Innovative Technology for Computer Professionals

DIG

N

S

DECEMBER 2012

http://www.computer.org

CS PRESIDENT'S MESSAGE, P. 6 COMPUTING CONVERSATIONS: VINT CERF, P. 10 IAAS CLOUD ARCHITECTURE, P. 65



Computer society



allrecipes

Cook up the next great app.

Opportunity doesn't just knock. In the Windows Store it swipes, taps and clicks, too. See how Allrecipes and others are building immersive apps for the new Windows experience and learn how you can put your app in the hands of new users everywhere.

Build for the new Windows Store. Open for business at **windowsstore.com**



Computer Professionals ter

Editor in Chief Ron Vetter

University of North Carolina Wilmington vetterr@uncw.edu

Associate Editor in Chief

Sumi Helal University of Florida helal@cise.ufl.edu

Area Editors

Computer Architectures David H. Albonesi **Cornell University** Tom Conte Georgia Tech Steven K. Reinhardt AMD Grea Byrd North Carolina State University **Graphics and Multimedia Oliver Bimber** Johannes Kepler University Linz **High-Performance Computing** Vladimir Getov University of Westminster Information and **Data Management** Naren Ramakrishnan Virginia Tech **Internet Computing** Simon Shim San Jose State University Multimedia Savitha Srinivasan IBM Almaden Research Center Networking Ahmed Helmy University of Florida Ying-Dar Lin National Chiao Tung University **Security and Privacy Rolf Oppliger** eSECURITY Technologies Software Renée Bryce University of North Texas **Robert B. France** Colorado State University

Associate Editor in Chief, Research Features Kathleen Swigger University of North Texas kathy@cs.unt.edu

Associate Editor in Chief, Special Issues

Bill N. Schilit Google schilit@computer.org

Jean-Marc Jézéquel University of Rennes David M. Weiss Iowa State University

Column Editors

Computing Conversations Charles R. Severance University of Michigan **Discovery Analytics** Naren Ramakrishnan Virginia Tech Education Ann E.K. Sobel Miami University **Entertainment Computing** Kelvin Suna University of Washington, Bothell **Forward Slash** David A. Grier George Washington University Green IT Kirk W. Cameron Virginia Tech **Identity Sciences** Karl Ricanek University of North Carolina, Wilmington In Development **Chris Huntley** Fairfield University Invisible Computing Albrecht Schmidt University of Stuttgart **Out of Band** Hal Berghel University of Nevada, Las Vegas

Computing Practices Rohit Kapur Synopsys rohit.kapur@synopsys.com

Perspectives Bob Colwell bob.colwell@comcast.net

Security Jeffrey M. Voas NIST Social Computing John Riedl University of Minnesota Software Technologies Mike Hinchey Lero—the Irish Software Engineering Research Centre 32 & 16 Years Ago

Neville Holmes University of Tasmania

Advisory Panel

Carl K. Chang Editor in Chief Emeritus Iowa State University Jean Bacon University of Cambridge Hal Berghel University of Nevada, Las Vegas **Doris L. Carver** Louisiana State University **Rick Mathieu** James Madison University Naren Ramakrishnan Virginia Tech **Theresa-Marie Rhyne** Consultant Alf Weaver University of Virginia

Publications Board

Multimedia Editor

Charles R. Severance

2012 IEEE Computer

Society President

iohnwalz@ameritech.net

csev@umich.edu

John W. Walz

Thomas M. Conte (chair), Alain April, David Bader, Angela R. Burgess, Greg Byrd, Jim Cortada, Koen De Bosschere, Hakan Erdogmus, Frank E. Ferrante, Jean-Luc Gaudiot, Linda I. Shafer, Per Stenström, George Thiruvathukal

Magazine Operations Committee

Jean-Luc Gaudiot (chair), Erik R. Altman, Isabel Beichl, Nigel Davies, Lars Heide, Simon Liu, Dejan Milojicic, Michael Rabinovich, Forrest Shull, John Smith, Gabriel Taubin, Ron Vetter, John Viega, Fei-Yue Wang

Editorial Staff

Judith Prow Managing Editor jprow@computer.org Chris Nelson Senior Editor Contributing Editors Camber Agrelius Christine Anthony Lee Garber Bob Ward Staff Multimedia Editors Brian Brannon Ben Jones Design and Production Larry Bauer

Design

Olga D'Astoli

Cover Design

Kate Wojogbe

Jennie Zhu

n Administrative Staff

Products and Services Director Evan Butterfield Senior Manager, Editorial Services Lars Jentsch Manager, Editorial Services Jennifer Stout Senior Business Development Manager Sandy Brown Senior Advertising Coordinator Marian Anderson

Circulation: Computer (ISSN 0018-9162) is published monthly by the IEEE Computer Society. **IEEE Headquarters**, Three Park Avenue, 17th Floor, New York, NY 10016-5997; **IEEE Computer Society Publications Office**, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314; voice +1 714 821 8380; fax +1 714 821 4010; **IEEE Computer Society Headquarters**, 2001 L Street NW, Suite 700, Washington, DC 20036. IEEE Computer Society membership includes \$19 for a subscription to Computer magazine. Nonmember subscription rate available upon request. Single-copy prices: members \$20; nonmembers \$99.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Canadian GST #125634188. Canada Post Corporation (Canadian distribution) publications mail agreement number 40013885. Return undeliverable Canadian addresses to PO Box 122, Niagara Falls, ON L2E 6S8 Canada. Printed in USA. **Editorial**: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *Computer*

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *Computer* does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space.

1

Innovative Technology for Computer Professionals

www.computer.org/computer

iter



COVER FEATURES

25 Advancing the Science of Digital Forensics

Gary C. Kessler

Digital forensics, the branch of forensic science that focuses on the recovery and investigation of digital data, has applications in many contexts outside the courtroom, including research, policy enforcement, and intelligence gathering..

28 Distinct Sector Hashes for Target File Detection

Joel Young, Kristina Foster, Simson Garfinkel, and Kevin Fairbanks

Using an alternative approach to traditional file hashing, digital forensic investigators can hash individually sampled subject drives on sector boundaries and then check these hashes against a prebuilt database, making it possible to process raw media without reference to the underlying file system.

36 Network Forensics: An Analysis of Techniques, Tools, and Trends

Ray Hunt and Sherali Zeadally

Researchers in the growing fields of digital and network forensics require new tools and techniques to stay on top of the latest attack trends, especially as attack vectors shift into new domains, such as the cloud and social networks.

44 SCADA Systems: Challenges for Forensic Investigators

Irfan Ahmed, Sebastian Obermeier, Martin Naedele, and Golden G. Richard III When security incidents occur, several

For more information on computing topics, visit the Computer Society Digital Library at www.computer.org/csdl.

challenges exist for conducting an effective forensic investigation of SCADA systems, which run 24/7 to control and monitor industrial and infrastructure processes.

52 Smartphone Security Challenges

Yong Wang, Kevin Streff, and Sonell Raman

Because of their unique characteristics, smartphones present challenges requiring new business models that offer countermeasures to help ensure their security.

COMPUTING PRACTICES

60 Using Tracing to Solve the Multicore System Debug Problem

Aaron Spear, Markus Levy,

and Mathieu Desnoyers

The common trace format overcomes deficiencies in traditional software tools to optimize modern multicore designs by providing open source formats to monitor the state and interaction of concurrent systems over time.

RESEARCH FEATURE

65 IaaS Cloud Architecture: From Virtualized Datacenters to Federated Cloud Infrastructures

Rafael Moreno-Vozmediano, Rubén S. Montero, and Ignacio M. Llorente

As a key component in a modern datacenter, the cloud operating system is responsible for managing the physical and virtual infrastructure, orchestrating and commanding service provisioning and deployment, and providing federation capabilities for accessing and deploying virtual resources in remote cloud infrastructures.

ABOUT THIS ISSUE

The cover features in this special issue offer a snapshot of four interesting, varied, and relevant areas of research activity in the digital forensics field: computer forensics, network forensics, control system vulnerabilities, and mobile device security.

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

COLUMNS

10 Computing Conversations Vint Cerf: A Brief History of Packets

Charles Severance

13 Computing and the Law Ď

Employment Issues Confronting Start-Up Companies

Brian M. Gaff, Timothy P. Van Dyck, and Elizabeth A. Peters

16 32 & 16 Years Ago

Computer, December 1980 and 1996 Neville Holmes

73 Discovery Analytics

Factorizing Event Sequences Naren Sundaravaradan, Naren Ramakrishnan, and David A. Hanauer

76 Invisible Computing

Digital Fabrication Manfred Lau, Jun Mitani, and Takeo Igarashi

80 Security

Atomic-Level Security for Web Applications in a Cloud Environment

Arnold Brown, Benjamin Apple, James **Bret Michael, and Michael Schumann**

84 Software Technologies

Awareness in Software-Intensive Systems **Emil Vassev and Mike Hinchey**







SUSTAINABLE FORESTRY INITIATIVE

ed wron containing

soy and/or vegetable oils

Flagship Publication of the IEEE Computer Society

December 2012, Volume 45, Number 12

90 Entertainment Computing

Building a Virtual World: The Pipeline and Process **Brad Hallisey**

116 Forward Slash 🛱

David Alan Grier and Erin Dian Dumbacher

NEWS

18 Technology News

Have Java's Security Issues Gotten out of Hand? Lee Garber

22 News Briefs Lee Garber

MEMBERSHIP NEWS

- **6** CS President's Message
- **93** Report to Members: **2012 Eelection Results**
- **96** IEEE Computer Society Connection
- **99** Call and Calendar

DEPARTMENTS

- **4** Elsewhere in the CS
- 59 Computer Society Information
- **102** Career Opportunities

ONLINE

2012 Annual Index www.computer.org/computer/12index

2012 Reviewer Thanks www.computer.org/computer/12reviewers

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

Reuse Rights and Reprint Permissions: Educational or personal use of this material is permitted without fee, provided such use: 1) is not made for profit; 2) includes



this notice and a full citation to the original work on the first page of the copy; and 3) does not imply IEEE endorsement of any third-party products or services. Authors and their companies are permitted to post the accepted version of their IEEE-copyrighted material on their own Web servers without permission, provided that the IEEE copyright notice and a full citation to the original work appear on the first screen of the posted copy. An accepted manuscript is a version which has been revised by the author to incorporate review suggestions, but not the published version with copyediting, proofreading and formatting added by IEEE. For more information, please go to: http://www.ieee.org/publications_standards/publications/rights/paperversionpolicy.html.

Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or pubs-permissions@ieee.org. Copyright © 2012 IEEE. All rights reserved.

Abstracting and Library Use: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy for private use of patrons, provided the per-copy fee indicated in the code at the bottom of the first page is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit www.ieee.org/ web/aboutus/whatis/policies/p9-26.html.

Work/Plav

Computer Highlights Society Magazines

he IEEE Computer Society's lineup of 12 peer-reviewed technical magazines cover cutting-edge topics in computing, including scientific applications, Internet computing, machine intelligence, pervasive computing, security and privacy, digital graphics, and computer history. Select articles from recent issues of Computer Society magazines are highlighted below.

Söftware

In "The 10-Minute Test Plan" in Software's September/ October issue, James A. Whittaker describes a method for writing test plans that he developed as an engineering director at Google. Whittaker wanted to reconcile the utility of test planning with the often voluminous test plans that are "written, reviewed, and left to rot in place as the hustle and bustle of software development progresses. ..." To answer the question, he adopted "a very industry-oriented technique," giving the group of people who reported to him a nebulous task with unrealistic time constraints. The first task: "Build a test plan for Google App Engine; you have 10 minutes." Subsequent tasks changed the application target but kept the same time constraint. Out of this exercise, three requirement categories emerged that boiled test planning down to its essence, creating a process for documenting them in 10 minutes or less.

SECURITY PRIVACY

"Integrating User Customization and Authentication: The Identity Crisis" is the Building Security In department article in *S&P*'s September/October issue. Authors Željiko Obrenović and Bart den Haak of Backbase describe four patterns for integrating advanced personalization mechanisms with existing security infrastructures. "Many people have come to expect user customization," they write, "but creating good designs isn't trivial." The authors offer their patterns as relevant beyond personalization to "any situation that requires mapping user identities to application-specific data."

Internet Computing

"Real-Time Communications in the Web: Issues, Achievements, and Ongoing Standardization Efforts" is the Standards department article in *IC*'s September/October issue. Authors Salvatore Loreto of Ericsson Research and Simon Pietro Romano of the University of Napoli Federico II look at Internet Engineering Task Force and Worldwide Web Consortium activities to enable a Web application running on any device to send and receive real-time media and data in a peer-to-peer fashion between browsers. The IETF's RTCWeb working group is focused on the communication protocols for setting up and managing a reliable communication channel between any pair of nextgeneration browsers. The W3C's WebRTC working group is defining an API that lets browsers and scripting languages interact with media devices and transmission functions.

lntelligent Systems

Taking cues from *The Matrix* and IBM's Watson, the authors of "An Artificial Player for a Language Game" introduce Ottho—short for "On the tip of my thought"—in *IS*'s September/October issue. Ottho solves language games through a knowledge-infusion process that adopts naturallanguage processing techniques to build a knowledge base in its memory, which a reasoning mechanism exploits to play the game. Borrowing its approach from the adaptivecontrol-of-thought theory about human long-term memory, Ottho crawls the Web and extracts knowledge that it stores as cognitive units in an interconnected network.

Professional TECHNOLOGY SOLUTIONS FOR THE ENTERPRISE

"Deploying Mobile Data Services: An Australian Case Study" is one of five cover features in *IT Pro*'s September/ October theme issue on mobile and wireless technologies. Researchers from the University of Southern Queensland present an analytical model to help decision makers successfully deploy mobile data services. They base their model on a review of current mobile-business models and a case study of six Australian organizations that implemented mobile data services.

micro

In *Micro's* September/October special issue on energyaware Computing, guest editors Thomas F. Wenisch of the University of Michigan and Alper Buyuktosunoglu of IBM's Thomas J. Watson Research Center introduce five articles that examine energy efficiency and its implications across scales—from sensors and embedded systems to data warehouses and from circuits through software. "The march of Moore's law continues to provide every more transistors," they write, "but unfortunately Dennard scaling—the concomitant reduction of CMOS threshold and supply voltages—has come to an end. ... In a world of Dennard scaling, energy efficiency is the new performance."

Computer Graphics

In *CG*&*A*'s November/December issue, the Spatial Interfaces department presents "Beaming: An Asymmetric Telepresence System," describing a European Commission project funded under the FP7 Information and Communication Technologies Work Programme. Beaming—an abbreviated term for Being in Augmented Multimodal Naturally Networked Gatherings—focuses on recreating a real meeting space in 3D virtual reality to support collaborative work with remote participants who use mobile, self-calibrating, and dynamically reconfigurable VR display technologies.

Computing

The theme of *CiSE*'s November/December issue is modern programming languages. "The proliferation of new languages and paradigms can at first appear confusing," write guest editors Massimo Di Pierro of DePaul University and David Skinner of the Lawrence Berkeley National Laboratory, "but it shows a clear trend toward increasing the expressiveness and readability of languages while decoupling the coding of the algorithm from parallelization and concurrency optimizations." They introduce articles on three languages—Clojure, Erlang, and Haskell—and compare performance results from the authors of each article on solving the same problem—namely, parallelizing a naïve solver for a 1D Poisson equation.

MultiMedia

Question-answering (QA) research addresses informationoverload problems by leveraging advanced domain knowledge and analytic techniques to return precise answers to natural-language queries. So far, the research has focused mostly on text, but *MultiMedia*'s October-December issue features "Multimedia Question Answering," a survey of preliminary work to extend text-based QA research to multimedia and of the challenges that remain.



In *PvC*'s October-December theme issue on healthcare, guest editors Maria Ebling of IBM's Watson Research Center and Joseph Kannry of the Mount Sinai Medical Center introduce four articles that describe technological capabilities that pervasive computing offers future healthcare providers and recipients. These range from advanced cardiovascular monitoring systems to large-scale sensing in support of community health. Privacy concerns come front and center in the final theme article, "Aging, Privacy, and Home-Based Computing: Development of a Framework for Design."

Athenals of the History of Computing

In 1949, a team at the US Army's Ballistic Research Laboratory in Aberdeen, Maryland, used the ENIAC to compute the decimal expansion of pi out to 2,035 places, more than double the previous record of 808 digits. The ENIAC wasn't designed to perform this type of calculation; it could only store 200 decimal digits. In *Annals'* July-September issue, Brian J. Shelburne of Wittenberg University reconstructs how the team did it in "The ENIAC's 1949 Determination of π ." The calculation took 70 hours, and its 2,000-plus digit computation held the record until 1954, when the IBM-built Naval Ordnance Research Calculator took 13 minutes to compute 3,000-plus digits.

Education Column: Online Only 💢

Computer's Education column editor Ann Sobel interviews Jennifer Dalby, of Seattle University, who describes her experiences and expertise with open online learning sessions and discusses the potential effects of massive open online courses (MOOCs); http://youtu.be/B2-NhImkDMo.

"Is There a MOOC in Your Future?," a plenary session at the 2012 IEEE Frontiers in Education Conference held in Seattle, examines current trends and future prospects in online education; www.youtube.com/watch?v=EG1JbDovI-w.

Implementing SP7: A Review of Progress in 2012



John W. Walz IEEE Computer Society 2012 President



A review of 2012 activities from the president's perspective reveals significant progress in implementing the Society's strategies and achieving its goals.

n my president's message in *Computer*'s January 2012 issue ("Creating Our Future," pp. 6-7), I provided an overview of the Computer Society's Strategic Plan (SP7), the three-year plan that outlines strategies reflecting the Society's values and core competencies and offers clearly defined measures and targets delineating the various ways volunteers can assist in achieving the Society's strategic objectives.

This year-end message gives me the opportunity to provide a review of our work to date to implement these activities.

2012 STATUS

To further the work required to implement our strategies and achieve our goals, the Board of Governors adopted a revised agenda format focusing on our key strategies. BoG members are involved with our program boards and our future technologies strategy. Volunteer leaders' new project initiation forms (PIFs) were revised to focus on supporting one strategy or substrategy. Similarly, the Staff Operations Plan (OpPlan) has been realigned to reflect key strategies. The future technologies (FTs) and Special Technical Communities (STCs) strategies have representatives working across two or more program boards.

The three Society presidents—the current president, president elect, and past-president (P3)—are working with other IEEE societies and councils (S/C) to increase membership and emphasize knowledge creation. The increase in memorandums of understanding (MOUs) with other IEEE societies shows a willingness to partner within and outside IEEE, rather than "going it alone."

Some committees are moving away from expensive and infrequent face-to-face meetings to self-managed, social media and Web-based community discussions focused on sharing and collaborative knowledge creation with a high degree of transparency. The Society has supported these efforts by acquiring a license for installation of the Liferay enterprise content management system as well as a license for Google Sites.

STRATEGIC GOALS DEPLOYMENT

The SP7 strategies address four overarching CS objectives:

- Operate as the premiere professional organization in computing technology. In 2012, we have maintained our position in all measures: periodical subscribers, conference attendees, IP downloads, and IP references/citations.
- Increase the number of members. Work is in progress to attract higher-grade members and affiliates with paid memberships. Although we may have lost some industry practitioners, we now have more student members as well as more engaged nonmembers from the community.
- Develop the next generation of volunteer leadership. Work is under way to develop a self-nomination system for all levels of committee positions for our more than 40,000 volunteers.
- *Create a sustainable financial model.* Efforts are being made to achieve this objective by reducing costs to the bone, realigning staff, seeking further integration with IEEE staff, and making investments in new products and services.

Major strides were made in 2012 to deploy the SP7 strategies,

which encompass five areas: future technologies, knowledge creation, education and professional development, outreach and engagement, and STCs.

Future technologies

Each of our FT teams will use the STC organic, Web-based, organization method to build a working group that includes experts, authors, and knowledge consumers. This structure allows the FT community to work in one or more of our Society's program boards, conferences, publications, standards efforts, professional, and educational opportunities.

P3 appointed the following volunteers, who report to the president, as chairs for specific FT strategies:

- Irena Bojanova, Cloud Computing (CC). For now, this group focuses on developing Webinars and training courses, and will contribute to the development of the recently approved *Cloud Magazine*, set to launch in 2014. It has also has been canvassing authors to solicit submissions for the new *Transactions on Cloud Computing*.
- **Roger Fujii,** Smart Grid (SG). The smart grid vision (SGV) project, led by Dan McCaugherty and Steve Widergren, is creating an SG technology vision document to guide future Society SG conferences and standards.
- Xue-Wen Chen, Life Science (LS). I'm currently working with Chen and IEEE's new initiative team to increase the value of all assets under the IEEE brand. The One IEEE initiative includes a Web portal, a newsletter, a LinkedIn group, and an annual IEEE Grand Challenges conference.
- **Hironori Kasahara,** Multicore. Plans are under way for developing a Society-sponsored 2014 conference of worldwide multicore experts supported by an active STC.

FT program manager Jennifer Schopf, who works out of the Society's Washington, DC, offices, was hired to support the FT chairs. In her role, Jennifer supports the FT volunteer team by helping to articulate requirements for staff resources and providing recommendations for opportunities to engage our members and the FT community. She ensures that the CC, SG, and LS FTs are aligned with related IEEE new initiatives to avoid duplication and share funding for new activities such as SGV.

IEEE committees to improve guidelines and develop controls to ensure the submission and publication of high-quality conference papers.

Education and professional development

Every FT needs training courseware, and Don Shafer's CC education group has taken the first step to meet this need by deploying CC training at IEEE Metro Area Workshops this year.

Chuck Walrad's IT Professional Activities committee selected the SFIA

Each of our FT teams will use the STC organic, Webbased organization method to build a working group that includes experts, authors, and knowledge consumers.

Knowledge creation

Activities for this strategy include attracting authors through both IEEE top-down initiatives and the Society's bottom-up STCs.

As a first step, IEEE surveys of periodical and conference authors will facilitate developing better policies and procedures to help promote high-quality knowledge creation.

In addition, the Society is communicating with authors interested in contributing open access (OA) author-pays submissions to our first OA journal, *IEEE Transactions on Emerging Topics in Computing* (TETC).

Other activities include helping OA authors connect with Society periodicals and informing periodicals editors and reviewers about the increasing demand for an OA option from potential authors.

In addition to *Cloud Magazine* and TETC, the CS Publications Board has recommended several other new periodicals this year for IEEE approval

Computing Now editor in chief Dejan Milojičić has been working with his team to build up the CS portal with more content and enhancements to attract viewers.

Paul Croll's Technical and Conference Activities Board is working with

Foundation for IT Competency Model, which will drive our educational courseware and future certification scheme.

Dick Fairley's team has completed a review of the *Body of Knowledge* and Curriculum to Advance Systems Engineering (BKCASE), creating the Systems Engineering Body of Knowledge (SEBoK), v.1, which will drive new courseware. They will finish the Graduate Reference Curriculum for Systems Engineering (GRCSE) in 2013 so that qualified universities can use it to upgrade their systems engineering curriculum.

Phillip A. Laplante's Software Engineering Licensure Exam committee completed questions sets for professional software engineering licensure in the US that starts in 2013.

The joint Project Management Institute (PMI) and IEEE CS committee finished the draft software extension to the *Project Management Body of Knowledge*, with Dick Fairley representing the CS in this effort, which will help in preparing our upcoming software project management courses.

Liz Burd's TryComputing.org went live this fall to encourage precollege students and their parents to consider computing careers and to

PRESIDENT'S MESSAGE

provide information about the high school preparation that colleges are expecting. Through TryComputing. org, teachers can share lesson plans to assist them in teaching essential computing skills, and counselors can keep up with computing technology to help students make informed choices.

Outreach and engagement

To support its outreach and engagement goal, the Society will use the STC model to gather and organize potenalong with CS chapters at the conference venues. Discussions with conference stakeholders, including sponsors, TC leaders, organizers, keynote speakers, program board experts, authors, and attendees have showcased the importance of the dynamic knowledge creation process supported by our conferences. In several conferences, the colocated periodical editorial board meeting revealed established authors who publish in both the conference pro-

Discussions with conference stakeholders have showcased the importance of the dynamic knowledge creation process supported by our conferences.

tial members such as public policy experts, joint members of our national sister societies such as the Computer Society of India, and our members interested in groups such as the IEEE Nanotechnology Council. The STCs will identify interested members and recruit them to serve in leadership roles in the community.

STCs

Martin Arlitt and Ishfaq Ahmad are cochairs of the STC on Sustainable Computing, our largest grassroots STC, with more than 300 members, including 137 on LinkedIn and 15 ExCom leaders. This STC has published 10 newsletters this year.

Carlos Jiménez began the work on the new e-Government (eGov) STC, and under his leadership, this STC is now organizing interested parties across IEEE, using a LinkedIn group for IEEE members only.

STAKEHOLDER COMMUNICATIONS

The Society has numerous stakeholders: paid memberships, volunteers, authors, reviewers, conference attendees, subscribers, and partners such as libraries.

As president, I have visited more than a dozen of our key conferences

ceedings and a related periodical. Thus, conference program committees should consider working with periodical editorial boards in the same technical domain.

Going forward, the Society and IEEE need to allow the organizers to be agile in managing the creation, growth, merger, splitting off/pruning, and sun-setting of conferences. Program committees need an infrastructure to attract submissions from new and established authors to ensure timely, high-quality, and relevant submissions and published papers from their technical community.

These sample conference visits showcased some technologies of particular relevance to the Society, including software engineering, visualization and graphics, and computer engineering encompassing computer architectures, microprocessors, VLSI design and test technologies, parallel and distributed processing, and dependable and fault-tolerance computing. Some of these technologies are ready for standards development activities with strong industry support.

During this year, I have recorded video blogs on specific topics, using social media to reach our large community. I have enjoyed these opportunities to refocus and articulate my goals to a large, general, and unknown audience. I have found this media fits well between the short elevator conversation and an introductory speech to a body with a common interest. I have often reused sections of a video blog at subsequent face-to-face meetings for validation of our goals and activities.

At the IEEE level, I served as cochair of an ad hoc committee focused on increasing IEEE attractiveness to industry practitioners. Our work helped define the problem of the declining membership of industry practitioners, shared best practices for the turnaround of membership in technical societies, and identified possible IEEE operational changes needed to support the various societies that comprise the IEEE.

