Department of Computing (COMP) of The Hong Kong Polytechnic University was amongst the first in Hong Kong to offer education in Computing and Information Technology. Over the past decades, the department has gained its reputation and international recognition through its high quality teaching and cutting edge research which not only attracted the dedicated professionals but also inspired generations of young people. Today, COMP has emerged with its competent courses, competitive projects and comprehensive services for the benefits of society. More information about COMP is available at <http://www.comp.polyu.edu.hk/>.

Associate Professor / Assistant Professor in Software Engineering / Cloud Computing / Social Computing / E-business / Business Intelligence / Knowledge Engineering (two posts) (Ref. 11111707)

The appointees will be required to (a) initiate innovative research projects and thus obtain competitive external grants; (b) engage in research collaborations both internally and internationally; (c) teach at both undergraduate and postgraduate levels, and supervise research students; and (d) participate in professional services to the academic community and in promotional activities.

Applicants should have a PhD degree with an excellent track record of research in one of the above areas relevant to the strategic development of the department. Applicants should also show strong evidence of commitment to research that leads to top quality publications with high impacts.

Remuneration and Conditions of Service

A highly competitive remuneration package will be offered. Initial appointments for Assistant Professor will be on a fixed-term gratuity-bearing contract. Re-engagement thereafter is subject to mutual agreement. An appropriate term will be provided for appointments of Associate Professor. Applicants should state their current and expected salary in the application.

Application

Please submit application form together with detailed curriculum vitae and provide names and addresses of three referees via email to hrstaff@polyu.edu.hk; by fax at (852) 2364 2166; or by mail to **Human Resources Office, 13/F, Li Ka Shing Tower, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong**. Application forms can be obtained via the above channels or downloaded from <http://www.polyu.edu.hk/hro/job.htm>. The closing date for application is **10 February 2012**. Details of the University's Personal Information Collection Statement for recruitment can be found at <http://www.polyu.edu.hk/hro/jobpics.htm>.

For further details about the University, please visit www.polyu.edu.hk

In terms of applications, areas of interest include computational science, engineering, or the social sciences. We encourage applications from candidates whose research examines computational issues raised by very large data sets or massively parallel processing.

Candidates are required to have a PhD or an equivalent terminal degree, or to be able to certify that they will receive the degree within one year of the expected start date. In addition, we seek candidates who have an outstanding research record and a strong commitment to undergraduate teaching and graduate training.

Required documents include a CV, a statement of research and teaching interests, up to three representative papers, and names and contact information for at least three references.

The Computer Science and Applied Mathematics programs at Harvard University benefit from outstanding undergraduate and graduate students, an excellent location, significant industrial collaboration, and substantial support from the Harvard School of Engineering and Applied Sciences. Information about Harvard's current faculty, research, and educational programs is available at <http://www.seas.harvard.edu>.

Applications will be reviewed as they are received. The review of applications will begin on December 15, 2011, and applicants are strongly encouraged to submit applications by that date; however, applications will continue to be accepted at least until January 15, 2012. Harvard is an Equal Opportunity/ Affirmative Action Employer. Applications from women and minority candidates are strongly encouraged.

Kettering University
Department of Computer Science
Assistant Professor of Computer Science

Kettering University's Department of Computer Science, located in Flint, MI invites applications for a tenure-track Assistant Professor position. Requirements include a Ph.D. in Computer Science in the areas of Virtual Reality, Computer Simulation or Computer Gaming; demonstrated or exceptional promise of excellence in teaching a broad range of undergraduate courses; evidence or exceptional promise of research in the areas of Virtual Reality, Simulation or Gaming. Job duties: Teach undergraduate courses in computer science; advise students; conduct research and publish results in peer-reviewed journals & at conferences; actively pursue research funding opportunities; and provide service to the department, university, profession and community. Application review will begin December 1, 2011 and will continue until the position is filled, to start July 2012. For more information or to apply, visit jobs.kettering.edu/applicants/Central?quickFind=50991 for application instructions. AA/EOE

Lawrence Technological University
Director of Video Game Development

The Department of Mathematics and Computer Science at Lawrence Tech anticipates the appointment of a Director of Video Game Development.

The ideal candidate will have at least a Masters degree in computer science, at least 5 years of ex-

perience as a game designer, project management experience, be committed to the development of a nationally recognized program, be primarily committed to the development of undergraduate and professional graduate computer science students through teaching, applied projects and scholarship, be able to work effectively in interdisciplinary teams, and believe strongly in the value of both theory and application. Applicants should email a cover letter, curriculum vitae, statement of teaching philosophy and research interests, and three letters of recommendation via email.

Contact Person: Computer Science Search Committee
 Email Address: cssearch@ltu.edu

Max Planck Institute for Software Systems (MPI-SWS)
Tenure-track openings

Applications are invited for tenure-track and tenured faculty positions in all areas related to the study, design, and engineering of software systems. These areas include, but are not limited to, data and information management, programming systems, software verification, parallel, distributed and networked systems, and embedded systems, as well as cross-cutting areas like security, machine learning, usability, and social aspects of software systems.

A doctoral degree in computer science or related areas and an outstanding research record are required. Successful candidates are expected to build a team and pursue a highly visible re-



ADVERTISING IN CAREER OPPORTUNITIES

How to Submit a Classified Line Ad: Send an e-mail to acmm mediasales@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: \$325.00 for six lines of text, 40 characters per line. \$32.50 for each additional line after the first six. The MINIMUM is six lines.

Deadlines: 20th of the month/2 months prior to issue date. For latest deadline info, please contact: acmm mediasales@acm.org

Career Opportunities Online: Classified and recruitment display ads receive a free duplicate listing on our website at: <http://jobs.acm.org>

Ads are listed for a period of 30 days. For More Information Contact:

**ACM Media Sales
 at 212-626-0686 or
acmm mediasales@acm.org**



TEMASEK RESEARCH FELLOWSHIP (TRF)

The Nanyang Technological University (NTU) and the National University of Singapore (NUS) invite outstanding young researchers with a PhD Degree in science or technology to apply for the prestigious TRF awards.

The TRF scheme provides selected young researchers an opportunity to conduct and lead research that is relevant to defence. It offers:

- 3-year research grant that commensurate with the scope of work, with an option to extend up to a further 3 years.
- Possible tenure-track academic appointment with the university at the end of the TRF.
- Attractive and competitive remuneration.

