

COMMUNICATIONS

OF THE

ACM

CACM.ACM.ORG

11/2012 VOL.55 NO.11



Association for
Computing Machinery





27th

IEEE INTERNATIONAL PARALLEL & DISTRIBUTED PROCESSING SYMPOSIUM

www.ipdps.org



*Hyatt Regency Cambridge
Boston, Massachusetts USA
20 - 24 May 2013*

IPDPS 2013 will be held along the banks of the Charles River at the Hyatt Regency Cambridge with splendid views of the Boston skyline. The Boston-Cambridge venue offers an exciting multicultural setting in the midst of world renowned educational institutions and landmarks of America's founding, all navigable by foot and public transit. The IPDPS event offers a rich program of workshops, a PhD Forum, and the technical program of contributed papers, invited speakers, and tutorials.

GENERAL CO-CHAIRS

Martin Herbordt (Boston University, USA)
Charles C. (Chip) Weems (University of
Massachusetts at Amherst, USA)

PROGRAM CHAIR

Sanjay Ranka (University of Florida, USA)

WORKSHOPS CHAIR

Ümit V. Çatalyürek (Ohio State University, USA)

LOCAL PARTICIPATION & COMMERCIAL CHAIR

Ayşe Coskun (Boston University, USA)

IPDPS extends a special invitation to Boston area companies to consider participation in IPDPS 2013 to share in the benefits of associating with an international community of top researchers and practitioners in fields related to parallel processing and distributed computing. Visit the IPDPS Web for details.

PhD FORUM CO-CHAIRS

Luc Bougé (ENS Cachan, France)
Bo Hong (Georgia Institute of Technology, USA)

The PhD Forum offers graduate students an opportunity to present a research poster describing their ongoing dissertation to the entire IPDPS conference audience and to submit a short paper for publication. The deadline for submissions is January 14, 2013. See details at www.ipdps.org.

IPDPS 2013 WORKSHOPS

Held on the first and last day of the conference, workshops are a major part of the IPDPS week-long family of events, providing the IPDPS community an opportunity to explore special topics and present work that is more preliminary or cutting-edge than the more mature research presented in the main symposium. Each workshop has its own requirements and schedule for submissions and all are linked from the workshops page on the IPDPS Website. Note that the submission date for most workshops falls sometime after author notification date of December 23rd for regular conference papers.

HCW	Heterogeneity in Computing
RAW	Reconfigurable Architectures
HIPS	High-Level Parallel Programming Models & Supportive Environments
NIDISC	Nature Inspired Distributed Computing
HiCOMB	High Performance Computational Biology
APDCM	Advances in Parallel and Distributed Computing Models
CASS	Communication Architecture for Scalable Systems
HPPAC	High-Performance, Power-Aware Computing
HPGC	High-Performance Grid and Cloud Computing
ASHES	Accelerators and Hybrid Exascale Systems
STDN	Security and Trust of Distributed Networking Systems
EduPar	NSF/TCPP -Parallel and Distributed Computing Education
PDSEC	Parallel and Distributed Scientific and Engineering Computing
DPDNS	Dependable Parallel, Distributed and Network-Centric Systems
MTAAP	Multi-Threaded Architectures and Applications
LSPP	Large-Scale Parallel Processing
PCO	Parallel Computing and Optimization
ParLearning	Parallel & Distributed Computing for Machine Learning & Inference Problems
HPDIC	High Performance Data Intensive Computing
CloudFlow	Workflow Models, Systems, Services and Applications in the Cloud
JSSPP	Job Scheduling Strategies for Parallel Processing
LSDSS	Large Scale Distributed Service-oriented Systems
PLC	Multicore and GPU Programming Models, Languages and Compilers
ViPES	Virtual Prototyping of Parallel and Embedded Systems

Sponsored by

IEEE  computer society
Technical Committee on Parallel Processing

In cooperation with

 ACM SIGARCH
 IEEE Computer Society Technical Committee on Computer Architecture
 IEEE Computer Society Technical Committee on Distributed Processing

Visit the new Website for *ACM Inroads*

The magazine for computing
educators worldwide



<http://inroads.acm.org>

Paving the way toward excellence in computing education

Departments

- 5 **Editor's Letter**
Will MOOCs Destroy Academia?
By Moshe Y. Vardi
-
- 7 **From the President**
Why Is Accessibility So Hard?
By Vinton G. Cerf
-
- 8 **Letters to the Editor**
**When Predicting,
Start With Humility**
-
- 10 **BLOG@CACM**
New Opportunities for New SQL
Michael Stonebraker expects
a substantial increase in the number
of New SQL engines using a variety
of architectures in the near future.
-
- 12 **Major Update to ACM's
Computing Classification System**
By Bernard Rous
-
- 27 **Calendar**
-
- 110 **Careers**

Last Byte

- 120 **Puzzled**
Weighed in the Balance
By Peter Winkler

News



- 13 **Software on Mars**
With the AEGIS system, the Mars
Exploration Rovers can
autonomously select, capture, and
analyze images using onboard logic.
By Gregory Goth
-
- 16 **Control Without Controllers**
Disney's Touché project could
transform every conductive surface
into a touch-control surface.
By Tom Geller
-
- 19 **On the Digital Trail**
Forensics experts increasingly
use data from digital devices
to solve crimes, fight lawsuits,
and unravel accidents.
By Samuel Greengard

Viewpoints

- 22 **Privacy and Security**
**The Research Value
of Publishing Attacks**
Security research can be improved
by more effectively sharing
what is learned from attacks on
information systems.
By David Basin and Srdjan Capkun
-
- 25 **Legally Speaking**
**Oracle v. Google:
Are APIs Copyrightable?**
Assessing the first phase of the
trial based on claims that Google's
Android platform infringes Oracle's
Java-related copyrights and patents.
By Pamela Samuelson
-
- 28 **Economic and Business Dimensions**
**Decentralization versus
Centralization in IT Governance**
It's not as simple as you might think.
By Kristina McElheran
-
- 31 **Education**
**Learning to Teach Computer
Science: The Need
for a Methods Course**
A multipronged approach to
preparing computer science
teachers is critical to success.
By Aman Yadav and John T. Korb
-
- 34 **Computing Ethics**
**Societal Implications of
the Emerging Smart Grid**
Seeking solutions to concerns
that go beyond the engineering
of the smart grid.
By Timothy Kostyk and Joseph Herkert
-
- 37 **Viewpoint**
Keeping Technology Promises
Considering new models for
educational technology and methods.
By Richard A. DeMillo



Association for Computing Machinery
Advancing Computing as a Science & Profession


Practice



48

- 40 **Resilience Engineering: Learning to Embrace Failure**
A discussion with Jesse Robbins, Kripa Krishnan, John Allspaw, and Tom Limoncelli.
- 48 **Weathering the Unexpected**
Failures happen, and resilience drills help organizations prepare for them.
By Kripa Krishnan
- 53 **Disks from the Perspective of a File System**
Disks lie. And the controllers that run them are partners in crime.
By Marshall Kirk McKusick

Q Articles' development led by **acmqueue**
queue.acm.org



About the Cover:
Affetto, the baby robot created at Osaka University, is one of the latest examples of bio-inspired soft robotic technologies designed to study the principles underlying human and animal behavior and how to transfer them to robots. Affetto will be used to examine child development issues.
Illustration by

Brian Greenberg/Andrij Borys Associates; based on photographs courtesy of Asada Laboratory, Osaka University.


Contributed Articles



56

- 56 **Functional Encryption: A New Vision for Public-Key Cryptography**
Decryption keys allow users to learn a specific function of the encrypted data and nothing else.
By Dan Boneh, Amit Sahai, and Brent Waters
- 65 **Cheminformatics**
Open-source chemistry software and molecular databases broaden the research horizons of drug discovery.
By Joerg Kurt Wegner, Aaron Sterling, Rajarshi Guha, Andreas Bender, Jean-Loup Faulon, Janna Hastings, Noel O'Boyle, John Overington, Herman van Vlijmen, and Egon Willighagen

Review Articles



76

- 76 **The Challenges Ahead for Bio-Inspired 'Soft' Robotics**
Soft materials may enable the automation of tasks beyond the capacities of current robotics.
By Rolf Pfeifer, Max Lungarella, and Fumiya Iida

Research Highlights

- 89 **Technical Perspective**
Open Platforms for Computational Photography
By Richard Szeliski
- 90 **The Frankencamera: An Experimental Platform for Computational Photography**
By Andrew Adams, David E. Jacobs, Jennifer Dolson, Marius Tico, Kari Pulli, Eino-Ville Talvala, Boris Ajudin, Daniel Vaquero, Hendrik P.A. Lensch, Mark Horowitz, Sung Hee Park, Natasha Gelfand, Jongmin Baek, Wojciech Matusik, and Marc Levoy
- 99 **Technical Perspective**
The Realities of Home Broadband
By Henning Schulzrinne
- 100 **Measuring Home Broadband Performance**
By S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè



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

Vinton G. Cerf

Vice-President

Alexander L. Wolf

Secretary/Treasurer

Vicki L. Hanson

Past President

Alain Chesnais

Chair, SGB Board

Erik Altman

Co-Chairs, Publications Board

Ronald Boisvert and Jack Davidson

Members-at-Large

Eric Allman; Ricardo Baeza-Yates;

Radia Perlman; Mary Lou Soffa;

Eugene Spafford

SGB Council Representatives

Brent Hailpern; Joseph Konstan;

Andrew Sears

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

Margaret Gray

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.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) 626-0686

F (212) 869-0481

Director of Media Sales

Jennifer Ruzicka

jen.ruzicka@hq.acm.org

Media Kit acmmidasales.com

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

elc@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; 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; 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 *queue*.

CONTRIBUTED ARTICLES

Co-Chairs

Al Aho and Georg Gottlob

Board Members

Robert Austin; Elisa Bertino; Gilles Brassard;

Kim Bruce; Alan Bundy; Peter Buneman;

Erran Carmel; Andrew Chien;

Peter Druschel; Carlo Ghezzi; Carl Gutwin;

James Larus; Igor Markov; Gail C. Murphy;

Shree Nayyar; Bernhard Nebel; Lionel M. Ni;

Sriram Rajamani; Marie-Christine Rousset;

Avi Rubin; Krishan Sabnani;

Fred B. Schneider; Abigail Sellen;

Ron Shamir; Yoav Shoham; Marc Snir;

Larry Snyder; 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; Sanjeev Arora; Dan Boneh;

Andrei Broder; Stuart K. Card; Jon Crowcroft;

Alon Halevy; Monika Henzinger;

Maurice Herlihy; Norm Jouppi;

Andrew B. Kahng; Xavier Leroy;

Mendel Rosenblum; Ronitt Rubinfeld;

David Salesin; Guy Steele, Jr.; David Wagner;

Alexander L. Wolf; Margaret H. Wright

WEB

Chair

James Landay

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@cacm.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



Printed in the U.S.A.



DOI:10.1145/2366316.2366317

Moshe Y. Vardi

Will MOOCs Destroy Academia?

“Thy destroyers and they that made thee waste shall go forth of thee,” wrote the prophet Isaiah. This phrase has been popping into my mind as I have been following

the recent raging discussions over the topic of MOOCs.

For those readers who paid no attention to recent developments, a MOOC is *massive open online course*; it is a tuition-free course taught over the Web to a large number of students. While online education has a long history, the current wave started in the fall of 2011 when about 450,000 students signed up for three computer-science courses offered by Stanford University. Since then, MOOCs have become the hottest topic of discussion in higher education in the U.S. Within months of the Stanford experiments, several start-up companies debuted, including one that immodestly claims to be “the first elite American university to be launched in a century.” Many leading U.S. universities now offer MOOCs, either on their own or in partnership with some of these companies, even though no business model has emerged for MOOC-based education. Some describe the current environment as “MOOC panic” or “MOOC mania.” John Hennessy, Stanford’s president, describes the phenomenon as a “tsunami.”

Early rhetoric about the educational value of MOOCs was quite lofty, talking about the goal of reaching the quality of individual tutoring, but it is difficult to reconcile such rhetoric with massiveness as an essential feature of MOOCs. A more honest comment from one of the early MOOC pioneers was: “We were tired of delivering the same lectures year after year, often to a half-empty classroom because our classes were being videotaped.” In fact, the absence of serious pedagogy in MOOCs is rather

striking, their essential feature being short, unsophisticated video chunks, interleaved with online quizzes, and accompanied by social networking.

The bitter truth, however, is that academic pedagogy has never been very good. It is well established that a professorial soliloquy is an ineffective way of teaching. We do know what works and what does not work when it comes to teaching. Much has been written in the last few years about “active learning,” “peer learning,” “flipping the lecture,” and the like, yet much of academic teaching still consists of professors monologuing to large classes. We could undoubtedly improve our teaching, but MOOCs are not the answer to our pedagogical shortcomings.

To understand the real significance of MOOCs you must consider the financial situation in which U.S. colleges and universities have found themselves in the aftermath of the Great Recession. The financial crisis dealt a severe blow to U.S. higher education. Private institutions saw their endowments take significant hits, while public institutions saw state support, which was already shrinking, decline even faster. While outstanding student debt has exceeded the \$1T mark, students are facing a highly constrained job market, challenging their ability to repay their debt. After years of college tuition escalating faster than inflation, the very value of college education is being seriously questioned; an Internet entrepreneur is even offering a skip-college fellowship. In this environment, the prospect of higher edu-

cation at a dramatically reduced cost is simply irresistible.

It is clear, therefore, that the enormous buzz about MOOCs is not due to the technology’s intrinsic educational value, but due to the seductive possibilities of lower costs. The oft-repeated phrase is “technology disruption.” This is the context for the dismissal (and later reinstatement) last summer of Theresa A. Sullivan, University of Virginia’s president, because she was not moving fast enough with online education. The bigger picture is of education as a large sector of the U.S. economy (over \$1T) that has so far not been impacted much by information technology. From the point of view of Silicon Valley, “higher education is a particularly fat target right now.” MOOCs may be the battering ram of this attack.

My fear is the financial pressures will dominate educational consideration. In his recent book *What Are Universities For?*, Stefan Collini, a Cambridge don, describes universities as “perhaps the single most important institutional medium for conserving, understanding, extending and handing on to subsequent generations the intellectual, scientific and artistic heritage of mankind...we are merely custodians for the present generation of a complex intellectual inheritance which we did not create, and which is not ours to destroy.” If I had my wish, I would wave a wand and make MOOCs disappear, but I am afraid that we have let the genie out of the bottle.

Moshe Y. Vardi, EDITOR-IN-CHIEF



Association for
Computing Machinery

Advancing Computing as a Science & Profession

membership application & digital library order form

Priority Code: AD13

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/about/code-of-ethics>

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 _____



Vinton G. Cerf

DOI:10.1145/2366316.2366341

Why Is Accessibility So Hard?

I sometimes think that, of all the disciplines, ours ought to be the most effective at adapting to the varied needs of users, including those that are challenged to interact with computing

systems in one way or another. From low to no vision, deafness or hearing loss to carpal tunnel syndrome and various other physical limitations, we really should be able to configure our software to adapt. And in many cases, some very useful, clever, and general-purpose software adaptations have been achieved. But the problem persists, and it is still not the case that one can hold high expectations of accessible adaptation for a random application that happens to become necessary or, at least, of high interest.

I think I understand some of the problem, but this column is an attempt to begin a dialogue about improving the state of accessibility in our field. This is not only important from the purely ethical perspective, but it is also pragmatic given the demographics of our society and the increasing incidence of need for accessible applications. We are an aging society and we are welcoming home many wounded warriors with the need for assistive response, to mention only two obvious beneficiary groups. One reason this seems to be so hard is that software has unlimited variations and interfaces to applications can take virtually any form. Moreover, we are extending the modalities of interaction to include speech, gestures, mice, touchscreens, other “pointers,” keyboards, and so on. We have Web-based applications that take advantage of a wide range of presentation and interaction choices. Not all applications take into account the need to offer distinct and

configurable user interfaces and even when some or many such adaptations are offered, some work a lot better than others. The other side of this equation is that the users also manifest unlimited variations in their abilities and it seems unlikely that programmers can be fully cognizant of the nuances of each.

Another theme is the proliferation of platforms through which we may interact with computer-based services and applications. It becomes increasingly difficult to design in detail every mode of interaction, including accessibility variations, for every platform and application imaginable. And even if our imaginations were that good, someone is bound to invent a new application requiring assistive responses that we have not thought about before.

One popular tactic has been to try to create general-purpose tools such as screen readers to assist blind users or automatic captions to help deaf users. Another tactic is to “parameterize” the interfaces so users can pick and choose the variations best suited to their needs. In my experience, the range of parameters devised is fairly large and it is easy to get lost in selecting configurations or even anticipating how well they will fit user needs. Still, these tactics seem important for practitioners to apply when appropriate. The challenges strike me as fundamental, given the range of needs and potential interface designs.

This is by no means a new problem. There cannot be much debate that programmers and user interface (UI) and

experience (UX) experts need to think fairly broadly and deeply about potential use cases before settling on an interface design. While the use of libraries intended to confer “accessibility” on arbitrary applications may be helpful, it seems to me that no amount of automatic adapting will make a poorly designed interface accessible. For some of the same reasons that security ought to be “built in” to the initial design, so should accessibility. If UI designers had to try their designs while blindfolded or use their applications with the sound off, they might gain insights into the nuanced demands that accessibility places on good design.

One feature of good interface design is anticipating what the user is likely to need to do next and to prepare for that. A similar notion might inform thinking about accessibility. One is struck by the seemingly impossible challenge faced by blind users and UI designers for them. In the Web-based world, two-dimensional displays, touchscreens, pop-up windows, drop-down menus, color highlighting, and other signals seem utterly out of reach. One must think how a user interface will behave when it is serialized for audible presentation. In addition, consistency of format and audio feedback from screen to screen also seems like a helpful philosophy.

I would like very much to hear from ACM members, SIGs interested in this space, UX design experts, as well as users of accessibility features about their experiences and their ideas.^a Somehow we must find ways to approach this problem with a richer combination of design principles, pragmatic tactics, and artful implementations than we have in hand today. **□**

^a I also recommend ACM's *Transactions on Accessible Computing* as a valuable resource.

Vinton G. Cerf is Vice President and Chief Internet Evangelist at Google Inc. and the president of ACM.

© 2012 ACM 0001-0782/12/11 \$15.00

When Predicting, Start With Humility

IN HIS VIEWPOINT “Don’t Feel Bad If You Can’t Predict the Future,” Peter J. Denning (Sept. 2012) wrote: “Make sure your models are validated and that their assumed recurrences fit the world you are forecasting. Ground your speculations in observable data...” Hmm... Who validates the models? Physics-based models can be validated, in light of, say, their ability to predict/replicate the results of observable phenomena (such as gravity and inertia); experts in the discipline agree that the assumptions, calculations and/or algorithms, and predicted results match what is seen in the real world. On the other hand, social models rely on assumptions about human behavior, both individual and en masse, that cannot be measured or demonstrated and on predictions that can never be more than “face-validated.” That is, “I can’t tell you we got the right answer for the right reason; the best I can say is the predicted behavior corresponds to what is observed in real life x% of the time.”

This inability to validate the quantification of variables is seen in efforts to model military interactions, as well as social, economic, and political phenomena; for example, no version of either the Lanchester model reflecting the relative strengths of a predator/prey pair or of the many “expanded” Lanchester variants is capable of predicting the outcome of the Battle of Rorke’s Drift depicted in the 1964 movie *Zulu* between British troops and Zulu warriors in South Africa in 1879. Tank on tank, we can predict the odds; add human crews, and things get dicey; witness the dramatically uneven results of combat in Operation Desert Storm when a U.S.-led coalition reversed Iraq’s 1991 invasion and nominal annexation of Kuwait. Similarly, one has only to open the newspaper to understand the degree to which we have so far failed to model the American economy sufficiently to suggest effective measures to relieve the ongoing recession. As

Denning pointed out, predicting the future is difficult and fraught with danger. Be humble...

Joseph M. Saur, Atlanta, GA

I very much agree with Peter J. Denning (Sept. 2012) that one should be humble when predicting anything, especially if the prediction depends on some future human action or decision. Unlike atoms and molecules, humans have free will. More than 60 years ago, the economist and philosopher Ludwig von Mises explored this idea in his monumental book *Human Action*. More recently, Walter Isaacson’s biography of Steve Jobs and Malcolm Gladwell’s book *Outliers: The Story of Success* only reinforced the impossibility of predicting human behavior. Historian J. Rufus Fears wrote: “Nations and empires rise and fall not because of anonymous social and economic forces but because of decisions made by individuals” in the description of his course *Wisdom of History*. As for Jobs, predicting even the next five minutes would have been futile. Any given human action or even random event might have yielded a totally different technological (or economic or political) world from the one we have today.

Per Kjeldaas, Monroe, LA

“I can’t tell you we got the right answer for the right reason; the best I can say is the predicted behavior corresponds to what is observed in real life x% of the time.”

Depends Who Pays

Moshe Y. Vardi’s Editor’s Letter “Why ACM?” (Sept. 2012) explored the publishing dichotomy of “reader pays” vs. “writer pays.” When the author of an article pays for publication, the reader gets thinly veiled advertising. Consider all the free publications we get (as “qualified professionals”) but never read because we realize the content is really nothing more than marketing.

Robert Wilkens, Levittown, NY

Pearl’s Grace and Science

For those unaware of the life, interests, and tragic death of the brilliant *Wall Street Journal* reporter Daniel Pearl, a suitable place to begin would be his Wikipedia page (http://en.wikipedia.org/wiki/Daniel_Pearl). Most readers of *Communications* know of his murder in Pakistan in 2002 by anti-U.S. militants, so read especially the “Aftermath” and “Legacy” sections, marveling how the empathetic world reacted, with the most creditable the formation of the Daniel Pearl Foundation (<http://www.danielpearl.org/>), seeking to heal the bitterness and distrust that has led to so much violence since 9/11.

Among those helped by the Foundation are a number of talented young Pakistani journalists. Another aspect of the Foundation’s work from which Pakistan can continue to benefit are the transcripts and videos of the Annual Daniel Pearl Lecture Series at the University of California, Los Angeles and at Stanford University, with links through the Foundation’s ++Web site; subjects covered are open-ended and reflect Pearl’s diverse personal interests.

Almost certainly unknown to Pakistanis and others outside the computer science community are the outstanding scientific and scholarly achievements of Judea Pearl, Daniel’s father and the Foundation’s prime mover, whose smiling face looks out from the cover of *Communications* (June 2012), which included an interview “A Sure Thing” (<http://www.tinyurl.com/94qfqs>) cov-

ering his work in artificial intelligence and his winning the ACM A.M. Turing Award, the equivalent of a Nobel prize in computer science.

The interviewer's final question and Pearl's response should be of interest to all: "Does your research inform your work at the Daniel Pearl Foundation, especially in conducting interfaith dialogues?" to which he replied, "I have an advantage over my dialogue partners in that I'm an atheist, and I understand religious myths are just metaphors, or poetry, for genuine ideas we find difficult to express otherwise. So, yes, you could say I use computer science in my religious dialogues, because I view religion as a communication language. True, it seems futile for people to argue if a person goes to heaven from the East Gate or the West Gate. But, as a computer scientist, you forgive the futility of such debates, because you appreciate the computational role of the gate metaphor."

It is indeed gratifying that the venerated Pakistani public figure Abdul Sattar Edhi is on the Foundation's board. It is through him I would urge the government of Pakistan to invite Pearl as a state guest to speak to our media and on our campuses.

Q. Isa Daudpota, Islamabad, Pakistan

On Proving Continuity of Programs

Swarat Chaudhuri et al.'s article "Continuity and Robustness of Programs" (Aug. 2012) said: "The most basic reason why software systems can violate continuity is conditional branching" but ignored a more fundamental cause, namely that program variables have a limited number of states. Computer representation of real numbers is inexact, and only a finite subset of the integers can be represented exactly. Consequently, in computer arithmetic, equations (such as $(x+y) + (z-y) = x+z$) need not hold and can introduce discontinuity. The article ignored the problem by both declaring, "...our reals are infinite-precision" and not specifying upper and lower bounds for integers. These assumptions are common in mathematics but not valid for computer programs. Some programs can be shown to be continuous by Chaudhuri's method but will exhibit discontinuity when executed.

Models often have properties real mechanisms do not have, and it is possible to verify the correctness of a model of a program even if the actual program will fail.

tinuous behavior when executed.

The article also ignored real problems by proposing metrics that are based on data types alone, saying: "The metric over arrays of reals or integers of the same length is the L^∞ -norm: $d(A_1, A_2) = \max_i \{|A_1[i] - A_2[i]|\}$ We define $d(A_1, A_2) = \infty$ if A_1 and A_2 have different sizes." As illustrated in the following example, to get a relevant definition of continuity, the nature of the application must be considered: The inappropriateness of the article's metric for some applications can be seen by considering a "data mining" application that identifies a family through the children's Social Security numbers. If the article's metric is applied to three records

A: 101234567 104432769

B: 101234567 104432769 106222444

C: 101234568 104432768

the distance between A and B would be infinite and A and C would be very close. However, record B, an extension of record A, describes the family described by A after the birth of the third child. Record C describes a different family. An appropriate metric would consider A close to B and far from C.

Moreover, the article described its examples as "everyday programs," but these programs were typical textbook algorithms and not typical of the software we use every day. For example,

the article proved that programs that compute the length of the shortest path between two nodes are continuous. However, widely used route-finding software outputs a path, not just length. A small change in one arc's length could change the output drastically by suggesting a completely different route. One reason software continues to replace analog devices is that users often require discontinuous behavior from those devices. Software with continuous behavior will always be rare.

Proving programs correct has been a goal of computer scientists for half a century; the article reflected how far we still are from achieving that goal. Rather than verify properties of an actual program, it examined models of programs. Models often have properties real mechanisms do not have, and it is possible to verify the correctness of a model of a program even if the actual program will fail.

The article's approach is useful because attempting to prove a model of a program correct can reveal subtle errors. However, when a correctness proof is obtained, it must be taken with a grain of salt.

David Longe Parnas, Ottawa, Canada

Authors' Response:

From a purely mathematical perspective, any function between discrete spaces is continuous, so all computer programs are continuous. But this fact does not carry any useful information. In practice, some programs behave robustly and some do not, and infinite-precision models of programs offer a good way to predict whether a program is robust.

Also, our framework extends to programs that operate on values ranging over finite sets. Continuity is not a good robustness property for such programs, but, say, Lipschitz continuity is. The example programs in our article remain robust under these definitions; we also have evidence that our robustness analysis can be adapted to this context.

Swarat Chaudhuri, Houston, TX, and
Sumit Gulwani, Redmond, WA

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 0001-0782/12/11 \$15.00

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 publish selected posts or excerpts.



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

DOI:10.1145/2366316.2366319

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

New Opportunities for New SQL

Michael Stonebraker expects a substantial increase in the number of New SQL engines using a variety of architectures in the near future.



Michael Stonebraker
**"New SQL:
An Alternative to
NoSQL and Old SQL
for New OLTP Apps"**

[http://cacm.acm.org/
blogs/blog-cacm/109710](http://cacm.acm.org/blogs/blog-cacm/109710)
June 16, 2011

Historically, Online Transaction Processing (OLTP) was performed by customers submitting traditional transactions (order something, withdraw money, cash a check, etc.) to a relational DBMS. Large enterprises might have dozens to hundreds of these systems. Invariably, enterprises wanted to consolidate the information in these OLTP systems for business analysis, cross selling, or some other purpose. Hence, Extract-Transform-and-Load (ETL) products were used to convert OLTP data to a common format and load it into a data warehouse. Data warehouse activity rarely shared machine resources with OLTP because of lock contention in the DBMS and because business intelligence (BI) queries were so resource-heavy that they got in the way of timely responses to transactions.

This combination of a collection of OLTP systems, connected to ETL, and

connected to one or more data warehouses is the gold standard in enterprise computing. I will term it "Old OLTP." By and large, this activity was supported by the traditional RDBMS vendors. In the past I have affectionately called them "the elephants"; in this posting I refer to them as "Old SQL."

As noted by most pundits, "the Web changes everything," and I have noticed a very different collection of OLTP requirements that are emerging for Web properties, which I will term "New OLTP." These sites seem to be driven by two customer requirements:

The need for far more OLTP throughput. Consider new Web-based applications such as multiplayer games, social networking sites, and online gambling networks. The aggregate number of interactions per second is skyrocketing for the successful Web properties in this category. In addition, the explosive growth of smartphones has created a market for applications that use the phone as a geographic sensor and provide location-based services. Again, successful applications are seeing explosive growth in transaction requirements. Hence, the Web and smartphones are driving the volume of

interactions with a DBMS through the roof, and New OLTP developers need vastly better DBMS performance and enhanced scalability.

The need for real-time analytics. Intermixed with a tidal wave of updates is the need for a query capability. For example, a Web property wants to know the number of current users playing its game, or a smartphone user wants to know "What is around me?" These are not the typical BI requests to consolidated data, but rather real-time inquiries to current data. Hence, New OLTP requires a real-time query capability.

In my opinion, these two characteristics are shared by quite a number of enterprise non-Web applications. For example, electronic trading firms often trade securities in several locations around the world. The enterprise wants to keep track of the global position for each security. To do so, all trading actions must be recorded, creating a fire hose of updates. Furthermore, there are occasional real-time queries. Some of these are triggered by risk exposure—i.e., alert the CEO if the aggregate risk for or against a particular security exceeds a certain monetary threshold. Others come from humans, e.g., "What is the current position of the firm with respect to security X?"

Hence, we expect New OLTP to be a substantial application area, driven by Web applications as the early adopters. These applications will be followed by more traditional enterprise systems. Let's look at the deployment options.

1. Traditional OLTP. This architec-

ture is not ideal for New OLTP for two reasons. First, the OLTP workload experienced by New OLTP may exceed the capabilities of Old SQL solutions. In addition, data warehouses are typically stale by tens of minutes to hours. Hence, this technology is incapable of providing real-time analytics.

2. NoSQL. There have been a variety of startups in the past few years that call themselves NoSQL vendors. Most claim extreme scalability and high performance, achieved through relaxing or eliminating transaction support and moving back to a low-level DBMS interface, thereby eliminating SQL.

In my opinion, these vendors have a couple of issues when presented with New OLTP. First, most New OLTP applications want real ACID. Replacing real ACID with either no ACID or “ACID lite” just pushes consistency problems into the applications where they are far harder to solve. Second, the absence of SQL makes queries a lot of work. In summary, NoSQL will translate into “lots of work for the application”—i.e., this will be the full employment act for programmers for the indefinite future.

3. New SQL. Systems are starting to appear that preserve SQL and offer high performance and scalability, while preserving the traditional ACID notion for transactions. To distinguish these solutions from the traditional vendors, we term this class of systems “New SQL.” Such systems should be equally capable of high throughput as the NoSQL solutions, without the need for application-level consistency code. Moreover, they preserve the high-level language query capabilities of SQL. Such systems include Clustrix, NuoDB, SQLFire, MemSQL, and VoltDB. (Disclosure: I am a founder of VoltDB.)

Hence, New SQL should be considered as an alternative to NoSQL or Old SQL for New OLTP applications. If New OLTP is as big a market as I foresee, I expect we will see many more New SQL engines employing a variety of architectures in the near future.

Disclosure: Michael Stonebraker is associated with four startups that are either producers or consumers of database technology.

Readers' Comments

RavenDB, a document database of the NoSQL genre, has great ETL, OLTP

support, and is backed by ACID storage. See <http://t.co/fVxvmSV> (and it has evolved a lot since then).

—Anonymous

I was hoping to know the characteristics of New SQL. I am actually not convinced with your arguments about NoSQL databases as you have presented no arguments. Problems with Old SQL are known, promises of NoSQL are also known. What is New SQL?

—Anonymous

In the context of transaction processing, I would define a New SQL DBMS as one having the following five characteristics:

1) SQL as the primary mechanism for application interaction; 2) ACID support for transactions; 3) A nonlocking concurrency control mechanism so real-time reads will not conflict with writes and thereby cause them to stall; 4) An architecture providing much higher per-node performance than available from the traditional “elephants”; 5) A scale-out, shared-nothing architecture capable of running on a large number of nodes without bottlenecking.

—Michael Stonebraker

About point 3, isn't this achieved in most traditional databases by MVCC?

About point 4, most databases (Clustrix, Akiban, NimbusDB) that are New SQL candidates talk only about better query performance using distributed query or a kind of object storage. I am not sure if they have anything in better for DML performance.

VoltDB is an exception; I am not sure if it is much better than TeraData or Greenplum, which are based on Old RDBMS architecture.

About point 5, yes, this is a new feature. If I understand correctly, it means scaling the performance by adding new nodes without interrupting existing users.

—Anonymous

My previous comment suggested five criteria that defined a New SQL DBMS. I would like to stress three points that I made previously. First, my posting focused on DBMSs for new OLTP applications. Data warehouse vendors, such as TeraData and Greenplum, are focused on a completely different market, and are not designed to perform high-velocity transactions. Hence, they are not considered as New SQL vendors.

Second, most of the Old SQL vendors use standard two-phase locking, although there are exceptions. Hence, there are Old SQL engines that satisfy some of my five criteria.

Third, one of the big problems with Old SQL engines is their mediocre per-node performance. One of my criteria for New SQL vendors is much better per-node performance. The proof of this is via performance on standard benchmarks. Hence, whether any particular vendor satisfies this criteria would have to be determined experimentally. As such, the list of vendors who satisfy the five criteria may well change over time.

—Michael Stonebraker

In case any database has scale-out architecture, why is it most necessary that per-node performance also should be very high? Anyway, if performance is getting better by adding more nodes, it will be achieved.

My main focus is to understand why can't some existing database like PostgreSQL be considered an option similar to New SQL if we enhance it to support scale-out architecture in it. Saying this doesn't mean I have a clear idea about how to achieve it.

It will have benefits to existing customers as well even though the performance is somewhat less than New SQL engines. It can save a lot of effort to change applications to suit to new engines.

—Anonymous

In round numbers, New SQL engines are a factor of 50 or more faster on New OLTP than Old SQL. That means that an Old SQL engine would require 500 nodes to do a task that can be accomplished by a New SQL engine with 10 nodes.

The downside of 500 nodes is increased hardware cost, increased power consumption, increased system administration complexity, increased database administration complexity, high availability complexity (if a node fails once a month, then New SQL fails every third day), while Old SQL fails once an hour, and less flexibility (if load doubles, have to add 500 more nodes, rather than 10).

—Michael Stonebraker

Michael Stonebraker is an adjunct professor at the Massachusetts Institute of Technology.

© 2012 ACM 0001-0782/12/11 \$15.00

Major Update to ACM's Computing Classification System

ACM HAS COMPLETED a major update of its Computing Classification System (CCS), which has served as the de facto standard for classifying the computing literature since 1964. NYU professor Zvi Kedeem served as Editor-in-Chief of this special CCS Update Project, leading a group of 120 computing specialists—a third of whom are ACM Fellows^a—to collaborate with ACM staff and Silverchair Information Systems, Virginia-based specialists in ontology development and semantic markup. The 2012 version of the CCS reflects a year-long team effort that included two review stages and many iterations. The vendor created first drafts using *inter alia* user search logs from the ACM Digital Library; machine analysis of DL texts and of author-supplied keyword

occurrences; and manual examination of extant computer science taxonomies. ACM's domain experts used these drafts as their starting point.

The CCS is used in the DL to index content for subject-oriented searching; to find *similar* documents; to create author expertise profiles; to identify strong research areas in Institutional Profiles; and to create the topical tag clouds found in aggregated SIG and conference views.

Outside the DL, researchers and institutions use the CCS in their own applications and research projects. The 2012 CCS was developed as a semantic ontology and is available in SKOS format (Simple Knowledge Organization System), providing better support for linked data for those building semantic Web applications.^b

Even as the merits of full-text versus taxonomic indexing are debated,

ACM's 2012 CCS taxonomy provides a modern cognitive map of the field of computing in all its breadth. A new visual presentation of the CCS in the ACM Digital Library (<http://dl.acm.org/ccs.cfm>), designed by Wayne Graves, ACM's Director of IS, facilitates scoped navigation of the field while providing a mechanism for community review at every level of the taxonomy. Regular CCS updates utilizing this feedback are planned to maintain its currency. All feedback will be evaluated for the next periodic update, planned for 2014.


The old CCS scheme has been mapped to the new, and both the 1998 and 2012 terms are available on Citation Pages of all indexed articles. The 2012 concepts are presented in diagrammatic displays that help contextualize a work within the field, as shown in the accompanying example.

Fuller integration of the 2012 CCS in the Digital Library will take place in the coming months. Author and Institutional Profile pages will incorporate the new version; topical tag clouds for ACM Special Interest Groups and conferences will be expressed using the new concepts; and an Advanced Subject Search will be developed around the 2012 CCS. Tools to facilitate author application of the new index terms are now being built. When complete, authors can begin indexing their articles using the 2012 CCS, anticipated by January 2013.^c

a See <http://www.acm.org/about/class/2012?pageIndex=2>

b See <http://dl.acm.org/ccs/skos>

An example of ACM's new CCS as applied to a recent ISI-KDD Proceedings paper "Quantitative Analysis for Privacy Leak Software with Privacy Petri Net."



c In addition to the DL display and the SKOS format, the 2012 CCS is also available in Word: <http://dl.acm.org/ccs/word> and HTML: <http://dl.acm.org/ccs/html>

Bernard Rous (rous@hq.acm.org) is the director of ACM's Office of Publications.


Software on Mars

With the AEGIS system, the Mars Exploration Rovers can autonomously select, capture, and analyze images using onboard logic.

SINCE JANUARY 2004, scientists have been following the activities of two spacecraft on Mars with the objective of obtaining data from the red planet's soil and rocks that could offer clues about the presence of water there. The missions, launched by the U.S. National Aeronautics and Space Administration (NASA), have also been testbeds by necessity for computer vision and autonomous analysis capabilities.

Among the most innovative applications NASA computer scientists have contributed to the mission is Autonomous Exploration for Gathering Increased Science (AEGIS), which analyzes images gathered by the rovers' navigation cameras to identify features of interest, typically rocks with certain preprogrammed characteristics, without needing synchronous communication with scientists on earth. This capability saves significant time and bandwidth on the deep space network. In fact, NASA considers AEGIS so innovative it won the agency's Software of the Year award for 2011.

The Mars Exploration Rovers (MERs), named Spirit and Opportunity, have served as interplanetary geologists; their onboard instrumentation includes panoramic cameras, numerous spectrometers, magnets, and mi-



NASA's ill-fated Spirit used its front hazard-avoidance camera to record this forward view of its arm and surroundings during the rover's 2,052nd day on Mars (Oct. 11, 2009).

croscopic tools. The rovers, as implied by their names, have also traversed significant distances for unmanned small vehicles; Spirit traveled more than 7.7 kilometers before going silent in March 2011, and Opportunity continues to travel, having logged almost 35 kilometers since it landed. Opportunity has struck some significant scientific finds, including evidence that what-

ever crashed into Mars and created the Endeavour crater led to an impact that released heated, underground water that deposited zinc in that rock. The discovery is among those that NASA calls important discoveries about wet environments on ancient Mars that may have been favorable for supporting microbial life. The Endeavour site is about 21 kilometers from the rover's

previous location, Victoria crater, a distance that took the rover almost three years to traverse.

This mobility, however, comes at a cost.

“Communications bandwidth has not grown as fast as rover traverse range,” NASA computer scientists explained in a paper presented at the 10th International Symposium on Artificial Intelligence, Robotics and Automation in Space in Sapporo, Japan, in 2010. “As this trend in increased mobility continues, the quantity of data that can be returned to Earth per meter traversed is reduced. Thus, much of the terrain the rover visits on a long traverse may never be observed or examined by scientists.”

Autonomy Introduced

To meet the scientific goals of the MER missions, the engineers and scientists, who work at the Jet Propulsion Laboratory (JPL) in Pasadena, CA, built a science platform called Onboard Autonomous Science Investigation System (OASIS), which enables the rovers to autonomously perform image and data analysis, planning, execution, and interaction with robotic control without real-time human direction.

“We have a hard time getting all the data down quickly,” says Tara Estlin, a senior member of the JPL artificial intelligence group that developed OASIS. “Communication bandwidth is a restricted and precious resource and sometimes it can be several days, if not longer, for data from a certain area to come down. It could come down after the rover has left an area, maybe days later.”

AEGIS, which is a subset of OASIS, was uploaded to Opportunity in December 2009. AEGIS uses onboard data-analysis techniques to seek out scientist-defined high-quality targets with no human in the loop. Prior to AEGIS, images were transmitted from the rover to the operations team on Earth; scientists manually analyzed the images, selected geological targets for the rover’s remote-sensing instruments, and generated a command sequence to execute the new measurements.

The new approach, Estlin says, is a boon to the overall goals of the MER mission “because we can’t have the rover stay in every area long enough

to look around and take images of every spot. So this gives Opportunity the ability to do reasoning onboard about what’s interesting.”

In a seven-step process, AEGIS locates likely targets of high scientific value (typically rocks), prioritizes targets within a wide-field image, and then further analyzes properties of prioritized rocks, such as brightness, shape, and size, using the rover’s narrow field-of-vision tools. The target parameters are sent to the rover by the craft’s sequencing team at JPL, based on information supplied in advance by the MER science team.

“During a day when we’re going to be planning a drive, the scientists fill out the request for AEGIS to be run,” Estlin says. The scientists typically specify what point of the drive they want it to be run in, what space around the rover they want it to look at, and what would make an interesting target.

“The way they do that is they have a number of different ways to specify a rock’s properties,” she says, “everything from the size of the rock to the shape, reflectance, or the brightness of the image, and they can choose one or two parameters to emphasize. That information goes up with the sequence along with everything else that is going to be done that day or the next several days, and everything is done automatically.”

Thus far, Estlin says, AEGIS has been used to collect targeted, 13-color

OASIS enables the Mars rovers to autonomously perform image and data analysis, planning, execution, and interaction with robotic control without real-time human direction.

filter, panoramic camera images on a number of different terrain features including rock outcrop, crater ejector, boulders, and cobbles. These color images identify different properties of Mars surface materials including physical, mineralogic, and photometric properties, and have contributed to determining the geologic and aqueous history of Mars.

AEGIS is based on the principles of the Canny edge detector algorithm, which finds likely edges through a process of calculating intensity gradients in an image; essentially, the algorithm in a Canny-derived detector employs hysteresis to denote a minimum threshold of differentiation between pixels in a selected range of an image that marks an edge.

The AEGIS algorithm is called Rockster, and the JPL team conserved on computational resources by employing Canny techniques such as image smoothing in a preprocessing mode. This reduces the total number of edge elements detected.

Benjamin Bornstein, the project’s software development lead, says he has been very satisfied with the trade-off the team needed to accept in order to both deliver useful data and conserve resources. The team reported a small number of false positives, including cable ties on the rover deck that had extended into the field of view, and during an experiment in which AEGIS was allowed to consider targets below 25 pixels in size.

“We have the ability to adjust some of the parameters of the algorithm that indicate how aggressive it will be at trying to find rocks,” Bornstein says. “We’ve been very happy so far with our particular implementation and the choice of parameters we typically run.”

Best Practices Confirmed

While AEGIS’s function marks a pronounced shift in the way interplanetary science will be conducted, Estlin and Bornstein say there were no revolutionary methodologies or technologies used in its creation. In fact, they say, the limitations of space platforms actually compelled them to “think old.”

“Processors are typically several generations behind what you may find on your desktop,” Estlin says. “Things that might take a few seconds to run on

your desktop could easily take minutes to hours on a flight computer.”

MER’s processor, for instance, is a 25 megahertz RAD6000, a radiation-hardened version of the IBM RISC single-chip CPU, which features 128 megabytes of RAM and 256 megabytes of flash memory, “several orders of magnitude slower than what you might expect,” Estlin says.

AEGIS also has a four-megabyte cap of available RAM, and the fact that it often processes images of more than one megabyte each dictated that the developers employ various conservation techniques, such as bit packing and representing data as bit images. The AEGIS software was written in C, which Bornstein says proved ideal for the operations with which his team was charged.

“With languages like Java, or especially C++,” Bornstein says, “abstractions can be convenient, but there are a lot of implicit operations that happen, such as when copy constructors are invoked, or destructors or assignment operators, or if you have any sort of operator overloading. Those implicit sort of function calls, unless you’re an absolute expert in the code base and know exactly how everything was designed, can actually create real problems when reasoning about a piece of code. Whether you’re looking at it with an automated tool or with human eyes, trying to determine exactly what a particular line of code is doing, we want to keep things as simple as possible.”

The AEGIS software also underwent painstaking testing, including 348 unit tests, automated regression testing, and extensive run-throughs in JPL’s onsite “Mars yard,” a simulated Martian landscape, on rover hardware. Bornstein says the code was also examined line-by-line by members of the AEGIS team, JPL machine vision and artificial intelligence experts who were not on the team, and JPL experts familiar with the other code onboard the MERs with which the AEGIS software would interface. However, there was no novel or unique testing regimen simply because the software was destined for use millions of miles from Earth.

“I wish we could say there was a nice little reinvention story here, but there’s no panacea,” Bornstein says. “You layer good practice on top of good practice, each layer adding insurance and catching problems along the way. We had

AEGIS is based on the principles of the Canny edge detector algorithm, which finds likely edges through a process of calculating intensity gradients in an image.

very standard development practices, or at least what I would hope would be standard development practices.”

In fact, while it may seem counterintuitive to think that developing vision software for space applications may be simpler than for terrestrial platforms, JPL computer vision supervisor Larry Matthies says there are sound reasons for it.

“In many ways, the terrain is less complex,” Matthies says. “You have no vegetation or water; you have basically desert. You have some dust storms, but where we’re operating there is basically no weather. You have effectively no shadows, because the only thing casting shadows is the rover and we can arrange things so shadows aren’t a problem. And the rovers are moving pretty slowly, so even with very limited computation power, we can get by with simpler algorithms.”

Yet Matthies says the research on various components of the Mars missions has yielded benefits for Earth-bound projects, such as breakthroughs in stereo vision, visual odometry to help autonomous vehicles cope with slippage over uneven terrain, and hazard detection and obstacle avoidance during landing.

For example, he says, through 1990, stereo vision was considered to be very expensive computationally and error prone. As a result, researchers focused on algorithms that computed 3D information at high-contrast places in images, like corners of objects or distinctive edges. “Unfortunately, this produced pretty sparse 3D information, and it wasn’t adequate for Mars,” Matthies

says. He created a much faster algorithm that produced 3D information at many more points in the image with high reliability in 1990 at JPL, which is still being used on the Curiosity rover.

While a graduate student at Carnegie Mellon University in 1986, Matthies also discovered a new algorithm, which is being used to improve navigation. “It could estimate where a robot moved, much more accurately than previously, by using onboard stereo cameras to track distinctive points in the environment, which essentially served as local landmarks,” he says. “The heart of this innovation was a better understanding of statistical measurement errors. This class of algorithm is now called visual odometry.”

Estlin says AEGIS will be uploaded onto the newest rover, Curiosity, which landed on Mars in early August, during its first year of operation. Although Curiosity’s capabilities exceed those of its predecessors, Bornstein says AEGIS will have to share the new hardware.

“There are a lot of other things running and consuming resources,” he says. “There will be an improvement, but maybe not as dramatic as we would like it to be.” **C**

Further Reading

Canny, J.

A computational approach to edge detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8, 6, June 1986.

Castano, R., et al.

OASIS: Onboard autonomous science investigation system for opportunistic rover science, *Journal of Field Robotics* 24, 5, May 2007.

CMU Robotics Institute

AEGIS Automated Targeting for the Mars Exploration Rover Mission, <http://www.youtube.com/watch?v=X9ortg6NTiU>, Nov. 15, 2010.

Estlin, T.A., et al.

AEGIS automated targeting for the MER Opportunity rover, 10th International Symposium on Artificial Intelligence, Robotics, and Automation in Space, Sapporo, Japan, Aug. 29–Sept. 1, 2010.

Matthies, L., et al.

Computer vision on Mars, *International Journal of Computer Vision* 75, 1, Oct. 2007.

Gregory Goth is an Oakville, CT-based writer who specializes in science and technology.

© 2012 ACM 0001-0782/12/11 \$15.00

Control Without Controllers

Disney's Touché project could transform every conductive surface into a touch-control surface.

IF “THE MOST profound technologies are those that disappear,” as Xerox PARC manager Mark Weiser wrote in 1991, then most human-computer interfaces are still relatively shallow. While the touchscreen-enabled pocket device has “disappeared” the computer in many contexts, a truly vanished device would never have to leave your pocket, or the hem of your clothes, or its spot embedded in your jewelry: Its interface would be your complex touch on the world’s existing surfaces.

Scientists at Disney Research in Pittsburgh have advanced this vision with their Touché sensor project, which uses a new Swept Frequency Capacitive Sensing technique to distinguish several kinds of touch on any of five types of surface. These sensor-enabled surfaces could be used to control any kind of electronic device. Perhaps Touché’s most interesting mode turns one’s own body into a controller: In a demonstration video, a user controls the music player in his pocket by touching his hands and arms to play a song, change volume, and switch tracks.

The Disney paper “Touché: Enhancing Touch Interaction on Humans, Screens, Liquids, and Everyday

Objects” has sparked enthusiasm among researchers and analysts, partly because it could transcend traditional computing devices. Engineer and human-computer interaction researcher Darren Leigh believes Touché’s design is best suited for “interaction that’s much more subtle, and perhaps much more natural” than existing solutions. “Previous touch technologies say where you were touching a device,” he says. “Touché is different because it can tell *how* you’re touching it.”

But the project raises as many questions as it answers. It is unclear how much user training would be needed for acceptable results; radio interference caused by the frequency sweep could incite regulatory agencies to act; and the project’s paper refers several times to the need for “future work” in various areas. Further, the market is glutted with more than a dozen touch technologies including those based on resistive, optical, and other types of capacitive phenomena.


Using Capacitive Profiles

Technically, Touché is unusual in that it measures capacitance in a range of frequencies—from 1KHz to 3.5MHz, cycled 33 times a second—rather

than relying on a single frequency, as is most common. This in itself is not new: Antenna builders commonly use a network analyzer in this way to optimize signal and minimize interference. But Touché goes further by taking the resulting curve, which its researchers call a “capacitive profile,” and comparing it to others using machine learning, thereby distinguishing among touches.

The results are impressive, if imperfect. Touché was able to distinguish among five types of doorknob touch—one-finger, pinch, in a circle, grasping, and absent—at nearly 97% accuracy. Dropping the most troublesome one, the circle touch, raised accuracy to 99% with per-user training and 96% without. Similar results came from the other four modes studied, with users touching a table; a handheld device; their own bodies, as in the music-player example; and water in an aquarium.

Rather than try to predict the capacitive profile that would result from each touch, Touché’s researchers fed their experimental data curves into a machine learning system. One of the paper’s authors, Disney Research’s Senior Research Scientist Ivan Poupyrev, described their approach as “brute



Touché can sense a variety of human touches and hand postures, including grasping a doorknob and fiddling with a handheld device.

force.” “We just took all our data and fed it into a machine-learning system,” he says. “We did identify some of the distinctive features, such as peaks and slopes. But we didn’t try to understand the exact reasons why data looked that way or do much fine tuning.” Fortunately, their research methods yielded an enormous amount of data to work with, and they have already started secondary analysis. “We didn’t just take the curve, but also multiple derivatives,” says Poupyrev. “Now it’s time to go back and look at those curves and understand why they work this way. Then we can design specific algorithms to build something better.” Such optimization could improve “walk-up” accuracy significantly.

Jonathan Westhues, a self-employed engineer and designer of capacitive systems, points to other ways of extracting information from the project’s data. “It would be cool to see how much that additional info [Touché’s Swept Frequency Capacitive Sensing system] buys you over single-frequency sensing,” Westhues says. “You could train a classifier using just one scalar capacitance, then train a classifier using the full set of data over frequency. If you compared those, and saw that the recognizer trained with the full set of data performed better, you’d have mathematical proof that there’s extra information in there. The researchers are in an excellent position to test it, possibly using their existing data.”

Even without follow-up optimization and analysis, Poupyrev believes the current results could be good enough for some purposes. He points out that “the sweet spot for accuracy depends on the application. For example, when you’re buying something from a vending machine, you don’t want to learn how to use it. You just want the stuff, so it has to be 99.9% accurate right away. At the other extreme are home controls—your TV or home stereo and such. People have incentive to do it right, and therefore they can quickly learn the exact gestures that would get them what they want.”

Market Possibilities

The Touché paper is short on specific applications, focusing instead on general categories of interaction. But

Rather than try to predict the capacitive profile that would result from each touch, Touché’s researchers fed their experimental data curves into a machine learning system.

the payoff could be high for creating touch interface applications that can compete with existing ones. Jennifer Colegrove, vice president of emerging display technologies at DisplaySearch, says the market for touchscreens has been on the upswing since 2009 when she sized it at \$4 billion. It jumped to \$7 billion in 2010, then to \$12 billion in 2011; Colegrove predicts it will exceed \$20 billion in 2017.

Although Touché would not be appropriate for all types of applications in the touchscreen space, it could capture some parts of the market, particularly those that require instrumenting nontraditional surfaces such as medical devices and tabletops.

But substantial roadblocks stand in the way. According to Colegrove, competing touch technologies are cheaper, easier to get, and backed by known manufacturers. “The current touchscreen market is about price competition,” she says, “and I don’t see Touché as a threat in that market within three years. Also, the market depends on its suppliers, and Touché is still in the early stage, with only one institute doing it. Other touch technologies have 20 or 80 suppliers. Among them, Intel is trying to grab market share with projected capacitive [technology].”

Another barrier could come in the form of government regulation. “Because Touché radiates radio frequencies, anyone who tries to commercialize it will have problems with Federal

Security

Infected USB Drives

As networks become more resilient against attack, removable media—especially USB flash drives—have become the method of choice to bypass defenses and infect systems with malware.

That’s why the Honeynet Project, a nonprofit research organization that raises awareness about computer security threats, is supporting the development of the Ghost USB Honeytrap Project. Based on the research of Sebastian Poeplau, a student at Bonn University’s Institute of Computer Science, Ghost is a malware detection component that simulates the connection of a USB flash drive.

“If malware that propagates via USB sticks resides on a system,” says Christian Seifert, Honeynet Project CEO, “the malware will attempt to copy itself onto this simulated USB stick therefore allowing Ghost to raise an alert.”

Ghost can also serve as an analysis tool in a research setting to categorize malware. With Ghost—in combination with the analysis platform CuckooBox—one can characterize the extent to which USB flash drives are a propagation method for malware.

“Reports about worms and viruses that spread via USB drives go back until at least 2006, but the method became popular when the Conficker worm introduced it as an additional infection vector in early 2009,” says Tillmann Werner, senior security researcher at CrowdStrike and a member of the Honeynet Project. “Stuxnet is a more recent example where the method was used to infect machines that weren’t reachable over the network.”

The next version of Ghost is expected to make deployment in production environments more convenient.

One challenge, says Seifert, is that “if Ghost is widely adopted, new malware will learn to detect a Ghost USB drive and just won’t bite. Our job will be to make Ghost more resilient against detection and evasion.”

—Paul Hyman

Communications Commission [FCC] certification,” says Leigh. “It sweeps from audio frequencies [starting at 1kHz] through the AM radio frequencies, and then into the shortwave bands [up to 3.5MHz]. The FCC controls everything over 9kHz. The AM band is already a very noisy part of the spectrum, and AM radios are incredibly sensitive devices.”

But if industry can work out those issues, a new world of as-yet unimaginable applications could appear. So says Brandon Taylor, an engineer at Samsung Electronics who developed a “graspable” user interface that could, for example, let players vary pitch selections in a baseball video game based on how they grasp the ball. “In my work, I asked, In what types of interactions is fine manipulation really important?” he says. “If Touché produces a rich set of data about how a user’s hand is actually manipulating an object, you could start to understand some of the differences between a skilled artist in a field and a beginner. That could provide a different insight into how people learn these skills and how they’re really manipulating things. But that means going from ‘We can determine there are two fingers in water’ to ‘We can understand why someone is a virtuoso on the piano.’”

Poupyrev believes Touché could also be broadly valuable for increasing accessibility. “We don’t talk about this in the paper, but users can define their own gestures,” he says. “That’s very interesting, because we design our world assuming that people all have five fingers and can use them. But there are a lot of people with injuries or disorders who are often left

The Touché paper is short on specific applications, focusing instead on general categories of interaction, but the payoff could be high for creating touch interface applications that can compete with existing ones.

out. So we usually have to design something special for them that often could be ugly, expensive, heavy, and sets them apart from everybody else. But people who are in these situations don’t want to look disabled or wear funny devices—they want the same ones as everyone else. With our technology we might make devices easily adaptable to any special needs.”

Touché is an example of research that combines existing technologies—capacitive sensing, network analysis, and machine learning—rather than introducing truly new ones. Taylor again recommended that follow-on research look beyond traditional computer applications.