ENGAGING OUR VOLUNTEER RESOURCES

A rough census conducted this year revealed that there are more than 40,000 Computer Society volunteers. To better use this wealth of talent, one presidential initiative is to recruit and promote volunteers based on their interests. Developing this initiative, which will take a couple of years with P3 support, starts with communications, such as the recent election video blog, a bottom-up self-nomination Web form, and top-down Executive Committee calls for candidates for CS committee officers. In addition, we need to create an administrative workflow to help volunteer leaders make appointments, followed by volunteer directory publication, volunteer satisfaction surveys, and volunteer recognition.

Because our technical community is larger than our paid memberships, new efforts, processes, and benefit packages were designed and deployed for the retention and recruiting of full members and affiliates. In moving our members from transactional usage to becoming engaged contributors to the organization, the enrollment process now helps connect them so they can join one of our more than 30 TCs. I offer my personal thanks to David Alan Grier, who assisted in these efforts in his role as chair of the membership working group.

As one-half of our members are industry practitioners, we are getting excellent product and services feedback and strategy recommendations from our Industry Advisory Board. This board was instrumental in coaching the developers of our Corporate Affiliate Program (CAP), which provides high-quality, affordable educational training to corporations.

To better engage and utilize our external worldwide Sister Society partners, past-president Sorel Reisman redesigned our application and MOU forms, which have been used in successful trial implementations with the Computer Society of India and China Computer Federation.

In recognition of the portion of our members with interests in technical public policy, we are organizing our experts in this area into an STC to work with the IEEE Government Relations Council, through our appointment of Jim Moore, who will request input on technical white papers for distribution to congressional members, their staff, and our key Washington allies.

To address all aspects of our numerous fields of interest, we are now focusing the Computing Now (CN) portal content on six hot technical areas: cloud computing, high-performance computing, mobile computing, networking, security, and software engineering. Our annual member survey revealed that most of our members work in four broad technical areas, for which we will be producing individual technical newsletters: software, information and communication technologies, security and privacy, and computer engineering.

CONTINUITY WITH PAST PRESIDENTS' INITIATIVES

The advantage of the P3 structure is that it facilitates the conservation and sharing of organization knowledge over time so that the CS leadership has a better understanding of what works and what doesn't. As this review of the work we have undertaken in 2012 indicates, implementation of 2011 president Sorel Reisman's SP7 is ongoing, and efforts to support the concepts introduced last year are continuing this year.

In addition, 2010 president Jim Isaak is serving as chair of the International Public Policy community; 2009 president Susan (Kathy) Land is working on broad IEEE and TAB issues that affect our Society; 2008 president Rangakar Kasturi helped move the affiliate membership dues issue to a financial conclusion: 2004 president Carl K. Chang has upgraded COMPSAC to be the Society's flagship conference, with past Society presidents serving as general cochairs; and 2003 president Steve Diamond is leading the IEEE Cloud Computing Initiative, in which our Society is the strongest player.

THANKS TO THE 2012 TEAM

In addition to past-president Sorel Reisman and president-elect David Alan Grier, the 2012 Executive Committee (ExCom) included seven vice presidents: Tom Conte, first VP (and VP of publications); André Ivanov, second VP (and BoG secretary); Liz Burd, VP of educational activities; Paul Croll, VP of technical and conference activities; Paul Joannou, VP of professional activities; Sattupathu Sankaran, VP of member and geographic activities; Charlene (Chuck) Walrad, VP of standards activities; and Jim Moore, treasurer, who also serves as Division V Director. Nonvoting ExCom members include Kathy Land, Division VIII Director; Roger Fujii, Division VIII Director-Elect; and Angela Burgess, the Society's executive director.

Each of the six program boards has made outstanding progress this year on board-specific and cross-Society projects. At the four ExCom meetings, each program board VP reported on future planning, work with staff on funding requests, progress on various projects, deliverables, gaps addressed, and issues resolved.

A special thanks goes to Sankaran, who has rendered significant service to the Society with excellence and dedication in his numerous contributions during his tenure in the membership area.

The volunteer-staff partnership serves as the foundation for all of our efforts within the Society. I extend my appreciation to all who have worked together this year to shape the future of our organization.

s I conclude my term as the Society's 2012 president, I extend my sincere congratulations to president-elect David Alan Grier. I look forward to working with him and his ExCom while serving as a member of P3 in 2013.

John W. Walz retired from Lucent/ AT&T with more than 25 years of software and systems engineering and management leadership experience. Contact him at j.walz@computer.org.



Vint Cerf: A Brief History of Packets

Charles Severance



TCP/IP evolved from 20 years' research that sought a way to move from a telephone-style circuit-switched infrastructure to a packetswitched infrastructure.

n the late 1980s and early 1990s, academics, governments, and companies around the world had built and deployed our current shared Internet infrastructure using the TCP/IP protocol. Having a welldeveloped open protocol implemented on a wide range of computers was an essential prerequisite for the Internet's rapid success and growth.

The simple answer to the question, "Where did TCP/IP come from?" is that its basis was ARPANET, an earlier small-scale research network. I recently spoke about this technology's emergence with Vint Cerf, who is recognized as one of the "fathers of the Internet" and was a cofounder of the Internet Society, which was established in 1992 to provide leadership in establishing Internet-related standards, education, and policy.

To view the full video of this interview, visit www.computer.org/ computingconversations.

FROM THE GROUND UP

Of course, the real answer to the question of TCP/IP's origins is that it ultimately emerged from 20 years' research into a wide array of topics that explored moving from a telephone-style, circuit-switched infrastructure that required dedicated resources for every active pair of users to a packet-switched infrastructure where all active users dynamically shared all resources:

Leonard Kleinrock specifically studied the packet-switching concept at MIT in 1961. Kleinrock focused on message switching and did a brilliant dissertation on the use of queuing theory to analyze what networks of queues would look like with a messageswitching approach. Although he never used the word "packet," Kleinrock's analysis is as applicable to packet switching as it is to message switching.

One advantage of a packetswitched network is that it can dynamically route data around partial network outages as might be experienced during wartime or perhaps due to a severe large-scale weather event, such as a hurricane or typhoon. While working at RAND in 1962, Paul Baran created an extensive design for a resilient packet-switched voice network. As Cerf describes it,

Before the existence of integrated circuits or anything else, [Paul] is saying, "We really should be digitizing and packetizing voice and using pole-

mounted radios to create a highly connected environment so that if holes are knocked out by nuclear explosions, information can still get from one end to the other." You chop up the speech into little 20-millisecond pieces and dynamically route it like a "hot potato"—if you get something, you get rid of it as fast as you can. That's around 1962 and is documented in an 11-volume series called "On Distributed Communications," but he can't sell it to anybody. The traditional telcos—AT&T in particular-and the people who were in the Defense Communications Agency laughed him out of the room and say it's a silly idea that can't possibly work, so he should just go away. He never gets anywhere with the idea in spite of all his documentation.

But academics were starting to do research on how to break data into packets and send those packets over the traditional telephone network or "local area" wired networks:

In the 1964-1965 timeframe, Donald W. Davies of the National Physical Laboratory in London also gets the packet-switching bug and tries to get money from the science research commission in England, but he only gets enough to build one node. He builds a network and invents the term "packet" to describe what these objects are—and it works. He has a bunch of terminals and other things hanging off this one node so in a funny way, he builds a local-area network, but it's based on physical wires. In 1966, Larry Roberts along with Thomas Merrill does a point-to-point experiment to test packet switching. It's between the ANS FQ-7 machine at System Development Corporation in Santa Monica and the TX-2 machine at MIT Lincoln Laboratories, which is where Larry is. They demonstrate on a 2,400-bit per second modem that can move packets back and forth.

While the research into the theoretical and technical aspects of packet switching was under way, another thread of inquiry explored what people might do with a ubiquitous always-on network:

J.C.R. Licklider is a psychologist at MIT, but he's convinced in the early 1960s that computing is important to non-numeric processing: it will allow people to collaborate in ways they never could before. He starts the Information Processing Techniques Office at ARPA and encounters Douglas Engelbart at SRI International. The two bond because Engelbart's oN-Line System (NLS) is all about nonnumeric computing and the ability of people to build up a superstructure of communication and documents and interact with each other. Engelbart has a "world wide web" in a box at SRI, and Licklider sends out slightly tongue-incheek notes to his community of people about this "intergalactic network." Licklider really gets credit for having put this meme of communication and collaboration in place at ARPA.

MAKING CONNECTIONS

In addition to the growing notions of how people could collaborate via a shared digital network, there was also the practical consideration of deploying increasing numbers of computers and computer terminals to meet the military's informationprocessing needs:

Robert W. Taylor comes along to run the IPTO after Licklighter and is all hacked-off because he has three terminals in his office at the Pentagon connected to three different machines: "Why can't there be one terminal talking to all three? We need a network!" As he's pursuing this idea with Charlie Herzfeld, the head of ARPA at the time. Charlie hands him \$1 million over a 20-minute conversation, and now Taylor has to figure out who will actually do this. Taylor isn't a technologist, he's a psychologisttype, so he recruits Larry Roberts from Lincoln Laboratories, who did that earlier packet experiment over the 2.400-baud modem.

Another thread of inquiry explored what people might do with a ubiquitous alwayson network.

With \$1 million to spend on research into a shared packetswitched network infrastructure, the IPTO wrote a Request for Quotation (RFQ) to solicit research proposals. Cerf wrote one of many responses to the ARPA RFQ for network research:

I help to write one of them with my colleague Steve Crocker while we're still at UCLA as graduate students and consulting with a company called Jacobi Systems in Santa Monica, but our proposal is not funded. But the next thing we know, Len Kleinrock, who wrote the original dissertation work on packet-switched networking at MIT, has come to UCLA to teach and explore queuing theory. Len is a close compatriot of Larry Roberts because they were both at Lincoln Labs together, so Len gets the network measurement center piece of the ARPANET project. Cerf, Crocker, and Jon Postel—all UCLA graduate students who had attended the same high school in California's San Fernando Valley were recruited to join Kleinrock's lab to help design and build the technology that would be used in ARPANET:

I'm the principal programmer, Steve Crocker takes responsibility for managing and leading the network working group, which led to the documents describing the host-to-host protocols, and Jon Postel becomes the keeper of the documentation. He's the Request for Comments editor, the guy who becomes the numbers czar to keep track of the address spaces and allocations.

ARPANET eventually evolved to the point that, in 1972, it worked well enough for a demonstration:

The first demonstration of ARPANET happens in the basement of the Washington Hilton Hotel in October 1972. A bunch of people from the packet-switched networking community attend, not only from the US but also from France, England, Italy, Germany, and elsewhere. That group of about 25 or 30 people convenes, sees ARPANET in operation, and then forms this international network working group modeled after the working group that Steve Crocker managed. At this point, I become the chairman of that group because Steve's busy at ARPA doing artificial intelligence.

EVOLVING PROTOCOLS

At the end of 1972, Cerf graduated from UCLA and became a faculty member at Stanford. Bob Kahn moved from BB&N to ARPA and took the ARPANET project to the next level:

In 1973, Bob comes out from ARPA and says that ARPA is working on networking capabilities beyond the original ARPANET for the military. If we're serious about putting computers in command and control, they have to be mobile: we need mobile radio and satellite, in addition to the fixed wire systems represented by ARPANET.

Bob's brilliant idea is not to build one network with all those technologies embedded in it—instead, he breaks them apart and says let's build a packet satellite network that takes into account that it has a halfsecond of round-trip time. Let's build a packet-radio network that optimizes a system whose connectivity is changing with time as things move around and you get variable delay and interference.

Expanding from the original telephone-line based ARPANET to a network architecture that included many cooperating networks with points of interconnection required a new design:

We decide to build a gateway that today we call a "router" and also introduce other things like how to refer to another network. Each network thinks it's the only network in the universe. At that time, you didn't have a vocabulary that said, "Take this packet and move it to another computer on another network somewhere else that you might not even be connected to." So we have to invent an "Internet" address space to solve that problem. We have to find a way to allow packet losses in this path to be recovered. TCP now becomes a manager of reliability on an end-to-end basis instead of relying on each underlying network to be reliable. ARPANET was built on the assumption that you could build a reliable underlying network. The Internet was based on the assumption that no network was necessarily reliable and you had to do end-to-end retransmissions to recover.

This work on a "network of networks" began in earnest as a research project funded by ARPA. Cerf and his colleagues did all of their design work in the open and shared it within the academic and commercial communities:

Bob and I get the first paper written and published in IEEE Transactions on Communications in May 1974. Nobody pays much attention to it. Meanwhile, ARPA is funding us to make this actually work. At Stanford, I'm working with my graduate students, some at Xerox PARC and some at Stanford, on detailed specifications of TCP/IP. We publish it in December 1974, and it's the first time the word "Internet" shows up in print anywhere.

From its earliest days, ARPANET focused on connecting people, information, and technology.

From 1973 to 1978, the research team designed and implemented four complete iterations of the Internet protocols as they found and solved new challenges. One innovation was to separate the TCP layer into the TCP and IP layers, which let real-time applications use the Internet without the error correction added by TCP:

For the next five years, we do everything we can to get TCP/IP implemented on every operating system we can find. It goes onto IBM machines, Digital machines, HP, Unix. We sent a Unix version built by BB&N out to Berkeley to the BSD release guys, and Bill Joy says, "I don't like that code," so he writes his own and puts it into BSD 4.2, the version of Unix that carries TCP/IP to the academic world. Around the same time, Sun Microsystems comes along and builds these fantastic workstations; it wants to use open protocols, so it adopts Unix, and TCP/IP comes along with it. Ethernet connects the workstations together. Sun drives the academic community to TCP/IP.

As the number of workstations and mainframe computers that could support TCP/IP and be connected to ARPANET grew and an increasingly rich set of networked applications were developed, more universities and research labs wanted to be connected to ARPANET, and its infrastructure started to groan under the network traffic load:

All of this places huge demands on the ARPANET backbone, which is only running at 50 kilobits per second, and eventually leads to the need for higher speed. NSF jumps into the fray, seeing how valuable all this is for the academic community, and concludes that it should build a network that runs even faster. It does, and the result is NSFNet.

hat's amazing about the story of ARPANET is that from its earliest days, it focused on connecting people, information, and technology. It's also a story of having patience and being willing to take the time to build it right technically, even if doing it right meant starting over from time to time. A relatively small group of well-funded researchers worked collectively for more than 20 years, starting over and redesigning their systems multiple times as new use cases presented themselves. It's fortunate that by the mid-1980s, TCP/IP was a well-developed technology that provided a solid basis for the Internet revolution of the 1980s and the Web revolution of the 1990s. C

Charles Severance, Computing Conversations column editor and Computer's multimedia editor, is a clinical associate professor and teaches in the School of Information at the University of Michigan. Follow him on Twitter @drchuck or contact him at csev@umich.edu.

CN Selected CS articles and columns are available for free at http://ComputingNow.computer.org.

Employment Issues Confronting Start-Up Companies

Brian M. Gaff, Timothy P. Van Dyck, and Elizabeth A. Peters Edwards Wildman Palmer LLP



The twelfth in a series of articles providing basic information on legal issues facing people and businesses that operate in computing-related markets discusses employment law issues that new companies encounter.

mployment law encompasses all the rights and obligations that define a company's relationship with its employees, as well as job applicants, and former employees. Understanding and complying with these sometimes overlapping state and federal laws can be overwhelming, especially to a first-time entrepreneur. New companies often need help identifying and understanding the employment laws and regulations that apply to them. This article provides an overview of some of the more significant legal challenges faced by new employers.

Be sure to check the IEEE Computer Society's website for the podcast that accompanies this article (www.computer.org/portal/web/ computingnow/computing-and-the-law).

QUESTIONS TO AVOID IN INTERVIEWS

Interviews with job applicants allow employers and prospective

employees to learn more about one another and about the available job to determine whether the applicant would be a good fit for the position. While there are many things an employer might like to know about a prospective employee, employers must be cautious about what subjects they broach in an interview.

Employers are prohibited from asking questions that might be interpreted as discriminatory—even if that is not their intent. This means that employers can't ask questions related to a candidate's religion, race, ethnicity, national origin, citizenship, age, marital or family status, disability, sexual orientation or gender, or genetic information.

Even if a candidate volunteers information relating to these topics, it's wise for the employer to simply move on to the next question. The best way to avoid any appearance of discrimination is to stick to issues related to the job. Asking someone whether they have children might seem like chit chat to you, but if it's completely unrelated to the candidate's qualifications, it could be interpreted as discriminatory.

FLSA EXEMPT AND NONEXEMPT CLASSIFICATIONS

The Fair Labor Standards Act (FLSA) establishes standards for minimum wage, overtime pay, recordkeeping, and child labor. Some employees are exempt from FLSA's minimum wage and overtime pay provisions. This includes executive, administrative, professional, outside sales, and computer employees. These exemptions are narrowly defined, so employers should carefully check the requirements for each potential exemption before classifying an employee as FLSA-exempt.

For example, to qualify for the computer employee exemption, today the employee must be compensated

at a rate of not less than \$455 per week or \$27.63 per hour, and must be employed as a computer systems analyst, computer programmer, software engineer, or other similarly skilled worker in the computer field. The employee's primary duties must consist of: (1) the application of systems analysis techniques and procedures, including consulting with users, to determine hardware, software, or system functional specifications; (2) the design, development, documentation, analysis, creation, testing, or modification of computer systems or programs, including prototypes, based on and related to user or system design specifications; (3) the design, documentation, testing, creation, or modification of computer programs related to machine operating systems; or (4) a combination of these duties, the performance of which requires the same level of skills.

Employers should be sure to properly classify workers as FLSA-exempt or nonexempt at the commencement of their employment. Misclassification of employees as FLSA-exempt could lead to significant liability on the employer's part in both unpaid wages and monetary penalties.

INDEPENDENT CONTRACTORS VERSUS EMPLOYEES

New companies might be tempted to hire workers as "independent contractors" instead of as employees to avoid the complications and obligations of being an employer. The proper classification of workers as independent contractors or as employees is one of the most important issues facing employers today. Proper classification determines a company's obligations with respect to issues ranging from minimum wage and overtime requirements, payroll taxes, benefit plan eligibility, workers' compensation, unemployment benefits, and potential exposure under workplace discrimination laws.

Government agencies have varied

interpretations of the proper classification standards for independent contractors, but the factors tend to fall under three general requirements. First, to be properly classified as an independent contractor, the worker must be free from the company's control or direction in the performance of the services, both under the terms of the contract and in fact. Second, the services provided by the independent contractor must be outside the usual course of business of the company. Third, the worker must be engaged in an independently established

In recent years, government agencies have been cracking down on employers who misclassify employees as independent contractors.

business. It's important for employers to realize that the existence of an independent contractor agreement, by itself, is not determinative of a worker's status.

In recent years, government agencies have been cracking down on employers who misclassify employees as independent contractors. New employers should be careful to properly classify their employees.

CONFIDENTIALITY AND NONCOMPETITION AGREEMENTS

If employees will have access to the employer's confidential information or trade secrets, it is imperative that the employer require them to sign a stand-alone confidentiality and nondisclosure agreement on the commencement of employment. The agreement should provide a general description of the categories of confidential information that are protected, and prohibit the employee from disclosing that confidential information both during and after the employee's employment with the company.

Also, employers might want employees to sign noncompetition or nonsolicitation agreements on the commencement of employment. Such agreements are generally enforceable in most states (California is one major exception), as long as they're reasonably tailored to project a legitimate business interest of the employer. Legitimate business interests might include, for example, the protection of confidential information and trade secrets and long-term customer relationships. The noncompetition or nonsolicitation restrictions should be limited in time, and tailored in scope to fit the geographic region where the employee works for the employer or the customers with whom the employee has contact.

Employees are increasingly mobile. Companies should take steps to protect their confidential and proprietary information to minimize the chances that a former employee takes this kind of information with him to use at a new job.

EMPLOYEE PRIVACY ISSUES

Employers understandably don't want employees surfing inappropriate websites, shopping, or playing games while on the clock, let alone selling trade secrets or using company computers to harass their coworkers. The law generally allows employers to monitor employees' Internet use and communications while on the job and within reason, as long as the monitoring does not infringe upon the employees' privacy rights.

Employers are generally allowed to monitor the websites that employees visit on their work computers. Employers are also generally permitted to monitor employees' use of their company-provided email accounts, as long as the employer has not informed the employee that work emails will be private or confidential.

Employees' personal email accounts are afforded more protection. Employers are generally prohibited from using employees' private passwords to hack into their personal email accounts. However, when employees access their personal email accounts from their companyowned computers, information might be stored on the computers, which would allow the employer to view the personal emails the employee accessed on the work computer. This is somewhat of a gray area, but employers might be permitted to view such emails if they have a strong policy in place that makes absolutely clear to employees that their work computers should not be used for personal email, and employees have no expectation of privacy in any activity on their work computers.

Employers should be reasonable about how and when they monitor their employees' communications. Courts are less likely to find an employer liable for violating an employee's right to privacy if the employer has a legitimate workrelated reason for monitoring the communications. To better protect themselves, employers should adopt strong computer, phone, Internet, and social media use policies to make clear to employees what activities are allowed and what activities are prohibited, and to warn employees that they do not have any expectation of privacy in any activities on their work computers.

mployment law is complicated, especially in view of different requirements imposed by individual states in the US, and how those requirements interrelate with applicable US federal law. Consequently, it can be easy to inadvertently run afoul of one or more regulations. Some oversights can result in large penalties. To minimize the chances of your company having to contend with this, work closely with your employment lawyer to ensure that your human resources policies and procedures are fully compliant with all applicable laws.

Brian M. Gaff is a senior member of IEEE and a partner at the Edwards Wildman Palmer LLP law firm. Contact him at bgaff@edwardswildman.com.

Timothy P. Van Dyck is a partner at the Edwards Wildman Palmer LLP law firm and is Co-Chair of the Firm's Labor and Employment Group. Contact him at tvandyck@edwardswildman.com.

Elizabeth A. Peters is an associate at the Edwards Wildman Palmer LLP law firm and is a member of the Firm's Labor and Employment Group. Contact her at epeters@edwardswildman.com.

The content of this article is intended to provide accurate and authoritative information with regard to the subject matter covered. It is offered with the understanding that neither IEEE nor the IEEE Computer Society is engaged in rendering legal, accounting, or other professional services or advice. If legal advice or other expert assistance is required, the services of a competent professional person should be sought.



handles the details so you don't have to!

- Professional management and production of your publication
- Inclusion into the IEEE Xplore and CSDL Digital Libraries
- Access to CPS Online: Our Online Collaborative Publishing System

Choose the product media type that works for your conference: Books, CDs/DVDs, USB Flash Drives, SD Cards, and Web-only delivery!

Contact CPS for a Quote Today!

www.computer.org/cps or cps@computer.org



IEEE **()** computer society

32 & 16 YEARS AGO

DECEMBER 1980

www.computer.org/csdl/mags/co/1980/12/index.html

PRESIDENT'S MESSAGE (p. 3) "The first thing that comes to my mind is the society membership, which has continuously and significantly increased during the past two years—from 38,701 (or less than one fifth of the total IEEE membership) at the beginning of 1979 to the projected end-of-year figure of 51,000 (about one quarter of the total IEEE membership). ..."

USABILITY (p. 6) "In the interests of human engineering and so we won't appear to be a complete set of nits, I wish to suggest that for our next Computer Society election we print the 'punch-out' boxes as the leftmost column on the ballot, followed closely by the names of the candidates. This would eliminate the myriad problems in alignment and registration experienced in the recent election by those like myself with minimal motor skills and visual acuity."

VIDEOTEX (p. 8) "... second-generation systems are being planned and discussed, and viable videotex services are rapidly evolving into public information utilities. These utilities will offer a variety of information services and transactions, such as retrieval from multiple independent data bases, messaging, electronic mail, conferencing, banking, teleshopping, and interest matching. In secondgeneration systems the emphasis is more on videotex as a communication medium rather than as a simple information retrieval system. Future systems will move toward two-way communication among users as well as between users and information providers."

INSTRUMENTING PROGRAMS (p. 17) "Symbolic traces increase error-detection capabilities of program tests and indicate the extent of their coverage. This instrumentation system generates traces automatically upon program execution."

PROGRAMMING SOLO (p. 24) "The unique situation of the individual programmer warrants a specialized guide to verification and testing. The programmer who works alone performs all management tasks, and lacks independent internal or external quality assurance groups. However, his tasks are usually of an intellectually manageable size, and he does not face many of the problems encountered in larger systems—coordination among programmers and massive integration, for example."

TUTORIAL ON PROCESS CONTROL (p. 35) "Progress in electronic hardware alone will not ensure success in measurement and control of nonelectronic processes. Nevertheless, the future for stored program controllers has never seemed brighter."

WORKSHOP ON VLSI DESIGN (p.

51) "Fault-tolerant techniques are needed in VLSI design. But, which of the familiar ones are appropriate? Should they be applied on-chip or in multiple-chip systems?

What areas require the development of new techniques? These and similar problems were discussed at the Workshop on Fault-Tolerant VLSI Design, ..."

COMMUNICATION (p. 73) "Software people do not communicate well with colleagues from more established disciplines. Our noses are too close to the software grindstone for us to begin to understand our place in the broader world. When we do attempt to communicate, the information is often too cryptic for outsiders to understand or simply wrong, especially when the subject of the communication is a time estimate for a software job."

WHAT IS SOFTWARE? (pp. 73-74) "We no longer ask *whether* we shall use software to solve a given problem—software is the way of the digital-computer-controlled future. The question has become, How can we better understand what software is and is not? ..."

CAMERA COUPLING (p. 85) "Magicam's Programmable Motion Control system is an outgrowth of a system used in the filming of Carl Sagan's *Cosmos, Star Trek: The Motion Picture*, and scores of commercials. Magicam's underlying techniques are being extended to applications in computeraided manufacturing, numerical control, model board and hybrid simulators, microwave and tracking cameras, animation stands, optical printers, and repeat-pass photographic systems."

COMPCOM SPRING 81 (p. 87) "Nobel Laureate Herbert A. Simon, professor of computer science and psychology at Carnegie-Mellon University, will present the special address, 'Prometheus or Pandora: The Influence of Automation on Society.' Simon states that automation, made possible by rapidly advancing computer technology, promises to be the major source of increased productivity we need in our economy. But he goes on to ask, 'How can we realize these benefits without incurring the heavy social cost of unemployment, alienation from work, and loss of privacy and freedom?'"