Fellows may lead and conduct research, and publish in these areas:

1. Advanced Protective Materials
2. Bio-mimetic Aerodynamics
3. Cyber Security
4. Sensemaking Technology
5. Sensor Systems and Signal Processing

Other fundamental areas of science or technology, where a breakthrough would be of interest to defence and security, will also be considered.

Singapore is a globally connected cosmopolitan city-state with a supportive environment and vibrant research culture. For more information and application procedure, please visit

NTU – http://www3.ntu.edu.sg/trf/index_trf.html

NUS – <http://www.nus.edu.sg/dpr/funding/trf.htm>

Closing date: 10 February 2012 (Friday)

Shortlisted candidates will be invited to Singapore to present their research plans, meet local researchers and identify potential collaborators in May 2012.

search agenda, both independently and in collaboration with other groups.

Senior candidates must have demonstrated leadership abilities and recognized international stature.

MPI-SWS, founded in 2005, is part of a network of eighty Max Planck Institutes, Germany's premier basic research facilities. MPIs have an established record of world-class, foundational research in the fields of medicine, biology, chemistry, physics, technology and humanities. Since 1948, MPI researchers have won 17 Nobel prizes. MPI-SWS aspires to meet the highest standards of excellence and international recognition with its research in software systems.

To this end, the institute offers a unique environment that combines the best aspects of a university department and a research laboratory:

a) Faculty receive generous base funding to build and lead a team of graduate students and post-docs. They have full academic freedom and publish their research results freely.

b) Faculty supervise doctoral theses, and have the opportunity to teach graduate and undergraduate courses.

c) Faculty are provided with outstanding technical and administrative support facilities as well as internationally competitive compensation packages.

MPI-SWS currently has 11 tenured and tenure-track faculty, and is funded to support 17 faculty and about 100 doctoral and post-doctoral positions. Additional growth through outside funding is possible. We maintain an open, international and diverse work environment and seek applications from outstanding researchers regardless of national origin or citizenship. The working language is English; knowledge of the German language is not required for a successful career at the institute.

The institute is located in Kaiserslautern and Saarbrücken, in the tri-border area of Germany, France and Luxembourg. The area offers a high standard of living, beautiful surroundings and easy access to major metropolitan areas in the center of Europe, as well as a stimulating, competitive and collaborative work environment. In immediate proximity are the MPI for Informatics, Saarland University, the Technical University of Kaiserslautern, the German Center for Artificial Intelligence (DFKI), and the Fraunhofer Institutes for Experimental Software Engineering and for Industrial Mathematics.

Qualified candidates should apply online at <http://www.mpi-sws.org/application>. The review of applications will begin on January 3, 2012, and applicants are strongly encouraged to apply by that date; however, applications will continue to be accepted through January 2012.

The institute is committed to increasing the representation of minorities, women and individuals with physical disabilities in Computer Science. We particularly encourage such individuals to apply.

Milcord LLC **R&D Software Engineer**

Develop innovative state-of-the-art knowledge management software applications. Lead the design, prototyping, and development. Must have MS degree in Information Systems Management plus 6 months experience as a Software Engineer. Send resume to: Director, HR, Milcord LLC, 303 Wyman Street, Suite 300, Waltham, MA 02451

Milwaukee School of Engineering **Software Engineering Faculty**

The **Milwaukee School of Engineering** invites applications for a full-time faculty position in our Software Engineering program beginning in the fall of 2012. Rank will depend on qualifications and experience of the candidate. Applicants must have an earned doctorate degree in software engineering, computer engineering, computer science or closely related field, as well as relevant experience in software engineering practice.

The successful candidate must be able to contribute in several areas of software engineering process and practice while providing leadership in one of the following: software architecture and design, software requirements, or real-time systems.

MSOE expects and rewards a strong primary commitment to excellence in teaching at the undergraduate level. Continued professional development is also expected.

Our ABET-accredited undergraduate software engineering program had its first graduates in spring 2002. Founded in 1903, MSOE is a private, application-oriented university with programs in engineering, business, and nursing. MSOE's 15+ acre campus is located in downtown Milwaukee, in close proximity to the Theatre District and Lake Michigan. Please visit our website at www.msOE.edu.

Submit all application material via email in pdf format to work@msOE.edu. Applicants should include a letter of application, curriculum vitae, statement of teaching interests, and names (with email and physical addresses) of at least three references.

MSOE is an Equal Opportunity/Affirmative Action Employer

National Taiwan University **Professor – Assistant** **Professor - Associate, Professor**

The Department of Computer Science and Information Engineering has faculty openings at all ranks beginning in August 2012. Highly qualified candidates in all areas of computer science/engineering are invited to apply. A Ph.D. or its equivalent is required. Applicants are expected to conduct outstanding research and be committed to teaching. Candidates should send a curriculum vitae, three letters of reference, and supporting materials before January 31, 2012, to Prof. Ai-Chun Pang, Department of Computer Science and Information Engineering, National Taiwan University, No 1, Sec 4, Roosevelt Rd., Taipei 106, Taiwan.

Oregon Health & Science University **Neuroengineering/Ubiquitous** **Computing Faculty Position**

The Dept of Biomedical Engineering in the School of Medicine at Oregon Health & Science University invites applications at all ranks for a faculty position in neuroengineering starting 9/12. The position extends our programs in technologies for early detection and remediation of cognitive decline, functional disability, and chronic disease through unobtrusive, continuous, in-home technologies by adding a researcher with a well-developed program in behavioral and predictive modeling. Faculty appointees in BME

maintain independent, extramurally funded research programs and teach supporting courses. Apply at www.ohsujobs.com, position IRC35204 by 1/31/2012.

Oregon State University **School of Electrical Engineering and** **Computer Science** **Assistant/Associate/Full Professor tenure-track position in Computer Science**

The School of Electrical Engineering and Computer Science at Oregon State University invites applications for up to four tenure-track positions in Computer Science. We seek strong candidates with a commitment to quality teaching and with research strengths in the areas of programming languages, algorithms with a focus on optimization or probabilistic reasoning, systems-focused HCI, databases with a focus on very large data systems (excluding data mining), and computer security and privacy. Exceptionally strong candidates in other areas of computer science will also be considered. Applicants should demonstrate a strong commitment to collaboration with other research groups in the School of EECS, with other departments at Oregon State, and with other universities.