“The real root of this sort of re-

search isn’t in designing another way to right-click on a mobile phone,” says Taylor. “Instead it’s ‘If we broaden our scope of thinking, what does this allow us to do that isn’t already being done in another way?’ Maybe these sensors could figure out what you’re doing and react appropriately. Are you busy writing something? Should I route things to voicemail because you’re engaged in a behavior that I can recognize as being occupied? That’s where we can see a real breakthrough with these types of interactions.” **C**

Further Reading

Poupyrev, I., Yeo, Z., Griffin, J.D., and Hudson, S. Sensing human activities with resonant tuning, *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, Atlanta, GA, April 10–15, 2010.

Sato, M., Poupyrev, I., and Harrison, C. Touché: enhancing touch interaction on humans, screens, liquids, and everyday objects, *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems*, Austin, TX, May 5–10, 2012.

Taylor, B.T., and Bove, V. Graspables: Grasp-recognition as a user interface. *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, Boston, MA, April 4–9, 2009.

Touché
<http://www.touche-sensor.com>

Zimmerman, T.G., Smith, J.R., Paradiso, J.A., Allport, D., and Gershenfeld, N. Applying electric field sensing to human-computer interfaces, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Denver, CO, May 7–11, 1995.

Tom Geller is an Oberlin, OH-based science, technology, and business writer.

© 2012 ACM 0001-0782/12/11 \$15.00

Milestones

Congressman Launches Golden Goose Awards

U.S. congressman Jim Cooper (D-TN) has launched the Golden Goose Awards in an effort to “demonstrate the human and economic benefits of federally funded research by highlighting examples of seemingly obscure studies that have led to major breakthroughs and resulted in significant societal impact.”

The name of the awards is a play on the Golden Fleece awards, which were presented between 1975 and 1988 by former senator William Proxmire (D-WI) on federally funded research that he considered a waste of government money, such as a National Institute on Alcohol Abuse and Alcoholism study on whether fish who had

drunk too much tequila are more aggressive than fish who are sober.

Among the inaugural Golden Goose honorees is physicist Charles Townes, whose government-supported research on amplifying waves of radiation into a single, continuous stream during the 1950s was met with much doubt and derision, but led

to the invention of lasers. He won a Nobel Prize in Physics for this work in 1964.

The Golden Goose Awards will be announced three to four times a year, and an annual awards banquet will be held in Washington, D.C. to honor awardees.

—Jack Rosenberger

On the Digital Trail

Forensics experts increasingly use data from digital devices to solve crimes, fight lawsuits, and unravel accidents.

OVER THE COURSE of several years, a nefarious character who came to be known as Mr. Swirl left his indelible mark on the Internet. He sexually assaulted at least a dozen young boys throughout Southeast Asia and posted more than 200 photos of his sexual activities on the Web. In order to hide his identity, the man created a digital swirl to replace his face.

Beginning in 2004, investigators from Interpol began to search for Mr. Swirl. But the pedophile continued posting images and, using digital photo-editing software, altering his face so that it was unrecognizable. Experts had to find a way to unswirl the images and figure out who was behind the brutal sexual assaults. This task required reassembling millions of pixels.

Interpol called in German computer forensics experts who began examining the photos. Since the pixels in the digital images were lossless—they were moved but not altered—the task was to create an algorithm to reverse the swirl. Eventually, the forensics team cracked the photos and identified Mr. Swirl partly by examining objects in the photos and tracing the IP address of the computer from which the images were sent. In 2007, a Canadian citizen named Christopher Paul Neil was arrested in Thailand and sentenced to prison, where he currently resides.

Digital forensics has moved into mainstream society. As more and more devices record our movements, actions, and activities, there is a growing focus on using the data to solve crimes, fight lawsuits, and unravel accidents. Smartphones, automated teller machines, electronic tollbooths, credit and debit cards, and server logs all comprise a growing body of data that provide a window into numerous everyday events.

“We have moved beyond computer forensics and into the age of digital forensics,” says James Robertson, a



Interpol released these Mr. Swirl photos in October 2007 in an appeal for help to identify Christopher Paul Neil, left, who sexually abused children in photos posted on the Internet.



professorial fellow and director of The National Centre for Forensic Studies at the University of Canberra. “There is a growing array of software and hardware tools used to record what people are doing and where they are doing it.”

Beyond Computing

Digital forensics has evolved far beyond a way to examine a hard drive for metadata, time stamps, and deleted files. It's now used to unravel everything from international business

espionage to cyberspying and cyberwars. For example, an extensive examination of the Stuxnet code used to cripple an Iranian nuclear facility in 2010 eventually pointed to a joint American-Israeli operation, according to *The New York Times*.

Forensic techniques are also being used by individuals to document events that might have flown under the radar in the past. In July, University of Toronto computer science professor Steve Mann claimed he was assaulted and forcibly removed from a McDonald's in Paris after employees objected to an augmented reality headset device he was wearing. McDonald's conducted an investigation and claimed there was no physical contact. However, a video that Mann captured with the augmented reality headset appears to have recorded contact, including him being allegedly pushed out of the McDonald's and onto the street.

There's also the case of a Burger King employee in Ohio who posted a photo of himself on Internet image board 4chan with his feet in trays of lettuce and boasted “This is the lettuce

“We have moved beyond computer forensics and into the age of digital forensics,” says James Robertson.

you eat at Burger King.” At that point, other 4chan users, including members of hacker group Anonymous, began conducting their own forensic investigation. They grabbed GPS data on the photo and used a barcode on a box captured in the photo to track the exact location of the restaurant within 15 minutes. Three employees were subsequently fired.

David Billard, a professor at the University of Applied Sciences in Geneva, Switzerland and a lecturer at the Institute for Scientific Police, points out that digital technology now touches nearly every part of our lives. Cameras snap photos with time and GPS stamps, cellular towers track our movements on highways and byways, RFID readers record the precise time we pass through a tollbooth, and electronic financial transactions display a fingerprint of what we have bought and where we have been. What’s more, event data recorders (EDRs)—once limited to commercial aircraft—are now embedded in many automobiles.

In fact, these EDRs can reveal a number of things, including how fast a vehicle was traveling at the time of an incident or collision, how a driver was steering, braking, and accelerating, and which passengers were wearing seat belts. Although EDRs were designed to collect data that could be used to improve safety standards, they are increasingly used as evidence in court. Moreover, when they are combined with text and phone logs—and possibly credit card receipts—it is sometimes possible to gain a remarkably complete picture of what was taking place around the time of a collision.

Almost every court case now includes some digital evidence, Billard notes. In many instances, particularly divorce cases, understanding a chain of events is as simple as sifting through text messages, credit card receipts, and phone logs. More sophisticated types of crime—or those where a computer is used to commit the offense—may require an examination of a hard drive or an analysis of network traffic flows and data packets, adds Cal Waits, technical manager of operations at the CERT Digital Intelligence and Investigation Directorial for Carnegie Mellon University.

The Mr. Swirl case is a prime example of how digital forensics helps solve crimes. Once investigators created an algorithm to unswirl Neil’s face, they still faced the onerous task of tracking him down. An IP address indicated the computer was most likely located in a suburb of Vancouver, Canada, but law enforcement agencies could not identify the exact location. Only after investigators publicly released the images in 2007 did a Canadian teacher in South Korea recognize Neil and report him to police.

By then, however, Neil knew authorities were after him. He fled South Korea but an airline ticket revealed that he had traveled to Bangkok, Thailand. There, authorities caught his image on a surveillance camera and knew they were closing in on him. Thai police eventually tracked him down using a variety of high-tech surveillance systems focused on Neil’s transvestite lover. This included monitoring his lover’s movements through his mobile phone. A few weeks later, police arrested Neil in Khorat, a small village located about 150 miles from Bangkok.

The most complex cases, like Mr. Swirl, involve data from multiple sources and an array of systems or devices. Besides the sheer volume of digital data that now exists, forensics experts must extract the evidence without destroying the underlying system or device. “It’s not unlike DNA evidence,” Billard notes. “When you use a fragment you destroy a bit of the evidence.” Analyzing a mobile phone, for example, requires a forensics expert to modify the state of the device. “There is no way to capture the contents of the

memory without modifying the device state,” says Billard. As a result, forensics specialists must approach investigations methodically and, even then, they risk destroying valuable evidence.

It’s something of a cat and mouse game too. Tech-savvy individuals and crooks are increasingly turning to encryption, cloaking techniques, anonymizer software, and other tools to make forensics more challenging. Even advances in technology create new challenges. For instance, solid-state computer drives and flash memory make it more difficult for experts to find and extract data. For one thing, the data is stored in smaller 2KiB or 4KiB blocks rather than traditional 512 byte blocks. For another, these drives completely erase data pages rather than storing deleted data on the drive even after it is erased. This usually results in a far more complex and lengthy process, with a lower likelihood of finding the desired data.

Likewise, cloud computing complicates an array of issues, including who owns data and which country’s laws take precedence. In many cases, data might reside on multiple virtualized servers or travel across servers and change locations on a regular basis. In addition, the company that owns the data may not own the infrastructure. Consequently, a person or company under investigation could migrate its data to different servers. Establishing a chain of custody and authenticating the data can prove daunting.

The Laws of Data

Despite marked advances in digital forensics, police and courts are struggling to keep up with all the changes. “The digital age is only beginning to hit courts around the world,” says Robertson, who served as chief of forensics for the Australian Federal Police for 20 years. “Across countries, there are radically different abilities to handle both the volume and complexity of digital data.”

Steven Hunter, a partner in business litigation at Quarles & Brady, says digital forensics is increasingly used to investigate corporate data theft, determine whether a person who leaves one company and goes to work for another is taking along trade secrets, and address international business and trade

Event data recorders—once limited to commercial aircraft—are now embedded in many automobiles.

disputes. He points out that as economies and companies become more digital and global, resolving disputes and handling e-discovery becomes more complex. "Countries have very different privacy and data protection laws," he says.

For example, Hunter notes that many countries—particularly in Europe—view data privacy as a fundamental right and impose restrictions on how electronically stored information can be gathered, processed, used, and transmitted beyond borders. In 2011, Russia amended its data privacy laws to require written consent to transfer any "personal data." China also strengthened its protection of "personal information" last year, apparently to protect against the loss of corporate and state secrets. In the U.S., where privacy laws are weaker, e-discovery is more advanced than in many other parts of the world.

All of this is leading some companies—and government agencies—to focus heavily on where data is actually stored on a server and, in some cases, avoid the cloud unless there is certainty about the specific physical location data is stored. Not surprisingly, some cloud providers now guarantee that data will remain in a specified country. Billard says these issues can cut both ways: They can protect organizations but complicate international crime investigations. "Police must comply with national laws, which may limit their ability to collect information," he says.

The stakes continue to grow. Today, spouses increasingly use digital forensics tools to spy on partners, banks have entire forensics departments set up to spot fraud, audit companies pore over financial transactions for major companies, and law enforcement agencies chase hackers and cyberspies through the wormholes of the virtual world. Governments, too, are turning to systems that enable digital forensics. For instance, in Greece, Italy, and Spain, there is now a push to limit cash transactions for larger purchases ranging from €1,000 to €2,500. This could be a step toward eliminating cash altogether, and to possibly help thwart crime and tax evasion.

At some point, society will have to define the limits of how EDRs and other devices can be used—and where the

Event data recorders can reveal how fast a car was traveling at the time of an accident, how a driver was steering, and which passengers were wearing seat belts.

boundaries between reasonable privacy, fair use of data, and unreasonable search and seizures collide with governments' desire to monitor citizens and protect against perceived or real threats.

In the end, Waits says that as computers and digital systems become more sophisticated, society must think through the consequences—and the unintended consequences—of compiling vast stores of digital data. "There's a need to balance privacy with sophisticated tools used to understand complex events ranging from accidents to crimes," he says. "Digital forensics is often the key to unlocking complex mysteries." **C**

Further Reading

Casey, E. *Digital Evidence and Computer Crime, Third Edition: Forensic Science, Computers, and the Internet*. Academic Press, Waltham, MA, 2011.

Garfinkel, S.L. Digital forensics research: The next 10 years, *Proceedings of the Tenth Annual DFRWS Conference 7*, supplement, Portland, OR, August 24, 2010.

Golden, R.G. III and Roussev, V. Next-generation digital forensics, *CACM* 49, 2, Feb. 2006.

Real Crime *The Hunt for Mr. Swirl, Part 1*, <http://www.youtube.com/watch?v=wSw4zN-7UA>, Nov. 23, 2010.

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

© 2012 ACM 0001-0782/12/11 \$15.00

ACM Member News

IAN FOSTER WINS HPDC'12 AWARD



Ian T. Foster, professor of computer science at the University of Chicago, recently won

the first annual Achievement Award by HPDC, the International ACM Symposium on High-Performance Parallel and Distributed Computing.

Foster has worked for almost 20 years on problems relating to large-scale resource federation and collaborative discovery, focusing on the concerns of large scientific projects. The term "grid computing" is often used to refer to this work.

CERN director general Rolf Heuer has acknowledged grid computing as being "essential to the discovery of the Higgs mechanism" because it enabled the Large Hadron Collider community to connect thousands of computers at more than 200 institutions worldwide. The HPDC award was recognition of that work.

"What we used to call grid computing is now called cloud computing," he says, "and it has taken off in an immensely exciting way."

Foster predicts the next big opportunity will be leveraging grid computing's latest developments to deliver "powerful discovery services to researchers, educators, students, and industry. We need to work out what those services should be, how to build and use them, and ways of paying for them."

Foster recently launched the Globus Online project that aims to outsource complex and time-consuming research management processes to software-as-a-service providers. "The goal," he says, "is to make the discovery potential of massive data, exponentially faster computers, and deep interdisciplinary collaboration accessible to every one of the millions of professional researchers worldwide—and to the billions of potential citizen scientists—not just a select few 'big science' projects."

—Paul Hyman

Privacy and Security

The Research Value of Publishing Attacks

Security research can be improved by more effectively sharing what is learned from attacks on information systems.

INFORMATION SECURITY IS booming. Companies are making money selling fear and countermeasures. The research community is also extremely active, churning out papers featuring attacks on systems and their components. This includes attacks on traditional IT systems as well as IT-enhanced systems, such as cars, implantable medical devices, voting systems, and smart meters, which are not primarily IT systems but have increasing amounts of IT inside. Moreover, any new paper on analysis methods for critical systems is now considered incomplete without a collection of security-relevant scalps on its belt. Pretty much every system imaginable, critical or not, is now a target of attacks.

There are good reasons for this trend. Fear sells! Headlines are good for conference attendance, readership, and tenure cases. Moreover, negative messages about successful attacks are simple and understandable by the general public, much more so than other research results. And security and insecurity are, after all, two sides of the same coin.

Seek and Ye Shall Find

Systems have bugs and large, complex systems have many bugs. In their recent analysis of open source projects, Covert² used a static analysis tool to find 16,884 defects in approximately 37.5 million lines of source code from well-managed open source projects, which is approximately 0.45 bugs per 1,000 lines of code. These were medium to high-risk defects, including typical security-critical vulnerabilities such as memory corruption problems and API usage errors. For large-scale projects, developers cope with the seemingly infinite number of bugs in their products by employing triage processes to classify which bugs they work on and which they ignore. There are simply too many to address them all.

This should not come as a surprise. Complexity is at odds with security. Moreover, economic factors are often at play, where timeliness and functionality are more important than security. But there are other reasons why insecurity is omnipresent.

To begin with, systems undergo constant evolution. There has been

a recent surge in attacks where once-closed systems, like medical devices and cars, open up and are enhanced with new communication interfaces (for example, see Francillon et al.,³ Halperin et al.,⁴ and Rouf et al.⁶). The problem here is that the extended capabilities were usually not anticipated in the original design, often resulting in vulnerabilities that are easy to exploit. Not surprisingly, adding wireless communication without measures to ensure the confidentiality and authenticity of transmitted data results in a system vulnerable to eavesdropping and spoofing. This problem is particularly acute for products manufactured by traditional industries that did not previously require expertise in information security.

Systems not only interface with the outside world, they also interface with each other. For their composition to be secure, the assumptions of one subsystem must match the guarantees of the other. However, economics and market availability often dictate the choices made, especially for hardware components where manufactur-



ing one's own components is often not an option.

Finally, even when a system's security is carefully analyzed, this analysis depends on the deployment scenarios considered, in particular, the associated adversary model. What kind of an adversary should the system withstand? A system secure against a network attacker may be completely insecure against one with a screwdriver and physical access to the server. Many IT-enhanced systems have been developed using proprietary protocols and communication technology, leading to the belief that it was difficult for outsiders to interface with them. However, for wireless communication, the increasingly widespread availability of tools and equipment, such as Universal Software Radio Platforms, has made it easy and inexpensive for nonspecialists to communicate with even the most exotic systems, thus dramatically changing the adversary's capabilities. As scenarios and adversaries change over time, so do the possible attacks.

While publishing attacks has been controversial in the past, it has become common to publish attack papers. Today, there are, in fact, markets in vul-

nerabilities, with companies as well as governments participating in them.

Summing up, it is not surprising to see so many system attacks reported, in particular on IT-enhanced systems. But what makes attacks worthy of scientific publication? Are all these attacks of the "yet another buffer overflow" variety? Is there any point in publishing research papers that feature attacks on systems that were not designed to resist attacks, not used as they were designed, or used in scenarios for which they were not designed?

Learning from Attacks

A hallmark of good research is the generality of the insights gained. In securi-

ty, these are insights into the problem and countermeasures.

Increasing awareness is a common argument for publishing attack papers and has its merits. In particular, a heightened awareness of problems and their severity may lead to the system in question being withdrawn from service; alternatively, others can follow up with designs that solve the documented problems. Such attacks have, in the past, raised awareness among policy makers of the immaturity of existing technologies and the associated risks. This is particularly valuable for new systems and technologies. Here, the novelty of the kind of attack is less relevant than the novelty of the system and the impact of its compromise.

Although raising awareness is important, it can backfire as too much sensationalism numbs the readers' sensitivity to what the real problems are. And there is usually limited research value in just showing how standard problems can be exploited in yet another setting. It is clear that unauthenticated communication opens the door to spoofing attacks, whether we are talking about cars, medical implants, or personal robots. The same

As scenarios and adversaries change over time, so do the possible attacks.

holds for standard, well-studied, software vulnerabilities. In contrast, a paper that refines an existing attack, demonstrates a novel kind of attack, or contributes to new attacker models can have tremendous research value.

One benefit of studying attacks is a better understanding of the cause of the underlying vulnerability, for example, whether it is the result of a design or implementation error, the unavailability of solutions on the market, improper system usage, or an oversight in the risk analysis. This last reason occurs surprisingly frequently; systems are often left unprotected because the designers simply do not believe they need to be protected or assume the systems are sufficiently closed or obscure and therefore unlikely to be reverse-engineered by attackers (or determined researchers). As recent attacks on medical devices and modern cars show, these assumptions are incorrect.

An attack paper can also explicate what is required for a successful attack. Is the exploitation of a vulnerability straightforward or only possible by well-funded, technically sophisticated attackers? The devil is in the details! A good attack paper can show how to construct an exploit and the cost of doing so. Moreover, it can help refine the conditions under which the attack may succeed and its success probability. An attack might be conditioned not only on the attacker's computational capabilities but also on its physical location, antenna size, transmission power, and other factors. For example, the success of spoofing attacks on Global Positioning System receivers strongly depends on the locations and characteristics of the attackers' antennas.

To expand on our last point, what makes security special is the role of the adversary. A system's security can only be evaluated with respect to a model of the adversary, that is, a description of his capabilities. Thus, in our view, the most important reason for studying attacks is that they can help refine this model for the domain at hand. Here, we give two examples of this from the domain of security protocols and relay attacks.

In 1978, Needham and Schroeder proposed one of the first authentication protocols. Their protocol used

As our physical and digital worlds become more tightly coupled, the incidence of attacks will increase as well as their consequences.

public key cryptography to achieve mutual authentication between two principals in the presence of an attacker who can eavesdrop and spoof messages. Eighteen years after its publication, Lowe⁵ showed that the protocol could be attacked by a man-in-the-middle, who executes the protocol as an insider in two interleaved sessions. This attack sensitized the security protocol community to the importance of considering adversaries who have insider capabilities. Later, motivated by attacks on long-term keys stored in memory, weak random number generators, and the ability of adversaries to read out part of an agent's session state, cryptographers developed a host of more refined adversarial models and security definitions reflecting these enhanced capabilities. These new models have led to improved protocols as well as methods and tools for reasoning about the security of protocols and systems, with respect to these refined adversarial models; for example, see Basin and Cremers.¹

Second, more recent examples are Relay, Mafia-Fraud, and Wormhole attacks where the attackers simply relay messages, unmodified, between the two communicating parties. Such attacks have been recently used to compromise entry and start systems in cars³ and payment systems that rely on near-field communication. These attacks showed that the success of relay attacks on such systems strongly depends on the speed that attackers can process signals. They further demonstrated that existing technology enables attackers to build relays that have practically undetectable processing delays. This was particularly im-

portant in the case of entry and start systems for cars; the attacks revealed that these systems could only detect relays that introduce delays longer than several microseconds. This led to refined attacker models and also motivated new security solutions, for example distance bounding protocols.

Conclusion

As our physical and digital worlds become more tightly coupled, the incidence of attacks will increase as well as their consequences. Many of these attacks will be newsworthy, but most will not be research-worthy. This does not mean papers featuring attacks on highly visible systems should not find their way into research conferences; having had such papers published, the authors of this column do appreciate that the community accepts results of this kind. However, as researchers we should have high aspirations. With every attack paper there is an opportunity to truly contribute to the community with new insights into both systems and their vulnerabilities, and adversaries and their capabilities. We believe that one should take this opportunity and, after discovering an attack, take a step back and reflect on what can be learned from it, and afterward present it to the community. **C**

References

1. Basin, D. and Cremers, C. Modeling and analyzing security in the presence of compromising adversaries. In *Computer Security—ESORICS 2010*, volume 6345 of *Lecture Notes in Computer Science*. Springer, 2010, 340–356.
2. *Coverity Scan: 2011 Open Source Integrity Report*. Coverity, Inc. San Francisco, CA, 2011.
3. Francillon, A., Danev, B., and Capkun, S. Relay attacks on passive keyless entry and start systems in modern cars. In *Proceedings of the Network and Distributed System Security Symposium*, 2011.
4. Halperin, D. et al. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy, SP '08*, IEEE Computer Society Washington, D.C., 2008, 129–142.
5. Lowe, G. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. *Software—Concepts and Tools* 17, 3 (1996), 93–102.
6. Rouf, I. et al. Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study. In *Proceedings of the 19th USENIX Conference on Security, USENIX Security'10*. USENIX Association, Berkeley, CA, 2010, 21.

David Basin (basin@inf.ethz.ch) is a professor in the Department of Computer Science at ETH Zurich and the founding director of the Zurich Information Security and Privacy Center.

Srdjan Capkun (srdjan.capkun@inf.ethz.ch) is an associate professor in the Department of Computer Science at ETH Zurich and the director of the Zurich Information Security and Privacy Center.

Copyright held by author.



Legally Speaking

Oracle v. Google:

Are APIs Copyrightable?

Assessing the first phase of the trial based on claims that Google's Android platform infringes Oracle's Java-related copyrights and patents.

ARE APPLICATION PROGRAM interfaces (APIs) of computer programs protectable by copyrights in software that embodies them? *Oracle v. Google* is the most definitive ruling yet that addresses this question because the judge took pains to understand exactly what Java APIs are, how and why Google implemented them, and how the copyright dispute between these two software giants meshed with software copyright precedents.

The judge rejected Oracle's claim of copyright in Java APIs and his ruling suggests that APIs are uncopyrightable more generally. Oracle will appeal, but Judge William Alsup of the U.S. District Court of Northern California did a very careful job in analyzing the issues. I predict affirmance.

Origins of the Lawsuit

In 2005, Google began making plans to create a platform for mobile phones. It held some discussions with Sun Microsystems about a possible license to use, adapt, and open source Java for mobile devices. These negotiations proved inconclusive, so Google went ahead with its Android project without Sun's involvement. Google decided to use the Java language to design its own virtual machine and to write its own implementations of 37 of the 166 Java APIs.

The Android platform was released in 2007. Android-based mobile phones went on the market the next year. More than 300 million of these phones have



Courtroom sketch depicting Google attorney Robert Van Nest questioning Oracle CEO Larry Ellison during *Oracle v. Google*.

been activated since then. Google provides the Android platform for free to smartphone manufacturers and other developers. It makes money on smartphone ads and search.

Although Sun knew that Google was using Java without a license, it did not sue. Shortly after Oracle acquired Sun in 2010, Oracle sued Google for copyright infringement in a California federal court. (There were patent claims as well, which the jury rejected.)

Oracle's main copyright claim was that Google's Android platform infringed copyright because it copied the "structure, sequence, and organization" (SSO) of 37 Java APIs without permission. As the new owner of Java assets, Oracle had the right to sue anyone who it thought infringed its rights.

In August 2011, Google tried to avoid going to trial about its use of Java APIs by moving to dismiss that part of Oracle's lawsuit. If APIs are not pro-

tectable by copyright law, as Google believed, a trial was unnecessary.

By denying Google's motion, the judge seemed to give some credence to Oracle's claim. He seemed to give it further credence when he sent the API claim to the jury, asking it to decide two things: first, whether, *assuming that APIs were copyrightable*, had Google infringed by copying them, and second, whether this copying was fair use. The jury answered the first question in the affirmative and split 9-3 on the fair use question.

In the week or so after the judge dismissed the jury, everyone interested in the *Oracle v. Google* case was on tenterhooks. Because he was taking so long to rule on the legal question, the judge gave the impression that Oracle's claim might have some merit.

Yet in view of his eventual ruling, it now appears the judge was just being careful. The trial gave him a chance to understand what APIs were, what uses Google had made of the Java APIs, and why these uses were important to interoperability. Once he understood this, Oracle's claim was doomed.

Oracle's Arguments

Here is why Oracle thought its API copyright claim might succeed. In addition to making much ado about a Google engineer's belief that Google needed a license to implement Java, Oracle relied upon some appellate court decisions upholding SSO infringements.

In *Whelan v. Jaslow*, for instance, a small software company sued the owner of a dental lab because the latter wrote a program for managing dental lab operations similar to that which he had commissioned from this firm. Jaslow's program was written in a different programming language and used different algorithms. But the overall structure of his program was substantially similar to Whelan's and five subroutines were substantially similar in operations.

Because copyright law considers computer programs to be "literary works" and because copyright protects the SSO of conventional literary works, the *Whelan* court reasoned that the SSO of programs should be protectable also. The court thought that if copyright was only protected against exact copying of another's

code, as Jaslow claimed, this would undermine incentives to invest in software development.

A key issue in *Whelan* concerned the proper interpretation of section 102(b) of U.S. copyright law: "In no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation, concept, principle or discovery, regardless of the form in which it is described, explained, illustrated, or embodied in such work."

Jaslow argued that he had only copied the unprotectable methods and procedures of Whelan's program, not any expressive aspects of that software. Whelan argued that because Jaslow had copied detailed structures from her program, he had copied expression. Section 102(b), in Whelan's view, merely restated the rule that copyright does not protect abstract ideas, but only expressions of ideas. The appellate court agreed with Whelan's interpretation. Several other cases relied on *Whelan* in affirming infringement rulings based on program SSO similarities.

Another factor that seemed to support API copyrights was that developing APIs requires some creativity. Oracle asserted that Java APIs were creative enough to meet copyright's originality standard. APIs are also valuable aspects of programs and software companies have been known to recoup at least some of the costs of program development by licensing interface specifications and SSO to other companies. Moreover, if one took section 102(b) literally, computer programs would not be copyrightable, for they are by their very nature functional processes. Yet Congress clearly intended that programs should be protected.

Google had ample ammunition with which to counter Oracle's legal arguments.

Oracle analogized Java API names and name groupings to taxonomy names and groupings that had been held copyright-protectable in some non-software cases. Oracle pointed out that it was unnecessary for Google to use the same function names as Java or to arrange the names in the same way.

Because the Android platform was not fully interoperable with Java, Oracle challenged Google's compatibility defense. Moreover, by implementing only 37 of the 166 Java APIs, Oracle complained, Google was contributing to fragmentation of Java and undermining the "write once, run everywhere" goal of Java.

Google's Arguments

Google had ample ammunition with which to counter Oracle's legal arguments. The *Whelan* case, on which Oracle so heavily relied, has been discredited by a later much-cited appellate court decision, *Computer Associates v. Altai*. *Altai* criticized *Whelan* for having an outmoded conception of computer programs and as relying too heavily on metaphysical distinctions rather than practical considerations. Because *Whelan* did not involve copying of APIs, its precedential value was weak on that score as well.

Altai ruled that if a subsequent programmer's design choices were constrained by external factors, such as the hardware or software with which the program was designed to interoperate, copyright law would not penalize similarities attributable to this. Although the *Altai* court did not specifically mention APIs, it observed that parameter list similarities between Computer Associates's and *Altai*'s scheduling programs were noninfringing because they were due to constraints imposed by the IBM operating system programs with which both litigants' programs were designed to interoperate.

Also supporting Google's API defense were two other appellate court decisions, *Sega v. Accolade* and *Sony v. Connectix*. Both ruled that reverse engineering another firm's software for a legitimate purpose such as seeking access to information necessary for interoperability was non-infringing.

Accolade was free to use information obtained by reverse engineering to adapt its video games to run on

Sega's Genesis platform. The court in *Sega* characterized interface information as "functional requirements for achieving compatibility with other programs," which was unprotectable under section 102(b). Connectix could lawfully develop an alternative software platform on which consumers could play video games designed for the Sony PlayStation.

The Court's Analysis in *Oracle v. Google*

Although none of these cases specifically mentioned APIs, their import for the API copyrightability question in the *Oracle* case was clear.

The judge was also influenced by the *Lotus v. Borland* decision, which rejected Lotus' argument that its choice of command names and the arrangement of those commands in a hierarchy was "expressive." The court held that Lotus' menu command hierarchy, which was a fundamental part of the functionality of the Lotus macro system, was an unprotectable method of operation within the meaning of section 102(b).

Based on these cases and his understanding of Java APIs implemented in Android, the judge concluded that the Java APIs were unprotectable methods: "[A]nyone is free under the Copyright Act to write his or her own code to carry out exactly the same function or specification of any methods used in the Java API. It does not matter that the declaration or method header lines are identical. Under the rules of Java, they must be identical to declare a method specifying the same functionality—even when the implementation is different."

As for the Java names and name groupings, the judge observed that "the names are more than just names—they are symbols in a command structure...Each command calls into action a pre-assigned function." The Java command structure, the judge went on to say, was an unprotectable system under section 102(b), duplication of which "is necessary for interoperability." Hence, Google had not infringed copyright by its use of the Java APIs.

Oracle's concerns about fragmentation of Java were, he ruled, irrelevant to the copyright issues in the case.

When one company buys another, it is not bound by the legal positions previously taken by the acquired company.

Conclusion

There is irony in Oracle's API copyright claim against Google. When Sun was a major player in the computer industry, it was among the most vigorous proponents of pro-interoperability intellectual property rules. Sun, for instance, submitted amicus curiae briefs in the *Altai* and *Lotus* cases taking exactly the opposite position from Oracle in the *Google* case. Sun also supported an exclusion of software interfaces from copyright protection in Europe.

When one company buys another, it is not bound by the legal positions previously taken by the acquired company. Yet it should be more judicious in making contrary claims than Oracle was. Former Sun employees and lawyers may be among those smiling at Oracle's having been hoisted by its own petard. The *Oracle* API copyright ruling is a big victory for Google, but an even bigger one for competition and ongoing innovation in the software industry.

The month of May this year was a good one for interoperability. In addition to the Oracle API ruling, the Court of Justice of the EU rendered its decision in the *SAS Institute v. World Programming Language* case (about which I wrote in my March 2012 *Communications* column). It ruled that copyright protection for software does not extend to the functional behavior of programs, to programming languages, or to interfaces necessary for interoperability.

As a result, it is now safer than ever to develop interoperable programs on both sides of the Atlantic Ocean. 

Pamela Samuelson (pam@law.berkeley.edu) is the Richard M. Sherman Distinguished Professor of Law and Information at the University of California, Berkeley.

Copyright held by author.

Calendar of Events

November 22–23

Annual Workshop on Network and Systems Support for Games, Venice, Italy,
Contact: Maha Abdallah,
Email: maha.abdallah@lip6.fr

November 26–30

The 24th Australian Computer-Human Interaction Conference, Melbourne, Australia,
Contact: Vivienne Farrell,
Email: vfarrell@swin.edu.au

November 28–29

Conference on Visual Media Production, London, U.K.,
Contact: Oliver Grau,
Email: oliver.grau@bbc.co.uk

December 2–6

HILT 2012: High Integrity Language Technology: ACM SIGAda Annual, Boston, MA
Sponsored: SIGAda,
Contact: Benjamin M. Brosgol,
Phone: 646-375-0730,
Email: brosgol@gnat.com

December 3–7

13th International Middleware Conference, Montreal, Canada,
Contact: Bettina Kemme,
Phone: 514-398-8930,
Email: kemme@cs.mcgill.ca

December 10–13

Conference on Emerging Networking Experiments and Technologies, Nice, France
Sponsored: SIGCOMM,
Contact: Renata Cruz Teixeira,
Email: renata.teixeira@lip6.fr

December 18–20

Third Kuwait Conference on e-Services and e-Systems, Kuwait,
Contact: Saleh Kassern,
Email: kasserns@yahoo.com

January 6–8

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

Economic and Business Dimensions Decentralization versus Centralization in IT Governance

It's not as simple as you might think.

IN 2006, AIRBUS experienced an unexpected two-year delay in the development of its A380 megajet. When preassembled bundles of hundreds of miles of cabin wiring were delivered from its factory in Hamburg, Germany, to the assembly line in Toulouse, France, they failed to fit into the planes. Assembly ground to a standstill. The entire wiring system had to be redesigned and rebuilt. The price tag: \$6.1 billion. The reason given for one of the most expensive missteps in the history of commercial aerospace: incompatible design software.¹

Industry experts frequently blame disappointing IT performance on a lack of centralized control over IT purchases. While the Airbus example may be an extreme case, many believe decentralization of IT governance has led to ever-growing costs of technology ownership. Yet evidence also exists that centralized IT solutions fail to address the full range of needs within large multidivisional firms. A poor alignment between the technology and local business needs forces business units to reengineer their processes to fit the technology or simply work around it. A Harvard Business School case on IT at Cisco describes a “shadow IT system” that grew up around unmet



local demands for customized data and reports.²

How should firms decide whether to centralize or decentralize the selection and purchasing of critical IT infrastructure? This question has persisted in the information systems community at least since the 1970s, with periods of renewed attention whenever significant technical change reshuffled the board. The diffusion of the commercial Internet in the mid-1990s represents one such tectonic realignment. My recent study forthcoming in the *Journal*

of Economics and Management Strategy provides new evidence that, the advice of IT consultants notwithstanding, there is no single best way to structure this process.³ Firms are, in fact, more diverse in their approach than is widely recognized. The best choice depends critically on the *specific business context* in which the IT will be deployed.

Reality Is Complicated

This multi-industry study focused on approximately 3,000 firms in the U.S. manufacturing sector to understand

how they allocated authority for IT purchases. The data is from 1998, a time of vigorous investment in IT, but a period as-yet unaffected by the bursting of the dot-com bubble. The IT purchases in question included network equipment, servers, terminals, and enterprise software applications that arguably work best when they are interoperable throughout the firm. The IT governance data come from the Harte Hanks Computer Intelligence Technology Database, which surveys establishments with more than 100 employees on their use of information technology. In particular, they ask whether the authority to make IT purchases resides with local managers or with the corporate parent. This data was merged with the U.S. Census of Manufactures to gain insight into the local business environment of individual plants and the broader firm and industry context in which they operate.

Because of the wide diversity of firms in the sample, one would expect to see a variety of purchasing regimes—particularly among firms in different industries or of different sizes or vintages. A surprising finding, however, was the variation that existed *within firms*. The accompanying figure shows that firms tended neither to be fully decentralized nor fully centralized, but had local purchasing authority at some locations and centralized control at others. This is in sharp contrast to widespread advice to high-level firm managers to centralize IT purchasing for large organizations.

Adaptation vs. Coordination in Firms

Why was this the case? Were the instances of delegation simply “mistakes” to be corrected? To better understand this surprising diversity of approaches, the study leveraged an important model from the economics literature to explore whether trade-offs between “adaptation” and “coordination” might explain the patterns in the data. According to this type of model, both decentralization and centralization are reasonable outcomes for profit maximizing firms (that is, neither is a “mistake”). Instead, the former makes sense in the presence of particularly compelling local business needs that demand high levels of adaptation; the latter is desirable when coordination is of paramount impor-

How should firms decide whether to centralize or decentralize the selection and purchasing of critical IT infrastructure?

tance. Ultimate outcomes depend on the magnitudes of these competing needs throughout the firm.

This is, indeed, consistent with what the data indicate. For instance, if a plant is a major contributor to total firm sales, getting the IT “just right” at that location trumps the benefits of firmwide coordination. Adaptation is also important if a plant produces products outside the main focus of the firm. In this case, local managers may know better than those at headquarters what the plant needs. More-diversified firms also tend to delegate more, presumably because the diversity of IT needs grows with the range of firm products. The presence of legacy systems or growth through mergers and acquisitions also is associated with delegation.

These findings are in line with the notion that centralized managers have too much to contend with in the face of changing IT solutions, upgrades, and


customization requests across large, distributed firms. The benefit to the firm of letting local managers make well-adapted choices in economically sensitive or less-common circumstances may outweigh the costs of degraded coordination, overall.

This is not true in all cases, however. In firms where the value of IT interoperability is particularly high, the likelihood of centralized purchasing is much higher. For instance, firms that must coordinate production across locations—and whose IT must be able to communicate across plants—are significantly less likely to delegate. Also, firms apparently do not give up on monitoring what local managers do. In firms that are large enough that central managers may have trouble “keeping tabs” on plant-level purchasing managers, delegation is somewhat less likely. Finally, variation in IT governance structures holds across the range of IT spending; even the biggest IT budgets are often left to local decision-makers.

One “Size” Does Not Fit All

Firms must weigh these many factors in deciding whether to keep the purchasing authority for IT systems at headquarters or farm it out to individual plants. The analysis provides evidence that one approach is not necessarily better than the other in all circumstances. In addition, it points to the need to align the IT governance approach with other key choices at the firm such as its growth and diversification strategy, reliance on legacy IT systems, and the distribution of critical

Distribution of IT purchasing delegation in U.S. manufacturing firms.



ACM's Career & Job Center

Looking for your next IT job?

Need Career Advice?

Visit ACM's Career & Job Center at:

<http://jobs.acm.org>

Offering a host of career-enhancing benefits:

- 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 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!

<http://jobs.acm.org>



Association for
Computing Machinery

Advancing Computing as a Science & Profession

operations within the firm. These findings raise key questions for managers to keep in mind when making this important decision:

- ▶ Are the IT needs of this business unit consistent with what exists throughout the firm? Is this a “special case” that needs a different approach?

- ▶ Can our central IT division manage the coordination challenges in a firm this size and/or scope? If not, would strategic delegation to some units make better sense?

- ▶ If coordination is absolutely critical, do we have the centralized IT resources to manage this well? How can we be sensitive to local needs while adhering to unified standards throughout the firm?

- ▶ How important is this one unit to overall firm performance—if a compromise must be made in terms of technology fit, is this the place to make it?

Getting the IT purchasing decision right is a critical one for firm performance. Anecdotal evidence points to the pitfalls that arise from mistakes in both directions. Firms that fail to coordinate when interoperability is critical risk Airbus' fate. Yet, forcing inappropriate standards onto idiosyncratic—and economically important—internal operations is also an unattractive option. Examples of “too much coordination” may make the news less often than a single \$6 billion coordination failure, yet the costs of daily workarounds and poor IT fit may be quite costly in the long-run, as Cisco can attest.

Weighing the costs and benefits of a monolithic approach for the entire firm is unlikely to be a fruitful endeavor. Instead, managers should consider a case-by-case approach that may better accommodate the diversity of business contexts that exist within their particular firms. **C**

References

1. Matlack, C. Airbus: First, blame the software. *Businessweek*. (Oct. 5, 2006).
2. McAfee, A., McFarlan, F.W., and Wagonfeld, A.B. Enterprise IT at Cisco (2004). Harvard Business School Case #9-605-015, 2007.
3. McElheran, K. Delegation in multi-establishment firms: Evidence from I.T. purchasing. *Journal of Economics and Management Strategy*. Forthcoming.

Kristina McElheran (kmcclheran@hbs.edu) is the Lumry Family Assistant Professor of Business Administration in the Technology and Operations Management Unit at Harvard Business School, Cambridge, MA.

Copyright held by author.

Education

Learning to Teach Computer Science: The Need for a Methods Course

A multipronged approach to preparing computer science teachers is critical to success.

COMPUTER SCIENCE IS a crucial driver of innovation and productivity in our technology-rich society and there is a strong demand for computationally educated workers. Yet, there is a shortage of computer science undergraduates and computationally educated students. Furthermore, computer science plays only a minor role in high school curricula in the United States. Exposing students to computer science in K–12 education could be a natural pipeline for the field, but the U.S. is facing a very different reality: only 15,000 high school students take the Advanced Placement Computer Science (APCS) exam annually, less than 10% of U.S. high schools offer an APCS course, and nationally there are only 2,000 teachers qualified to teach the APCS course.

The National Science Foundation has started an ambitious effort, the CS10K Project, to have computer science taught in 10,000 schools by 10,000 highly qualified teachers by 2016. However, teaching quality computer science courses not only requires substantial understanding of the specific subject matter, but also a solid background of pedagogy. These two areas of knowledge, referred to as “content knowledge” and “pedagogical knowledge” are important to someone who hopes to become an effective teacher. However, knowledge of just the content and just the pedagogy is insufficient. It

is also important to understand ways of teaching the particular subject matter, that is, to have *pedagogical content knowledge*. *Pedagogical content knowledge* is a kind of knowledge that “goes beyond the knowledge of subject matter per se to the dimension of subject matter knowledge for teaching.... [and it includes] the ways of representing and formulating the subject that make it comprehensible to others.”¹³

Prospective teachers gain knowledge in these three areas from three kinds of courses: content courses in their discipline (for example, mathematics and physics), pedagogical courses that provide broad educational training (such as learning theories and classroom management), and methods courses (pedagogical approaches to teaching a specific discipline). For example, a mathematics methods course weaves together knowledge of mathematics with knowledge about how children learn mathematics, how the classroom environment influ-

ences that learning, and the teachers’ role in student learning.¹

In this column, we argue that in order to meet the CS10K Project goals, high school CS teachers must have in-depth computer science knowledge as well as strong pedagogical content knowledge, developed through a computer science methods course.

Learning to Teach

Learning to teach can be conceptualized around four main ideas—learning to *think* like a teacher, learning to *know* like a teacher, learning to *feel* like a teacher, and learning to *act* like a teacher.⁷ These knowledge systems are developed with a comprehensive understanding of the subject matter to be taught as well as ways of teaching that subject matter, that is, pedagogical content knowledge. Teachers with in-depth pedagogical content knowledge understand ways of representing and formulating the subject matter—using powerful analogies, illustrations, examples, explanations, demonstrations, and so forth—to make it understandable to students.¹³ These teachers also know which topics students find easy or difficult to learn, which ideas (often misconceptions) students bring with them to the classroom, and how to transform those misconceptions. In addition, teachers understand how students develop and learn as well as how to teach diverse learners.

ARISTOTLE

**“Teaching is
the highest form
of understanding.”**



A methods course is typically where prospective teachers are introduced to this skill set and learn about “pedagogical ways of doing, acting, and being a teacher.”¹ This knowledge is developed within the context of learning and teaching a particular subject area. Transforming Ball’s statement about mathematics to computer science implies that a computer science methods course is about how computer science is learned and taught, and about how classrooms can provide an environment for learning computer science.

A Computer Science Methods Course: Learning to Teach CS

Recently, there has been discussion within the CS education community concerning the need for teacher licensure programs to address the critical need to adequately prepare high school CS teachers.⁸ At Purdue University, we offer a CS Methods Course as part of our computer science supplemental licensure program.^a The methods course builds upon existing course work in computer science and education to prepare secondary STEM education majors to be effective high school computer science teachers. Prospective teachers take the course as a final step before their student teaching experience.

^a The Purdue Computer Science Teaching Supplemental Licensure Program prepares education majors to teach computer science in secondary schools. For more information and related coursework see <http://cs4edu.cs.purdue.edu/license>.

Our CS methods course develops students’ pedagogical content knowledge through experiences that allow them to think and act like computer science teachers. Specifically, the methods course trains prospective computer science teachers to combine pedagogical principles with computer science content to improve the learning experience for their students. Prospective teachers are expected to understand computer science concepts as well as to have a basic understanding of how children learn and how to promote effective learning strategies.

The specific activities and assignments in the methods course help prospective teachers understand and address computer science classroom situations. The course involves reading, discussing, and reflecting on papers that describe pedagogical practices of teaching computer science principles.

Teachers with in-depth pedagogical content knowledge understand ways of representing and formulating the subject matter.

Reading materials for the course are selected from a variety of sources, including ACM SIGCSE, *Reflections on the Teaching of Programming*,³ and *Guide to Teaching Computer Science*.⁹ For example, we use Wu’s Castle⁶ and image processing⁴ to illustrate ways of teaching arrays and loops.

The course includes a laboratory component in which students review and evaluate various pedagogical tools, such as DrJava, Alice, Scratch, Light-Bot, Greenfoot, and *CS Unplugged*² activities. The lab also includes a review of topics covered in the AP Computer Science (Java) course, including GridWorld.^b

Prospective CS teachers develop lesson plans for teaching computer science topics, chosen in consultation with the course instructor. The lesson plans detail student learning goals, how the teacher (and students) will accomplish those goals, a corresponding laboratory activity, and methods of assessing whether or not students have met the learning goals. The prospective teachers use one of their lesson plans for a micro-teaching experience during which they “teach” their peers/course instructors (who role-play as students) and practice their teaching skills in a supportive environment. In addition, students in the course also complete a “theory in practice” component (that is, field experience visits), which involves observing local high-school computer science classrooms.

Students keep two types of journals: a learning journal and a reflection journal. In their learning journal, students record their reactions to the assigned readings and class activities using a set of guided questions. In their reflection journal, they think and reflect upon their visits to the school as well as their micro-teaching experiences. These prospective CS teachers benefit from their reflections because it helps them develop a clear and more complete understanding and insight into computer science teaching. Specifically, reflections have been found to be very effective in developing a teacher’s thinking, since “having to express oneself to others, so that others truly understand one’s ideas, reveals both the strengths

^b See http://apcentral.collegeboard.com/apc/public/courses/teachers_corner/4483.html.

and the holes in one's thinking."¹²

The multifarious experiences built into our methods course, such as classroom observations and micro-teaching combined with reflections provide opportunities for prospective CS teachers to *think, know, feel, and act* like a teacher. The activities used in our methods course also meet NCATE (National Council for Accreditation of Teacher Education) requirements of planning and delivering lesson plans, observing secondary CS teachers, and evaluating their own practice of teaching. The many experiences (field experience, observation, micro-teaching, reflections, and so forth) provided in our methods course equip prospective CS teachers with the ability to use a broad range of pedagogical approaches and computing tools when they begin teaching a computer science class (ISTE[®]).

We have presented one framework for teaching a computer science methods course; however, this model is not meant to be prescriptive. Other researchers have also been involved in teaching computer science methods courses and have presented a similar framework for their methods course (see, for example, Lapidot and Hazzan¹¹). CS educators interested in teaching a methods course must be aware of contextual factors (for example, student background, state teaching standards, and teaching licensure requirements) that are relevant to their course. Regardless of the framework, a quality methods course builds on students' comprehensive understanding of computer science concepts and a thorough knowledge of pedagogy.

Conclusion and Future Directions

To meet the CS10K goal, we need a multipronged approach to preparing computer science teachers, both at the preservice and inservice levels. A CS methods course is a necessary component to prepare not only preservice, but also inservice teachers to teach rigorous computer science courses at the high school level. A methods course geared toward the needs of inservice teachers could help them meet the challenges of a CS classroom, such as creating and maintaining a productive lab environ-

Preparing CS teachers is an important step to increase the role, availability, and recognition of computer science in high schools.

ment and assessing open-ended student work (programs), while giving them opportunities to grow their pedagogical content knowledge and enhance their overall computer science knowledge.

Typically, high schools cannot afford a full-time CS teacher (for example, due to low enrollment), but instead must rely on teachers whose primary licensure is in another content area, such as mathematics. This situation can lead to teachers with insufficient pedagogical content knowledge to teach CS courses. Furthermore, these teachers face isolation in their school because there are no other teachers with whom to brainstorm ideas to meet the challenges of a CS classroom. The combination of rapidly changing computer technology and busy workday schedules of teaching multiple content areas also makes it difficult for teachers to stay current in the field.

Historically, inservice teachers have gained knowledge through university courses as well as professional development opportunities, including workshops. However, professional development approaches that rely heavily on "one and done" workshops are known to have limited success.⁵ Hence, there is a critical need to provide inservice CS teachers with opportunities for in-depth and continual training so they can gain sufficient knowledge necessary to offer rigorous computer science courses.

These approaches to professional development need to be sustainable, scalable, and fit well within teachers' busy schedules. One potential way to meet this need and provide an accessible means for inservice teach-

ers to expand knowledge and develop skills to teach computer science (that is, pedagogical content knowledge) is through online courses. An online methods course would allow teachers to continue their learning and communicate with their peers about topics and techniques they use when teaching computer science. We believe such a course has the potential to support inservice CS teachers and develop their pedagogical content knowledge.

Preparing CS teachers is an important step to increase the role, availability, and recognition of computer science in high schools. And given that CS teachers need various kinds of knowledge and skills to be successful in the classroom, a CS methods course is a necessary and important step to meet this goal. **C**

References

1. Ball, D.L. Breaking with experience in learning to teach mathematics: The role of a preservice methods course. *For the Learning of Mathematics* 10, 2 (Feb. 1990), 10–16.
2. Bell, T., Witten, I.H., and and Fellows, M. *Computer Science Unplugged: An Enrichment and Extension Programme for Primary-Aged Children*. Computer Science Unplugged, Canterbury, New Zealand, 2006.
3. Bennedsen, J., Caspersen, M.E., and Kolling, M. *Reflections on the Teaching of Programming*. Springer, 2008.
4. Burger, K.R. Teaching two-dimensional array concepts in Java with image processing examples. In *Proceedings of the 34th SIGCSE Technical Symposium on Computer Science Education*. ACM, New York, 2003, 205–209.
5. Dickinson, D. and Caswell, L. Building support for language and literacy in preschool classrooms through in-service professional development: Effects of the Literacy Environment Enrichment Program (LEEP). *Early Childhood Research Quarterly* 22, (2007), 243–260.
6. Eagle, M. and Barnes, T. Wu's castle: Teaching arrays and loops in a game. *SIGCSE Bulletin* 40, 3 (Mar. 2008), 245–249.
7. Feiman-Nemser, S. Teacher learning: How do teachers learn to teach? In *Handbook of Research on Teacher Education: Enduring Questions in Changing Contexts (Third Edition)* M. Cochran-Smith, S. Feiman-Nemser, and D.J. McIntyre, Eds. Routledge, New York, 2008, 697–705.
8. Gal-Ezer, J. and Stephenson, C. The current state of computer science in U.S. high schools: A report from two national surveys. *Journal for Computing Teachers*, Spring 2009.
9. Hazzan, O., Lapidot, T., and Ragonis, N. *Guide to Teaching Computer Science*. Springer, 2011.
10. Horstmann, C. *Java Concepts*. Wiley, Hoboken, NJ, 2010.
11. Lapidot, T. and Hazzan, O. Methods of teaching computer science course for new prospective teachers. *Inroads—The SIGCSE Bulletin* 35, 4 (2003), 29–34.
12. Rodgers, C. Defining reflection: Another look at John Dewey and reflective thinking. *Teachers College Record* 104, 4 (Apr. 2002), 842–866.
13. Shulman, L.S. Those who understand: Knowledge growth in teaching. *Educational Researcher* 15, 2 (Feb. 1986), 4–31.

Aman Yadav (yadav0@purdue.edu) is an associate professor of Educational Psychology in the Department of Educational Studies and also has a courtesy appointment in the Department of Computer Science at Purdue University, West Lafayette, IN.

John T. Korb (jtk@purdue.edu) is the assistant head of the Department of Computer Science at Purdue University, West Lafayette, IN.

Copyright held by author.

c See http://www.iste.org/Libraries/PDFs/NCATE_ISTE_csed_2002.sflb.ashx.


Computing Ethics

Societal Implications of the Emerging Smart Grid

Seeking solutions to concerns that go beyond the engineering of the smart grid.

LIKE OTHER COMPONENTS of the nation's infrastructure the U.S. electrical power grid is deteriorating; the annual number of large power outages has been increasing since the late 1990s.¹ Though not as catastrophic as the recent blackouts in India, the increasing numbers, duration, and impact of power failures across the U.S. due to the degradation of the grid have severe implications for the energy-intensive way of life, economic stability, and even national security. The cost of neglect is high; a report commissioned by the Edison Foundation estimated that to retrofit the U.S. electricity infrastructure including new generators and new power delivery systems will require approximately \$1.5 trillion over 20 years⁴ with more than half of this investment going to transmission and distribution facilities.

The proposed solution is widely known as the "smart grid." The increasing occurrences of outages and instances of cyber intrusions between 2000 and 2008 were considered so threatening to U.S. economic viability and security that the federal government, as part of the American Recovery and Reinvestment Act of 2009, earmarked more than \$3.3 billion in smart grid technology development grants and an additional \$615 million for smart grid storage, monitoring, and technology viability as an initial investment in building the smart grid. In addition, utilities have begun to mount demonstration projects and government and professional societies have begun the devel-



opment of smart grid standards.

The smart grid will be comprised of three fundamental structural elements: replacement of aging core physical infrastructure items including transmission lines and switching equipment with more efficient and reliable newer technologies; two-way distributed and loosely coupled supply and demand connectivity to the grid, which allows consumers to supply electricity through technologies such as photovoltaic cells and wind power; and, most importantly, highly optimized two-way information and communication technology (ICT) sys-

tems architectures and networks that control the grid through process- and rule-based programs to match power demand with supply in order to improve efficient use of energy resources.

The fundamental differences between the existing grid and the smart grid are the ICT and distributed connectivity capabilities. While these innovative features of the smart grid hold great potential for improved energy efficiency through better management of consumer demand and improved stewardship of energy resources including greater utilization of renewable generation, they also pose a number of social

and ethical challenges including: protecting the privacy of consumer usage information; securing the grid from attacks by foreign nations, terrorists, and malevolent hackers; and ensuring social justice both in terms of access and cost of electric power service. As with many new technologies the engineers engaged in developing the smart grid often overlook such issues or only turn to considering them once the technical standards and specifications have been settled. Failure to address these issues in a timely manner, however, may result in delays in establishing the smart grid and undermine its potential. Engineers and others involved in developing the smart grid need to examine ways to address organizational, social, and ethical dimensions that distributed generation and more extensive efforts to influence consumer usage patterns will raise. The cost of doing so would amount to an insignificant fraction of the projected necessary investments.

Privacy and Security Issues

As is the case for many other modern ICT applications such as the Internet and geographical positioning system (GPS), ensuring consumer privacy will be a challenge for the smart grid. Up until now our personal energy usage had been recorded by simple consumption metrics such as kilowatt hours measured using a conventional meter attached to a home or business. In the initial transition to a smart grid, utilities have begun to install “smart meters” that can provide feedback to the utility and customers on such factors as time of use of electricity. Since every appliance has a unique “load signature,” smart meter data can be analyzed to determine the types of appliances and other equipment consumers are using.³ In the future, as more demand-side technologies are developed, the smart grid could have the capability to monitor and control the usage of every plugged-in electrical device, which would allow the electric utility to turn the device off during times of peak demand to balance load across the grid. For the privilege of acquiring data and controlling consumer electrical devices utility companies may charge a reduced rate. Alternatively, rate structures that vary by time of day or fuel source (coal vs. wind, for example) may be instituted

The fundamental differences between the existing grid and the smart grid are the ICT and distributed connectivity capabilities.

in order to influence consumer energy usage behaviors.