Editor: Neville Holmes; holmeswn@yahoo.com.au

www.computer.org/csdl/mags/co/1996/12/index.html

VIDEODISCS (p. 14) "Digital videodisc (DVD) technology, which has been touted as the next hot consumer electronics technology, has apparently overcome problems involving format standards and security, and is scheduled to debut this Christmas."

GOLDEN CHIPS (p. 17) "Purdue University researchers are coating microchips with gold in an attempt to develop nanotechnology that would permit the manufacture of ultrasmall computer components.

"The researchers have created a thin gold film coating that conducts electricity by inducing electrons to hop from one gold cluster to another one at a time. (Each cluster consists of 2,400 gold atoms.) Using this technology would cause less heat buildup than using conventional silicon-based circuits, which have continuous currents flowing through them."

TRUE SEABORN (p. 24) "This is the last issue of *Computer* to appear during my tenure [as IEEE CS publisher], the last during the Computer Society's 50th anniversary year, and the last before we launch a redesigned *Computer*. After 23 years on the job, the timing seems right to make a few remarks about how *Computer* got where it is today and what we hope to achieve next year."

MULTIPROCESSORS (p. 29) "Shared memory multiprocessors have been used mostly to increase the throughput of independent jobs, rather than to speed up the execution of an entire application. The main reason their use has been restricted is the lack of mature compiler technology, which has made parallel programming a privilege for experts only. However, the wide availability of these systems over the last couple of years is changing this situation ..."

MAKING DECISIONS (p. 33) "Probabilistic inference is becoming an integral part of decision-making systems, but it is so computationally intensive that it is often impractical. The authors report on the effectiveness of speeding up this technique by exploiting its parallelism."

VIZUALIZATION (p. 42) "Visualization applications—such as flight simulators and virtual reality environments—use geographic information systems to represent actual terrain. Applications like these impose stringent restrictions on acceptable performance and response time. Sequential methods do not meet these requirements, but parallel methods can."

MASSIVE PARALLELISM (p. 50) "Can scientific programming on a distributed-shared-memory multiprocessor architecture be made as easy and efficient as it is on vector supercomputers? Maybe, but we're not there yet."

PARALLEL PERFORMANCE (p. 57) "In this article, we present the performance of 14 applications on the [MIT] Alewife machine, including both coarse- and fine-grain applications. Not surprisingly, Alewife's mechanisms support the good performance of traditional coarse-grain applications But we also show that Alewife provides excellent communication mechanisms for fine-grain applications"

PARALLEL COMPILING (p. 78) "Parallel programming tools are limited, making effective parallel programming difficult and cumbersome. Compilers that translate conventional sequential programs into parallel form would liberate programmers from the complexities of explicit, machineoriented parallel programming. Polaris, an experimental translator of conventional Fortran programs that target machines such as the Cray T3D, is the first step toward this goal."

PARALLEL EFFICIENCY (p. 84) "Locating parallelism is just the first step in producing efficient multiprocessor code. Achieving high performance also requires effective use of the memory hierarchy, and multiprocessor systems have more complex memory hierarchies than typical vector machines: They contain not only shared memory but also multiple levels of cache memory."

COPYRIGHT (p. 111) "The Trumpet Winsock case is important to software owners, Internet service providers, and others who distribute software. The court ruled that shareware distributed on the Internet has copyright protection and that no one can imply a general license that forces shareware owners to abandon copyright protection and retain only those rights expressly reserved in the software license."

ADDRESSES (p. 112) "The other IETF proposal, called IP Version 6 or IP Next Generation, would establish 128-bitlong addresses. This would allow far more addresses than the 32-bit system, so it would better accommodate Internet growth. However, this plan would also require development of new TCP/IP software versions, which could be an expensive and time-consuming process."

PARALLEL PERFORMANCE (p. 152) "... The most important impediment may be the mathematical way we represent problems. This may not be the most effective way to represent problems for solution by parallel machines. Therefore, if we want to realize parallel processing's tremendous potential, we may need to focus our attention on problem representation."

Have Java's Security Issues Gotten out of Hand?



Lee Garber

In the past year, security experts have found many vulnerabilities, some critical, in Java. This represents a serious trend because Java is so widely used.

ava is among the world's most popular software platforms. It offers the ability to develop and run applications that can be used across operating systems.

However, some experts say Java is now also a major security risk. Researchers have found many flaws that enable hackers to bypass security measures, take over computing systems, steal data, and cause other problems.

"They've found dozens in 2012 and dozens in 2011," said Atif Mushtaq, senior staff scientist with security vendor FireEye.

One recent vulnerability affects every version of Java issued since 2004, including those used in most of today's smartphones. To make matters worse, online exploit collections for hackers and penetration-testing applications for security professionals now include the code for easy-toimplement Java attacks, as the sidebar "Exploit Toolkits" explains.

Many experts say that if you don't absolutely need Java, uninstall it; if you do need it for some applications, take precautions. In response to Java's problems, Apple decided to stop bundling it with OS X and has taken other protective measures. The Mozilla Foundation has blocked older, unpatched Java plug-ins from running on its Firefox browser.

As the "Java Backgrounder" sidebar discusses, Oracle, gained control of Java when it merged with Sun Microsystems in 2010. The company did not respond to multiple requests for comment.

SPATE OF PROBLEMS

Over time, Java has faced many security problems, which have gained a higher profile this year.

Researchers are discovering unique vulnerabilities throughout the code base that exist only in the Java language and the Java virtual machine (JVM), said Dan Guido, CEO of security vendor Trail of Bits. Because these vulnerabilities are specific to Java, traditional exploit mitigations don't help, he noted.

Security Explorations, a Polish consultancy, has been particularly active in investigating Java problems. It analyzed binary and source code for Java implementations from Oracle, IBM and Apple, and found multiple critical security vulnerabilities, said company CEO Adam Gowdiak.

Just since 2 April 2012, Security Explorations has uncovered many flaws, including critical sandbox bypasses that let hackers attack and gain control of systems.

"We found 50 security issues in total: 31 were reported to Oracle, including 17 complete JVM sandboxbypass exploits;" Gowdiak said, "two were reported to Apple, including one sandbox-bypass exploit; and 17 were reported to IBM, including 10 sandbox bypass exploits."

Oracle has been eliminating the recent Java problems this year and says it will fix the last two in February 2013, according to Gowdiak. He noted that hackers have exploited some of the Oracle and Apple flaws but none of the IBM vulnerabilities.

IBM, which didn't respond to requests for comment, plans to release patches this month, Gowdiak said. Apple, which also didn't respond to requests for information, has been working on patches, too. "Most of the bugs have their origin in insecure implementation of Java code," Gowdiak explained. "That naturally makes them platform independent. This means that the vulnerabilities could be easily exploited on all operating system platforms supporting vulnerable Java SE versions." They could also work on both 32- and 64-bit systems.

Two flaws are particularly troubling and exemplify the recent Java security issues.

Java 5, 6, and 7 flaws

Security Explorations found and showed how to exploit a sandboxbypass flaw in Java Standard Edition 5, 6, and 7, which are all of the versions issued since 2004 and which are found in most smartphones. This problem puts hundreds of millions of Java users at risk. The company tested the flaw on the Chrome, Firefox, Internet Explorer, Opera, and Safari browsers running on a fully patched 32-bit computer running Windows 7.

"This is bigger than any previous issue we found as part of our Java security research project," Gowdiak said. The flaw lets attackers violate the JVM's type-safety security, which is supposed to keep code from running on parts of a system it isn't authorized to access.

Such a problem, Gowdiak explained, "can lead to unrestricted access to Java classes, their fields and methods, [regardless] of Java security access controls. In the most obvious attack scenario, such a condition could lead to full privilege elevation."

A hacker could set the global value of Java's SecurityManager to null, disabling all security checks in the target JVM and enabling the hacker to cause problems such as installing malware or backdoors, as well as stealing, changing, or deleting data.

Gowdiak said that he provided Oracle with detailed information about the vulnerability and that it hasn't been exploited in the wild yet

EXPLOIT TOOLKITS

A factor that has made Java flaws more dangerous in recent years has been the inclusion by malicious groups and security developers of ready-to-launch attacks in exploit kits for hackers and penetration-testing kits for security professionals.

Major Java attacks have been found in kits such as Metasploit, Blackhole, VulnDisco SA Canvas, Eleonore, and Crimepack.

Security developers and hacker groups generally sell the kits online, which make the attacks available in an easy-to-implement form to a large number of people, including unskilled hackers who otherwise couldn't exploit the Java flaws, noted Gary McGraw, chief technology officer of software-security consultancy Cigital.

The incorporation of Java problems into exploit kits started in the third quarter of 2010, said Dan Guido, CEO of security vendor Trail of Bits. A recent company study of the 15 most popular toolkits showed that 11 had at least one Java exploit and two-thirds had at least two Java exploits. Now, Guido said, "almost all exploits kits available include at least one recent Java exploit."

Atif Mushtaq, senior staff scientist with security vendor FireEye, said, "We're seeing hundreds of groups exploiting this vulnerability. I guess people are not really updating their software."

JAVA BACKGROUNDER

S un Microsystems began work on Java as an internal project in 1990 and released it publicly in 1995. Oracle acquired Java when it merged with Sun in 2010.

Java enables the development and deployment of applications on multiple computing platforms. The key is the use of a Java virtual machine (JVM)—unique to each platform— that executes the same set of Java bytecode on all supported systems. The technology is used in many environments, such as embedded and mobile devices, PCs, enterprise servers, websites, and even supercomputers.

Java includes various forms of security, including cryptography, authentication, authorization, and a public-key infrastructure. A key part of the security model is use of a sandbox to keep code from escaping the Java environment and affecting the rest of the host system.

On its website, Oracle says that Java runs on 850 million PCs and billions of other devices worldwide.

Numerous companies—including Apple, Hewlett-Packard, and IBM—have developed their own proprietary JVMs. There are also many open source JVMs.

except for one zero-day attack in August.

Java 7 vulnerability

Several security investigators found another critical problem affecting Java SE 7, which hundreds of hackers have taken advantage of. Security researchers traced one of the early attacks to servers in China.

Researchers said the Java 7 vulnerability is highly exploitable, is not easily detectable by security software, and doesn't have effects that users would readily notice.

"It's not a bug, it's a design flaw," said FireEye's Mushtaq.

The Java Runtime Environment (JRE) vulnerability in Java 7 let hackers execute arbitrary code via an applet that exploits the flaw. "The applet must be specially crafted and use the vulnerable function in a certain way," noted Mushtaq.

Attackers use the Java 7 flaw to arbitrarily change the software's security settings, allowing malware to read, write, and execute code on an infected system.

"This method lets hackers obtain privileged references to private fields of arbitrary classes," Gowdiak explained. "This means attackers could obtain references to any field from the Java SE class environment and also set its value to an arbitrary value."

In essence, Mushtaq noted, this lets hackers change security settings and bypass Java's SecurityManager restrictions, enabling them to run code with full privileges. "It gives

TECHNOLOGY NEWS

the hacker access to a system. The first group of hackers used the flaw to deliver the Poison Ivy remoteaccess toolkit to infected systems. Now, hackers could use it to install and execute malware on the system." They could also steal information.

Investigators have identified phishing campaigns with e-mails—purportedly from major companies—that contain links that direct users to websites that launch the Java 7 exploits.

The SANS Institute, a security research and education organization, found an exploit that used a fake Microsoft e-mail—built with the template of a real message from the company—about an actual change to Microsoft's terms of service. The e-mail included a link that sent victims to a compromised website, which took advantage of the Java 7 vulnerability to deliver the Zeus Trojan. Zeus steals data from a victim's computer.

Another attack, which security vendor Websense uncovered, uses what is purportedly an e-mail from Amazon. The message asks the user to click on a link to verify an order. Clicking on the link takes the victim to a webpage containing a Java exploit. So far, researchers have found exploits on about 100 websites.

Security Explorations privately notified Oracle of the problem in April 2012, and FireEye did the same in August. However, Oracle reportedly didn't issue a patch—which also addressed two other Java issues until 30 August, outside its normal quarterly cycle of Java-related fixes.

Gowdiak said his team investigated the patch and found that it created yet another security issue that enabled exploitation of some bugs that Oracle hadn't addressed yet.

BEHIND THE PROBLEMS

A study by security vendor Qualys based on its BrowserCheck application—which scans browsers and plug-ins for problemsshowed that about 80 percent of the computers analyzed have Java enabled and 40 percent run versions with critical vulnerabilities. According to Qualys, this makes it the most vulnerable browser plug-in.

Referring to the sandbox-related and other Java problems that have cropped up this year, said Gary McGraw, chief technology officer of software-security consultancy Cigital."They look exactly like the Java vulnerabilities from 10 years ago. It's déjà vu all over again."

Java lets untrusted code run and assumes its sandbox will contain

Some experts say that Java now poses a major security risk for users.

any problems, McGraw said. "This assumes [Java's] sandbox is working properly and doesn't have implementation errors. But it's had a history of implementation errors."

"Java is beginning to show its age," he stated. "There are many newer platforms that might be better from a security perspective, such as Ruby on Rails, HTML5, and .NET."

Java is a good target for hackers because it is on so many computers, McGraw said. The software has come bundled with some operating systems, and many people have downloaded it, he explained. In addition, it works across platforms, so it is on many different types of machines.

Hackers look to attack a large, widely deployed code base like Java, noted FireEye's Mushtaq. Java has been around a long time and past flaws may still be in systems, at least in legacy code and older versions.

Some organizations and individuals use applications requiring older Java versions, which, if unpatched, could contain flaws. Even after updates, vulnerable legacy versions of the JRE usually remain on a computer.

Meanwhile, some users don't patch their software regularly, even when they could implement automatic updates. Others aren't even aware that their browsers have Java plug-ins enabled by default.

And now that vendors have made OSs and browsers safer, hackers are increasingly attacking browser plug-ins.

WHAT LIES AHEAD

To cope with possible problems, Mushtaq said, "If you don't need Java, you should just disable it or simply uninstall it."

Many programs for which Java is supposedly required will run as well or almost as well without the technology. However, disabling or deleting Java isn't convenient for all users. For example, Adobe's popular Creative Suite of graphic design, video editing, and Web development applications have required users to run a JRE.

Organizations that need the technology could utilize one browser that has Java disabled or deleted for general browsing and a second with Java enabled for websites that require it.

"Java wasn't really designed for modern [security] guidelines," Mushtaq said. "If Oracle wants its software to be popular, it needs to improve its code and design process. They need to put more effort into security. Now, they're putting more effort into features."

McGraw said, "It's not clear to me that Oracle's going to invest in pushing Java ahead. I hope they do. This latest round of problems has certainly shone the spotlight on Java in a negative way. Sometimes that spurs companies to do a better job." For Java to be more secure, he stated, they have to build security in from the beginning of the development process, which hasn't always been done in the past.

"Oracle needs to learn from Microsoft how to handle vulnerabilities," Mushtaq said. "A quarterly patch cycle is inadequate. I hope that they've learned their lesson."

With few exceptions, Oracle rolls out Java patches several times a year. For example, next year, the patches are slated for 19 February, 18 June, and 15 October.

In a dynamically changing security world, Oracle should issue patches more often and fix serious problems right away, said Security Exploration's Gowdiak. The failure to do so in the past, he contended, "speaks for itself."

Trail of Bits' Guido agreed that Oracle's response to Java problems has been inadequate. "There is not much they can do to remove all of these vulnerabilities from their codebase," he said. "They need to shift to patching quickly to drive down the value an attacker can derive from exploiting these flaws in the wild. Their quarterly patch schedule creates an enormous window of opportunity for an attacker."

Also, Gowdiak noted, "They could invest resources in developing a real sandbox, one that operates at a lower level of privilege so that when the JVM is exploited, attackers are still not in the position to do anything malicious to the exploited computer."

According to Gowdiak, "The education of software engineers should help decrease the number of new bugs introduced to the code. Engineers need to be aware of common Java security pitfalls. Code-review efforts should help catch security bugs prior to final product release."

uido said Java might not present such a security problem in the future because browser vendors are increasingly moving to protect their users and preventing Java and other plug-ins from running without users enabling it. "This will decrease the rates of success of attackers using exploits for Java and force them to move to other vectors to achieve the same effect." he stated.

FireEye's Mushtaq, on the other hand, said "Overall, I'm not really optimistic about the future of Java in the browser because of its vulnerabilities." And users don't need it because websites don't use it much any more.

The key to Java becoming safer, Cigital's McGraw said, is whether people care enough about security

SUBMIT TODAY!

Publishing in 2013

wider dissemination of information.

to push for it. He said he was pessimistic, adding, "I think we can make a prediction that [Java security] will [still] be a major issue in 2028."

Lee Garber is the IEEE Computer Society's senior news editor. Contact him at lgarber@computer.org.



IEEE TRANSACTIONS ON

EMERGING TOPICS

IN COMPUTING

Selected CS articles and columns CN are available for free at http:// ComputingNow.computer.org.

IEEE Transactions on Emerging Topics in Computing publishes papers on emerging aspects of computer science, computing technology, and computing applications not currently covered by other IEEE Computer Society Transactions. TETC is an open access journal which allows for



Submit your manuscript at: www.computer.org/tetc. Submissions are welcomed on any topic within the scope of TETC. Some examples of emerging topics in computing include:

- IT for Green
- Synthetic and organic computing structures and systems
- Advanced analytics
- Social/occupational computing
- Location-based/client computer systems
- Morphic computer design
- Electronic game systems
- Health-care IT
- Computer support for peer tutoring and learning via discovery or project work or field or lab work
- Creation and management of learning objects.

IEEE

IEEE (computer society

NEWS BRIEFS

New Supercomputer Is the World's Fastest

A 17.59-petaflops Cray Inc. supercomputer recently unveiled at the US Department of Energy's Oak Ridge National Laboratory has been recognized as the world's most powerful.

The Top500 project (www.top500. org), in which several academic and research experts list the world's fastest nondistributed supercomputer systems twice annually, ranked the Titan system in the top spot as of November 2012.

Titan's 17.59 petaflops—17.59 quadrillion (20×10^{15}) floating-point operations per second—makes it faster than the 16-petaflops IBM Sequoia computer at the US Lawrence Livermore National Laboratory. The Top500 project rated Sequoia as the world's fastest in its June 2012 rankings.

To create Titan, Cray upgraded its Jaguar machine—once the world's most powerful—at Oak Ridge. Cray replaced Jaguar's 224,256 CPUs with 299,008 faster AMD chips and added 18,688 Nvidia GPUs, which act as CPU accelerators. This lets Titan offer nine times Jaguar's performance with just one-third more CPUs. Running at just 2.3 petaflops, Jaguar required 7 megawatts of energy, enough to power 7,000 homes. Expanding the supercomputer without making it more energy efficient wouldn't have been practical, according to Oak Ridge. However, the new AMD chips are five times more energy efficient than Jaguar's CPUs. Thus, Titan requires only 8.2 megawatts of power, making its energy costs only a bit more than Jaguar's.

Oak Ridge plans to use Titan which has 710 terabytes of memory for computer simulations and other types of research in areas such as climate change, biofuels, nuclear energy, and new materials.

By 2016, the Department of Energy plans to upgrade the system so that it will perform 200 petaflops.

Researchers Turn the Tables on Sophisticated Hackers

Researchers with the country of Georgia's Computer Emergency Response Team planted malware on the computers of hackers who had been placing the same malicious software on numerous governments' computers to steal important national-security documents.



Cray's new 17.59-petaflops Titan supercomputer, used at the US Oak Ridge National Laboratory, is now the world's fastest system.



The researchers even used the malware to activate an embedded camera in one of the targeted machines and take photos of one of the hackers.

Investigators say their counterstrike enabled them to electronically infiltrate the group behind the Georbot Botnet, which used sophisticated techniques to break into the systems of government ministries, legislative bodies, banks, and nongovernmental organizations in Canada, China, France, Georgia, Germany, Russia, Ukraine, and the US in 2011 and 2012.

The hackers attacked Georgian computers by exploiting software vulnerabilities and by placing malicious hyperlinks on news-related and other webpages they hoped their intended victims would visit. When someone visited these pages, their browsers automatically executed a hidden script that uploaded the TrojanDownloader:JS/SetSlice malware installation tool. The tool used known Windows vulnerabilities to secretly download a malicious executable.

Once on a machine, the malware scanned Word, PDF, Excel, text, rich-text, and PowerPoint files for keywords such as "NATO" and "CIA" to identify documents worth stealing. The hackers also compromised machines' microphones and webcams to eavesdrop on victims.

Georgia's CERT investigators began investigating the attacks with help from the US Federal Bureau of Investigation, the US Computer Emergency Readiness Team, and various Eastern European cyberemergency-response groups.



The investigators installed malware on one of the hackers' computers by placing on one of their own machines an infected ZIP file named "Georgian-NATO Agreement," which they figured—correctly—the attackers would try to steal. The malware activated a camera on one of the hackers' computers and took photos of him. The investigators also were able to identify information such as his home city, ISP, and email address.

The Georgian CERT team said they tracked the recent attacks to the Russian Ministry of Internal Affairs. The two countries are on bad terms and fought a brief war in 2008.

Alliance Certifies Technology to Improve Wi-Fi Performance

The Wi-Fi Alliance, an international trade organization, has begun certifying technology designed to create direct links between mobile devices, bypassing access points (APs) and making wireless communications more efficient.

The efficiency could improve wireless activities such as media streaming, data backup, printing, and file transferring, and could extend devices' limited battery life.

The Wi-Fi Alliance has started certifying devices that comply with Tunneled Direct Link Setup (TDLS) technology, which is based on the IEEE 802.11z standard.

Typically, mobile devices communicate via an AP in a star topology. TDLS defines mechanisms that let Wi-Fi networks automatically set up a direct link between the devices while they are still connected to the AP. The technology does this by tunneling the protocol messages inside data frames.

MAN CLIMBS SKYSCRAPER STAIRWAY USING BRAIN-CONTROLLED PROSTHETIC LEG

A Washington state man who lost his lower leg in a motorcycle accident has climbed the 2,100 stairs of one of the world's tallest buildings using a prosthetic leg he controlled with his thoughts.

Zac Vawter of Yelm, Washington, put the limb through its paces by climbing the 103 floors of Chicago's Willis Tower, formerly called Sears Tower and once the world's tallest building.

Vawter—accompanied by 2,700 climbers from 38 states and seven countries ascended the stairs as part of the Rehabilitation Institute of Chicago's annual SkyRise Chicago fund-raising event. He is receiving treatment at the facility.

The surgeon who amputated Vawter's lower leg reattached to his hamstring muscle the nerves that previously carried signals past his knee. When Vawter thinks about climbing, signals from those nerves that formerly made his lower leg move are redirected to equipment in the 10-pound prosthetic limb. The equipment then activates and coordinates the leg's two motors, belts, and chains.

Brain-controlled prosthetic arms have been around for several years. Recently, the Rehabilitation Institute has been focusing on leg amputees, who outnumber those who have lost arms or hands.

Institute officials wanted Vawter to use the leg on the Willis Tower climb so that they could test it under extreme conditions. They say researchers still must refine the prosthetic leg and don't expect to release it commercially for several years.

The \$8 million project is funded by the US Department of Defense and includes researchers from Vanderbilt University, MIT, the University of Rhode Island, and the University of New Brunswick.



Zac Vawter, who lost a leg in a motorcycle accident, used a braincontrolled prosthetic limb to climb 2,100 stairs to the top of Chicago's 103-story Willis Tower.

Eliminating the AP connection makes communications more direct and bypasses AP congestion. In addition, TDLS lets devices communicate via the fastest Wi-Fi version and with the most security available, even if it's more than what the AP supports.

Devices could also measure network signal strength and decide whether a direct link would be better. The process starts when one device sends a discovery request to another via the network to determine whether it is also TDLS-compliant. If so, the target device sends a response detailing its capabilities.

The TDLS certification program will be available for TVs, smartphones,

tablet computers, cameras, printers, PCs, projectors, and gaming devices. Companies such as Broadcom, Marvell, Ralink Technology, and Realtek Semiconductor are already participating in the program.

Wi-Fi Direct technology also lets wireless devices connect directly, but it works via a limited wireless AP embedded in the devices and requires some user intervention.

US ISPs Adopt Plan to Curb Copyright Violations

After four years of preparation, some of the US's biggest ISPs are instituting a plan to discourage online users from violating copyrights held by content owners.

NEWS BRIEFS

The Copyright Alert System, supported by movie studios and music distributors, would impose escalating mitigation measures for repeat violations.

The Center for Copyright Information—which the Motion Picture Association of America and the Recording Industry Association of America founded to work on the new system—anticipates the plan will take effect later this year. ISPs AT&T, Cablevision Systems, Comcast, Time Warner Cable, and Verizon have agreed to participate.

Until now, many ISPs simply forwarded copyright-infringement notices from content owners to subscribers. There hasn't been a common approach for effectively dealing with ongoing violations.

With the Copyright Alert System, when an ISP receives notices from a copyright holder about infringements, the provider will notify the subscribers involved that someone used their account for content theft and inform them that this is illegal and that penalties could result. Failure to comply could lead to popup notifications, educational messages, Internet-speed reductions, or redirection to a landing page until the subscribers contact their ISP.

The Copyright Alert System doesn't call for ISPs to filter content that infringing users access or to terminate their service. However, federal law requires that ISPs have a termination policy for repeat copyright violators.

Subscribers accused of infringement who pay a \$35 filing fee to an arbitrator could ask for a review of their situation before an ISP imposes mitigation measures.

The Center for Copyright Information says content theft in the US annually costs \$58 billion in revenue for companies, 373,000 jobs, \$16 billion in employee earnings, and \$2.6 billion in tax revenue for federal, state, and local governments.

Smart Meters Are Not Always So Smart

Some smart utility meters are transmitting unencrypted information that hackers could intercept to determine, for example, if a building is occupied, according to University of South Carolina researchers. This could leave homes vulnerable to burglary or other problems.



Articles

IEEE Software seeks practical, readable articles that will appeal to experts and nonexperts alike. The magazine aims to deliver reliable information to software developers and managers to help them stay on top of rapid technology change. Submissions must be original and no more than 4,700 words, including 200 words for each table and figure.

Author guidelines: www.computer.org/software/author.htm Further details: software@computer.org www.computer.org/software Automatic-meter-reading devices are an early smart-meter technology found in about one-third of US homes and businesses. Some of these AMR devices automatically collect consumption, diagnostic, and other information from water, electric, and gas meters and send it via a network to a database for billing and analysis. Utility workers can also use the devices to gather information via a handheld collection wand.

The University of South Carolina team found that one type of AMR meter broadcasts information every 30 seconds, unlike those that send information only when requested by the host utility system.