The School of EECS supports a culture of energetic collaboration, and the faculty are committed to excellence in both education and research. With 40 tenure/tenure-track faculty, we enroll 185 Ph.D., 125 MS and 1200 undergraduate students. OSU is recognized for its "very high research activity" by the Carnegie Foundation for the Advancement of Teaching. The School of EECS is housed in the Kelley Engineering Center, a green building designed to support collaboration among faculty and students across campus.

Corvallis is a college town renowned for its high quality of life.

https://jobs.oregonstate.edu/applicants/jsp/shared/Welcome_css.jsp

Princeton University Computer **Science Department** **Assistant Professor** **Tenure-Track Positions**

The Department of Computer Science at Princeton University invites applications for faculty positions at the Assistant Professor level. We are accepting applications in all areas of Computer Science.

Applicants must demonstrate superior research and scholarship potential as well as teaching ability. A PhD in Computer Science or a related area is required. Successful candidates are expected to pursue an active research program and to contribute significantly to the teaching programs of the department. Applicants should include a CV and contact information for at least three people who can comment on the applicant's professional qualifications.

There is no deadline, but review of applications will start in December 2011. Princeton University is an equal opportunity employer and complies with applicable EEO and affirmative action regulations. You may apply online at: <http://jobs.cs.princeton.edu/>.

Requisition Number: 0110422

Princeton University
Computer Science Department
PostDoctoral Research Associate

The Department of Computer Science at Princeton University is seeking applications for post-doctoral or more senior research positions in theoretical computer science. Candidates will be affiliated with the Center for Computational Intractability (CCI) or the Princeton Center for Theoretical Computer Science. Candidates should have a PhD in Computer Science, a related field, or on track to finish by August 2012.

Candidates affiliated with the CCI will have visiting privileges at partner institutions NYU, Rutgers University, and The Institute for Advanced Study. Review of candidates will begin on Jan 1, 2012, and will continue until positions are filled. Applicants should submit a CV and research statement, and contact information for three references. Princeton University is an equal opportunity employer and complies with applicable EEO and affirmative action regulations.

Apply to: <http://jobs.princeton.edu/> requisition# 0110698

Rose-Hulman Institute of Technology
Faculty Positions in Computer Science and Software Engineering

The Department of Computer Science and Software Engineering at Rose-Hulman Institute of Technology invites applications for tenure-track and visiting faculty positions beginning August 2012. Rank and salary will be commensurate with qualifications. The department offers B.S. degrees in software engineering and computer science, as well as an M.S. in software engineering. Both undergraduate programs are ABET accredited and the graduate program is based on the new Graduate Software Engineering Reference Curriculum.

Ranked #1 among undergraduate engineering colleges by US News and World Report for the 13th year running, Rose-Hulman is a highly selective, coeducational, primarily undergraduate college of engineering, science and mathematics (www.rose-hulman.edu). Rose-Hulman professors are outstanding teachers.

Detailed information is available from Cary Laxer, Head of Computer Science and Software Engineering. laxer@rose-hulman.edu (812) 877-8429 www.csse.rose-hulman.edu/faqs. Applicants should submit a cover letter, a resume, a statement on their teaching philosophy, and a statement on their professional development goals to <https://jobs.rose-hulman.edu>. Three recommendation letters should be submitted by the applicant's references through the same website. Applications will be accepted until the positions are filled. EEO/AA.

Southern Methodist University
Position #51048
Department of Computer Science and Engineering
Faculty Position in Computer Science and Engineering

The Department of Computer Science and Engineering in the Lyle School of Engineering at Southern Methodist University invites applications for a faculty position in computer science

and engineering beginning Fall 2012. Individuals with experience and research interests in all areas of computer science and engineering are encouraged to apply. Priority will be given to individuals with expertise and research interest in algorithms, data mining, informatics, and related areas. The search is focused at the tenure-track assistant professor level. The successful candidates must have or expect to have a Ph.D. in computer science, computer engineering, or a closely related area by date of hire. Successful applicants will demonstrate a deep commitment to research activity in computer science and engineering and a strong potential for excellence in teaching.

The Dallas/Fort Worth area, one of the top three high-tech industrial centers in the country, has the largest concentration of telecommunications corporations in the US, providing abundant opportunities for industrial research cooperation and consulting. Dallas/Fort Worth is a multifaceted business and high-tech community, offering exceptional museums, diverse cultural attractions, and a vibrant economy.

The CSE Department resides within the Bobby B. Lyle School of Engineering and offers BS, MS, and Ph.D. degrees in Computer Engineering and Computer Science, the Doctor of Engineering in software engineering, and the MS in Security Engineering and Software Engineering. The department currently has 15 faculty members with research concentrations in security engineering, software engineering, computer networks, telecommunications, data mining, database systems, VLSI and digital systems, and computer arithmetic. Additional information may be found at: www.lyle.smu.edu/cse.

Interested individuals should send a complete resume and names of three references, including a one-page statement of research interests and accomplishments to:

csesearch@lyle.smu.edu

or

CSE Faculty Search Position #51048
Department of Computer Science and Engineering
SMU
Dallas, TX 75275-0122

The committee will begin its review of the applications on or about February 1, 2012. To ensure full consideration, applications must be time and date stamped before February 1, 2012. However, the committee will continue to accept applications until all positions are filled. Hiring is contingent upon the satisfactory completion of a background check.

SMU will not discriminate on the basis of race, color, religion, national origin, sex, age, disability, or veteran status. SMU is committed to nondiscrimination on the basis of sexual orientation.

Stanford University
Graduate School of Business
Faculty Positions in the Stanford Institute for Innovation in Developing Economies

The Graduate School of Business at Stanford University seeks to fill three positions for faculty members whose research and teaching are in areas that will support the nascent Stanford Institute for Innovation in Developing Economies ("SEED"). We seek faculty who are established leaders in related fields to provide intellectual energy and leadership

to this effort. We believe that the launch of the Institute provides us with an unprecedented opportunity to attract academic leaders in the study of poverty alleviation to Stanford. We seek full time appointments in the existing areas of the school (Accounting; Economics; Finance; Marketing; Operations, Information and Technology; Organizational Behavior; and Political Economy) of faculty members who will be affiliated with the Institute. Any faculty member appointed through this process should have significant knowledge of and interest in gaining broader on the ground experience in the developing world. We expect these to be senior appointments, although we will consider junior faculty who are on a clear trajectory to be leaders in the field and who demonstrate the capability for institutional leadership. The candidate's research and teaching interests must be related to the core mission of the institute, in particular poverty alleviation through entrepreneurship and management of growing enterprises in the developing world.