As we move from theory to design, the emerging smart grid will become a vast ICT network populated with a diverse set of data acquisition devices capable of tracking the source, ownership, performance, and behavioral characteristics of each connected component. The smart grid technologies with the potential to be privacy invasive include “smart” power meters, energy monitoring and control software programs, and monitoring chips built into devices that consume electricity. In addition to control and monitoring functions, however, the smart grid will have the ability to collect, aggregate, and store individual consumer usage data such as the temporal pattern of electricity usage and the number, type, and usage of electrical appliances and electronic devices. Analysis of this data could reveal such information as home occupation patterns, the number of occupants, and the manufacturer and usage of individual devices—valuable to utility planners but additionally to marketing agencies, insurance companies (property, health, and life) and, potentially, criminals (for example, outsiders may be able to tell when a home is occupied, determine the type of security system, and learn other sensitive information). Much like the data acquired by supermarket bar-code scanners and loyalty cards, data on specific devices in homes and consumers’ patterns of energy use will become a prized resource. Electric utilities or third-party vendors may sell personal

data to other organizations to defray costs or simply to increase profits.

The PowerMeter application being developed by Google is an example of how third-party vendors may become involved in the management of smart grid data. An Internet-based application, PowerMeter receives information from utility smart meters and energy management devices and provides customers with access to their home electricity consumption on their personal iGoogle home page. Google is only one of many data-hungry organizations racing to develop smart grid monitoring equipment and data systems.

Of course, like supermarket loyalty cards, utility customers may be willing to give up some of their personal data if they think it is being used benignly and if they are getting something in return (such as reduced prices or rates). Up to now, however, utilities have not had to deal with consumer energy usage data on this scale; they and the public utility commissions that regulate them may be unwilling to incur the added expense of protecting consumer data from illegitimate uses or reassuring consumers that this data is protected adequately. The implications have not escaped the privacy watchdogs or even high-ranking U.S. federal government officials. Indeed, former Commerce Secretary Gary Locke warned that privacy concerns might be the “Achilles’ heel” of the smart grid. Achieving public acceptance of the smart grid may prove difficult if privacy concerns are not addressed in a proactive manner.

Unsurprisingly, many security aspects of the smart grid look like those of the Internet. Although the Internet has not been designated as the primary source of ICT communications, the smart grid will more than likely mature into a system that will utilize the Internet as its backbone. To secure both the informational and power-carrying capacity of the smart grid two important features must be addressed: the physical security of power and ICT networks and equipment and the security of huge databases and computers that analyze the data. The smart grid of the future will integrate both these networks creating the ability for either one to cause disruption to the other. Examples abound where highly automated systems have been brought to a halt or

damaged by failures or security breaches in their ICT backbones (such as failures in automated securities trading, cyber warfare damage to Iran's centrifuges for nuclear fuel enrichment, and malevolent hacking resulting in infiltration and shutdown of corporate and government Web sites).

Security breaches in the smart grid could lead to brownouts or even blackouts, and could cause serious, long-term damage to power generation, transmission, and distribution equipment. With the integration of power and ICT networks, power delivery components and even everyday power devices (such as appliances) will become nodes on the Internet. In the future, cyber attacks such as denial-of-service or virus attacks could cause outages in the smart grid and limit electricity supplies, including critical services such as infrastructure and public safety. These attacks could originate anywhere in the world and could start as easily as introducing false data regarding energy usage across many nodes.

What do these concerns mean for the development of security mechanisms, policies, and practices to secure the smart grid? There will be pressure to introduce a wider range of surveillance technologies; such technologies are already at the forefront of many heated debates regarding the intrusion of local, state, and federal governments, and also corporations, into the daily lives of individuals. Security and surveillance systems bring their own data needs, which promise to further erode personal freedoms, including privacy.

Pricing and Access

Though not as obvious as privacy and security issues, the smart grid also poses potential problems for equitable pricing and access to electric power service. The nature of these impacts will depend on whether consumer energy usage is left under utility control or consumers are allowed to make their own usage decisions under variable pricing schemes. The former case would limit consumer autonomy. One utility, for example, has already proposed that it be permitted to control customers' thermostats. Variable pricing, on the other hand, would place an energy management burden on all residential consumers. Those with lower educational levels, limited

Unsurprisingly, many security aspects of the smart grid look like those of the Internet.

Internet access or computer skills, medical or cognitive impairments, or those who simply lack time, resources, or motivation to manage their usage patterns could be at a disadvantage. Both cases will require innovative ratemaking and oversight by public utility commissions and greater coordination and standardization within and among retail service areas. Though smart meter experiments are just in the beginning stages, there have already been regulatory and legal controversies over such issues as required prepaid service plans for low-income consumers and alleged price gouging under mandatory switches to smart meters.

Conclusion

Achieving the smart grid's potential while tending to privacy, security, and equity concerns should begin with the realization that the smart grid is a complex sociotechnical system that requires solutions that go beyond the engineering of the grid. Solutions must include thoughtful deliberation by federal and state regulatory agencies, flexible utility responses in addressing consumer concerns and, most importantly, an engineering culture that recognizes and addresses the societal implications of the smart grid upstream in the R&D process and as standards are being developed.

For example, while The National Institute of Standards (NIST) highlighted privacy concerns in a recent report,⁷ the U.S. federal government has yet to enact any smart grid privacy legislation or regulations. On the other hand, The California Public Utilities Commission's (CPUC) 2011 decision on protecting privacy and security of consumer data is a landmark ruling that should provide a strong template for other state commissions.⁵

One solution for addressing customer's concerns regarding the smart grid is to provide opt-out options, such as Pacific Gas and Electric's proposal to permit customers worried about the environment, health, and safety effects of smart meter wireless radio signals to request that the signals be shut off (albeit with a charge for conventional meter reading).² Willingness to provide such options may be necessary to ensure public trust of utilities as the smart grid develops.

As in the case of the human genome project and nanotechnology, where the U.S. federal funding agencies earmarked a percentage of research funds to examine such issues,⁶ there is an urgent need to examine the societal implications of the smart grid concurrent with its development. Failure to do so will further threaten civil liberties in the information age and is likely to pose substantial barriers to public acceptance. 

References

1. Amin, M. and Schewe, P.F. Preventing blackouts. *Scientific American* 296, 5 (May 2007), 60–67.
2. Barringer, F. Pacific Gas offers solution to turn off smart meters. *The New York Times* (Mar. 24, 2011).
3. Bleicher, A. Privacy on the smart grid. *IEEE Spectrum*, online edition; <http://spectrum.ieee.org/energy/the-smarter-grid/privacy-on-the-smart-grid/>.
4. Chupka, M. et al. Transforming America's Power Industry: The Investment Challenge 2010–2030. Prepared by the Brattle Group for The Edison Foundation, November 2008.
5. CPUC. Decision Adopting Rules To Protect The Privacy And Security Of The Electricity Usage Data Of The Customers Of Pacific Gas And Electric Company, Southern California Edison Company, And San Diego Gas & Electric. Decision 11-07-056; http://docs.cpuc.ca.gov/PUBLISHED/FINAL_DECISION/140369.htm.
6. Mills, K. and Fleddermann, C. Getting the best from nanotechnology: Approaching social and ethical implications openly and proactively. *IEEE Technology and Society Magazine* 24, 4 (Winter 2005), 18–26.
7. NIST. *Guidelines for Smart Grid Cyber Security: Vol. 2, Privacy and the Smart Grid*. NISTIR 7628, August 2010.

Timothy Kostyk (Timothy.Kostyk@asu.edu) is a Ph.D. student in Human and Social Dimensions of Science and Technology at Arizona State University in Tempe, AZ. He has 25 years of experience as an enterprise architect working for companies including Sprint, Carlson, and IBM along with working with the Open Group in the development of TOGAF 9.

Joseph R. Herkert (joseph.herkert@asu.edu) is the Lincoln Associate Professor of Ethics and Technology at Arizona State University in Tempe, AZ.

The authors thank Rachelle Hollander and two anonymous reviewers for providing helpful comments on earlier drafts of this column.

Development of the material from which this column was derived was supported by grants from the U.S. National Science Foundation (Awards SES-0921806 and SES-1032966). The views expressed in this column are those of the authors and do not necessarily represent the views of the National Science Foundation or the U.S. government.

Copyright held by author.

Viewpoint

Keeping Technology Promises

Considering new models for educational technology and methods.

THERE IS A collapse of confidence under way in U.S. colleges and universities. It is a collapse that has been documented in what seems like a steady stream of recent reports and books,^{1,3,7,11} including my own.⁵ Amid the many dire warnings there is one bright thread: advances in information technology are often viewed as a pathway to rebuilding public confidence in higher education by reducing costs, expanding access, improving outcomes, and increasing financial transparency. If technology could help rebuild public confidence, higher education would be better off for it, but without more engagement from the research community in attacking the problems facing the nation's colleges and universities I am not optimistic that will happen.

It would not be the first time that technologists have promised to improve education. The historical intersection of computing research and education is filled with examples that were more about computing and less about education. The result: a stream of educational technology that—at great expense—missed the mark, ultimately making promises that could not possibly be kept. Educational institutions were for the most part unfazed, but this time is different. The scale and size of the underlying problems are enormous. The pace of change is frightening, and there is genuine fear that higher education is an economic bubble that is about to burst. This time the system cannot withstand the shock of another generation of unkept technology promises.



I have found myself in recent months increasingly involved in strategic planning sessions, media conversations, and public debate of measures that might help stem the tide of bad news about high tuition, student debt, educational quality, and low completion rates. It is a national conversation in the U.S. that mainly engages economists, sociologists, and professionals who specialize in university administration. It is an important discussion and although there is a compelling argument for rapid innovation to disrupt the status quo, academic computer science has been on the sidelines for much of it.

Sentiment is strong for sweeping changes in higher education. Everything from financial transparency and greater scrutiny of intercollegiate athletics to clearer productivity and accountability standards is on the table. There is no shortage of “Big Fix” solutions that—improperly formulated—run the risk of doing massive damage to the great U.S. system of public and private universities.

Technology is what I would call a “Small Fix” solution. Technology’s disruptive power does not necessarily require a Big Fix as a prerequisite. The right innovations—and the innova-

tions they in turn spawn—could push the discussion in a very different direction. Let's take the issue of reducing classroom costs as an example. In the public mind, technology is the surest path toward reducing overall costs.

It is a promising idea, but technology has had remarkably little impact on classrooms. For the past millennium classrooms have consisted of spaces for a teacher to stand, facing rows of seated students. Chalkboards did not make a classroom appearance until 1801. They were an immediate hit. They were inexpensive, easy to use, and they did not require much upkeep. Despite a constant flow of gadgets and renewed technology promises, the blackboard was the last invention that had such obvious pedagogical value that it became a ubiquitous classroom fixture.

A recurring technology promise is to reduce costs by replacing human teachers with automata. It is an elusive goal, but that has never kept us from designing computers to make live classrooms more efficient. The most spectacular attempt was called PLATO. Backed by Control Data Corporation, whose CEO Robert Norris predicted that most of the company's revenues would come from PLATO and related products and

services, the total R&D investment in PLATO soon topped a billion dollars, a cost that CDC tried to recover by unrealistic pricing to universities. By the time Norris stepped down as Control Data CEO in 1986, the company was looking for an exit strategy for PLATO and the education market.

In the wake of PLATO, dozens of projects made more determined attempts to marry technology with traditional classrooms—in effect, to define the classroom of the future. But, aside from some minor tweaks to the blackboard's user interface, classrooms have remained virtually unchanged. Underneath it all, after few of the technology promises were kept, the classroom of the future had little to do with education. When I asked Classroom 2000 project director Gregory Abowd—who reluctantly shuttered the doors to his laboratory in 2002—about the apparent resistance of classrooms to change, he disagreed with my characterization. “I don't think the classrooms had been immune to technology,” he told me. *There was lots of technology, but much of it was in the aid of the presenter of the material and not for the students who were struggling to keep pace with the increased flow of information.*⁵

And it is not only classroom technology that failed to live up to its promise. A parade of Learning Management Systems has given us a glimpse of how administrators would like to stitch together content, back-office infrastructure, and classroom delivery into the kind of enterprise-quality software behemoths that keep large corporations humming in compliance with hundreds of business and market constraints. It is an idea that has been rejected in certain terms by radical innovators like education technologist and self-described EduPunk Jim Groom at the University of Mary Washington in Fredericksburg, Virginia. He would like to see a very different approach: *The whole idea is a reaction to the over-engineered, badly designed, and intellectually constraining technology that has been foisted onto the American higher education system as a substitute for deep reflection about what the universities should be evolving into.*⁵

What are universities evolving into? Nobody knows for sure, but we know what is not working today. We in fact

know a lot about the Failed Assumptions of our current system:

- ▶ A group-oriented vision of an instructor broadcasting to a classroom of pupils, passive except for recitations and exams.

- ▶ A factory model of efficiency in which 18–24-year-old cohorts with uniform interests and abilities are co-located and experience education in lockstep fashion.

- ▶ A language and culture of assessment that seems borrowed from a century in which a fascination with quality on the factory floor seeped into the administration of universities and their programs.

It is a common meme among non-technologists that technology is responsible for depersonalizing and sterilizing education. An impersonal, sterile learning experience is one of the failures of education, but it is difficult to blame technology for that. It is much more likely the real fault lies with the Failed Assumptions.⁴ They are certainly what provoke rage among traditionalists like Humanities professor Laurie Fendrich: *Outcomes-assessment practices in higher education are grotesque, unintentional parodies of both social science and “accountability.” No matter how much they purport to be about “standards” or student “needs,” they are in fact scams run by bloodless bureaucrats who, steeped in jargon like “mapping learning goals” and “closing the loop,” do not understand the holistic nature of a good college education. For all the highfalutin pronouncements accompanying the current May Day parade of outcomes assessment, in the end they boil down to a wholesale abandonment of the very idea of higher education.*⁵

Here is a set of principles, the basis for a set of assumptions for educational technology. It is not a complete list, but it has been enough to start a discussion at Georgia Tech, where the Office of the Provost has organized to place the newly chartered Center for 21st Century Universities at the center of a new ecosystem. It is an attempt to inject engineering-style experimentation into educational innovation by actively identifying, promoting, and supporting many—often competing—approaches to change.¹⁰ Each of the principles summarizes a movement in higher education, and together they constitute a technology-driven change agenda.

Coming Next Month in COMMUNICATIONS

*Moving Beyond
the Turing Test*

Alan Turing Remembered

*Theory of Algorithmic
Self-Assembly*

SPDYing Up the Web

Q&A with Sanjeev Arora

*An Introduction to Data
Representation Synthesis*

*Natural Algorithms
and Influence Systems*

Also, the latest news in quantum computing, zoomable user interfaces, and disruptive education.

► **Open CourseWare and Open Certification:** Universities have over the last decade lost their stranglehold as gatekeepers. Traditional universities that hold content too closely will find their value eroded.

► **Open and Democratic Systems:** Universities cannot beat the economies of scale of a global market and will have to adapt to whatever technology their stakeholders use to connect to courses, professors, and learning networks.

► **Digital Identities:** New college students approach their colleges with existing digital identities, and it will be the role of the university to recognize, preserve, extend, and enhance those identities.

► **Ascendance of Learning Communities:** Web-based delivery, new social theories of capital formation and flow, and the explosive growth of both students and schools in a world that has been flattened by economics and politics enables and rewards global learning communities. These communities challenge the exclusive authority of traditional campuses.

► **Transformative Power of Technology on Content:** It is a unique capability of information technology to act on itself, to discover hitherto hidden patterns, or even to accelerate the creation of new ideas, theories, and ways of thinking about the world.

Georgia Tech is unusual in its institutional embrace of disruptive change, but it is hardly alone. There is already a hotbed of innovation surrounding some of these principles. For example, MIT's intention to offer inexpensive credentials for satisfactory completion of its online offerings will certainly require business models and platform technologies that do not yet exist.⁸

The departure⁶ of Stanford faculty whose Massive Open Online Course (MOOC) drew tens of thousands of students raises profound questions about how technology hollows out the value proposition of traditional institutions. Startups like OpenStudy use the value systems of online games and social networks to redefine the idea of a scholarly community.⁹

Still others aim to replace expensive, process-heavy learning management systems with lightweight open publishing models. The language of federated identities and intelligent tutoring

What are universities evolving into? Nobody knows for sure, but we know what is not working today.

seems to be everywhere. These are not ideas that are aimed at new gadgets or quick fixes. They threaten the status quo. They are important, but they do not represent a new wave of research.

Most innovation is not taking place in computer science. In fact, areas like e-textbooks—low-hanging fruit for dramatic transformation—have been remarkably untouched by computing researchers. This is especially important for publications like *Communications*, because you would expect fundamental advances that transform content to be first visible here. There is a vanguard of change in scientific publishing but computing is not yet part of it.

UCSD Pharmacology professor and winner of the 2007 Microsoft Research Jim Gray e-Science award Phil Bourne has been at the head of an Open Science movement that has used new publication technology models to transform the pace of scientific discovery. A founder of PLoS, the Public Library of Science, Bourne has a deep belief that the idea of scientific text as a static object is already obsolete.

Bourne refers to this as “unleashing the full power of the Internet to transform research” by transforming the way science is reported and communicated to students² Bourne, for example, is a founder of SciVee.tv, a Web 2.0 platform for synchronizing written text and video in what bioinformaticists call PubCasts. This simple idea fundamentally alters the workflow of scientists, but it also places new burdens on authors who have to adapt to an unfamiliar way of authoring text.

PubCasting is an example of the transformative power of technology on content and makes fundamental use of open content and open system, but, to achieve its full potential, educators will have to teach students how to com-

municate using these new models. Our publications will have to embrace the new technologies.

How fitting it would be for *Communications* to become transformative for educational technology. It might be the seed for innovation that would move computing to the center of the U.S. national debate about the fate of colleges and universities and for once keep technology's promise.

Addendum

In the 12 months since this column was written, higher education has been rocked by computer scientists at top research universities. Stanford spin outs, Coursera and Udacity, and edX, a Harvard/MIT joint venture, are the kinds of experiments I called for in my original column submission and their principals have taken seats at the very tables I cited. These are important experiments, but they do not come close to scratching the surface of what computing technology might accomplish. It remains my hope that computing researchers will engage in the process of redefining higher education. 

References

1. Arum, R. and Roksa, J. *Academically Adrift, Limited Learning on College Campuses*. University of Chicago Press, 2011.
2. Bourne, P.E. and Fink, J.L. Reinventing scholarly communication for the electronic age. *CT Watch Quarterly* 3, 3 (Aug. 2007), 26–31.
3. Christensen, C. and Eyring, H. *The Innovative University: Changing the DNA of Higher Education from the Inside Out*. Jossey-Bass, 2011.
4. Christensen, C. and Horn, M. *Disrupting Class: How Disruptive Innovation will Change the Way the World Learns*. McGraw-Hill, 2008.
5. DeMillo, R.A. *Abelard to Apple: The Fate of American Colleges and Universities*. MIT Press, 2011.
6. DeSantis, N. Tenured professor departs Stanford U., hoping to teach 500,000 students at online start-up. *Chronicle of Higher Education* (Jan. 23, 2012); <http://chronicle.com/blogs/wiredcampus/tenured-professor-departs-stanford-u-hoping-to-teach-500000-students-at-online-start-up/35135>.
7. Is college worth it? College presidents, public assess value, quality and mission of higher education. *Social and Demographic Trends*. Pew Research Center, Washington, D.C.; <http://pewsocialtrends.org>.
8. Lewin, T. MIT expands its free online courses. *New York Times* (Dec. 19, 2011).
9. Perry, M. Startup aspires to make the world one big study group. *Chronicle of Higher Education* (Sept. 8, 2010); <http://chronicle.com/blogs/wiredcampus/start-up-aspires-to-make-the-world-one-big-study-group/26780>.
10. Selinger, J. If engineers were to rethink higher education. *Chronicle of Higher Education* (Sept. 27, 2011); <http://chronicle.com/blogs/next/2011/09/27/if-engineers-were-to-solve-higher-eds-future/>.
11. *Time is the Enemy*. Complete College America, 1250 H Street, Washington D.C.

Richard A. DeMillo (rad@gatech.edu) is a Distinguished Professor of Computing at Georgia Institute of Technology in Atlanta and the director of the Center for 21st Century Universities (C21U).

Copyright held by author.

Article development led by **acmqueue**
queue.acm.org

**A discussion with Jesse Robbins, Kripa Krishnan,
John Allspaw, and Tom Limoncelli.**

Resilience Engineering: Learning to Embrace Failure

IT IS VERY nearly the holiday shopping season and something is very wrong at a data center handling transactions for one of the largest online retail operations in the country. Some systems have failed, and no one knows why. Stress levels are off the charts while teams of engineers work around the clock for three days trying to recover.

The good news is that it is not a real disaster—though it could have been. Instead, it is an exercise designed to teach a company how to adapt to the inevitable: system failure. Things break; disaster happens; failure is real. Although no one—perhaps least of all software developers and systems engineers—likes to believe they cannot prevent failure, the key to preparing for it is first to accept it.

Many operations are turning to resilience engineering not in hopes of becoming impervious to failure, but rather to become better able to adapt to it when it occurs. Resilience engineering is a familiar concept in high-risk industries such as aviation and health care, and now it is being adopted by large-scale Web operations as well.

In the early 2000s, Amazon created GameDay, a program designed to increase resilience by purposely injecting major failures into critical systems semi-regularly to discover flaws and subtle dependencies. Basically, a GameDay exercise tests a company's systems, software, and people in the course of preparing for a response to a disastrous event. Widespread acceptance of the GameDay concept has taken a few years, but many companies now see its value and have started to adopt their own versions. This discussion considers some of those experiences.

Participants include **Jesse Robbins**, the architect of GameDay at Amazon, where he was officially called the Master of Disaster. Robbins used his training as a firefighter in developing GameDay, following similar principles of incident response. He left Amazon in 2006 and founded the Velocity Web Performance and Operations Conference, the annual O'Reilly meeting for people building at Internet scale. In 2008, he founded Opscode, which makes Chef, a popular framework for infrastructure automation. Running GameDay operations on a slightly smaller scale is **John Allspaw**, senior vice president of technological operations at Etsy. Allspaw's experience includes stints at Salon.com and Friendster before joining Flickr as engineering manager in 2005. He moved to Etsy in 2010. He also recently took over as chair of the Velocity conference from Robbins.

Google's equivalent of GameDay is run by **Kripa Krishnan**, who has been with the program almost from the time it started six years ago. She also



“

TOM LIMONCELLI

You said something quite extraordinary when you mentioned how quickly the executives at Amazon embraced the notion that failure is unavoidable.

”

works on other infrastructure projects, most of which are focused on the protection of users and their data.

Moderating this discussion is a Google colleague, **Tom Limoncelli**, who is a site reliability engineer. Well known in system administrator circles, he started out at Bell Labs and has written four books, most notably *Time Management for System Administrators* and *The Practice of System and Network Administration*.

TOM LIMONCELLI: Jesse, you've probably been involved in more GameDay exercises than anybody. What's the most important lesson you've taken away from that?

JESSE ROBBINS: More than anything else, I've learned that the key to building resilient systems is accepting that failure happens. There's just no getting around it. That applies to the software discipline, as well as to the systems management and architectural disciplines. It also applies to managing people.

It's only after you've accepted the reality that failure is inevitable that you can begin the journey toward a truly resilient system. At the core of every resilience program—whether it's what you find at Google, Facebook, Etsy, Flickr, Yahoo, or Amazon—is the understanding that whenever you set

out to engineer a system at Internet scale, the best you can hope for is to build a reliable software platform on top of components that are completely unreliable. That puts you in an environment where complex failures are both inevitable and unpredictable.

KRIPA KRISHNAN: We've learned a few things as well. The most important of those lessons is that an untested disaster recovery plan isn't really a plan at all. We also know now that if doing something is hard, repetition is going to help make it easier. At Google scale, even if there's only a fraction of a 1% chance of a failure occurring, that means it's a failure likely to occur multiple times. Our plan is to preemptively trigger the failure, observe it, fix it, and then repeat until that issue ceases to be one. In our most recent GameDay exercise, we retested some things that caused serious failures two to three years ago and were pleased to find they can now be resolved effortlessly.

We've also learned that real success doesn't come from just running a GameDay test once a year but instead from getting teams to test their services internally all year round. That said, GameDay gives us an opportunity to test some less-exercised links. For example, we design tests that require engineers from several groups



JESSE ROBBINS

You want to make these drills seem as real as possible in order to expose those special systems and backdoor boxes that some system administrators have been holding onto in case of emergency.



who might not normally work together to interact with each other. That way, should a real large-scale disaster ever strike, these people will already have strong working relationships established.

Another point of emphasis is that none of the functions people are asked to perform on GameDay is significantly different from what they would do on any other day. That way, there won't be anything unusual to remember or panic about should a real failure occur.

TL: Speaking of panic, what do these exercises look like from a software engineer's or system administrator's perspective? What do they actually go through during one of these exercises?

JR: The program I designed at Amazon began with a series of company-wide briefings advising everyone that we were about to do an exercise of a certain scale—say, something on the order of a full-scale data-center destruction. They didn't know which data center was going to be taken offline, and they didn't know exactly when it was going to happen, but they usually had a few months to make sure their systems were capable of surviving the sudden loss of a significant amount of capacity. The expectation was that they would use that time to make sure they had eliminated all the

single points of failure they could find.

Once the event actually came to pass, people would typically receive a notification as part of our standard incident-management process, often including some information specific to their sites. For example, we might tell them that certain servers were no longer responding or that their service had lost some amount of capacity. Hopefully, their own standard monitoring tools would have picked all that up, but there have been instances where the GameDay exercise ended up taking out those monitoring capabilities. That isn't exactly what you would like to see, but still it's one of the classic defects these exercises are really good at exposing.

Sometimes you also expose what we call "latent defects"—problems that appear only because of the failure you've triggered. For example, you might discover that certain monitoring or management systems crucial to the recovery process end up getting turned off as part of the failure you've orchestrated. You would find some single points of failure you didn't know about that way. But, as Kripa said, as you do more and more of these exercises in a progressively intense and complicated way, it does start to become just a regular, ordinary part of doing business.

TL: So as an engineer, I might be sitting at my desk when I get a page telling me that one of our data centers has just gone down. Would that be a page about something completely fictional, or are we talking about a situation where someone has actually disconnected a cable, so to speak?

JR: In the exercises I designed, we used real failures. We would literally power off a facility—without notice—and then let the systems fail naturally and allowed the people to follow their processes wherever they led. In one of those exercises, we actually drew on my fire-service background to concoct a simulated fire. I wrote out the timing to the minute for that according to when we would expect certain things to happen as part of a full-scale fire response. Then, posing as some of the facilities guys, we called people in operations to update them on what was happening.

My view is that you want to make these drills seem as real as possible in order to expose those special systems and backdoor boxes that some system administrators have been holding onto in case of emergency. Not everything in these exercises can be simulated, and once you start powering down machines or breaking core software components, the problems that surface are real. Still, it's important to make it clear that the "disaster" at the core of the exercise is merely simulated so people on the periphery don't freak out. Otherwise, what happens in the course of that exercise ought to feel just as real as possible.

JOHN ALLSPAUGH: Yes, and the exercise should also make people feel a little uncomfortable. The truth is that things often break in ways we can't possibly imagine. In the course of responding to those surprises, you get a chance to learn from your mistakes, of course. Ultimately, you also get an opportunity to gain confidence that the system you've built and the organization that's been built up around it are actually pretty resilient.

TL: How does this work at Google, Kripa?

KK: We normally give people three to four months' notice, telling them the event is set to take place within some particular week or month. The GameDay event itself is generally a

round-the-clock, 72- to 96-hour exercise. Even though people are unaware of the exact timing, they know there's a period coming when there will be a lot of disruptions and they're going to be expected to respond to each of those disruptions as if it were a real event.

The disruptions we orchestrate range on the one hand from technical failures such as powering down a full data center or purposely corrupting data on our back ends to exercises that test the human element—for example, creating a situation where an entire team is rendered incommunicado for 72 hours and other teams are forced to work around them to simulate what might happen following an earthquake. Tests involving less-severe issues might last only a few hours, whereas the larger-scale tests tend to run a few days. The idea is always to discover as much as we can about how the company performs under stress with reduced capacity over an extended period of time.

Just to give you a sense of scale, we typically have hundreds of engineers working around the clock for the duration of one of these exercises. We usually have one set of engineers who respond to the test and another set who act as "proctors" monitoring the test. The proctors keep a close eye on the communications flying back and forth across IRC (Internet relay chat) channels. We also staff "war rooms" to make sure things don't get so out of control that they end up making the situation worse or impacting the production environment.

Roughly speaking, during the first 24 hours of the test, it's all about the initial response. The big problems generally surface then. Between 24 and 48 hours, a lot of routine team-to-team testing takes place. Independent engineering teams write tests for their counterparts in other locations, and they end up doing a lot of bidirectional testing. Then, by the 72-hour mark, signs of exhaustion really start to show. And it turns out exhaustion and other human factors are an important part of what we test. That's because, in a real emergency, you might not have the option of handing off work at the end of your shift.

You asked earlier about whether

the paging that gets sent out is fictional or real. We actually do page people for real, but we also try to make it clear the page relates to a GameDay test. The last thing we want is to give a test priority over a real production issue. Still, we've had situations where we learned about the "latent defects" Jesse was talking about, such as when we discovered the paging infrastructure was located in a facility we'd brought down, which ended up causing a pretty serious ruckus.

TL: I think everybody goes through that. You know you're conducting a successful exercise when the first thing you learn is that some critical part of your response framework is in scope for the test.

KK: Yes, but the test is really successful only if, when you repeat it later, you're able to ensure that everything works fine. In this particular case, the offending test caused tons of pages to queue up in the infrastructure, with none of them getting out to the right people. Then, when we finally managed to undo the test, a lot of pagers got blasted with thousands of pages—with no way to figure out which ones were real and which ones were GameDay-related. So you know that got fixed!

Introducing GameDay scenarios into some of these Web-scale companies has initiated a difficult cultural shift from a steadfast belief that systems should never fail—and if they do, focusing on who's to blame—to actually forcing systems to fail. Rather than expending resources on building systems that don't fail, the emphasis has started to shift to how to deal with systems swiftly and expertly once they do fail—because fail they will.

Failure can be a hard sell, but converts may indeed be won over gradually when they see how much they can learn from GameDay exercises. Much of the value in running such an exercise comes from changing the collective mind-set of the engineers who designed and built the systems. It's not easy for them to watch as their systems fail and to see what consequences come of that. Over time, however, they start to gain confidence that the systems and practices they rely upon are actually pretty resilient.

Companies that buy into the GameDay philosophy also hope to invoke a more just culture—one in which people can be held accountable without being blamed, or punished, for failure.

TL: You have managed to establish yourself in the business of orchestrating disasters, but traditionally management has considered any and all outages to be completely unacceptable. Does this suggest a major cultural shift?

KK: We definitely went through a major cultural shift over the first couple of years of our disaster-recovery testing program. What really worked in our favor, though, was that we had a solid sponsor—our VP, Ben Treynor—who strongly believed the only way you can be sure something works as expected is to test it. But even with strong support like that, we still got some significant pushback. Some nonengineering departments in particular saw only the risks and the investments that would be involved, and some operations teams just couldn't see the advantage of GameDay-style testing over the continuous testing programs they already had under way. In fact, the most important predictor of an organization's willingness to cooperate proved to be how the people in that organization had handled previous failures. If those earlier outages had resulted in engineering investigations that didn't affix blame to individuals but instead just looked for root causes, the organization almost invariably proved eager to participate. On the other hand, if those earlier outages had led to hunts for the guilty, then the organization generally proved to be more reluctant.

Today, however, GameDay is embraced across the board. Our most recent exercise involved 20 times as many teams writing tests as we had five years ago, with participation coming from both technical and business groups. Now, when we find things that are broken, no one feels ashamed. They just accept learning about the problem as an opportunity to go back and fix it. Which is to say that everyone now seems to grasp that the whole point of the exercise is to find issues so remedies or course corrections can be implemented proactively.

JR: I had a somewhat different experience at Amazon back in 2003–2004 when the idea of horizontal scalability across unreliable hardware wasn't yet a fully formed concept—in fact, we were pretty much out at the hairy edge and feeling our way. We did recognize, however, that the scale and complexity of our failures were growing as our sites grew in size and number. That's when it dawned on us that there were probably only two approaches we could take: one would be to spend as much money as necessary to make things more reliable; and the other was that, as an organization, we could choose to embrace the idea that failure happens.

I enjoyed great executive support early on since it was fairly clear we weren't going to solve this challenge just by throwing money at it. So right away our executive team bought into the notion that triggering failures in a controlled manner represented an opportunity for us to learn some really important big lessons at a much lower cost and with far less disruption than would be the case if we just waited for problems to surface on their own.

That became the fundamental basis for running a GameDay-type exercise. That is, you can't choose whether or not you're going to have failures—they are going to happen no matter what—but you can choose in many cases when you're going to learn the lessons. The thing is, you really want to be able to expose the defects in your systems on your terms and at the time of your choosing, and there's just no other way to find some of those problems than to trigger failures.

Once an organization buys into that thinking, the culture tends to change pretty quickly. People who go through the process, difficult as it might be, find something valuable, and many of them soon become powerful advocates. It generally doesn't take long before you've got a fair number of people who accept that failure happens and so become frankly quite willing to have the fun of breaking stuff in order to expose the problems they know are lurking in there somewhere. Out of that comes a new operational culture, but it's one that can be built only through a series of exercises.

You can't do just one of these exer-

cises and then be done with it. This is one of the great lessons we took from the fire service, where it's understood you have to keep training and drilling regularly to develop the operating competencies that can only come over time.

TL: I just have to point out that you said something quite extraordinary when you mentioned how quickly the executives at Amazon embraced the notion that failure is unavoidable. Traditionally executives have been taught that anytime a major system failure happens, the first thing they should do is to fire some people just to show everyone they're in control of the situation.

JR: We probably had a particular advantage at Amazon in that much of the organization already was experienced with process optimization around fulfillment centers. Basically, we already had something of a plant-operations mentality.

It also helped that around that same time we had a few big outages. As a result, people were especially open and receptive to doing something different. At companies where there hasn't been an outage for a while, it might prove a little harder to sell the GameDay concept because, you know, complacency builds. What I always recommend to people at companies that haven't accepted the wisdom of embracing failure is that they should keep an eye out for any failures that come up, along with any other issues that might give them the ammunition they need to argue for a small GameDay experiment. Of course, then, should the experiment prove to be at all successful, they ought to publicize the heck out of it by pointing out the problems they managed to address in a controlled manner instead of just waiting around for disaster to strike.

JA: There are some lessons in what you've just said for any organization that aims to be resilient—particularly those that rely on complex systems. And by resilient, I mean the ability to sustain operations before, during, and after an unexpected disturbance—something that's outside the bounds of what the designer had anticipated.

Erik Hollnagel, a pioneer in the area of resilience engineering, has noted that the four cornerstones of



JOHN ALLSPAW

In addition to looking hard at why you ended up suffering some particular failure, you might want to turn that around and ask why you aren't having failures all the time.



resilience are: (1) know what to expect (anticipation); (2) know what to look for (monitoring); (3) know what to do (how to respond); and (4) know what just happened (learning). That last one is where things like postmortems come into play.

The view Hollnagel articulates is a little different from the traditional one, which holds that screw-ups and outages result from mistakes made by certain individuals. When you're talking about a complex product, however, it's likely that any flaws actually derive from some failure in the overall design process or something that went wrong in one of the many different product-development steps along the way. Which is to say the responsibility for any problems that surface later really should be borne collectively. So when all the blame ends up being put on an individual—most likely an engineer—well, that's just ridiculously reductionist. It's like saying, "Here's the guy. He made the mistake, and everything came crumbling down because of that. So what's our problem here? He's our problem." That's nothing more than a classic illustration of hindsight bias.

The trick, of course, is to get people throughout the organization to start building their anticipation muscles by thinking about what might possibly

go wrong. You already see an aspect of this in the current craze for continuous deployment, since that focuses on how to keep systems up while failed components are swapped out for repairs or replacement. If one part of your Web site is broken, it shouldn't bring down the whole site, right? They don't shut down the Brooklyn Bridge just because one lane is kaput.

TL: That covers the anticipation aspect. What about some of the other cornerstones of resilience?

JA: One of the most important aspects—and one that gets far too little attention—is learning from failure. The traditional response has been: guy makes a mistake...site goes down...have a meeting...fire that dude. He must have been careless or negligent, or maybe there was even a willful violation of procedures.

But I'd like to believe our professional culture is going to come around to embracing the wisdom in holding people accountable without necessarily blaming them. The medical profession and the aviation industry have both benefited from that attitude. Beyond the obvious issues of basic fairness, believing that "human error" is the single root cause of some problem is just pure folly. First of all, when you're dealing with something as complex and inherently flawed as the Web



KRIPA KRISHNAN

We repeatedly emphasize to people that they should file bug reports about every single broken thing they've found. That applies to people processes, as well as to bugs in the code.



infrastructure, the very notion of having only a single root cause is laughable. Then to pass the problem off as something that might be fixed by firing someone means you really haven't learned anything—and that's sure not going to help you become more resilient in the future.

The only way you can learn something useful from a failure is to find out all you can about the actions that led up to the outage—and that means finding out why it seemed to make sense to take those particular steps at that particular time. There are some human-factors tests that can help with this. One that we use at Etsy is called the substitution test, and it comes in handy in figuring out why somebody decided to run a command that ended up bringing down the site. We'll grab another engineer who had no involvement in the problem situation, and we'll fill him in on the context and all the particulars known to the operator at the time. Then we'll ask that engineer, "What would you have done in that situation?" Almost every single time, he will tell us he would have run exactly the same command.

The problem therefore isn't due to a lack of training, a lack of intelligence, or any other personal failing. The solution isn't just to fire the guy. In fact, you want to keep that guy and

drill down into what led to the mistake. Find out why he thought it made sense to do what he did.

I've already announced publicly on several occasions that I'm never going to fire anybody for taking down a site I'm responsible for. I don't think a lot of other engineering leaders are willing to go that far. In fact, I think it makes a couple of people in my own organization pretty nervous.

JR: Yes, but in many cases, you can only learn from those mistakes.

JA: The flipside to this is that some see resilience engineering as being not only about looking at the future and trying to anticipate failures, but also about deconstructing scenarios and then putting the pieces back together again to better understand how we've responded in the past so we can think about how we might respond better in the future. That's just another way of bolstering our powers of anticipation.

The really mind-blowing thing whenever you do that is to consider all the times when your site didn't fail, even though it probably should have. In addition to looking hard at why you ended up suffering some particular failure, you might want to turn that around and ask why you aren't having failures all the time.

GameDay exercises start out by break-

ing things in ways that can be easily imagined—in fact, in ways that are typically scripted—but then things often end up going off that script in some rather significant and unexpected ways. This adds a level of complexity that people must learn to manage. Why did these things happen and how should they be dealt with? The volume of information collected from a GameDay exercise can be staggering. Then it becomes important to come up with ways to process all that information and distribute it to the appropriate people so it can be put to use when needed.

Human factors are also a complication in exercises of this magnitude. GameDays can go on for a few days, so keeping the participants focused can be a challenge. You don't want people to become complacent, so there's a lot of value in keeping the pressure on, to a reasonable degree. Yet people make mistakes in stressful circumstances, so you have to compensate for the exhaustion factor by rotating people in at appropriate intervals. Often people don't want to give up until they have solved the problem, though their decision-making abilities may have become degraded by that point.

TL: You've all said that flushing out latent defects is one of the primary motivations for triggering failures. Is there ever a time when you can declare victory?

KK: Not really. One reason is that for organizations at our scale, systems are constantly evolving as new layers are added, and just taking inventory of all those services alone can require considerable effort. Complexities are introduced as new capabilities are developed. And even more complexity is introduced whenever acquisitions require you to merge in new code bases, which, of course, only makes it harder to anticipate where something might break. It gets progressively harder to see where our dependencies are and what might lead to cascading failures. The only way to find those latent defects is to run exercises where we can trigger actual failures.

A few examples might be the best way to illustrate this. We've had situations where we brought down a network in, say, São Paulo, only to find

that in doing so we broke our links in Mexico. That seems totally bizarre on the face of it, but as you dig down, you end up finding some dependency no one knew about previously. And we still wouldn't know about it to this day if we hadn't caused that network to fail.

We also had a case once where we turned off a data center only to find a good percentage of our machines there wouldn't come back online even when we tried to power-cycle them. A full night and a day later, we figured out that we had run out of DHCP (Dynamic Host Configuration Protocol) leases. Testing tends to surface silly stuff like that as well, but even silly stuff can have dire consequences.

JR: For the most part, we've been talking about a class of failures you learn about only in the course of one of these outages, but we also see problems and failures crop up as we're working to recover from the outages. Since leaving Amazon, I've helped a number of other organizations and have run into a number of recovery issues related to scale and complexity. As organizations grow, the tools they have for configuration management generally don't keep up. They find themselves in situations where they need to deploy code to 1,000 new boxes, but that turns out to be something their software deployment system hasn't been sized for.

TL: How do you manage to process all that you've learned over the course of one of these exercises?

KK: Before, during, and after the test, we repeatedly emphasize to people that they should file bug reports about every single broken thing they've found. That applies to people processes, as well as to bugs in the code. Those of us who manage the exercise also file our own set of bug reports. We have a rotating team of about 50 people in war rooms around the world—predominantly volunteers—each focusing on different parts of the operation. We're all constantly taking down notes on Post-its and whiteboards as we find things that are broken.

Following the exercise, we spend a couple of weeks assimilating all that information into a report and sorting out a punch list of fixes. This is

possibly the hardest part of the job. Hundreds of services generate a thousand issues. This gets processed, and more importantly, prioritized. Complex companywide issues need buy-in from several organizations, whereas service-specific bugs just need to be filed against each team. Really, that's all there is to it. Once the right information has been put in front of the right people, we can—with trivial effort on our part—close out 80% to 85% of the issues raised by each GameDay exercise.

No matter how those changes end up being made, the most important point is for organizations to realize how important it is to constantly test their systems and the integration points between different teams, as well as between all the underlying technology. That's particularly important for Internet-based companies, since the infrastructure and the software they rely upon are continuously changing. There's just no way of knowing what your altered system is going to be able to sustain without putting it to the test, and that means triggering actual failures through GameDay-style exercises.

JA: It's also helpful for all of us to remember that the field of Web engineering is only about 11 years old. Meanwhile, some fields of engineering have many decades of experience in building and maintaining complex systems, so there's something to be gained from considering some of those other domains to see what we might learn in terms of design principles and best practices that can be adapted for our own purposes. GameDay exercises demonstrate that we're starting to do just that. 

Related articles on queue.acm.org

Scale Failure

George Neville-Neil

<http://queue.acm.org/detail.cfm?id=2147781>

Automating Software Failure Reporting

Brendan Murphy

<http://queue.acm.org/detail.cfm?id=1036498>

Fault Injection in Production

John Allspaw

<http://queue.acm.org/detail.cfm?id=2353017>

Article development led by **acmqueue**
queue.acm.org

Failures happen, and resilience drills help organizations prepare for them.

BY KRIPA KRISHNAN

Weathering the Unexpected

WHETHER IT IS a hurricane blowing down power lines, a volcanic-ash cloud grounding all flights for a continent, or a humble rodent gnawing through underground fibers—the unexpected happens. We cannot do much to prevent it, but there is a lot we can do to be prepared for it. To this end, Google runs an annual, companywide, multi-day Disaster Recovery Testing event—DiRT—the objective of which is to ensure that Google’s services and internal business operations continue to run following a disaster. DiRT was developed to find vulnerabilities in critical systems and business processes by intentionally causing failures in them, and to fix them before such failures happen in an uncontrolled manner. DiRT tests both Google’s technical robustness by breaking live systems, and our operational resilience

by explicitly preventing critical personnel, area experts and leadership from participating. Where we are not resilient but should be, we try to fix it. (See the accompanying sidebar by Tom Limoncelli.)

For DiRT-style events to be successful, an organization first needs to accept system and process failures as a means of learning. Things *will* go wrong. When they do, the focus must be on fixing the error instead of reprimanding an individual or team for a failure of complex systems.

An organization also needs to believe that the value in learning from events like DiRT justifies the associated costs. These events are not cheap—they require a sizable engineering investment, are accompanied by considerable disruptions to productivity, and can cause user-facing issues or revenue loss. DiRT, for example, involves the work of hundreds of engineering and operations personnel over several days; and things do not always go according to plan. DiRT has caused accidental outages and in some cases revenue loss. Since DiRT is a company-wide exercise, however, it has the benefit of having all the right people available at a moment’s notice to contain such events should they arise.

However, to benefit the most from such recovery events, an organization also needs to invest in continuous testing of its services. DiRT-style, large, companywide events should be less about testing routine failure conditions such as a single-service failovers or on-call handoffs, and more about testing complex scenarios or less-tested interfaces between systems and teams. Complex failures are often merely a result of weaknesses in smaller parts of the system. As smaller components of the system get tested constantly, failures of larger components become less likely.

A simple example is testing an organization’s ability to recover from the loss of a data center. Such a loss may be simulated by powering down the facility or by causing network links to



fail. The response would theoretically involve a sequence of events, from redirecting traffic away from the lost data center to a series of single-service failovers in some specific order. All it would take to choke the recovery process, however, is the failure of a single instance of a core infrastructure service—such as DNS (Domain Name System) or LDAP (lightweight Directory Access Protocol)—to failover. Testing the failover of such a service can and should happen continuously and should not have to wait for a DiRT event.

Growing the Program

A good way to kick off such an exercise is to start small and let the exercise evolve. It is quite easy to make this a large and complex affair right from

the start, which in turn will probably come with unexpected overhead and complications.

Starting small applies to not only the number of teams involved in the exercise, but also the complexity of the tests. A few easy-to-remember rules and a simple, repeatable exercise format go a long way toward engaging teams quickly. If not all teams buy in, then work with the few that do; and as the exercise proves itself useful, more teams will participate.

Google's experience serves as an example: DiRT in its original form focused only on testing critical user-facing services. The initial bar was that all major user-facing teams wrote tests and that the tests were safe and caused no disruption, although we did realize that some of the tests were not very

useful. This got teams “playing.” Over a few iterations, the exercise attracted many more teams and tolerated fewer low-quality/low-value tests.

The same can be said for test designs. While the quality of tests matters a lot and directly affects the value of the exercise, DiRT events do not have to begin with overly complicated tests or the perfect set of tests (they do not exist). DiRT started with individual groups testing failure scenarios specific to their service. The overarching “disaster” was merely theoretical. In a subsequent DiRT exercise, the first major outage tested was that of our primary source-control management servers, which exposed several non-replicated critical functions dependent on this system. As each piece was fixed, we progressed to a larger disaster

Google DiRT: The View from Someone Being Tested

There's no telling where the zombies might strike next.

BY TOM LIMONCELLI

This is a fictionalized account of a Google Disaster Recovery Testing (DiRT) exercise as seen from the perspective of the engineers responsible for running the services being tested. The names, location, and situation have been changed.

[Phone rings]

Me: Hello?

Mary: Hi, Tom. I'm proctoring a DiRT exercise. You are on call for [name of service], right?

Me: I am.

Mary: In this exercise we pretend the [name of service] database needs to be restored from backups.

Me: OK. Is this a live exercise?

Mary: No, just talk me through it.

Me: Well, I'd follow the directions in our operational docs.

Mary: Can you find the doc?

[A couple of key clicks later]

Me: Yes, I have it here.

Mary: OK, bring up a clone of the service and restore the database to it.

Over the next few minutes, I make two discoveries. First, one of the commands in the document now requires additional parameters. Second, the temporary area used to do the restore does not have enough space. It had enough space when the procedure was written, but the database has grown since then.

Mary files a bug report to request the document be updated. She also files a bug report to set up a process to prevent the disk-space situation from happening.

I check my email messages and see the notifications from our bug database. The bugs are cc'ed to me and are tagged as being part of *DiRT2011*. Everything with that tag will be watched by various parties to make sure it gets attention over the next few months. I fix the first bug while waiting for the restore to complete.

The second bug will take more time. We will need to add the restore area to our quarterly resource estimation and allocation process. Plus, we will add some rules to our monitoring system to detect whether the database size is nearing the size of the restore area.

Me: OK, the service's backup has been read. I'm running a clone of the service on it, and I'm sending you an instant message with a URL you can use to access it.

[A couple of key clicks later]

Mary: OK, I can access the data. It looks good. Congrats!

Me: Thanks!

Mary: Well, I'll leave you to your work. Oh, and I'm not supposed to tell you this, but at 2 P.M. there will be some... fun.

Me: You know my on-call shift ends at 3 P.M., right? If you happen to be delayed an hour...

Mary: No such luck. I'm in California and 3 P.M. your time is when I'll be leaving for lunch.

A minute after the exercise is over I receive an email message with a link to a post-exercise document. I update it with what happened, links to the bugs that were filed, and so on. I also think of a few other ways of improving the process and document them, filing feature requests in our bug database for each of them.

At 2 P.M. my pager doesn't go off, but I see on my dashboard that there is an outage in Georgia. Everyone in our internal chat room is talking about it. I'm not too concerned. Our service runs out of four data centers around the world, and the system has automatically redirected Web requests to the other three locations.

The transition is flawless, losing only the queries that were "in flight," which is well within our SLA (service-level agreement).

A new email message appears in my inbox explaining that zombies have invaded Georgia and are trying to eat the brains of the data-center technicians there. The zombies have severed the network connections to the data center. No network traffic is going in or out. Lastly, the message points out that this is part of a DiRT exercise and no actual technicians have had their brains eaten, but the network connections really have been disabled.

[Again, phone rings]

Mary: Hi! Having fun yet?

Me: I'm always having fun. But I guess you mean the Georgia outage?

Mary: Yup. Shame about those technicians.

Me: Well, I know a lot of them and they have big brains. Those zombies will feed for hours.

Mary: Is your service still within SLA?

I look at my dashboard and see that with three data centers doing the work normally distributed to four locations the latency has increased slightly, but it is within SLA. The truth is I don't need to look at my dashboard because I would have gotten paged if the latency was unacceptable (or growing at a rate that would reach an unacceptable level if left unchecked).

Me: Everything is fine.

Mary: Great, because I'm here to proctor another test.

Me: Isn't a horde of zombies enough?

Mary: Not in my book. You see, your SLA says that your service is supposed to be able to survive two data-center outages at the same time.

She is correct. Our company standard is to be able to survive two outages at the same time. The reason is simple. Data centers and services need to be able to be taken down occasionally for planned maintenance. During this window of time another data center might go down for unplanned reasons (such as a zombie attack). The ability to survive two simultaneous outages is called N+2 redundancy.

Me: So what do you want me to do?

Mary: Pretend the data center in Europe is going down for scheduled preventive maintenance.

I follow our procedure and temporarily shut down the service in Europe. Web traffic from our European customers distributes itself over the remaining two data centers. Since this is an orderly shutdown, zero queries are lost.

Me: Done!

Mary: Are you within the SLA?

I look at the dashboard and see the latency has increased further. The entire service is running on the two smaller data centers. Each of the two down data centers is bigger than the combined, smaller, working data centers; yet, there is enough capacity to handle this situation.

Me: We're just barely within the SLA.

Mary: Congrats. You pass. You may bring the service up in the European data center.

I decide to file a bug, anyway. We stayed within the SLA, but it was too close for comfort. Certainly we can do better.

I look at my clock and see that it is almost 3 P.M. I finish filling out the post-exercise document just as the next on-call person comes online. I send her an instant message to explain what she missed.

I also remind her to keep her office door locked. There is no telling where the zombies might strike next.

Tom Limoncelli is a site reliability engineer in Google's New York office.

ter involving a major “earthquake” in the Bay Area.

We simulated the earthquake by taking down a data center in the area that housed a number of our internal systems. While the outage uncovered several services that were singly homed, it also exposed other interesting dependencies. For example, to avoid being affected by the outage, some teams decided to failover services from the data center to their workstations. Since the “earthquake” occurred near Google headquarters in Mountain View, the testing team disconnected the Mountain View campus as well—which meant all these failovers had failed. Also, what many did not anticipate was that the data-center outage caused authentication systems to fail in unexpected ways, which in turn locked most teams out of their workstations.

When the engineers realized that the shortcuts had failed and that no one could get any work done, they all simultaneously decided it was a good time to get dinner, and we ended up DoS’ing our cafes. In keeping with the DiRT goals, several of these issues were fixed by the next test.

Today, production and internal systems, network and data-center operations, and several business units such as HR, finance, security, and facilities test during DiRT. In the most recent DiRT exercise, we brought down several data-center clusters, infrastructure hubs, and offices without notice. Most of the scenarios were resolved painlessly.

It is very important to mention that well before Google even considered the concept of DiRT, most operations teams were already continuously testing their systems and cross-training using formats of popular role-playing games. As issues were identified, fixes got folded into the design process. For many of these teams, DiRT merely provided a safe opportunity to test riskier failure conditions or less-tested interactions with other systems and teams.

What to Test

There are several angles to consider when designing tests for DiRT. The simplest case, as described earlier, is service-specific testing. This category tests that a service and its components

are fault-tolerant. These tests are usually contained, needing only the immediate team to respond, and they uncover technical and operational issues including documentation gaps, stale configurations, or knowledge gaps in handling critical emergencies. Ideally, these tests become part of the service’s continuous testing process.

More involved technical test cases create scenarios that cause multiple system failures in parallel. Examples include data-center outages, fiber cuts, or failures in core infrastructure that manifest in dependent services. Such tests have a lot more value if the team that designs them is cross-functional and incorporates technical leads and subject-matter experts from various areas in the company. These are the people who understand the intricacies of their services and are in excellent positions to enumerate dependencies and failure modes to design realistic and meaningful scenarios.

The goal of this category of tests is to identify weaknesses in the less-tested interfaces between services and teams. Such scenarios can be potentially risky and disruptive, and they may need the help of several teams to resolve the error condition. DiRT is an excellent platform for this category of testing since it is meant to be a companywide exercise and all teams necessary for issue resolution are available on demand.

An often-overlooked area of testing is business process and communications. Systems and processes are highly intertwined, and separating out testing of systems from testing of business processes isn’t realistic: a failure of a business system will affect the business process, and conversely a working system is not very useful without the right personnel.

The previous “earthquake” scenario exposed several such examples, some of which are described here. The loss of the Bay Area disconnected both people and systems in Mountain View from the world. This meant that teams in geographically distributed offices needed to provide round-the-clock on-call coverage for critical operations. The configuration change that was needed to redirect alerts and pages to these offices, however, depended on a system that was affected by the outage.

Even for these teams with fully global expertise, things did not go smoothly as a result of this process failure.

A more successful failover was an approvals-tracking system for internal business functions. The system on its own was useless, however, since all the critical approvers were in Mountain View and therefore unavailable. Unfortunately, they were the same people who had the ability to change the approval chain.

In the same scenario, we tested the use of a documented emergency communications plan. The first DiRT exercise revealed that exactly one person was able to find the plan and show up on the correct phone bridge at the time of the exercise. During the following drill, more than 100 people were able to find it. This is when we learned the bridge would not hold more than 40 callers. During another call, one of the callers put the bridge on hold. While the hold music was excellent for the soul, we quickly learned we needed ways to boot people from the bridge.

As another example, we simulated a long-term power outage at a data center. This test challenged the facility to run on backup generator power for an extended period, which in turn required the purchase of considerable amounts of diesel fuel without access to the usual chain of approvers at headquarters. We expected someone in the facility to invoke our documented emergency spend process, but since they didn’t know where that was, the test takers creatively found an employee who offered to put the entire six-digit charge on his personal credit card. Copious documentation on how something should work doesn’t mean anyone will use it, or that it will even work. The only way to make sure is through testing.

Of course, tests are of almost no value if no effort is put into fixing the issues that the tests surface. An organizational culture that embraces failure as a means of learning goes a long way toward getting teams both to find and to resolve issues in their systems routinely.

Risk Mitigation

DiRT tests can be disruptive, and failures should be expected to occur at

any point. Several steps can be taken to minimize potential damage.

At minimum, all tests need to be thoroughly reviewed by a cross-functional technical team and accompanied by a plan to revert should things go wrong. If the test has never before been attempted, running it in a sandbox can help contain the effects. The flip side to sandboxing, though, is that sometimes these environments may have configurations that are significantly different from those in production, resulting in less realistic outcomes.

There are ways of testing without disrupting services: at Google, we “whitelist” services we already know won’t be able to survive certain tests. In essence, they have already failed the test and there is no point in causing an outage for them when the failing condition is already well understood. While services can “prefail” and exempt themselves, there is no concept of “prepassing” the test—services have to make it through to “pass.”

A centrally staffed command center that understands and monitors all tests going on at any given time makes DiRT a safer environment for testing. When the unforeseen happens, the team in the command center (comprised largely of technical experts in various areas) jumps in to revert the test or fix the offending issue.

The Team

At DiRT’s core are two teams: a technical team and a coordination team.

The technical team is responsible for designing all major tests and evaluating all tests written by individual teams for quality and impact. The technical team is also responsible for actually causing the larger outages and monitoring them to make sure things do not go awry in the process. This is also the team that handles unforeseen side effects of tests.