Using a basic software-defined radio, an amplifier, and open source radio software, the researchers discovered that they could intercept meter signals, after learning that they could predict the various frequencies that the devices use to transmit signals.

Within a few days, after reviewing documentation on the Internet and information the meter makers publicly provide, the team determined how the devices' proprietary technologies worked.

The researchers then captured transmitted packets from 106 electric meters in the area they tested and identified the data. The information was sent in plaintext and included the device's identification number which could help identify the metered structure's address—and its utilityusage reading. They said they could determine a household's average power usage and identify periods of time when no one is likely to be home.

Newer smart-meter technology encrypts information that is transmitted wirelessly. However, the 46 million AMR meters in the US will require considerable time to replace.

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

GUEST EDITOR'S INTRODUCTION



Advancing the Science of Digital Forensics

Gary C. Kessler, Embry-Riddle Aeronautical University

Digital forensics, the branch of forensic science that focuses on the recovery and investigation of digital data, has applications in many contexts outside the courtroom, including research, policy enforcement, and intelligence gathering.

igital forensics combines methods from science, technology, and engineering to acquire and interpret information stored on digital devices for use in answering questions in court. Of course, these same methods allow for the acquisition of data for use in many contexts outside the courtroom, such as pure and applied research, policy enforcement, information security incident response, and intelligence gathering.

EARLY BEGINNINGS

My first foray into anything even remotely related to what we do today in computer forensics occurred in 1981. At that time, I was a programmer and coordinator of academic computing at a small college in Vermont, and our computer was an IBM System/34. Due to some catastrophic failure during shutdown the prior evening (yes, we shut the system down every night), the computer would not boot up the next morning. We found out later that this was the result of a corrupted volume table of contents (VTOC), the rough equivalent of today's file allocation table (FAT) or \$Bitmap.

A consultant systems programmer came in to show us how to recover the files by reconstructing the VTOC based upon the prior morning's routine printout of the hard drive contents (yes, we made such a printout every day or two). Since we couldn't use the computer without overwriting the files in what was now, essentially, unallocated space and had no PC-class systems at the time, we did the hex conversions by hand—on paper. It took us three days to reconstruct the VTOC and get back online.

That was the beginning of the computer forensics process and that was our environment: using a hex editor to get down to the bare metal of the hard drive and file system. And that's how it was for most of the next 15 years—hackers (when the term was implicitly White Hat and, indeed, noble, before Black Hat hackers hijacked the term) with an interest in investigations, most often in the law enforcement community, building rudimentary tools for use in looking deep into the computer and its file system.

By the late 1990s, computer science departments began taking serious notice of computer forensics, and academic programs in digital forensics were introduced in the early 2000s. And yet, it was not until 2009 that the American Academy of Forensic Sciences adopted digital forensics as a science.

Forensic sciences are largely based on Locard's exchange principle: every contact leaves a trace—if one

person hits another on the head with a tree branch, part of the tree branch stays on the victim's head and part of the head stays on the tree branch. This is as true in cyberspace as it is in real space. The challenge with digital forensics is to find the traces, interpret them correctly—and place a person's fingers on the keyboard.

A primary difference between digital forensics and the other forensic sciences is that practitioners advanced the field before the computer science community generally got involved with research and education. Thus, although digital forensics has been around for several decades, it is still a young science, and the body of peer-reviewed, academic literature that is essential for every science is currently relatively small—but it is growing.

IN THIS ISSUE

The cover features in this special issue are not intended to provide a survey of the digital forensics field, but rather to offer a snapshot of four interesting, varied, and relevant areas of research activity: computer forensics, network forensics, control system vulnerabilities, and mobile device security.

Computer forensics

Many computer science applications use hashing to build a data structure for use in mapping one set of table entries to another, such as a variable name to an address in memory. For these applications, the hash values tend to be short, and hash collisions—that is, two different entries having the same hash value—are to be expected. Cryptographic hashing has a different function, namely attempting to provide data integrity and a unique identifier for a data item, such as a file on a hard drive.

Although hashes are not unique over the entire universe of possible files, hash collisions are rare in practice. Thus, hashes can be used as the basis for searching and filtering files in a computer forensics examination to identify known contraband and malware as well as known trusted files. While using hash sets to assist in identifying files of interest in an examination streamlines the process, this approach has severe limitations. If as little as a single bit in the file is altered due to system error or deliberate user action, the file's hash is very different than the expected value.

In "Distinct Sector Hashes for Target File Detection," Joel Young and his colleagues from the Naval Postgraduate School and Johns Hopkins University describe a method for employing hashes on a per-sector basis rather than per-file to identify known files and discuss the efficacy of using this approach with various file systems.

Network forensics

Today, it is unusual to find a computer that is not connected to the Internet. Just as investigators need to

understand computer operating systems and file systems to get the most out of an examination of a computer, they also need knowledge of network applications and protocols when investigating a network.

While so-called hacker tools have become essential elements in a security officer's toolkit, these same tools can help in a network-based examination during a criminal or civil investigation, incident response analysis, or intelligence-gathering operation. Indeed, knowledge of network components, communication protocols, operating system utilities, the Transmission Control Protocol/ Internet Protocol (TCP/IP) suite, the Internet, application software (including malware, browsers, and peer-to-peer clients), and cloud applications (including social networks and file-sharing sites) is essential to understanding the network artifacts found on computers.

A primary difference between digital forensics and the other forensic sciences is that practitioners advanced the field before the computer science community generally got involved with research and education.

Network forensics, a specialty within the digital forensics field, requires its own set of processes. "Network Forensics: An Analysis of Techniques, Tools, and Trends" by Ray Hunt of the University of Canterbury and Sherali Zeadally of the University of the District of Columbia provides an overview of the network forensics space, a review of state-of-the-art tools and methodologies, and a glimpse into the future.

Control system vulnerabilities

Supervisory control and data acquisition (SCADA) systems are employed to monitor and manage industrial control systems and processes. SCADA systems can be as simple as a temperature-sensing device used to turn a heater on and off or as complex as a radiological-sensing device that manages the position of control rods. Such systems are used extensively in critical infrastructures as varied as chemical plants, oil refineries, utility distribution systems, waterway and dam management systems, transportation systems, and manufacturing plants.

The information security vulnerabilities of SCADA systems have been studied extensively, and the vulnerable nature of these systems is well-known. But in the case of a security breach, what are the computer forensics ramifications? What tools and techniques are available to the investigator? How could the process be improved? Indeed, what training is available? "SCADA Systems: Challenges for Forensic Investigators" by Irfan Ahmed and Golden G. Richard III from the University of New Orleans and Sebastian Obermeier and Martin Naedele from the ABB Corporate Research Center, Switzerland, explores these issues and describes the response from the digital forensics research community.

Mobile device security

During the past decade, mobile phones have evolved from being cool tech toys to become ubiquitous personal necessities. And, not surprisingly, cell phones—particularly, but not exclusively, those with cameras—are increasingly becoming the record keeper, victim, or instrument of criminal activity. Indeed, smartphones are essentially portable Internet terminals that, arguably, contain more probative data per byte examined than computers. At the same time, mobile phones have become a favored target of criminal hackers and a growing body of malware apps. Thus, mobile device forensics is a rapidly growing subspecialty of digital forensics.



Unrestricted access to today's groundbreaking research via the IEEE *Xplore*[®] digital library

IEEE offers a variety of open access (OA) publications:

- Hybrid journals known for their established impact factors
- New fully open access journals in many technical areas
- A multidisciplinary open access mega journal spanning all IEEE fields of interest
- Discover top-quality articles, chosen by the IEEE peer-review standard of excellence.

Learn more about IEEE Open Access www.ieee.org/open-access



he authors and I thank *Computer* for contributing to the efforts to advance the science of digital forensics by publishing this special issue.

Gary C. Kessler is an associate professor of homeland security at Embry-Riddle Aeronautical University, Daytona Beach, Florida; an adjunct associate professor at Edith Cowan University, Perth, Australia; and a member of the Northern Florida Internet Crimes Against Children (ICAC) Task Force. He received a PhD in computing technology in education from Nova Southeastern University, Fort Lauderdale, Florida. Contact him at gck@garykessler.net.

Selected CS articles and columns are available for free at http://ComputingNow.computer.org.

IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS

Special Issue on '-Omics'-Based Companion Diagnostics for Personalized Medicine

Manuscripts describing original research as well as reviews of emerging directions are solicited for this special issue, covering a range of circuits and systems topics including but not limited to

- DNA, RNA, proteins and small molecule sensors for companion diagnostics;
- technologies for '-omics' measurements;
- micro/nanofluidics technologies related to omics;
- healthcare and social impact of –omics circuits and systems;
- innovative circuit/system designs using –omics theories and principles, such as gene circuits and selfassembling DNA circuits, and biochemical network modules;
- circuit-based modeling and simulation of –omics systems such as gene regulatory and signaling networks;
- novel molecular sensing and imaging techniques for on-the-spot molecular diagnosis;
- portable devices for companion diagnostics; and
- other –omics methodologies and applications in personalized care delivery.

All manuscripts will be peer-reviewed and must follow the standard guidelines for manuscript preparation and submission posted on the IEEE TBioCAS website at www. ieee.org/tbiocas. Select the –Omics special issue, rather than Regular Issue, when uploading your manuscript on the https://mc.manuscriptcentral.com/tbcas submission site. **Manuscript submission deadline is 30 April 2013.**

Distinct Sector Hashes for Target File Detection

Joel Young, Kristina Foster, and Simson Garfinkel, Naval Postgraduate School Kevin Fairbanks, Johns Hopkins University

Using an alternative approach to traditional file hashing, digital forensic investigators can hash individually sampled subject drives on sector boundaries and then check these hashes against a prebuilt database, making it possible to process raw media without reference to the underlying file system.

orensic examiners frequently search disk drives, cell phones, and even network flows to determine if specific known content is present. For example, a corporate security officer might examine a suspicious employee's laptop for unauthorized documents; law enforcement officers might search a suspect's home computer for illegal pornography; and network analysts might reconstruct Transmission Control Protocol streams to determine if malware was downloaded. In these and many other cases, examiners typically identify files by computing their cryptographic hash—often with MD5 or SHA1 hash algorithms—and then searching a database for the resulting hash value.

Use of hash values for file identification is pervasive in digital forensics—every popular forensics package has built-in support. One of the most widely used databases is the National Software Reference Library (NSRL) Reference Data Set (RDS). Version 2.36, released in March 2012, contains 25,892,924 distinct file hashes (www.nsrl.nist.gov). Other databases are available to customers of specific companies and to law enforcement organizations. There are many limitations when using file hashes to identify known content. Because changing just a single bit of a file changes its hash, pornographers, malware authors, and other miscreants can evade detection simply by changing a comma to a period or appending a few random bytes to a file. Likewise, hash-based identification will not work if sections of the file are damaged or otherwise unrecoverable. This is especially a problem when large video files are deleted and the operating system reuses a few sectors for other purposes: most of the video is still present on the drive, but recovered video segments will not appear in a database of file hashes.

SECTOR HASHING

We are developing alternative systems for detecting target files in large disk images using cryptographic hashes on sectors of data rather than entire files. Modern file systems align the start of most files with the beginning of a disk sector. Thus, when a megabyte-sized video is stored on a modern hard drive, the first 4 kibibytes are stored in one disk sector, the second 4 KiBytes are stored in another disk sector, typically the adjacent one, and so on. (In our work, we distinguish between power-of-two-based sizes of digital artifacts, such as kibibytes, and power-of-ten-based sizes, such as kilobytes. See the "Decimal versus Binary Prefixes" sidebar for more details.) Furthermore, by sampling randomly chosen sectors from the drive, it is only necessary to read a tiny fraction of the drive to determine with high probability if a target file is present. This enables rapid triage of drive images.

We compare drive sector hashes to a hash database of fixed-sized file fragments, which we call *blocks*. The terms "sector" and "block" are often used incorrectly as syn-

DECIMAL VERSUS BINARY PREFIXES

oday there are two standards for representing sizes of files, storage systems, and memory banks: SI (International System of Units) decimal prefixes and IEC (International Electro-technical Commission) binary prefixes. SI decimal prefixes are commonly used to represent metric quantities. For example, the SI prefix gigamultiplies the value that follows by 10° ; thus, a gigabyte (Gbyte) is $10^{\circ} = 1,000,000,000$ bytes. In contrast, the IEC prefix gibi- multiplies the value that follows by 2^{30} ; a gibibyte (GiByte) is thus $2^{30} = 1,073,741,824$ bytes.

The confusion over prefixes dates back to the early days of computing, when K and M meant 1,024 and 1,048,576 when describing memory systems but 1,000 and 1,000,000 when describing storage systems. The difference in terminology resulted from the way that these systems were addressed. Memory was addressed by a series of binary lines, while electromechanical drums and disks were addressed by specifying a head, a track, and sector numbers: such numbers only map to even powers of two when the number of heads, tracks, and sectors are also even powers of two, and this is rarely the case due to manufacturing concerns.

For much of computing history, the fact that 1K sometimes meant 1,000 and sometimes 1,024 was not a major problem, as the correct size could be inferred from context and, in any event, the difference between 1,000 and 1,024 is not that great. However, the distinction became an issue in the 1990s as memory capacity mushroomed and

commonly used prefixes went from Ks to Ms and then Gs, resulting in a larger divergence between the power-of-two measurement and the corresponding power-of-ten measurement. The IEC accordingly proposed binary prefixes in 1996 and standardized their use in 1999. In 2008, the International Organization for Standardization adopted the IEC standard with the addition of prefixes for describing exbi- (2⁶⁰), zebi- (2⁷⁰), and yobi- (2⁸⁰) byte quantities.

Parto 4.6.

Despite this standardization effort, we live in a world in which 4-Gbyte memory sticks sold as system RAM can store 4,294,967,296 bytes of data but 4-Gbyte microSD (Secure Digital) cards for cell phones are only warranted to store 4,000,000,000 bytes of data. However, since those 4 billion bytes are organized in 512-byte logical sectors, the microSD card typically stores 7,812,500 (or more) sectors, a number that does not make much sense technically but makes a great deal of sense when the design of flash-based storage systems is considered. That is, flash systems contain more physical memory than they advertise, with the system removing bad blocks from service as the device ages. Thus, a "4-Gbyte" microSD card might actually have 8 million or even 9 million physical sectors, but those extra physical sectors are invisible to the operating system.

We expect use of IEC binary prefixes to increase with time. We use them here to describe block size and sector size, as they are typically multiples of 512 (2°). We use SI decimal prefixes to describe disk sizes, since that is the way they are sold by manufacturers.

onyms. For clarity, we use "sector" and "block" to refer to chunks of data extracted from drive images and files or file systems, respectively. Our approach depends on the existence of file blocks that only occur in a single distinct file. Experiments show that such *distinct blocks* comprise the vast majority of both executable files and user-generated content. Matches against block hashes shown not to occur elsewhere are strong evidence that a corresponding target file is or was present.

As the "Previous Work" sidebar describes, little work has been done on the use of sector hashes for file identification. However, sector hashing has numerous advantages over file hashing in forensic analysis. In many cases, using sector hashing with full media analysis—comparing every sector of the drive to an appropriate database—can detect a single block from a file that was once present. Alternatively, sector hashing can be combined with random sampling, making it possible to scan a terabyte-sized drive for the presence of select data in just a few minutes.

While sector hashing offers advantages when used for file detection in a forensic context, it also presents technical difficulties.

BLOCK SIZE AND HASH ALGORITHM

Two important design choices for using sector hashes are the block size and the hash algorithm.

Clearly, the block size must be small enough so that file blocks will align with drive sectors. The easiest way to assure this is to use a block size of 512 bytes, the sector size of most mass storage systems from the 1970s until quite recently. When presented with a device that has a larger sector size—for example, 2 KiBytes in CD-ROMs or 4 KiBytes in modern drives—the sectors could be divided into 512-byte blocks and hashed accordingly.

However, 512 bytes might be smaller than necessary. Many file systems use a 4-KiByte allocation size (NTFS has a default cluster size of 4 KiBytes for drives smaller than 16 Tbytes). In addition, using a 4-KiByte block size would reduce the hash value database's size by a factor of eight. The danger with a 4-KiByte block size is that a file system with a 4-KiByte allocation size might be used to write to a device with 512-byte sectors. If the blocks are not aligned on an eight-sector boundary, there is a risk that each set of eight sectors hashed would contain part of one block and part of another. The result is that no distinct blocks would be found.

This problem can be avoided in devices with 512-byte sectors by reading 15 sectors at a time, producing eight hashes: the first from sectors 0-7, the second from sectors 1-8, and so on. While multiple hashing does increase the computational costs of both hashing and database operations, the need for such hashing will decrease over time as 512-byte-sector devices are phased out of use.

We chose the MD5 hash algorithm, which is widely used within the forensic community and computationally fast. Although MD5 is no longer collision resistant, our technique relies on using hashes to match adversary data to target content—in fact, collisions actually facilitate the process.

COVER FEATURE

b

PREVIOUS WORK

While at the US Department of Defense Cyber Crime Center, Nick Harbour developed the dcfldd disk imaging tool (http:// dcfldd.sourceforge.net), based on GNU dd, that would compute a hash on a disk image as it was created. Harbour subsequently modified dcfldd to compute hashes over segments of the disk image so that if it was inadvertently modified, a chain of custody could be maintained for at least part of the image. He called this *piecewise hashing*. Jesse Kornblum's md5deep (http://md5deep.sourceforge. net) extended piecewise hashing to multiple files.

As part of his solution to the 2006 Digital Forensics Research Workshop (DFRWS) Data Carving Challenge, Simson Garfinkel introduced a new technique dubbed "the MD5 trick."¹ After finding the original challenge documents based on text fragments from the challenge description, Garfinkel computed the MD5 hash of 512byte file blocks and searched the challenge drive for matching 512-byte sectors. Using this technique, he identified all of the challenge files including a fragmented Microsoft Word file.

Three years later, Naval Postgraduate School researchers released frag_find, a tool that automates this process.² Sylvain Collange and colleagues called this approach *hash-based data carving*³ and explored the use of GPUs to speed the hashing load. They found that, with a powerful enough GPU, it is possible to simultaneously hash a block of data on subsector boundaries—for example, 1,024 bytes of data can be hashed in 512-byte chunks on 4-byte boundaries, creating 128 distinct hash values—although doing so dramatically increases pressure on the database.

In 2009, Simon Key developed the File Block Hash Map Analysis (FBHMA) EnScript, a dual-purpose tool that creates a hash map of file blocks from a master file list and searches selected areas of a target drive for the blocks.⁴ Like frag_find, however, FBHMA EnScript does not support billion-block hash databases or sufficiently fast lookup speeds to use sector hashing in full media analysis or random sampling.

References

Dem

- 1. S.L. Garfinkel, "DFRWS 2006 Challenge Report," 2006; http://sandbox. dfrws.org/2006/garfinkel/part1.txt.
- S. Garfinkel, "Announcing frag_find: Finding File Fragments in Disk Images Using Sector Hashing," 1 Mar. 2009; http://tech.groups.yahoo. com/group/linux_forensics/message/3063.
- S. Collange et al., "Using Graphics Processors for Parallelizing Hash-Based Data Carving," Proc. 42nd Hawaii Int'l Conf. System Sciences (HICSS 09), IEEE CS; http://hal.archives-ouvertes.fr/docs/00/35/09/62/PDF/ ColDanDauDef09.pdf.
- S. Key, "File Identification and Recovery Using Block Based Hash Analysis," lab presented at the annual Computer Enterprise and Investigations Conf. (CEIC), 2012; www.ceicconference.com/AJAX/ courseScheduleLightbox.aspx?id=1000018721.

UNDERSTANDING DISTINCT BLOCKS

Identifying files with sector hashes relies on the presence of distinct file blocks. A distinct block is one that does not occur anywhere more than once except as a block in a copy of the original file. Using distinct blocks as a forensic tool leverages two hypotheses:

• if a block of data from a file is distinct, then a copy of that block found on a data storage device is evidence that the file is or was once present; and

• if the blocks of that file are shown to be distinct with respect to a large and representative corpus, then those blocks can be treated as if they are universally distinct.¹

The first hypothesis is trivially true if we could know that a particular block is indeed distinct. Unfortunately, it is impossible to know this. On the other hand, we can determine the frequency of blocks in large collections of real files.

The frequency of distinct blocks

We counted the number of blocks in several million-file corpora that occurred once, twice, or more frequently. We call these *singleton*, *paired*, and *common* blocks, respectively. If paired and common blocks are extremely unusual, then it is reasonable to believe that singleton blocks are indeed universally distinct. Also, by examining the context of paired and common blocks, we might understand the root causes of their nondistinctness: a common method used to generate the data, an extrinsic process that created similar files, or some other kind of data-sharing mechanism.

For these experiments, we used three corpora modified to remove all duplicate files:

- Govdocs, a collection of 974,741 freely redistributable files downloaded from US government webservers (average file size: 493 KiBytes);²
- OpenMalware 2012, a collection of 2,998,898 malware samples (average file size: 417 KiBytes);³ and
- the 2009 NSRL RDS, a set of 12,236,979 block hashes for a collection of known, traceable software applications (average file size: 235 KiBytes).

To our knowledge, no previous studies have analyzed the co-occurrence of blocks across such a large number of files and file types. Using these corpora let us make some general conclusions about the frequency of distinct blocks.

We analyzed each corpus using both 512-byte and 4-KiByte blocks—the sector size of older and modern hard drives, respectively—except in the case of the 2009 NSRL RDS, for which 512-byte block hashes were not yet available. We also compared OpenMalware 2012 to the 2009 NSRL RDS to find the most common blocks across legitimate and malicious executables. Table 1 lists the incidence of singletons, pairs, and common sectors in the three corpora.

The vast majority of blocks in the corpora correspond to single, specific files. This is not surprising given that high entropy data approximates a random function. A truly random 512-byte block contains 4,096 bits of entropy. There are thus $2^{4.096} \approx 10^{1.200}$ possible different blocks, and all are equally probable. It is inconceivable that two

randomly generated blocks would have the same content. The randomness of user-generated content is less than 8 bits per byte, of course, but even for content that has entropy of 2 bits per byte, a 512byte block still contains 1,024 bits of entropy, again making it very unlikely that two blocks will be the same.

As Table 1 shows, all kinds of user-generated content, including word processing files, photos, and video, contain sectors that are not seen elsewhere that is, distinct blocks according to our definition. The frequency of distinct blocks in the OpenMalware 2012 and 2009 NSRL RDS datasets is significantly lower but still quite high. However, our experiments make it clear that it is impossible to assume a priori that a given singleton block is distinct.

Origin of nondistinct blocks

To better understand the root causes of nondistinct blocks, we analyzed the most common blocks from each corpus. Our original intuition was that blocks that had low entropy or that contained repeating byte patterns would occur frequently. We found that many of the common blocks indeed had these characteristics.

As expected, the block of all NUL (0×00) bytes was the most common block across all corpora. But we found other examples as well. For instance, there were more than 200,000 occurrences of an Adobe PDF internal data structure in the Govdocs corpus. Likewise, we found several common blocks that contained Microsoft Office internal structures.

Several high-entropy blocks were common in the OpenMalware 2012 dataset. We found that these blocks occurred in different files but always at the same byte offset. Further analysis revealed that the containing files were actually different variants of the same malware, as reported by several antivirus tools on VirusTotal.com. The repeated blocks did not appear in any legitimate files listed in the 2009 NSRL RDS corpus. Clearly, these blocks are unique to a specific malware family and not general executables or other system files.

Although traditional file identification techniques require each variant's hash, our findings show that shared blocks can identify some malware variants. We suspect that these common malware blocks are the result of handpatching existing malware and code reuse, or elementary attempts to change a file hash by adding bytes to the end of the file.

BLOCK HASH DATABASE

To develop a useful system for performing sector analysis, it is not enough to choose which or what size blocks should be used to capture a target dataset. It is necessary to, first, efficiently store the hashes for the target blocks

in three file corpora.											
No. of blocks	Govdocs	OpenMalware 2012	2009 NSRL RDS								
Block size: 512 bytes											
Singleton	911.4 M (98.93%)	1,063.1 M (88.69%)	N/A								
Pair	7.1 M (.77%)	75.5 M (6.30%)	N/A								
Common	2.7 M (.29%)	60.0 M (5.01%)	N/A								
Block size: 4 kibibytes											
Singleton	117.2 M (99.46%)	143.8 M (89.51%)	567.0 M (96.00%)								
Pair	0.5 M (.44%)	9.3 M (5.79%)	16.4 M (2.79%)								
Common	0.1 M (.11%)	7.6 M (4.71%)	7.1 M (1.21%)								
	1		1								

Table 1. Incidence of singleton, paired, and common sectors

and, second, check quickly enough to determine whether disk sectors are present in the dataset.

Performance requirements

Our goal is to create a database of one billion file block hashes that can be field deployed on a laptop. The database should be fast enough to support searches of hashes that are created by reading a consumer hard drive at the maximum I/O transfer rate (assuming that hashing is free). Given that it takes approximately 200 minutes to read the contents of a Tbyte-size hard drive, this translates to a database that can perform roughly 150,000 hash lookups per second. With a billion 512-byte block hashes, the database would allow identification of 512 gigabytes of known content, a number that is sufficient for many applications. Because hash values are evenly distributed, the database can be trivially parallelized using prefix routing.⁴ A cluster with 1,000 such databases could thus support 10¹² block hashes and address half a petabyte of known content.

Instead of hashing every sector of the drive, it is possible to conduct an exhaustive investigation sampling only one million randomly chosen sectors. Although the sample contains only 0.05 percent of the drive, there is a 98.17 percent chance of detecting 4 Mbytes of known content, provided that each of those 8,000 blocks is in the database.

This is an instance of the well-known "urn problem" in statistics, which describes the probability of pulling some number of red beans out of an urn that contains a mix of randomly distributed red and black beans. In this case, the red beans are distinct sectors, there are 8,000 (*C*) of them distributed randomly, there are two billion beans in total (*N*), and one million (*n*) are selected randomly. The probability *p* of not finding even a single red bean in *n* draws is

$$p = 1 - \prod_{i=1}^{n} \frac{\left((N - (i-1)) - C \right)}{(N - (i-1))} .$$

Applying this equation to 500,000 and 250,000 randomly selected sectors, we find that the chance of detecting

COVER FEATURE

4 Mbytes of known content, provided each of the 8,000 blocks is in the database, is 86.47 percent and 63.21 percent, respectively.