Applicants are strongly encouraged to submit their applications electronically by visiting the web site <http://www.gsb.stanford.edu/recruiting> and uploading their curriculum vitae, research papers and publications, and teaching evaluations, if applicable, on that site. Alternatively, all materials may be sent by e-mail to faculty_recruiter@gsb.stanford.edu, or by postal mail (non-returnable) to Office of Faculty Recruiting, Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA 94305-7298. However, submissions via e-mail and postal mail can take 4-6 weeks for processing. Applicants may include three letters of recommendation e-mailed to the preceding e-mail address, or sent via postal mail. The search is open until filled.

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 members of minority groups, as well as others who would bring additional dimensions of diversity to the University's research, teaching and clinical missions.

Swarthmore College
Computer Science Department
Tenure Track and Two Year Leave Position

Applications are invited for a tenure track assistant professor position and for a two-year leave replacement position at the assistant professor level. Both positions begin Fall semester 2012. Swarthmore College is a small, selective liberal arts college located in a suburb of Philadelphia. The Computer Science Department offers majors and minors in computer science at the undergraduate level. The teaching load is 3 courses and 3 labs per year. Applicants must have teaching experience and should be comfortable teaching a wide range of courses at the introductory and intermediate level.

For the tenure track position, all areas of CS will be considered, but we are particularly interested in areas that complement our current offerings, including databases, networking, and other systems and algorithms areas. Candidates should additionally have a strong commitment to involving undergraduates in their research. A Ph.D. in CS by or near the time of appointment is required.

For the leave replacement position, all areas of CS will be considered. A Ph.D. in CS by or near the time of appointment is preferred (ABD required).

See <http://cs.swarthmore.edu/jobs> for appli-

cation submission information and more details about both positions. We expect to begin interviewing by mid January 2012. Applications will be accepted until the positions are filled.

Swarthmore College has a strong institutional commitment to excellence through diversity in its educational program and employment practices and actively seeks and welcomes applications from candidates with exceptional qualifications, particularly those with demonstrable commitments to a more inclusive society and world.

Texas Christian University **Assistant Professor of Computer Science**

The Department of Computer Science at Texas Christian University seeks candidates at the Assistant Professor level for a tenure-track position in Computer Science to start on or after August 1, 2012. Applicants are expected to have an earned doctorate in Computer Science or related field, must have excellent verbal and written communication skills, and a strong commitment to teaching and research. Responsibilities include teaching undergraduate computing courses in the Department's programs in Computer Information Technology (CITE) and Computer Science (COSC); advising and mentoring majors; developing computing curricula in both programs; conducting research and engaged in scholarship in computing fields; serving on college and university committees; and serving the computing community of educators, researchers, and professionals. Preferred areas of specialization include: operating systems, mobile and cloud computing, and/or parallel architectures. Salary is commensurate with qualifications and experience.

TCU only accepts online applications. Interested candidates are directed to: https://tcu.igreentree.com/CSS_faculty. All required documents must be attached to the application during the input process. Do not send documents directly to the department.

Applicants must upload a cover letter, vita and representative publications and have three letters of recommendation sent directly to Jacquelyn Curry in TCU Human Resources at j.curry@tcu.edu. The candidate should not send these confidential recommendations.

Review of applications will begin immediately and continue until the position is filled.

Toyota Technological Institute Chicago **Computer Science at TTI Chicago** **Faculty Positions at All Levels**

Toyota Technological Institute at Chicago (TTIC) is a philanthropically endowed degree-granting institute for computer science located on the University of Chicago campus. Applications are being accepted in all areas, but we are particularly interested in machine learning, speech processing, computational linguistics, Computer vision, computational biology and optimization. Positions are available at all ranks, and we have a large number of three year limited term positions currently available. For all positions we require a Ph.D. Degree or Ph.D. candidacy, with the degree conferred prior to date of hire. Submit your application electronically at: <http://ttic.uchicago.edu/facapp/>

Toyota Technological Institute at Chicago is an Equal Opportunity Employer

University College London (UCL) **Department of Computer Science** **Faculty Positions**

The Department of Computer Science at University College London (UCL) invites applications for faculty positions in the areas of Networking, Systems, Bioinformatics, and Programming Languages. Appointments will be made at the rank of Lecturer, Senior Lecturer, or Reader (equivalent to Assistant Professor, junior Associate Professor, and senior Associate Professor, respectively, in the US system), commensurate with qualifications. In Networking and Systems, particular interests include operating systems, systems security, distributed systems, networking, and their intersection, with an emphasis on experimental system-building. In Bioinformatics, particular interests include novel algorithm development in genomics (particularly for Next Generation Sequencing and other high-throughput data methodologies), novel algorithm development for other cutting-edge biological problems, and the analysis of biological networks. In Programming Languages, interests include both theory and practice.

Applicants must hold an earned Ph.D. in Computer Science or a closely related field (or in the case of Bioinformatics applicants, in the life sciences) by the time they begin their appointment. In either case, 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 program of research. They must also demonstrate a zest for teaching at the graduate and undergraduate levels.

Since 1973, when UCL CS became the first ARPAnet node outside the United States, the department has been a leading centre for networking research, as demonstrated by its long-standing strong presence in the SIGCOMM conference proceedings. UCL also has a long history of world-leading excellence in bioinformatics and computational biology, and is home to over 20 groups actively working in the area, as well as internationally recognised resources such as the PSIPRED server and the CATH database.

UCL, the third oldest university in England, and the first to admit students without regard for their race or creed, sits a few blocks from the British Museum in central London, a vibrant metropolis of 7.1 million that offers an ever-renewing panoply of cultural and culinary choices, and is a budget airfare away from all of continental Europe.

Further details about UCL CS, the posts, and how to apply may be found at:

<http://www.cs.ucl.ac.uk/vacancies>

Our application procedure should be read before you apply, as we have specific requirements regarding applying online and reference letters arriving by the closing date.

All application materials must reach UCL by the **20th of January, 2012**.