The coordinators handle a large part of the planning, scheduling, and execution of tests. They work very closely with the technical team to make sure the tests do not conflict with each other and that preparation work (such as setting up sandboxes) for each of these tests is done ahead of DiRT.

Both teams populate the DiRT



Copious documentation on how something should work doesn’t mean anyone will use it, or that it will even work. The only way to make sure is through testing.



command center. At the helm is usually someone with a sufficiently large Rolodex. When not much is going on, the command center is filled with distractions; it houses very smart people with short attention spans who are low on sleep and high on caffeine. When things go wrong, however—and they do—they are alert, on target, and fully focused on firefighting and getting the error communicated, resolved, or rolled back—and, furthermore, filed for fixing.

The command center is also home to the person with one of the most fun 20% projects at Google: the storyteller who concocts and narrates the disaster, ranging from the attack of the zombies to a bizarre psychological thriller featuring an errant fortune-teller.

Conclusion

Whatever its flavor, disaster recovery testing events are an excellent vehicle to find issues in systems and processes in a controlled environment. The basic principle is to accept that failures happen and that organizations need to be prepared for them. Often, a solid executive sponsor and champion is instrumental in setting the right tone for the exercise. In Google’s case, VP of operations Ben Treynor has championed both learning from continuous testing and preemptively fixing failures.

It is true that these exercises require a lot of work, but there is inestimable value in having the chance to identify and fix failures before they occur in an uncontrolled environment. 

Related articles on queue.acm.org

Fault Injection in Production

John Allspaw

<http://queue.acm.org/detail.cfm?id=2353017>

Thinking Clearly about Performance

Cary Millsap

<http://queue.acm.org/detail.cfm?id=1854041>

Improving Performance on the Internet

Tom Leighton

<http://queue.acm.org/detail.cfm?id=1466449>

Kripa Krishnan is a technical program manager at Google who has been running the company’s disaster recovery program (DiRT) for six years. She also leads the Google Apps for Government effort. She is currently working on privacy and security infrastructure initiatives for Google Apps. Prior to Google, Kripa worked with the Telemedicine Program of Kosovo to set up telemedicine infrastructure and a virtual education network in the region.

© 2012 ACM 0001-0782/12/11 \$15.00



Disks lie. And the controllers that run them are partners in crime.

BY MARSHALL KIRK MCKUSICK

Disks from the Perspective of a File System


MOST APPLICATIONS DO not deal with disks directly, instead storing their data in files in a file system, which protects us from those scoundrel disks. After all, a key task of the file system is to ensure the system can always be recovered to a consistent state after an unplanned system crash (for example, a power failure).

While a good file system will be able to beat the disks into submission, the required effort can be great and the reduced performance annoying. This article examines the shortcuts that disks take and the hoops that file systems must jump through to get the desired reliability.

While the file system must recover to a consistent state, that state usually reflects the one that the file system was in some time before the crash. Often data written in the minute before the crash may be lost. The reason for this loss is that the file system has not yet had the opportunity to write that data to disk. When an application needs to ensure that data can be recovered after a crash, it does an `fsync` system call on the file(s) that contain the data in need of long-term stability. Before returning from the `fsync` system call, the file system must

ensure that all the data associated with the file can be recovered after a crash, even if the crash happens immediately after the return of the `fsync` system call.

The file system implements `fsync` by finding all the dirty (unwritten) file data and writing it to the disk. Historically, the file system would issue a write request to the disk for the dirty file data and then wait for the write-completion notification to arrive. This technique worked reliably until the advent of track caches in the disk controllers. Track-caching controllers have a large buffer in the controller that accumulates the data being written to the disk. To avoid losing nearly an entire revolution to pick up the start of the next block when writing sequential disk blocks, the controller issues a write-completion notification when the data is in the



track cache rather than when it is on the disk. The early write-completion notification is done in the hope the system will issue a write request for the next block on the disk in time for the controller to be able to write it immediately following the end of the previous block.

This approach has one seriously negative side effect. When the write-completion notification is delivered, the file system expects the data to be on stable store. If the data is only in the track cache but not yet on the disk, the file system can fail to deliver the integrity promised to user applications using the `fsync` system call. In particular, semantics will be violated if the power fails after the write-completion notification but before the data is written to disk. Some vendors eliminate this problem by using non-volatile memory for the track cache and providing microcode restart after power failure to determine which operations need to be completed. Because this option is expensive, few controllers provide this functionality.

Newer disks resolve this problem with a technique called *tag queueing*, in which each request passed to the disk driver is assigned a unique numeric tag. Most disk controllers supporting tag queueing will accept at least 16 pending I/O requests. After each request is finished—possibly in a different order than the one in which they were presented to the disk—the tag of the completed request is returned as part of the write-completion notification. If several contiguous blocks are presented to the disk controller, it can begin work on the next one while notification for the tag of the previous one is being returned. Thus, tag queueing allows applications to be accurately notified when their data has reached stable store without incurring the penalty of lost disk revolutions when writing contiguous blocks. The `fsync` of a file is implemented by sending all the modified blocks of the file to the disk and then waiting until the tags of all those blocks have been acknowledged as written.

Tag queueing was first implemented in SCSI disks enabling them to have both reliability and speed. ATA disks, which lacked tag queueing, could be run either with their write cache enabled (the default) to provide speed at the cost of reliability after a crash or with the write cache disabled, which provided the reliability after a crash but at a 50% reduction in write speed.

To escape this conundrum, the ATA specification added an attempt at tag queueing with the same name as that used by the SCSI specification: Tag Command Queueing (TCQ). Unfortunately, in a deviation from the SCSI specification, TCQ for ATA allowed the completion of a tagged request to depend on whether the write cache was enabled (issue write-completion notification when the cache is hit) or disabled (issue write-completion notification when media is hit). Thus, it added complexity with no benefit.

Luckily, serial ATA (SATA) has a new definition called Native Command Queueing (NCQ) that has a bit in the write command that tells the drive if it should report completion when media has been written or when cache has been hit. If the driver correctly sets this bit, then the disk will display the correct behavior.

In the real world, many of the drives targeted to the desktop market do not implement the NCQ specification. To ensure reliability the system must either disable the write cache on the disk or issue a cache-flush request after every metadata update, log update (for journaling file systems), or `fsync` system call. Both of these techniques lead to noticeable performance degradation, so they are often disabled, putting file systems at risk if the power fails. Systems for which both speed and reliability are important should not use ATA disks. Rather, they should use drives that implement Fibre Channel, SCSI, or SATA with support for NCQ.

Another recent trend in rotating media has been a change in the sector size on the disk. Since the time of their first availability in the 1950s until about 2010, the sector size on disks has been 512 bytes. Beginning in 2010, disk manufacturers began producing disks with 4,096-byte sectors.

As the write density for disks has increased over the years, the error rate per

bit has risen, requiring the use of ever-longer correction codes. The errors are not uniformly distributed across the disk. Rather, a small defect will cause the loss of a string of bits. Most sectors will have few errors, but a small defect can cause a single sector to experience many bits needing correction. Thus, the error code must have enough redundancy for each sector to handle a high correction rate even though most sectors will not require it. Using larger sectors makes it possible to amortize the cost of the extra error-correcting bits over longer runs of bits. Using sectors that are eight times larger also eliminates 88% of the sector start and stop headers, further reducing the number of nondata bits on the disk. The net effect of going from 512-byte to 4,096-byte sectors is a near doubling of the amount of user data that can be stored on a given disk technology.

When doing I/O to a disk, all transfer requests must be for a multiple of the sector size. Until 2010, the smallest read or write to a disk was 512 bytes. Now the smallest read or write to a disk is 4,096 bytes.

For compatibility with old applications, the disk controllers on the new disks with 4,096-byte sectors emulate the old 512-byte sector disks. When a 512-byte write is done, the controller reads the 4,096-byte sector containing the area to be written into a buffer, overwrites the 512 bytes within the sector that is to be replaced, and then writes the updated 4,096-byte buffer back to the disk. When run in this mode, the disk becomes at least 50% slower because of the read and write required. Often it becomes much slower because the controller has to wait nearly a full revolution of the disk platter before it can rewrite a sector that it has just read.

File systems need to be aware of the change to the underlying media and ensure they adapt by always writing in multiples of the larger sector size. Historically, file systems were organized to store files smaller than 512 bytes in a single sector. With the change in disk technology, most file systems have avoided the slowdown of 512-byte writes by making 4,096 bytes the smallest allocation size. Thus, a file smaller than 512 bytes is now placed in a 4,096-byte block. The result of this change is that it takes up to eight times as much



File systems must be aware of the disk technology on which they are running to ensure they can reliably deliver the semantics they have promised.



space to store a file system with predominantly small files. Since the average file size has been growing over the years, for a typical file system the switch to making 4,096 bytes the minimum allocation size has resulted in a 10%–15% increase in required storage.

Some file systems have adapted to the change in sector size by placing several small files into a single 4,096-byte sector. To avoid the need to do a read-modify-write operation to update a small file, the file system collects a set of small files that have changed recently and writes them out together in a new 4,096-byte sector. When most of the small files within a sector have been rewritten elsewhere, the sector is reclaimed by taking the few remaining small files within it and including them with other newly written small files in a new sector. The now-empty sector can then be used for a future allocation.

The conclusion is that file systems must be aware of the disk technology on which they are running to ensure they can reliably deliver the semantics they have promised. Users need to be aware of the constraints that different disk technology places on file systems and select a technology that will not result in poor performance for the type of file-system workload they will be using. Perhaps going forward they should just eschew those lying disks and switch to using flash-memory technology—unless, of course, the flash storage starts using the same cost-cutting tricks. ■

Related articles on queue.acm.org

Building Systems to be Shared, Securely

Poul-Henning Kamp, Robert Watson

<http://queue.acm.org/detail.cfm?id=1017001>

The Five-Minute Rule 20 Years Later: and How Flash Memory Changes the Rules

Goetz Graefe

<http://queue.acm.org/detail.cfm?id=1413264>

GFS: Evolution on Fast-forward

Marshall Kirk McKusick, Sean Quinlan

<http://queue.acm.org/detail.cfm?id=1594206>

Marshall Kirk McKusick writes about and teaches classes on Unix- and BSD-related subjects, and provides expert-witness testimony on software-patent, trade-secret, and copyright issues. He has been a developer on the FreeBSD Project since its founding in 1994. While at the University of California at Berkeley, he implemented the 4.2BSD fast file system and was the research computer scientist at the Berkeley Computer Systems Research Group overseeing the development and release of 4.3BSD and 4.4BSD.

© 2012 ACM 0001-0782/12/11 \$15.00

DOI:10.1145/2366316.2366333

Decryption keys allow users to learn a specific function of the encrypted data and nothing else.

BY DAN BONEH, AMIT SAHAI, AND BRENT WATERS

Functional Encryption: A New Vision for Public-Key Cryptography

ENCRYPTION IS A method for users to securely share data over an insecure network or storage server. Before the advent of public-key cryptography, a widely held view was that for two users to communicate data confidentially they would have to first establish a mutually held secret key k . While acceptable, perhaps, for some small or tight-knit organizations, such a solution is clearly infeasible for larger networks (such as today's Internet). More than 30 years ago, Diffie and Hellman^{11,12} introduced the concept of public-key cryptography, where two parties securely communicate with each other without having a prior mutual secret, radically challenging the conventional wisdom of the time.

Today, public-key encryption is invaluable, ubiquitous in securing Web communication (such as HTTPS and SSH), voice traffic, and storage systems. However, within the technical community, there is an ingrained view that:

- ▶ Access to the encrypted data is “all or nothing”; one either decrypts the entire plaintext or learns nothing about the plaintext (other than a bound on its length); and

- ▶ Encryption is a method to encode data so a single secret key can decrypt that data.

However, for many applications, this notion of public-key encryption is insufficient; for example, the encryptor may want to encrypt data so anyone satisfying a certain policy can then decrypt it. Consider encrypting a message to a company so the only users who can decrypt it are employees in, say, the accounting or sales departments whose office is in the company's main building. Realizing this application using existing public-key encryption raises several questions:

- ▶ How do we discover the public keys of all individuals who satisfy this policy?;

- ▶ What if someone joins the system or receives certain credentials well after the data is encrypted and stored?;

- ▶ What if we want to give someone a partial view of the plaintext depending on their credentials?; and

- ▶ Should a given user even be al-

» key insights

- Unlike traditional encryption, where decryption is all or nothing, in a functional encryption system decryption keys may reveal only partial information about the plaintext; for example, decrypting an encrypted image with a cropping key will reveal a cropped version of the image and nothing else.
- Many advances in public-key encryption over the past decade can be viewed as special cases of functional encryption.
- Current functional encryption systems are quite expressive, yet much work remains in expanding the set of functionalities supported by existing constructions.

```

<mainCommandSet />
<id="browser">
  <browser>
    <content flex="1">
      <otero-pane-stack fullScreen=

```

lowed to learn the identities of all individuals who have certain credentials?

Functional encryption. It is time to adopt a new broad vision of encryption that takes into account such concerns. To this end, we advocate the concept of “functional encryption” where a decryption key enables a user to learn a specific function of the encrypted data and nothing else. Briefly, in a functional-encryption system, a trusted authority holds a master secret key known only to the authority. When the authority is given the description of some function f as input, it uses its master secret key to generate a derived secret key $sk[f]$ associated with f . Now anyone holding $sk[f]$ can compute $f(x)$ from an encryption of any x . In symbols, if $E(pk; x)$ is an encryption of x , then decryption accomplishes

given $E(pk; x)$ and $sk[f]$,
decryption outputs $f(x)$.

Note it is $f(x)$ that is made available to the secret key holder, even though x was the value that was encrypted. A functional-encryption system can support a variety of functions this way. Intuitively, the security of the functional-encryption system should guarantee the secret-key holder can learn nothing else about x beyond $f(x)$. We thus envision functional encryption as analogous to secure computation^{18,33} but with the critical difference that functional encryption is completely non-interactive once a recipient obtains the recipient’s own secret key.

Consider what is possible if functional encryption would be realized for a broad set of functions:

Spam filtering on encrypted mail. A user wishes to leverage a partially trust-

ed proxy to filter out all encrypted email messages identified as spam according to the user’s criteria. The user wants to achieve the seemingly conflicting goals of hiding the message’s contents from the proxy while allowing the proxy to determine if the message is spam according to some arbitrary criteria. The user can achieve these goals by setting up a functional-encryption system, then giving the proxy a key $sk[f]$ where f is the user-specified program that outputs 1 if the plaintext is spam and 0 otherwise. The proxy can use $sk[f]$ to test if an encrypted message is spam without learning anything more about the plaintext (see the figure here).

One can naturally consider generalizations of this idea; for instance, the proxy might selectively send important email messages (as deemed by the function f) to the user’s mobile device. Taking things further we can imagine the destination of a packet is encrypted, and the secret key $sk[f]$ allows a router to learn the next hop and nothing more.

Expressive access control. In large organizations a user will often think of sharing data according to some access policy. In addition to our corporate example, this might also occur in other domains (such as health care, insurance companies, government institutions, and universities). Bridging the gap between how a user thinks of sharing data and discovering the public keys of all other users who match or will match such sharing can be difficult and is subject to the problems outlined earlier; for example, a system might try to encrypt data separately to the public key of every user matching a certain policy. However, as also noted, this user-specific encryption requires


identification of each user, as well as the overhead of encrypting to each one individually. Moreover, this user-specific encryption does not cover users who do not meet the criteria today but will in the future.

Using functional encryption a user can directly express how the user (or organization) wishes to share the data in the encryption process. In particular, the user can encrypt $x = (P, m)$ where m is the data the user wishes to share, and P is the access policy that describes how the user wants to share it. The user’s secret-key function $sk[f]$ will then check whether the user’s credentials or attributes match the policy and reveal only m in this case. Corresponding to the example of an accounting or sales department with an office in the company’s main building, P could embed the policy (“ACCOUNTING” OR “SALES”) AND “MAIN BUILDING.” A recipient’s function f would embed the attributes of the particular user and check if they satisfy the formula and if so return m .

Mining large datasets. Data mining is used in medical research, social networks, network security, and financial fraud detection. Administrators often want to give users the ability to mine datasets for certain types of queries but not let them learn anything else. Consider a medical researcher who wants to test if there is a link between a genotype and a type of cancer in a particular ethnic group. If the administrator has data consisting of patient gene sequences and medical history, the administrator would like to give the researcher the ability to test for this linkage, without revealing the details of all patients’ medical conditions.

Note that in practice, an administrator typically does not know the queries that will be of interest until well after the data is created and stored. Functional encryption provides an elegant solution. When data is created it is simply encrypted in a functional-encryption system. Later, a user requests to be allowed to learn a certain query or function f of the data. If data access is authorized, the user is given $sk[f]$ and can apply this key to (attempt to) decrypt existing or future encrypted data. Thus, in a functional-encryption system supporting a class of functions $\mathbf{F}$ a user could be given the ability to

The email recipient, who has a master secret key sk , gives a spam-filtering service a key $sk[f]$ for the functionality f ; this f satisfies $f(x) = 1$ whenever message x is marked as spam by a specific spam predicate, otherwise $f(x) = 0$. A sender encrypts an email message x to the recipient, but the spam filter blocks the message if it is spam. The spam filter learns nothing else about the contents of the message.



compute any function from this class on the dataset.

These three examples of functionality motivate the research agenda we put forward here—to create functional-encryption systems supporting the richest possible families of functions and understand what limitations are inherent for functional-encryption systems.

Functional Encryption

Recall that public-key encryption systems (such as RSA and El-Gamal) consist of three algorithms:

Setup. Outputs a secret key denoted sk and a public key pk ; anyone can encrypt message using pk , but only the secret key holder is able to decrypt;

Encryption E . Takes a public key pk and a message as input and outputs a ciphertext; and

Decryption D . Takes a secret key sk and a ciphertext as input and outputs a message.

A functional-encryption system includes the same three algorithms but also includes a fourth algorithm called *KeyGen*. Here, the secret key output by the *Setup* algorithm is called the master key, denoted by mk . The *KeyGen* algorithm takes as input mk and the description of some function f . It outputs a key that is specific to the function f and denoted $sk[f]$. More precisely, if c is the result of encrypting data x with public key pk , then

$$D(sk[f]; c) \text{ outputs } f(x)$$

We emphasize that $sk[f]$ does not fully decrypt c , outputting only a function f of the full decryption. To fully decrypt c the authorized user can use a secret key $sk[g]$, where g is the identity function, namely $g(x) = x$ for all x .

Informally, security of a functional-encryption system means an attacker with a set of secret keys $sk[f_1], \dots, sk[f_e]$ can learn nothing about the decryption of some ciphertext c other than what is revealed by the keys at the attacker's disposal.

To illustrate the power of functional encryption, the following section covers how it naturally captures many advanced encryption concepts in cryptography. First, it should be clear that traditional public-key encryption is a very special case of functional encryption, where the only supported



A filtering server can use the user's functional secret key to test if an encrypted email message is spam without learning anything else about the plaintext.



function is the identity function; the decryptor learns either the complete decryption or nothing at all.

Identity-based encryption. A more advanced public-key concept called “identity-based encryption,” or IBE, is an encryption system where any string can serve as a public key; a user's email address, a date, an IP address, a location, or even the numbers 1, 2, and 3 are all potential public keys. IBE public keys are often called “identities” and denoted by id . To obtain the secret key for a particular identity the user communicates with an authority holding a master key. The authority verifies the user is authorized to receive the requested secret key, and, if so, it generates the secret key using its master key. IBE was proposed by Shamir²⁹ in 1984, and the first implementations of IBE were proposed by Boneh and Franklin⁶ and Cocks¹⁰ in 2001; notable constructions include Agrawal et al.,¹ Boneh and Boyen,⁴ Gentry,¹⁵ Gentry et al.,¹⁶ Waters,³⁰ and Waters.³¹

Using the terminology of functional encryption the IBE problem can be recast as an equality testing functionality. Let pk and mk be the output of the functional encryption setup algorithm. To encrypt a message m to identity id the encryptor calls the encryption algorithm as

$$E(pk; (id, m))$$

and obtains a ciphertext c . Note the data being encrypted is the pair (id, m) .^a A recipient with identity id^* obtains a secret key for id^* by asking the authority for a secret key $sk[f_{id^*}]$ where the function f_{id^*} is defined^b as

$$f_{id^*}((id, m)) := \begin{cases} m & \text{if } id = id^*, \\ \perp & \text{otherwise} \end{cases}$$

The authority generates $sk[f_{id^*}]$ using its functional master key mk . Using this secret key the user can decrypt messages intended for identity id^* but learns nothing about messages encrypted for other identities.

Recall that IBE systems reduce reliance on certificate directories needed

^a Using our earlier notation we would have $x = (id, m)$, though we omit x for readability.

^b We use the $\perp$ symbol as a special symbol to denote failure to decrypt.

for traditional public-key encryption; to encrypt to identity id , the encryptor needs only the global public key pk and the recipient's identity id . General functional encryption systems have the same property: they require no on-line certificate directory. An encryptor needs only the global public key pk and the payload x to be encrypted and no other information about the intended recipient(s).

Attribute-based encryption. Another encryption concept called “attribute-based encryption,” or ABE, lets the encryptor specify more abstractly who is authorized to decrypt a specific ciphertext. ABE was proposed by Sahai and Waters²⁸ and later refined by Goyal et al.¹⁹ into two different formulations: key-policy ABE and ciphertext-policy ABE.

In the ciphertext-policy system the encryptor specifies a policy φ on recipient attributes that determines who can decrypt the ciphertext; for example, the encryptor can encrypt messages to anyone who is a

“U.S. citizen” and “female”) or (“over 30”)

which is a Boolean formula φ on three variables. To encrypt a message m with decryption policy φ the encryptor calls

$$E(pk, (\varphi, m))$$

and obtains a ciphertext c .

Now, consider a recipient who wants to decrypt the ciphertext. The recipient has a number of attributes, say,

“U.S. citizen,” “Rhodes Scholar,” “female,” “under 30”

Let n be the total number of attributes, and we represent the set of user attributes as a Boolean vector of length n ; the vector is 1 at positions that correspond to attributes that are true and 0 everywhere else. With this setup each user has an attribute vector $\mathbf{u}$ in $\{0, 1\}^n$.

A recipient with attribute vector $\mathbf{u}$ obtains a secret key for his attribute vector by asking the authority for a secret key $sk[f_{\mathbf{u}}]$ where the function $f_{\mathbf{u}}$ is defined as

$$f_{\mathbf{u}}((\varphi, m)) := \begin{cases} m & \text{if } \varphi(\mathbf{u}) = 1, \\ \perp & \text{otherwise} \end{cases}$$

A tantalizing question is whether techniques from lattices, which have been so useful in the context of fully homomorphic encryption, can help achieve greater functionality for functional encryption.

The authority generates $sk[f_{\mathbf{u}}]$ using its functional master key mk . Using this secret key the user can decrypt ciphertexts where the user's attributes satisfy the decryption policy but learns nothing about the decryption of other ciphertexts.

A related concept called “key-policy attribute-based encryption” places the access policy φ in the key and the vector $\mathbf{u} \in \{0, 1\}^n$ in the ciphertext. The secret key $sk[f_{\varphi}]$ decrypts all encryptions $E(pk, (u, m))$ for which $\varphi(\mathbf{u}) = 1$.

Security

Here we turn to constructing functional-encryption systems but first explain what it means for a functional system to be secure. The full definition is a bit technical, and we give only high-level intuition; for more, see Boneh et al.⁷

Roughly speaking, a functional-encryption system is secure if an attacker with a set of secret keys $sk[f_1], \dots, sk[f_t]$ can learn nothing about the decryption of some ciphertext c other than what is revealed by the keys at the attacker's disposal. If c is the encryption of some data x , then the attacker can use the attacker's secret keys to learn $f_1(x), \dots, f_t(x)$. However, the attacker must be unable to learn anything else about x ; for example, if the attacker has secret keys that reveal the first three bits of x , then clearly the attacker can learn these bits, given an encryption of x but would be unable to learn anything about the remaining bits of x .

To give a bit more detail about security requirements, let A be a polynomial-time adversary that takes as input three things: the public key pk , a set of secret keys $sk[f_1], \dots, sk[f_t]$ for functions $f_1, \dots, f_t$ of its choice, and a ciphertext $c = E(pk, x)$. This A might output some information about the decryption of c (such as the least significant bit of x). We say the system is secure if for every such A there is another polynomial-time algorithm B , called a simulator, that, given pk and $f_1(x), \dots, f_t(x)$ but not given c is able to output the same information about x that A output. Since B never got to see c it must have deduced the information about x strictly from $f_1(x), \dots, f_t(x)$. Since A and B output the same information about x , the existence of B means the only information A can learn about x from the cipher-

text c is information it can learn from $f_1(x), \dots, f_t(x)$ but cannot learn anything else about x . Hence, A can learn from c whatever is revealed by the secret keys at its disposal but nothing else.^c

Challenge: Preventing collusion attacks. Attacks on functional encryption using multiple functional secret keys are called collusion attacks, and preventing them is the main obstacle to constructing secure functional systems. To illustrate the problem consider again the functionality described earlier and suppose the encryptor wishes to encrypt a message m to this policy

“U.S. citizen” and “over 30”

A simple implementation is to associate a public key pk_1 with the attribute “U.S. citizen” and a public key pk_2 with the attribute “over 30” and double-encrypt the message m as

$$c = E(pk_1, E(pk_2, m))$$

where here $E(\cdot, \cdot)$ is a regular public-key-encryption algorithm. To decrypt c the recipient—call her Alice—must possess both secret keys corresponding to pk_1 and pk_2 , implementing the conjunction policy specified by the encryptor.

Now, suppose another user, Bob, has attributes “U.S. citizen” and “male” where the attribute “male” is associated with a public key pk_3 . He would be given the secret keys corresponding to pk_1 and pk_3 , letting him decrypt messages encrypted for some policies (such as (“U.S. citizen” and “male”)). In addition, suppose Alice has the attribute “over 30” and is then given only the secret key corresponding to pk_2 . Thus, she cannot decrypt the message associated with the original policy on her own.

The problem is Alice and Bob can collude to combine their secret keys and create new secret keys neither one should have; for example, Alice and Bob working together can decrypt ciphertexts intended for policy “over 30” and “male,” even though neither can decrypt the ciphertext by themselves. In this example, collusion enabled Alice and Bob to decrypt a ciphertext to

which neither should have access.

Secure constructions. Secure constructions for complex functionalities must prevent collusion attacks. Collusion attacks are prevented by binding together all secret keys for a set of attributes, so mixing the keys given to distinct users does not help. As a visual metaphor, one can imagine that all the keys given to Alice are colored blue, while all the keys given to Bob are colored red. Decryption succeeds only when the decryptor uses a set of keys of the same color. The colors ensure Alice and Bob cannot combine their keys to decrypt ciphertexts they should not be able to decrypt.

In practical terms, the colors are implemented through randomization values. All the keys given to Alice are blinded by the same random value, while all the keys given to Bob are blinded by a different random value. Decryption with keys blinded by the same randomizer produces the correct decrypted message. Decryption with keys blinded by different randomizers results in a random value unrelated to the correct decryption.

State of the Art

The state of the art in functional encryption can be understood by considering what information about the plaintext x is exposed by the ciphertext to all participants. We refer to this information as the result of the “empty functionality” denoted $f_\epsilon(\cdot)$; for example, it is inherent in any encryption scheme that the empty functionality exposes some information about x (such as a bound on the size of the plaintext). When the exact plaintext length is leaked by the ciphertext we write $f_\epsilon(x) = |x|$ to indicate that anyone can learn the plaintext length from the ciphertext.

Public index: ABE. In general, we can consider the problem of functional encryption where the data to be encrypted is decomposed into two parts $x = (\text{ind}, m)$, where ind denotes a public index the encryptor does not mind revealing to all participants in the system. That is, we define the empty functionality as $f_\epsilon(\text{ind}, m) = (\text{ind}, |m|)$.

Now consider the specific case of ABE, where the access policy φ is now considered a public index. In it, where access policy does not require protection, we have fairly broad and efficient

constructions of secure ABE schemes; secure ABE schemes exist that support any access policy φ that can be expressed as a Boolean formula over the attributes (as in the earlier examples).^{3,19,22,23,25,26,32} Going beyond policies expressible as Boolean formulas remains a vexing open problem for researchers, with the ultimate goal of supporting policies expressible as arbitrary Boolean circuits or Turing Machines.

Non-public index. A more challenging setting arises where we insist the empty functionality reveals as little as possible, namely $f_\epsilon(x) = |x|$. Here, our current understanding of functional encryption is extremely limited. The state of the art is limited to the inner-product functionality over prime fields.^{2,21,22,25} Because this functionality is somewhat technical (and before we describe it more formally), we briefly discuss some applications: First, consider the question of searching on encrypted data, where the data is encrypted based on a public key and stored on a public server.⁵ The security challenge in this setting is to hide the specific nature of the search query from the public server while still allowing the public server to send back only data entries that match the search query. The inner-product functionality we describe in the following paragraphs allows a user to perform such a search based on disjunction queries and more generally searches defined by CNF and DNF formulae or by checking whether a univariate search polynomial evaluates to zero on a particular input value.

The functionality we consider is defined over a prime field F_p where p is a large prime chosen randomly during the setup of the functional-encryption scheme. Messages and keys will correspond to vectors of a fixed arbitrary dimension n over F_p . Let us denote the message by the vector $\mathbf{v}$ and the vector underlying a secret key by $\mathbf{u}$. We then have

$$f_u(\mathbf{v}) := \begin{cases} 1 & \text{if } \sum_{i=1, \dots, n} \mathbf{u}_i \cdot \mathbf{v}_i = 0, \\ \perp & \text{otherwise} \end{cases}$$

To see how this functionality can be applied, consider again the example of disjunction queries: Suppose a ciphertext is meant to encrypt a single keyword we hash down to a value a in our finite field F_p . Then to encrypt this val-

^c Note our security model does not rely on any assumption about trusted hardware or online servers needed during decryption.

ue a , the system actually encrypts the vector $\mathbf{v} = (1, a, a^2, \dots, a^{n-1})$. Now, suppose we have to create a key corresponding to a disjunction query “ a_1 OR a_2 OR a_3 ”. We do this by first considering the polynomial $p(x) = (x - a_1)(x - a_2)(x - a_3)$, writing it out in standard form as $p(x) = c_0 + c_1x + c_2x^2 + c_3x^3$, where the c_i are the appropriate coefficients. We then issue a key for the vector $\mathbf{u} = (c_0, c_1, c_2, c_3, 0, \dots, 0)$. Glancing at the functionality, we see our key will indeed match the ciphertext for value a if and only if $p(a) = 0$; that is, if the value a is a root of our polynomial $p(x)$, which was designed to have roots only at the three values a_1, a_2, a_3 in our desired disjunction. Other special cases of inner products, including conjunctions and range testing functionalities, were considered in Boneh and Waters.⁸

Unfortunately, the exact cryptographic mechanisms by which the results work in Katz et al.,²¹ Lewko et al.,²² and Okamoto and Takashima²⁵ are too technically involved to describe here; we encourage all to look into these sources for further technical detail.

Current limitations. Current functional-encryption schemes, especially in non-public index settings, are limited. From a technical standpoint, current techniques for building functional-encryption schemes are all based on elliptic-curve groups equipped with efficiently computable bilinear pairings that map into the multiplicative structure of a finite field. At a very high level of design abstraction a pairing operation allows for a single multiplication between the exponents of two “source” group elements. However, the result of a pairing operation is a “target” group for which the operation cannot be repeated.

The reason we can handle inner products of two vectors is because this operation requires only one parallel call to the multiplication operation, which is all that bilinear maps provide. A tantalizing question is whether techniques from lattices, which have been so useful in the context of fully homomorphic encryption,¹⁴ can help achieve greater functionality for functional encryption.

Efficiency. The efficiency of functional-encryption systems varies significantly with specific cryptographic constructions. However, we can offer

an approximate sense of the efficiency of the ABE where the ciphertext is associated with any access policy φ that can be expressed as a Boolean formula over attributes. In current systems, the size of the ciphertext scales with the size of the Boolean formula φ ; for example, in Waters,³² a ciphertext consisted of two group elements for every leaf node of φ , and encryption took three exponentiations for every leaf node. Decryption requires two of the aforementioned pairing operations for each attribute used in the formula. While difficult to predict how future functional-encryption systems might evolve, developers could expect that the number of public-key operations required will scale with the complexity of the functionality.

Functional Encryption vs. Fully Homomorphic Encryption

Fully homomorphic encryption (FHE) is arguably the most impressive development in cryptography over the past few years, enabling one to compute on ciphertexts in the following sense: Given a public key pk , encryptions of messages $x_1, \dots, x_t$ under pk , and the description of a function f as input, any user can construct an encryption of the message $f(x_1, \dots, x_t)$; see Gentry¹³ for a detailed discussion. A more restricted version of FHE, called univariate FHE, allows any user to construct an encryption of $f(x)$ from an encryption of x for all univariate functions f .

While both FHE and functional encryption support some form of computation on ciphertexts, it is not known how to construct functional encryption from FHE; FHE does not even seem to imply basic functionalities (such as identity-based encryption). The reason for this limitation is that the output of an FHE computation on encrypted data is an encrypted result; in contrast, the output of a functional-encryption computation is available in the clear.

To further illustrate the difference between FHE and functional encryption recall the spam-filtering example discussed at the beginning of the article. In it, the spam filter was given a secret key $sk = sk[f]$, where f is a function that outputs 1 if an email is spam and 0 otherwise. The key sk lets the spam filter run the spam predicate f on encrypted email messages and

block encrypted spam. With FHE, the spam filter can likewise run the spam predicate f on encrypted email messages, but the filter learns only the encrypted output of the predicate; it does not and cannot learn whether an encrypted email message is spam. In particular, with FHE, the filter can tag an encrypted email message with only an encrypted tag indicating “spam” or “not spam” but cannot block spam email messages for the end user. This example illustrates the potential power of functional encryption over FHE. However, constructing a fully functional encryption scheme is still an open problem, whereas FHE constructions exist.

Generalizations

Here, we cover a few generalizations, variants, and extensions of functional encryption that are motivated in practice:

Delegating keys. Users might sometimes want to delegate a limited set of their capabilities to another user or device; for example, a medical researcher with a secret key able to decrypt raw medical records might want to distribute to a grad student a key that can output only certain statistics (such as averages over the data). As another example, suppose users are planning to travel with their mobile devices but are concerned the devices might be lost or stolen; they might then want to copy a key to the devices that decrypts only the data that was encrypted during the travel time or restrict the key to capabilities related only to the purpose of the trip.

A simple approach is for users with a key $sk[f]$ to query the authority for a more restrictive key $sk[f']$ anytime they wish to delegate a key for a more restrictive function f' . However, involving the authority in every delegation is cumbersome, exposes an online authority to more risk, and will not work if the authority is unreachable. Therefore, we would like the delegation operation to be autonomous. Roughly, a user with $sk[f]$ can create $sk[f']$ if f' is more limited than the function f ; whatever we can learn from f' we can learn from f .

The concept of delegation arose in identity-based encryption in Gentry and Silverberg¹⁷ and in Horwitz and

Lynn^{17,20} and can be realized in attribute-based encryption.¹⁹

Functionality over multiple authorities. In a standard functional-encryption system, one authority is responsible for issuing private keys, though some systems might require more flexibility. Returning to the example of (ciphertext-policy) attribute-based encryption, in a standard system, one authority is responsible for both determining what attributes/credentials to issue to each user and creating the keys.

While a single-authority solution is likely workable for smaller organizations, in many applications a user might want to create policies spanning many trust domains; for instance, suppose we wish to encrypt a document for all military personnel who are also members of the ACM, asking who should manage the system? Using a central authority creates several problems; for one, no single party is always able to speak authoritatively for multiple trust domains or organizations. Indeed, a user might wish to create a policy that spans organizations that are not even aware of one another. Another core limitation is that a central authority creates a central performance bottleneck and consolidates trust in one entity. Are any two different organizations able to agree who to trust in this role?

Recent work in decentralized attribute-based encryption^{9,24} has sought to overcome these limitations so users are able to encrypt according to an ABE policy issued as a formula over attributes issued from different authorities. An interesting direction is to see what other functionalities beyond ABE might arise from the use of multiple authorities in functional-encryption systems.

Functional encryption with public-key infrastructure. Finally, we consider how ideas from functional encryption can be applied to other scenarios; specifically, consider a scenario involving the following conditions:

Per-user infrastructure. There exists a per-user public-key infrastructure where every user u obtains a secret key $sk_u[f]$ for some function f_u appropriately chosen by the user and also establishes a public key pk_u unique to the user; this public key should also not leak information about the function f_u . Such a



While existing functional-encryption systems are remarkably expressive, the central challenge is to construct a system that supports creation of keys for any function in both public and non-public index settings.



key is established through interaction between an authority and the user;

Targeting a specific key. Encryptions are always targeted at a specific user's public key pk_u . However, the encryptor does not know the function f_u corresponding to the user hidden by the public key pk_u . At the same time, if a user u obtains an encryption of x under the user's public key pk_u , then decryption allows the user to learn $f_u(x)$ but nothing else. Users should also not be able to obtain additional capabilities by combining secret keys corresponding to different public keys; and

Misbehaving central authority. A misbehaving central authority must not be able to decrypt encryptions intended for honest users in the system.

We stress this scenario is quite different from the functional-encryption scenario we have considered here. One of the key properties of functional encryption is it does not require public-key directories, thus enabling a variety of applications (such as secure storage in the cloud and secure searching on encrypted data). At the same time, the support comes at the cost of users needing to trust a key-generation authority (or set of such authorities) capable of breaking the security of ciphertexts.

This scenario was considered in recent work²⁷ where it was shown that in this setting, called “worry-free encryption,” the system can support functions (in the non-public index setting) specified by any arbitrary polynomial-size circuit—significantly beyond what is possible with standard functional encryption today. However, it must be stressed that this setting does not cover motivating applications of functional encryption (such as secure storage in the cloud and searching on encrypted data); see Sahai and Seyalioglu²⁷ for more detail on this setting.

Future of Functional Encryption

What will functional encryption look like in 10 years? While existing functional-encryption systems are remarkably expressive, the central challenge is to construct a system that supports creation of keys for any function in both public and non-public index settings. If we could create such systems we could imagine embedding anything from arbitrarily complex spam filters

to image-recognition algorithms into encryption systems. Imagine an encryption system that lets a user view only an image if a facial-recognition algorithm matches a picture of the user to a face in the encrypted image. Moreover, the output of the decryption could show the area immediately surrounding the identified user and blur out the rest of the image.

Current progress on building functional encryption systems is dominated by the tool of groups with bilinear maps mentioned earlier. However, as also mentioned earlier there are reasons to suspect there might be fundamental barriers to realizing more advanced functional encryption systems from this tool.

Cryptography researchers need to search further out, though a reason for optimism is the recent dramatic leap in what we can achieve in homomorphic encryption systems. Hopefully, such a leap will be achieved in the not-too-distant future (perhaps using related techniques) in the realm of functional encryption.

Finally, more applied research is needed to build functional encryption into real-world systems, as well as to specify formats for attribute spaces and languages for expressing access policies. Due to the expressive power of these systems we hope to see real-world deployments of functional encryption over the next decade. The end result is much greater flexibility in specifying who can and cannot access protected data.

Acknowledgments

Dan Boneh is supported by the National Science Foundation, the Defense Advanced Projects Agency PROgramming Computation on EncryptEd Data (PROCEED) program, the Air Force Office of Scientific Research under the Multidisciplinary University Research Initiative award for “Collaborative policies and assured information sharing” (Project PRESIDIO), a Google Faculty Research Award, Samsung, and the Packard Foundation.

Amit Sahai is supported by a Defense Advanced Research Projects Agency/Office of Naval Research PROCEED award, NSF grants 1228984, 1136174, 1118096, 1065276, 0916574, and 0830803, a Xerox Foundation

Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based on work supported by DARPA through the U.S. Office of Naval Research under Contract N00014-11-1-0389.

Brent Waters is supported by NSF CNS-0915361 and CNS-0952692, Air Force Office of Scientific Research Grant No: FA9550-08-1-0352, DARPA PROCEED, DARPA N11AP20006, Google Faculty Research award, an Alfred P. Sloan Fellowship, Microsoft Faculty Fellowship, and the Packard Foundation.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Department of Defense or the U.S. government. 

References

- Agrawal, S., Boneh, D., and Boyen, X. Efficient lattice (H)IBE in the standard model. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2010, 553–572.
- Agrawal, S., Freeman, D.M., and Vaikuntanathan, V. Functional encryption for inner product predicates from learning with errors. In *Proceedings of ASIACRYPT, Lecture Notes in Computer Science*, Springer, 2011, 21–40.
- Bethencourt, J., Sahai, A., and Waters, B. Ciphertext-policy attribute-based encryption. In *Proceedings of the IEEE Symposium on Security and Privacy*, IEEE Computer Society, 2007, 321–334.
- Boneh, D. and Boyen, X. Efficient selective-id secure identity-based encryption without random oracles. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2004, 223–238.
- Boneh, D., Crescenzo, G.D., Ostrovsky, R., and Persiano, G. Public-key encryption with keyword search. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2004, 506–522.
- Boneh, D. and Franklin, M.K. Identity-based encryption from the weil pairing. In *Proceedings of Crypto, Lecture Notes in Computer Science*, Springer, 2001, 213–229.
- Boneh, D., Sahai, A., and Waters, B. Functional encryption: Definitions and challenges. In *Proceedings of TCC, Lecture Notes in Computer Science*, Springer, 2011, 253–273.
- Boneh, D. and Waters, B. Conjunctive, subset, and range queries on encrypted data. In *Proceedings of TCC, Lecture Notes in Computer Science*, Springer, 2007, 535–554.
- Chase, M. Multi-authority attribute-based encryption. In *Proceedings of TCC, Lecture Notes in Computer Science*, Springer, 2007, 515–534.
- Cocks, C. An identity-based encryption scheme based on quadratic residues. In *Proceedings of the Institute of Mathematics and Its Applications, Lecture Notes in Computer Science*, Springer, 2001, 360–363.
- Diffie, W. and Hellman, M.E. Multiuser cryptographic techniques. In *Proceedings of AFIPS National Computer Conference*, AFIPS Press, 1976, 109–112.
- Diffie, W. and Hellman, M.E. New directions in cryptography. *IEEE Transactions on Information Theory* 22 (1976), 644–654.
- Gentry, C. Computing arbitrary functions of encrypted data. *Commun. ACM* 53, 3 (Mar. 2010), 97–105.
- Gentry, C. Fully homomorphic encryption using ideal lattices. In *Proceedings of STOC 2009*, ACM Press, New York, 2009, 169–178.
- Gentry, C. Practical identity-based encryption without random oracles. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2006, 445–464.
- Gentry, C., Peikert, C., and Vaikuntanathan, V. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of STOC*, ACM Press, New York, 2008, 197–206.
- Gentry, C. and Silverberg, A. Hierarchical id-based cryptography. In *Proceedings of ASIACRYPT 2002, Lecture Notes in Computer Science*, Springer, 2002, 548–566.
- Goldreich, O., Micali, S., and Wigderson, A. How to play any mental game or a completeness theorem for protocols with honest majority. In *Proceedings of STOC*, ACM Press, New York, 1987, 218–229.
- Goyal, V., Pandey, O., Sahai, A., and Waters, B. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the ACM Conference on Computer and Communications Security*, ACM Press, New York, 2006, 89–98.
- Horwitz, J. and Lynn, B. Toward hierarchical identity-based encryption. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2002, 466–481.
- Katz, J., Sahai, A., and Waters, B. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2008, 146–162.
- Lewko, A.B., Okamoto, T., Sahai, A., Takashima, K., and Waters, B. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2010, 62–91.
- Lewko, A.B., Sahai, A., and Waters, B. Revocation systems with very small private keys. In *Proceedings of the IEEE Symposium on Security*, IEEE Computer Society, 2010, 273–285.
- Lewko, A.B. and Waters, B. Decentralizing attribute-based encryption. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2011, 568–588.
- Okamoto, T. and Takashima, K. Fully secure functional encryption with general relations from the decisional linear assumption. In *Proceedings of CRYPTO, Lecture Notes in Computer Science*, Springer, 2010, 191–208.
- Ostrovsky, R., Sahai, A., and Waters, B. Attribute-based encryption with non-monotonic access structures. In *Proceedings of the ACM Conference on Computer and Communications Security*, ACM Press, New York, 2007, 195–203.
- Sahai, A. and Seyalioglu, H. Worry-free encryption: Functional encryption with public keys. In *Proceedings of the ACM Conference on Computer and Communications Security*, ACM Press, New York, 2010, 463–472.
- Sahai, A. and Waters, B. Fuzzy identity-based encryption. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2005, 457–473.
- Shamir, A. Identity-based cryptosystems and signature schemes. In *Proceedings of CRYPTO, Lecture Notes in Computer Science*, Springer, 1984, 47–53.
- Waters, B. Efficient identity-based encryption without random oracles. In *Proceedings of EUROCRYPT, Lecture Notes in Computer Science*, Springer, 2005, 114–127.
- Waters, B. Dual-system encryption: Realizing fully secure ibe and hibe under simple assumptions. In *Proceedings of CRYPTO, Lecture Notes in Computer Science*, Springer, 2009, 619–636.
- Waters, B. Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization. In *Proceedings of the Public Key Cryptography Conference, Lecture Notes in Computer Science*, Springer, 2011, 53–70.
- Yao, A.C.-C. Protocols for secure computations (extended abstract). In *Proceedings of FOCS*, IEEE Computer Society, 1982, 160–164.

Dan Boneh (dabo@cs.stanford.edu) is a professor of computer science and electrical engineering at Stanford University.

Amit Sahai (sahai@cs.ucla.edu) is a professor of computer science at the University of California, Los Angeles.

Brent Waters (bwaters@cs.utexas.edu) is an assistant professor of computer science at the University of Texas at Austin.

© 2012 ACM 0001-0782/12/11 \$15.00

Open-source chemistry software and molecular databases broaden the research horizons of drug discovery.

BY JOERG KURT WEGNER, AARON STERLING, RAJARSHI GUHA, ANDREAS BENDER, JEAN-LOUP FAULON, JANNA HASTINGS, NOEL O'BOYLE, JOHN OVERINGTON, HERMAN VAN VLIJMEN, AND EGON WILLIGHAGEN

Cheminformatics

NOVEL TECHNOLOGIES in the life sciences produce information at an accelerating rate, with public data stores (such as the one managed by the European Bioinformatics Institute <http://www.ebi.ac.uk>) containing on the order of 10PB of biological information. For nearly 40 years, the same was not so

for chemical information, but in 2004 a large public small-molecule structure repository (PubChem <http://pubchem.ncbi.nlm.nih.gov>) was made freely available by the National Library of Medicine (part of the U.S. National Institutes of Health) and soon followed by other databases. Likewise, while many of the foundational algorithms of cheminformatics have been described since the 1950s, open-source software implementing many of them have become accessible only since the mid-1990s.¹⁰

Why is chemical information important? Why should chemists and com-

puter scientists care about its public availability? And how does chemical information relate to the field of computer science? Though cheminformatics is used in fields from agrochemical research to the design of novel materials, here we use drug discovery as our context due to its relevance to human well-being. The art and science of drug discovery focuses on small molecules, with a core component of methods involving techniques and algorithms to handle, analyze, and interpret chemical-structure information. Unfortunately, drug-discovery research is slow, expensive, and prone to failure. Public availability of chemical information and related tools is important, as the more information available to each researcher, the better the chances of avoiding the many causes of attrition. Computer-science research is highly relevant to managing the volume and complexity of chemical information. A single database (such as PubChem) contains more than 34 million chemical-structure records, along with an even larger number of annotations (such as synonyms, or different names for the same molecule, with Pub-

» key insights

- **Molecules with similar physical structure tend to have similar chemical properties.**
- **Open-source chemistry programs and open-access molecular databases allow interdisciplinary research opportunities that did not previously exist.**
- **Cheminformatics combines biological, chemical, pharmaceutical, and drug patient information to address large-scale data mining, curation, and visualization challenges.**

Chem storing 50.1 million synonyms for 19.6 million compounds), known targets (drugs with which a molecule is known to interact), mode of action (how the molecule interacts with its targets), and status within the regulatory-approval process.

While bioinformatics often deals with sequences, the domain of cheminformatics is chemical structures. In the former, information is frequently represented as 1D strings that are relatively easy to handle computationally. In the latter, chemical structures are complex graphs that may include rings and branches, along with multiple valid representations of the same molecule to be considered for particular algorithms (such as tautomers, where hydrogen atoms are positioned at different places in the molecule). Hence,

chemical structures are more difficult to standardize, and algorithms dealing with them must take this lack of standardization into account; in Figure 1, the top-left graph contains much implicit information, since unlabeled nodes represent carbons with valencies satisfied by hydrogens, and the top-right graph is fully labeled; that is, all atoms are explicitly labeled with their symbols. The former is what is usually exchanged informally between chemists, but actually mining chemical structures computationally requires something like the latter; for example, graph-mining algorithms can be applied to labeled graph representations to discover patterns in large-scale datasets,³⁷ information that can also be determined through simulations (such as of molecular dynamics).

Though graphs implicitly encode the 3D structure of a molecule (when combined with knowledge about favored bond angles and distances), many different low-energy 3D structures, or conformers, may be consistent with the same graph structure. Moreover, the 3D structure may have a “secondary” geometrical arrangement of features (such as a right-handed helix) that cannot be encoded in the graph. Chemists thus have 3D representations (see Figure 1) that make explicit the 3D arrangement of atoms and bonds; the representation at the bottom right goes further, showing a molecular surface and some property that varies over the surface (such as lipophilicity, or the ability to dissolve in non-polar solvents). Even though some structure representations contain more explicit information than others, all are equally valid. When searching for substructures, a 2D representation suffices, but when exploring protein binding, a 3D structure is preferred; see the online Appendix for a notable subtlety in which the number and types of atoms and their connectivity may not uniquely define a structure.

A principle of cheminformatics is that similar molecules exhibit similar properties;¹⁶ the choice of representation is key in determining how such similarities are evaluated and thus the effectiveness of subsequent analyses. But there is a further challenge in balancing computational costs with the utility of a representation; for example, a full 3D description of a molecule accounting for all possible conformers would allow more accurate prediction of many properties, though such conformer predictions would also depend on good-quality force fields and algorithms. But the size of the representation and the time required for its evaluation would be prohibitive. So how can accurate similarity predictions be obtained from a subset of conformers? Or how can comparable accuracy be obtained through a 2D representation? Moreover, if it can be, what type of labels are required? Many such questions are answered today by trial and error through definition of an objective function (usually root mean square error or percentage correct) and iterative adaptation of descriptors and modeling approaches to optimize the objec-

Figure 1. A hierarchy of chemical structure representations, increasing in information content from left to right and top to bottom.

A chemist's traditional 2D depiction (top left); a labeled graph (top right, with “A” corresponding to aromatic atoms, “O” to oxygens, “On” to oxygens or nitrogens, and the implicit “C” to unlabeled nodes); a 3D conformation (bottom left); and a surface representation (bottom right). All show differing, valid aspects of the same molecule, each suited for different purposes, with the choice of representation guided by intended use.





Figure 2. Standard chemical graph illustration and SMILES InChI and InChIKey codes for lipoic acid; note SMILES is much more human-readable than the others.



tive function. No unifying theory is yet able to explain or even suggest optimal approaches in all cases.

Cheminformatics aims to support better chemical decision making by storing and integrating data in maintainable ways, providing open standards and tools allowing application and data use across heterogeneous platforms and mining the many chemical property spaces in a time- and space-efficient way. However, to flourish, cheminformatics needs closer collaboration between chemists and computer scientists, with the former able to pose their problems in a way that is relevant for practical applications and the latter to devise ways of capturing, storing, and analyzing chemical data to achieve optimal balance of space, performance, and complexity; for a detailed introduction to cheminformatics, written for computer scientists, see Brown.⁶ Cheminformatics is broadly divided into three areas: capturing data using lab notebooks or potentially using such formats as the Chemical Markup Language for publications; storing data (such as designing database schemas and devising ontologies); and mining data (such as for predicting biological activity of compounds).

Bridging Cheminformatics and Computer Science

Here, we highlight current topics in cheminformatics, presenting computational and algorithmic problems and how computer science can contribute to their solution; for more, see the Appendix. We use as example cheminformatics methods “risk minimization” in drug discovery, minimizing the chances of a small molecule failing due to poor physical, chemical, or biological properties during research and development as a drug candidate; for example, a molecule must be soluble and show a certain degree of bioavailability to be considered a drug candidate. In cases where these properties are poor, a cheminformatics approach can suggest replacement of certain functional groups (connected sets of atoms affecting the characteristics of the chemical reactions of the molecule) to maintain potency but improve the solubility and bioavailability; see the Appendix for example properties a drug can-



A principle of cheminformatics is that similar molecules exhibit similar properties; the choice of representation is key in determining how such similarities are evaluated and thus the effectiveness of subsequent analyses.



didate must satisfy to be considered therapeutically useful and the role of cheminformatics at each stage of drug development.

Representing and searching structures. Most cheminformatics applications rely on large databases of chemical structures and their properties and relationships to, say, biological targets. Organizing and maintaining them, as well as searching and clustering similar structures together, are essential for many scientific applications. However, each of these areas poses computer-science challenges; for example, chemical and bioactivity databases (such as ChEMBLdb and PubChem) are freely available and contain possibly millions (ChEMBLdb) and tens of millions (PubChem) of data points. Integration of this disparate data is essential for researchers to gain the fullest possible perspective on what is presently known, tracking ongoing advances in science as they become available. Integration of data across chemical databases is a challenge due to sheer data volume and to the difficulties in normalization of chemical and bioactivity data. Molecular graphs must be captured in a machine-readable fashion. It is also necessary that life scientists be able to search chemical data (such as multi-labeled graphs where graph labels can change from database to database). Just as labels can differ between databases chemical graphs can be encoded in multiple ways, depending how the nodes are ordered, resulting in multiple representations of the same molecule. The challenge for database search increases when considering protonation states (how the molecule changes when one proton is added to it) and tautomer states (how the molecule can change when a proton migrates from one part of the molecule to another). The need for a unique representation that is invariant with respect to atom ordering arises due to the expense of graph isomorphism (checking whether two structures represent the same molecule). Morgan²³ described the first such “unique representation algorithm,” or canonicalization algorithm, allowing chemists to generate unique string representations of chemical graphs and compare structures through string compari-

sons. The Simplified Molecular-Input Line-Entry System, or SMILES, format defined by Weininger in 1988³⁹ is an example of a representation that can be canonicalized. Since the original canonicalization algorithm was proprietary, multiple implementations of the format have become available, each employing a different canonicalization algorithm, usually based on the Morgan algorithm; see Warr³⁸ for an extensive discussion on chemical-structure representations.

Each database using its own algorithm for molecular encoding hinders the automated exchange of data between different databases. As more and more data has been made available online, the use of unique structure-based identifiers became a pressing need, resulting in development of the International Union of Pure and Applied Chemistry's International Chemical Identifier (InChI), a non-proprietary, structured textual identifier for chemical entities.³² InChI identifiers are not intended to be read and understood by humans but are useful for computational matching of chemical entities. For quick database lookups, the InChIKey is a hashed key for the InChI with an invariant length of 14 characters; Figure 2 outlines the SMILES, InChI, and InChIKey for lipoic acid.

InChI is widely used for matching identical chemical structures but is still limited; for example, it cannot differentiate between certain types of stereoisomers (informally, molecules that are "3D mirror images" of one another) and is sensitive to tautomeric chemical forms; see the Appendix for a discussion of two stereoisomers for which the generated InChI is the same.

InChI and other identity-mapping algorithms allow for exact searching; two other practically relevant algorithms for scientific discovery based on chemical databases are substructure searching and similarity searching required to generalize from the search molecule to other, related molecules. In substructure searching, the database is searched for a specified wholly contained part of the search structure; in similarity searching, structures are retrieved that are similar (in some structure or property space) to the provided search structure. Chemical search packages are often implement-



For chemists it is important to appreciate that chemical space (occupied by all possible chemical structures) is, in principle, infinite.



ed and optimized for a given database technology; for example, the OrChem package is an open-source chemical search package for the Oracle database application.²⁷

Graph substructure matching is a variant of graph isomorphism, widely viewed as computationally intractable;⁸ executing a graph isomorphism search across a full chemical database of thousands or millions of structures is simply not feasible.⁴⁰ Speedups can be obtained through structural fingerprint filters; fingerprints encode characteristic features of a given chemical structure, usually in a fixed-length bit-map. Fingerprints fall broadly into two categories: structure keys and hashed keys. In structure keys, each bit position corresponds to a distinct substructure (such as a functional group); examples are MACCS and PubChem keys. In hashed keys substructural patterns are represented as strings and then hashed to a random bit position; as a result, a given position can encode multiple substructures. The advantage of fingerprints is they can cover an arbitrarily large collection of substructures (such as "paths of length N ," or circular environments); examples are "daylight fingerprints" ("folded" to optimize information density and screening speed) and "extended connectivity fingerprints," or ECFPs, that use local topological information. Given a binary fingerprint, chemists can first pre-screen a database to ignore molecules that cannot possibly match the query by requiring all bits in a query fingerprint to also be present in the target fingerprint. Since the target fingerprints are pre-computed, performing this check on modern hardware is quick. As a result, chemists apply the actual isomorphism test on only those molecules that pass the screen. Fingerprints can also be used to quickly search databases for similar molecules, using a similarity metric (such as the Tanimoto coefficient) to compare the query and target fingerprints. Additional heuristics³⁴ further speed up similarity searches.