Note that the 4 Mbytes might be a single high-resolution JPEG or 40 medium-resolution JPEGs—the key issue is that there are 8,000 distinct blocks stored on sectors of the drive, and each random choice represents another chance to find one of them. Furthermore, because each sample is random, the distribution of the sectors on the drive is irrelevant—the chance of finding them with a random search is the same whether they are randomly distributed or clustered in a single location.

A 7,200-rpm hard drive can perform approximately 300 seeks per second. If the million randomly chosen sectors are sorted in advance, most systems could read all of them in 30 minutes; it is possible to read more data in the same time by increasing the read size to 8, 64, or even 128 sectors, although the statistical calculation becomes more complicated because many of the samples are now correlated, not strictly random. Thus, for the random sampling application, a database lookup of a few thousand transactions per second might be sufficient.

Our goal is to create a database of one billion file block hashes that can be field deployed on a laptop.

Designing the database

Neither conventional SQL databases such as MySQL, PostgreSQL, and SQLite nor NoSQL databases such as MongoDB have sufficient performance to support even high-speed random sampling. Using recent versions of each database on a Dell R510 server equipped with Dual Xeon E5620 2.4-GHz processors (each with 16 cores, a 12-mibibite cache, and 128-gibibite main memory), we got less than 1,000 lookups per second for databases containing one billion hashes.

To achieve better performance, we created our own purpose-built key-value pair store, where the key is a cryptographic block hash and the value identifies the source file and offset. We tested various custom-built solutions using hash maps, B-trees, red/black trees, and sorted vectors. In keeping with our goal for field deployment, the database is precomputed, finalized, and distributed to the client as a single file.

When looking for known content, we expect few of the sector hashes from a subject drive to actually be present in the database. We leverage this by checking a Bloom filter⁵ before checking the database. Bloom filters facilitate efficient probabilistic set-membership checking with a zero false-negative rate and a false-positive rate dependent on

the filter's parameters—the number of bits used in each hash (*M*) and the number of hash functions used (*k*).

When storing an item in the Bloom filter, we first hash the item k times, yielding k M-bit integers. We then set the corresponding bits in the filter. To test membership, we repeat the process, but instead of setting the bits, we check them, and if one or more bits are not set, the item cannot be in the filter. Note, however, that if all k bits are set, the item might or might not be present, as the bits might be aliases set for other items.

As we are storing the 128-bit MD5 hash values for the block, we do not need to compute *k* new hashes, but instead can partition the MD5 hash into *M* bit chunks. The resulting Bloom filter consumes 2^M bits or $2^M/8 =$ 2^{M-3} bytes. When M = 32, for example, the result is only k = 128/32 = 4 hashes, and the Bloom filter occupies 512 MiBytes of disk space. The theoretical false-positive rate of such a filter with a billion items is 13.48 percent, approximated by $P_{\rm fp} = (1 - e^{-kn/m})^k$. Doubling the size of the filter lowers the false-positive rate to 1.92 percent.⁶

One typically implements red/black trees, flat maps (essentially sorted vectors), and hash maps as in-memory data structures. To achieve persistence, we developed a data structure based on the boost::interprocess library, which allows transparently placing Boost C++ container implementations into memory-mapped files. For the B-tree back end, we selected Beman Dawes' proposed boost::btree library (https://github.com/Beman/Boost-Btree) and adapted our framework to support this back end. We used the Naval Postgraduate School's Bloom filter implementation.

After the user finalizes the database, the framework packs the data structures and releases extra space. In addition, it rewrites the B-tree with fully packed nodes at maximum density, enabling it to preload part of the tree into RAM.

Finally, the framework supports sharding the database into multiple chunks by the high-order bits in the key type.

Our key type is the 128-bit MD5 hash, and the record type is 64 bits partitioned to represent a file identifier and an offset yielding a 24-byte-per-record minimum cost. The flat map and B-tree back ends are most efficient, using less than 25 bytes per element, while the red/black tree and hash maps are less efficient, using 64 and 61 bytes per element, respectively. The red/black tree overhead comes from the tree nodes, while the hash map overhead results from unused buckets.

As cryptographic hashes are designed to be unpredictable, there is no similarity from one hash to the next, thus there is little locality of reference that the operating system can exploit when building the database. System RAM thus becomes the dominant factor in determining the time required. On one Intel Xeon E5620-based server (2.4 GHz, 12-MiByte L2 cache) with 32 GiBytes of RAM, it took 29

Table 2. Total transactions per second (TPS) for best execution.											
Bloom filter		Database		TPS at 1 M lookups		TPS at 1,200 seconds					
k	м	Size	Strategy	Size	Present	Absent	Present	Absent			
100 million records											
3	31	257 MiBytes	B-tree (preload)	2.3 GiBytes	35.3 K	49.5 K	161.3 K	1.8 M			
3	31	257 MiBytes	B-tree	2.3 GiBytes	11.6 K	565.8 K	156.8 K	2.3 M			
3	31	257 MiBytes	Hash map	5.3 GiBytes	13.9 K	656.9 K	641.9 K	3.0 M			
3	31	257 MiBytes	Flat map	2.2 GiBytes	28.2 K	746.9 K	356.4 K	2.6 M			
3	31	257 MiBytes	Red/black tree	6.0 GiBytes	12.9 K	694.5 K	187.0 K	2.7 M			
1 billion records											
3	34	2.1 GiBytes	B-tree (preload)	23 GiBytes	2.2 K	6.1 K	3.6 K	23.1 K			
3	33	1.1 GiBytes	B-tree	23 GiBytes	2.6 K	85.8 K	3.7 K	114.9 K			
3	33	1.1 GiBytes	Hash map	57 GiBytes	-	-	0.3 K	3.1 K			
3	34	2.1 GiBytes	Flat map	22 GiBytes	-	_	0.4 K	4.0 K			
3	33	1.1 GiBytes	Red/black tree	60 GiBytes	-	-	0.1 K	1.4 K			

Dashes indicate that 1 million queries were not completed in the 1,200 seconds allowed.

days to create a billion-record hash map, while it took less than four hours on a slower, AMD Opteron 6174-based system (2.2 GHz, 512-KiByte L2 cache) with 256 GiBytes of RAM.

We found that creating some locality by first building the database as a flat map and then converting to either a B-tree or hash map was faster than generating the B-tree or hash map directly. Likewise, we found that tuning the Linux operating system parameters dirty_ratio, dirty_background_ratio, and dirty_expire_centisecs to allow dirty pages to stay in memory longer improved performance by helping the OS use the disk cache more efficiently.

When fielding systems using the block hash database, system memory and I/O speed are the prime drivers. A drive triage system must be able to read disk sectors as fast as possible from a subject drive and test hashes of those sectors against the database. Large RAM allows caching more of the database, reducing I/O pressure. The database should be stored on a solid state drive (SSD) to further speed I/O, since every lookup will require one or more random seeks within the database file.

For systems supporting fixed sites, such as a customs and immigration checkpoint, a large memory server or cluster can maintain the entire database in RAM and support several triage stations over a gigabit network.

Back-end testing

We performed back-end testing with databases containing 100 million and 1 billion records. The tests were done on a laptop with 8 GiBytes of RAM, a 2.67-GHz processor, and a 250-Gbyte SSD attached via eSATA and USB2 drives. We performed additional testing on a desktop system with 24 GiBytes of RAM and spinning media. All runs were performed with 50/50 random blends of database hits and misses, which might be unrealistically pessimistic. To guarantee that no part of the database was already loaded in memory, we directed the OS to stop caching all disk files by syncing the disks and then writing a "3" into /proc/sys/ vm/drop_caches between each run.

Table 2 shows the read transactions per second against the 100 million and one billion record databases after one million lookups (2-384 seconds, depending on the row) and at 1,200 seconds, obtained with the four back-end strategies and B-tree with and without preload. Performance graphs for all of the runs are available at http://domex. nps.edu/deep.

The hash map offered the best performance at 100 million records, followed in order by the red/black tree, the flat map, and the B-trees. There was a factor-of-eight difference for queries that were present, but only a 40 percent spread for queries that were not present. In all cases, we observed that database misses were dramatically faster than hits, a result of prefiltering with the Bloom filter. The back-end performance is still relevant for misses, however, due to the false positives. We also observed that very large Bloom filters negatively impacted speed because of increased memory pressure. At one billion records, we obtained the best performance with M = 33 for the no-preload B-tree. Note that while the hash map outperformed the other strategies at 100 million records, B-trees overall dominated all other strategies by a factor of almost 30 (300 times better than the classic databases). The USB2 drive was roughly half the speed of the eSATA drive.

In sum, for billion-record hash databases, the B-tree is the best choice. For smaller datasets, the hash map

COVER FEATURE

is the fastest, but the flat map offers the best compromise between space and time (being among the smallest and tied for second place in speed), while the B-tree offers the poorest (requiring the most space and being the slowest). In either case, database lookups can be performed faster than sectors can be read from a drive being triaged.

SECTOR HASHING IN DIFFERENT FILE SYSTEMS

An advantage of sector hashing is the ability to process raw media without reference to the underlying file system. Doing so requires aligning the file data on sector boundaries. Fortunately, most file systems in use today align files in data units that consist of multiple disk sectors. These allocation units are variously called clusters, blocks, or sectors.

Sector hashing can aid in revealing the presence of encrypted files, provided they have not been re-encrypted.

Current file systems

The FAT (File Allocation Table) file system, introduced by Microsoft with DOS, has become the de facto format for storage devices such as thumb drives, external hard disks, and Secure Digital (SD) cards. All three FAT variants (12, 16, and 32) block-align data.

Microsoft developed exFAT (extended FAT) to address FAT's file size and performance limitations. It lacks the NTFS security features, but it can support file sizes greater than 4 GiBytes. Like its predecessors, exFat has a blockaligned data region.

NTFS (New Technology File System) is the default file system for the current generation of Windows. It uses a master file table (MFT) that has an entry for every file and directory. NTFS block-aligns large files but not files smaller than 1,024 bytes, which can be contained entirely in the MFT. The advent of 4-KiByte physical sector drives raises an issue, as they are not supported by products prior to Windows 8 and Server 2012. Instead, NTFS uses an "emulation mode" 512e to return a logical sector. While emulation can occur transparently, it also can induce file system clusters to cross physical sector boundaries, causing every physical sector to contain parts of different contiguous clusters. The techniques for working with 512-byte physical sector sizes also address alignment problems.

Ext4 is the default file system for most Linux distributions, including newer versions of Android. A major difference between Ext4 and Ext3 is that the former uses extents, while the latter employs a block pointer system. Despite this, Ext3 and Ext4 base their allocation on blocks, so file data is invariably aligned with the underlying storage media.

Next-generation file systems

Newer file systems handle data storage quite differently than their predecessors. Differences include data and metadata integrity mechanisms, copy-on-write transactions instead of journals, and built-in support for snapshots. However, sector hashing should work on these systems.

ZFS is the most mature next-generation file system, and the only one used in production environments. ZFS "blocks" are dynamically sized extents consisting of multiple sectors. If a file requires more space than the maximum block size, the system allocates multiple blocks. Since blocks are always aligned with the underlying storage media, there is no impact on sector hashing.

The B-tree file system (Btrfs) is poised to become the file system of choice for Linux. Although Btrfs uses extents of blocks to store large files, it can pack small files into the leaf block of the B-tree used to store file attributes. Thus, the system might not sector-align files smaller than 4 KiBytes.

The Resilient File System (ReFS) is Microsoft's upcoming innovation. Like Btrfs, ReFS makes extensive use of B-trees and extents. ReFS will block-align data, but whether it will pack small files into the B-trees is currently unknown.

Encrypted file systems

If an application such as Adobe Acrobat encrypts a file and transfers it to a different system, the encrypted data blocks will remain the same on the target media. Thus, sector hashing can aid in revealing the presence of encrypted files, provided they have not been re-encrypted.

Encrypted file systems, in contrast, present a significant problem. BitLocker for NTFS and ReFS and FileVault 2 for Apple's HFS+ encrypt data blocks as they are written to the storage medium and decrypt them when they are read back. Because each drive is encrypted with a different key, the same data will be encrypted differently on different drives. Thus, sector hashing will not work with these drives unless the block device is read through the file system after the decrypted.

uickly detecting documents or images of interest in digital media is critical to the forensic investigation process. Given a large disk or set of disks, an investigator requires an efficient triage process to determine if known bad or illegal files or previously unseen files that require additional analysis are present. Traditionally, forensic investigators use file-hashing tools to analyze the file system. However, file hashing has several shortcom-
ings: it does not work with files that have been modified in any way, it requires files to be intact, and it requires the ability to extract both allocated and deleted files from the subject media in a forensically sound manner.

Our approach for forensic identification of data searches disk sectors for distinct file blocks, rather than searching the file system for distinct files. Our method is agnostic to the file system and file type, and can analyze all portions of the media including unallocated space, metadata, and encrypted content. Sector hashing can also be parallelized since each sector is processed independently of all others.

Using sector hashes presents several challenges. The first is choosing an appropriate file block size that balances the ability to identify distinct chunks of files with the amount of data that needs to be stored and analyzed. A block hash database's large size makes it necessary to design a custom data store for this application. It is also useful to identify disk sectors that are likely to be nondistinct to minimize the number of queries made to the database without missing critical distinct file blocks.

A potential critique of our approach is that an attacker could defeat it by adding or removing semantically empty data to files, thereby changing the sector alignment. In response, we note that it would be easier to encrypt the entire drive and that many people are still not doing this.

Although sector hashing will be a powerful tool for media forensics, the size of the block hash database will surely hamper widespread adoption. Consequently, sector hashing is more likely to appeal to large organizations searching for stray copies of their own files and new variants of malware that they have already encountered, rather than by small organizations seeking to match their media against a database distributed by a vendor or the US government.

Acknowledgments

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the US government. The US government is authorized to reproduce, distribute, or authorize reprints for any reason notwithstanding any copyright annotations thereon.

References

- S. Garfinkel et al., "Using Purpose-Built Functions and Block Hashes to Enable Small Block and Sub-File Forensics," *Digital Investigation*, Aug. 2010, pp. S13-S23; www.dfrws. org/2010/proceedings/2010-302.pdf.
- S. Garfinkel et al., "Bringing Science to Digital Forensics with Standardized Forensic Corpora," *Digital Investigation*, Sept. 2009, pp. S2-S11; www.dfrws.org/2009/proceedings/ p2-garfinkel.pdf.
- D. Quist, "State of Offensive Computing," blog, 7 July 2012; www.offensivecomputing.net/?q=node/1868.

- 4. E.M. Bakker, J. van Leeuwen, and R.B. Tan, "Prefix Routing Schemes in Dynamic Networks," *Computer Networks and ISDN Systems*, Dec. 1993, pp. 403-421.
- 5. B.H. Bloom, "Space/Time Trade-Offs in Hash Coding with Allowable Errors," *Comm. ACM*, July 1970, pp. 422-426.
- P. Farrell, S.L. Garfinkel, and D. White, "Practical Applications of Bloom Filters to the NIST RDS and Hard Drive Triage," *Proc. Ann. Computer Security Applications Conf.* (ACSAC 08), IEEE CS, 2008, pp. 13-22.

Joel Young is an assistant professor in the Department of Computer Science at the Naval Postgraduate School. His research interests include computer forensics, algorithm design, and machine learning. Young received a PhD in computer science from Brown University. He is a member of the Association for the Advancement of Artificial Intelligence. Contact him at jdyoung@nps.edu.

Kristina Foster is a student at the Naval Postgraduate School. Her research interests include computer forensics and computer security. Foster received an MS in engineering, electrical engineering, and computer science from the Massachusetts Institute of Technology. Contact her at kmfoster@nps.edu.

Simson Garfinkel is an associate professor in the Department of Computer Science at the Naval Postgraduate School. His research interests include computer forensics, security visualization, and information policy. Garfinkel received a PhD in computer science from the Massachusetts Institute of Technology. He is a member of IEEE and ACM. Contact him at slgarfin@nps.edu.

Kevin Fairbanks is a cybersecurity research engineer in the Applied Physics Laboratory at Johns Hopkins University. His research interests include digital forensics and computer security. Fairbanks received a PhD in electrical and computer engineering from the Georgia Institute of Technology. He is a member of IEEE. Contact him at kevin. fairbanks@jhuapl.edu.

CN Selected CS articles and columns are available for free at http://ComputingNow.computer.org.

Intelligent Sustems The #1 Artificial Intelligence Magazine!

IEEE Intelligent Systems delivers the latest peer-reviewed research on all aspects of artificial intelligence, focusing on practical, fielded applications. Contributors include leading experts in

Intelligent Agents
 The Semantic Web
 Natural Language Processing

Robotics
 Machine Learning

Visit us on the Web at www.computer.org/intelligent



Network Forensics: An Analysis of Techniques, Tools, and Trends

Ray Hunt, University of South Australia Sherali Zeadally, University of the District of Columbia

Researchers in the growing fields of digital and network forensics require new tools and techniques to stay on top of the latest attack trends, especially as attack vectors shift into new domains, such as the cloud and social networks.

igital forensics is a science concerned with the recovery and investigation of material found in digital artifacts, often as part of a criminal investigation.¹⁻³ Digital artifacts can include computer systems, storage devices, electronic documents, or even sequences of data packets transmitted across a computer network.

Network forensics is a branch of digital forensics that focuses on the monitoring and analysis of network traffic. Unlike other areas of digital forensics that focus on stored or static data, network forensics deals with volatile and dynamic data. It generally has two uses. The first, relating to security, involves detecting anomalous traffic and identifying intrusions. The second use, relating to law enforcement, involves capturing and analyzing network traffic and can include tasks such as reassembling transferred files, searching for keywords, and parsing human communication such as emails or chat sessions.

A GROWING FIELD

The evolution of network security, as well as its associated forensic processes and related toolsets, is largely driven by recent advances in Internet technologies. As more aspects of our daily lives migrate to online systems and databases—where they are subject to criminal activity the need for sophisticated analysis tools is increasing accordingly. Some commonly stated reasons for using network forensics include

- analyzing computer systems belonging to defendants or litigants;
- gathering evidence for use in a court of law;
- recovering data in the event of a hardware or software failure;
- analyzing a computer system after a break-in;
- gaining information about how computer systems work for the purposes of debugging them, optimizing their performance, or reverse engineering them;
- collecting and analyzing live data packets to detect and potentially prevent a malicious attack; and
- learning more about zero-day attacks, particularly through the use of honeypots and honeynets.

This list merely scratches the surface of what network forensics can do as part of risk assessment and data recovery; the following example demonstrates the vital role this technology can play in an investigative process.

The TCP/IP family of Internet protocols carries most of today's online traffic information, and attackers can manipulate these protocols to spoof addresses or embed malware. In particular, they can embed data in unexpected places such as the options field in an Internet Control Message Protocol packet. ICMP messages are used to communicate error information, such as a requested service's unavailability or a host that cannot be reached, or to indicate congestion, such as a downstream router's lack of buffering capacity. There is no expectation that ICMP packets will carry application data, so most firewalls and intrusion-detection/prevention systems do not examine their contents, resulting in a concealed channel that most network security systems simply cannot see.

Some intrusions can be difficult to detect and subsequently analyze—for example, a simple port scan might hide a serious stealthy attack on a crucial system resource. Intrusion analysis and the collection of forensically sound data thus seek answers to the following questions:

- Who generated the (incoming) intrusion or (outgoing) data transfer?
- What equipment and services were involved in gaining entry?
- Where did the intrusion come from, and what parts of the infrastructure were affected?
- Was the attack made possible because by limitations or weaknesses in incoming or outgoing security mechanisms?

This real-time analysis process involves collecting, storing, and tracing data and then recovering the system, all while continuously scanning traffic and logs. As Figure 1 shows, the recovery process starts with security and then moves into forensic analysis—who perpetrated the attack and from where—followed by getting the system going again.

A CONTINUING EVOLUTION

Researchers in the growing fields of digital and network forensics require new tools and techniques to stay on top of the latest attack trends, especially as attack vectors shift into new domains, such as the cloud and social networks.

Several open source tools are available for general forensic analysis of open ports, mapped drives, and open or mounted encrypted files on live computer systems. The currently available open source tools include Sleuth Kit (www. sleuthkit.org), Scalpel (www.digitalforensicssolutions. com/Scalpel), and DEFT Linux (Digital Evidence & Forensics Toolkit, www.deftlinux.net); well-known commercial products include EnCase (www.guidancesoftware.com), FTK (Forensic Toolkit, www.accessdata.com), ProDiscover (www.techpathways.com), and Helix (www.e-fense.com/ products.php).

Some important differences

Traditionally, researchers performed computer forensics on stored or static data—for example, the contents of files or images on hard drives. This dead or postevent analysis is also referred to as reverse engineering. But in recent years, there has been an increased emphasis on live system analysis, examining network traffic as it arrives.

Recent network forensics work has taken this one step further, focusing on live packet capture because packets



Figure 1. Real-time detection, recovery, and forensic analysis process. The process collects, stores, and traces data and uses it to perform real-time recovery while carrying out forensic analysis to determine the source of an attack.

are not normally stored upon arrival at their destination. Other types of live capture focus on attacks that leave no trace on the computer's hard drive because the attacker only exploits information in the computer's volatile memory, including encryption keys.

Network forensics is concerned with monitoring network traffic to see if anomalies exist and whether they indicate an attack or could lead to one. The objective is to determine the attack's nature and then capture, store in a forensically sound manner, analyze, and, finally, present some visual form of it. Because an attacker might have erased all the log files on a compromised host, networkbased evidence might be the only material available for forensic analysis.

Unlike digital forensics, which retrieves information from a computer's disks or other storage devices, network forensics retrieves both traffic and information about which ports it used to access the network. Frequently, investigators and adversaries use the same tools: one using the tools to cause an incident and the other using them to investigate it. Current examples include Wireshark, TCP-Dump, the NetScanTools Pro toolkit (www.netscantools. com/nstpromain.html), and the HENPA framework.⁴ NetScanTools includes tools for network information gathering and security testing; IP/MAC address ranges and locations; visible, hidden, and writable shared folders; TCP/ UDP port and DHCP analysis; SMTP and SNMP activity; and conventional packet viewers.

It might be possible to trace an attack back to its source—or at least to the ISP that carried the attack—while the attack is in progress, but in many cases, this type of analysis happens after the event. An essential aspect of live network forensics is the ability to collect data from the network fast enough so that no information is lost, which requires very fast processors and I/O devices as well as significant storage capacity. The best way to capture the data is to use a moving window of hours, bounded by the time by which an attack would be expected to be discovered. Sustained attacks of even 10 Gbps make significant demands on both the storage and processing of network forensic data, so, for example,

- 10-Gbps traffic flow with a two-hour sliding window requires 10 Tbytes of storage, and
- 20-Gbps traffic flow with a 12-hour sliding window requires 1 Pbyte of storage.

Because of the sheer sizes involved, only a sample of packets can be stored for subsequent analysis. The processing of network forensic data in real time demands large-scale distributed and parallel processing engines as well as the flexibility to customize the process. Even a sliding window of a few hours covering the duration of real-time traffic of interest could require terabytes of storage. The largest distributed denial-of-service (DDoS) attack on an ISP was recorded in 2010 and reached nearly 100 Gbps.¹ DDoS attacks of this size represent a hundredfold increase over the past 10 years, so current-generation network forensic analysis can require the implementation of parallel processing using supercomputers or Beowulf cluster computing.

It is unlikely that a single tool will suffice for any investigation—more than likely, investigators will use a combination of tools.

A suitable tradeoff between security and performance is also important. Complex tools and techniques could significantly affect the system and have serious consequences—for example, a disruption in communications induced by a network forensic tool's complexity could interrupt the infrastructure's fundamental functionalities due to their strong interrelationships.

Originally, digital and network forensics were viewed as closely related technologies, but in reality, the two are quite different. Digital forensics is driven largely by law enforcement organizations and the need to gain sound evidence to resolve criminal activities. Network forensics has evolved in response to the hacker threat and has strong links with security architecture, including firewalls, port blocking and filtering, threat assessment and surveillance, intrusion detection, and data loss prevention.

In digital forensics, the investigator and the attacker are at two different skill levels, with the investigator supposedly at a higher level. In network forensics, the investigator and the attacker theoretically have the same skill levels. The network forensics specialist uses many of the same tools and engages in the same set of practices as the person being investigated.

Common tools and techniques

Tools to assist with network forensics come in a variety of forms: some are merely packet sniffers, whereas others might focus on fingerprinting, mapping, location identification, email traffic, URLs, traceback services, and honeypots. Table 1 summarizes some of the tools more commonly used to support network forensic investigations, along with their properties.

It is unlikely that a single tool will suffice for any investigation—more than likely, investigators will use a combination of tools. For example, if the focus is on traffic analysis, and the investigators already understand the malware traffic's nature, basic Unix utilities such as Ngrep, TCPDump, or Omnipeek/Etherpeek might be sufficient. But when the investigation merits using a traffic analysis engine, tools such as Wireshark, NetMiner, Driftnet, or Xplico might be required. For commercial organizations, tools such as NetWitness offer a powerful range of analysis options for network monitoring or assessing insider threats, zero-day exploits, and targeted malware.

Cloud computing challenges

To date, although many systems are moving into the cloud, little research has been performed on the tools, processes, and methodologies necessary to obtain legally defensible forensic evidence in that domain.⁵ Most investigations require evidence retrieval from physical locations, so cloud network forensic must be able to physically locate data with, for example, a given time-stamp and trace network forensic data at a given time period, taking into account the authority at different locations.

Although the live and dead forensics categories still exist, cloud models present new challenges because network data is often difficult to locate, thus acquisition might be challenging or even impossible. Analysis without acquiring network data is impossible, so network forensic tools must evolve yet again, forming an amalgam of current live and dead collection and analysis methods, as well as incorporating the intelligence to find and predict artifacts based on forensic heuristics.