Please note that UCL CS will **not** request your letters; it is your responsibility to ensure your letter-writers send their letters to UCL CS in time to arrive by the above deadline.

Questions about these posts may be directed to Brad Karp (Head of Networks Research Group, bkarp@cs.ucl.ac.uk), David Jones (Head of Bioinformatics Research Group, d.jones@cs.ucl.ac.uk), and John Shawe-Taylor (Head of Depart-

ment of Computer Science, jst@cs.ucl.ac.uk).

We particularly welcome female applicants and those from an ethnic minority, as they are under-represented within UCL at this level.

UCL Taking Action for Equality

United States Military Academy **Assistant Professor or** **Senior Associate Professor of** **Computer Science**

THE UNITED STATES MILITARY ACADEMY SEEKS AN ASSISTANT PROFESSOR OR SENIOR ASSOCIATE PROFESSOR OF COMPUTER SCIENCE

The United States Military Academy (USMA) seeks applicants for an Assistant Professor or senior Associate Professor of Computer Science. USMA, the oldest engineering school in the nation, is situated 50 miles north of New York City in the picturesque Hudson Valley. We offer a Tier I undergraduate program devoted to educating leaders of character committed to the values of Duty, Honor, and Country. USMA's rigorous programs, quality faculty, and student-to-faculty ratio of 12:1 attract a student body of 4,500 of the nation's best. The Department of Electrical Engineering and Computer Science includes a team of 65 civilian and military faculty members dedicated foremost to teaching and student development. Faculty and students have excellent opportunities for research, supported by a ubiquitous campus computing environment and laboratory facilities that are among the best in the nation. The department offers ABET accredited majors in Computer Science, Information Technology and Electrical Engineering, two minors, and two required freshman and junior level courses for all students. Between 50 and 80 majors graduate from our programs each year. Applicants must be senior educators holding an advanced degree (Ph.D.) in Computer Science or closely related field. Extensive teaching experience in mainstream computer science is desired; relevant industry experience is beneficial. We are looking for well-rounded professionals capable of designing, teaching, and overseeing a variety of courses and programs, conducting research in mainstream computer science, serving as mentors to junior faculty members, supporting students in and outside the classroom, and contributing to department and academy governance. USMA offers renewable contracts in lieu of tenure. Salaries are competitive. Applications will be accepted until the position is filled. Applications and questions may be directed to COL Barry L. Shoop, United States Military Academy, Department of Electrical Engineering and Computer Science, MADN-EECS, West Point, New York, 10996-1787, (845) 938-6980, barry.shoop@usma.edu. The United States Military Academy is an Equal Opportunity Employer.

United States Naval Academy **Computer Science Department** **Assistant Professor**

The U.S. Naval Academy's Computer Science Department invites applications for one or more tenure track positions at the rank of Assistant Professor. These positions are anticipated to begin in the Fall of 2012. A Ph.D. in Computer Science or closely related field is required.

Applicants will be considered from all areas of Computer Science, and must have a dedication to

teaching, broad teaching interests, and a strong research program. Applicants with research interests in computer security, network security, and information assurance are especially encouraged to apply.

The Computer Science Department offers majors in Computer Science and Information Technology. All faculty teach courses in both majors. We currently have 85 CS majors, 100 IT majors and a faculty of 21. The department is housed in a state of the art building overlooking the scenic Severn River. Our spaces provide outstanding office, laboratory, and research facilities for both students and faculty, including specialized labs for robotics, networking, and information assurance in addition to three micro-computing labs and two high performance computing labs. In addition to computer science and information technology courses for the majors, we also teach a course on computer security to the entire freshman class.

The Naval Academy is an undergraduate institution located in historic downtown Annapolis, Maryland on the Chesapeake Bay. Roughly half of its faculty are tenured or tenure track civilian professors with Ph.D.s who balance teaching excellence with internationally recognized research programs. The remaining faculty are active duty military officers with Masters or Doctoral degrees. Each year the academy graduates roughly 1000 undergraduate students with majors in the sciences, engineering, and humanities. More information about the department and the Academy can be found at <http://www.usna.edu/cs/> and <http://www.usna.edu/>.

Applicants should send a cover letter, teaching and research statements, curriculum vitae, and arrange for three letters of recommendation that address both teaching and research abilities to be sent to cssearch@usna.edu.

Review of applications will begin in December, continuing until the position is filled.

The United States Naval Academy is an Affirmative Action/Equal Opportunity Employer. This agency provides reasonable accommodations to applicants with disabilities. This position is subject to the availability of funds.

University of California, Riverside **Tenure-track Faculty Positions**

The Department of Computer Science and Engineering, University of California, Riverside invites applications for **tenure-track** faculty positions beginning in the 2012/2013 academic year with research interests in **Systems**. Candidates in Operating/Distributed Systems, Cyber-security, Cyber-physical systems and High performance computing are especially encouraged to apply. Exceptional candidates in all areas will be considered. Positions require a Ph.D. in Computer Science (or in a closely related field) at the time of employment. Junior candidates must show outstanding research, teaching and graduate student mentorship potential. Senior candidates with exceptional research, teaching and graduate student mentorship records may be considered. Salary level will be competitive and commensurate with qualifications and experience. Details and application materials can be found at www.engr.ucr.edu/facultysearch. Full consideration will be given to applications received by January 6, 2012. Applications will continue to be received

until the positions are filled. For inquiries and questions, please contact us at search@cs.ucr.edu. EEO/AA employer.