Molecules in their biological context. In the quest to discover novel therapeutic agents, studying molecules in a biological context is essential, as metabolites, cofactors, hormones, and signaling molecules. Developers of


chemical databases must thus represent and organize chemical data across biological databases (such as those with pathways, protein information, and biological samples). Integration and processing data from such disparate domains underlie systems-level biological research. This yields additional challenges standard chemical structural representation cannot address: first, the representation of classes of chemical entities, since, in many cases, a class of compounds behaves in a certain context in a biological system rather than as a single molecule with a fully specified structure and, likewise, chemical concepts (such as groups, or parts of molecules); second, the need to represent non-structural groupings of compounds of interest, since compounds may bind to the same or similar targets or have similar biological functions (such as acting as an anti-neoplastic, preventing development of tumors). With multiple biological databases having to refer to these different chemical entities in an organized, standardized fashion to enable cross-database integration for a whole-system perspective, ontologies and other semantic technologies are used to support annotation, classification, and semantic cross-domain querying of chemical entities within biological systems. The most widely used ontology for biologically relevant chemical entities—ChEBI⁹—contained (as of February 2012) approximately 27,100 entities in an interlinked semantic graph and was used for annotation in dozens of biological databases. Ontologies are based on logical languages of varying levels of expressivity, accompanied by sophisticated reasoning algorithms. They can be used to construct various sorts of semantic similarity (such as in function and in application) that complement traditional structural-similarity methods.¹¹ An open challenge is how to integrate graph-based chemical-structure information with wider logic-based ontology information to allow combined automated reasoning and querying over the different domains encoded in a growing number of interlinked ontologies. Enabling automated inference in such heterogeneous linked systems brings chemists closer to a systems-level understanding of small-molecule activity.²⁶

Activity mining and prediction.

The basis of predictive modeling in cheminformatics is that the biological activity of a molecule is a function of the molecule's chemical structure. Together with the similar-property principle¹⁶ mentioned earlier, the goal of any modeling approach is capturing and characterizing correlations between structural features and observed biological activity. Such approaches must also describe the likelihood of error when using the models for decision making. A variety of approaches can be employed to assess the error in (or conversely, the reliability or confidence of) a prediction, ranging from statistical approaches to more empirical approaches (such as defining an applica-

account for both ligand and receptor structures. Proteochemometric methods are an extension of statistical QSAR methods to simultaneously model the receptor and ligand in a “system chemistry” sense, as first reported by Lapinsh et al.²⁰

The first step in predicting biological activities is to generate molecular descriptors, or features, that are numerical representations of structural features; for example, labeled graphs and their associated characterizations are easily accessible to computer scientists yet miss significant physicochemical features (such as surface distributions and 3D pharmacophores). Chemists can also have difficulty objectively quantifying many chemical



bility domain, the region of input that can be predicted reliably, usually delineated by similarity to the training set).

In cases of receptor-mediated activity, the activity of a small molecule is due to its interaction with a receptor. Traditionally, quantitative structure–activity relationship (QSAR) model^{12,14} approaches do not consider receptor features, focusing instead on only small-molecule features, thus losing valuable information on ligand–receptor interactions. As a result, techniques (such as docking, which predicts how molecules fit together), structure-based pharmacophore modeling (a 3D approach to capturing protein–ligand interactions), and proteochemometric methods have been designed to

aspects of a molecule, such that the resultant descriptors are suitable for predictive modeling. Choosing a chemical descriptor should by no means be viewed as a “solved” problem; for more, see two comprehensive textbooks, one by Faulon and Bender,¹⁰ the other by Todeschini and Consonni.³⁶

Molecular graphs can be transformed into numerical vector representations ranging from counts of elements to eigenvalues of the Laplacian matrix. Alternatively, molecular graphs can be compared directly through kernel methods, where a kernel on graphs G and G' provides a measure of how similar G is to G' or a kernel on a single graph compares measured similarities between the nodes of the two graphs.

In cases where graph kernels have been defined, a chemist would, rather than compute vector representations, operate directly on the graph representations. Each method involves advantages and disadvantages; for example, the vector approach requires a computer scientist to identify a subset of relevant (to the property being modeled) descriptors; the feature selection problem is well covered in the data mining literature. A kernel approach does not require feature selection, but a computer scientist would face the problem of having to evaluate a dataset in a pairwise fashion, identifying an appropriate kernel. How to perform this evaluation is an important challenge, as the kernel must be se-

lected to satisfy Mercer's condition (a well-known mathematical property in machine learning that makes it easier to make predictions about a set of observations), and satisfying Mercer's condition is not always possible with traditional cheminformatics-based kernels (such as those based on multiple common substructures). These challenges can make kernel-based methods computationally prohibitive on larger datasets.

Having settled on a numerical representation and a possible class of model types, a computer scientist would have to address the goal of the model. Is the chemist looking for pure predictive power, with no interest in explanatory features or more interest


in decent predictive power, as well as some explanation as to why a molecule is predicted to be, say, toxic or nontoxic? The former scenario is common in virtual-screening settings, where the chemist might require a high degree of accuracy and fast predictions but not really care why one molecule is active and another inactive. Such models can be black boxes (such as neural networks) and algorithmic models⁵ (such as random forests). The latter is more common in exploratory and optimization settings, where a computer scientist might hope the output of the model will guide chemists in chemical modifications to improve the property being modeled. In this case chemists must understand the effect a certain

structural feature has on observed potency, expecting the model to provide insight. Such models are generally distributional (such as linear regression, partial least squares, and naïve Bayes), though some algorithmic approaches can also provide insight. Having chosen a modeling approach, the chemist and computer scientist must address model reliability closely tied to the concept of model applicability, or the reliability of the prediction of the model for a new object. This issue has become more important with increasing use of predictive models in regulatory settings. Misprediction (due to the model not being applicable to a certain class of inputs) can have significant financial

repercussions. Various methods have been developed that attempt to characterize the model's domain, determining not only whether models are applicable but whether additional biological experiments are required to reduce the prediction error on certain compound classes.

A key challenge faced in predictive modeling is the fact that small molecules are not static and do not exist in isolation. Though, traditionally, predictive models have focused on a single structure for a small molecule and ignore the receptor, small molecules can exist in multiple tautomeric forms and conformations. Enhancing the accuracy of predictions ideally requires taking into account the 3D geometries of the molecule and the receptor as much as possible. Though it is possible to generate reasonable low-energy conformations *ab initio*, the "biologically relevant" conformation might differ significantly (in terms of energetics) from the lowest-energy conformation of the molecule considered in isolation, necessitating conformational search. Multi-conformer modeling was addressed by the 4D-QSAR methodology described by Hopfinger et al.¹ More recent techniques (such as multiple-instance learning) are also applied to the multi-conformer problem.

With the advent of high-throughput screening technologies, large libraries of compounds can now be screened against multiple targets in an efficient manner. Such panel assays provide a broad, systems-level view of small-molecule activities. Models developed on such data afford the opportunity to identify targets and characterize off-target effects. However, most approaches to the problem of screening multiple targets tend to develop multiple individual models,⁷ leading to multiple, independent predictions for a given input molecule. Alternatively, a system chemist (chemogenomics scientist) might imagine an approach that takes into account the covariance structure of multiple observed activities and structural descriptors within a single model. Such an approach could lead to more robust predictions for panel assays, or which battery of tests would be most useful. Finally, this approach also better reflects clinically relevant compound profiles¹⁹ and



the “personalized medicine” concept (such as drug-drug interaction profiles).

Expanding chemical space. Enumerating molecules is a combinatorial problem that has fascinated chemists, computer scientists, and mathematicians alike for more than a century. Indeed, many fundamental principles of graph theory and combinatorics were developed by British mathematician Arthur Cayley, Hungarian mathematician George Pólya, and others in the context of counting isomers of paraffin. In the 1960s, Joshua Lederberg of Rockefeller University, Carl Djerassi of Stanford University, and others developed algorithms to enumerate structures based on spectral data, leading to DENDRAL, widely viewed as the first expert system.²¹

From a risk-reduction point of view, efficient methods for enumerating structures allow chemists to explore new regions of chemical space (such as to bypass patents), generate structures that exhibit desired properties, and identify molecules matching experimental data. For chemists it is important to appreciate that chemical space (occupied by all possible chemical structures) is, in principle, infinite. Even considering molecules for just 30 heavy atoms, the size of this space is on the order of 10^{60} heavy atoms.⁴ Any enumeration method would face a combinatorial explosion if implemented naïvely.

A key application of structure enumeration is the elucidation of structures based on spectral data,¹⁷ especially for identifying metabolites, or small molecules that are the by-products of metabolic processes and thus provide insight into an organism’s biological state (such as diseased and fasting). A chemist gathers spectral data (such as nuclear magnetic resonance, mass, and liquid chromatography–mass spectrometry), and an algorithm would ideally provide a list of structures that give rise to the observed spectra. Some commercial products (such as MOLGEN <http://molgen.de>) are able to perform this task quickly.

Another application of structure enumeration concerns Markush structures and searches;² for example, a Markush claim used in a patent would cite multiple “functionally equivalent”



Given that any experiment could become the core of a patent, ELNs allow organizations to efficiently define and implement audit trails.



chemical entities, where the user specifies a pattern (such as “...an aromatic ring with two alkyl groups attached to it”). Such a pattern is very general: an alkyl group can be only methyl, ethyl, or any chain of n carbons, with three possible positions for them on the ring. However, even this simple definition involves $3n^2$ possible structures. More complex Markushes can involve billions of possible structures. Explicit enumeration is not feasible. Analysis of Markush structures thus faces a number of challenges, primarily the ability to search (based on structural or property similarity) through the implicit enumerated space. A number of commercial vendors, including Digital Chemistry and ChemAxon, offer toolkits for these problems.

Structure enumeration plays a fundamental role in molecular design, or the design of compounds that optimize some physical, chemical, or biological property or activity.³⁰ A key challenge is how to combine enumeration algorithms with efficient property prediction and is closely related to methods in predictive modeling of chemical properties discussed earlier. This approach is also termed inverse QSAR, where the chemist must devise a set of features that describe a molecule with a specified property and subsequently reconstruct the molecule from those features. Moreover, the reconstruction of molecules from feature sets can be viewed as a cryptography problem. Such reconstruction is relevant to the pharmaceutical industry, as it may be necessary to share data on molecules without also sharing explicit structures. Feature sets that allow easy reconstruction of molecular structures are thus undesirable. A number of chemistry researchers have investigated the problem of identifying such one-way molecular descriptors, as well as methods for reconstructing molecules from descriptors.²² New methods that generate such one-way molecular features but still correlate with physicochemical properties would be valuable.

While molecular-graph enumeration is a challenge, an alternative to enumerating molecular structures based on spectral data is to sample these structures.¹³ Sampling procedures based on metropolis or genetic

algorithms help elucidate compounds from NMR data. However, even in the face of the complexity of molecular enumeration, computational chemists have developed and successfully used enumeration tools to generate large chemical libraries (such as GDB-13 <http://www.cbligand.org/gdb13/>) of almost one billion chemical structures³ considering only molecules up to 13 heavy atoms and using only carbon, nitrogen, oxygen, sulphur, and chlorine. However, current enumeration software products do not generally produce stereoisomers or tautomers that require specific enumeration procedures and are thus still the subject of cheminformatics research.

Another relevant application of enumeration methods is how to generate chemical-reaction networks. The problem here consists of enumerating all possible compounds that can be produced by applying reaction rules to a set of initial molecules. By reversing the reaction rules chemists are also able to find sets of starting compounds necessary for producing a given target, a process called retrosynthesis. Designing new drugs and chemicals, understanding the kinetics of combustion and petroleum refining, studying the dynamics of metabolic networks, and applying metabolic engineering and synthetic biology to produce heterologous compounds (compounds from different species) in microorganisms all involve enumeration of reaction networks. As reviewed in Faulon and Bender¹⁰ several network-enumeration techniques have been developed but generally suffer from a combinatorial explosion of product compounds. One way to limit the number of compounds being generated is to simulate the dynamics of the network while it is being constructed and remove compounds of low concentration.

Following this idea, methods have been developed based on the Gillespie Stochastic Simulation Algorithm (http://en.wikipedia.org/wiki/Gillespie_algorithm) to compute on-the-fly species concentrations. Chemical-reaction-network enumeration and sampling is an active field of research, particularly in the context of metabolism, for studying biodegradation or proposing metabolic-engineering strategies to biosynthesize compounds of

commercial interest. However, a difficulty with metabolic network design is that, in addition to network generation based on reactions, a system chemist or biologist must also verify there are possible enzymatic events to enable reaction catalysis; enzymes must be present to reduce the energy required for some reactions to take place. This additional task requires computer scientists include both chemical structures and protein sequences and develop tools at the interface between cheminformatics and bioinformatics.

Knowledge management. Scientific-knowledge management is increasingly relevant not only to reduce risk in current research but to enable new collaboration and innovation opportunities with internal partners and a growing number of external (public) partners. In cheminformatics and chemistry, scientists switched years ago from paper lab notebooks to online collaboration platforms called electronic lab notebooks, or ELNs. So, what finally drove chemists and pharmaceutical companies to adopt “enterprise 2.0” social online-collaboration culture? Before 2000, many chemists still used paper lab notebooks and multiple compound-registration and search tools. The overall architecture was too inflexible to adapt to fast-changing data standards, and scientists spent too much time on administrative work; moreover, chemical data quality was inconsistent, and alignment with other working groups was inefficient (such as for running analytical experiments). Legal disputes would require manual searches through many paper lab notebooks, and data synchronization was painful, hindering large-scale collaboration.

While ELNs were available as long ago as the 1990s, chemists only began adopting them in large numbers in 2004 and 2005, and the market continues to grow.³⁵ A 2006 report (http://www.atriumresearch.com/library/Taylor_Electronic_laboratory_notebooks.pdf) said, “The initial drive for the development of ELNs came from the field of chemistry, perhaps driven by the early adoption by chemists of computer technologies [...] The pharmaceutical industry shifted from a position of ‘if we get an ELN’ to ‘when we get an ELN’.”³⁵ The growth of ELNs is driven

by many factors, including ease of use (rapid search across organizationwide data sources and easy sharing of domain knowledge) and regulatory compliance (time-stamped experimental data, experiment and data approvals, and quality-control information). Given that any experiment could become the core of a patent, ELNs allow organizations to efficiently define and implement audit trails.

“Electronic lab notebooks have replaced the traditional paper lab book across the pharmaceutical industry”¹⁸ but are not yet common in academia due to their cost. Worth mentioning here is that chemistry is traditionally a field accustomed to bookshelves laden with reactions and compounds. Chemists have switched to online ELNs not only to increase efficiency but because they allow easy collaboration with trusted colleagues. Third-party paper catalogues and academic publications proved inefficient and, critically, did not account for compounds from trusted colleagues. The change of management to ELNs, as with enterprise 2.0, has required delivering on the promise of greater interconnectivity and collaboration opportunities, made legally possible in the U.S. due to the U.S. Food and Drug Administration’s regulation Title 21 CFR Part 11, permitting electronic signatures in chemical documents. Leading players include CambridgeSoft’s E-Notebook and Accelrys’s Symyx Notebook, but at least 35 different companies were producing ELNs as of February 2011²⁸ for chemistry, biology, quality control, and other research domains. Another knowledge-management tool designed specifically for chemists is the Web-based service Reaxys (<https://www.reaxys.com/info/>) launched in 2009 to provide “a single, fully integrated chemical workflow solution.”

However, using external (public) databases with chemical and bioactivity data is a challenge due to differences in identifiers, synchronization, curation, and error-correcting mechanisms, as well as the need to provide efficient substructure and similarity search within complex data types. A collaboration with external parties (such as contract research) poses further challenges, including compound duplication and efficient data synchronization. Some

external partners may lack sufficient IT resources themselves and thus rely on external services. Cloud services could help provide a service infrastructure for all parties involved, as well as the required private-public interface. Furthermore, many data sources are still not being indexed properly; for example, chemistry patents are often cryptic, and chemical image and text mining is a challenge though is being addressed in academic and industrial research.^{15,29} The closed Chemical Abstract Service (CAS <http://cas.org>) is one highly trusted source, and public chemical and bioactivity databases must improve their quality and interconnectivity to compete. SciFinder (<http://www.cas.org/products/scifinder/>) is another chemical-abstract service, with a version, SciFinder Scholar, marketed to universities. Open-source efforts like ChEMBL and PubChem-BioAssay are on the right track, though, unlike some commercial tools, do not abstract reactions. Still, improving data quality and standards between public and closed sources is critical for ensuring growth, use, and collaboration by private and public parties alike.

Novel Algorithms

Essential for researchers interested in new algorithms is a software library that handles chemical structures and related data; a variety of such chemical toolkits are available, both proprietary (possibly involving free academic licenses) and open source.


Due to the importance of cheminformatics to the pharmaceutical industry, numerous commercial vendors provide libraries, applications, and database cartridges, including Accelrys, BioSolveIT, Chem-Axon, Chemical Computing Group, Daylight, Open Eye, Schrodinger, Tripos, and Xemistry. Since about 1995, open-source cheminformatics software has emerged, providing opportunities for rapid development and implementation of novel algorithms that build on the existing open-source ecosystem. Here, we do not debate open source, focusing instead on open-source tools to help explore cheminformatics problems.

Open-source software in cheminformatics has lagged open-source software in bioinformatics. However, recent years have seen a rise in open-

source cheminformatics software. One notable organization is the Blue Obelisk group²⁵ (<http://blueobelisk.org>) of chemists seeking to create interoperable open-source chemistry software and open-source chemical toolkits, including the Chemistry Development Kit (CDK),³³ Open Babel,²⁴ RDKit, and Indigo, written in Java or C++, though those in C++ have, via <http://swig.org>, bindings to a variety of languages; for example, such toolkits are used to read/write chemical files in various formats, manipulate chemical structures, measure similarity of molecules, search for substructures, and generate 2D depictions. The underlying algorithms include graph algorithms (such as maximal

robust support for 3D structure generation and optimization and supports the interconversion, or conversion of one file format to another, of a large number of chemical file formats; the Appendix includes a comparison of features offered by the open-source toolkits. Given there is ample scope for software engineering (such as algorithm implementation, code quality and analysis, and building systems), CDK and Open Babel are both open to new contributions, driven by public mailing lists and tracking systems.

Certain challenges faced by nearly all cheminformatics toolkits stem from the graph representation of chemicals, that, while succinct and amenable to computation, only approximates real-



common substructure and canonical labelling of graphs), geometrical methods (such as Kabsch alignment), and vector manipulation (such as converting coordinates in various systems and 3D structure generation). Besides being subject to programmatic use by cheminformaticians, many applications rely on these toolkits to handle chemical data; for example, the molecular viewer Avogadro (<http://avogadro.openmolecules.net>) uses Open Babel, and the molecular workbench Bioclipse³¹ uses CDK.

Though the toolkits share many features they also have certain distinguishing features; for example, CDK implements a large collection of molecular descriptors. Open Babel has

ity, with edge cases abounding. Areas of interest include enumeration of colored graphs accounting for symmetry and symmetry detection itself, which not only derives from the chemical graph but typically also from the 3D geometry of the chemical. Chemists must realize, though the sophistication of chemical representations has increased over the years, many chiral molecules and non-standard bonding cases (such as organometallic compounds) cannot be handled by current representation systems. A further challenge is the chemical representation of a molecule would need to capture non-static graphs to account for delocalization and the tautomeric phenomena molecules undergo in real biological contexts.

Of equal importance to development of new algorithms is freely available data on which to train and test new methods. As noted earlier, large structure and bioactivity data collections are available, enabling much more robust validation of cheminformatics methodologies, as well as large-scale benchmarking of algorithms. Benchmarking is especially relevant for the data-mining techniques in cheminformatics. Open data focuses primarily on structure and activity data types, with a notable lack of “open” textual data (such as journal articles). While PubMed abstracts (<http://www.ncbi.nlm.nih.gov/pubmed/>) are a proxy for journal articles, text-mining methods in cheminformatics are hindered by not being able to mine the full text of many scientific publications. Patent information is publicly accessible, supporting these efforts. Open data does not explicitly address data quality or problems integrating data sources. In many cases, manual curation is the only option for maintaining high-quality databases; fortunately, open-data facilitates curation by allowing independent access to anyone with a computer.

Conclusion

Cheminformatics, the computer science of chemical discovery, is an industrial and academic discipline with roots dating to the 19th century and a flowering in the 1960s along with modern computing technologies. While many key cheminformatics techniques have been available in the literature since the 1960s, most cheminformatics tools and implementations have been proprietary; likewise, most data sources have been proprietary (usually with restrictive licensing policies) until recently. Companies look to protect their intellectual property for many reasons primarily involving profit, competitive intelligence, and intellectual property relevant to the pharmaceutical industry, as well as to other chemical-related industries. Unlike much of bioinformatics, these issues are where the data and tools have been freely available since the field's inception. The disparity between cheminformatics and bioinformatics can be attributed to the fact that the outcomes from cheminformatics



Keeping in mind cheminformatics is fundamentally a practical field, serving experimental chemistry, the key challenges require an understanding by scientists of the underlying chemical systems.



methods and software have a more direct effect on profits, in terms of identifying lead-like compounds and improving the properties of drug candidates, while bioinformatics software is found in upstream areas (such as target identification) and is perhaps less directly related to possible profits as a candidate small molecule. Moreover, acquiring chemical data (such as structure and activity) is more difficult and when done on a large scale can involve much time and effort, whereas acquiring bioinformatics data (such as sequences) is much easier.

While both fields have theoretical components, it is possible that the free availability of a large amount of bioinformatics data drove development of related publicly available tools. In contrast, the proprietary nature of chemical data would imply that tools needed to process and analyze it are primarily of interest to the owners. As a result, there has been little incentive to make cheminformatics software publicly available, and since users are primarily industrial, commercial software is the norm.

However, the cheminformatics landscape is shifting, with free open-source cheminformatics toolkits, applications, and open databases with tens of millions of compounds and experimental bioactivity data, while modern experimental techniques (such as high-throughput sequencing) have made generation of large amounts of structure-activity data much more accessible.

Prospects for academic and industrial collaboration between chemists and computer scientists are bright, encouraging more computer scientists to participate in cheminformatics research; for many specific open research questions, see the Appendix.

Significant domain knowledge is required to address problems in cheminformatics. Indeed, many issues faced by chemists are not as clean an abstraction as “algorithms on a string,” (substring matching is not very useful on multilabeled graphs) the way many bioinformatics algorithms can be abstracted into a computer-science framework. Hence, while chemistry researchers can contribute to cheminformatics simply by considering structures as graphs, they are inevitably lim-

ited to somewhat abstract problems. Keeping in mind cheminformatics is fundamentally a practical field, serving experimental chemistry, the key challenges require an understanding by scientists of the underlying chemical systems.

Nevertheless, due to increasing availability of tools and data, the barrier to entry for non-chemists and non-cheminformaticians is significantly lower than a decade ago. Many questions that would benefit from computer science can now be addressed; for example, for theorists, graph-theoretic questions of 3D enumeration; for database designers, more effective search algorithms and ontologies; for practical programmers, expansion of the many open-source cheminformatics projects. If more chemists would think algorithmically and more computer scientists chemically, the pharmaceutical industry and associated industries would be much better positioned to deal with not only “simple” chemicals (only one identifier and isomer possible per molecule) but also their complex relationships, transformations, and combinatorial challenges, bringing cheminformatics closer to the goal of supporting primary translational research in the wider context of chemical biology and system chemistry.

Acknowledgments

We would like to thank Danny Verbinen of Janssen Pharmaceutical Companies for sharing his insight into ELNs and John Van Drie of Van Drie Consulting for his insight into cheminformatics. This article was written while Aaron Sterling was visiting the Department of Electrical Engineering and Computer Science at Northwestern University where he was supported in part by National Science Foundation grant CCF-1049899. Andreas Bender thanks Unilever for funding. Janna Hastings thanks the European Union project EU-OPEN-SCREEN for funding.

References

- Albuquerque, M., Hopfinger, A., Barreiro, E., and de Alencastro, R. Four-dimensional quantitative structure-activity relationship analysis of a series of interphenylene 7-oxabicycloheptane oxazole thromboxane a₂ receptor antagonists. *Journal of Chemical Information and Computer Sciences* 38, 5 (Oct. 1998), 925–938.
- Barnard, J. A comparison of different approaches to Markush structure handling. *Journal of Chemical Information and Computer Sciences* 31, 1 (Feb. 1991), 64–68.
- Blum, L.C. and Raymond, J.-L. 970 million druglike small molecules for virtual screening in the chemical universe database GDB-13. *Journal of the American Chemical Society* 131, 25 (July 2009), 8732–8733.
- Bohacek, R., McMartin, C., and Guida, W. The art and practice of structure-based drug design: A molecular-modeling perspective. *Medicinal Research Reviews* 16, 1 (Jan. 1996), 3–50.
- Breiman, L. Statistical modeling: The two cultures. *Statistical Science* 16, 3 (Aug. 2001), 199–215.
- Brown, N. Chemoinformatics—An introduction for computer scientists. *ACM Computing Surveys* 41, 2 (Feb. 2009), 1–38.
- Chen, B. and Wild, D.J. PubChem BioAssays as a data source for predictive models. *Journal of Molecular Graphics and Modelling* 28, 5 (Jan. 2010), 420–426.
- Cordella, L., Foggia, P., Sansone, C., and Vento, M. An improved algorithm for matching large graphs. In *Proceedings of the Third IAPR-TC15 Workshop on Graph-Based Representations in Pattern Recognition* (Ischia, Italy, 2001), 149–159.
- De Matos, P., Alcántara, R., Dekker, A., Ennis, M. et al. Chemical entities of biological interest: An update. *Nucleic Acids Research* 38 (Jan. 2010), D249–D254.
- Faulon, J.-L. and Bender, A. *Handbook of Cheminformatics Algorithms*. CRC Press, Boca Raton, FL, 2010.
- Ferreira, J. and Couto, F. Semantic similarity for automatic classification of chemical compounds. *PLoS Computational Biology* 6, 9 (Sept. 2010), e1000937.
- Free, S. and Wilson, J. A mathematical contribution to structure activity studies. *Journal of Medicinal Chemistry* 7 (July 1964), 395–399.
- Goldberg, L. and Jerrum, M. Randomly sampling molecules. *SIAM Journal on Computing* 29, 3 (Dec. 1999), 834–853.
- Hansch, C., Maloney, P., Fujita, T., and Muir, R. Correlation of biological activity of phenoxyacetic acids with Hammett substituent constants and partition coefficients. *Nature* 194 (Apr. 1962), 178–180.
- Jessop, D.M., Adams, S.E., Willighagen, E.L., Hawizy, L., and Murray-Rust, P. OSCAR4: A flexible architecture for chemical text mining. *Journal of Cheminformatics* 3, 1 (Oct. 2011), 41–53.
- Johnson, M. and Maggiora, G. *Concepts and Applications of Molecular Similarity*. John Wiley & Sons, New York, 1990.
- Kind, T. and Fiehn, O. Advances in structure elucidation of small molecules using mass spectrometry. *Bioanalytical Reviews* 2, 1–4 (Dec. 2010), 23–60.
- King, A. Waving goodbye to the paper lab book. *Chemistry World* 8, 11 (Nov. 2011), 46–49.
- Kuhn, M., Campillos, M., Letunic, I., Jensen, L.J.J., and Bork, P. A side-effect resource to capture phenotypic effects of drugs. *Molecular Systems Biology* 6 (Jan. 2010), 20.
- Lapinsch, M. Development of proteo-chemometrics: a novel technology for the analysis of drug-receptor interactions. *Biochimica et biophysica acta* 1525, 1–2 (Feb. 2001), 180–190.
- Lindsay, R., Buchanan, B., Feigenbaum, E., and Lederberg, J. DENDRAL: A case study of the first expert system for scientific hypothesis formation. *Artificial Intelligence* 61, 2 (June 1993), 209–261.
- Masek, B.B., Shen, L., Smith, K.M., and Pearlman, R.S. Sharing chemical information without sharing chemical structure. *Journal of Chemical Information Modelling* 48, 2 (Feb. 2008), 256–261.
- Morgan, H.L. The generation of a unique machine description for chemical structures: A technique developed at Chemical Abstracts Service. *Journal of Chemical Documentation* 5, 2 (May 1965), 107–113.
- O’Boyle, N.M., Banck, M., James, C.A., Morley, C., Vandermeersch, T., and Hutchison, G. R. Open Babel: An open chemical toolbox. *Journal of Cheminformatics* 3, 33 (Oct. 2011), 33–47.
- O’Boyle, N.M., Guha, R., Willighagen, E.L., Adams, S.E. et al. Open data, open source and open standards in chemistry: The Blue Obelisk five years on. *Journal of Cheminformatics* 3, 1 (Oct. 2011), 37–53.
- Oprea, T.I., Tropsha, A., Faulon, J.-L., and Rintoul, M.D. Systems chemical biology. *Nature Chemical Biology* 3, 8 (Aug. 2007), 447–450.
- Rijnbeek, M. and Steinbeck, C. An open source chemistry search engine for Oracle. *Journal of Cheminformatics* 1, 17 (Oct. 2009).
- Rubacha, M., Rattan, A.K., and Hosselet, S.C. A review of electronic laboratory notebooks available in the market today. *Journal of Laboratory Automation* 16, 1 (Feb. 2011), 90–98.
- Sayle, R. Foreign-language translation of chemical nomenclature by computer. *Journal of Chemical Information and Modelling* 49, 3 (Mar. 2009), 519–530.
- Schneider, G. and Fechner, U. Computer-based de novo design of drug-like molecules. *Nature Reviews Drug Discovery* 4, 8 (Aug. 2005), 649–663.
- Spieth, O., Alvarsson, J., Berg, A., Eklund, M. et al. Bioclipse 2: A scriptable integration platform for the life sciences. *BMC Bioinformatics* 10, 397 (Dec. 2009).
- Stein, S., Heller, S., and Tchekhovskoi, D. An open standard for chemical structure representation: The IUPAC chemical identifier. In *Proceedings of the 2003 International Chemical Information Conference* (Nimes, France, Oct.). Infonortics, Malmesbury, England, 2003, 131–143.
- Steinbeck, C., Han, Y., Kuhn, S., Hortlacher, O. et al. The Chemistry Development Kit (CDK): An open-source Java library for chemo- and bioinformatics. *Journal of Chemical Information and Computer Sciences* 43, 2 (Mar. 2003), 493–500.
- Swamidass, S.J. and Baldi, P. Bounds and algorithms for fast exact searches of chemical fingerprints in linear and sublinear time. *Journal of Chemical Information and Modelling* 47, 2 (Feb. 2007), 302–317.
- Taylor, K.T. The status of electronic laboratory notebooks for chemistry and biology. *Current Opinion in Drug Discovery and Development* 9, 3 (May 2006), 348–353.
- Todeschini, R. and Consonni, V. *Handbook of Molecular Descriptors, Volume 11 of Methods and Principles in Medicinal Chemistry*. Wiley-VCH, New York, 2000.
- Van der Horst, E., Okuno, Y., Bender, A., and Ijzerman, A. Substructure mining of GPCR ligands reveals activity-class specific functional groups in an unbiased manner. *Journal of Chemical Information and Modelling* 49, 2 (Feb. 2009), 348–360.
- Warr, W. Representation of chemical structures. *Interdisciplinary Reviews of Computational Molecular Science* 1, 4 (Mar. 2011), 557–579.
- Weininger, D. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of Chemical Information and Modelling* 28, 1 (Feb. 1988), 31–36.
- Weininger, D. Fingerprints: Screening and similarity; <http://www.daylight.com/dayhtml/doc/theory/theory.finger.html>

Joerg Kurt Wegner (jwegner@its.jnj.com) is a senior scientist in the Integrative System Biology Department of Janssen Pharmaceutical Companies of Johnson & Johnson, Beerse, Belgium.

Aaron Sterling (sterling@iastate.edu) is a Ph.D. candidate in the Department of Computer Science of Iowa State University, Ames, IA.

Rajarshi Guha (guhar@mail.nih.gov) is a research scientist of the National Institutes of Health, Bethesda, MD.

Andreas Bender (andreas.bender@cantab.net) is a lecturer in cheminformatics in the Unilever Center for Molecular Informatics of the University of Cambridge, Cambridge, U.K.

Jean-Loup Faulon (Jean-Loup.Faulon@univ-evry.fr) is a professor in the Biology Department of Evry University, UniverSud Paris and director of the Institute of Systems and Synthetic Biology of the French National Centre for Scientific Research, Paris, France.

Janna Hastings (hastings@ebi.ac.uk) is a member of the Cheminformatics and Metabolism Team of the European Bioinformatics Institute, Hinxton, U.K.

Noel O’Boyle (baoileach@gmail.com) is a Health Research Board postdoctoral fellow in the School of Pharmacy of the University College Cork, Ireland.

John Overington (jpo@ebi.ac.uk) is a group leader at the European Bioinformatics Institute, Hinxton, U.K., of the European Molecular Biology Laboratory.

Herman van Vlijmen (hvvlijme@its.jnj.com) is Professor of Computational Drug Discovery in the University of Leiden, the Netherlands, and Senior Director of the Molecular Sciences Department of Janssen Pharmaceutical Companies of Johnson & Johnson, Beerse, Belgium.

Egon Willighagen (egon.willighagen@maastrichtuniversity.nl) is a computational chemist in the Department of Bioinformatics of Maastricht University, Maastricht, the Netherlands.

Soft materials may enable the automation of tasks beyond the capacities of current robotic technology.

BY ROLF PFEIFER, MAX LUNGARELLA, AND FUMIYA IIDA

The Challenges Ahead for Bio-Inspired ‘Soft’ Robotics

THERE ARE MANY different kinds of robots: factory automation systems that weld and assemble car engines; machines that place chocolates into boxes; medical devices that support surgeons in operations requiring high-precision manipulation; cars that drive automatically over long distances; vehicles for planetary exploration; mechanisms for powerline or oil platform inspection; toys and educational toolkits for schools and universities; service robots that deliver meals, clean floors, or mow lawns; and “companion robots” that are real partners for humans and share our daily lives. In a sense, all these robots are inspired by biological systems; it’s just a matter of degree. A driverless vehicle imitates animals moving autonomously in the world,

a factory automation system is intended to replace humans in tasks that are dull, dirty, or dangerous. The term “robot” itself is anthropomorphic as it is derived from the Czech word “robota,” which is generally translated as “drudgery” or “hard work,” suggesting the analogy to people.

However, if we look inside these robots, we find that for the better part, they function very differently from biological creatures: they are built from metal and plastic, their “brains” are microprocessors, their “eyes” cameras, their “ears” microphones, and their “muscles” electrical motors that sit in the joints. Humans and other animals, by contrast, are built from biological cells; they have muscles made of fiber-like material that pull tendons anchored to the bones of the head, arms, fingers, and legs; they have a soft skin covering the entire body; their sense of sight relies on a retina that spatially encodes visual information and performs a lot of processing right at the periphery; and their brains are

» key insights

- Recent developments in the field of bio-inspired robotics have been centered on the idea that behavior is not only controlled by the brain, but is the result of the reciprocal dynamical coupling of brain (control), body, and environment.
- Future generations of robots will be bio-inspired, have soft bodies composed of soft materials, soft actuators and sensors, and will be capable of soft movements and soft and safe interaction with humans.
- Progress in bio-inspired robotics can only occur when various technologies—computation, sensors, actuators, materials—are integrated and can be made to smoothly cooperate to achieve desired behaviors.
- Because part of the control in bio-inspired soft robotic systems is outsourced to morphological and material properties, novel design principles for “orchestrating” behavior must be developed.
- Bio-inspired soft robotics technologies might entail a quantum leap in the engineering of robots with complex skill sets capable of dexterous manipulation and safe cooperation with humans.

made up of an extremely intricate network of neurons.

Biological inspiration does not imply that we attempt to copy nature. Rather, the goal is to understand the principles underlying the behavior of animals and humans and transfer them to the development of robots. For example, when we walk, our muscles constantly change their stiffness: they are loose when the leg is swinging forward; they stiffen up when we put the foot on the ground. This idea can be employed on robots without having to apply the same “technology” as biological muscles. The important principle is the smooth change in stiffness achieved, for example, with tunable springs, that is, actuators where the spring constant can be dynamically varied.

The degree to which robots resemble biological systems can be used as one of the “dimensions” for characterizing the “robot space” in which we can position the different kinds of robots that exist (see Figure 1). The other dimension we use to characterize the robot space is behavioral diversity or complexity. Behavioral diversity designates the variety of actions a robot is capable of executing. Take a stuffed animal, a Teddy (Figure 1c), as an extreme case: It does not move, so its behavioral diversity is very low, zero in fact. The possible actions of a vacuum cleaning robot such as Roomba are also quite restricted: it can turn left or right, it can turn on and off its vacuuming device, and it can dock onto the charging station if the battery charge starts getting low. By contrast, animals and humans can do many different kinds of things: they can walk, run, look around, find their way to a food source and back to the nest or homes, handle objects such as food and materials for building a nest or house, they can hunt, eat, engage in courtship behavior, and they can reproduce. Note that behavioral diversity is different from technological (or biological) complexity: The autonomous car Stanley (Figure 1b), which won the 2005 DARPA Grand Challenge, has low behavioral diversity—it can,

in essence, speed up, slow down, turn left and turn right—but with its sophisticated sensory systems and its advanced computational facilities, it has extremely high technological complexity.

Diversity is also expressed in the environments in which robots operate. In a factory, for example, the environment is well known so that the robots’ movements can, in principle, be preprogrammed down to the last detail. By contrast, service and companion robots have to function in the real everyday world, such as a city street, a shopping center, a soccer field, a school, or an amusement park where everything changes rapidly, with people rushing about so that individual movements can no longer be planned ahead of time. Humans can function easily and without effort in such real-world situations because they have evolved to cope with environments where predictability is very limited. Moreover, the real world is full of different kinds of objects such as plates, glasses, eggs, tomatoes, bags, hammers and nails, pets, children, and cars that require very diverse sets of skills to interact with and manipulate: Glasses are rigid but fragile and if they are filled with drinks, they must be handled with care so the liquid does not spill. Especially challenging is the treatment of eggs whose shells are very brittle and making a fried egg requires special abilities.

Affetto, the robot baby created by researchers in the Asada Laboratory at Osaka University, is being used to study cognitive robotics. For more on Affetto, visit <http://projectaffetto.blogspot.jp/p/project.html>

Project Affetto is a five-year research project financed by the Japanese Society of the Promotion of Science (grant #24000012) and supported by JSPS KAKENHI (grant #22220002, 09J00755).





Figure 1. Robot space.

A bio-inspired “soft” robot is: **(a)** capable of exhibiting substantial behavioral diversity (not only walk, run, and grasp, but also smile, cook, and cooperate with humans). In contrast, robots today exhibit a rather limited behavioral diversity: autonomous cars **(b)** can “only” speed up, slow down, or turn; toys **(c)** cannot move as flexibly as animals; and most of our research platforms can only do specific tasks such as climbing **(d)**, or swimming in water and walking on land **(e)**. There are also robots with a somewhat higher but still limited behavioral diversity, such as Asimo **(g)** that can walk, run, wave, open doors, recognize objects and people, and interact in basic ways with humans; or the “dog robot” Aibo **(h)** that can also walk, interact with people, and play soccer. Industrial robots, such as CNC machines, robot arms **(f)**, and automated factories **(i)** can conduct many variations of movement, but they are designed for well-defined environments and lack bio-inspired mechanisms.

The imitation of most of these capabilities has to date defied automation.

The kinds of next-generation robots we discuss in this article are those with a high behavioral diversity and a high degree of bio-inspiration^a (see Figure 1a). Similar to humans, such robots will, in addition to hard components such as bones, have soft bodies made out of soft materials, soft actuators (muscles, tendons, and ligaments) and sensors (soft, deformable skin with touch and temperature receptors), and will be capable of soft movements and soft interaction with people. It is our contention that fast, efficient and robust behavior can be achieved by adequately exploiting material properties in particular “softness:” tails of fish and wings of birds passively adapt to the environment during locomotion, the elasticity of

the muscle tendon system supports coping with impact and moving over rough ground, and the deformable, soft tissue in the hand and on fingertips enormously simplifies the task of grasping hard objects. In other words, soft materials may enable us to automate tasks that are beyond the capacities of current robotic technology.

However, introducing softness into the design of robots leads to design issues that differ completely from the classical ones known from “hard” type engineering, and we will thus need to elaborate a novel set of design principles that can act as a bridge between hard and soft type engineering. The main ideas are borrowed from biology and from the concept of embodiment used to characterize the role of the body in intelligent behavior. We therefore claim that with deep understanding of the far-reaching implications of embodiment, we will greatly accelerate the development of soft robotic systems.

In this article we briefly review a selection of recently developed bio-in-

spired robots and derive some of the enabling technologies and principles that have been, explicitly or implicitly, incorporated into their design. We then look at the evolution of the field and identify some of its great challenges. Finally, we argue that we might in fact be on the way to a new industrial revolution.

Biologically Inspired Robots

The biological world is immensely diverse—the total number of named species populating our planet is almost two million. Such variability of life forms provides an exceptional source of inspiration for scientists and engineers, which is reflected in the myriad of robots that have been developed. Bio-inspiration has been derived from fungus-like organisms such as the slime mold⁴² all the way to primates such as monkeys and human beings.^{18,21,31} Biology contains especially rich knowledge for robotics in disciplines such as neuroscience (in particular, computational neuroscience), biomechanics (the science of movement), and systems biology.

Building on the seminal work of Rodney Brooks^{5,6} at MIT in the 1980s, which was the starting point for the field of “embodied intelligence” in artificial intelligence and robotics, a striking variety of bio-inspired robots have been built over the last 20 years. Bio-inspiration has driven research and applications on robot locomotion (crawling, walking, running, climbing, jumping, swimming, and flying), navigation and orientation behaviors, spatial memory formation, exploration, environmental monitoring, manipulation, imitation, and cooperation, among others.

Many of these robots have been used to study and test models of natural neural information processing; for instance, to explore the switching between swimming and walking observed in salamanders²⁰ or to investigate adaptive dynamic walking on irregular terrain³⁰ (for recent surveys, refer to Floreano and Mattiussi,¹⁴ Meyer and Guillot,³² and Pfeifer et al.³⁶). Much attention has also been devoted to emulating navigation and orientation behavior. Examples abound and include visual homing inspired by how bees or wasps find their way back to their nests, cricket phonotaxis (how

a Sometimes “bio-inspired” is taken to designate a very broad class of robots whereas “bio-mimetic” is used to suggest a close resemblance to the biological system. Generally, however, bio-inspired bio-mimetic, and bionic are used interchangeably.

female crickets move toward the mating sounds of males in highly rugged and noisy environments⁴⁵), and spatial memory formation by modeling place fields and head-direction cells that account for the remarkable navigational skills of rodents.⁴⁸ Soccer robots are a special breed of robots that are bio-inspired in a broad sense, where at least until recently, the capabilities for rapid movement, kicking skills, perception, cooperation, and strategy^b are of greater interest than external shape and appearance.

A recent trend in bio-inspired robotics is to simplify the typically computationally intensive neurally inspired control through clever morphological design and use of functional materials.^{9,12,19} A case in point is aquatic locomotion. The key to the control of underwater robots, which mostly have a multi-segmented structure, is the translation of computational activity into torques propagating through the individual segments, so the resulting forces lead to forward movement. An alternative strategy to building such robots is to under-actuate them, that is, to drive only some of the joints, leaving the others passive, and “outsource” as much as possible to the morphology and bio-inspired materials. Simply moving a tail back and forth gives rise to surprisingly lifelike movements if the material of the tail fin is chosen appropriately; such materials allow the “tuning” of the mechanical properties of the fin in a way that optimally distributes the hydrodynamic forces over the fish’s body during propulsion and maneuvering^{13,50} (see Figure 2i).


As for swimming, also for flying one possible avenue might be the exploitation of passive dynamics and the morphology of bio-inspired materials^{8,44} (see Figure 2j). Take an insect wing during hovering flight. Its material properties in terms of resilience, stiffness, and deformability are essential to generate adequate lift in the absence of any forward velocity. It has been observed that the shape of the wing changes greatly when moving back and forth through the stroke plane. Although such change in shape could in principle be actively con-

trolled, it is more efficient and faster if the intrinsic material characteristics are exploited and control is outsourced to the morphological and material properties of the wing. An additional advantage of this solution is that the wings can be made much lighter because less actuation is required.

Finally, materials can also be exploited for climbing, as beautifully

showcased by the uncanny climbing skills of geckos that can dash up smooth walls and walk across ceilings with great ease. The geckos owe their sticky feet to the structural properties of their toes, which are covered with millions of nanoscale hair-like stalks branching into hundreds of tiny endings. The use of micropatterned fibrillar dry adhesives inspired by gecko


Figure 2. Recent bio-inspired robots.



- (a) Scratchbot, a mobile robot equipped with two 3×3 arrays of actively controlled rat-like whiskers used to test models of action arbitration by model basal ganglia.³⁴
- (b) This insect-like miniature jumping robot is able to perform a jump in a given direction, land, and then jump again. It only weighs 14g, has a size of 18cm, and can jump up to 62cm.²²
- (c) The Festo Bionic Handling Assistant is inspired by elephant trunks and due to its intrinsic compliance can safely interact with humans. Because of its mechanical construction, the gripper passively adapts to the shape of the object.
- (d) GoQBot is a soft material robot capable of rapidly curling its body into a wheel-like structure for a ballistic rolling locomotion inspired by caterpillars. The robot has a composite body consisting of several mixtures of silicone rubbers, and shape memory alloy is used to control the body shapes.²⁶
- (e) This robot arm inspired by the octopus is extremely soft, flexible, and continuous. Inside the arm there are 18 transverse actuators (shape memory alloys) and 12 cables anchored along the arm. This enables the arm to elongate, shorten, and bend in any direction, as well as to stiffen in different parts of the arm (picture courtesy Matteo Cianchetti, SSSA, Pisa, Italy).
- (f) Boston Dynamics’ fuel-powered four-legged robot BigDog³⁷ is dynamically stable, 0.9m long, and 0.7m tall, and weighs 110kg. It can traverse rough terrain at a speed of 6.5km/h while carrying 150kg.
- (g) The child humanoid robot CB² is designed for studies on robot learning, development and human-robot interaction.¹³³ It is covered with a fully sensorized soft skin, has flexible joints, and is actuated by 51 pneumatic cylinders.
- (h) The self-organizing modular robotic system, Slimy II, moves through local interaction dynamics.⁴² Although no module can move on its own, by using neural oscillators as drivers for the actuators and through the physical coupling between the units, a coordinated global wave of activation can be induced, which leads to forward movement, even though there is no global control.
- (i) The goal of the FILOSE robot is to gain a better understanding of the principles underlying fish locomotion and sensing.¹³ It has a continuous flexible body whose morphology and material characteristics are such that a minimal set of control parameters are sufficient to excite various modes of locomotion.
- (j) Phoenix is a fully computerized bird-like flapping robot with passively adaptable wings. It has a wingspan of 2m, weighs only 1kg, and is capable of perching.⁹

^b See <http://www.robocup.org>.

Figure 3. The anthropomimetic humanoid robot ECCE (Embodied Cognition in a Compliantly Engineered Robot).



To achieve large behavioral diversity and soft interactions in uncertain environments, many bio-mimetic technologies have been implemented including exteroceptive sensors (visual and haptic), flexible skeletal structures (rigid limbs connected through passive joints), deformable materials to enable flexible body movements, energy-efficient body actuation with series-elastic actuators, proprioceptive sensors (length and force sensors to measure body movement), and an under-actuated control architecture to allow soft interactions with humans.³¹

foot morphology is bound to lead to impressive advances in the construction of robots that can climb vertical or inverted surfaces of all kinds.⁴³

A field that has also received considerable attention over the past decade is humanoid robotics. Substantial efforts have been directed toward engineering robots with high behavioral diversity capable of performing a large variety of “human-oriented” tasks such as assisting the elderly (healthcare, physical support, shopping, leisure), doing household chores (washing the dishes, cooking dinner, and ironing), helping workers on assembly lines, as well as surveillance, and entertainment. Much progress has been made in the study of basic abilities such as bipedal locomotion;¹⁹ manipulation;^{12,24} understanding the surrounding environment, including the recognition of objects, people, and other robots;¹ and social interaction^{3,33} (Figure 2g). Because such robots need to operate in environments built for humans, a reasonable morphology of choice is humanoid or anthropomorphic.

Humanoid robots often have highly

sophisticated sensory-motor systems, and as a consequence they are confronted with the difficult problem of processing potentially large amounts of information in real time. Although much research has been conducted on learning in the real world, especially in the fields of artificial intelligence and cognitive robotics,⁴¹ the tasks and environments have, for the most part, been of relatively limited complexity. Potential reasons might be the size of the search spaces for learning optimal decision policies in realistic scenarios and the slow operating speed of real robots, which implies that the direct transfer of traditional algorithm-based machine-learning techniques to robots is not straightforward. A possible solution to the problem might be imitation learning, in which robots learn from humans or other robots.^{1,3,11} This idea has a special appeal because imitation is a powerful mechanism for reducing the search spaces associated with learning in the real world, which could lead to robots that will require only a minimal amount of programming. It is important to note

that, in the context of learning by imitation, the embodiment has to be carefully considered because such learning would not work as efficiently if the mechanical properties of robot bodies were completely different from the ones of the teacher (humans).

Let us now draw some conclusions concerning technologies that have in fact enabled the construction of these robots, and about principles that have—explicitly or implicitly—been employed in their design and construction.

Enabling Technologies

As we move toward more highly complex anthropomorphic robots, a number of enabling technologies are increasingly important for progress to take place. In a complex robot such as ECCE (Figure 3) a wide range of technologies, which are complex in themselves, have to be integrated and they must smoothly cooperate to achieve the desired behaviors: sensors, actuators, materials for body limbs and skin, and computational resources. This is especially true for bio-inspired soft robotics, where the physical properties of the body, for instance the elasticity of the actuators and the deformability of the surface of the hands and the feet, are instrumental to cope with impact in walking and for safe interaction between man and machine. Moreover, because for the better part, bio-inspired robots are mobile, they require some kind of portable power source, and if they are to interact closely with humans in everyday life, they should not be too big. This implies on the one hand the need for efficient power sources and actuators, and on the other that the machines should be light and if possible also compact.

Let us start by reviewing recent developments in biologically inspired sensory systems. Due to space constraints, we focus on visual and tactile sensing and do not review other sensor modalities (for example, hearing, smell, taste, and pain).

The field of computer vision has a long history and has produced a large number of applications, for example in surveillance, quality control for electronic circuit boards, image search, automatic driving, traffic enforcement cameras, face recognition systems, among others.⁴⁰ Biologically inspired

vision attempts to replicate—and even go beyond—some of the amazing features of the human perceptual system, such as its extremely high reaction speed, its capacity for recognizing objects at varying distances, orientations, and partial occlusions, its ability for reliable interpretation of the environment, for localization, for providing a stable image while moving, and its adaptivity to dramatic changes in lighting conditions (from bright sunlight to near darkness). All these competences are still unmatched by today's vision-based robots.


One potential avenue toward truly bio-inspired vision systems is neuromorphic engineering—a discipline that aims at realizing artificially engineered systems inspired by the physical foundations, function and structure of biological nervous systems.

Novel (neuromorphic) silicon retinas offer advantages over conventional cameras, such as a higher range of sensitivity, speed, color detection, and the discounting of shadows—an ability that is currently under development²⁷ (see Figure 4c). These silicon retinas mimic the computational principles of biological vision and rely on a continuous stream of asynchronous events from individual pixels—equivalent to the spikes delivered from the retina to the brain via the optic nerve—yielding activity that represents, for instance, scene contrasts and contrast changes. Such activity-driven, event-based systems have been tested in many applications where rapid responses are required such as in tracking systems, soccer playing robots, and pole balancing tasks.⁷ At this point in time, vision systems with artificial retinas

are still laboratory prototypes but because of their many desirable characteristics, we expect the technology to expand rapidly in the near future. Similar neuromorphic technologies have also been successfully applied to silicon cochleae.²⁷

One of the most important organs of humans (and many animals) is the skin, and there is a vast amount of evidence demonstrating its significance in survival and intelligent behavior. The skin is equipped with touch and temperature sensors (and pain receptors, which we will not discuss further in this article) that are distributed over the entire organism. Their distribution is not homogeneous but varies greatly depending on their position on the body: in the hand and on the fingertips, in the face and on the lips, their density is very high (about 250

Figure 4. Bio-inspiration and enabling technologies.



(a) Muscle-like motor functions can be realized, for example, with pneumatic artificial muscles or electro-active polymers. These actuation technologies are used in many bio-inspired robots.

(b) To equip robots with a sense of touch, a number of different types of pressure-sensitive artificial skins relying on various transduction mechanisms (force-sensitive resistors as used on the hand; MEMS-based pressure sensors), type of substrate, and spatial resolutions have been developed. Touch remains hard to mimic.

(c) Bio-inspiration has also been applied to analog computing in visual sensors as well as a space-variant distribution of photoreceptors—typically a high density in the center, and a low density at the periphery.

(d) Different approaches of distributed computation are being exploited in robotics projects, for example, sensor networks and distributed control for multi-robot systems.

mechanoreceptors per sq. cm in the fingertips alone), whereas on the back, for example, it is low (see Figure 4b). Skin is soft and deformable, but at the same time robust and waterproof. Because of the central importance of tactile sensing for learning, manipulation, and tool use, improvements in skin technology are likely to lead to a quantum leap in soft robotics. Various technologies have been suggested (for a detailed review of tactile sensing in robotics, see Dahiya et al.¹⁰). The approach to mimicking the pressure sensors in the skin on the fingertips shown in Figure 4b is based on force-sensitive resistors built into a flexible fabric. The fabric can be freely bent without obstructing the function of the sensors that is, of course, essential because the skin has to be placed everywhere on the hand (and perhaps other parts of the robot). In CB², the child robot^{1,33} (see Figure 1g), the entire body is covered with skin, which is important because skin is a prerequisite for forming an image of one's own body while growing up.¹⁷

Let us now switch to actuator technologies, keeping in mind that sensor mechanisms are often integrated with movement systems, for example, in the human muscle-tendon complex, which not only actuates the body but also incorporates sensing for force and length of muscles via muscle spindles and Golgi tendon organs. Proper actuation for bio-inspired robots has turned out to be a notorious issue and has so far been a decisively limiting factor.¹⁹ Biological muscles have many desirable properties: they have a high contraction ratio, they are energy efficient, they are intrinsically compliant, and their stiffness can be varied smoothly and dynamically. In order to deal with impact, the stiffness is increased, as pointed out earlier. Most of the robots shown in Figure 2 can, to some extent, change the compliance of their actuators, which they exploit to take over some of the control functions (coping with impact, smooth manipulation, and interaction). One type of actuator that dynamically changes its compliance is the pneumatic artificial muscle or fluidic air muscle (Figure 4a). This actuator is a contractile device that consists of an internal bladder surrounded by a braided



Because of the central importance of tactile sensing for learning, manipulation, and tool use, improvements in skin technology are likely to lead to a quantum leap in soft robotics.



mesh shell attached at either end to fittings. When the internal bladder is pressurized, the actuator shortens and its stiffness increases. In other words, pneumatic artificial muscles are intrinsically compliant, and their compliance can be controlled via the global parameter of air pressure. Another type of compliant drive system used, for example, in the ECCE-Robot (Figure 3) or in MIT's Domo,²¹ is the series-elastic actuator, where an elastic element (for example, a spring) is placed between the output of a motor's gear train and the load. Many actuation technologies have been proposed, but none has achieved the performance level of biological muscles. Note that using standard electrical motors in the joints the change in compliance can only be achieved at a high cost in terms of electronics and computation.

Highly complex robots such as ECCE require elaborate middleware technology, that is, computer software that acts as a bridge between the multitude of heterogeneous and interconnected hardware and software modules composing them (sensors, actuators, software frameworks, programming languages, among others).^c Robotics middleware is an active field of research and there are great hopes that it will allow robot designers to make the necessary abstractions to come to grips with the increasing complexity (and messiness) of robotic systems.