When conventional network forensic tools work, the only aspect that a cloud tool changes is the collection method. For situations in which acquisition is difficult, new network forensic tools will need to visualize physical and logical data locations in a way that indicates both obtainable and unobtainable data and metadata. In addition to visualization, forensic tools will need to use the cloud as a discovery engine for network forensic analysis. So, for example, a network forensic compilation that contains unobtainable data will need to be submitted to a cloud environment for heuristic and signature-based analysis. This is similar to the way network forensics investigators use antivirus engines to converge collections of incomplete data into reliable presentations as the number of submissions increases.⁶

Table 1. Tools commonly used to support a variety of network forensics investigations.					
ΤοοΙ	Features and advantages	Website	Attributes		
TCPDump, Windump	Command-line network packet analyzer that supports network forensic analysis	www.tcpdump.org; www.backtrack-linux.org/ backtrack-5-release	F		
Ngrep	Simple, low-level network traffic debugging tool	http://ngrep.sourceforge.net	F		
Wireshark	Widely used network traffic analysis tool; forms basis of network forensics studies	www.wireshark.org	F		
Driftnet	Listens to network traffic and picks out images; used in Backtrack v5	http://linux.softpedia.com/progDownload/ Driftnet-Download-15905.html	F		
NetworkMiner	Network forensic analysis tool that can be used as a passive network sniffer/packet-capturing tool	www.netresec.com/?page=NetworkMiner	F		
Airmon-ng,Airodump- ng, Aireplay-ng, Aircrack-ng	Widely used suite of low-level traffic analysis tools for wireless LANs; used in Backtrack v5	www.backtrack-linux.org/backtrack-5-release	F, L, R, C		
Kismet	Network detector, network packet sniffer, and intrusion-detection system for wireless LANs	www.kismetwireless.net	F		
NetStumbler	Widely used wireless LAN analysis tool for devices and network traffic analysis	www.netstumbler.com	F		
Xplico	Network forensic analysis tool that allows for data extraction from traffic captures; used in Backtrack v5	http://packetstormsecurity.org/search/?q=Xplico	F		
DeepNines	Provides real-time identity-based network defense for content and applications, along with basic network forensics	www.deepnines.com	F		
Argus	Used for network forensics, nonrepudiation, detect- ing very slow scans, and supporting zero-day attacks	www.qosient.com/argus	F, L		
Fenris	Suite of tools for code analysis, debugging, protocol analysis, reverse engineering, network forensics, diagnostics, security audits, vulnerability research	http://lcamtuf.coredump.cx/fenris/whatis.shtml	F		
Flow-Tools	Software package for collecting and processing NetFlow data from Cisco and Juniper routers	www.splintered.net/sw/flow-tools	F, L		
EtherApe	Graphical network monitor for capturing network traffic	http://etherape.sourceforge.net	F		
Honeyd	Improves cybersecurity by providing mechanisms for traffic monitoring, threat detection, and assessment	www.citi.umich.edu/u/provos/honeyd	F		
Snort	Widely used, popular tool for network intrusion detection and prevention, as well as for network forensic analysis	www.snort.org	F		
Omnipeek, Etherpeek	Low-level traffic analyzer for network forensics	www.wildpackets.com	F, L, R		
Savant	Appliance for live forensic analysis, surveillance, network analysis, and critical infrastructure reporting	www.intrusion.com	F, R		
Forensic and Log Analysis GUI	Log file analysis combined with network forensics; Python implementation	http://sourceforge.net/projects/pyflag	L		
Dragon IDS	Provides network, host intrusion detection and network forensic capture analysis	www.enterasys.com; www.intrusion- detection-system-group.co.uk/dragon.htm	F, R, L, C		
Infinistream, nGenius	Appliance for network forensics, incident analysis combined with session reconstruction and playback	www.netscout.com/products/enterprise/nSAS/ ngenius_analysis/Pages/nGenius_Forensic_ Intelligence.aspx	F, R, C		
RSA EnVision	Provides live network forensics analysis, log management, network security surveillance, data leakage protection	www.emc.com/security/rsa-envision.htm	F, L, R, C, A		
NetDetector	Appliance for network forensic analysis, network security surveillance, signature-based anomaly detection	www.niksun.com	F, R, C, A		
NetIntercept	Appliance for network forensics, monitoring, and analysis	www.niksun.com/sandstorm.php	F, R, C, A		
NetWitness	Addresses network forensic analysis, insider threat, data leakage protection, compliance verification, designer malware, and 0-day detection	www.netwitness.com; www.rsa.com	F, L, R, C, A		
Solera DS	Appliance for live network forensics, application	www.soleranetworks.com/products/appliances	F, R, C, A		



Figure 2. Real-time adaptive security incorporating network intrusion detection and forensics logging.

New frontiers in network intrusion

Intrusion detection systems (IDSs) monitor network and system activity for malicious behavior or policy violations. Some systems might attempt to stop such an intrusion, but work on developing the ability to dynamically modify firewall rules in the face of an attack is still in its infancy. The combination of network forensics and intrusion detection might be adequate for a user's home system, when manual intervention is appropriate, but most intrusion-detection or prevention systems focus only on identifying possible incidents, logging information, and reporting such attempts. Therein lies the problem: any system of realistic scope or size that supports sensitive client data must include an automated combination of intrusion analysis with network forensic log analysis as well as dynamic feedback to modify access rules in the face of real-time attacks.

Some attackers explore a victim's network prior to launching an attack. A sophisticated IDS might be able to correlate data obtained from the attacker's reconnaissance—possibly along with additional log data—to either forecast the attack or to obtain better forensic evidence during or after the attack. However, although some progress has been made recently with distributed IDS architectures,⁷ many IDSs cannot detect complex intrusions and distributed or coordinated attacks. Figure 2 shows the components required to provide a forensically sound intrusion-detection and prevention system. The combination of such a system with reactive firewalls, traffic storage, and subsequent analysis provides a powerful forensic security architecture.

APPLYING NETWORK FORENSICS IN CRITICAL INFRASTRUCTURES

The critical infrastructures that attackers seek to launch their strikes against include not just the traditional areas associated with cybersecurity attacks, such as the water supply, traffic systems, and power and gas plants, but also any network system that could be considered critical to electronic commerce operations. The secure operations of, for example, banking, airline, communications, weather forecasting, and a host of other business enterprises depend almost entirely on a safe and secure network, which implies significant security issues for the ISPs and telecom operators that provide network infrastructures for these organizations.

Botnets

The environment in which an organization's user base operates continues to grow more hostile with the release of sophisticated and polymorphic malware such as Conficker, Koobface, and Zbot. DDoS attacks from botnets are a particularly serious global threat.¹ Botnets are now available for hire from criminal syndicates and can be used to mount DDoS attacks as well as to harvest identities and financial credentials. Additional attack methods include DNS spoofing and cache poisoning, Border Gateway Protocol (BGP) route hacking, and VoIP infrastructure flooding.

The network forensic process must be able to detect scans and probes outside the firewall and then use this data to inform a security information event management (SIEM) system that includes network forensic analysis tools. Although several SIEM engines are available, only a few include a logging system from which such data can be used later as evidence. A progressive threat assessment requires software monitors to trigger an alert when unusual time-based IP address patterns occur inside the secure perimeter, indicating a potential botnet intrusion.



Figure 3. Monitoring and analyzing worm propagation using a sinkhole. In this example, an infected host is scanning for others targets to infect. Because it sucks in any internally originated traffic destined for detailed network forensic analysis, the sinkhole can detect a worm's scanning activity.

Network forensics can play a pivotal role in botnet attack threat assessment because the SIEM system not only handles log files in a forensically sound manner, but it also stores a moving window of log data as evidence for potential future activity. Real-time adaptive feedback resulting from this analysis could potentially avert or minimize a real-time attack via firewall rule adaptation.

Wireless networks

Wireless forensics, a subdiscipline of network forensics, provides the methodology and tools required to collect and analyze wireless network traffic. This new area has some techniques in common with fixed networks, along with some differences. Evaluating wireless networks from a forensic computing perspective helps to understand the current state of wireless misuse as well as the various tools and techniques used for identification, containment, and analysis. This research reveals the limitation of current tools and procedures for forensic computing investigations on wireless devices and networks, and highlights various forms of misuse that might escape detection by forensic investigations. Some commercial players in fixed network forensics also claim wireless capabilities, at least for WLANs. Wireless network forensics requires these tools to analyze 802.11 headers and corresponding protocol data flows. From an open source perspective, there are no well-known, dedicated network forensicsanalysis tool alternatives.

Sinkholes

A sinkhole is a security tool that has the potential to accept, analyze, and forensically store attack traffic. Originally, ISPs used sinkholes to draw attack traffic away from a customer; more recently, they have used them to monitor attacks, detect scanning activity from infected machines, perform a forensic analysis, and generally monitor for malicious activity. Figure 3 shows how the sinkhole gateway router can be used to forward attack traffic to a sinkhole target router via a switch for basic Wireshark and TCP-Dump sniffing, intrusion detection, and forensic analysis.

Figure 3 also shows how a sinkhole can be used to monitor internally generated worm propagation. In this example, an infected host is scanning for other computers to infect. It pulls in any internally originated traffic destined for either bogon addresses or dark IP address

space—bogon is unallocated address space, and dark IP space is allocated but unused. Consequently, the worm's scanning activity can be detected at the sinkhole. Monitoring the dark IP address space is essential because future worms might be written to purposely ignore such address blocks.

Additionally, a sinkhole can remove other noise from the network, such as reflector or backscatter traffic, which often indicates the start of a worm or DDoS attack. Backscatter traffic can occur as the result of large-scale DDoS attacks that use spoofed source addresses. A high increase in backscatter traffic could be the first sign of a new worm's release. Forensically sound event logs and network traffic storage of this traffic is therefore crucial.

EMERGING NETWORK FORENSICS AREAS

Network forensics has important roles to play in new and developing areas related to social networking, data mining and digital imaging, and data visualization.

Social networks

Social networking sites such as Google+, Facebook, Twitter, and YouTube have expanded astronomically in recent years, but because the success of such sites depends on the number of users they attract, there is pressure on developers to design systems that encourage behavior that increases both the number of users and their connections. Security has not been a high priority, leading to the emergence of inevitable security risks.

Obviously, there is a need for network forensic tools that address such an important area of usage, but to date, only traditional digital and network forensic tools are available.^{8,9}

Data mining

Forensic profiles can be created using data mining technology, which provides a way to discover relevant patterns, thus generating profiles from large quantities of data. Although there has been significant work in the areas of extracting and analyzing digital evidence from physical devices such as hard disks, less work has been reported on data mining in portable storage devices such as flash drives, cell phones, digital cameras, radio frequency identification devices, compact disks, and iPods.¹⁰

The extraction of historical data from supervisory control and data acquisition (SCADA) systems, which are widely used to monitor and control equipment in various industries such as oil and gas refining, water and waste control, and transportation, is an important area that draws on the combination of data mining and network forensics.

There is currently no generic model for understanding the processes necessary to gather digital evidence from SCADA systems. However, such a model is needed to enable incident response, intelligence gathering, digital evidence collection, and legal action against system intruders. There is a distinct difference between the process of network forensics-based data mining investigations (where time-based data is analyzed to detect potential malware intrusion) and incident recovery and response (where the key purpose is to respond to an alarm and implement recovery).

Some work has been done to incorporate the use of decision trees as well as naive Bayesian, a priori, and neural network techniques.¹¹ Recently proposed architectures also incorporate mechanisms for monitoring process behavior, analyzing trends, and optimizing plant performance.¹²

Digital imaging and data visualization

Researchers have developed numerous state-of-the-art tools to assist in conducting digital crime investigations. However, digital investigations are increasingly complex and time-consuming due to the amount of data involved. The visualization of data obtained from such investigations is a new and developing area and has the potential to display significant volumes of data where the dimensionality, complexity, or volume prohibits manual analysis.

Data visualization is the graphical interpretation of high-dimensional data, which is particularly appropriate for obtaining an overall view and locating important aspects within a dataset. This is useful in network forensics because the data encountered in digital investigations is often significant in size, multidimensional, and complex. Consequently, obtaining an overall view can help digital investigators obtain a better understanding of the data and identify important aspects to assist in the recovery of appropriate digital evidence.¹³

ell-funded hackers, criminals, and terrorists are hiding data in new ways. Antiforensics tools are now as sophisticated as the tools they endeavor to defeat—Metasploit, for example, has developed three tools that have the potential to devastate automated forensic analysis tools.

Law enforcement agencies strive to both prevent such attacks and catch the perpetrators using the latest security and forensics tools. However, this work requires the design and implementation of a secure and forensically sound architecture. Resource limitations are a problem, and the process of developing innovative solutions will need to include computer software manufacturers, security tools providers, antimalware organizations, forensic tool providers, ISPs, and telecommunications companies. It will also require the dedication and diligence of users themselves.

Regardless of the exact tools used, developers must build forensics capabilities into security systems. Botnet attacks, for example, generate traffic logs, and tracing them in real time requires progressive threat assessment as attacks move through the system. Determining how the attacker gained access or how information leaked out of the organization while simultaneously quantifying the scale and impact of an attack as it happens are the very foundation of cohesive security and forensic processes.

Acknowledgments

We thank the anonymous reviewers for their constructive comments, which helped us improve the article. We also express our gratitude to Gary Kessler and Ron Vetter for their support and encouragement throughout the preparation of this article.

References

- 1. D. Anstee, "Worldwide Infrastructure Security Report," vol. 7," Arbor Networks, Feb. 2012; www.arbornetworks. com/report.
- 2. NIST Information Testing Laboratory, "Computer Forensics Tool Testing Program," 2012; www.cftt.nist.gov.
- 3. NIST, 2012; "Guide to Integrating Forensic Techniques into Incident Response," http://csrc.nist.gov/publications/nistpubs/800-86/SP800-86.pdf.
- 4. J. Broadway, B. Turnbull, and J. Slay, "Improving the Analysis of Lawfully Intercepted Network Packet Data Captured for Forensic Analysis," Proc. 3rd Int'l Conf. Availability, Reliability, and Security (ARES 08), IEEE CS, 2008, pp. 1361-1368.
- 5. K. Ruan et al., "Cloud Forensics: An Overview," Proc. 7th IFIP Conf. Cloud Computing, Centre for Cybercrime Investigation, Univ. College Dublin, 2012; http:// cloudforensicsresearch.org/publication/Cloud_Forensics_ An_Overview_7th_IFIP.pdf.
- 6. S. Zimmerman and D. Glavach, "Cyber Forensics in the Cloud," IA Newsletter, vol. 14, no. 1, 2011, pp. 4-7; http://iac. dtic.mil/iatac/download/Vol14_No1.pdf.
- 7. C. Zhou, C. Leckie, and S. Karunasekera, "A Survey of Coordinated Attacks and Collaborative Intrusion Detection," Computers & Security, vol. 29, no. 1, 2010, pp. 124-140.
- 8. J. Cheng et al., "Forensics Tools for Social Network Security Solutions," Pace Univ., May 2009; http://csis.pace. edu/~ctappert/srd2009/a4.pdf.
- 9. H.V. Zhao et al., "Behavior Modeling and Forensics for Multimedia Social Networks: A Case Study in Multimedia Fingerprinting," IEEE Signal Processing Magazine, Jan. 2009, pp. 118-139.
- 10. V.H. Bhat, "A Novel Data Generation Approach for Digital Forensic Application in Data Mining," Proc. 2nd Int'l Conf. on Machine Learning and Computing (ICMLC 10), IEEE, 2010, pp. 86-90.
- 11. F. Camastra, A. Ciaramella, and A. Staiano, "Machine Learning and Soft Computing for ICT Security: An Overview of Current Trends," J. Ambient Intelligence and Humanized Computing, Oct. 2011; doi:10.1007/s12652-011-0073-z.
- 12. T. Kilpatrick et al., "An Architecture for SCADA Network Forensics," Proc. IFIP Int'l Conf. Digital Forensics (IFIP 06), Nat'l Center for Forensic Science, 2006, pp. 273-285.

13. B. Fei, "Data Visualisation in Digital Forensics," Univ. of Pretoria, 2007; http://upetd.up.ac.za/thesis/submitted/etd-03072007-153241/unrestricted/dissertation.pdf.

Ray Hunt is an associate professor at the University of Canterbury, New Zealand; an adjunct associate professor at the University of South Australia; an honorary associate professor at Deakin University, Melbourne; and an adjunct associate professor at Edith Cowan University, Perth. His research interests include firewalls and security architectures, intrusion-detection systems, networking protocols, quality of service in wireless and mobile networks, broadband wireless technologies, and policy-based management security in heterogeneous mobile networks. Hunt received a PhD in computer science from the University of South Australia, Adelaide. Contact him at ray.hunt@canterbury.ac.nz.

Sherali Zeadally is an associate professor in the Department of Computer Science and Information Technology at the University of the District of Columbia, Washington, DC. His research interests focus on computer networks, including wired/wireless networks, network/system/ cyber security, mobile computing, ubiquitous computing, multimedia, and performance evaluation of systems and networks. Zeadally received a PhD in computer science from the University of Buckingham, UK. He is a Fellow of the British Computer Society and a Fellow of the Institute of Engineering Technology, England. Contact him at szeadally@udc.edu.

Selected CS articles and columns are available for free at http://ComputingNow.computer.org.



publications and activities wherever you are.

@ComputerSociety twitter @ComputingNow facebook.com/IEEEComputerSociety facebook facebook.com/ComputingNow



SCADA Systems: Challenges for Forensic Investigators

Irfan Ahmed, University of New Orleans Sebastian Obermeier and Martin Naedele, ABB Corporate Research Golden G. Richard III, University of New Orleans

When security incidents occur, several challenges exist for conducting an effective forensic investigation of SCADA systems, which run 24/7 to control and monitor industrial and infrastructure processes.

n industrial automation and control system is a set of devices that regulate the behavior of physical processes. For example, a thermostat is a simple control system that senses the temperature and turns a heater on or off to maintain the temperature at a set point. These systems are used to monitor and control industrial and infrastructure processes such as chemical plant and oil refinery operations, electricity generation and distribution, and water management.

A control system that is spread over a wide area and can supervise its individual components is often called a supervisory control and data acquisition (SCADA) system.¹ However, here we use the term SCADA to refer to all kinds of control systems that share a common key characteristic: they are connected to physical processes and thus need to be continuously available and able to respond within a deterministic time bound.

Early SCADA systems were intended to run as isolated networks, not connected to the Internet, and thus did not require any specific cybersecurity mechanisms. These systems consisted of simple I/O devices that transmitted the signals between master and remote terminal units. In recent years, SCADA systems have evolved to communicate over public IP networks.² Some are also connected to a corporate intranet or directly to the Internet to seamlessly integrate SCADA data with external information such as corporate email or weather data.

The integration of SCADA systems within a much wider network brings threats that were unimagined at the time these systems were conceived. During the past decade, vendors, asset owners, and regulators recognized this growing concern and began to address it through new laws and various security mechanisms, processes, and standards.³

The discoveries in the wild of Stuxnet in June 2010 and Flame in May 2012 were additional eye-openers for SCADA owners and operators. Stuxnet, the first known malware designed to target automation systems, has infected 50,000 to 100,000 computers worldwide,⁴ while Flame is a cyberespionage tool an order of magnitude more sophisticated than Stuxnet.⁵

SCADA ARCHITECTURE

As Figure 1 shows, a typical SCADA system for controlling infrastructures for utilities such as power, gas, oil, or water generally consists of a control center and numerous field sites. The sites are distributed over a wide geographical area and are connected to the control center by different communication media such as satellites, wide



Figure 1. Simplified logical view of a typical supervisory control and data acquisition (SCADA) architecture.

area networks (WANs), and radio, microwave, or cellular networks. Field sites are equipped with devices such as programmable logic controllers (PLCs) or remote terminal units (RTUs) that control the on-site machines and periodically send information about the state of the field equipment to the control center.

The control center is the SCADA system's hub. Its major components include a human-machine interface (HMI), the database management system (historian), and the server or master terminal unit (MTU). The MTU initiates all communication with field sites and receives the data sent from the field devices. If necessary, it then preprocesses the data and sends it to the historian for archiving. The HMI presents information to the human operator.

FORENSICS FOR SCADA SYSTEMS

Digital forensics is an aspect of cyberdefense that becomes essential in the event of a security breach.⁶ It can generally be defined as the collection and analysis of digital data from different sources such as computer systems, storage devices, and network streams to investigate the causes and consequences of an intrusion or some other incident. If investigators find traces of a crime such as unauthorized network access or theft of a digital file, they can present such data as evidence in a court of law. Digital forensics is also commonly used in internal corporate investigations to help limit the possibility of an incident occurring again in the future.

The recent attacks against SCADA systems by powerful malware such as Stuxnet and Flame highlight the need for

forensic investigations to improve cyberdefenses against both internal and external perpetrators with malicious intent and to thwart entities that try to sabotage a country's critical infrastructure.⁷ In addition to playing a vital role in developing a protection strategy for SCADA systems and assisting in the identification and prosecution of attackers, digital forensics can help deal with nonmalicious but harmful events such as malfunctioning hard disks or other hardware by performing a deep analysis of the underlying SCADA IT system.

A forensic investigation can be the most effective, if not the only, way to answer many questions about an incident. For example, consider a scenario in which malware attacks a SCADA system, causing it to malfunction:

- A virus scan revealed that the Java cache contains a known exploit. Was the exploit successful? What payload does it have? Is this what compromised the system?
- How can the operator clean the system and reliably return it to a known good state without having to shut down the complete system?
- An operator has installed a suspicious, untrusted application downloaded from the Internet. Did that application change components that are important for the system's stable operation?

From a forensics perspective, a SCADA system can be seen in different layers based on the connectivity of the various components with each other as well as with other networks such as the Internet.¹ In Figure 2, layer 0 contains



Figure 2. SCADA system layers. Most forensic analysis involves layers 0-2, which contain the components that control the underlying industrial processes.

the individual field devices connected via a bus network. Layer 1 has controllers that receive input signals from the field devices and other controllers upon which they perform operations to steer the field devices by sending output signals to them. Layer 2 consists of the supervisory network, typically a local network connected to the lower layers for specific operations such as showing the current monitoring state at the HMI. Layer 3 is the operation DMZ, where historians, domain controllers, and application servers are located. Layers 4 and 5 correspond to the enterprise IT networks, in which the enterprise desktops and business servers operate.

Most forensic analyses of SCADA systems involve layers 0-2, as they contain the components that control the un-

derlying industrial processes. However, the analysis can extend to layers 3-5 if needed. Here, we focus on the first three layers.

LIVE FORENSICS

Because a SCADA system must be continuously operational, a forensic investigator cannot turn it off to capture and analyze data.⁸ In this case, *live forensics* is a viable solution for a digital investigation.⁹ A relatively new and emerging field in digital forensics, live forensics involves performing data acquisition and analysis on a running system. However, the critical nature of SCADA systems and their 24/7 availability requirement dictate that forensic investigators spend as little time as possible on a live SCADA system. Consequently, the investigators typically acquire live data and subsequently analyze it offline.

Live data acquisition includes both volatile data such as the contents of physical memory and nonvolatile data such as data stored on a hard disk. This differs from traditional dead data acquisition, which involves first taking the system offline, losing all volatile data. However, volatile data plays a significant role in an effective forensic investigation. For instance, volatile data in physical memory contains information about the system's current state, such as the number of open network connections, process information, and encryption keys.

Live data acquisition challenges

Because volatile data changes continuously on a running system, capturing live data presents two key challenges for forensic investigators.

Early data acquisition after an incident. Live data must be acquired as quickly as possible after an incident to capture any of the incident's traces before processes or services on the running system overwrite useful volatile data—for example, data about recently unloaded kernel modules or drivers.

Digital evidence validity. Digital data might not be admissible in court if its integrity is violated. The intention is to prevent the malicious manufacturing of evidence against an innocent person and to avoid errors while handling evidence in the course of an investigation. Forensic investigators normally prove the integrity of evidence by computing a cryptographic hash of the actual evidence on the compromised system and its acquired copy, which is used for all examinations and analysis. If, however, the compromised system remains live, the data's state might change between the copying and the hash calculation, rendering hashing ineffective as an integrity check.

This also creates an inconsistent data image that does not accurately represent the state when data acquisition starts or after it ends, which can cause difficulty in analyzing the acquired data. For example, due to data inconsistency, sometimes an operating system in the disk image cannot boot for experimental analysis.

Live data acquisition on SCADA systems

It is still unclear how to acquire live data on a SCADA system in a way that minimizes risk to the system's services. To the best of our knowledge, no guidelines for accomplishing this are currently available. However, safe data acquisition should be possible under many circumstances.

Specifically, SCADA systems typically have a primary and a backup system. When the primary system is broken or malfunctioning, operators switch to the backup system.¹⁰ Forensic investigators could leverage this capability by switching the system to the backup and performing live data acquisition on the infected system without worrying about the availability of SCADA services. However, this approach might not be feasible when the malware that has infected the primary system has also infected the backup system. That scenario might even demand immediate recovery if the SCADA owners and operators decide that the incident can jeopardize the system's normal functionality. This usually results in flushing all the infected system components and bringing them back to their normal state, which would not allow sufficient time for the investigator to perform data acquisition.

FORENSIC CHALLENGES IN SCADA SYSTEMS

Beyond the challenges of live data acquisition, forensic investigators must deal with various problems arising from SCADA systems' unique features, which prevent directly applying contemporary forensic tools and techniques.

Deterministic network traffic

Network traffic in SCADA systems is deterministic in that a system component communicates with other system components in a predefined manner. This contrasts with office IT systems, in which desktop machines and servers communicate based on requests in a nondeterministic way.¹¹

It is still unclear how to acquire live data on a SCADA system in a way that minimizes risk to the system's services.

Based on this deterministic behavior, administrators can apply stringent rules to harden the system's security, with any nondeterministic behavior flagged as an anomaly. For instance, an intrusion detection system might be configured to consider a specific communication pattern as normal.¹² Security tools that expect such deterministic behavior might raise false alarms or prevent forensic tools from operating properly. For example, a firewall might have strict rules that allow communication between specific SCADA components but disallow communication between the investigator's machine and SCADA components during data acquisition.

Customized operating system kernels

A SCADA system can have customized kernels running on its components to achieve better performance, support critical applications, and so on, despite the fact that updating such kernels is difficult. For example, PatriotSCADA (www.sage-inc.com/cgi-bin/products_scadasentry.php) is a firewall solution for SCADA networks that uses a customized

Linux kernel to enforce access control and role-based security for every request in the kernel.

However, data acquisition tools might not run on a customized kernel unless they are compatible with each other. For example, the DD disk copy tool might require loading the fmem kernel module (in Linux) to access the physical memory through the device /dev/fmem (which the module creates) if the regular /dev/mem device in Linux is not accessible. Until the module is compiled with the customized kernel, the module might not load into the kernel.