University of California, San Diego **Department of Computer Science and Engineering** **Faculty Positions**

The UCSD Department of Computer Science and Engineering invites applications at all levels (tenure-track/tenured) in all areas of computer science and engineering, with particular interest in the areas of human-computer interaction, robotics, software systems for massive parallelization, data intensive systems, quantitative and systems biology, energy-efficient computing systems, and sustainable/green computing. However, exceptional candidates in all areas will be seriously considered. The department is looking for applicants with outstanding research credentials. Successful applicants are expected to lead a vigorous research program and to have a strong commitment to teaching. We encourage candidates to send applications as soon as possible. Faculty applications received by January 16, 2012 will be given full consideration. However, positions remain open until filled. To apply, candidates should follow instructions posted at the website <https://csefacapp.ucsd.edu/applicant>. UCSD is an Equal Opportunity/Affirmative Action Employer with a strong institutional commitment to the achievement of excellence and diversity among its faculty and staff: <http://diversity.ucsd.edu>

University of Colorado, Boulder **Assistant Professor**

The Department of Computer Science at the University of Colorado Boulder (CU) seeks outstanding candidates in human-centered computing (HCC) at the level of Assistant Professor, although senior candidates at higher ranks may be considered. The position will help shape HCC initiatives at CU. Candidates working in human-computer interaction, computer-supported cooperative work, information visualization, and/or social computing are invited to apply. Candidates should have an orientation to computer science as their primary faculty research and teaching home. The University of Colorado Boulder is an Equal Opportunity Employer. Applications must be submitted on-line at www.jobsatcu.com/applicants/Central?quickFind=66808

University of Connecticut **Computer Science & Engineering Department** **Assistant/Associate Professor**

The University of Connecticut invites applications for a nine-month, tenure-track faculty position for teaching, research, service and outreach in the Computer Science & Engineering Department in the School of Engineering. Applicants must hold a Ph.D. in computer science or a related discipline by the time of appointment and possess an excellent record of research accomplishments. While applicants in all areas will be considered, the department seeks to strengthen its coverage of core areas of computing such as

Systems, Algorithms, and High Performance Computing. It is preferred that candidates possess the ability to contribute through research, teaching and/or public engagement to the diversity and excellence of the learning experience. The department will give preference to applicants with relevant teaching experience. For a full job description and application instructions please visit www.jobs.uconn.edu. The University of Connecticut is an EEO/AA Employer. (Search # 2012121)

University of Houston **Department of Computer Science** **Faculty Position**

The Department of Computer Science at the University of Houston (www.cs.uh.edu) invites applications for one tenure-track faculty position starting in August 2012. The current search is focused on the following areas: software design/engineering, computer security, human-computer interaction, and theory. However, outstanding candidates in other areas will also be considered. Candidates should hold a Ph.D. in Computer Science, Computer Engineering, or a closely related field. Candidates at all levels are invited to apply. The full advertisement is available at <http://www.cs.uh.edu/faculty-search>.

The University of Houston is an equal opportunity/affirmative action employer. Minorities, women, veterans, and persons with disabilities are encouraged to apply.

University of Illinois at Chicago **Faculty Position** **Department of Computer Science, UIC**

The Computer Science Department at the University of Illinois at Chicago invites applications for tenure-track positions at the rank of Assistant Professor (exceptional candidates at other ranks may also be considered). Candidates in the following areas are especially encouraged to apply: Computer Security, Software Engineering, Machine Learning and Computer Systems (Mobile/Ubiquitous computing and other experimental subareas).

The University of Illinois at Chicago (UIC) ranks among the nation's top 50 universities in federal research funding. It is the largest research university in the Chicago area, and is one of the most diverse universities in the country. The Computer Science department has 27 tenure-track faculty representing major areas of computer science, and offers BS, MS and PhD degrees. Two of our faculty members are ACM Fellows and eight members are recipients of NSF CAREER awards. Our annual research funding has averaged \$6.5M over the last five years and includes major funding from NSF, DARPA, DoD and NASA, including two NSF IGERT awards, nine Trustworthy Computing awards and several other research and instrumentation grants; awards from state agencies such as the Illinois Department of Transportation, and from companies such as Google, Yahoo! and Motorola. Our department is home to many pioneering and discipline-defining efforts in the areas of virtual reality (CAVE), software engineering (Petri Nets, Model Checking), Data Management and Mining, and Computational Transportation. We have growing research programs

in areas such as computational biology, learning technologies, mobile and distributed systems, and security and privacy. At UIC, there are plenty of opportunities for interdisciplinary work—UIC houses the largest medical school in the country, and our faculty are engaged with several cross-departmental collaborations with faculty from health sciences, social sciences and humanities, urban planning and the business school.

Chicago is the third most populous city in the USA. Located by the shore of Lake Michigan, the city offers an outstanding array of cultural and culinary experiences. As the birthplace of the modern skyscraper, Chicago boasts one of the world's tallest and densest skylines, combined with an extensive system of parks and public transit. Its primary airport is the second busiest in the world, with frequent non-stop flights to virtually anywhere. Yet the cost of living, whether in an 85th floor condominium downtown or on a tree-lined street in one of the nation's finest school districts, is surprisingly low.

Applications must be submitted at <https://jobs.uic.edu/>. Please include a resume, teaching and research statements, and names and addresses of at least three references in the online application. Applicants needing additional information may contact the Faculty Search Chair at search@cs.uic.edu.

Application processing will commence on Nov 15th. We will continue to accept and process applications after that date until all the positions are filled. The University of Illinois at Chicago is an Affirmative Action/Equal Opportunity Employer.

University of North Carolina at Charlotte

Executive Director

Technology Solutions Office (TSO)

College of Computing and Informatics (CCI)

The University of North Carolina at Charlotte seeks an **Executive Director** to lead the Technology Solutions Office (TSO) within the College of Computing and Informatics (CCI). The TSO provides technology solutions to support research, education, and operations for the CCI, and other clients in academia, industry, government and nonprofits. The Director should be a highly skilled technology expert, strategist, team-builder and leader able to leverage new and emerging technologies to create unique solutions. The Director should be a skilled customer-service provider, able to develop and serve a diverse client base. Responsibilities include: provide vision, expertise, leadership, and coordination to support teaching and research and to develop a leading R&D consultancy group; hire, train and manage personnel, including student workers; collaborate with faculty to identify needs; liaise with University IT Services to ensure integrated support; develop and manage client relationships, contracts and performance; Serve as the CCI Data Security Officer. The successful candidate will hold bachelors and masters degrees with at least one degree in a computing or technology related field. Preference will be given to candidates with at least seven years experience in research, development, and/or technical support; with at least three of those years in R&D technology project management, customer service, and/or marketing. A detailed job description is at <http://cci.uncc.edu/?q=content/employment-opportunities>. Apply online at <https://jobs.uncc.edu> (Position # 4469). Review of applications will begin November 14, 2011 and continue until the position is filled. The University of North Carolina at Charlotte is an EOE/AA employer and an NSF ADVANCE Institution.