Perhaps more than any other factor, the lack of good power sources has limited the progress and the infiltration of mobile robots into our daily lives. Although computer processors and sensors have become cheaper and more powerful by the year, power generation and storage are still inefficient and, in the case of batteries, heavy, and slow to recharge, leaving engineers and scientists to dream of a day when they will have the juice to instill long lives into their creations. The BigDog robot (see Figure 2f) is driven by a combustion engine, which is convenient in terms of weight and fuel efficiency, but also very noisy and can only be used outdoors. The energy demands of legged robots are especially high compared to wheeled robots. But while wheels

^c See <http://www.ros.org/wiki/>.

are effective in their energy use, they are limited to hard, smooth surfaces like roads. Recently, there have been some fascinating approaches to tackle the energy-efficiency issue in walking robots. One of the champions in this league is the somewhat odd looking Cornell Ranger that can go for 40.5 miles (65.2km) untethered on one battery charge alone in about 30 hours and 49 minutes.

Because most of the bio-inspired sensor and actuator technologies only exist in the form of laboratory prototypes and cannot be bought off the shelf, they often have to be manufactured individually for each robot. Rapid prototyping tools, such as 3D printers and freeform fabrication technologies represent a substantial enabling factor for this kind of engineering research. If, in addition, the 3D machines are capable of printing active, multi-material components, this will enable the automatic integration of sensors, logic gates, actuators, and power sources into complex structures, rendering the prototyping process even more effective.^{9,29,47}

Due to the many unknowns, especially in the context of soft robotics, we also need powerful simulation tools so that ideas can be quickly tested before the systems are actually constructed in real hardware. In recent years, simulation methods that can be applied to robotics such as artificial evolution and physics-based modeling, have enormously improved (once again profiting from Moore's Law) and are now standard development tools.^{4,14} This is not to say that the simulation of soft robots is an easy task; in fact, quite the opposite is the case and a lot of theoretical advances will be necessary before soft robots can be adequately simulated.

Progress in any of these technologies will lead to progress in bio-inspired robotics.

Design Principles


From the examples reviewed earlier, we can now derive a number of theoretical design principles that are beginning to make their way into the robotics community. These principles, which are compatible with current work in biomechanics, neuroscience, engineering, and embodied intelligence,^{34,35} can

be seen as an extension of the work by Brooks mentioned earlier. We feel they represent an excellent starting point for future theoretical developments.

Physical embedding and task distribution. First, the behavior of a system is not merely the outcome of some internal control structure (such as the brain or a microprocessor), but it is also shaped by the environment in which the system is physically embedded, and by its morphological and material characteristics. For example, the elasticity of muscles can help to cope with the unevenness of the ground in walking, or the distribution of the receptor cells in the human eye already provides spatial information that substantially reduces processing cost in perception.²⁵ Because of this "task distribution," which is common to all humans and animals, a fundamental

rethinking of our classical notion of control is required. There is a kind of trading space where computation is "outsourced," so to speak, to the physical aspects of the agent, which implies that often the central control or computational effort can be reduced by orders of magnitude and the reaction times are dramatically shortened. In the jumping robot (Figure 2b), dealing with impact is delegated to the springy actuator system, in the Bionic Handling Assistant (Figure 2c), adaptation to the shape of the object is taken over by the mechanics of the gripper, which is based on the so-called fin-ray effect known from fish, and in the FILOSE robot (Figure 2i) finding the optimal shape for the tail fin is taken over by its material properties that, depending on the task, can be dynamically adjusted. In all these cases, very little central

Figure 5. An illustration of the design principles for bio-inspired robots.



- (a) Behavior is not the outcome of an internal control structure only; computation is "outsourced" to body morphology and material properties.
 - (b) All components of the agent are coupled; movement capabilities have to match those of the sensory systems.
 - (c) There is a direct link between embodiment and sensory information; each action leads to patterns of sensory stimulation.
 - (d) The behavior of agents is characterized in terms of attractor states and transitions between them.
- (Cartoon by Shun Iwasawa, adapted from Pfeifer and Bongard.³⁵)

computational control is necessary.

Physical dynamics and information processing. Second, a direct link exists between embodiment and information: Coupled sensory-motor activity and body morphology induce statistical regularities in the sensory inputs, which simplifies the information processing task of the brain (or generally, the controller^{28,36}). This information structure, which is the result of a physical interaction with the world, is at the core of how we can learn something about the environment. It depends strongly on the agent's shape, the physical characteristics and the distribution of the different sensory systems, on the particular action and, of course, on the environment itself. If the robot shown in Figure 2a stimulates its whiskers by moving past an object, the sensory patterns induced will depend on the robot's motion and on the object's shape and surface characteristics. Or, when we walk, the environment travels across the visual field, a phenomenon called optic flow. It turns out that optic flow cannot only be easily calculated but is also extremely useful in navigation tasks. In order to avoid obstacles, optic flow is sufficient because nearby objects induce more optic flow than those further away. Note that the optic flow is induced through the agent's own movement, it is not passively perceived. The walking robot, BigDog (Figure 2f) is continuously generating patterns of stimulation in the force sensors in its legs and its pressure receptors on its feet, which can be exploited for gait stabilization. The FILOSE fish robot (Figure 2i) creates stimulation in its flow sensors that delivers valuable information about how well its actions translate into forward movements.

Self-organization and emergence. Third, because robots are embodied, they can be viewed as complex dynamical systems that enable us to apply concepts of self-organization and emergence, rather than top-down control. Results from biology and bio-inspired robotics suggest that stable movement patterns, for example, can be productively characterized as attractor states. Agents display self-organization and emergence at multiple levels: induction of sensory stimulation, movement generation, exploitation of shape and material properties and interaction be-

tween individual modules and entire agents. For example, because both the swimming (Figure 2i) and the flying robot (Figure 2j) are under-actuated, that is, not all the degrees of freedom are driven, their bodies or wings will self-organize into the proper movements, namely, they will do the right thing even though they are not directly controlled. The movement of the Slimebot⁴² (see Figure 2h) is the result of a decentralized control in which each individual module is actuated independently from the others and whose global, emergent activity is coordinated through the physical interaction with the environment—yet another beautiful illustration of self-organization.

Complete agents. Finally, embodiment implies that all the components of the agent continuously interact and influence each other, and this needs to be kept in mind during the design process. For example, the movement capabilities must match those of the visual system: a snail with a powerful vision system such as the one of humans would not make any sense because, even if it had the brain to detect an approaching bird (which it has not), it could not do anything about it. The four legs of the BigDog robot are not only coordinated via internal connections, but because they are part of the same physical agent, each leg movement will instantaneously influence the values in the force and pressure sensors in the other legs. Also, induction of optic flow happens because the visual system is part of a complete agent that moves around in the world.

Although a certain consensus on these principles is emerging in the robotics and related communities, the new field of bio-inspired soft robotics is still lacking a firm foundation like control theory for traditional robotics and factory automation. The theory needs to be further developed, and we must work toward a better understanding of how behavior is “orchestrated” rather than controlled. This requires, among other issues, an elaboration of the trading space outlined here and implies a quantitative approach to morphology that, alas, is still missing.

In order to get a better idea of where the field of bio-inspired robotics might be going, let us briefly sketch its major challenges.

Robotics: Tough Challenges

To be sure, there has been enormous progress in robotics over the last 50 years—as can also be inferred from our review of some of the recent bio-inspired robots. In 1961, the first industrial robot, called Unimate, was commercialized and joined the assembly line of a General Motors plant; about 10 years later, there were already 3,000 robots operational worldwide; in 2009 the operational stock of industrial robots amounted to more than one million. If we look at mobile service robots, the progress is even more impressive. There were all of about 10 in the 1970s, mostly at universities and research institutes, whereas in 2009 almost 10 million had been sold worldwide, many of them used daily in homes for vacuum cleaning and lawn mowing, for entertainment and education, in hospitals, on disaster sites, or for surveillance purposes. According to the projections of the International Robotics Federation the sales for service robots for personal use are expected to increase sharply in the near future.

In spite of this quite impressive history, robot engineers have been jealously watching the rapid progress in computer technology and they have been asking why it is that microprocessor technology excels continuously—by now Moore's Law has been in effect for several decades—whereas progress in robotics has been comparatively slow. The way most robots walk today is still very unnatural, their perceptual and manipulation skills are extremely limited, and they get easily confused if an unexpected situation occurs. We can only speculate about the reasons but this may help us to identify some of the major challenges.

First, robotics has become strongly interdisciplinary and involves many different fields of expertise, as well as researchers and engineers with a broad set of skills. Robots are highly complex systems. For instance, take humanoids such as ECCE, Domo, or Kojiro.³⁸ They have a head, arms, hands, many joints, muscles (or motors), tendons, “eyes,” a skin with touch sensors, sensors to measure accelerations of individual body parts, force and length sensors for the tendons, and, of course, microprocessors with programs for perception, for con-

trolling the robot's movements, decisions, and learning behavior. In order for such robots to behave flexibly and efficiently, advances in sensors, actuators, energy, and propulsion technologies will be required. Moreover, and this represents a big theoretical challenge, the interaction among these components will have to be coordinated and orchestrated to achieve the desired behaviors. One of the reasons (biological) humans can walk and move so stably is that they are equipped with an enormous number of smoothly integrated sensors and a highly redundant muscle-tendon system. The walking robot BigDog has over 50 different sensors,³⁷ but a humanoid, in order to competently walk, run, and manipulate objects, will need many more.

There is an additional challenge, especially for mobile machines, that we alluded to earlier: Because robots are physical systems, they require energy to move, which implies there is a hard limit to how much their energy consumption can be reduced. Likewise, in the commercial airline industry, progress has also been slow, at least much slower than in microprocessor technology, not only because many technologies have to come together to make it work, but because there is a minimum amount of energy required to transport humans and the airplane itself, both physical systems. Power consumption for microprocessors is also an issue but so far the theoretical limits in terms of energy required for processing and storing information have not been reached. In robotics, energy consumption and storage is a notorious problem.

So far, we have been stressing the fact that focusing on computational aspects alone will not be sufficient if we are to build soft, real-world machines. However, computation has formed, and will form in the future, a core component of any intelligent machine: perception, decision making, problem solving, and action generation will always require sophisticated kinds of computation. What the new theoretical insights from the area of embodied intelligence are bringing into the discussion is, as introduced earlier, the idea that certain aspects of the computation can be off-loaded to morphological



Given the enormous challenges we are up against when engineering robots for the real world, why should we do it in the first place?



and material components. This not only simplifies computation but also increases reaction speed—physical processes are often much faster than computational ones. Given the experience accumulated within computer science over a half century, it seems natural to join forces in an attempt to devise a novel concept of computation—which has also been called “morphological computation”—and see how the existing concepts could be transferred or extended to capture the additional phenomena in soft robotics, or to what extent new ideas are required. Undoubtedly, researchers will strongly benefit from this rich experience. We are convinced that computational thinking, as outlined in Jeannette Wing's manifesto,⁴⁶ will also apply to this challenge: it is not primarily about algorithms or programs, but a way of thinking about the world.

A New Industrial Revolution?

Given the enormous challenges we are up against when engineering robots for the real world, why should we do it in the first place? There are a number of key driving forces behind all these developments: scientific, social, demographic, and economic.

First off, scientists are continuously striving to overcome the limitations of current, “hard” robotics in order to move to the next level of behavior diversity. As argued throughout this article, this implies using a bio-inspired soft robotics approach. The ultimate scientific goal (and dream) is, of course, to have a robot with the skills of a human. Throughout the history of mankind, many attempts have been made to imitate humans and their abilities by exploiting the technologies and scientific insights at the time. From a scientific perspective, bio-inspired robotics is of great interest because with the soft technologies the area has become extremely interdisciplinary—it might in fact be one of the most interdisciplinary scientific fields around. The increased level of complexity of the subject matters requires novel ways of cooperation, supported by state-of-the-art collaboration technologies.

At the social level, mobile communication and computing technologies are rapidly spreading and increas-

ingly penetrating our daily lives. This is now also beginning to happen with robot technology, for example, pet robots, service robots, health monitoring and support technologies for the elderly, toy robots, educational robots, and finally, it will be household and companion robots. As Bill Gates forcefully requested in his famous *Scientific American* article,¹⁶ “a robot in every home;” just as he demanded a computer in every home more than 20 years ago! Thanks to the ubiquitous use of mobile devices, technology in general is becoming more and more accepted as a close interaction partner. It is likely that the next step will be the integration of autonomous systems into our daily lives—a development that is already under way in Japan and Korea. In order to be able to perform tasks similar to those humans do, for example, shopping, cooking, cleaning the dishes, tidying up the children’s rooms, or pouring a glass of beer, they will have to be of the “soft” type.

Demographically, one of the core drivers is a rapidly aging population: Technologies must be developed to enable individuals to live autonomously for as long as possible. This includes service robots of all types, assistive technologies to compensate for lost bodily functions, “helpers” in everyday tasks, and monitoring systems for physiological conditions. Again, because Japan was the first to be strongly confronted with this issue, it is the country where the pertinent technologies are most advanced. Because similar tendencies are on their way in Europe and the U.S., massive R&D efforts in this direction backed by large investment commitments have taken off in these two regions.

This demographic development suggests two scenarios. The first and most commonly outlined is the companion robot—a robot capable of performing most or all of the tasks that humans can, as well as hold intelligent and personal conversations, and give medical, financial, and entertainment advice, and so on. Obviously, it must be strongly biologically inspired and must have extremely high behavioral diversity. Whether we will ever have such robots populating our homes is an open question, especially



We humans can do many things, but for most sensory-motor tasks, there are—or will be—machines that perform the task faster, cheaper, and more precisely.



if we think of all the challenges we have pointed out here. Typically, such super-complex systems tend to be extremely expensive and fragile in their operation. Ultimately, it will be the markets that will decide on the viability of this concept. But there is an even more important caveat. Let us look at real biological humans for a moment. We humans can do many things, but for most sensory-motor tasks, there are—or will be—machines that perform the task faster, cheaper, and more precisely. For example, driving a screw the way we do is probably the most inefficient and slowest possible way: it would be much better to have a screwdriver motor that can rotate continuously. So why develop a machine, a companion robot that can perform many different things but nothing really well?

This implies an alternative scenario, one in which we will have many specialized machines specifically geared toward particular tasks such as vacuum cleaners, lawn mowers, dish-handling machines, shopping assistants on wheels, automated waiters, among others. While these robots will for the better part not be humanoid, they will, depending on the particular tasks, employ soft robotics principles and technologies. In other words, what we have learned from building super-complex humanoids can be exploited and transferred to specialized machines. Festo’s bionic handling assistant or the BioRob arm,²⁴ designed for manipulation and safe interaction with humans are not humanoid but incorporate soft robotics concepts. But, as suggested earlier, the jury is still out on which scenario will in fact materialize.

The last driving force, and a very strong one, is economic. Recently, there has been a lot of talk in the media that in certain areas in China a shortage of qualified labor is beginning to emerge and workers are asking for higher salaries and better working conditions, and they even risk going on strike. The country’s one-child policy, for the better or the worse, is beginning to show its effects. Thus, it can no longer be the default strategy for European and U.S. companies to outsource manufacturing and assembly tasks that are beyond the current level of factory automation to

low labor cost countries. Foxconn, for example, produces and packages iPhones and iPads for Apple and components for many other electronic equipment manufacturers. Western enterprises are starting to ponder the idea of re-insourcing some of the production back into their own countries—in particular, by empowering small- and medium-sized companies with next-level low cost but highly refined automation solutions. However, in order to remain competitive, the initial capital requirements must be reduced and the degree of automation increased, such that tasks currently beyond what existing manufacturing methods can achieve in terms of price and speed, can be automated as well. Moreover, there is the additional requirement of flexibility and versatility: The tasks must be switchable quickly without the need for a lot of tedious reprogramming.

In spite of enormous advances in robotic tools to support task switching to date, the rapid adaptability of humans to new tasks remains unparalleled. And this is where the soft approach to robotics might in fact come into play, because it is likely to lead to robots with complex skill sets capable of dexterous manipulation and safe interaction and cooperation with humans, for example, where humans could teach tasks to robots by demonstration. Rethink Robotics and Adept located in the U.S. and pi4_robotics from Germany are companies specializing in next-level factory automation. So, what we have learned and are still learning by building full-fledged, super-humanoid robots might in fact have spinoffs that could ultimately lead to a new industrial revolution. **C**

References

- Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., Ogino, M. and Yoshida, C. Cognitive developmental robotics: A survey. *IEEE Trans. on Autonomous Mental Development* 1, 1 (2009), 12–34.
- Ayers, J. and Witting, J. Biomimetic approaches to the control of underwater walking machines. *Phil. Trans. A Math. Phys. Eng. Sci.* 365, 1850 (2007), 273–295.
- Billard, A., Calinon, S., Dillmann, R. and Schaal, S. Robot programming by demonstration. *Handbook of Robotics*. B. Siciliano and O. Khatib, eds. Springer-Verlag, Berlin, Heidelberg, 2008, 1372–1394.
- Bongard, J.C. Morphological change in machines accelerates the evolution of robust behavior. *Proc. Nat. Acad. Sciences USA* (2011).
- Brooks, R.A. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2 (1986), 14–23.
- Brooks, R.A. *Cambrian Intelligence: The Early History of the New AI*. MIT Press, Cambridge, MA, 1999.
- Conradt, J., Cook, M., Berner, R., Lichsteiner, P., Douglas, R.J. and Delbruck, T. A pencil-balancing robot using a pair of AER dynamic vision sensors. In *Proc. of Int. Conf. on Circuits and Systems*, (2009), 781–784.
- Cory, R. *Supermaneuverable Perching*. Ph.D. Thesis. MIT, Cambridge, MA, June 2010.
- Cutkosky, M.R. and Kim, S. Design and fabrication of multi-material structures for bioinspired robots. *Phil. Trans. R. Soc.* 367, (2009), 1799–1813.
- Dahiya, R.S., Metta, G., Valle, M. and Sandini, G. Tactile sensing—From humans to humanoids. *IEEE Trans. on Robotics* 26 (2010), 1–20.
- Dillmann, R., Asfour, T., Cheng, G. and Ude, A. Toward cognitive humanoid robots. Special Issue: *Int. J. of Humanoid Robotics* 5, 2 (2008).
- Dollar, A.M. and Howe, H.D. The highly adaptive SDM hand: Design and performance evaluation. *Int. J. of Robotics Research* 29, 5 (2010), 585–597.
- Fiazza, C., Salumae, T., Listak, M. et al. Biomimetic mechanical design for soft-bodied underwater vehicles. In *Proc. of IEEE Oceans* (2010), 1–7.
- Floreano, D. and Mattiussi, C. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. MIT Press, Cambridge, MA, 2008.
- Floyd, S. and Sitti, M. Design and development of the lifting and propulsion mechanism for a biologically inspired water running robot. *IEEE Trans. on Robotics* 24, 3 (2008): 698–709.
- Gates, B. A robot in every home. *Scientific American* (Jan. 2007) 58–65.
- Hoffmann, M., Marques, H.G., Arieta, A., Sumioka, H., Lungarella, M. and Pfeifer, R. Body schema in robotics: A review. *IEEE Trans. on Autonomous Mental Development* 2, 4 (2010), 304–324.
- Hosoda, K., Sakaguchi, Y., Takayama, H. and Takuma, T. Pneumatic-driven jumping robot with anthropomorphic muscular skeleton structure. *Autonomous Robots* 28, 3 (2009), 307–316.
- Iida, F. Biologically inspired motor control for underactuated robots—Trends and challenges. *Robot Motion and Control*, K.R. Kozlowski, ed. *LNCIS* 396, (2009), 145–154.
- Ijspeert, A.J., Crespi, A., Ryzcko, D. and Cabelguen, J.M. From swimming to walking with a salamander robot driven by a spinal cord model. *Science* 315, 5817 (2007), 1416–1420.
- Kemp, C.C., Edsinger, A. and Torres-Jara, E. Challenges for robot manipulation in human environments. *IEEE Robotics and Automation Magazine* 14 (2007), 20–29.
- Kovac, M., Schlegel, M., Zufferey, J.-C. and Floreano, D. Steerable miniature jumping robot. *Autonomous Robot* 28 (2010), 295–306.
- Laschi, C., Mazzolai, B., Cianchetti, M., Margheri, L., Follador, M. and Dario, P. A soft robot arm inspired by the octopus. *Advanced Robotics* 26, 7 (2012), 709–727.
- Lens, T., Kunz, J., Trommer, C., Karguth, A. and von Stryk, O. BioRob-Arm: A quickly deployable and intrinsically safe, light-weight robot arm for service robotics applications. In *Proc. of 41st Intl. Symp. on Robotics*, (2010), 905–910.
- Li, Z., Wei, S. and Shi, Z. Simulated distribution of the retinal photoreceptors for space variant resolution imaging. *Information Technology Journal* 8, 5 (2009), 717–725.
- Lin, H.-T., Leisk, G.G. and Trimmer, B. GoQBot: A caterpillar-inspired soft-bodied rolling robot. *Bioinspiration and Biomimetics* 6, 2 (2011), 026007.
- Liu, S.-C. and Delbruck, T. Neuromorphic sensory systems. *Current Opinion in Neurobiology* 20 (2010), 1–8.
- Lungarella, M. and Sporns, O. Mapping information flow in sensorimotor networks. *PLoS Computational Biology* 2, 10 (2006), e144.
- Malone, E., Berry, M. and Lipson, H. Freeform fabrication and characterization of Zn-air batteries. *Rapid Prototyping Journal* 14, 3 (2008), 128–140.
- Manoonpong, P., Woegotter, F. and Pasemann, F. Biological inspiration for mechanical design and control of autonomous walking robots: Towards life-like robots. *Int. Journal of Applied Biomedical Engineering* 3, 1 (2010), 1–12.
- Marques, H. et al. ECCE1: The first of a series of anthropomorphic musculoskeletal upper torsos. In *Proc. of IEEE Conf. on Humanoid Robotics* (2010), 391–396.
- Meyer, J.-A. and Guillot, A. Biologically inspired robotics. *Handbook of Robotics*, B. Siciliano and O. Khatib, eds. Springer-Verlag, Berlin, Heidelberg, 1395–1418, 2008.
- Minato, T., Yoshikawa, Y., Noda, T., Ikemoto, S., Ishiguro, H. and Asada, M. CB2: A child robot with biomimetic body for cognitive developmental robotics. In *Proc. of IEEE-RAS Int. Conf. on Humanoid Robots* (2007), 557–562.
- Pearson, M., Pipe, A., Melhuish, C., Mitchinson, B. and Prescott, T. Whiskerbot: A robotic active touch system modeled on the rat whisker sensory system. *J. of Adaptive Behavior* 15, 3 (2007), 223–240.
- Pfeifer, R. and Bongard, J.C. *How the Body Shapes the Way We Think—A New View on Intelligence*. MIT Press, Cambridge, MA, 2007.
- Pfeifer, R., Lungarella, M. and Iida, F. Self-organization, embodiment, and biologically inspired robotics. *Science* 318 (2007), 1088–1093.
- Raibert, M., Blankenspoor, K., Nelson, G., Playter, R. et al. BigDog, the rough-terrain quadruped robot. In *Proc. of 17th World Congress of the Int. Fed. of Automatic Control*, (2008), 10822–10825.
- Sodeyama, Y., Nishino, T., Namiki, Y., Nakanishi, Y., Mizuuchi, I. and Inaba, M. The designs and motions of a shoulder structure with a spherical thorax, scapulas, and collarbones for humanoid “Kojiro”. In *Proc. of Int. Conf. on Intelligent Robots and Systems* (2008), 1465–1470.
- Spenko, M.J., Haynes, G.C., Saunders, J.A., Cutkosky, M.R., Rizzi, A.A., Full, R.J. and Koditschek, D.E. Biologically inspired climbing with a hexapedal robot. *Journal of Field Robotics* 25, 4–5 (2008), 223–242.
- Szeliski, R. *Computer Vision: Algorithms and Applications*. Springer Verlag, London, 2011.
- Thrun, S., Burgard, W. and Fox, D. *Probabilistic Robotics*. MIT Press, Cambridge, MA, 2005.
- Umedachi, T., Takeda, K., Nakagaki, T., Kobayashi, R. and Ishiguro, A. Fully decentralized control of soft-bodied robot inspired by true slime mold. *Biological Cybernetics* 102, 3 (2010), 261–269.
- Unver, O. and Sitti, M. Tankbot: A palm-size, tank like climbing robot on rough and smooth surfaces. *Int. J. of Robotics Research* 29, 14 (2010), 1761–1777.
- Van Breugel, F., Regan W. and Lipson, H. From insects to machines: a passively stable, untethered flapping-hovering micro air vehicle. *IEEE Robotics and Automation Magazine* 15, 4 (2008), 68–74.
- Webb, B., Consi, T.R. *Biorobotics*, MIT Press, Cambridge, MA, 2001.
- Wing, J. Computational thinking. *Commun. ACM* 49, 3 (Mar. 2006), 33–35.
- Wood, R.J. The first biologically inspired at-scale robotic insect. *IEEE Trans. on Robotics* 24, 2 (2008), 341–347.
- Wyeth, G. and Milford, M. Spatial cognition for robots. *IEEE Robotics Automation Magazine* 16, 3 (2009), 24–32.
- Zhou, C. and Low, K.-H. Better endurance and load capacity: An improved design of manta ray robot. *J. of Bionic Engineering* 7, 1 (2010), 137–144.
- Ziegler, M., Hoffmann, M., Carbajal, J.P. and Pfeifer, R. Varying body stiffness for aquatic locomotion. In *Proc. of Int. Conf. on Robotics and Automation* (2011), 2705–2712.

Rolf Pfeifer (pfeifer@ifi.uzh.ch) is a professor of computer science and director of the Artificial Intelligence Lab in the Department of Informatics at the University of Zurich, Switzerland.

Max Lungarella (lunga@ifi.uzh.ch) is a senior researcher in the Artificial Intelligence Lab at the University of Zurich and CTO of Dynamic Devices AG, Switzerland.

Fumiya Iida (fumiya.iida@mavt.ethz.ch) is an assistant professor and director of the Bio-Inspired Robotics Laboratory in the Department of Mechanical and Process Engineering at the Swiss Federal Institute of Technology, Zurich.

research highlights

P. 89

Technical Perspective **Open Platforms for Computational Photography**

By Richard Szeliski

P. 90

The Frankencamera: An Experimental Platform for Computational Photography

By Andrew Adams, David E. Jacobs, Jennifer Dolson, Marius Tico, Kari Pulli, Eino-Ville Talvala, Boris Ajdin, Daniel Vaquero, Hendrik P.A. Lensch, Mark Horowitz, Sung Hee Park, Natasha Gelfand, Jongmin Baek, Wojciech Matusik, and Marc Levoy

P. 99

Technical Perspective **The Realities of Home Broadband**

By Henning Schulzine

P. 100

Measuring Home Broadband Performance

By S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè

Technical Perspective

Open Platforms for Computational Photography

By Richard Szeliski

COMPUTATIONAL PHOTOGRAPHY IS AN emerging discipline that enables the creation of enhanced-quality photographs through novel combinations of digital images, algorithms, optics, and sensors.^{2,5} The field lies at the intersection of image processing, computer vision, and computer graphics, and has spawned its own workshops and conferences. It has also engendered many new features used in digital cameras and smartphones.

While scientists have applied image analysis and enhancement techniques to images for decades, the application of sophisticated algorithms to consumer photography started in the mid-1990s. Early examples of such algorithms include stitching multiple images into seamless panoramas, merging multiple exposures to create and display high dynamic range (HDR) images, and combining flash and no-flash images to provide better details in dark regions without harsh shadows.

As with most of computing, computational photography algorithms were originally developed and deployed on professional workstations and desktop personal computers. Unfortunately, the inability to deploy these algorithms inside cameras has severely limited real-world experimental validation and the percolation of these scientific advances into consumer products.

The migration of these algorithms into hardware and firmware has been hampered by a number of factors.¹ For example, digital image processing algorithms used by cameras are protected by patents and trade secrets. Vendors also tightly control the user experience, rather than taking the more open approach embraced by the app development community.

An even more fundamental impediment to the widespread development and deployment of in-camera algorithms is the lack of a clean open architecture for controlling camera features and writing the correspond-

ing real-time processing and viewing algorithms. The following article by Adams et al. is the first to address this problem, and it does so in a beautiful and elegant fashion.

The need for real-time processing and immediate feedback requires cameras to perform many different tasks in parallel. For example, cameras need to determine the optimal exposure time, aperture, analog gain, and focus settings for each picture.

Coming up with an elegant, programmable architecture and the APIs that support the deployment of sophisticated computational photography algorithms is a challenging architectural design problem. The authors show that in order to achieve this, the architecture must allow the specification of parameter sets (called shots) that control (or suggest) how individual images should be taken.

Because setting up these parameters can take time, the architecture keeps the desired and actual parameters tightly coupled with raw (unprocessed) images returned to the image processor. The complete architecture proposed in the paper therefore consists of shots (desired parameter sets), sensors that capture either individual, burst, or continuous streams of shots, frames that return the captured images and metadata, and devices such as lenses and flash units that can be controlled by the program.

To demonstrate the utility and generality of their approach, the authors built a custom-made experimental Frankencamera from commercial imaging parts and also reprogrammed an existing Nokia N900 smartphone. They then developed a collection of useful and compelling computational photography algorithms.

Since its original publication at SIGGRAPH 2010, the Frankencamera paper and associated hardware/firmware systems have had a dramatic impact on computational photogra-

phy research and teaching, as well as consumer-level photography devices. The Frankencamera devices and software have been used in the Stanford CS 448A course on Computational Photography⁴ as well as computational photography courses at other universities. Numerous computational photography apps can now be found for smartphones, and ideas inspired by the paper are also being incorporated into upcoming versions of smartphone operating systems and libraries.

One additional ingredient needed to make computational photography algorithms easy to develop is a high-level language and compiler tailored to such programs. Fortunately, a SIGGRAPH 2012 paper describing a system called *Halide* promises to do just that by enabling programmers to write high-level array-like descriptions of algorithms and then giving hints to the compiler about the desired levels of tile-based caching, parallelism, pipelining and reuse.³

Computational photography is blossoming as both a research field and a vibrant application area affecting all aspects of digital photography. The following paper provides an elegant example of how well-designed architectures in computer science can facilitate and accelerate the adoption of new technologies and expose novel capabilities to new generations of students. **C**

References

1. Levoy, M. Experimental platforms for computational photography. *IEEE Computer Graphics and Applications* 30, 5 (2010), 81–87.
2. Nayar, S. K. Computational cameras: Redefining the image. *Computer* 39, 8, (2006), 30–38.
3. Ragan-Kelley, J., Adams, A., Paris, S., Levoy, M., and Amarasinghe, and Durand, F. Decoupling algorithms from schedules for easy optimization of image processing pipelines. *ACM Transactions on Graphics* 31, 4 (2012).
4. Stanford CS 448A: Computational Photography; <http://graphics.stanford.edu/courses/cs448a-10/>.
5. Szeliski, R. *Computer Vision—Algorithms and Applications*. Springer, 2010.

Richard Szeliski (szeliski@microsoft.com) is a distinguished scientist at Microsoft Research, Redmond, WA.

© 2012 ACM 0001-0782/12/11 \$15.00

The Frankencamera: An Experimental Platform for Computational Photography

By Andrew Adams, David E. Jacobs, Jennifer Dolson, Marius Tico, Kari Pulli, Eino-Ville Talvala, Boris Ajdin, Daniel Vaquero, Hendrik P.A. Lensch, Mark Horowitz, Sung Hee Park, Natasha Gelfand, Jongmin Baek, Wojciech Matusik, and Marc Levoy

Abstract

Although there has been much interest in computational photography within the research and photography communities, progress has been hampered by the lack of a portable, programmable camera with sufficient image quality and computing power. To address this problem, we have designed and implemented an open architecture and application programming interface (API) for such cameras: the Frankencamera. It consists of a base hardware specification, a software stack based on Linux, and an API for C++. Our architecture permits control and synchronization of the sensor and image processing pipeline at the microsecond timescale, as well as the ability to incorporate and synchronize external hardware like lenses and flashes. This paper specifies our architecture and API, and it describes two reference implementations we have built. Using these implementations, we demonstrate several computational photography applications: high dynamic range (HDR) viewfinding and capture, automated acquisition of extended dynamic range panoramas, foveal imaging, and inertial measurement unit (IMU)-based hand shake detection. Our goal is to standardize the architecture and distribute Frankencameras to researchers and students, as a step toward creating a community of photographer-programmers who develop algorithms, applications, and hardware for computational cameras.

1. INTRODUCTION

Computational photography refers broadly to sensing strategies and algorithmic techniques that enhance or extend the capabilities of digital photography. Representative techniques include high dynamic range (HDR) imaging, flash/no-flash imaging, coded aperture and coded exposure imaging, panoramic stitching, digital photomontage, and light field imaging.¹⁸

Although interest in computational photography has steadily increased among graphics and vision researchers, few of these techniques have found their way into commercial cameras. One reason is that cameras are closed platforms. This makes it hard to incrementally deploy these techniques, or for researchers to test them in the field. Ensuring that these algorithms work robustly is therefore difficult, and so camera manufacturers are reluctant to add them to their products. For example, although HDR imaging has a long history,^{5,13} the literature has not addressed the question of automatically deciding

which exposures to capture, that is, metering for HDR. As another example, while many of the drawbacks of flash photography can be ameliorated using flash/no-flash imaging,^{7,15} these techniques produce visible artifacts in many photographic situations.⁶ Since these features do not exist in actual cameras, there is no strong incentive to address their artifacts.

Particularly frustrating is that even in platforms like smartphones, which encourage app creation and have increasingly capable imaging hardware, the programming interface to the imaging system is highly simplified, mimicking the physical interface of a point-and-shoot camera. This is a logical interface for the manufacturer to include, since it is complete for the purposes of basic camera operations and stable over many device generations. Unfortunately, it means that in these systems it is not possible to create imaging applications that experiment with most areas of computational photography.

To address this problem, we describe a camera architecture and application programming interface (API) flexible enough to implement most of the techniques proposed in the computational photography literature. We believe that the architecture is precise enough that implementations can be built and verified for it, yet high-level enough to allow for evolution of the underlying hardware and portability across camera platforms. Most importantly, we have found it easy to program for.

In the following section, we review previous work in this area, which motivates an enumeration of our design goals at the beginning of Section 3. We then describe our camera architecture in more detail. Our two reference implementations are shown in Figure 1. The first is the F2, which is composed of off-the-shelf components mounted in a laser-cut acrylic case. It is designed for extensibility. Our second platform is a Nokia N900 with a custom software stack. While less customizable than the F2, it is smaller, lighter, and readily available in large quantities. It demonstrates that current smartphones often have hardware components with more capabilities than their APIs expose. With these implementations in mind, we describe how to program for our architecture in Section 4. To demonstrate the capabilities of the architecture and API, we show several computational

The original version of this paper was published in *ACM Trans. Graph.* 29, 4 (2010).

Figure 1. Two implementations of the Frankencamera architecture: The custom-built F2 (left)—portable and self-powered, best for projects requiring flexible hardware; and the Nokia N900 (right) with a modified software stack—a compact commodity platform best for rapid development and deployment of applications to a large audience.



photography applications that cannot easily be implemented on current cameras (Section 5).

2. PRIOR WORK

A digital camera is a complex embedded system, spanning many fields of research. We limit our review of prior work to camera platforms rather than their constituent algorithms, to highlight why we believe a new architecture is needed to advance the field of computational photography.

2.1. Consumer cameras

Although improvements in the features of digital single-lens reflex cameras (DSLRs) have been largely incremental, point-and-shoot camera manufacturers are steadily expanding the range of features available on their cameras. Unfortunately, the camera software cannot be modified, and thus no additional features can be explored by the research community. Software development kits (SDKs) by manufacturers such as Canon and Nikon require tethering the camera to a computer, and provide no more control than the normal user interface.

Though the firmware in these cameras is always proprietary, several groups have successfully reverse-engineered the firmware for some Canon cameras. In particular, the Canon Hack Development Kit⁴ nondestructively replaces the original firmware on a wide range of Canon point-and-shoot cameras. Photographers can then script the camera, adding features such as custom burst modes, motion-triggered photography, and time-lapse photography. Similarly, the Magic Lantern project¹² provides enhanced firmware for Canon 5D Mark II DSLRs. While these projects remove both the need to attach a PC to the camera and the problem of latency, they yield roughly the same level of control as the manufacturer SDKs: the lower levels of the camera are still a black box.

2.2. Smartphones

Smartphones are programmable cell phones that allow and even encourage third-party applications. The newest smartphones are capable of capturing still photographs and videos with quality comparable to point-and-shoot cameras. These models contain numerous input and output devices (e.g., touchscreen, audio, buttons, GPS, compass, accelerometers), and are compact and portable. While these systems seem like an ideal platform for a computational camera, they provide limited interfaces to their camera subsystems. Neither Android nor Apple's iOS devices allow application control over absolute exposure time, or retrieval of raw sensor data—much less the ability to stream full-resolution images at the maximum rate permitted by the sensor. In fact, they typically provide less control of the camera than a DSLR SDK. This lack of control makes these devices useful for only a narrow range of computational photography applications. Despite these limitations, the iPhone app store has several hundred third-party applications that use the camera. This confirms our belief that there is a great interest in extending the capabilities of traditional cameras, an interest we hope to support and encourage with our architecture.

2.3. Smart cameras

Smart cameras are image sensors combined with local processing, storage, or networking, and are generally used as embedded computer vision systems.^{3, 22} These cameras provide fairly complete control over the imaging system, with the software stack implementing frame capture, low-level image processing, and vision algorithms such as background subtraction, object detection, or object recognition. Example research systems are the CMUcam,²⁰ Cyclops,¹⁶ MeshEye,⁸ and the Philips wireless smart camera motes.¹¹ Commercial systems include the National Instruments 17XX, Sony XCI-100, and the Basler eXcite series.

The main limitation of these systems is that they are not complete cameras. Most are tethered, few support synchronization with other I/O devices, and none contain a viewfinder or a shutter button. Augmenting these devices with a separate display complicates the system and introduces additional latency.

Our Frankencamera platforms attempt to provide everything needed for a practical computational camera: full access to the imaging system like a smart camera, a full user interface with viewfinder and I/O interfaces like a smartphone, and the ability to be taken outdoors, untethered, like a consumer camera.

3. THE FRANKENCAMERA ARCHITECTURE

Informed by our experiences programming for (and teaching with) smartphones, point-and-shoots, and DSLRs, we propose the following set of requirements for a Frankencamera:


1. Is handheld, self-powered, and untethered. This lets researchers take the camera outdoors and face real-world photographic problems.
2. Has a large viewfinder with a high-quality touch-

screen to enable experimentation with camera user interfaces.

3. Is easy to program. To that end, it should run a standard operating system, and be programmable using standard languages, libraries, compilers, and debugging tools.
4. Has the ability to manipulate sensor, lens, and camera settings on a per-frame basis at video rate, so we can request bursts of images with unique capture parameters for each image.
5. Labels each returned frame with the camera settings used for that frame, to allow for proper handling of the data produced by requirement 4.
6. Allows access to raw pixel values at the maximum speed permitted by the sensor interface. This means uncompressed, undemosaicked pixels.
7. Provides enough processing power in excess of what is required for basic camera operation to allow for the implementation of nearly any computational photography algorithm from the recent literature, and enough memory to store the inputs and outputs (often a burst of full-resolution images).
8. Allows standard camera accessories to be used, such as external flash or remote triggers, or more novel devices, such as GPS, inertial measurement units (IMUs), or experimental hardware. It should make synchronizing these devices to image capture straightforward.

Figure 2 illustrates our model of the imaging hardware in the Frankencamera architecture. It is general enough to cover most platforms so that it provides a stable interface to the application designer, yet precise enough to allow for the low-level control needed to achieve our requirements. It encompasses the image sensor, the fixed-function imaging

Figure 2. The Frankencamera abstract architecture. The architecture consists of an application processor, a set of photographic devices such as flashes or lenses, and one or more image sensors, each with a specialized image processor. A key aspect of this system is that image sensors are pipelined. The architecture does not dictate the number of stages; here we show a typical system with four frames in flight at a time.



pipeline that deals with the resulting image data, and other photographic devices such as the lens and flash.

3.1. The image sensor

One important characteristic of our architecture is that the image sensor is treated as stateless. Instead, it is a pipeline that transforms requests into frames. The requests specify the configuration of the hardware necessary to produce the desired frame. This includes sensor configuration like exposure and gain, imaging processor configuration like output resolution and format, and a list of device actions that should be synchronized to exposure, such as if and when the flash should fire.

The frames produced by the sensor are queued and retrieved asynchronously by the application. Each one includes both the actual configuration used in its capture, and also the request used to generate it. The two may differ when a request could not be achieved by the underlying hardware. Accurate labeling of returned frames (requirement 5) is essential for algorithms that use feedback loops like autofocus and metering.

As the manager of the imaging pipeline, a sensor has a somewhat privileged role in our architecture compared to other devices. Nevertheless, it is straightforward to express multiple-sensor systems. Each sensor has its own internal pipeline and abstract imaging processor (which may be implemented as separate hardware units, or a single time-shared unit). The pipelines can be synchronized or allowed to run independently. Simpler secondary sensors can alternatively be encapsulated as devices (described later), with their triggering encoded as an action slaved to the exposure of the main sensor.

3.2. The imaging processor

The imaging processor sits between the raw output of the sensor and the application processor, and has two roles. First, it generates useful statistics from the raw image data, including a small number of histograms over programmable regions of the image, and a low-resolution sharpness map to assist with autofocus. These statistics are attached to the corresponding returned frame.

Second, the imaging processor transforms image data into the format requested by the application, by demosaicking, white-balancing, resizing, and gamma correcting as needed. As a minimum we only require two formats: the raw sensor data (requirement 6) and a demosaicked format of the implementation's choosing. The demosaicked format must be suitable for streaming directly to the platform's display for use as a viewfinder.

The imaging processor performs both these roles in order to relieve the application processor of essential image processing tasks, allowing application processor time to be spent in the service of more interesting applications (requirement 7). Dedicated imaging processors are able to perform these roles at a fraction of the compute and energy cost of a more general application processor.

Indeed, imaging processors tend to be fixed-functionality for reasons of power efficiency, and so these two statistics and two output formats are the only ones we require in

our current architecture. We anticipate that in the longer term image processors will become more programmable, and we look forward to being able to replace these requirements with a programmable set of transformation and reduction stages. On such a platform, for example, one could write a “camera shader” to automatically extract and return feature points and descriptors with each frame to use for alignment, or structure-from-motion applications.

3.3. Devices

Cameras are much more than an image sensor. They also include a lens, a flash, and other assorted devices. In order to facilitate use of novel or experimental hardware, the requirements that the architecture places on devices are minimal.

Devices are controllable independently of a sensor pipeline by whatever means are appropriate to the device. However, in many applications the timing of device actions must be precisely coordinated with the image sensor to create a successful photograph. The timing of a flash firing in second-curtain sync mode must be accurate to within a millisecond. More demanding computational photography applications, such as coded exposure photography,¹⁷ require even tighter timing precision.

To this end, devices may also declare one or more actions that they can take synchronized to exposure. Programmers can then schedule these actions to occur at a given time within an exposure by attaching the action to a frame request. Devices declare the latency of each of their actions, and receive a callback at the scheduled time minus the latency. In this way, any event with a known latency can be accurately scheduled.

Devices may also tag returned frames with meta-data describing their state during that frame’s exposure (requirement 5). Tagging is done after frames leave the imaging processor, so this requires devices to keep a log of their recent state.

Some devices generate asynchronous events, such as when a photographer manually zooms a lens, or presses a shutter button. These are time-stamped and placed in an event queue, to be retrieved by the application at its convenience.

3.4. Discussion

While this pipelined architecture is simple, it expresses the key constraints of real camera systems, and it provides fairly complete access to the underlying hardware. Current camera APIs model the hardware in a way that mimics the physical camera interface: the camera is a stateful object, which makes blocking capture requests. This view only allows one active request at a time and reduces the throughput of a camera system to the reciprocal of its latency—a fraction of its peak throughput. Streaming modes, such as those used for electronic viewfinders, typically use a separate interface, and are mutually exclusive with precise frame level control of sensor settings, as camera state becomes ill-defined in a pipelined system. Using our pipelined model of a camera, we can implement our key architecture goals with a straightforward API.

4. PROGRAMMING THE FRANKENCAMERA

Developing for a Frankencamera is similar to developing for any Linux device. One writes standard C++ code, compiles it with a cross-compiler, and then copies the resulting binary to the device. Programs can then be run over ssh, or launched directly on the device’s screen. Standard debugging tools such as gdb and strace are available. To create a user interface, one can use any Linux UI toolkit. We typically use Qt and provide code examples written for Qt. OpenGL ES 2.0 is available for hardware-accelerated graphics, and regular POSIX calls can be used for networking, file I/O, synchronization primitives, and so on. If all this seems unsurprising, then that is precisely our aim.

Programmers and photographers interact with our architecture using the “FCam” API. We now describe the API’s basic concepts illustrated by example code.

4.1. Shots

The four basic concepts of the FCam API are *shots*, *sensors*, *frames*, and *devices*. We begin with the shot. A *shot* is a bundle of parameters that completely describes the capture and post-processing of a single output image. A shot specifies sensor parameters such as gain and exposure time (in microseconds). It specifies the desired output resolution, format (raw or demosaicked), and memory location into which to place the image data. It also specifies the configuration of the fixed-function statistics generators by specifying over which regions histograms should be computed and at what resolution a sharpness map should be generated. A shot also specifies the total time between this frame and the next. This must be at least as long as the exposure time and is used to specify frame rate independently of exposure time. Shots specify the set of actions to be taken by devices during their exposure (as a standard STL set). Finally, shots have unique ids auto-generated on construction, which assist in identifying returned frames.

The example code below configures a shot representing a VGA resolution frame, with a 10 ms exposure time, a frame time suitable for running at 30 frames per second, and a single histogram computed over the entire frame:

```
Shot shot;
shot.gain = 1.0;
shot.exposure = 10000;
shot.frameTime = 33333;
shot.image = Image (640, 480, UYVY);
shot.histogram.regions = 1;
shot.histogram.region[0] = Rect (0, 0, 640, 480);
```

4.2. Sensors

After creation, a *shot* can be passed to a *sensor* in one of the two ways—by *capturing* it or by *streaming* it. If a sensor is told to *capture* a configured shot (by calling `sensor.capture(shot)`), it pushes that shot into a request queue at the top of the imaging pipeline (Figure 2) and returns immediately.

The sensor manages the entire pipeline in the background. The shot is issued into the pipeline when it reaches the head of the request queue, and the sensor is ready to

begin configuring itself for the next frame. If the sensor is ready, but the request queue is empty, then a bubble necessarily enters the pipeline. The sensor cannot simply pause until a shot is available, because it has several other pipeline stages; there may be a frame currently exposing and another currently being read out. Bubbles configure the sensor to use the minimum frame time and exposure time, and the unwanted image data produced by bubbles is silently discarded.

Bubbles in the imaging pipeline represent wasted time and make it difficult to guarantee a constant frame rate for video applications. In these applications, the imaging pipeline must be kept full. To prevent this responsibility from falling on the API user, the sensor can also be told to *stream* a shot. A shot to be streamed is copied into a holding slot alongside the request queue. Then whenever the request queue is empty, and the sensor is ready for configuration, a copy of the contents of the holding slot enters the pipeline instead of a bubble. Streaming a shot is done using `sensor.stream(shot)`.

Sensors may also capture or stream vectors of shots, or *bursts*, in the same way that they capture or stream shots. Capturing a burst enqueues those shots at the top of the pipeline in the order given and is useful, for example, to capture a full high-dynamic-range stack in the minimum amount of time. As with a shot, streaming a burst causes the sensor to make an internal copy of that burst, and atomically enqueue all of its constituent shots at the top of the pipeline whenever the sensor is about to become idle. Thus, bursts are atomic—the API will never produce a partial or interrupted burst. The following code makes a burst from two copies of our shot, doubles the exposure of one of them, and then uses the sensor's stream method to create frames that alternate exposure on a per-frame basis at 30 frames per second. The ability to stream shots with varying parameters at video rate is vital for many computational photography applications, and hence was one of the key requirements of our architecture. It will be heavily exploited by our applications in Section 5.

```
std::vector<Shot> burst(2);
burst[0] = shot;
burst[1] = shot;
burst[1].exposure = burst[0].exposure*2;
sensor.stream(burst);
```

To update the parameters of a shot or burst that is currently streaming (e.g., to modify the exposure as the result of a metering algorithm), one merely modifies the shot or burst and calls *stream* again. Since the shot or burst in the internal holding slot is atomically replaced by the new call to *stream*, no partially updated burst or shot is ever issued into the imaging pipeline.

4.3. Frames

On the output side, the sensor produces *frames*, retrieved from a queue of pending frames via the `getFrame` method. This method is the only blocking call in the core API. A frame contains image data, the output of the statistics generators, the precise time at which the exposure

began and ended, the actual parameters used in its capture, and the requested parameters in the form of a copy of the shot used to generate it. If the sensor was unable to achieve the requested parameters (e.g., if the requested frame time was shorter than the requested exposure time), then the actual parameters will reflect the modification made by the system.

Frames can be identified by the `id` field of their shot. Being able to reliably identify frames is another of the key requirements for our architecture. The following code displays the longer exposure of the two frames specified in the burst above, but uses the shorter of the two to perform metering. The functions `displayImage` and `metering` are hypothetical functions that are not part of the API.

```
while (1) {
    Frame frame = sensor.getFrame();
    if (frame.shot().id == burst[1].id) {
        displayImage(frame.image);
    } else if (frame.shot().id == burst[0].id) {
        unsigned newExposure = metering(frame);
        burst[0].exposure = newExposure;
        burst[1].exposure = newExposure*2;
        sensor.stream(burst);
    }
}
```

In simple programs, it is typically not necessary to check the ids of returned frames, because our API guarantees that exactly one frame comes out per shot requested, in the same order. Frames are never duplicated or dropped entirely. If image data is lost or corrupted due to hardware error, a frame is still returned (possibly with statistics intact), with its image data marked as invalid.

4.4. Devices

In our API, each device is represented by an object with methods for performing its various functions. Each device may additionally define a set of *actions*, which are used to synchronize these functions to exposure, and a set of *tags* representing the metadata attached to returned frames. While the exact list of devices is platform-specific, the API includes abstract base classes that specify the interfaces to the lens and the flash.

The *lens* can be directly asked to initiate a change to any of its three parameters: focus (measured in diopters), focal length, and aperture, with the methods `setFocus`, `setZoom`, and `setAperture`. These calls return immediately, and the lens starts moving in the background. For cases in which lens movement should be synchronized to exposure, the lens defines three actions to do the same. Each call has an optional second argument that specifies the speed with which the change should occur. Additionally, each parameter can be queried to see if it is currently changing, what its bounds are, and its current value. The following code moves the lens from its current position to infinity focus over the course of 2s.

```
Lens lens;
float speed = (lens.getFocus() - lens.
farFocus()) / 2;
lens.setFocus(lens.farFocus(), speed);
```

A lens tags each returned frame with the state of each of its three parameters during that frame. Tags can be retrieved from a frame like so:

```
Frame frame = sensor.getFrame();
Lens :: Tags *tags = frame->tags(&lens);
cout << "The lens was at: " << tags->focus;
```

The *flash* has a single method that tells it to fire with a specified brightness and duration, and a single action that does the same. It also has methods to query bounds on brightness and duration. Flashes with more capabilities (such as the strobing flash in Figure 3) can be implemented as subclasses of the base flash class. The flash tags each returned frame with its state, indicating whether it fired during that frame, and if so with what parameters.

The following code example adds an action to our shot to fire the flash briefly at the end of the exposure (second-curtain sync). The results of a similar code snippet run on the F2 can be seen in Figure 3.

```
Flash flash;
Flash :: FireAction fire(&flash);
fire.brightness = flash.maxBrightness();
fire.duration = 5000;
fire.time = shot.exposure - fire.duration;
shot.actions.insert(&fire);
```

Other devices can be straightforwardly incorporated into the API, allowing easy management of the timing of their

Figure 3. The Frankencamera API provides precise timing control of secondary devices like the flash. Here, two Canon flash units were mounted on an F2, one set to strobe and one to fire once at end of the exposure.



actions. One merely needs to inherit from the Device base class, add methods to control the device in question, and then define any appropriate actions, tags, and events. This flexibility is critical for computational photography, in which it is common to experiment with novel hardware that affects image capture.

4.5. Implementation

In our current API implementations, apart from fixed-function image processing, FCam runs entirely on the ARM CPU in the OMAP3430, using a small collection of user-space threads and modified Linux kernel modules. Our system is built on top of Video for Linux 2 (V4L2)—the standard Linux kernel video API. V4L2 treats the sensor as stateful with no guarantees about timing of parameter changes. To provide the illusion of a stateless sensor processing stateful *shots*, we use several real-time-priority threads to manage updates to image sensor parameters, readback of image data and metadata, and device actions synchronized to exposure.

Our image sensor drivers are standard V4L2 sensor drivers with one important addition. We add controls to specify the time taken by each individual frame, which are implemented by adjusting the amount of extra vertical blanking in sensor readout.

4.6. Discussion

Our goals for the API were to provide intuitive mechanisms to precisely manipulate camera hardware state over time, including control of the sensor, fixed-function processing, lens, flash, and any associated devices. We have accomplished this in a minimally surprising manner, which should be a key design goal of any API. The API is limited in scope to what it does well, so that programmers can continue to use their favorite image processing library, UI toolkit, file I/O, and so on. Nonetheless, we have taken a “batteries included” approach, and made available control algorithms for metering and focus, image processing functions to create raw and JPEG files, and example applications that demonstrate using our API with the Qt UI toolkit and OpenGL ES.

Implementing the API on our two platforms required a shadow pipeline of in-flight shots, managed by a collection of threads, to fulfill our architecture specification. This makes our implementation brittle in two respects. First, an accurate timing model of image sensor and imaging processor operation is required to correctly associate output frames with the shot that generated them. Second, deterministic guarantees from the image sensor about the latency of parameter changes are required, so that we can configure the sensor correctly. In practice, there is a narrow time window in each frame during which sensor settings may be adjusted safely. To allow us to implement our API more robustly, future image sensors should provide a means to identify every frame they produce on both the input and output sides. Setting changes could then be requested to take effect for a named future frame. This would substantially reduce the timing requirements on sensor configuration. Image sensors could then return images tagged with their frame id (or even the entire sensor state), to make association of image data with sensor state trivial.

5. APPLICATIONS

We now describe a number of applications of the Frankencamera architecture and API to concrete problems in photography. Most run on either the N900 or the F2, though some require hardware specific to one platform or the other. These applications are representative of the types of in-camera computational photography our architecture enables, and several are also novel applications in their own right. They are all either difficult or impossible to implement on existing platforms, yet simple to implement under the Frankencamera architecture.

5.1. IMU-based lucky imaging

Long-exposure photos taken without use of a tripod are usually blurry, due to natural hand shake. However, hand shake varies over time, and a photographer can get “lucky” and record a sharp photo if the exposure occurs during a period of stillness (Figure 4). Our “Lucky Imaging” application uses an experimental Nokia three-axis gyroscope affixed to the front of the N900 to detect hand shake. Utilizing a gyroscope to determine hand shake is computationally cheaper than analyzing full resolution image data, and will not confuse blur caused by object motion in the scene with blur caused by hand shake. We use an external gyroscope because the internal accelerometer in the N900 is not sufficiently accurate for this task.

To use the gyroscope with the FCam API, we created a device subclass representing a three-axis gyroscope. The gyroscope object then tags frames with the IMU measurements recorded during the image exposure. The application streams full-resolution raw frames, saving them to storage only when their gyroscope tags indicate low motion during the frame in question. The ease with which this external device could be incorporated is one of the key strengths of our architecture.

This technique can be extended to longer exposure times where capturing a “lucky image” on its own becomes very

unlikely. Indeed, Joshi et al.⁹ show how to deblur the captured images using the motion path (as recorded by the IMU) as a prior.

5.2. Foveal imaging

CMOS image sensors are typically bandwidth-limited devices that can expose pixels faster than they can be read out into memory. Full-sensor-resolution images can only be read out at a limited frame rate: roughly 12 fps on our platforms. Low-resolution images, produced by downsampling or cropping on the sensor, can be read at a higher-rate: up to 90 fps on the F2. Given that we have a limited pixel budget, it makes sense to only capture those pixels that are useful measurements of the scene. In particular, image regions that are out-of-focus or oversaturated can safely be recorded at low spatial resolution, and image regions that do not change over time can safely be recorded at low temporal resolution.