Resource-constrained devices

The availability of SCADA services also depends on the adequacy of system resources: CPU, memory, I/O, and so on. Some system components run on legacy/proprietary hardware and operating systems that might have been deployed for more than 10 years, have moderate computing capabilities compared to modern systems, and have limited or no vendor support.¹³ Moreover, field devices such as RTUs and PLCs are generally resource constrained. SCADA systems with such limited resources demand lightweight data acquisition tools.

The forensic process can be improved in SCADA systems through preparedness and the selection of appropriate tools.

Inadequate logging

Collecting logs of events soon after an incident is crucial for successful forensic investigation. However, SCADA systems' logging capabilities are geared toward process disturbances, not security breaches, and thus are often inadequate.¹³ In such cases, administrators must improve historical visibility in SCADA system components.

Extensive lower-layer data

Capturing and analyzing data on lower layers in SCADA systems is challenging due to the large amount of data that individual sensors generate. In electricity grids, for example, sensors can carry out up to 4,000 measurements per second.¹⁴

FORENSIC TOOLS AND METHODOLOGIES

The forensic process can be improved in SCADA systems through preparedness and the selection of appropriate tools.

Data acquisition plan

To help capture the most relevant data related to an incident, forensic investigators should craft a data acquisi-

tion plan that accurately documents the SCADA system's design, its unique features, the application data flow, and temporary and permanent data storage locations. The plan should also specify what data to acquire for different types of incidents.

Mark Fabro and Eric Cornelius¹⁵ proposed guidelines for creating such a plan in three phases. The first phase involves identifying the system environment and its unique characteristics, including whether the system has modern computing capabilities, is still fully supported by vendors, uses contemporary operating systems, and has continuing support for any open source components. The second phase consists of defining environment-specific requirements such as the impact of vendor solutions on operating systems. The third phase consists of the identification and collection of data, such as activity and transaction logs.

Data acquisition monitoring

During forensic acquisition, no matter how careful an investigator is when copying data, there is always a risk of upsetting the availability of SCADA services. However, this risk can be mitigated by monitoring the availability of system services during data acquisition so that the process can be stopped in case of any serious perturbation. A monitoring tool can facilitate this process by detecting the perturbation as soon as it occurs and automating the response, to avoid any serious damage to the system. EnCase Cybersecurity (www.guidancesoftware.com/ encase-cybersecurity.htm) is an exemplar of a data acquisition monitoring tool that can be integrated into management systems and configured to respond automatically to alerts or events.

Lightweight data acquisition

Data acquisition tools should have a minimal impact so that adequate system resources are available for SCADA services to work properly.

To get a preliminary idea of how resource intensive data acquisition tools are, we ran three well-known versions of the DD tool—WinDD (www.moonsols.com/windows-memorytoolkit), George Garner's DD (www.gmgsystemsinc.com/ fau), and DD for Linux variants—to acquire the entire physical memory and hard-disk data of a computer and recorded the computer's resource consumption during data acquisition for analysis. We acquired the data using Garner's Netcat tool over a 100-Mbps network, a preferred method for forensic investigations.

To emulate a resource-constrained system, we used a PC with an Intel Celeron 1.7-GHz CPU, 384 Mbytes of RAM, and a 40-Gbyte hard drive running at 7,200 rpm. We used two different operating systems, Windows XP Service Pack 2 (SP2) and an Intel Centos 4, for our initial experiments. We kept the machine idle to leave all possible system resources available for the data acquisition tools so that they could

exploit the resources at their full capacity without any constraints. For data acquisition over the network, we directly connected the PC through a crossover cable with the investigative machine, where the data was transferred, to avoid the overhead of packet switching or routing.

The investigative machine was a modern computer with an Intel Core 2 Duo CPU, 4 Gbytes of RAM, and a 300-Gbyte hard drive running at 15,000 rpm, which is unlikely to have caused any performance bottleneck in the data acquisition tools.

Table 1. Resource consumption of data acquisition tools.								
Tool	Operating system	Device	CPU idle time (percent) ¹	Free physical memory (percent) ²	Disk queue length ³			
WinDD	Windows XP (SP2)	Physical memory	90.72	75.60	0.03			
Garner's DD	Windows XP (SP2)	Hard drive	27.49	74.01	0.72			
DD (on Linux variants)	Centos 4	Physical memory	51.98	79.69	0.00			
DD (on Linux variants)	Centos 4	Hard drive	0.646	71.14	0.805			

¹ CPU idle time: average percentage of time during data acquisition that the CPU was idle

² Free physical memory: average percentage of free physical memory during data acquisition

³ Disk queue length: average number of (read and write) requests outstanding on the hard disk during data acquisition

As Table 1 shows, the tools did not exhaust the system resources for data acquisition per se, and might consume less with better hardware than we used. However, the results do not guarantee that the tools are compatible with a particular SCADA environment and would not significantly impact services during the data acquisition process until they are run and tested on that environment or its equivalent, such as a production environment testbed. Moreover, the tools not included in the experiments might not necessarily show a similar performance impact.

Plug-ins for forensic analysis tools

To the best of our knowledge, state-of-the-art forensic analysis tools do not support the unique features of diverse SCADA environments, including protocols and numerous applications' proprietary log formats. Researchers must develop plug-ins or modules for contemporary forensic tools to augment analysis of SCADA systems.

RESEARCH CHALLENGES AND TRENDS

The heightened focus of governments worldwide on protecting their critical infrastructures has led to increased research funding for this purpose. However, the critical nature of SCADA systems also imposes limitations on the research community.

Research challenges

While a security incident in an office environment might lead, at worst, to significant monetary loss or service disruption, breaches of SCADA systems can have dangerous consequences for both human life and the environment.¹⁵⁻¹⁷ In addition, performance requirements for SCADA protection systems have an impact on some security features. For example, in certain use cases, the overhead induced by asymmetric cryptography is intolerable.¹⁸

Thus, research in this domain should be practical and conclusive, which requires the availability of SCADA systems for research purposes. However, building real systems is expensive. To deal with this problem, the SCADA research community usually opts for the following approaches, each of which has its own merits and limitations.

Using simulators. Some commercial simulators, such as Opal Software's simSCADA (www.opalsoftware.com.au/ index.php?option=com_content&view=article&id= 35&Itemid=67), provide a virtual environment for studying SCADA systems. They are mostly used to imitate the network traffic between field devices and MTUs and are effective at reducing hardware purchase and installation costs. However, simulators are subject to errors and thus typically do not provide the same level of confidence that a real system would.

Building small-scale SCADA systems. Government and academic researchers use commercial hardware and software to build laboratory-scale testbeds of some SCADA systems such as industrial blowers, gas pipelines, power grids, and petroleum storage tanks. For example, Mississippi State University has a testbed for studying and learning about multiple industrial control systems.¹⁹ The Idaho National Laboratory has a testbed of a full-scale electrical grid that is dedicated to control system cybersecurity assessment, standards improvements, and training (www.inl.gov/research/national-supervisory-control-and-data-acquisition-test-bed).

Industry collaboration. When applying for project funding, researchers try to engage SCADA owners and operators as industrial partners. The terms of agreement for a project usually involve technical assistance, facility access (at least to the testbed the operators use for testing application patches from vendors), and financial support. Industrial collaboration provides close access to real-world SCADA systems and the technical personnel who actually experience the problems and understand the limitations of their particular system.

However, industry collaboration is often difficult to achieve due to the critical nature of SCADA systems, which discourages owners and operators from cooperating with the research community to prevent information leakage.

This creates a gap between the research community's efforts and resolving the problems that SCADA owners and operators face.

In this situation, governments are often in the best position to play a mediator role and help reduce this gap. For example, the Australian government regularly organizes community-of-interest meetings to provide a platform for discussions among SCADA owners and operators, SCADA vendors, and researchers from academia.

Research trends

Thus far, the research community has mainly focused on SCADA system security. However, there has also been limited work on the forensic investigation of SCADA systems.

Tim Kilpatrick and colleagues^{11,20,21} proposed an architecture for capturing and subsequently analyzing sensor data and control actions in a SCADA network. Agents placed at strategic locations within the network capture local traffic and forward a relevant portion of packets, called a synopsis, to a data warehouse. After analyzing a synopsis, the data warehouse creates its digital signature and stores it with the synopsis in the agent's designated storage area. A relational database and query mechanisms support forensic investigations. The modular agent design and configurable synopsis engines accommodate diverse SCADA protocols, some of their implementation variations, and subsets of standard or proprietary protocols. The researchers developed a prototype of the architecture based on the Modbus TCP protocol using two control devices and one HMI station.

Craig Valli²² created a framework that produces forensically verified signatures for the Snort intrusion detection system (IDS) for known and published vulnerabilities of SCADA and control systems, enabling investigators to trace exploits during analysis. Valli first looked for vulnerability announcements or traces at Black Hat, hacker, vendor, CERT (Community Emergency Response Teams), and other relevant sites and reproduced the vulnerability scenarios. He then examined the vulnerabilities of SCADA communication protocols such as Modbus and DNP3.

Valli conducted experiments involving an attacker, a victim/target machine running SCADA software, the Snort IDS, and a network sniffer that captures all network traffic in a tcpdump binary capture file for analysis to generate Snort rules. He analyzed the exploit's modus operandi and used this to create a rule set to reduce or stop the attack. Valli later included the rule set in the Snort configuration to test its resilience under sustained attack.

According to Jill Slay and Elena Sitnikova,²³ a generic approach to forensics in SCADA systems requires a big picture view of the process that encompasses a range of technical and procedural issues at the government, industry, and academic levels.

ecause of the underlying industrial processes they Π control, performing a forensic investigation of SCADA systems is fundamentally different from investigating corporate or home networks. The critical nature of SCADA systems demands that investigators be well trained and thoroughly understand the requirements to manage such systems. Engaging investigators early so they can become accustomed to a particular environment is highly recommended. It is also desirable to encourage SCADA system owners and operators to initiate steps that can facilitate an investigation if needed-for example, by maintaining a data acquisition plan and regularly testing data acquisition tools to ensure that they will not affect the availability of SCADA system services.

SCADA-focused forensic research is essential to address the unique challenges associated with these systems. The forensic research community must engage SCADA owners and operators and those actively working on SCADA systems to highlight research problems and develop solutions. Governments must take a more active role in organizing these efforts and helping to provide researchers with resources and suitable access to SCADA systems.

Acknowledgment

This work was supported in part by NSF grant CNS #1016807.

References

- 1. D. Bailey and E. Wright, *Practical SCADA for Industry*, Newnes, 2003.
- R. Kalapatapu, "SCADA Protocols and Communication Trends," *Proc. 2004 ISA Industrial Network Security Symp.* (ISA Expo 04), Instrumentation, Systems, and Automation Soc., 2004; www.isa.org/journals/intech/ TP04ISA048.pdf.
- M. Brändle and M. Naedele, "Security for Process Control Systems: An Overview," *IEEE Security & Privacy*, Nov./Dec. 2008, pp. 24-29.
- T.M. Chen and S. Abu-Nimeh, "Lessons from Stuxnet," Computer, Apr. 2011, pp. 91-93.
- 5. G. Keizer, "Development Timeline Key to Linking Stuxnet, Flame Malware," *Computerworld*, 30 May 2012; www. computerworld.com/s/article/9227580/Development_ timeline_key_to_linking_Stuxnet_Flame_malware.
- 6. K. Mandia, C. Prosise, and M. Pepe, *Incident Response and Computer Forensics*, 2nd ed., McGraw-Hill/Osborne, 2003.
- R.N. Charette, "Gone Missing: The Public Policy Debate on Unleashing the Dogs of Cyberwar," blog, 4 June 2012; http://spectrum.ieee.org/riskfactor/telecom/security/gonemissing-the-public-policy-debate-on-unleashing-thedogs-of-cyberwar.
- M. Naedele, "Addressing IT Security for Critical Control Systems," *Proc. 40th Hawaii Int'l Conf. System Sciences* (HICSS 07), IEEE CS, 2007; doi:10.1109/HICSS.2007.48.
- 9. F. Adelstein, "Live Forensics: Diagnosing Your System without Killing It First," *Comm. ACM*, Feb. 2006, pp. 63-66.
- K. Stouffer, J. Falco, and K. Scarfone, *Guide to Industrial* Control Systems (ICS) Security, NIST special publication 800-82, Nat'l Inst. Standards and Technology, 2011; http://

csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final. pdf.

- T. Kilpatrick et al., "An Architecture for SCADA Network Forensics," *Advances in Digital Forensics II*, M.S. Olivier and S. Shenoi, eds., Springer, 2006, pp. 273-285.
- H. Hadeli et al., "Leveraging Determinism in Industrial Control Systems for Advanced Anomaly Detection and Reliable Security Configuration," *Proc. 14th Int'l Conf. Emerging Technologies and Factory Automation* (ETFA 09), IEEE, 2009, pp. 1189-1196.
- M. Fabro and E. Cornelius, *Recommended Practice: Creating Cyber Forensics Plans for Control Systems*, tech. report INL/ EXT-08-14231, Idaho Nat'l Lab., 2008.
- H. Kirrmann, "Seamless Redundancy: Bumpless Ethernet Redundancy for Substations with IEC 61850," *ABB Rev.*, Aug. 2010, pp. 57-61.
- 15. D. Dzung et al., "Security for Industrial Communication Systems," *Proc. IEEE*, June 2005, pp. 1152-1177.
- F. Köster et al., "Collaborative Security Assessments in Embedded Systems Development—The ESSAF Framework for Structured Qualitative Analysis," *Proc. Int'l Conf. Security and Cryptography* (Secrypt 09), INSTICC Press, 2009, pp. 305-312.
- E. Levy, "Crossover: Online Pests Plaguing the Offline World," *IEEE Security & Privacy*, Nov./Dec. 2003, pp. 71-73.
- S. Fuloria et al., "The Protection of Substation Communications," *Proc. SCADA Security Scientific Symp.* (S4 10), 2010; www.cl.cam.ac.uk/~rja14/Papers/S4-2010.pdf.
- T. Morris, R. Vaughn, and Y.S. Dandass, "A Testbed for SCADA Control System Cybersecurity Research and Pedagogy," *Proc. 7th Ann. Cyber Security and Information Intelligence Research Workshop* (CSIIRW 11), ACM, 2011; doi:10.1145/2179298.2179327.
- T. Kilpatrick et al., "Forensic Analysis of SCADA Systems and Networks," *Int'l J. Security and Networks*, Feb. 2008, pp. 95-102.
- 21. R. Chandia et al., "Security Strategies for SCADA Networks," *Critical Infrastructure Protection*, E. Goetz and S. Shenoi, eds., Springer, 2008, pp. 117-131.
- 22. C. Valli, "SCADA Forensics with Snort IDS," *Proc. 2009 Int'l Conf. Security and Management* (SAM 09), CSREA Press, 2009, pp. 618-621.
- 23. J. Slay and E. Sitnikova, "The Development of a Generic Framework for the Forensic Analysis of SCADA and Process Control Systems," *Forensics in Telecommunications, Information, and Multimedia*, M. Sorrell, ed., Springer, 2009, pp. 77-82.

Irfan Ahmed is a postdoctoral research associate in the Department of Computer Science at the University of New Orleans. His research interests include industrial control system security, digital forensics, and malware detection and analysis. Ahmed received a PhD in computer science from Ajou University, South Korea. Contact him at irfan. ahmed@uno.edu.

Sebastian Obermeier is a principal scientist at ABB Corporate Research. His research interests include IT security for industrial control systems and database technology. Obermeier received a PhD in computer science from the University of Paderborn, Germany. Contact him at sebastian. obermeier@ch.abb.com.

Martin Naedele is the R&D program manager for Industrial Software Systems at ABB Corporate Research. His research interests include software engineering and IT security. Naedele received a PhD in computer engineering from ETH Zurich. He is a GIAC-certified security auditor and a member of IEEE, ACM, and the International Council on Systems Engineering. Contact him at martin.naedele@ ch.abb.com.

Golden G. Richard III is a professor of computer science and a university research professor at the University of New Orleans. His research interests include digital forensics, reverse engineering, and operating systems internals. Richard received a PhD in computer science from the Ohio State University. He is a member of IEEE, Usenix, ACM, and the American Academy of Forensic Sciences. Contact him at golden@cs.uno.edu.

Selected CS articles and columns are available for free at http://ComputingNow.computer.org.

IEEE ISM 2012

14th IEEE International Symposium on Multimedia

10-12 December 2012

Irvine, CA, USA

IEEE ISM 2012 is an international forum for researchers to exchange information regarding advances in the state of the art and practice of multimedia computing, as well as to identify the emerging research topics and define the future of multimedia computing.

Register today! http://ism.eecs.uci.edu





Smartphone Security Challenges

Yong Wang, Kevin Streff, and Sonell Raman, Dakota State University

Because of their unique characteristics, smartphones present challenges requiring new business models that offer countermeasures to help ensure their security.

martphones are quickly becoming the dominant device for accessing Internet resources. Sales of smartphones overtook PC sales in the global market in Q4 2010.¹ Shipments of smartphones surpassed those of feature phones in Western Europe in Q2 2011.² According to a May 2011 Nielsen survey, smartphones outsold feature phones in the US in this same period.³ Compared to 5.9 billion worldwide mobile phone subscribers, smartphone usage (835 million) is still steadily increasing.⁴ IDC predicts smartphone shipments will approach one billion in 2015.⁵

Smartphones offer many more functions than traditional mobile phones. In addition to a preinstalled mobile operating system, such as iOS, Android, or Windows Mobile, most smartphones also typically support carrier networks, Wi-Fi connectivity, and Bluetooth so that users can access the Internet to download and run various thirdparty applications. Most smartphones support Multimedia Message Service (MMS) and include embedded sensors such as GPS, gyroscopes, and accelerometers, as well as a high-resolution camera, a microphone, and a speaker.

Smartphones' increasing popularity raises many security concerns.⁶⁻⁹ Their central data management makes them easy targets for hackers. Since the first mobile phone viruses emerged in 2004, smartphone users have reported significant malware attacks. In the last seven months of 2011, malware attacks on the Android platform increased 3,325 percent.¹⁰ As the use of smartphones continues its rapid growth, subscribers must be assured that the services they offer are reliable, secure, and trustworthy.

SMARTPHONE THREATS AND ATTACKS

In a smartphone threat model, a malicious user publishes malware disguised as a normal application through an app store or website. Users will unintentionally download the malware to a smartphone, which carries a large amount of sensitive data. After infiltrating a smartphone, the malware attempts to control its resources, collect data, or redirect the smartphone to a premium account or malicious website.

This model divides a smartphone into three layers:

- The *application* layer includes all of the smartphone's apps, such as social networking software, email, text messaging, and synchronization software.
- The *communication* layer includes the carrier networks, Wi-Fi connectivity, Bluetooth network, Micro USB ports, and MicroSD slots. Malware can spread through any of these channels.
- The *resource* layer includes the flash memory, camera, microphone, and sensors within a smartphone. Because smartphones contain sensitive data, malware targets their resources to control them and manipulate data from them.

An attack forms a loop starting with the launch of the malware, moving through the smartphone's application, communication, and resource layers, on to premium accounts/malicious websites, and back to the malicious user. Figure 1 shows such an attack.



Figure 1. Smartphone threat model. In this attack, the smartphone user unwittingly downloads malware to a smartphone through social networking software via a carrier's network. The malware hijacks the smartphone's resources and sends Multimedia Message Service (MMS) messages to a premium account.

Affected services

Malware's impact can range from minor issues, such as degraded performance, spam messages, and slow operation, to more significant challenges, such as the user not being able to receive and make phone calls or incurring financial loss. The impact to any one smartphone user might be completely different from that experienced by other subscribers.

Jeopardized resources

Resources containing sensitive data are attractive to hackers. Once malware finds a way into the smartphone, it will try to gain privileges to access and control these resources.

For example, flash memory can be quickly reprogrammed with malware that cannot be removed until the user reprograms the flash memory. Some smartphones also include MicroSD memory cards. With a data cable or a card reader, a malicious user can easily disclose the memory card's content.

Sensors, such as GPS, gyroscopes, and accelerometers, also contain sensitive information. GPS, for example, can reveal a smartphone subscriber's location, which the subscriber might not want to disclose.

In addition, a user might not be aware that malware has turned the smartphone camera or microphone on. Malware with full control of a smartphone can thus transform it into a tapping device.

Moreover, data can be leaked when a user transfers data from a smartphone to a computer through a Wi-Fi or Bluetooth network.

Finally, smartphones depend on batteries to power on.

Battery exhaustion attacks can dissipate battery power faster than normal, disabling a smartphone's functions.

Malware

Smartphone malware falls into three main categories: viruses, Trojans, and spyware.¹⁰

Viruses are typically disguised as a game, security patch, or other desirable application, which a user downloads to a smartphone. Viruses can also spread through Bluetooth. Two Bluetooth viruses have been reported in smartphones:

- Bluejacking sends unsolicited messages over Bluetooth to a Bluetooth-enabled device within a limited range (usually around 33 feet).
- Bluesnarfing accesses unauthorized information in a smartphone through a Bluetooth connection.

Most smartphone Trojans are related to activities such as recording calls, instant messaging, finding a location via GPS, or forwarding call logs and other vital data. Smart Message System Trojans comprise a large category of mobile malware that run in an application's background and send SMS messages to a premium rate account owned by an attacker. HippoSMS, for example, increases users' phone charges by sending SMS messages to premium mobile accounts and blocks service providers' messages alerting users of the additional charges.

Spyware collects information about users without their knowledge. According to a 2011 report, spyware was the dominant malware affecting Android phones, accounting for 63 percent of the samples identified.¹⁰

Carrier IQ software, which runs hidden in the background and does not require authorized consent to function, is usually preinstalled in a smartphone to collect usage data intended to help carriers improve service. Mobile operators, device manufacturers, and application vendors can use this information to deliver high-quality products and services. However, smartphone subscribers usually are not aware of what data is being collected or how it is processed and stored.

Threats and attacks

Smartphone threats and attacks include sniffing, spam, attacker spoofing, phishing, pharming, vishing, and data leakage.

The WebKit engine, which most mobile platforms use, has a vulnerability that lets attackers crash user applications and execute malicious code.

Sniffing captures and decodes packets as they pass over the airwaves. There are various ways to sniff or tap a smartphone. In 2010, Karsten Nohl showed that A5/1, the Global System for Mobile Communications (GSM) encryption function for call and SMS privacy, could be broken in seconds.¹¹ Thus, all GSM subscribers are at risk for sniffing attacks. Further, as eavesdropping software continues to be available, smartphone subscribers using 3G or 4G networks are also at risk.

Spam can be carried through email or MMS messages. These messages can include URLs that direct users to phishing or pharming websites. MMS spam can also start a denial-of-service (DoS) attack. According to industry analyst Richi Jennings, the number of spam text messages generated in the US increased 45 percent in 2011 to 4.5 billion messages.¹²

Another threat involves an attacker spoofing a caller ID and pretending to be a trusted party. Researchers have spoofed MMS messages that appear to come from 611, which carriers use to send alerts or update notifications.¹³ Base stations can also be spoofed.

A phishing attack can masquerade as a trusted party to steal personal information, such as a username, password, or credit card account number. Many phishing attacks have occurred in social networking, email, and MMS messages. For example, a malicious application could include a "Share on Facebook" button that redirects users to a spoofed target application, which could request the user's secret credentials and steal the data.

Pharming attackers can redirect Web traffic on a smartphone to a malicious or bogus website. By collect-

ing the subscriber's smartphone information, a pharming attack can lead to other attacks. For example, when a user browses a website on a smartphone, the HTTP header usually includes information about the smartphone's operating system, browser, and version number. With this information, an attacker can learn the smartphone's security vulnerabilities and start other directed attacks.

Vishing is short for "voice phishing." In a vishing attack, malicious users try to gain access to a smartphone user's financial and other private information. By spoofing a caller ID, the attacker might look like a trusted party and fool the smartphone user into releasing personal credentials.

Data leakage is the unauthorized transmission of personal information or corporate data. Malicious software can steal personal information such as a contact list, location information, or bank information and send this data to a remote website. A smartphone's data leakage can put its owner at risk of identity theft. Business owners or classified users such as government and military personnel have even more concern about data leakage. ZitMo, a mobile version of Zeus, has been found in Symbian, Black-Berry, and Android devices. An attacker could use ZitMo to steal one-time passwords sent by banks to authenticate mobile transactions.

Web browsers are also vulnerable to smartphone attacks. The WebKit engine, which most mobile platforms use, has a vulnerability that lets attackers crash user applications and execute malicious code. CrowdStrike revealed that attackers could use the WebKit vulnerability to install a remote access tool to eavesdrop on smartphone conversations and monitor user locations. The vulnerability has been found in BlackBerry, iOS, and Android devices.

For various reasons, smartphones are also vulnerable to DoS attacks:

- Because they are based on radio communication technology, smartphones can incur an attack in which a jamming device is used to disrupt the communication between the smartphone and its base station.
- Flooding attacks can generate hundreds of text messages or incoming calls, thus disabling a smartphone.
- A battery exhaustion attack on a smartphone causes more battery discharge than is typically necessary.
- A malicious user could use a smartphone's blocking features to start a DoS attack. If a malicious user keeps calling a smartphone from a blocked phone number, the subscriber cannot use any of the smartphone's functions.

Many attacks operate in a stealth mode. Users might not notice these attacks for days or even months. In addition, a malicious user could plant malware in a smartphone but not use it until later.

MOBILE DEVICE FORENSICS

Mobile device forensics—which covers cell phones, smartphones, tablets, personal digital assistants, and GPS receivers—is a subspecialty of computer forensics, necessitated by the near-ubiquity of these devices in today's society. Because mobile devices are increasingly the instrument, target, or record keeper of criminal and other nefarious behaviors, they are of interest in criminal investigations, civil litigation, and intelligence collection. Some would argue that mobile devices contain more probative information per byte examined than traditional computers.

Because most smartphones now come with sophisticated applications, built-in cameras, lots of storage capacity, and high-speed network connectivity, a vast amount of computing power is readily available within the user's grasp. Indeed, smartphones are more properly thought of as portable Internet terminals than merely as phones.

Although mobile device forensics involves retrieving and examining data even if it might have been deleted, for both criminal and civil proceedings, the processes and tools are also used in applications outside the courtroom. Data that can be recovered from a mobile device includes call history, sent and received Short Message Service (SMS) and multimedia messages, contacts and phone numbers, emails, photos, videos, geolocation and GPS information, wireless network settings, Web browsing history, voicemail messages, social networking information, application histories and logs, and other data that might be retained within smartphone apps.