University of Pittsburgh School of Information Sciences Three Tenure-stream Faculty Positions

The School of Information Sciences (<http://www.ischool.pitt.edu>) at the University of Pittsburgh is seeking candidates for three tenure-stream faculty positions to start in the fall term of 2012. The positions are open to both individuals seeking their first academic position at the assistant professor level as well as more senior faculty who may already hold tenure at another institution. Appointments will be made at a rank commensurate with qualifications. Open positions include:

► Assistant Professor – Software engineering / systems

► Assistant/Associate/Full Professor – Library and information science

► Assistant/Associate Professor – Data-driven scholarship

This top-ranked information school (iSchool) offers degree programs in Information Science & Technology, Library & Information Science, and Telecommunications & Networking. Candidates are sought who have research and teaching interests in alignment with the School's signature strengths.

For a complete description, please visit www.ischool.pitt.edu/news/facultyopenings.php.

The University of Pittsburgh is an Equal Opportunity, Affirmative Action employer and strongly encourages women and candidates from under-represented minorities to apply.

University of Puerto Rico, Rio Piedras Tenure-track positions in Comp. Science

The Computer Science Department at the University of Puerto Rico, Rio Piedras invites applications for tenure-track positions in Computer Science. http://ccom.uprrp.edu/job_openings.php

University of Rochester Tenure Track Position in Computer Science

The University of Rochester Department of Computer Science seeks applicants for a tenure-track faculty position. We are particularly interested in researchers in human-computer interaction and machine learning, but will consider all outstanding candidates. See <http://www.cs.rochester.edu/recruit> for details. UR is an Equal Opportunity Employer.

University of Science and Technology of China School of Computer Science and Technology Faculty Positions

The School of Computer Science and Technology at University of Science and Technology of China (USTC) invites applications for tenure-track or tenured positions at all levels. Research areas of particular interest include programming languages and

compilers, formal verification, robotics, computational intelligence, machine learning, data mining, computer architectures, parallel and high-performance computing, network and distributed systems. For more information about the positions, please see http://en.cs.ustc.edu.cn/join_us.

University of Tennessee Assistant Professor

Tenure-track position. Teach a variety of computer science courses typical of a 4-year computer science program. Responsibilities include teaching, advising, curriculum development, assessment activities, university service, and scholarly activity. Effective pedagogic and communication skills are required.

The University of Tennessee is an EEO/AA/ Title IX/ Section 504/ADA/ADEA institution in the provision of its education and employment programs and services. All qualified applicants will receive equal consideration for employment without regard to race, color, national origin, religion, sex, pregnancy, marital status, sexual orientation, gender identity, age, physical or mental disability, or covered veteran status.

University of West Georgia Computer Science Assistant/Associate Professor- Tenure Track

Applications and nominations are invited for a tenure track faculty position for Fall 2012. We seek an energetic and enthusiastic individual who loves teaching and shares our dedication to excellence. For details, visit jobs.cs.westga.edu.

Wichita State University Assistant Professor

The Department of Electrical Engineering and Computer Science at Wichita State University has multiple open tenure-eligible faculty positions at the assistant professor level in Data Mining/Machine Learning/Artificial Intelligence, Computer Networking, Controls/Robotics, and Electronics. Duties and responsibilities of all positions include teaching undergraduate and graduate courses, advising undergraduate students, supervising MS and PhD students in their theses and dissertations, obtaining research funding, conducting an active research program, publishing the results of research, actively participating in professional societies, and service to the department, college and university. Complete information can be found on our website, www.eecs.wichita.edu.

To ensure full consideration, the complete application package must be submitted online at wichita.edu/wsujobs by January 17, 2012. Applications will be continuously reviewed after that date until determinations are made with regard to filling the positions. Offers of employment are contingent upon completion of a satisfactory criminal background check as required by Board of Regents policy.

Questions only (not applications) can be directed to the search chair, Ward Jewell, wardj@ieee.org.

Wichita State University is an equal opportunity and affirmative action employer.

[CONTINUED FROM P. 136] solemn, like he has a secret, laying out a sales pitch in his quiet synthesized voice. I guess it's how Google makes its money. Whitey didn't cost me anything, but he sure annoys me sometimes.

Also, the constant searching can be pretty distracting.

Last month, my ex-girlfriend was over and we were getting along for once. At a certain point, I reached into my bedside drawer only to find old Whitey standing there pawing through my junk, taking inventory while me and my girl were two feet away in a delicate position. Trust me, it was a real mood killer.

Google says it's a feature.

In this one old ad, a guy even Googles for his car keys. Bull. I've never Googled my keys once. I'd never even bother to ask Whitey. It's way quicker to just look around. I've got eyes and ears, for God's sake.

That's why I bought a Near Cloud from Intel.

It's a smooth beige box that sits in my closet. Full of processors. It's got routines to help Whitey walk and pick

The difference is my Near Cloud isn't talking to anybody out there.

things up and make the bed and stuff. The difference is my Near Cloud isn't talking to anybody out there. It's not reporting the location of every product in my house to a multinational company that desperately wants me to buy, buy, buy.

Anyway, enough complaining.

I pry off Whitey's back casing with a nail file. Pull out his little SIM card. Plug in the cord from my phone and start running the jailbreak program.

Simple as that.

* * *

Whitey is gone. Sort of. Things are different on the Near Cloud.

If I wasn't so lazy, I'd have saved up and got a real Japanese-made local domestic robot years ago. Like the Honda Pro-Asimo. Built-in personality. Onboard computation. No Far Cloud. No Near Cloud. No ads. No spying. And those things are quality, made of real metal and carbon fiber.

Maybe someday.

In the meantime, Whitey has gone quiet. Operating on the Near Cloud, he doesn't hit me up with ads anymore. Doesn't do that slow ceremonial approach right before trying to sell me a new kind of toothpaste.

Weird to say, but I kind of miss it.

This morning, I woke up and realized something. I'm getting pretty tired of Jimmy Dean sausages. With the old Whitey gone, I guess I'm starting to wonder what other kinds of breakfast sausage are out there. 

Daniel H. Wilson is the *New York Times* bestselling author of *Robocalypse*. Follow him on Twitter at [danielwilsonpdx](http://twitter.com/danielwilsonpdx) or visit his Web site <http://www.danielhwilson.com>.

© 2012 ACM 0001-0782/12/01 \$10.00

Announcing ACM's Newly Improved Career & Job Center!