Foveal imaging uses a streaming burst, containing shots that alternate between downsampling and cropping on the sensor. The downsampled view provides a 640×480 view of the entire scene, and the cropped view provides a 640×480 inset of one portion of the scene, analogously to the human fovea (Figure 5). The fovea can be placed on the center of the scene, moved around at random in order to capture texture samples, or programmed to preferentially sample sharp, moving, or well-exposed regions. For now, we have focused on acquiring the data, and present results produced by moving the fovea along a prescribed path. In the future, we intend to use this data to synthesize full-resolution high-framerate video, similar to the work of Bhat et al.²

Downsampling and cropping on the sensor is a capability of the Aptina sensor in the F2 not exposed by the base API. To access this, we use derived versions of the `Sensor`, `Shot`, and `Frame` classes specific to the F2 API implementation. These extensions live in a sub-namespace of the FCam API. In general, this is how FCam handles platform-specific extensions.

Figure 4. Lucky Imaging. An image stream and three-axis gyroscope data for a burst of three images with 0.5s exposure times. The FCam API synchronizes the image and motion data, and only the frames determined to have low motion are saved to storage.





Figure 5. Foveal imaging records a video stream that alternates between a downsampled view of the whole scene and full-detail insets of a small region of interest. In this example, the inset is set to scan over the scene, the region of interest moving slightly between each pair of inset frames.



5.3. HDR viewfinding and capture

HDR photography operates by taking several photographs and merging them into a single image that better captures the range of intensities of the scene.¹⁹ While modern cameras include a “bracket mode” for taking a set of photos separated by a preset number of stops, they do not include a complete “HDR mode” that provides automatic metering, viewfinding, and compositing of HDR shots. We use the FCam API to implement such an application on the F2 and N900 platforms.


HDR metering and viewfinding is done by *streaming* a burst of three 640×480 shots, whose exposure times are adjusted based on the scene content, in a manner similar to Kang et al.¹⁰ The HDR metering algorithm sets the long-exposure frame to capture the shadows, the short exposure to capture the highlights, and the middle exposure as the midpoint of the two. As the burst is streamed by the sensor, the three most recently captured images are merged into an HDR image, globally tone-mapped with a gamma curve, and displayed in the viewfinder in real time. This allows the photographer to view the full dynamic range that will be recorded in the final capture, assisting in composing the photograph.

Once it is composed, a high-quality HDR image is captured by creating a burst of three full-resolution shots, with exposure and gain parameters copied from the viewfinder burst. The shots are *captured* by the sensor, and the resulting frames are aligned and then merged into a final image using the *Exposure Fusion* algorithm.¹⁴ Figure 6 shows the captured images and results produced by our N900 implementation.

5.4. Panorama capture

The field of view of a regular camera can be extended by capturing several overlapping images of a scene and stitching them into a single panoramic image. However, the process of capturing individual images is time-consuming and prone to errors, as the photographer needs to ensure that all areas of the scene are covered. This is difficult since panoramas are traditionally stitched

Figure 6. HDR imaging. The high-speed capture capabilities of FCam allow capturing a burst of frames for handheld HDR with minimal scene motion. The final composite produced on-device is on the right.



off-camera, so that no on-line preview of this capture process is available.

In order to address these issues, we implemented an application for capturing and generating panoramas using the FCam API on the N900. In the capture interface, the viewfinder alignment algorithm¹ tracks the position of the current viewfinder frame with respect to the previously captured images, and a new high-resolution image is automatically captured when the camera points to an area that contains enough new scene content. A map showing the relative positions of the previously captured images and the current camera pose guides the user in moving the camera (top left of Figure 7). Once the user has covered the desired field of view, the images are stitched into a panorama in-camera, and the result can be viewed for immediate assessment.

In addition to in-camera stitching, we can use the FCam API’s ability to individually set the exposure time for each shot to create a panorama with extended dynamic range, in the manner of Wilburn et al.²¹ In this mode, the exposure time of the captured frames alternates between short and long, and the amount of overlap between successive frames is increased, so that each region of the scene is imaged by at least one short-exposure frame and at least one long-exposure frame. In the stitching phase, the long and

Figure 7. Extended dynamic range panorama capture. A Frankencamera platform allows for experimentation with novel capture interfaces and camera modes. Here we show a semiautomated panorama capture program, which produces an extended dynamic range panorama.



short exposure panoramas are generated separately, then combined¹⁴ to create an extended dynamic range result.

6. CONCLUSION

We have described the Frankencamera—a camera architecture suitable for experimentation in computational photography, and two implementations: our custom-built F2, and a Nokia N900 running the Frankencamera software stack. Our architecture includes an API that encapsulates camera state in the shots and frames that flow through the imaging pipeline, rather than in the photographic devices that make up the camera. By doing so, we unlock the underexploited potential of commonly available imaging hardware. The applications we have explored thus far are low-level photographic ones. With this platform, we now plan to explore applications in augmented reality, camera user interfaces, and augmenting photography using online services and photo galleries.

The central goal of this project is to enable research in computational photography. We are therefore distributing our platforms to students in computational photography courses, and are eager to see what will emerge. In the longer term, our hope is that consumer cameras and devices will become programmable along the lines of what we have described, enabling exciting new research and creating a vibrant community of programmer-photographers.

Acknowledgments

For this work, A. Adams was supported by a Reed-Hodgson Stanford Graduate Fellowship; E.-V. Talvala was supported by a Kodak Fellowship. S.H. Park and J. Baek acknowledge support by Nokia. D. Jacobs received support from a Hewlett Packard Fellowship, and J. Dolson received support from an NDSeg Graduate Fellowship from the United States Department of Defense. D. Vaquero was an intern at Nokia during this work. This work was partially done while W. Matusik was a Senior Research Scientist and B. Ajdin was an intern at Adobe Systems, Inc., and we thank David Salesin and the Advanced Technology Labs for support and feedback. Finally, M. Levoy acknowledges support from the National Science Foundation under award 0540872. 

References

- Adams, A., Gelfand, N., and Pulli, K. Viewfinder alignment. 2 (2008), 597–606.
- Bhat, P., Zitnick, C.L., Snavely, N., Agarwala, A., Agrawala, M., Cohen, M., Curless, B., Kang, S.B., Using photographs to enhance videos of a static scene. In (2007).
- Bramberger, M., Doblander, A., Maier, A., Rinner, B., Schwabach, H. Distributed embedded smart cameras for surveillance applications. , 2 (2006), 68–75.
- The CHDK Project, 2010.
- Debevec, P.E., Malik, J. Recovering high dynamic range radiance maps from photographs. In (New York, NY, USA, 1997), ACM Press/Addison-Wesley Publishing Co, 3 69–378.
- Durand, F. private communication, 2009.
- Eisemann, E., Durand, F. Flash photography enhancement via intrinsic relighting. , 3 (2004), 673–678.
- Hengstler, S., Prashanth, D., Fong, S., Aghajan, H. Mesheye: a hybrid-resolution smart camera mote for applications in distributed intelligent surveillance. In , 360–369.
- Joshi, N., Kang, S.B., Zitnick, C.L., Szeliski, R. Image deblurring using inertial measurement sensors. , 3 (Aug. 2010).
- Kang, S.B., Uyttendaele, M., Winder, S., Szeliski, R. High dynamic range video. In (2003), ACM, New York, NY, 319–325.
- Kleihorst, R., Schueler, B., Danilin, A., Heijligers, M. Smart camera mote with high performance vision system. In ACM SenSys 2006 Workshop on Distributed Smart Cameras (DSC 2006) (Oct. 2006).
- The Magic Lantern project, 2010.
- Mann, S., Picard, R.W. On being 'undigital' with digital cameras: extending dynamic range by combining differently exposed pictures. In (1995), 442–448.
- Mertens, T., Kautz, J., Reeth, F.V. Exposure fusion. In (2007).
- Petschnigg, G., Szeliski, R., Agrawala, M., Cohen, M., Hoppe, H., Toyama, K. Digital photography with flash and no-flash image pairs. In (2004), ACM, New York, NY, 664–672.
- Rahimi, M., Baer, R., Iroez, O., Garcia, J.C., Warrior, J., Estrin, D., Srivastava, M. Cyclops: in situ image sensing and interpretation in wireless sensor networks. In (2005), 192–204.
- Raskar, R., Agrawal, A., Tumblin, J. Coded exposure photography: motion deblurring using fluttered shutter. In (2006), ACM, New York, NY, 795–804.
- Raskar, R., Tumblin, J. Computational Photography: Mastering New Techniques for Lenses, Lighting, and Sensors. A K Peters, Natick, MA, 2010, in press.
- Reinhard, E., Ward, G., Pattanaik, S., Debevec, P. High Dynamic Range Imaging - Acquisition, Display and Image-based Lighting. Morgan Kaufman Publishers, San Francisco, CA, 2006.
- Rowe, A., Goode, A., Goel, D., Nourbakhsh, I. CMUcam3: An Open Programmable Embedded Vision Sensor. Technical Report RI-TR-07-13, Carnegie Mellon Robotics Institute, May 2007.
- Wilburn, B., Joshi, N., Vaish, V., Talvala, E.V., Antunez, E., Barth, A., Adams, A., Horowitz, M., Levoy, M. High performance imaging using large camera arrays. In (2005), ACM, New York, NY, 765–776.
- Wolf, W., Ozer, B., Lv, T. Smart cameras as embedded systems. (2002), 48–53.

Andrew Adams (abadams@csail.mit.edu),
Courier CSAIL, MIT.

David E. Jacobs, Mark Horowitz, Sung Hee Park, Jongmin Baek, Marc Levoy ([dejacobs, horowitz, shpark7, jbaek, levoy]@cs.stanford.edu), Stanford University.

Jennifer Dolson (jen.dolson@gmail.com),
Stanford University.

Marius Tico (mariustico@gmail.com),
Nokia Research Center.

Kari Pulli (karip@nvidia.com), NVIDIA
Research.

Eino-Ville Talvala (etalvala@google.com),
Stanford University (currently at Google Inc.).

Boris Ajdin, Hendrik P.A. Lensch ([boris.ajdin, hendrik.lensch]@uni-ulm.de),
Tübingen University.

Daniel Vaquero (daniel.vaquero@gmail.com),
University of California Santa Barbara.

Natasha Gelfand (ngelfand@gmail.com),
University of Labrador.

Wojciech Matusik (wojciech@csail.mit.edu), MIT CSAIL.

Technical Perspective

The Realities of Home Broadband

By Henning Schulzrinne

COMPARED TO OTHER consumer purchases, buying residential broadband services seems relatively simple: pick among a small number of plans, characterized by a one-speed number or maybe two, and then compare similar plans by price. (Assuming, of course, that there is a choice in plans and providers—a topic for another day...) If the Internet service seems slow, just upgrade to the next tier. Unfortunately, as the authors illustrate in this important contribution, reality is somewhat more complicated, with different technologies and providers delivering more or less of the promised “headline” speed, different packet latencies, and varying performance predictability.

Regulators have recognized that consumer choice requires more than the presence of multiple competing providers. Consumers also must have the ability to make informed choices. This is particularly important for residential broadband Internet access, as the expenditures are a non-trivial part of the family budget and it is often difficult and time-consuming to switch providers, with long-term contracts, installation fees, and, in some cases, having ISP technicians drill holes through home walls. In the U.S., this insight has partially motivated the FCC Open Internet regulatory proceeding, which has transparency as one of its key tenets.¹ The Open Internet order mandates the pre-sale disclosure of key performance metrics to allow consumers to pick the appropriate plan.

The data and report related to the performance measurements reported here also developed out of an FCC effort, the *Measuring Broadband America* program. This program had a rather interesting effect on the marketplace: An ISP whose measured performance met and exceeded the promised throughput ran prime-time TV commercials touting the findings compared to a competitor. A year later, the FCC published a follow-up report,² using the same basic

methodology, and two of the underperforming companies mentioned in the paper had significantly improved their actual performance compared to the advertised rates, illustrating the old “you manage what you measure” adage.

The work presented here can also help provide a quantitative foundation for two long-running, related discussions in the network research and policy realm: First, for about 20 years, a favorite panel and conference dinner topic has been whether differentiated or guaranteed quality-of-service (QoS) classes are necessary for the new application of the day, whether voice calls or video streaming, rather than offering only one “best effort” service. So far, ISPs have been offering only best effort as a residential Internet access service; however, they are also starting to use the same IP-based delivery for various specialized, or managed services, such as voice calling or IPTV, that is, the delivery of video content over IP. Independent observers must be able to gauge whether so-called over-the-top services that can only use best-effort services can indeed be seen as competing on an equal footing with these ISP-provided services.

The following paper provides an excellent example where open data, created with the resources and coordination that typically only a government entity can provide, combined with deep data analysis and additional experiments and measurements can yield insights into a topic of interest to policymakers, researchers, and consumers.

The FCC Measuring Broadband America program was founded on principles of openness and transparency, and the raw FCC-gathered measurement data is freely available to other researchers, so that others can both replicate the results reported and investigate other aspects. For example, the data contains several other active measurements that try to predict VoIP and video streaming performance.

The authors point out that much re-

mains to be done to improve our understanding of broadband services. They note that components such as buffers in home routers and WiFi networks as well as protocol effects can reduce the end-to-end performance. As raw access speeds increase from a few Mb/s to possibly a Gb/s, these impairments may well become dominant and will, among other problems, lead to disappointed consumers who find that paying more for a better service does not actually yield higher-quality video or faster-loading Web pages. The authors hint at the problem of network reliability. Even where raw speed is sufficient to meet home needs, the ability to consistently make phone calls that are equivalent to old “landline” quality or to work from home without worrying about losing connectivity during the day will depend upon reliability, and make for challenging future measurement projects. Also, wireless services, which are beyond the scope of the paper, raise many new challenges. The FCC and other national regulatory bodies are starting to gather data for those networks.

In the long run, consumers should not have to worry about the technical details of their Internet service, just like they do not worry (in most developed countries, at least) about whether their electricity service has enough amperage to power their newest gadget. Until that time, the work presented here will help consumers choose, policymakers protect consumers, and providers improve their services. 

References

1. FCC. *In the Matter of Preserving the Open Internet Broadband Industry Practices*. Report and Order, FCC 10-201, Dec. 2010; http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1_Rcd.pdf
2. FCC. *A Report on Consumer Wireline Broadband Performance in the U.S.* July 2012; <http://www.fcc.gov/measuring-broadband-america/2012/july>.

Henning Schulzrinne is a professor of computer science at Columbia University and currently serves as Chief Technology Officer of the FCC. The opinions are those of the author and do not necessarily reflect the view of the FCC, other commission staff members, or any commissioner.

© 2012 ACM 0001-0782/12/11 \$15.00

Measuring Home Broadband Performance

By S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè

Abstract

We present the results from the first study of Internet access link performance measured directly from home routers. In conjunction with the Federal Communication Commission's study of broadband Internet access in the United States, we investigate the throughput and latency of network access links from about 4000 routers across eight ISPs. Our findings provide a snapshot of access network performance across the United States, offer insights on how access network performance should be measured and presented to users, and inform various ongoing efforts to evaluate the performance of access networks around the world.

1. INTRODUCTION

Of nearly two billion Internet users worldwide, about 500 million are residential broadband subscribers.¹⁰ Broadband penetration will continue to increase, with people relying on home connectivity for day-to-day and even critical activities. Accordingly, the Federal Communication Commission (FCC) is developing performance-testing metrics for access providers.^{4, 9, 22} Policymakers, home users, and Internet Service Providers (ISPs) need better ways to benchmark broadband Internet performance.

Unfortunately, benchmarking home Internet performance is not as simple as running one-time “speed tests.” There exist countless tools to measure Internet performance.^{5, 16, 17, 20} Previous work has studied the typical download and upload rates of home access networks^{7, 14}; others have found that modems often have large buffers,¹⁴ and that DSL links often have high latency.¹⁵ These studies have shed some light on access link performance, but they have typically run measurements either from an end-host inside the home (from the “inside out”) or from a server on the wide-area Internet (from the “outside in”). Because these tools run from end-hosts, they cannot analyze the effects of confounding factors such as home network cross-traffic, the wireless network, or end-host configuration. Also, many of these tools run as one-time measurements and, without continual measurements of the same access link, we cannot establish a baseline performance level or observe how performance varies over time.

We measure broadband Internet performance directly from the router that is connected to the user's ISP. Measuring the access link from the home router offers several advantages over conventional methods. First, the home router is typically always on. Second, because it connects the home network to the ISP's network (as shown in Figure 1), taking measurements from this vantage point allows us to control the effects of many confounding factors, such as the home

wireless network and load on the measurement host. Table 1 summarizes the challenges involved in measuring access ISP performance, and how performing measurements from the home router addresses each of them.

We collected data from two complementary deployments. First, we gathered data from routers in over 4200 homes across the United States and eight different ISPs from a deployment sponsored by the Federal Communications Commission and operated by SamKnows. Second, we collected data from 16 homes in the BISmark deployment, spanning three ISPs in Atlanta. The SamKnows deployment provides a large user base, as well as diversity in ISPs, service plans, and geographical locations. We can access BISmark routers remotely and run repeated experiments to investigate the effect of factors that we could not study in a larger deployment. For example, to study the effect of modem choice on performance, we installed different modems in the same home and conducted experiments in the same controlled setting. Both deployments run a comprehensive suite of measurement tools that periodically measure throughput, latency, packet loss, and jitter. We use active measurement data from both deployments from December 14, 2010 to January 14, 2011. Table 2 lists the ISPs that we study and the number of measured access links for each of them.

Figure 1. The home router sits directly behind the modem in the home network. It takes measurements both to the last mile router (first non-NAT IP hop on the path) and to wide area hosts.




Table 1. Confounding factors and how we address them.

Factor	How we address it
Wireless effects	Use a wired connection to the modem
Cross-traffic	Measure cross-traffic and avoid it/account for it
Router load	Use a well-provisioned router
Server location	Choose a nearby server
End-to-end path	Focus on characterizing the last mile
Router configuration	Test configuration in practice and controlled settings

A previous version of this paper was published in the *Proceedings of SIGCOMM '11* (Toronto, Ontario, Canada, Aug. 15–19, 2011).

We characterize access network throughput (Section 3) and latency (Section 4) from the SamKnows and BISmark deployments. We explain how our throughput measurements differ from common “speed tests” and also propose several metrics that capture different aspects of latency. When our measurements cannot fully explain the observed behavior, we model the access link and test our hypotheses using controlled experiments. We find that the most significant factors affecting throughput are the access technology, ISPs’ traffic shaping policies, and congestion during peak hours. On the other hand, the quality of the access link, modem buffering, and cross-traffic within the home affect latency the most.

This study offers insights into both access network performance and measurement methods for benchmarking home broadband performance. Our findings include the following:

- Access link throughput is more variable during peak hours.
- ISPs use different policies and traffic shaping behavior that can make it difficult to compare measurements across them.
- Home network equipment and infrastructure can affect performance. For example, buffering in a user’s modem varies across models and can affect the latency that a user experiences during an upload. We found “bufferbloat,” or excessive buffering that adversely affects network performance, in many DSL modems.
- There is no “best” ISP for all users. Some ISPs have better short-term throughput, while others may have better sustained throughput, lower latency, or generally more consistent performance. Different users may prefer different ISPs depending on their usage profiles and how those ISPs perform along performance dimensions that matter to them.

As the first in-depth analysis of home access network performance, our study offers insights for users, ISPs, and policymakers. Users and ISPs can better understand the performance of the access link as measured directly from the router; ultimately, such a deployment could help an ISP differentiate performance problems within the home from those on the access link. Our study also informs policy by illustrating that a diverse set of network metrics ultimately affect the performance that a user experiences. The need

for a benchmark is clear, and the results from this study can serve as a principled foundation for such an effort.

2. ACCESS TECHNOLOGIES

We describe the two most common access technologies from our deployments: Digital Subscriber Line (DSL) and cable. A few users in our deployments have Fiber-To-The-Node (FTTN), Fiber-To-The-Premises (FTTP), and WiMax, but we do not have enough users to analyze these technologies.


DSL networks use telephone lines; subscribers have dedicated lines between their own DSL modems and the closest DSL Access Multiplexer (DSLAM). The DSLAM multiplexes data between the access modems and upstream networks, as shown in Figure 2a. The most common type of DSL access is asymmetric (ADSL), which provides different upload and download rates. In cable access networks, groups of users send data over a shared medium (typically coaxial cable); at a regional *headend*, a Cable Modem Termination System (CMTS) receives these signals and converts them to Ethernet, as shown in Figure 2b. The physical connection between a customer’s home and the DSLAM or the CMTS is often referred to as the *local loop* or *last mile*. Users buy a service plan from a provider that typically offers some *maximum* capacity in both the upload and download directions.

ADSL capacity. The ITU-T standardization body establishes that the achievable rate for ADSL 1¹¹ is 12 Mbits/s downstream and 1.8 Mbits/s upstream. The ADSL2+ specification¹² extends the capacity of ADSL links to at most 24 Mbits/s downstream and 3.5 Mbits/s upstream. Although the ADSL technology can theoretically achieve these speeds, many factors limit the capacity in practice. An ADSL modem negotiates the operational rate with the DSLAM (often called the *sync rate*); this rate depends on the quality of the local loop, which is mainly determined by the distance to the DSLAM from the user’s home and noise on the line. The maximum IP link capacity is lower than the sync rate because of the overhead of underlying protocols. The best service plan that

Table 2. The SamKnows and BISmark deployments. Active deployments are those that report more than 100 download throughput measurements over the course of our study.

ISP	Technology	SamKnows		BISmark
		Total	Active	Total
Comcast	Cable	864	560	4
AT&T	DSL/FTTN	787	335	10
TimeWarner	Cable	690	381	–
Verizon	DSL/FTTP	551	256	–
Cox	Cable	381	161	–
Qwest	DSL/FTTN	265	117	–
Charter	Cable	187	51	–
Cablevision	Cable	104	53	–

Figure 2. Access network architectures (a) DSL and (b) cable.



an ADSL provider advertises usually represents the rate that customers can achieve if they have a good connection to the DSLAM. Providers also offer service plans with lower rates and can rate-limit a customer's traffic at the DSLAM.

Modem configuration can also affect performance. ADSL users or providers configure their modems to operate in either *fastpath* or *interleaved* mode. In *fastpath* mode, data is exchanged between the DSL modem and the DSLAM in the same order that it is received, which minimizes latency but prevents error correction from being applied across frames. Thus, ISPs typically configure fastpath only if the line has a low bit error rate. *Interleaving* increases robustness to line noise at the cost of increased latency by splitting data from each frame into multiple segments and interleaving those segments with one another before transmitting them.

Cable capacity. In cable networks, the most widely deployed version of the standard is Data Over Cable Service Interface Specification version 2 (DOCSIS 2.0),¹³ which specifies download rates up to 42.88 Mbits/s and upload rates up to 30.72 Mbits/s in the United States. The latest standard, DOCSIS 3.0, allows for hundreds of megabits per second by bundling multiple channels. Cable providers often offer service plans with lower rates. An operator configures the service plan rate limit at the cable modem using a token bucket rate shaper. Many cable providers offer *PowerBoost*, which allows users to download (and, in some cases, upload) data at higher rates for an initial part of a transfer. The actual rate that a cable user receives will vary with the network utilization of other users connecting to the same headend.

3. UNDERSTANDING THROUGHPUT

We first explore how different techniques for measuring throughput can generate different results and offer guidelines on how to interpret them. We then investigate the throughput users achieve on different access links. Finally, we explore the effects of ISP traffic shaping and the implications it has for throughput measurement.

3.1. How to measure and interpret throughput

Users are often interested in the throughput that they receive on uploads or downloads, yet the notion of “throughput” can vary depending on how, when, and who is measuring it. For example, a run of www.speedtest.net in an author's home, where the service plan was 6 Mbits/s down and 512 Kbits/s up, reported a downlink speed of 4.4 Mbits/s and an uplink speed of 140 Kbits/s. Netalyzr¹⁴ reported 4.8 Mbits/s and 430 Kbits/s. Long-term measurements from the SamKnows router paint a different picture: the user achieves 5.6 Mbits/s down and 452 Kbits/s up. Both www.speedtest.net and Netalyzr measurements reflect transient network conditions, as well as other confounding factors.

There is no standard way to measure throughput. Bauer et al. list several notions of “broadband speed”: *capacity* is the total carrying capacity of the link; and the *bulk transfer capacity* is the amount of data that can be transferred along a path with a congestion-aware protocol like TCP.³ The SamKnows routers measure bulk transfer capacity using

three parallel HTTP transfers; this approach increases the likelihood of saturating the access link. The software first executes a “warm-up” transfer until throughput is steady to ensure that the throughput measurements are not affected by TCP slow start. The following download tests use the same TCP connection to exploit the “warmed up” session. The tests last for about 30s; the software reports snapshots of how many bytes were transferred for every 5-s interval. BISmark measures throughput by performing an HTTP download and upload for 15s using a single-threaded TCP connection once every 30min, regardless of cross-traffic. To account for cross-traffic, the router counts bytes transferred by reading from `/proc/net/dev` and computes the “passive throughput” as the byte count after the HTTP transfer minus the byte count before the transfer, divided by the transfer time. Table 3 summarizes the throughput measurements collected by the two deployments. Although measuring throughput may seem straightforward, our results demonstrate the extent to which different measurement methods can produce different results and, hence, may result in different conclusions about an ISP's performance.

Throughput measurement techniques—even commonly accepted ones—can yield variable results. We compare throughput measurement techniques in two locations that have both the SamKnows and BISmark routers (we use only two locations due to the logistical difficulty of deploying both routers in the same location). In both cases, the ISP is AT&T, but the service plans are different (6 Mbits/s down and 512 Kbits/s up; and 3 Mbits/s down and 384 Kbits/s up). We normalize the throughput with the advertised service plan so that we can compare the service plans. Figure 3 shows a CDF of the normalized throughput reported by the four methods from Table 3. Each data point represents a single throughput measurement. A value of 1.0 on the *x*-axis indicates that the throughput matches the ISP's advertised rate; no method achieves that value. This shortfall could be caused by many factors, including the sync rate of the modem to the DSLAM, layer-2 framing overhead on the line, and overhead from the measurement

Table 3. SamKnows and BISmark throughput measurements.

Parameter	Type	Prot.	Freq.	Comments
SamKnows				
Downstream throughput	Multi-threaded HTTP	TCP	2h	MLab, idle link
Upstream throughput	Multi-threaded HTTP	TCP	2h	MLab, idle link
BISmark				
Downstream throughput	Single-thread HTTP	TCP	30 min	curlget from Host/ proc/net/dev
	Passive throughput	N/A	30 min	
	Capacity	UDP	12h	ShaperProbe
Upstream throughput	Single-thread HTTP	TCP	30 min	curlput to Host/ proc/net/dev
	Passive throughput	N/A	30 min	
	Capacity	UDP	12h	ShaperProbe

techniques themselves. Multiple parallel TCP sessions nearly achieve the advertised throughput. UDP measurements also produce consistent measurements of throughput that are closer to the multi-threaded TCP measurement. A single-threaded TCP session may not achieve the same throughput, but accounting for cross-traffic with passive measurements does yield a better throughput estimate.

The behavior of single-threaded TCP measurements varies for different access links. We compare throughput for two BISmark users with the same ISP and service plan (AT&T; 3 Mbits/s down, 384 Kbits/s up) who live only a few blocks apart. Figure 4 shows that User 2 consistently sees nearly 20% higher throughput than User 1. One possible explanation for this difference is that the two users experience different loss rates: User 1 experiences four times more packet loss in both directions than User 2. The baseline latencies also differ by about 16ms (8ms vs. 24ms). We confirmed from the respective modem portals that User 1 has interleaving disabled and User 2 has interleaving enabled. Thus, User 2's connection recovers better from line noise. Single-threaded downloads suffer more from high packet loss rates than multi-threaded downloads; interleaving reduces the packet loss rate, and thus improves the performance of a single-threaded download. For the rest of the paper, we consider only multi-threaded TCP throughput.

Takeaway: Different throughput measurement techniques capture different aspects of throughput. A single-threaded TCP session is sensitive to packet loss. Augmenting this measurement with passive usage measurements improves

its accuracy. Multi-threaded TCP and the UDP capacity measurements measure the access link capacity more accurately and are more robust to loss.

3.2. Throughput performance

We investigate the throughput obtained by users in the SamKnows deployment and the consistency of these measurements.

What performance do users achieve? Figure 5 shows the average download and upload speeds for each router in the SamKnows dataset. Each point in the scatterplot shows the average performance that a single router in the deployment achieves. Clusters of points in the plot reveal common service plans of different ISPs, identified in the plot by labels. In general, these results agree with previous^{7,14} work, although our dataset also includes Verizon FiOS (FTTP) users that clearly stand out and other recent service offerings (e.g., AT&T U-Verse). Although there are some noticeable clusters around various service plans, there appears to be considerable variation, even within a single service plan. We seek to characterize both the performance variations and their causes.

Do users achieve consistent performance? We analyze whether routers in the SamKnows deployment consistently achieve their peak performance using the *Avg/P95* metric, which we define as the ratio of the average upload or download throughput obtained by a user to the 95th percentile of the upload or download throughput value obtained by the same user. A higher ratio reflects that a router's upload and download rates more consistently achieve performance that is closest to the highest rate; lower values indicate that performance fluctuates.

Figure 6 shows the CDF of the *Avg/P95* metric across users from each ISP. Most access links achieve throughput close to their 95th percentile value. Certain ISPs (e.g., Cox and Cablevision) achieve an average download throughput that is significantly less than their 95th percentile. Upload

Figure 3. Comparison of various methods of measuring throughput. (SamKnows and BISmark)




Figure 4. Users with the same service plan but different loss profiles see different performance. User 1 has higher loss and sees lower performance. (BISmark)





Figure 5. Average download rate versus the average upload rate obtained by individual users in the dataset. (SamKnows)



throughput is much more consistent, possibly because upload rates are typically much lower.

Why is performance sometimes inconsistent? One possible explanation for inconsistent download performance is that the access link may exhibit different performance characteristics depending on time of day. Figure 7a shows the Avg/P95 metric across the time of day. We obtain the average measurement reported by each user at that time of day and normalize it with the 95th percentile value of that user over all reports. Cablevision users see, on average, a 40% drop in performance

Figure 6. Consistency of throughput performance: The average throughput of each user is normalized by the 95th percentile value obtained by that user. (SamKnows)





Figure 7. Time of day is significant: The average download throughput for Cablevision and Cox users drops significantly during the evening peak time. Throughput is also significantly more variable during peak time. (SamKnows) (a) Peak and worst performance differ by up to 40%. (b) The standard deviation of throughput measurements increases during peak hours, most significantly for ISPs that see lower throughputs at peak hours.



in the peak evening hours; Cox users experience a 20% performance reduction on average. This effect exists for other ISPs to a lesser extent, confirming prior findings.⁷ Without knowing the service plan for each user, we cannot, in general, say whether the decrease in performance represents a drop below the service plans for those users (e.g., these users might see rates higher than their plan during off-peak hours). However, the FCC's report,¹ which analyzes performance in comparison to advertised rates, confirms that Cablevision users do see average performance significantly less than the advertised rates. Figure 7b shows how the standard deviation of normalized throughput varies depending on the time of day. All ISPs experience more variable performance during peak hours. Although most ISPs do not suffer an increase in loss rates during peak hours, Cox does. ISPs that exhibit poor performance during peak hours may be under-provisioned; they may be experiencing congestion, or they may be explicitly throttling traffic during peak hours.

Takeaway: Throughput performance is more variable during peak hours. A one-time "speed test" measurement taken at the wrong time could likely report misleading numbers that do not have much bearing on performance over time.

3.3. Effect of traffic shaping on throughput

ISPs shape traffic in different ways, which makes it difficult to compare measurements across ISPs, and sometimes even across users within the same ISP. We study the effect of PowerBoost across different ISPs, time, and users. We also model how Comcast implements PowerBoost.

Which ISPs use PowerBoost, and how does it vary across ISPs? Each SamKnows throughput measurement lasts 30s, and each report is divided into six snapshots at roughly 5-s intervals for the duration of the test. This technique highlights the evolution of throughput performance over time. On a link that is subject to traffic shaping, the throughput during the last snapshot will be less than the throughput during the first. For each report, we normalize the throughput in each period by the throughput reported for the first period. The normalized throughput on an unshaped link is close to one for all intervals. On the other hand, on an access link configured with PowerBoost, the throughput in the last 5s should be less than the throughput in the first 5s (assuming that PowerBoost lasts less than 30s). Figure 8 shows the progression of average throughput over all users in an ISP: the average normalized throughput decreases over time. Our data shows that most cable ISPs provide some level of PowerBoost for less than 30s, at a rate of about 50% more than the normal rate. Cablevision's line is flat; this suggests that either it does not provide PowerBoost, or it lasts well over 30s consistently (in which case the throughput test would reflect only the effect of PowerBoost). The gradual decrease, rather than an abrupt decrease, could be because PowerBoost durations vary across users or because the ISP changes PowerBoost parameters based on network state. In the case of uploads, only Comcast and Cox seem to deploy PowerBoost; in these cases, we observed a difference in throughput of about 20%. DSL ISPs do not appear to implement PowerBoost.

Do different users see different PowerBoost effects? We investigate Comcast's use of PowerBoost using the BISmark testbed. According to Comcast,⁶ their implementation of PowerBoost provides higher throughput for the first 10MB of a download and the first 5MB of an upload. We measure the shaped throughput for download and upload at the receiver using `tcpdump`. Because our tests are intrusive, we conducted them only a few times, but the results are consistent across traffic generators and ports.

Figure 9 shows the downstream throughput profiles for four users, each identified by their modem type. Although the modem does not affect burst rates, it does have different amounts of buffering, which can affect latency. All four users experience PowerBoost effects, but, surprisingly, each user experiences a different traffic shaping profile: The user with a D-LINK modem sees a peak rate of about 21 Mbits/s for 3s, 18.5 Mbits/s for a further 10s, and a steady-state rate of 12.5 Mbits/s. The Motorola user sees a peak rate of 21 Mbits/s for about 8s. The D-LINK profile can be modeled as a cascaded token bucket filter with rates of 18.5 Mbits/s and 12.5 Mbits/s, and buffer sizes of 10MB and 1MB, respectively, with a capacity of 21 Mbits/s. Upload profiles vary across different users as well, although the shaping profiles seem to indicate that only a single token bucket is applied on the uplink.

Takeaway: Many cable ISPs implement PowerBoost, which could distort speedtest-like measurements. In particular, any throughput measurement that lasts less than 35s will mainly capture the effects of PowerBoost. While some people may be interested only in short-term burst rates, others may be more interested in long-term rates. Any

Figure 8. The average throughput during the measurement decreases for the ISPs that enable PowerBoost. (SamKnows)





Figure 9. The level and duration of the burstiness are different for users with different modems, suggesting different shaping mechanisms or parameters. (BISmark)



throughput benchmark should aim to characterize both burst rates and steady-state throughput rates.

4. UNDERSTANDING LATENCY

We show how latency can drastically affect performance, even on ISP service plans with high throughput. We then study how various factors ranging from the user's modem to ISP traffic shaping policies can affect latency.

4.1. How (and why) to measure latency

Latency not only affects the throughput that users achieve, but it also affects the performance that users perceive, since it affects everything from DNS lookup time to the time to set up a TCP connection. Although measuring latency appears straightforward, arriving at the appropriate metric is a subtle challenge because our goal is to isolate the performance of the access link from the performance of the end-to-end path. End-to-end latency between endpoints is a common metric in network measurement, but it reflects the delay that a user experiences along a wide-area path. We use two metrics that are more appropriate for access networks. The first metric is the *last-mile latency*, which is the latency to the first IP hop inside the ISP's network. The last-mile latency captures the latency of the access link, which could affect gaming or short downloads. We measure last-mile latency in both of our deployments. The second metric we define is *latency under load*, which is the latency that a user experiences during an upload or download (i.e., when the link is saturated in either direction). For BISmark, we measure the last-mile latency under load; on the SamKnows platform, we measure end-to-end latency under load. Table 4 summarizes the latency measurements we collect.

We investigate the effect of last-mile latency on download times for popular Web pages. Figure 10 shows the download time for `www.facebook.com` and how it varies by both the user's throughput and baseline last-mile latency. Figure 10a plots the 95th percentile of each user's downstream throughput versus the average time it takes to download all objects from `www.facebook.com`. The average size of the download is

Table 4. SamKnows and BISmark latency and loss tests.

Parameter	Type	Prot.	Freq.	Comments
SamKnows				
Latency	End-to-end	UDP	600 pkts/h	MLab
	End-to-end	ICMP	5 pkts/h	MLab
	Last mile	ICMP	5 pkts/h	First IP hop
	Upstream load	ICMP	2 h	During upload
	Downstream load	ICMP	2 h	During download
Loss	End-to-end	UDP	600 pkts/h	MLab
Jitter	Bidirectional	UDP	1 h	500 pkts/30s
Web GET	HTTP	TCP	1 h	Alexa sites
BISmark				
Latency	End-to-end	ICMP	5 min	Host
	Last mile	ICMP	5 min	First IP hop
	Upstream load	ICMP	30 min	During upload
	Downstream load	ICMP	30 min	During download
Packet loss	End-to-end	UDP	15 min	D-ITG
Jitter	End-to-end	UDP	15 min	D-ITG

125KB. As expected, the download times decrease as throughput increases; interestingly, there is negligible improvement beyond a rate of 6 Mbits/s. Figure 10b plots download time against the baseline latency for all users whose downstream throughput (95th percentile) exceeds 6 Mbits/s. Minimum download times increase by about 50% when baseline latencies increase from 10ms to 40ms. The pronounced effect of latency on download time for Web objects underscores the influence of baseline latency.

4.2. Last-mile latency

We obtain the last-mile latency by running `traceroute` to a wide-area destination and extracting the first IP address along the path that is not a NAT address. Note that we are measuring the latency to the first network-layer hop, which

may not in fact be the DSLAM or the CMTS, because some ISPs have layer-two DSLAMs that are not visible in traceroute. The possibility of measuring slightly further than the DSLAM or CMTS should not materially affect our results, since the latency between hops inside an ISP is typically much smaller than the last-mile latency.

How does access technology affect last-mile latency? Last-mile latency is generally quite high, varying from about 10ms to nearly 40ms (ranging from 40–80% of the end-to-end path latency). Last-mile latency is also highly variable. One might expect that variance would be lower for DSL ISPs, since it is not a shared medium like cable. Surprisingly, the opposite is true: AT&T and Verizon have high variance compared to the mean. Qwest also has high variance, though it is a smaller

Figure 10. Effect of downstream throughput and baseline latency on fetch time from facebook.com. (SamKnows) (a) Fetch time stabilizes above 6 Mbits/s. (b) Baseline latency affects fetch times.





Figure 11. The baseline last-mile latency for each user is computed as the 10th percentile of the last-mile latency. Most users see latencies less than 10ms, but there are a significant number of users with the last-mile latency greater than 10ms. (SamKnows)



fraction of the mean. To understand this variance, we divide different users in each ISP according to their baseline latency, as shown in Figure 11. Most users of cable ISPs are in the 0–10ms interval. On the other hand, a significant proportion of DSL users have baseline last-mile latencies more than 20ms, with some users seeing last-mile latencies as high as 50 to 60ms. Based on discussions with network operators, we believe DSL ISPs may be enabling an interleaved local loop for these users. We also analyzed loss rate and jitter. We found that the average loss rates are small, but variance is high for all ISPs, suggesting bursty loss. Jitter has similar characteristics; although the average jitter is low, the variation is high, especially on the upstream connection.

Takeaway: Cable users tend to have lower last-mile latency and jitter, while for DSL users it may vary significantly based on physical factors such as distance to the DSLAM or line quality.

4.3. Latency under load

We now turn our attention to latency under load, a characteristic that is often affected by “bufferbloat,” or excessive buffering in network devices.¹⁹

Problem: Bufferbloat. Buffers on DSL and cable modems are too large. Buffering affects latency during periods when the access link is loaded; during these periods, packets can see substantial delays as they queue in the buffer. The capacity of the uplink also affects the latency that buffering introduces. For a given buffer size, the queuing delay will be lower on access links with higher capacities because the draining rate is higher. We study the effect of buffering on access links by measuring latency when the access link is saturated, under the assumption that the last mile is the bottleneck.

How widespread is bufferbloat? Figure 12 shows the average ratios of latency under load to baseline latency for each user across different ISPs for the SamKnows data. The histogram shows the latencies when the uplink and the downlink are saturated separately. This plot confirms that bufferbloat affects users across all ISPs, albeit to different extents. The factor of increase when the uplink is saturated is much higher than when the downlink is saturated. One plausible explanation is that the downlink usually has more capacity than the uplink, so buffering on the ISP side is lower. The home network is often better provisioned than the downlink, so downstream traffic experiences less buffering in the modem. The high variability in the latency under load can be partly explained by the variety in service plans; for example, AT&T offers plans ranging from 768 Kbits/s to 6 Mbits/s for DSL and up to 18 Mbits/s for UVerse and from 128 Kbits/s to more than 1 Mbit/s for upstream. In contrast, Comcast offers fewer service plans, which makes it easier to design a device that works well for all service plans.

How does modem buffering affect latency under load? To study the effects of modem buffers on latency under load, we conduct tests on AT&T and Comcast modems using BISmark. We ran tests on the best AT&T DSL plan (6 Mbits/s down; 512 Kbits/s up). We first started ICMP ping to the last

mile hop. After 30s, we flooded the uplink (at 1 Mbits/s for AT&T and at 10 Mbits/s for Comcast, using *iperf*’s UDP measurement). After 60s, we stopped *iperf*, but let ping continue for another 30s. The ping measurements before and after the *iperf* test established the baseline latency.

Figure 13 shows the latency under load for three different DSL modems. In all cases, the latency skyrockets when flooding begins and plateaus when the buffer is saturated. This latency plateau indicates the size of the buffer, since we know the uplink draining rate. Surprisingly, we observed more than an order of magnitude of difference in buffering in different modems. The 2Wire modem introduces the lowest worst-case latency of 800ms, the Motorola modem about 1600ms, while the Westell introduces more than 10s of latency! Comcast users experienced as much as 350ms of latency under load. Because modems are usually the same across service plans, we expect that latency under load may be even worse for users with slower plans (and, hence, slower drain rates).

We perform experiments in Emulab⁸ to model modem buffering; the topology has two end-hosts and one router. We configured a token bucket filter using *tc* with the buffer size as 512 Kbits/s times the maximum latency that the modem introduces. This calculation yields 640KB for Westell, 100KB for Motorola, and 55KB for 2Wire. This simple setup almost perfectly captures the latency profile that the actual modems exhibit. We observed little difference in throughput for the three buffer sizes. We also emulated other buffer sizes. For a 512 Kbits/s uplink, we observed that the modem buffers exceeding 20KB do little for throughput, but cause a linear increase in latency under load. Our

Figure 12. Latency under load: the factor by which latency goes up when the upstream or the downstream is busy translates to significant real latencies, often in the order of seconds. (SamKnows)





Figure 13. Different buffer sizes across modems lead to wide disparities in latencies when the upstream link is busy. (BISmark)




experiments confirm that buffer sizes in all three modems are too large for the uplink.

Can data transfer be modified to improve latency under load?

There has been recent work done on active queue management techniques to mitigate the effect of large buffers.¹⁹ In this section, we explore how large “bulk” flows and delay-sensitive flows can coexist without interfering with one another. We compare the effects of a 50MB download on a G.711 VoIP call in three different conditions: (1) not applying any traffic control, (2) sending intermittent traffic at capacity on 10.8s ON and 5.3s OFF cycle, and (3) shaping using the WonderShaper²³ tool. Figure 14 shows the result of this experiment. Without traffic control, the transfer takes 25.3s; immediately after the PowerBoost period, however, the VoIP call starts suffering high latency and loss until the end of the transfer. In the second scenario, traffic is sent in pulses, and the download takes 26.9s. In the third case, traffic is sent at just under the long-term rate and the download takes 32.2s. Sending intermittent traffic and shaping the traffic with Wondershaper do not increase latency much, because they do not ever fully deplete the tokens. The appropriate ON/OFF periods for intermittent transfers depend on the token bucket parameters^a and the size of the file to be transferred. Both approaches achieve similar long-term rates but yield significant latency benefit. These approaches require the user to properly tune the traffic shaping parameters according to the access link.

Takeaway: Modem buffers are too large. The smallest buffers we see induce nearly 1-s latency under load for AT&T

Figure 14. It is possible to maintain low latency by modifying data transfer behavior. (BISmark) (a) Throughput (b) Latency



^a If ρ_r is the rate we want to reserve for real-time applications, and ρ_t the token rate, the condition to be satisfied is: $(\rho_b + \rho_r - \rho_t) \times \tau_{on} \leq \tau_{off} \times (\rho_t - \rho_r)$, where ρ_b is the sending rate during the pulse, and τ_{on} and τ_{off} are the ON and the OFF times, respectively.

and 300ms for Comcast. Buffering degrades both interactivity and throughput. Transferring data in shorter bursts or shaping traffic using tools like WonderShaper can mitigate buffering problems.

5. LESSONS LEARNED

We conclude with some high-level lessons and suggestions for future research directions. One significant takeaway for users, policymakers, ISPs, and researchers is that *understanding access network performance requires continual measurement from the home router*. Existing speed tests and end-to-end latency measurements do not reflect access network performance over an extended period of time, and they neglect confounding factors within the home network. Our study of broadband networks yields several lessons:

Lesson 1 (One Measurement Does Not Fit All) *Different ISPs use different policies and traffic shaping behaviors that make it difficult to compare measurements across ISPs.*

There is no single number that characterizes performance, or even throughput. Certain ISP practices such as PowerBoost can distort benchmarking measurements; ISPs might even design their networks so that widely used performance tests yield good performance. Developing a benchmarking suite for ISP performance that users can understand (e.g., in terms of the applications they use) is critical; the measurements we develop in this paper may be a good starting point for that. Along these lines, more work is needed to understand the performance of specific applications, such as how video streaming performance compares across ISPs. The Netflix study on ISP streaming performance¹⁸ is a good start, but more such performance benchmarks are needed.

Lesson 2 (One ISP Does Not Fit All) *There is no “best” ISP for all users. Different users may prefer different ISPs depending on their usage profiles and how those ISPs perform along performance dimensions that matter to them.*

Different ISPs may be “better” along different performance dimensions, and the service plan that a user buys is only part of the picture. For example, we saw that, above a certain throughput, latency is the dominant factor in determining Web page load time. Similarly, a gamer might require low latency or jitter, while an avid file swapper may be more interested in high throughput. An imminent technical and usability challenge is to summarize access network performance data so that users can make informed choices about the service plans that are most appropriate for them (akin to a “performance nutrition label”²²). Our recent work proposes some first steps in this direction.²¹

Lesson 3 (Home Network Equipment Matters) *A user’s home network infrastructure can significantly affect performance.*

Modems can introduce latency variations that are orders of magnitude more than the variations introduced by the ISP. Other effects inside the home that we have not yet studied, such as the wireless network, may also ultimately affect the

user's experience. More research is needed to understand the characteristics of traffic inside the home and how it affects performance.

Acknowledgments

We thank the participants in the SamKnows and BISmark studies, and Walter Johnston at the FCC for help and access to the data from the SamKnows study. This project is supported by the National Science Foundation through awards CNS-1059350, CNS-0643974, a generous Google Focused Research Award, the European Community's Seventh Framework Programme (FP7/2007-2013) no. 258378 (FIGARO), and the ANR project C'MON.

References

1. Measuring Broadband America: A Report on Consumer Wireline Broadband Performance in the U.S. http://www.fcc.gov/cgb/measuringbroadbandreport/Measuring_US_-_Main_Report_Full.pdf.
2. Does broadband need its own government nutrition label? <http://arstechnica.com/tech-policy/news/2009/10/does-broadband-needs-its-own-government-nutrition-label.ars>, Ars Technica (Oct. 2010).
3. Bauer, S., Clark, D., Lehr, W. Understanding broadband speed measurements. In *38th Research Conference on Communication, Information and Internet Policy* (Arlington, VA, 2010).
4. Bode, K. FCC: One Million Speedtests and Counting. <http://www.dslreports.com/shownews/FCC-One-Million-Speedtests-And-Counting-109440>, July 2010.
5. Carlson, R. Network Diagnostic Tool. <http://e2epi.internet2.edu/ndt/>.
6. Comcast FAQ. <http://customer.comcast.com/Pages/FAQViewer.aspx?Guid=024f23d4-c316-4a58-89f6-f5f3f5dbdcf6>, Oct. 2007.
7. Dischinger, M., Haeberlen, A., Gummadi, K.P., Saroiu, S. Characterizing residential broadband networks. In *Proceedings of ACM SIGCOMM Internet Measurement Conference* (San Diego, CA, USA, Oct. 2007).
8. Emulab. <http://www.emulab.net/>, 2006.
9. National Broadband Plan. <http://www.broadband.gov/>.
10. Internet World Stats. <http://www.internetworldstats.com/dsl.htm>.
11. Asymmetric Digital Subscriber Line Transceivers. ITU-T G.992.1, 1999.
12. Asymmetric Digital Subscriber Line (ADSL) Transceivers - Extended Bandwidth ADSL2 (ADSL2Plus). ITU-T G.992.5, 2003.
13. Data-over-Cable Service Interface Specifications: Radio-Frequency Interface Specification. ITU-T J.112 2004.
14. Kreibich, C., Weaver, N., Nechaev, B., Paxson, V. Netalyzr: Illuminating the edge network. In *Proceedings of Internet Measurement Conference* (Melbourne, Australia, Nov. 2010).
15. Maier, G., Feldmann, A., Paxson, V., Allman, M. On dominant characteristics of residential broadband internet traffic. In *ACM Internet Measurement Conference* (2009).
16. Mathis, M. et al. Network Path and Application Diagnosis. <http://www.psc.edu/networking/projects/pathdiag/>.
17. Netalyzr. <http://netalyzr.icsi.berkeley.edu/>.
18. NetFlix Performance on Top ISP Networks. <http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html>, Jan. 2011.
19. Nichols, K., Jacobson, V. Controlling queue delay. *Queue* 10, 5 (May 2012), 20:20–20:34.
20. ShaperProbe. <http://www.cc.gatech.edu/~partha/diffprobe/shaperprobe.html>.
21. Sundaresan, S., Feamster, N., Teixeira, R., Tang, A., Edwards, K., Grinter, R., Chetty, M., de Donato, W. Helping users shop for ISPs with internet nutrition labels. In *ACM SIGCOMM Workshop on Home Networks* (2011).
22. Vorhaus, D. A New Way to Measure Broadband in America. <http://blog.broadband.gov/?entryId=359987>, Apr. 2010.
23. WonderShaper. <http://lartc.org/wondershaper/>, 2002.

Srikanth Sundaresan (srikanth@gatech.edu), Georgia Tech, Atlanta, GA.

Walter de Donato (walter.dedonato@unina.it), University of Napoli Federico II, Napoli, Italy.

Nick Feamster (feamster@cs.umd.edu), University of Maryland College Park, MD.

Renata Teixeira (renata.teixeira@lip6.fr), CNRS/UPMC Sorbonne University, Paris, France.

Sam Crawford (sam@samknows.com), SamKnows, London, UK.

Antonio Pescapè (pescapè@unina.it), University of Napoli Federico II, Napoli, Italy.

© 2012 ACM 0001-0782/12/09 \$15.00

World-Renowned Journals from ACM

ACM publishes over 50 magazines and journals that cover an array of established as well as emerging areas of the computing field. IT professionals worldwide depend on ACM's publications to keep them abreast of the latest technological developments and industry news in a timely, comprehensive manner of the highest quality and integrity. For a complete listing of ACM's leading magazines & journals, including our renowned Transaction Series, please visit the ACM publications homepage: www.acm.org/pubs.

ACM Transactions on Interactive Intelligent Systems



ACM Transactions on Interactive Intelligent Systems (TIIS). This quarterly journal publishes papers on research encompassing the design, realization, or evaluation of interactive systems incorporating some form of machine intelligence.

ACM Transactions on Computation Theory



ACM Transactions on Computation Theory (ToCT). This quarterly peer-reviewed journal has an emphasis on computational complexity, foundations of cryptography and other computation-based topics in theoretical computer science.

PLEASE CONTACT ACM MEMBER SERVICES TO PLACE AN ORDER
 Phone: 1.800.342.6626 (U.S. and Canada)
 +1.212.626.0500 (Global)
 Fax: +1.212.944.1318
 (Hours: 8:30am–4:30pm, Eastern Time)
 Email: acmhelp@acm.org
 Mail: ACM Member Services
 General Post Office
 PO Box 30777
 New York, NY 10087-0777 USA



Association for
Computing Machinery

Advancing Computing as a Science & Profession

www.acm.org/pubs

CAREERS

Auburn University **Assistant/Associate Professor**

The Department of Computer Science and Software Engineering (CSSE) invites applications for a tenure-track faculty position at the Assistant/Associate Professor level to begin in Fall 2013. Salary will be commensurate with the candidate's qualifications. *Women and minorities are encouraged to apply.* Responsibilities include research, graduate student supervision, graduate and undergraduate teaching, and service. For applications at the Associate Professor level, a record of success in securing external funding for research is expected, and potential for successfully obtaining external research funding will be considered at the Assistant Professor level position.

Applicants must have a Ph.D. in computer science, software engineering, or a closely related field; however, applicants who are ABD may apply if they reasonably expect to complete the terminal degree prior to August 2013. We encourage candidates from all areas of computer science and software engineering to apply. We are seeking candidates specializing in artificial intelligence, computer systems, cyber security, software engineering, and theory. Excellent communication skills are required.

Baylor University **Lecturer of Computer Science**

The Department of Computer Science seeks a dedicated teacher and program advocate for a lecturer position beginning August, 2013. The ideal candidate will have a master's degree or Ph.D. in Computer Science or a related area, a commitment to undergraduate education, effective communication and organization skills, and industry/academic experience in game development, especially with graphics and/or engine development. For position details and application information please visit: <http://www.baylor.edu/hr/index.php?id=81302>

Baylor, the world's largest Baptist university, holds a Carnegie classification as a "high-research" institution. Baylor's mission is to educate men and women for worldwide leadership and service by integrating academic excellence and Christian commitment within a caring community. Baylor is actively recruiting new faculty with a strong commitment to the classroom and an equally strong commitment to discovering new knowledge as Baylor aspires to become a top tier research university while reaffirming and deepening its distinctive Christian mission as described in *Pro Futuris* (<http://www.baylor.edu/profuturis/>).

Baylor is a Baptist university affiliated with the Baptist General Convention of Texas. As an AA/EEO employer, Baylor encourages minorities, women, veterans, and persons with disabilities to apply.

Baylor University **Assistant, Associate or** **Full Professor of Computer Science**

The Department of Computer Science seeks a productive scholar and dedicated teacher for a tenured or tenure-track position beginning August, 2013. The ideal candidate will hold a terminal degree in Computer Science or a closely related field and demonstrate scholarly capability and an established and active independent research agenda in one of several core areas of interest, including, but not limited to, game design and development, software engineering, computational biology, machine learning and large-scale data mining. A successful candidate will also exhibit a passion for teaching and mentoring at the graduate and undergraduate level. For position details and application information please visit: <http://www.baylor.edu/hr/index.php?id=81302>

Baylor, the world's largest Baptist university, holds a Carnegie classification as a "high-research" institution. Baylor's mission is to educate men and women for worldwide leadership and service by integrating academic excellence and Christian commitment within a caring community. Baylor is actively recruiting new faculty with a strong commitment to the classroom and an equally strong commitment to discovering new knowledge as Baylor aspires to become a top tier research university while reaffirming and deepening its distinctive Christian mission as described in *Pro Futuris* (<http://www.baylor.edu/profuturis/>).

Baylor is a Baptist university affiliated with the Baptist General Convention of Texas. As an AA/EEO employer, Baylor encourages minorities, women, veterans, and persons with disabilities to apply.

Boise State University **Four Assistant/Associate/Full Professors of** **Computer Science**

The Department of Computer Science at Boise State University invites applications for four open-rank, tenure/tenure-track positions. Applicants should have a commitment to excellence in teaching, a desire to make significant contributions in research, and an interest in collaborating with faculty and local industry to develop and sustain high-profile funded research programs.

The Department seeks candidates specializing in databases or software engineering. Salaries will be competitive. An earned Ph.D. in Computer Science or an equivalent field is required by the date of hire.

For more information, including application details, please visit us online at <http://coen.boisestate.edu/cs/opportunities>.

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.

Boston University **Department of Electrical &** **Computer Engineering (ECE)** **Faculty Positions**

The **Department of Electrical & Computer Engineering (ECE) at Boston University (BU)** is seeking candidates for anticipated faculty positions in Computer Engineering. All areas and ranks will be considered, with particular interest in entry-level candidates in software systems and cybersecurity. The Department is seeking to foster growth in the broad, interdisciplinary topics of energy, health, information systems, and cyberphysical systems. Candidates with research interests that transcend the traditional boundaries of ECE are strongly encouraged to apply. Joint appointments with other BU departments and with the Division of Systems Engineering are likely for candidates with appropriate experience and interests.