Are you looking for your next IT job? Do you need Career Advice?
Visit ACM's newly enhanced career resource at:
<http://jobs.acm.org>



The ACM Career & Job Center offers ACM members a host of benefits including:

- A highly targeted focus on job opportunities in the computing industry
- Access to hundreds of corporate job postings
- Resume posting keeping you connected to the employment market while letting you maintain full control over your confidential information
- An advanced Job Alert system that notifies you of new opportunities matching your criteria
- Career coaching and guidance from trained experts dedicated to your success
- A content library of the best career articles compiled from hundreds of sources, and much more!

The **ACM Career & Job Center** is the perfect place to
begin searching for your next employment opportunity!

Visit today at <http://jobs.acm.org>



Association for
Computing Machinery

Advancing Computing as a Science & Profession

Future Tense, one of the revolving features on this page, presents stories and essays from the intersection of computational science and technological speculation, their boundaries limited only by our ability to imagine what will and could be.

DOI:10.1145/2063176.2063205

Daniel H. Wilson

Future Tense The Near Cloud

Wish I never pulled the plug...

“JOHNSONVILLE BREAKFAST SAUSAGE has the variety to satisfy all your taste preference needs,” says a small plastic face, its black camera eyes peeking over my bedspread.

Groaning, I blindly shove the robot away. The child-size humanoid stumbles and falls softly to the carpet, then rolls clumsily onto its chest, unharmed. There isn’t much to this robot. Just a bunch of actuators and sensors and a radio. It’s durable and light and always getting smarter because it doesn’t do its own thinking. Its robot brain is out there somewhere. Nowhere.

In the Google Cloud.


Pushing it down isn’t the worst I’ve done to it, not by far. But, as usual, its bleached white casing provides protection. It’s like a little white knight in dull plastic armor.

That’s why I call him Whitey.

“Shall I run to the store for a pack of Vermont Maple Breakfast Sausage Links?” he intones, speaker muffled on my freshly vacuumed beige carpet. “Or perhaps Wisconsin Cheddar Breakfast—”

My pillow lands on Whitey, and he cuts it out. Every morning he hits me

**I’m not scared
of the government;
it’s the ads that
are killing me.**



with the advertising. Always for the opposite of whatever he’s found rummaging around in my refrigerator, bathroom, closet.

Johnsonville sausage, really? Whitey knows I’m a Jimmy Dean man.

I’ve finally had it. The Google corporation gave me Whitey five years ago. Today, I’m gonna jailbreak him and set us both free.

The robot is sitting on my coffee table between empty soda cans and old magazines. My smartphone has the jailbreak software loaded and ready to go. I’ve decided to keep my data here in the house from now on. It’s time to get Whitey’s head out of the Cloud.

And not just because of the spying.

Whitey isn’t a narc. Sure, his re-

cords could be subpoenaed. But the government is welcome to peek into my crummy apartment and watch me play video games while Whitey folds towels. The ads are right; if you’ve got nothing to hide, then you’ve got nothing to worry about. Right?

I’m not scared of the government; it’s the ads that are killing me.

Whitey has this compulsive need to peek into every drawer, pick things up and stare at them until he scans the brand name or UPC code or whatever he’s been programmed to look for. He’s built a brand catalogue of all the crap in my house. And trust me, most of it is crap.

Then he blurts out these ads for competing products, walking up to me real slow and [CONTINUED ON P. 135]



Distinguished Speakers Program

talks by and with technology leaders and innovators

Chapters • Colleges and Universities • Corporations • Agencies • Event Planners

A great speaker can make the difference between a good event and a WOW event!

The Association for Computing Machinery (ACM), the world's largest educational and scientific computing society, now provides colleges and universities, corporations, event and conference planners, and agencies – in addition to ACM local Chapters – with direct access to top technology leaders and innovators from nearly every sector of the computing industry.

Book the speaker for your next event through the ACM Distinguished Speaker Program (DSP) and deliver compelling and insightful content to your audience. **ACM will cover the cost of transportation for the speaker to travel to your event.** Our program features renowned thought leaders in academia, industry and government speaking about the most important topics in the computing and IT world today. Our booking process is simple and convenient. Please visit us at: www.dsp.acm.org. If you have questions, please send them to acmdsp@acm.org.

The ACM Distinguished Speaker Program is an excellent solution for:

Corporations Educate your technical staff, ramp up the knowledge of your team, and give your employees the opportunity to have their questions answered by experts in their field.

Colleges and Universities Expand the knowledge base of your students with exciting lectures and the chance to engage with a computing professional in their desired field of expertise.

Event and Conference Planners Use the ACM DSP to help find compelling speakers for your next conference and reduce your costs in the process.

ACM Local Chapters Boost attendance at your meetings with live talks by DSP speakers and keep your chapter members informed of the latest industry findings.

Captivating Speakers from Exceptional Companies, Colleges and Universities

DSP speakers represent a broad range of companies, colleges and universities, including:

IBM
Microsoft
BBN Technologies
Raytheon

Sony Pictures
McGill University
Tsinghua University
UCLA

Georgia Tech
Carnegie Mellon University
Stanford University
University of Pennsylvania

University of British Columbia
Siemens Information Systems Bangalore
Lawrence Livermore National Laboratory
National Institute of Standards and Technology

Topics for Every Interest

Over 300 lectures are available from nearly 100 different speakers with topics covering:

Software
Cloud and Delivery Methods
Emerging Technologies
Engineering

Web Topics
Computer Systems
Open Source
Game Development

Career-Related Topics
Science and Computing
Artificial Intelligence
Mobile Computing

Computer Graphics, Visualization
and Interactive Techniques
High Performance Computing
Human Computer Interaction

Exceptional Quality is Our Standard

The same ACM you know from our world-class digital library, magazines and journals is now putting the affordable and flexible Distinguished Speaker Program within reach of the computing community.

Microsoft®
Research
The DSP is sponsored
in part by Microsoft Europe



**Association for
Computing Machinery**

Advancing Computing as a Science & Profession

it's the experience

research

practice

experience

community

digital arts

DESIGN

games & entertainment

ENGINEERING

hci for kids

MANAGEMENT

health

USER EXPERIENCE

sustainability

keynotes

Margaret Stewart, YouTube

Hugh Herr, MIT Media Lab



CHI 2012
it's the experience!
AUSTIN • TEXAS
MAY 5-10, 2012

**30TH CONFERENCE ON
COMPUTER-HUMAN INTERACTION**

<http://chi2012.acm.org>