Qualified candidates must possess a relevant, earned PhD, and have a demonstrable ability to teach effectively, develop funded research programs in their area of expertise, and contribute to the tradition of excellence in research that is characteristic of the ECE Department. Self-motivated individuals who thrive on challenge and are eager to utilize their expertise to strengthen an ambitious program of departmental enhancement are desired. Women, minorities, and candidates from other underrepresented groups are especially encouraged to apply and help us continue building an exceptional 21st century university department.

ECE at BU is a world-class department with excellent resources that is steadily gaining national and international prominence for its exceptional research and education record. ECE is part of BU's rapidly growing and innovative College of Engineering, and currently consists of 40 faculty members, 200 graduate students, and 250 BS majors. Outstanding collaboration opportunities are available with nationally recognized medical centers and universities/colleges, nearby research centers, and industry throughout the Boston area.

Beyond its research and academic activities, BU has a lively, urban campus situated along the banks of the Charles River in Boston's historic Fenway-Kenmore neighborhood. The campus and surrounding areas offer limitless opportunities for recreational activities, from world-class art and performances to sporting events and fine dining.

Please visit <http://www.bu.edu/ece/facultysearch> for instructions on how to apply. Application deadline is December 31, 2012. The review of applications will begin on October 1, 2012. Therefore, applicants are encouraged to apply early. Boston University is an Equal Opportunity/Affirmative Action Employer.

Bowling Green State University THREE Tenure-track positions in CS

We are seeking to fill three tenure-track Assistant Professor positions in Computer Science to teach a variety of courses at the undergraduate and graduate levels and to be productive in scholarly research and sponsored projects. Preferred specializations are Big Data/high performance computing/visualization, computer and information security, software engineering. Applicants must hold a Ph.D. in CS (or closely related field) or complete it by the start date in August 2013, and be committed to excellence in teaching, scholarly research, and external funding. BGSU is located about an hour from Detroit airport, and the area offers excellent quality of life. BGSU is an AA/EOE. Email cssearch2013@cs.bgsu.edu a letter of interest, along with the name, title, email, and postal addresses for three professional references, curriculum vitae, official or unofficial transcripts by January 13, 2013. For finalists, three current letters of reference, official transcript of the highest degree, and background check are required. For details, go to <http://www.bgsu.edu/departments/compsci/jobs>

California State Polytechnic University, Pomona
Computer Science Department
Assistant Professor
<http://www.csupomona.edu/~cs/>

The Computer Science Department invites applications for a tenure-track position at the rank of As-

stant Professor to begin Fall 2013. We are particularly interested in candidates with specialization in Cloud Computing, Data Mining, or Computer Graphics and Animation. Cal Poly Pomona is 30 miles east of L.A. and is one of 23 campuses in the California State University. The department offers an ABET-accredited B.S. program and an M.S. program. Qualifications: Possess, or complete by September 2013, a Ph.D. in Computer Science or closely related area. Demonstrate strong English communication skills, a commitment to actively engage in the teaching, research, and curricular development activities of the department at both undergraduate and graduate levels, and ability to work with a diverse student body and multicultural constituencies. Ability to teach a broad range of courses, and to articulate complex subject matter to students at all educational levels. First consideration will be given to completed applications received no later than December 15, 2012. Contact: Faculty Search Committee, Computer Science Department, Cal Poly Pomona, Pomona, CA 91768. Email: cs@csupomona.edu. Cal Poly Pomona is an Equal Opportunity, Affirmative Action Employer. Position announcement available at: <http://academic.csupomona.edu/faculty/positions.aspx>. Lawful authorization to work in US required for hiring.

Colby College Visiting Faculty Position

Colby College invites applications for a one-year, full time visiting faculty position in computer science, beginning September 1, 2013. Teach-

ing responsibilities will include a mix of core CS courses and upper level electives.

For more information and required materials, see <http://cs.colby.edu/>. Review of applications will begin January 28, 2013 and continue until the position is filled.

Colby College is committed to equality and diversity and is an equal opportunity employer. We encourage inquiries from candidates who will contribute to the cultural and ethnic diversity of our college. Colby College does not discriminate on the basis of race, gender, sexual orientation, disability, religion, ancestry, or national origin, or age in employment or in our educational programs. For more information about the College, please visit our website: www.colby.edu.

Colorado State University Department of Computer Science Tenure Track Assistant or Associate Professor

Colorado State University is accepting applications for a tenure-track assistant or associate professor in Computer Science, beginning fall 2013. Only candidates in bioinformatics/ computational biology will be considered. This position is part of a university-wide effort to recruit additional faculty in biosciences and bioinformatics. More information may be viewed at <http://www.cs.colostate.edu>.

Applications must be received by **January 11, 2013**. Submit materials at <http://cns.natsci.colostate.edu/employment/Compsci/>. Application materials of semifinalist candidates, including letters of reference, will be made available for

Computer Science

Multiple senior faculty positions in computer science are available at Cornell's new CornellNYC Tech campus in New York City. Faculty hired in these positions will be in the Department of Computer Science, which will span the Ithaca and New York City campuses, but their teaching and research will be based in New York City. We will consider only candidates at the Associate and Full Professor level, but will consider candidates from all areas of computer science and related fields.

Candidates whose work fits into one of the three initial hubs at CornellNYC, Connective Media, Healthier Life, and Built Environment, are particularly encouraged. Candidates must hold a Ph.D., must have demonstrated an ability to conduct outstanding research, and must also have a strong interest in the technology commercialization and entrepreneurship mission of the campus. In addition, interest in international programs and/or pre-college (K-12) education is advantageous.

This search may include Cornell faculty positions that are part of the Technion-Cornell Innovation Institute.

To ensure full consideration, applications should be received by December 1, 2012, but will be accepted until all positions are filled. Candidates should submit a curriculum vitae, brief statements of research and teaching interests on-line at <https://academicjobsonline.org/ajo/jobs/1915>



Diversity and inclusion have been and continue to be a part of our heritage. Cornell University is a recognized EEO/AA employer and educator.

Electrical and Systems Engineering



Tenured/Tenure-Track Faculty Positions

The Department of Electrical and Systems Engineering of the School of Engineering and Applied Science at the **University of Pennsylvania** invites applications for tenured and tenure-track faculty positions at all levels. Candidates must hold a Ph.D. in Electrical Engineering, Systems Engineering, or related area. The department seeks individuals with exceptional promise for, or proven record of, research achievement, who will take a position of international leadership in defining their field of study, and excel in undergraduate and graduate education. Leadership in cross-disciplinary and multi-disciplinary collaborations is of particular interest. We are interested in candidates in all areas that enhance our research strengths in

1. Nanodevices and nanosystems (nanophotonics, nanoelectronics, integrated devices and systems at nanoscale),
2. Circuits and computer engineering (analog and digital circuits, emerging circuit design, computer engineering, embedded systems), and
3. Information and decision systems (communications, control, signal processing, network science, markets and social systems).

Prospective candidates in all areas are strongly encouraged to address large scale societal problems in energy, transportation, health, economic and financial networks, critical infrastructure, and national security. Diversity candidates are strongly encouraged to apply. Interested persons should submit an online application at <http://www.ease.upenn.edu/faculty-positions> including curriculum vitae, statement of research and teaching interests, and the names of at least four references. Review of applications will begin on December 1, 2012.

The University of Pennsylvania is an Equal Opportunity Employer. Minorities/ Women/Individuals with Disabilities/Veterans are encouraged to apply.

review by the entire faculty of the Department of Computer Science.

CSU is an EO/EA/AA employer. Colorado State University conducts background checks on all final candidates.

DePauw University
Department of Computer Science
Tenure-Track Position

DePauw University is seeking qualified candidates for a tenure-track position in computer science beginning August 2013. Rank and salary will be commensurate with experience. Visit <http://www.depauw.edu/offices/academic-affairs/open-faculty-positions/> for details.

Florida Institute of Technology
Professor (Open Rank)

The Department of Computer Sciences at the Florida Institute of Technology invites applications for two open faculty position beginning in August 2013. Required qualifications include an earned Ph.D. with a specialization in computer security, evidence of the ability to develop and sustain an active research program, and a sincere interest in quality teaching at the undergraduate and graduate level. All areas of computer security are of interest, but our preference is for a faculty member who conducts research that is both pragmatic and disruptive. New faculty will be expected to work within the Harris Institute for Assured Informa-

tion and support its overall mission and agenda by developing courses at the undergraduate and graduate level, helping to build our Ph.D. program, and executing their own research agenda. The department has significant active research funding from multiple government agencies and commercial companies and a significant cohort of faculty who work in cybersecurity. Florida Tech is an NSA/DHS designated Center for Academic Excellence in Information Assurance Research.

Florida Tech is located in Melbourne on Florida's Space Coast, one of the nation's fastest-growing high-tech areas. The campus is 5 minutes from the Indian River estuary, 10 minutes from the Atlantic Ocean and 50 minutes from Kennedy Space Center and Orlando. For more information on the Department of Computer Sciences please visit <http://cs.fit.edu>. Information on the Harris Institute is available at <http://harris-institute.fit.edu/>.

Applicants should send a letter of intent, curriculum vitae, research and teaching summary, and full contact information for at least three references to faculty-search@cs.fit.edu. Review of applications will begin immediately and continue until the position is filled. Florida Tech is an Equal Opportunity Employer.

Application review will begin in November and continue until the position is filled.

Florida State University
Tenure-Track, Assistant Professor

The Department of Computer Science at the Florida State University invites applications for

one tenure-track Assistant Professor position to begin August 15, 2013. Positions are 9-mo, full-time, tenure-track, and benefits eligible. We encourage strong applicants in all areas of Computer Science to apply. Preference may be given to applicants with research experience in the areas of Big Data/Databases and Software Engineering. Applicants should hold a PhD in Computer Science or closely related field, and have excellent research and teaching accomplishments/potential. The department offers degrees at the BS, MS, and PhD levels. The department is an NSA Center of Academic Excellence in Information Assurance Education (CAE/IAE) and Research (CAE-R).

FSU is classified as a Carnegie Research I university. Its primary role is to serve as a center for advanced graduate and professional studies while emphasizing research and providing excellence in undergraduate education. The department has experienced rapid growth in the major and new degree programs. Further information can be found at <http://www.cs.fsu.edu>

Screening will begin January 1, 2013 and will continue until the position is filled. Please apply online with curriculum vitae, statements of teaching and research philosophy, and the names of five references, at <http://www.cs.fsu.edu/positions/apply.html>

Questions can be e-mailed to Prof. Sudhir Aggarwal, Chair Search Committee, recruitment@cs.fsu.edu or to Prof. Robert van Engelen, Department Chair, chair@cs.fsu.edu.

The Florida State University is a Public Records Agency and an Equal Opportunity/Access/Affirmative Action employer, committed to diversity in hiring.



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

Computer Science

Multiple faculty positions are available at Cornell's Department of Computer Science, based in Ithaca, New York. Candidates are invited to apply at all levels including tenured, tenure-track, or lecturer, and from all areas of computer science and related fields.

Tenured and tenure track faculty must hold the equivalent of a Ph.D.; applicants for the position must have demonstrated an ability to conduct outstanding research. Lecturers must hold the equivalent of a Master's degree, with a Ph.D. preferred. To ensure full consideration, applications should be received by December 1, 2012, but will be accepted until all positions are filled.

Applicants should submit a curriculum vita, brief statements of research and teaching interests, and arrange to have at least three references letters submitted at:

<https://academicjobsonline.org/ajo/jobs/1917>



Diversity and inclusion have been and continue to be a part of our heritage. Cornell University is a recognized EEO/AA employer and educator.

Fordham University
Assistant Professor, Computer &
Information Science

The Department of Computer and Information Science (CIS) invites applications for a tenure-track Assistant Professor to begin in September 2013. A Ph.D. in Computer Science, Information Science, Informatics, or closely related field is required. The position requires excellence in teaching undergraduate and graduate courses, good communication skills, and demonstrated research potential with the ability to attract external research funding. We are interested in candidates with expertise in computational neuroscience, systems neuroscience, or other closely related areas such as neuroinformatics, brain and cognitive science, or cognitive computing and informatics.

The CIS department offers graduate and undergraduate programs at Fordham's Rose Hill campus in the Bronx, Lincoln Center campus in Manhattan, and Westchester campus in West Harrison, NY. For information about the department please visit <http://www.cis.fordham.edu>.

Review of applications will begin February 1st, 2013. Preferably submit your application electronically using the system at <https://secure.interfolio.com/apply/15923>. Alternatively you may send a letter of application, research summary, curriculum vitae, statement of teaching philosophy, and three letters of reference to faculty_search@cis.fordham.edu or to:

Faculty Search Committee Chair,
CIS Department
Fordham University, JMH 340
441 E. Fordham Road
Bronx, NY 10458

Fordham is an independent, Catholic University in the Jesuit tradition that welcomes applications from men and women of all backgrounds. Fordham is an Equal Opportunity/Affirmative Action Employer.

The Johns Hopkins University
Department of Computer Science
Multiple Tenure-Track Faculty Positions
(all levels)

With the anticipated opening of our new building, Malone Hall, the Department of Computer Science at **The Johns Hopkins University** is planning for substantial multi-year growth. We are currently seeking applications for **multiple tenure-track faculty positions at all levels**. The search is open to all areas of Computer Science. We particularly encourage candidates with research interests in algorithms, information security, natural language processing, and machine learning. We also welcome applicants who would enhance our institutional strengths with a focus on data-intensive computing, computational biology, and health-related applications of computing.

All applicants must have a Ph.D. in Computer Science or a related field and will be expected to establish a strong, independent, multidisciplinary, internationally recognized research program. Commitment to quality teaching at the undergraduate and graduate levels is required. The department is committed to building a diverse educational environment; women and minorities

are especially encouraged to apply.

A more extensive description of our search and additional supporting information can be found at <http://www.cs.jhu.edu/Search2013>. More information on the department is available at <http://www.cs.jhu.edu>.

Applicants should apply using the online application which can be accessed from <http://www.cs.jhu.edu/apply>. Applications should be received by Dec 15, 2012 for full consideration. Questions should be directed to fsearch@cs.jhu.edu. The Johns Hopkins University is an EEO/AA employer.

Faculty Search
Johns Hopkins University
Department of Computer Science
Room 224 New Engineering Building
Baltimore, MD 21218-2694
Fax: 410-516-6134
Phone: 410-516-8775
fsearch@cs.jhu.edu
<http://www.cs.jhu.edu/apply>

Massachusetts Institute of Technology
Faculty Positions

The Department of Electrical Engineering and Computer Science (EECS) seeks candidates for faculty positions starting in September 2013. Appointment will be at the assistant or untenured associate professor level. In special cases, a senior faculty appointment may be possible. Faculty duties include teaching at the graduate and undergraduate levels, research, and supervision of student research. We will consider candidates with backgrounds and interests in any area of electrical engineering and computer science. Faculty appointments will commence after completion of a doctoral degree.

Candidates must register with the EECS search website at <http://eeecs.mit.edu/ACM>, and must submit application materials electronically to this website. Candidate applications should include a description of professional interests and goals in both teaching and research. Each application should include a curriculum vita and the names and addresses of three or more individuals who will provide letters of recommendation. Letter writers should submit their letters directly to MIT, preferably on the website or by mailing to the address below. Please submit a complete application by December 15, 2012.

Send all materials not submitted on the website to:

Professor Anantha Chandrakasan
Department Head, Electrical Engineering
and Computer Science
Massachusetts Institute of Technology
Room 38-401
77 Massachusetts Avenue
Cambridge, MA 02139

M.I.T. is an equal opportunity/affirmative action employer.

Mississippi State University
Faculty Position in Computer Science
and Engineering

The Department of Computer Science and Engineering (<http://www.cse.msstate.edu>) is seeking to fill an open position for a tenure-track faculty mem-

ber at the Assistant/Associate Professor levels. Evidence of strong potential for excellence in research (including the ability to attract external funding) and teaching at the graduate and undergraduate levels is required. The primary research areas of interest for this position are artificial intelligence, bioinformatics, computer security, and computational science.

Mississippi State University has approximately 1300 faculty and 20,000 students. The Department of Computer Science and Engineering has 17 tenure-track faculty positions and offers academic programs leading to the bachelor's, master's and doctoral degrees in computer science and bachelor's degrees in software engineering and computer engineering. Faculty members and graduate students work with a number of on-campus research centers. Department research expenditures totaled approximately \$5.9M in FY11.

Candidates for this position are expected to hold a PhD in computer science or closely related field (ABDs may be considered). Level of appointment is commensurate with qualifications and experience.

Applicants must apply on-line <http://www.jobs.msstate.edu/> and complete a Personal Data Information Form. A letter of application, curriculum vita, teaching statement, research statement, and names and contact information of at least three references must also be submitted. Review of applications will not begin earlier than December 2012 and continue until the position is filled. MSU is an AA/EEO.

New York University
Faculty Openings

The department expects to have several regular faculty positions beginning in September 2013 and invites candidates at all levels. We will consider outstanding candidates in any area of computer science.

Faculty members are expected to be outstanding scholars and to participate in teaching at all levels from undergraduate to doctoral. New appointees will be offered competitive salaries and startup packages, with affordable housing within a short walking distance of the department. New York University is located in Greenwich Village, one of the most attractive residential areas of Manhattan.

The department has 34 regular faculty members and several clinical, research, adjunct, and visiting faculty members. The department's current research interests include algorithms, cryptography and theory; computational biology; distributed computing and networking; graphics, vision and multimedia; machine learning; natural language processing; scientific computing; and verification and programming languages.

Collaborative research with industry is facilitated by geographic proximity to computer science activities at AT&T, Google, IBM, Bell Labs, NEC, and Siemens.

Please apply at <https://cs.nyu.edu/webapps/facapp/register>

To guarantee full consideration, applications should be submitted no later than December 1, 2012; however, this is not a hard deadline, as all candidates will be considered to the full extent feasible, until all positions are filled. Visiting positions may also be available.

New York University is an equal opportunity/affirmative action employer.

Northern Arizona University Faculty Position in Computer Science

Applications for a tenure-track **Assistant Professor** in Computer Science (CS) at Northern Arizona University (NAU) will be accepted until the position is filled or closed, with screening beginning Dec. 1st.

Minimum qualifications for this position include an earned Ph.D. in CS or related specialty, conferred by Aug. 19, 2013.

Preferred qualifications include a strong background in one or more of the following specialty areas: Computer Security, Virtual Worlds and Gaming, Distributed Computing or Software Testing and Analysis; demonstrated skill in applied software development and application, particularly in the above specialty areas; an established record of scholarly success, including leadership in research funding efforts; demonstrated skill in or commitment to pedagogy and undergraduate education; demonstrated skill in or commitment to teaching core courses in Computer Science; outstanding oral and written communication skills; and, candidates with a demonstrated commitment supportive of the multicultural needs of Northern Arizona University and the surrounding area.

Job duties include teaching, research scholarship and funding, and service.

To apply, email: (1) a letter describing your qualifications for this position; (2) a curriculum vitae; and (3) contact information for three references to engineering@nau.edu with subject: CS Assistant Professor Search.

Please see www.nau.edu/hr for complete job description.

Ohio Northern University (ONU) Computer Science Computer Engineering Faculty Positions (Assistant, Associate, or Professor level) Computer Science and Computer Engineering Faculty Positions

Ohio Northern University (ONU) is proud to be a place where the learning, development, and welfare of its students are the highest priorities of the institution. ONU offers programs and experiences that prepare graduates to excel in a competitive global economy, while nurturing values and character traits that make graduates good citizens of a diverse world.

The T. J. Smull College of Engineering offers an environment which encourages the development of new initiatives and opportunities for its students, gives faculty and staff a sense of worth and job satisfaction, and offers faculty the chance to make a difference by educating the professionals of tomorrow. The mission of the college is to engage students through personal relationships and balanced educational experiences to maximize their success. If you would like to be part of a professional and collaborative team devoted to provide a dynamic environment where faculty members can make a significant and lasting impact on the lives of young engineers and computer scientists, the Department of Electrical & Computer Engineering and Computer Science at Ohio Northern University is pleased to invite you to apply for a **faculty position** to begin August 15, 2013.

There are two open positions; the appointment to either of the open positions may be made at the **Assistant, Associate, or Professor level**, commensurate with qualifications and experience:

Computer Science

Candidates with expertise in one or more of the following areas will be given preference: computer security, game development, mobile computing, and robotics. A Ph.D. in Computer Science by date of employment is required.

Computer Engineering

Candidates with expertise in one or more of the following areas will be given preference: embedded systems, computer architecture and/or VLSI. A Ph.D. in Computer Engineering or a related field by date of employment is required.

The applicant must be committed to teaching excellence in undergraduate education and must possess excellent verbal and written communication skills. Expectations include actively pursuing scholarly research and professional development opportunities.

All application materials must be submitted online at <https://jobs.onu.edu> and must include an application letter, vita, statements of teaching and research experience, transcripts, and the names and contact information of three professional references. The search will continue until the position is filled. Questions concerning the position should be referred to: Dr. Khalid Al-Olimat, P.E., Professor and Chair, ECCS Department, Ohio Northern University, Ada, OH 45810, k-al-olimat@onu.edu.

The University is located in Ada, a small community 15 miles east of Lima and approximately 75 miles from Toledo, Dayton, and Columbus. Further information about the University is available at <http://www.onu.edu>.

Ohio State University Tenure-Track Positions

The **Computer Science and Engineering Department at the Ohio State University** seeks faculty candidates for multiple tenured or tenure-track appointments **at the assistant, associate or full professor level**. The department is slated for significant growth as part of a multi-year expansion in size and scope, including targeted faculty hires in core areas as well as in areas that bridge CS with other disciplines. The specific searches being conducted this year include:

► One targeted position, open rank, in the area of big data, broadly defined. These include, but are not limited to, applicants that have core training in databases, systems, data analytics, or cloud computing, with experience in managing and analyzing big data stores.

► One targeted position, open rank, in the area of cybersecurity, broadly defined. We are specifically interested in applicants with research interests in network security, physical layer/information theoretic security, cyberphysical systems, data privacy, cryptography, or programming language security.

► One targeted position, open rank, at the interface of computer science and mathematics broadly defined. We are specifically interested in researchers working in theoretical computer sci-

ence, scientific computing, or complex systems.

► Multiple targeted positions, open rank, jointly searched with the Department of Biomedical Informatics. Priority consideration will be given to applicants whose research interests lie in the processing of biomedical text (e.g. NLP or text mining) as well as those interested in the multi-scale modeling and visualization of high-throughput and/or high-content molecular or image data. Outstanding applicants in other areas at the intersection of biology, clinical science, informatics, and computer science will also be considered.

► In addition to the above targeted searches the department is conducting a search for an additional position (open rank) that is open to all areas of computer science and engineering.

The department is committed to enhancing faculty diversity; women, minorities, and individuals with disabilities are especially encouraged to apply.

Applicants should hold or be completing a Ph.D. in CSE or a closely related field, have a commitment to and demonstrated record of excellence in research, and a commitment to excellence in teaching.

To apply, please submit your application via the online database. The link can be found at: <https://www.cse.ohio-state.edu/cgi-bin/portal/fsearch/apply.cgi>.

Review of applications will begin in November and will continue until the positions are filled.

Peking University School of EECS Tenure-Track Positions

The School of EECS at Peking University invites applications for tenure-track positions in the areas of energy efficient computing (including but not limited to energy-efficient architectures, communication, compilation, and system software) and applications (such as smart grid, mobile computing, sensor networks, and hardware acceleration of computing-intensive applications). These positions are associated with the Center for Energy-Efficient Computing and Applications (<http://ceca.pku.edu.cn>), which offers a new level of startup and compensation packages. Applications from distinguished candidates at senior levels are also encouraged. To apply, please email the resume, statements of research and teaching, and at least three names for references to ceca_recruiting@pku.edu.cn. Applications received by January 15, 2013 will be given full consideration. Early submissions are encouraged (the first set of interviews will be in December 2012).

Princeton University Computer Science 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 re-

lated area is required. Candidates should expect to receive their PhD before Fall, 2013. 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 be underway by December 2012. 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: 1200555.

Rensselaer Polytechnic Institute Faculty Positions - Computer Science

The Department of Computer Science at Rensselaer Polytechnic Institute, Troy NY seeks to hire two faculty members to join a strong and growing faculty.

The first opening, the Hamilton Chair in Computer Science, is intended for an associate professor in the area of "cyber risk," including but not limited to information/data security, privacy, accountability, trust, and forensics for computers, networks and cyber-physical systems. Applicants for this chaired position must demonstrate an outstanding record of research accomplishments as well as a strong commitment to teaching.

The second position, for an assistant professor, is focused on agent-based systems, including but not limited to multi-agent systems, agent learning, collective intelligence, agents in financial markets, agent-based risk assessment, and agent-based modeling, especially for networked and distributed systems.

Applicants for these positions must have a doctoral degree (or foreign degree equivalent) in computer science or in a related field. We seek highly collaborative applicants with strong technical vision and a focus on emerging 21st century technologies and challenges.

As part of the School of Science at Rensselaer, which is undergoing a broad expansion (<http://science.rpi.edu/>), the Department of Computer Science has strong undergraduate, master's and PhD programs involving over 600 total students. The department maintains strong interdisciplinary research efforts, bolstered by the Computational Center for Nanotechnology Innovations, the Data Science Research Center, the Network Science and Engineering Center, and the Tetherless World Constellation.

Qualified applicants must submit statements of research and teaching interests and a curriculum vitae including a list of publications to <https://cgi.cs.rpi.edu/faculty-apply/>. Applicants must also arrange for the submission of three letters of reference. Questions about these positions may be directed to Prof. Jim Hendler (hendler@cs.rpi.edu), Department Head, or Prof. Chuck Stewart (stewart@cs.rpi.edu), Chair of the Faculty Search committee, while questions about the process may be sent to Ms. Sharon Simmons, Administrative Coordinator (simmons2@cs.rpi.edu). Review of candidates is ongoing and applications will be accepted until the positions are filled. Applications received by December 1, 2012 are assured full consideration.

Successful candidates will have duties that include teaching graduate and undergraduate courses in the department, developing and maintaining robust programs of research and scholarship, as well as service to the department, the School of Science and to Rensselaer.

We welcome candidates who will bring diverse intellectual, geographical, gender and ethnic perspectives to Rensselaer's work and campus communities. Rensselaer Polytechnic Institute is an Affirmative Action/Equal Opportunity Employer.

St. Lawrence University Assistant Professor, Tenure Track

St. Lawrence University invites applications for a tenure-track position in computer science, at the rank of Assistant Professor, to begin in August 2013. We seek a colleague with a Ph.D. in computer science, an enthusiastic commitment to undergraduate teaching, and expertise in any research area.

The department of Mathematics, Computer Science & Statistics has 14 faculty members, three of whom are computer scientists. Our classes are small, taught in labs dedicated to CS education, and emphasize active student participation.

Review of applications will begin on November 15, 2012. Applications should include a CV, a letter of introduction, and a statement of teaching philosophy. Please arrange for three letters of reference, at least one of which should address teaching. Applications and letters may be sent via email to cssearch@stlawu.edu or via regular mail to: Computer Science Search Committee, Department of Math, CS & Stats, St. Lawrence University, Canton NY 13617.

For additional information about St. Lawrence, please visit SLU's homepage at <http://www.stlawu.edu>. SLU is an Affirmative Action/Equal Employment Opportunity employer.

Stanford University Department of Computer Science Faculty Openings

The Department of Computer Science at Stanford University invites applications for tenure-track faculty positions at the junior level (Assistant or untenured Associate Professor). We give higher priority to the overall originality and promise of the candidate's work than to the candidate's sub-area of specialization within Computer Science.

We are seeking applicants from all areas of Computer Science, spanning theoretical foundations, systems, software, and applications. We are also interested in applicants doing research at the frontiers of Computer Science with other disciplines, especially those with potential connections to Stanford's main multidisciplinary initiatives: Energy, Human Health, Environment and Sustainability, the Arts and Creativity, and the International Initiative. Interdisciplinary candidates whose research combines other fields of engineering or mathematics with computer science may be considered for a joint appointment in the Institute for Computational and Mathematical Engineering (<http://icme.stanford.edu/>).

Applicants must have completed (or be completing) a Ph.D., must have demonstrated the

ability to pursue a program of research, and must have a strong commitment to graduate and undergraduate teaching. A successful candidate will be expected to teach courses at the graduate and undergraduate levels, and to build and lead a team of graduate students in Ph.D. research. Further information about the Computer Science Department can be found at <http://cs.stanford.edu>. The School of Engineering website may be found at <http://soe.stanford.edu>.

Applications should include a curriculum vita, brief statements of research and teaching interests, and the names and contact information of at least four references. Please apply online at: <http://soe-apps.stanford.edu/FacultyApplyCS>

Questions should be directed to, Search Committee Chair, c/o Laura Kenny-Carlson, via electronic mail to search@cs.stanford.edu.

The review of applications will begin on November 16, 2012, and applicants are strongly encouraged to submit complete applications by that date for full consideration; however, applications will continue to be accepted until February 15, 2013.

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 to the university's research and teaching missions.

Université Catholique de Louvain (UCL) Professor - Large Scale and Cloud Computing

UCL seeks a full-time professor in Computer Science. The successful candidate will carry out research in the field of large scale and cloud computing, including but not limited to distributed computing, operating systems, data mining and optimization, software engineering and large scale software systems. Still, other areas of competence will also be considered, since qualifications take precedence over specialization.

Responsibilities include: build a high-quality research programme through acquisition of funding, publications, and supervision of Master and PhD students; deliver and develop innovative courses. Undergraduate-level courses are taught in French and graduate-level ones in English. PhD in Computer Science or equivalent.

UCL is Belgium's largest French-speaking university near Brussels, the capital of Belgium, in the heart of Europe. Application deadline is December 15th, 2012 via <http://www.uclouvain.be/en-icteam-jobs>

University at Buffalo, The State University of New York Faculty Positions in Computer Science and Engineering

The CSE Department invites excellent candidates in all core areas of Computer science and Engineering, especially software and hardware systems areas to apply for openings at the assistant professor level. The department is affiliated with successful centers devoted to biometrics, bioinformatics, biomedical computing, cognitive science, document analysis and recognition, high performance computing, and information assurance.

Candidates are expected to have a Ph.D. in Computer Science/Engineering or related field by August 2013, with an excellent publication record and potential for developing a strong funded research program.

Applications should be submitted by December 31, 2012 electronically via <http://www.ubjobs.buffalo.edu/>.

The University at Buffalo is an Equal Opportunity Employer/Recruiter.

University of California, Irvine
Department of Computer Science and
California Institute for Telecommunications
and Information Technology

The Department of Computer Science and the California Institute for Telecommunications and Information Technology (Calit2) at the University of California, Irvine (UC Irvine) have an opening for a tenure-track Assistant Professor in the area of computer networking. We are soliciting applications in all areas of networking, with particular interest in: (1) network security and privacy, (2) wireless and mobile networking, as well as (3) network gaming. Exceptionally qualified more senior candidates may also be considered.

The Department of Computer Science is the largest department in the Donald Bren School of Information and Computer Sciences, one of only a few such schools in the nation. The department has over 45 faculty members and over 200 graduate students. Faculty research is vibrant and broad, spanning networking, security, multimedia, distributed systems, operating systems, software, databases, embedded systems, theory, graphics, machine learning, artificial intelligence, and bioinformatics.

The California Institute for Telecommunications and Information Technology is a multi-disciplinary research institute at UC Irvine and UC San Diego. Calit2 conducts cutting-edge research in diverse fields to develop innovative information technology-based products and services that benefit society and ignite economic development in the region and state. Close ties with industry facilitate creative strategies to improve technology transfer, speeding downstream commercialization of research discoveries. More than 100 companies have become Calit2 partners.

One of the youngest UC campuses, UC Irvine is ranked 13th among the nation's best public universities by US News & World Report. Compensation is competitive with the nation's finest universities, and includes priority access to on-campus for-sale faculty housing. UC Irvine is located 4 miles from the Pacific Ocean and 45 miles south of Los Angeles. The area offers a very pleasant year-round climate, numerous recreational and cultural opportunities, and one of the highest-ranked public schools systems in the nation. Prospective applicants are invited to visit our webpages at <http://www.cs.uci.edu> and <http://www.calit2.uci.edu>.

Screening will begin immediately upon receipt of a completed application. Applications will be accepted until the position is filled, although maximum consideration will be given to applications received by January 1, 2013. Each application must contain: a cover letter, CV, up to 3 key publications, a statement of research and teaching interests, and 3-5 letters of recommen-

dation. All materials must be uploaded at <https://recruit.ap.uci.edu>.

UC Irvine is an equal opportunity employer committed to excellence through diversity and encourages applications from women, minorities, and other under-represented groups. UC Irvine is responsive to the needs of dual career couples, is dedicated to work-life balance through an array of family-friendly policies, and is the recipient of an NSF Advance Award for gender equity.

University of California, San Diego
Department of Computer Science and
Engineering
Tenure-track or Tenured Faculty
Lecturer
Visiting Professor

Tenure-track or Tenured Faculty: The UCSD Department of Computer Science and Engineering (CSE) seeks to fill multiple tenured or tenure-track faculty positions for Fall 2013. Exceptional candidates in all areas will be seriously considered. We especially encourage candidates working in the areas of parallel software, data-intensive systems, quantitative biology, medical informatics, energy-efficient systems, and sustainable computing to apply. The rank of the positions is at the Assistant Professor level; however, excellent candidates at all levels will be seriously considered. The department is looking for applicants with outstanding research credentials. A Ph.D. in computer science or a related area is desired.

Lecturer with (Potential) Security of Employment: The CSE Department also seeks applications for a Lecturer with Potential Security of Employment (Lecturer PSOE, which parallels a tenure-track assistant professor position) or Lecturer with Security of Employment (Lecturer SOE, which parallels a tenured professorial position). Successful candidates will be outstanding educators and should provide evidence of effective and innovative undergraduate teaching in computer science and engineering. In addition to teaching core courses, candidates are expected to lead the development and assessment of new educational initiatives, including applying for grants related to education. The successful candidate will provide guidance, leadership, and innovation for the CSE undergraduate programs. Candidates are expected to have a Ph.D. degree in computer science or a related area.

Visiting Professor: The CSE Department intends to appoint one or more Visiting (Assistant or Associate or Full) Professors, beginning Fall 2013 or earlier. Successful candidates are expected to teach undergraduate students and conduct research. These are two-year terminal positions, although candidates who prefer a one-year appointment will also be considered. Candidates at all ranks from new Ph.D.s to senior faculty, including sabbatical visitors, and in all areas of computer science, will be considered. Successful applicants should be able to provide evidence of excellence in teaching and of a promising research career.

The CSE Department is committed to building an excellent and diverse faculty, staff and student body. Women and minority applicants, veterans and persons with disabilities are encour-

aged to apply (see <http://diversity.ucsd.edu>). **In addition to the highest standards of scholarship, teaching, and professional activity, the preferred candidates for any position will have potential or demonstrated contributions to a climate that supports equity, inclusion, and diversity.**

We encourage candidates to send applications as soon as possible. Applications received by January 1, 2013 will be given full consideration. However, positions remain open until filled. To apply, follow the instructions at the website <https://csefacapp.ucsd.edu/applicant>

UCSD is an equal opportunity / affirmative action employer.

University of California, Santa Barbara
Faculty Position in Computer Science

The Department of Computer Science at the University of California, Santa Barbara, invites applications for a tenure-track position effective July 2013. We are particularly interested in outstanding candidates in the areas of applied cryptography and system security; however, exceptional candidates in all areas of computer science will be considered.

The Department of Computer Science has grown rapidly, both in size and stature, over the past 10 years, accompanied by a five-fold increase in extramural funding. The department, with 30 faculty and more than 100 doctoral students, is part of the College of Engineering, which is ranked among the top 20 in the Nation by the 2008 US News and World Report. The PhD program of the Department of Computer Science has recently been ranked among the top 10 departments in the nation by the National Research Council (NRC).

Additional information about the department and our graduate program can be found at <http://www.cs.ucsb.edu>. Applicants are expected to hold a doctoral degree in Computer Science or related fields, show outstanding research potential, and have a strong commitment to teaching.

Primary consideration will be given to applications received by December 14, 2012; however, the position will remain open until filled. Applications should be submitted electronically as PDF documents to: <https://www.cs.ucsb.edu/recruit/faculty>. Applications must include a detailed resume, research and teaching statements, and the names and addresses of four references.

The Department is especially interested in candidates who can contribute to the diversity and excellence of the academic community through research, teaching, and service. We are an Equal Opportunity/Affirmative Action employer.

University of California, Santa Barbara
Department of Electrical and Computer
Engineering
Tenure-track Faculty Position

The Electrical and Computer Engineering Department at the University of California, Santa Barbara invites applications for a tenure-track faculty position in the area of computer engineering with a start date of Fall quarter, 2013. For more details please visit - <http://www.ece.ucsb.edu/employment/>. An EO/AA employer.

The University of Chicago

Associate Professor

Req # 01455

The Department of Computer Science at the University of Chicago invites applications from exceptionally qualified candidates in the area of systems for faculty positions at the rank of Associate Professor. Systems is a broad, synergistic collection of research areas spanning systems and networking, programming and architecture, data-intensive computing and databases, graphics and visualization, and systems biology.

The University of Chicago has the highest standards for scholarship and faculty quality, and encourages collaboration across disciplines. We encourage strong connections with researchers across the campus in such areas as mathematics, natural language processing, bioinformatics, logic, molecular engineering, and machine learning, to mention just a few.

Applicants must be several years beyond the PhD and have an outstanding research record in a relevant field. The PhD should be in Computer Science or a related field such as Mathematics or Statistics.

The Department of Computer Science (cs.uchicago.edu) is the hub of a large, diverse computing community of two hundred researchers focused on advancing foundations of computing and driving its most advanced applications. Long distinguished in theoretical computer science and artificial intelligence, the Department is now building a strong Systems research group. This closely-knit community includes the Toyota Technological Institute, the Computation Institute, and Argonne's Mathematics and Computer Science Division.

The Chicago metropolitan area provides a diverse and exciting environment. The local economy is vigorous, with international stature in banking, trade, commerce, manufacturing, and transportation, while the cultural scene includes diverse cultures, vibrant theater, world-renowned symphony, opera, jazz, and blues. The University is located in Hyde Park, a Chicago neighborhood on the Lake Michigan shore just a few minutes from downtown on an electric commuter train.

All applicants must apply through the University's Academic Jobs website at

academiccareers.uchicago.edu/applicants/Central?quickFind=52485

A cover letter, curriculum vitae including a list of publications, a statement describing past and current research accomplishments and outlining future research plans, and a description of teaching experience must be uploaded to be considered as an applicant.

Candidates may also post a representative set of publications, as well as teaching evaluations, to this website. Three reference letters are required, one of which must address the candidate's teaching ability. The reference letters can be sent by mail to:

Chair, Department of Computer Science
The University of Chicago
1100 E. 58th Street, Ryerson Hall
Chicago, IL. 60637-1581

Or by email to: Recommend@mailman.cs.uchicago.edu (letters can be in pdf, postscript or Microsoft Word).

To ensure fullest consideration of your application all materials, including supporting letters, should be received by November 19. However, screening will continue until all available positions are filled.

The University of Chicago is an Affirmative Action / Equal Opportunity Employer.

University of Colorado Denver Computer Science and Engineering Department

Two Assistant Professor Positions

Two Assistant Professor positions, Computer Science and Engineering Department, University of Colorado Denver beginning Spring or Fall 2013. A Ph.D. in computer science, computer engineering or a closely related field is required. Successful candidates will develop a strong research and educational program at undergraduate, MS, and Ph.D. levels and will develop an externally funded research program. For details and to apply electronically, please visit www.jobsatcu.com posting #818654.

UMBC

University of Maryland Baltimore County An Honors University in Maryland Information Systems Department

The Information Systems Department at UMBC invites applications for two tenure-track faculty position at the **Assistant Professor** level in the area of human-centered computing as well as the area of intelligent information systems starting August 2013. Candidates must have earned a PhD in Information Systems or a related field no later than August 2013.

For the human-centered computing position, the primary research areas of interest are social computing, computer-supported cooperative work, social informatics, interaction design, and usability. For the area of intelligent information systems position, individuals should be engaged in state of the art work in machine learning, statistical natural language processing, personalization, mobile computing, knowledge representation, and information extraction, are especially encouraged to apply. Secondary research interests in Health IT or cybersecurity are desirable, but not mandatory for both positions. Ideal candidates will be engaged in research that spans two or more of these areas with preference given to those who can collaborate with current faculty. Candidates for both positions should have a strong potential for excellence in research, the ability to develop and sustain an externally funded research program, and the ability to contribute to our graduate and undergraduate teaching mission.

The Department offers undergraduate degrees in Information Systems and Business Technology Administration. Graduate degree programs, MS and PhD, are offered in both Information Systems and Human-Centered Computing, including an innovative online MS in IS program. Consistent with the UMBC vision, the Department has excellent teaching facilities, state-of-the-art laboratories, and outstanding technical support. UMBC's Technology Center, Research Park, and Center for Entrepreneurship

are major indicators of active research and outreach. Further details on our research, academic programs, and faculty can be found at <http://www.is.umbc.edu/>. Members of under-represented groups including women and minorities are especially encouraged to apply.

Applications will not be reviewed until the following materials are received: a cover letter, a one-page statement of teaching interests, a one to two-page statement of research interests, one or more sample research papers, and a CV. Applicants should also arrange to have three letters of recommendation sent to the department as soon as possible. Electronic submission of materials as PDF documents is preferred. Electronic copies should be sent to bmorris@umbc.edu. Copies can also be sent to: Dr. Arya Gangopadhyay, Chair of Faculty Search Committee, Information Systems Department, UMBC, 1000 Hilltop Circle, Baltimore, MD 21250-5398. For inquiries, please contact Barbara Morris at (410) 455-3795 or bmorris@umbc.edu.

Review of applications will begin immediately and will continue until the position is filled. This position is subject to the availability of funds.

UMBC is an Affirmative Action/Equal Opportunity Employer and welcomes applications from minorities, women and individuals with disabilities.

The University of Michigan, Ann Arbor Department of Electrical Engineering and Computer Science Computer Science and Engineering Division Faculty Positions

Applications and nominations are solicited for multiple faculty positions in the Computer Science and Engineering (CSE) Division. Highly qualified candidates from all areas of computer science and computer engineering will be considered for positions at all levels and are encouraged to apply. Particularly encouraged are applicants with research interests in the following areas.

1. Software systems, including databases, distributed systems, networking, and security.
2. Scalable parallel computing, including high-performance computing, compilers, programming languages, multi-core systems, and algorithms for big data.
3. Medical computing, including machine learning, big data, probabilistic reasoning, and visualization approaches to medicine and, more broadly, healthcare.
4. Computing and Media, including human-computer interaction, systems, and machine learning approaches to social networks and data, visualization, vision, music, video, and photography.

Qualifications include an outstanding academic record, a doctorate or equivalent in computer science or computer engineering or a discipline relevant to the above areas, and a strong commitment to teaching and research. Applications must be received by January 1, 2013.

To apply, please complete the form at:
<http://www.eecs.umich.edu/eecs/jobs/cse-jobs.html>

Electronic applications are strongly preferred, but you may alternatively send resume, teaching statement, research statement and names of three references to:

ACM Transactions on Reconfigurable Technology and Systems



This quarterly publication is a peer-reviewed and archival journal that covers reconfigurable technology, systems, and applications on reconfigurable computers. Topics include all levels of reconfigurable system abstractions and all aspects of reconfigurable technology including platforms, programming environments and application successes.

www.acm.org/trets
www.acm.org/subscribe



Professor Satinder Singh Baveja, Chair, CSE
Faculty Search
Department of Electrical Engineering and
Computer Science
University of Michigan
2260 Hayward Street
Ann Arbor, MI 48109-2121

The University of Michigan is a Non-Discriminatory/Affirmative Action Employer with an Active Dual-Career Assistance Program. The college is especially interested in candidates who can contribute, through their research, teaching, and/or service, to the diversity and excellence of the academic community.

University of Nevada Reno Computer Science and Engineering

CSE at UNR invites applications for a **tenure-track assistant professor** faculty position starting July 1, 2013. More information can be found www.cse.unr.edu. Candidates with interest and expertise in big data/cloud computing, embedded systems, or computer games may be given preference. Applicant should be strongly committed to quality research and teaching, expect to develop a robust externally funded research program, supervise MS and PhD students, and participate in service and outreach. Review of applications will begin on January 15, 2013. To apply <https://www.unrsearch.com/postings/11551>. CSE at UNR is also recruiting for a full time lecturer position starting on January 1, 2013 or later. <https://www.unrsearch.com/postings/11514>. EEO/AA

University of North Carolina Wilmington Computer Science Assistant Professor, Tenure-track Vacancy 13F008

Starts August 2013. Ph.D. in Computer Science or closely related area required. Emphasis in computer graphics, visualization, animation or closely related area. **Details at** <http://uncw.edu/hr/employment-epa.html> **Priority consideration date: January 2, 2013.** EEO/AA Employer. Women and Minorities encouraged to apply.

University of Rochester Faculty Positions in Computer Science: HCI and Big Data

The University of Rochester Department of Computer Science seeks applicants for multiple tenure track positions in **human-computer interaction (HCI)** and **big data research** (including machine learning and data mining, cloud computing, e-science applications, and very large databases) Candidates must have a PhD in CS or a related discipline. Applicants for the big data position will also be considered for a faculty search in that area by our Department of Electrical & Computer Engineering. Additional information and online application instructions appear at <http://www.cs.rochester.edu/dept/recruit>.

The Department of Computer Science is a research-oriented department with a distinguished history of contributions in systems,

theory, artificial intelligence, and HCI. We have a collaborative culture and strong ties to cognitive science, linguistics, and ECE. Over the past decade, a third of its PhD graduates have won tenure-track faculty positions, and its alumni include leaders at major research laboratories such as Google, Microsoft, and IBM.

The University of Rochester is a private, Tier I research institution located in western New York State. The University of Rochester consistently ranks among the top 30 institutions, both public and private, in federal funding for research and development. Teaching loads are light and classes are small. Half of its undergraduates go on to post-graduate or professional education. The university includes the Eastman School of Music, a premiere music conservatory, and the University of Rochester Medical Center, a major medical school, research center, and hospital system. The greater Rochester area is home to over a million people, including 80,000 students who attend its 8 colleges and universities.

The University of Rochester has a strong commitment to diversity and actively encourages applications from candidates from groups underrepresented in higher education. The University is an Equal Opportunity Employer.

University of Texas at Austin Department of Computer Science Tenured/Tenure-Track Faculty Positions

The Department of Computer Science of the University of Texas at Austin invites applications for tenure-track positions at all levels. Outstanding candidates in all areas of Computer Science will be considered, particularly in Formal Methods, Big Data, and Robotics. All tenured and tenure-track positions require a Ph.D. or equivalent degree in computer science or a related area at the time of employment.

Successful candidates are expected to pursue an active research program, to teach both graduate and undergraduate courses, and to supervise graduate students. The department is ranked among the top ten computer science departments in the country.

It has 42 tenured and tenure-track faculty members across all areas of computer science. Many of these faculty participate in interdisciplinary programs and centers in the University, including the Texas Advanced Computing Center (TACC), and those in Computational and Applied Mathematics, Computational Biology, and Neuroscience.

Austin, the capital of Texas, is a center for high-technology industry, including companies such as IBM, Dell, Freescale Semiconductor, Advanced Micro Devices, National Instruments, AT&T, Intel and Samsung. For more information please see the department web page: <http://www.cs.utexas.edu/>

The department prefers to receive applications online, beginning September 1, 2012. To submit yours, please visit <http://www.cs.utexas.edu/faculty/recruiting>

Applicants for an assistant professor position must have at least three (3) referees send letters of reference directly to the address provided. Applicants for a tenured position (associate or full professor) must have at least six (6) referees send letters of reference directly.

Inquiries about your application may be directed to faculty-search@cs.utexas.edu. For full consideration of your application, please apply by January 31, 2013. Women and minority candidates are especially encouraged to apply. The University of Texas is an Equal Opportunity Employer.

Vanderbilt University

All ranks considered with a preference for Assistant Professor level

THE DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE (EECS) AT VANDERBILT UNIVERSITY is seeking candidates for two tenured/tenure-track (T/TK) faculty positions in Computer Science. Appointments at all ranks will be considered, with a preference for candidates at the assistant professor level. The CS program has vibrant collaborative and interdisciplinary research efforts in artificial intelligence, computer animation and virtual environments, cyber-physical systems, distributed real-time and embedded middleware, human-systems integration, image processing, intelligent learning environments, robotics, model-integrated computing, software engineering, and trustworthy computing. Average T/TK faculty funding is ~\$700k per year from NSF, NIH, DARPA, DoD, NASA, and many industry sponsors. All junior faculty members hired during the past decade have received prestigious young investigator awards, such as NSF CAREER and DARPA CSSG.

Successful candidates are expected to teach at the undergraduate and graduate levels and to develop vigorous programs of externally funded research. We are seeking candidates who will strengthen the CS program's existing research areas and enhance the School of Engineering's strategic directions in health care and medicine, security, energy and natural resources, and entertainment.

Vanderbilt University is a private, internationally renowned research university located in vibrant Nashville, Tennessee. Its 10 schools share a single cohesive campus that nurtures interdisciplinary activities. The School of Engineering currently comprises 85 tenured and tenure-track faculty, operates with an annual budget of \$100 million including \$60 million of externally funded research, and teaches 1,300 undergraduate and over 400 graduate students. The School of Engineering over the past decade has been on a strong upward trajectory in national and international stature and prominence, and is in the process of building infrastructure to support a significant expansion in faculty size over the next five years. In the 2013 rankings of graduate engineering programs by *U.S. News & World Report*, the School ranks 5th among programs with fewer than 100 faculty. Vanderbilt University is an equal-opportunity, affirmative-action employer. We strongly encourage applications from women and members of under-represented minority groups. Applications should be submitted on-line at: <https://academicjobsonline.org/ajo?apply-243-1851>. For more information, please visit our web site: <http://eeecs.vuse.vanderbilt.edu>. Applications will be reviewed on a rolling basis, with an initial deadline of November 18, 2012, but will be accepted until all positions are filled.

Washington University in St. Louis Faculty Positions

The Department of Computer Science & Engineering at Washington University in St. Louis seeks outstanding tenure-track faculty in all areas of computer science and engineering at the assistant professor level. Exceptional candidates at the associate and full professor levels will also be considered. The department plans to grow its faculty size by 50% in the coming years. We seek multiple talented and highly motivated individuals who will build transformative research programs, both through work in the core disciplines of computer science and computer engineering and through interdisciplinary collaborations with researchers in areas such as biomedicine, engineering, and the sciences. Successful candidates must show exceptional promise for research leadership and a strong commitment to high-quality teaching. Candidates will be expected to publish their research in peer-reviewed journals, to teach, and to participate in department and University service. For full information about this search and application instructions, please visit: <http://cse.wustl.edu/aboutthedept/Pages/OpenFacultyPositions.aspx>. Applicants should hold a doctorate in Computer Science, Computer Engineering, or a closely related field. Washington University in St. Louis is an Equal Opportunity and Affirmative Action employer and invites applications from all qualified candidates. Employment eligibility verification required upon employment.

York University Software Engineering

York University, Toronto, Canada seeks an outstanding candidate in **Software Engineering** to commence on July 1, 2013. The position calls for research excellence in Software Engineering with an emphasis on mission critical systems, dependable safety critical systems, industrial strength formal methods for software systems, high assur-

ance business and mobile systems, and rigorous methods for verifying, validating and certifying software systems. Outstanding candidates in all areas of software engineering are invited to apply. The rank is open and commensurate with experience. This position will play a key role in the development of the software engineering program within the Lassonde School of Engineering.

York University offers a world-class, modern, interdisciplinary academic experience in Toronto, Canada's most multicultural city. York is at the centre of innovation, with a thriving community of 62,000 students, faculty and staff, as well as over 250,000 alumni worldwide. York's 11 Faculties and 28 research centres are committed to providing an engaged learning and research environment that cuts across traditional academic boundaries. The Lassonde School of Engineering currently offers fully accredited and innovative programs in Computer Engineering, Geomatics Engineering, and Space Engineering. We are currently expanding with new programs in Software Engineering, Electrical Engineering, Mechanical Engineering, Civil Engineering, and Chemical Engineering. The Software Engineering position will be in York's Department of Computer Science and Engineering (to be renamed Department of Electrical Engineering and Computer Science) which is a leading academic and research department in Canada with 45 research-active faculty members, offering a range of undergraduate programs in Computer Science, Computer Engineering, Software Engineering, Digital Media, and Computer Security, as well as research intensive MSc and PhD degrees in Computer Science and Engineering. Applications must be received by November 15, 2012 along with a CV, statement of contribution to research, teaching, and curriculum development, and three reference letters at: Chair, Search Committee for Software Engineering, Lassonde School of Engineering, York University, 4700 Keele Street, Toronto, ON, Canada M3J 1P3, Tel: (416) 650-8135, Email: eng@yorku.ca. For further details, please visit the URL <http://www.yorku.ca/acadjobs>. All York University positions are subject to final budgetary approval.



DEPARTMENT HEAD

Nominations and applications are being solicited for the position of Head of the Bradley Department of Electrical and Computer Engineering (ECE) at Virginia Tech. The Department Head's principal responsibility is to provide leadership and management of the programs, faculty, staff, and students. This entails leadership of departmental programs and administrative responsibility for planning, fiscal management, human resources, and communication within the department.

Applicants must have an earned doctorate in electrical or computer engineering or a closely related field with the experience and credentials appropriate for a tenured academic faculty appointment at the professor level at Virginia Tech.

Candidates should have demonstrated leadership and management skills to lead the department to become one of the premier ECE departments in the nation. Candidates must qualify for tenure at the rank of professor and should have achieved distinction in university-level teaching and research and have a record of superior scholarship, administrative ability, and leadership. Previous successful administrative and research leadership experience and Familiarity with ABET engineering accreditation processes preferred.

The ECE Department offers B.S., M.S., M.Eng., and Ph.D. degree programs in both Electrical Engineering and Computer Engineering with an enrollment of about 600 full-time undergraduate students and 500 graduate students. In addition to the main Blacksburg campus, an active research program and a full- and part-time graduate program are located in the National Capital Region (Washington, DC area), with 10 full-time ECE faculty who interact closely with faculty in Blacksburg.

Complete position information, application procedures, and the review process are available at <http://www.ece.vt.edu/>. Please apply online at <https://jobs.vt.edu>, posting **0122299**.

Virginia Tech is an equal opportunity/affirmative action institution.



DOI:10.1145/2366316.2366340

Peter Winkler

Puzzled

Weighed in the Balance


Many of us have pondered puzzles involving a set of n coins and a balance scale, the idea being typically to find the counterfeit coin and determine whether it is lighter or heavier than the others using the fewest possible weighings. Here we take a slightly different tack, but the equipment is familiar: a set of coins and a balance scale that can tell us which of two sets of coins is heavier or that they are of equal weight.

1. You have 13 coins with the property that any 12 of them can be split into two piles of six each that balance perfectly on the scale (see the figure here). Now prove all the coins have the same weight. (Advice: Try this for integer weights first, then rational, then—for the brave—arbitrary positive real weights.)

2. Eight coins have at most two different weights; now show that with three weighings, you can determine whether all the coins have the same weight.

3. Following the same rules as in the second puzzle, now solve it with 10 coins.

A balanced balance scale for Puzzle 1.



Readers are encouraged to submit prospective puzzles for future columns to puzzled@cacm.acm.org.

Peter Winkler (puzzled@cacm.acm.org) is William Morrill Professor of Mathematics and Computer Science at Dartmouth College, Hanover, NH.

Computing Reviews is on the move



Our new URL is

ComputingReviews.com



COMPUTING REVIEWS

A daily snapshot of what is new and hot in computing



SIGGRAPH
ASIA2012

The 5th ACM SIGGRAPH Conference and Exhibition on
Computer Graphics and Interactive Techniques in Asia



Embracing the **digital** convergence

Register now to save

SIGGRAPH Asia offers registration categories to meet your needs and budget.

Enjoy maximum savings by registering online now at <http://sa12registration.siggraph.org>

Register online by 20 November, 23:59 Singapore time to enjoy up to 10% off!

New: recruitment packages

New at the exhibition, SIGGRAPH Asia is offering recruitment packages.

Reach out to creative and talented individuals attending the exhibition through your booth, recruitment talk sessions, and other exclusive package entitlements!

To reserve exhibit space or enquire about recruitment packages, contact:

SIGGRAPH Asia 2012 Exhibition Management

Tel: +65 6500 6725

Email: exhibits_asia@siggraph.org

CONFERENCE 28 Nov - 1 Dec
EXHIBITION 29 Nov - 1 Dec
Singapore EXPO