Numerous commercial and open source products are available for extracting and analyzing mobile device data, ranging from camera kits that take screenshots to products that can acquire a file system (logical) or the entire memory (physical) to software that parses databases and hardware to physically examine the device's chips. Figure A shows two of the most widely used products for logical and physical data extraction and analysis: the Cellebrite Universal Forensics Extraction Device (UFED) Ultimate (www.cellebrite.com) and the Micro Systemation XRY Complete/XACT (www.msab.com).

Mobile device forensics requires a process and tools that can extract information from at least six different mobile operating systems (OSs), including iOS, Android, and Windows Mobile, and thousands of models of phones, tablets, and GPS devices.

Even if an examiner can physically acquire the device's memory, examination of that binary dump might still require good, oldfashioned analysis with hex editors, standard computer forensics tools, and regular expressions. As an example, Android phones have

SECURITY CHALLENGES AND IMPACT

As the "Mobile Device Forensics" sidebar describes, the near ubiquity of devices such as smartphones in today's society necessitates the development of this subspecialty of computer forensics. Many techniques used to secure desktop and laptop computers such as antivirus and antimalware software can be used for smartphone security. However, smartphones have some unique characteristics that make security extremely challenging.

Consumer products

The wide range of smartphone subscribers is matched by the variety of smartphone uses: communication, information, social networking, gaming, entertainment, business enterprise, and so on. Smartphones are conconsiderable fragmentation and variation in OS versions, which makes locating common data across OSs and devices difficult. Even Apple's iPhone returns different data depending on the OS version and whether the phone is jail-broken (freed from the limitations imposed on it by Apple).

Many feature phones and dumb phones in the marketplace store contacts and SMS data on the SIM card, while they usually store images or videos locally on the device. Tablets typically perform the same as handsets with regard to returned forensic data. Despite the availability of many tools, this is not push-button forensics.

Because of the incredible amount of personal and business information stored on mobile devices—or that can be inferred from information on them—security and privacy challenges abound. In addition to placing user information on the phone, the OS also stores information unbeknownst to the user. For example, in April 2011, Apple received considerable media attention when it became known that the iPhone had been recording a detailed history of user geographical locations in an unprotected file; with a simple extraction, a forensic examiner could create a geotagged map of all of the places that iPhone (and presumably its user) visited. The key lesson: sensitive data should always be encrypted in smartphones.



Figure A. Products for use in data extraction and analysis: (1) Cellebrite UFED and (2) Micro Systemation XRY Complete/ XACT.

sumer products, and different groups of people have different preferences, thus their needs for smartphone security also differ.

No single security tool is appropriate for all subscribers. Smartphone security tools should be configurable to meet the needs of different groups. For example, a business user is typically more concerned about smartphone security than a gamer and thus is willing to spend more to ensure device security.

Finally, most people do not expect to keep their smartphones for a long time. Instead, they expect them to be damaged or lost and to eventually need replacement. Thus, to merit the investment in its purchase, a security solution must be transferrable to a replacement device or it can be purchased in a replacement device at low cost. Demo



Figure 2. Comparison of desktop operating systems and smartphone operating systems market share.

Platform-oriented

Smartphones have a preinstalled mobile operating system. Unlike desktop operating systems, which are dominated by Microsoft Windows, as Figure 2 shows, Android, iOS, BlackBerry, Symbian, and Windows Mobile share the mobile operating system market.⁵

Each mobile operating system provides different applications, features, and interfaces. The variety lets consumers personalize their devices; however, this presents a challenge to the hardware vendors and smartphone application developers who must support these various mobile operating systems. Further, multiple versions of each mobile operating system might exist, especially for the Android OS.

The differences between these operating systems dictate the security software, as the smartphones must also be platform-oriented. Because operating systems are vulnerable to different security breaches, security software must specifically address each type of breach. Mobile security software developers must therefore customize the software for each mobile platform to deal with multiple operating system and version issues.

Multiple-entrance open system

Smartphones are multiple-entrance open systems, and each entrance is a potential back door for malware access. Each smartphone communication channel is a potential path for malware disguised as an application.

Because smartphones offer multiple entrances, an attack loop can consist of many combinations, but an attack loop cannot be formed if malware is detected, prevented, and removed from the smartphone. Securing a smartphone requires using one of many possible approaches to break the attack loop. For example, resource control could break the attack loop by preventing the malware from gaining access to the smartphone's resources to manipulate its data.

Central data management

Some applications cache users' secret credentials in the smartphone's storage units. This sensitive data might include personal information such as home address, phone numbers, photos, and contact lists; correspondence such as email, text messages, and call logs; credit card information, user names, and passwords; files on flash memory or a memory card; geographic location; and corporate data.

Disclosing this data can end in data leakage resulting in invasion of the smartphone owner's privacy or leading to financial loss.

In addition to using encryption techniques to protect data, migrating data from smartphones to the cloud is another option for securing information and reducing the risk of data theft.

Limited battery life

A smartphone is a resource-constrained device that is powered by a battery with a limited life and that must be recharged when drained. Any security solution must consider this limitation as enhanced security cannot sacrifice battery life.

Vulnerability to theft and loss

Among all potential security issues, loss and theft are two primary concerns for smartphone users. According to a report by Lookout, nine million smartphones were lost in the US in 2011—that's one phone every 3.5 seconds.¹⁴

Losing control of a smartphone, even temporarily—say, by loaning it to someone—can have catastrophic consequences. With some simple setup, a malicious user can reprogram a smartphone's firmware and flash memory, physically clone the memory card, or install spyware.

Some simple techniques can help protect against smartphone theft and loss. For example, the user can add a password or enable auto-lock. Antitheft technology that remotely deletes sensitive data when a smartphone leaves a secure zone is also available through third-party applications.¹⁵

Embedded sensors

Smartphones often contain many embedded sensors. Although these sensors enrich a smartphone's functions, they also increase risk. For example, researchers found a way to use accelerometers to decipher computer keystrokes. With a 58,000-word dictionary, they achieved 80 percent accuracy.¹⁶ Using this technique, it is easy to convert a smartphone into a tapping device. With more sensors planned for installation in smartphones, new threats and attacks might be explored and discovered.

However, smartphone owners are vulnerable to the abusive use of smartphone sensor data. For example, some smartphone applications disclose data to a third party. In addition, malware might be disguised as a normal application to request and receive access to GPS data. Finally, malware might jail break an iPhone or iPod Touch to allow the device to run code that is not authorized by Apple and thus gain control of its sensors.

Real-time resource monitoring can help reduce the risks of an attacker using a smartphone's embedded sensors. Further, smartphones could use real-time monitoring to detect and block illegal access of the embedded sensors.

Other concerns

As companies adopt smartphones for their businesses, the BYOD (bring your own device) concept is raising many security concerns for administrators and IT professionals.¹⁷ Although this concept lets employees easily use their own devices to access corporate applications and resources, auditing and enforcing security policies on a personal device is difficult.⁹

The smartphone security challenges that enterprises face include users' failure to back up and encrypt critical data. Further, employees might be unaware that the company has security policies for using smartphones.

Inevitably, most smartphones will include both personal and business data. However, this raises concerns about the security of corporate data. One solution is to develop tools that can distinguish between the two types of data and enforce a higher security level on corporate data.

Educating smartphone users can improve their security awareness. A company should enforce its security policies and regularly audit employee smartphones to ensure their security.

DESIRED SECURITY FEATURES

Confidentiality, integrity, and authentication are three of the most desirable security features in a smartphone.

Most smartphones support synchronization between the device and a computer. This function makes it possible for another user to access the smartphone file system. Thus, to keep data confidential, users should employ encryption techniques and avoid storing sensitive information in plaintext on a smartphone.

Integrity applies to both data and the system. App stores should verify software integration to avoid malicious modification. Further, smartphones should provide mechanisms to protect system integrity. They should also block unauthorized data access requests.

A smartphone authentication service could protect smartphone users against malware attacks that spoof caller IDs and MMS. Because femtocells help improve both coverage and capacity, authentication becomes important to validate a carrier's identity.

Separation of sensitive from nonsensitive data

A smartphone should separate sensitive from nonsensitive data and give users the flexibility to assign data as sensitive. Although sensitive data might be an easy target for hackers, having a clear target to protect instead of taking extra computational and battery power to protect the entire flash drive or memory card is advantageous. In addition, simple security techniques, such as encryption and steganography, can protect sensitive data.

Isolating sensitive data is also good for business. Smartphone users can identify corporate data as being sensitive and assign a higher security level to it.

Encryption

In addition to encrypting sensitive data stored in smartphones, users should encrypt their memory cards. Without a proper decryption key, the smartphone should not disclose the memory card's contents. However, public-key cryptography, such as RSA, usually requires additional computational power and should be used with caution to avoid draining the battery.

To keep data confidential, users should employ encryption techniques and avoid storing sensitive information in plaintext on a smartphone.

Migrating data from smartphones to the cloud is another way to protect sensitive data. Cloud-based intrusiondetection techniques could help detect misbehavior and protect sensitive data.¹⁸ This option would involve a cost for the cloud service. Migrating smartphone data to the cloud also involves background data and thus increases data usage.

any enterprises have started to explore security solutions for smartphones. Currently, smartphone subscribers are solely responsible for installing applications and ensuring that they are secure. However, security requires collaboration among mobile users, service providers, and industry partners. New business models for smartphone security are therefore highly desired.

Smartphone security is challenging and complicated. However, there are some easy ways to improve smartphone security:

• Increase security awareness. Like a desktop PC or a laptop, a smartphone can be hacked, infected, or phished. Smartphone users should be aware of potential threats and attacks when installing software¹⁹ or authorizing it to access flash memory or smartphone sensors.

- Apply password and auto-lock after a period of time. Most smartphones support these protective features.
- Do not store irreplaceable data in a smartphone. Smartphones can easily be lost or stolen.
- Backup smartphone data regularly. Sync your smartphone with a computer and always keep a backup of your data.
- Turn off Bluetooth when you are not using it. Viruses can spread through Bluetooth to your smartphone.
- Do not use unsecure Wi-Fi hotspots to connect to the Internet. Packet sniffer software such as Wireshark can disclose useful information from smartphone data traffic.
- Use a reliable and trusted security tool to secure your smartphone.
- Ask the smartphone vendor or service provider about antitheft technology such as "erase data" or "default smartphone" remotely.

Some subtle signs might indicate that a smartphone is under attack. For example, the phone's battery is warm even when the device has not been used; the phone lights up at unexpected times, including when it is not in use; or the phone unexpectedly beeps or clicks during phone conversations. If any of these occur, have a security professional check the smartphone.

References

- IDC, "Mobile Phone Market Grows 17.9% in Fourth Quarter, According to IDC," press release, 28 Jan. 2011; www.idc.com/about/viewpressrelease.jsp? containerId=prUS22679411.
- 2. IDC, "Smartphones Outstrip Feature Phones for First Time in Western Europe as Android Sees Strong Growth in 2Q11, Says IDC," press release, 9 Sept. 2011; www.idc. com/getdoc.jsp?containerId=prUK23024911.
- Nielsen, "In US, Smartphones Now Majority of New Cellphone Purchases," blog, 30 June 2011; http://blog.nielsen. com/nielsenwire/online_mobile/in-us-smartphones-nowmajority-of-new-cellphone-purchases.
- ITU, "Key Global Telecom Indicators for the World Telecommunication Service Sector," Nov. 2011; www.itu.int/ ITU-D/ict/statistics/at_glance/KeyTelecom.html.
- IDC, "Worldwide Smartphone Market Expected to Grow 55% in 2011 and Approach Shipments of One Billion in 2015," June 2011; www.idc.com/getdoc. jsp?containerID=prUS22871611.
- 6. N. Leavitt, "Mobile Security: Finally a Serious Problem?", *Computer*, June 2011, pp. 11-14.
- 7. W. Jeon et al., "A Practical Analysis of Smartphone Security," *Proc. Int'l Conf. Human Interface and the Management of Information—Part I*, Springer-Verlag, 2011, pp. 311-320.
- N. Husted, H. Saïdi, and A. Gehani, "Smartphone Security Limitations: Conflicting Traditions," Proc. 2011 Workshop on Governance of Technology, Information, and Policies, ACM, 2011, pp. 5-12.
- McAfee, Mobility and Security: Dazzling Opportunities, Profound Challenges, tech. report, May 2011; www.mcafee.

com/us/resources/reports/rp-cylab-mobile-security.pdf.

- Juniper Networks, 2011 Mobile Threats Report, tech. report, Feb. 2012; www.juniper.net/us/en/local/pdf/additionalresources/jnpr-2011-mobile-threats-report.pdf.
- K. Nohl, "Attacking Phone Privacy," *BlackHat Lecture Notes*, July 2010; http://media.blackhat.com/bh-us-10/ whitepapers/Nohl/BlackHat-USA-2010-Nohl-Attacking. Phone.Privacy-wp.pdf.
- O. Kharif, "Mobile Spam Texts Hit 4.5 Billion Raising Consumer Ire," *Bloomberg Businessweek*, 30 April 2012; www. businessweek.com/news/2012-04-30/mobile-spam-textshit-4-dot-5-billion-raising-consumer-ire.
- 13. Z. Lackey and L. Miras, "Attacking SMS," BlackHat, 2009.
- BusinessWire, "Lookout Projects Lost and Stolen Phones Could Cost U.S.Consumers over \$30 Billion in 2012," 22 Mar. 2012; www.businesswire.com/news/home/ 20120322005325/enLookout-Projects-Lost-Stolen-Phones-Cost-U.S.
- "Virginia Tech Cybersecurity Breakthrough Keeps Sensitive Data Confined in Physical Space, Engineering Team Says," Virginia Tech News, 17 Oct. 2011; www.vtnews.vt.edu/ articles/2011/10/101711-outreach-cybersecurephones.html.
- P. Marquardt et al., "(sp)iphone: Decoding Vibrations from Nearby Keyboards Using Mobile Phone Accelerometers," *Proc. 18th ACM Conf. Computer and Communications Security* (CCCS 11), ACM, 2011, pp. 551-562.
- J. Burt, "BYOD Trend Pressures Corporate Networks," *eWeek*, Sept. 2011, pp. 30-31.
- A. Houmansadr, S. Zonouz, and R. Berthier, "A Cloud-Based Intrusion Detection and Response System for Mobile Phones," *Proc. Dependable Systems and Networks Workshops* (DSN-W 11), IEEE, 2011, pp. 31-32.
- D. Barrera and P. Van Oorschot, "Secure Software Installation on Smartphones," *IEEE Security and Privacy*, May 2011, pp. 42-48.

Yong Wang is an assistant professor in the National Center for the Protection of the Financial Infrastructure at Dakota State University. His research interests include wireless networks, optical networks, smartphones, and related security and privacy issues. Wang received a PhD in computer science from University of Nebraska-Lincoln. He is a member of IEEE and the IEEE Communications Society. Contact him at yong.wang@dsu.edu.

Kevin Streff is the director of the National Center for the Protection of the Financial Infrastructure and founder of Secure Banking Solutions, a security consulting firm focused on improving security in community banks. His research interests include banking assurance and security risk management. Streff received a PhD in electronic business from Capella University. Contact him at kevin.streff@ dsu.edu.

Sonell Raman is a second-year graduate student at Dakota State University majoring in database management. His research interests focus on smartphone security with respect to mobile applications. Raman received a BS in computer science and engineering from Jawaharlal Nehru Technological University, Hyderabad, India. Contact him at sraman17633@pluto.dsu.edu.

IEEE (Computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEBSITE: www.computer.org **OMBUDSMAN:** To check membership status or report a change

of address, call the IEEE Member Services toll-free number, +1 800 678 4333 (US) or +1 732 981 0060 (international). Direct all other Computer Society-related questions—magazine delivery or unresolved complaints—to help@computer.org.

CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

AVAILABLE INFORMATION: To obtain more information on any of the following, contact Customer Service at +1 714 821 8380 or +1 800 272 6657:

- Membership applications
- Publications catalog
- Draft standards and order forms
- Technical committee list
- Technical committee application
- Chapter start-up procedures
- Student scholarship information
- Volunteer leaders/staff directory
- IEEE senior member grade application (requires 10 years
- practice and significant performance in five of those 10)

PUBLICATIONS AND ACTIVITIES

Computer: The flagship publication of the IEEE Computer Society, *Computer*, publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.

Periodicals: The society publishes 13 magazines, 18 transactions, and one letters. Refer to membership application or request information as noted above.

Conference Proceedings & Books: Conference Publishing Services publishes more than 175 titles every year. CS Press publishes books in partnership with John Wiley & Sons. **Standards Working Groups:** More than 150 groups produce IEEE standards used throughout the world.

Technical Committees: TCs provide professional interaction in more than 45 technical areas and directly influence computer engineering conferences and publications.

Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.

Certifications: The society offers two software developer credentials. For more information, visit www.computer.org/ certification.

NEXT BOARD MEETING 6–8 Feb. 2013, Long Beach, CA, USA

EXECUTIVE COMMITTEE

President: John W. Walz*

President-Elect: David Alan Grier*

Past President: Sorel Reisman*

VP, Standards Activities: Charlene (Chuck) Walrad[†]

Secretary: Andre Ivanov (2nd VP)* VP, Educational Activities: Elizabeth L. Burd*

VP, Member & Geographic Activities: Sattupathuv Sankaran[†]

- VP, Publications: Tom M. Conte (1st VP)*
- VP, Professional Activities: Paul K. Joannou*

VP, Technical & Conference Activities: Paul R. Croll[†] Treasurer: James W. Moore, CSDP*

2011–2012 IEEE Division VIII Director: Susan K. (Kathy) Land, CSDP[†] 2012–2013 IEEE Division V Director: James W. Moore, CSDP[†] 2012 IEEE Division Director VIII Director-Elect: Roger U. Fujii[†] *voting member of the Board of Governors †nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 2012: Elizabeth L. Burd, Thomas M. Conte, Frank E. Ferrante, Jean-Luc Gaudiot, Paul K. Joannou, Luis Kun, James W. Moore, William (Bill) Pitts

Term Expiring 2013: Pierre Bourque, Dennis J. Frailey, Atsuhiro Goto, André Ivanov, Dejan S. Milojicic, Paolo Montuschi, Jane Chu Prey, Charlene (Chuck) Walrad

EXECUTIVE STAFF

Executive Director: Angela R. Burgess Associate Executive Director; Director, Governance: Anne Marie Kelly Director, Finance & Accounting: John Miller Director, Information Technology & Services: Ray Kahn Director, Membership Development: Violet S. Doan Director, Products & Services: Evan Butterfield Director, Sales & Marketing: Chris Jensen

COMPUTER SOCIETY OFFICES

 Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928

 Phone: +1 202 371 0101 • Fax: +1 202 728 9614

 Email: hq.ofc@computer.org

 Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314

 Phone: +1 714 821 8380

 Email: help@computer.org

MEMBERSHIP & PUBLICATION ORDERS

Phone: +1 800 272 6657 • Fax: +1 714 821 4641 • Email: help@computer.org Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan Phone: +81 3 3408 3118 • Fax: +81 3 3408 3553 Email: tokyo.ofc@computer.org

IEEE OFFICERS

President and CEO: Gordon W. Day President-Elect: Peter W. Staecker Past President: Moshe Kam Secretary: Celia L. Desmond Treasurer: Harold L. Flescher President, Standards Association Board of Governors: Steven M. Mills VP, Educational Activities: Michael R. Lightner VP, Membership & Geographic Activities: Howard E. Michel VP, Publication Services & Products: David A. Hodges VP, Technical Activities: Frederick C. Mintzer IEEE Division V Director: James W. Moore, CSDP IEEE Division VIII Director: Susan K. (Kathy) Land, CSDP IEEE Division VIII Director-Elect: Roger U. Fujii President, IEEE-USA: James M. Howard



Using Tracing to Solve the Multicore System Debug Problem

Aaron Spear, VMware Markus Levy, Multicore Association Mathieu Desnoyers, EfficiOS



The common trace format overcomes deficiencies in traditional software tools to optimize modern multicore designs by providing open source formats to monitor the state and interaction of concurrent systems over time.

ebugging and optimizing modern multicore designs is increasing in difficulty proportional to the exponential growth of technology. For example, according to the Pew Research Center, in early 2012 the number of smartphones overtook low-end phones in the US, and 46 percent of American adults are smartphone owners (http://pewinternet.org/reports/2012/ smartphone-update-2012/findings.aspx).

Requirements on power consumption, external connectivity—for cellular, Wi-Fi, and Bluetooth—human interfaces, and time to market have resulted in general patterns for designing devices:

- the use of more cores (at often variable clock rates) to meet the same performance levels at increasingly lower power,
- dedicated hardware for performance-critical functionality,
- open source operating systems (OSs), drivers, and software stacks such as Linux and Android, and
- splitting of control and data planes into separate processors, with different OSs on each.

Embedded systems are doing more data processing for functions such as video analysis or packet routing, and they

are doing it in parallel. The net result is a set of requirements that are extraordinarily difficult to reconcile. Debugging an embedded system with a single core is difficult enough, but debugging multiple streams of software activity interacting with each other is even more challenging.

THE PROBLEM WITH SOFTWARE DEBUGGING

Developers can use the GNU Project Debugger (GDB), a traditional software debugger, to look at snapshots of one program at a time. The developer sets a breakpoint, runs the program, and then sequentially steps through the code from the breakpoint, verifying that state changes as expected. The developer sees the stack of the functions that the program called leading up to the next line of source code for execution. When concurrently debugging multiple threads of execution (whether for multiple or single cores), most software debuggers simply add more contexts to the same display. For example, Figure 1 shows the Eclipse C/C++ Development Tooling (CDT) with GDB, as it debugs three Linux processes with multiple threads.¹

Debugging a set of cores that are interacting with each other requires attaching a debugger to each core. Depending on the cores and tools, it is possible for the debugger to concurrently debug multiple cores. Sometimes hardware allows synchronous debugging that runs and stops all cores together. However, this use case tends to be the exception, not the rule.

In many cases, embeddedsystem designers must use decoupled tools to individually debug each core. As difficult as this description sounds, the real situation is in fact much worse because of the layering of OSs and application stacks. Android, for example, consists of three distinctly different software domains that require debugging:

~ - 0 🐐 🗱 🗈 🗉 🛋 🕺 🥱 🧑 🦛 🏂 Debug 🖾 50 ✓ ☑ Multi-process debugging [C/C++ Application] Thread [5] [core: 0] (Running) ✓ intersting/loopfirst [4966] [cores: 1] Thread [4] [core: 1] (Suspended : Breakpoint) Thread [3] 5005 [core: 1] (Running) thread_exec1() at /home/Imckhou/runtime-TestDSF/NonStop/src/NonStop.cpp:27 0x8048646 ≡ start thread() at 0xb7fbb96e ≡ clone() at 0xb7deda4e Thread [1] 4999 [core: 1] (Running) dbp 🛃 Figure 1. Multiprocess debugging with Eclipse CDT using GDB 7.2.

- Java applications and application frameworks at the top of the software stack,
- the Linux user process space and a unique Java virtual machine in the middle, and
- the kernel space (which includes the Linux kernel and drivers) at the bottom.

There are no open source tools that allow seamless debugging from Java down into a C/C++ user process, and then down into Linux kernel code. Debugging everything in the software stack requires three different debuggers for the core running Android, and then additional debuggers for other cores.

Problems that are impossible for a software debugger to solve further compound the extreme difficulty in simply setting up to debug. Debugging tools cannot stop systems with real-time deadlines. It is also increasingly common for designs to push significant portions of functionality into dedicated "hardware accelerators," on-chip resources that offload functionality such as networking, video, or graphics from the main processor. These blocks tend to be black boxes inaccessible to software debuggers.

Software debuggers are invaluable for debugging algorithms, but debugging a system of concurrently running hardware and software components requires a tool that can trace the state and interaction of software and hardware over time.

LOGGING BENEFITS AND LIMITATIONS

Programmers who have written any C code are well aware of the benefit of seeing state change over time. Since the 1950s, the standard method for logging has been to dump state information to standard output (stdout) streams or files via the venerable printf routine or one of its predecessors in older languages.

Having no alternative, developers add large amounts of instrumentation to software, providing visibility

into the system's behavior over time. Most commercial applications of any significant size have some sort of logging infrastructure, likely cobbled together by product implementers. These facilities can be invaluable for troubleshooting issues in the field after a product ships, for example.

Despite its benefit, this type of logging has practical limits:

- Logging is intrusive—it changes the system timing. If well architected, logging creates a small and deterministic intrusion. But, depending on the architecture, logging might also completely change the system's behavior. For example, consider two parallel threads that simultaneously insert a log message in the same buffer/file, which requires some form of mutual exclusion. Additionally, if the program uses string formatting to produce the logs, system performance might be a concern.
- Logs only cover instrumented locations. This type of logging requires inserting preprocessor macros or logging function calls into the source code and then building the application. This static instrumentation is extremely valuable, but the developer might not anticipate all the required information. Among other things, this technique also requires changing and rebuilding all the instrumented components, which might be difficult or impossible.
- *Correlation in time is difficult.* Correlating multiple contexts running in parallel requires including a time stamp to understand ordering. In the multicore world, this can mean correlating cores with unrelated clocks.
- Tools for multicore log analysis are not readily available. Different technologies often produce unique log formats, and developers sometimes create their own tools to analyze them because what they need simply does not exist in a usable form.

TRACING

A solution to the multicore system analysis problem must record a heterogeneous system's state over time and must also address logging's limitations. For several reasons, we use tracing instead of logging. The difference between these two terms is arguably semantic; however, the following attributes distinguish them:

- Tracing intrusion is minimized by discarding data, or overwriting the oldest data in the case of bufferfull conditions. Logging is primarily concerned with information completeness, and as such will block the application until I/O can complete.
- In tracing, performance is paramount, so any solution must minimize intrusion. Thus, methods typically employed in logging, such as string-formatted messages, might not be acceptable.
- Trace events tend to be low level and high volume, such as OS scheduler information.

Even with a capable trace analyzer that displays events over time and in context, adapting multiple formats is difficult.

- Developers can dynamically insert tracing into a system. Logging tends to be a static feature of a product.
- Tracers might have stringent robustness requirements. Some systems function in "flight recorder" mode, always collecting trace data as a part of normal operation. Doing so enables detailed analysis in the case of a failure or error that only occurs on production workloads.

Many companies and open source initiatives have created capable tracing mechanisms, including Dtrace, SystemTap, Ftrace, strace, and Linux Trace Toolkit Next Generation (LTTng). However, all of these mechanisms have one fundamental problem: they cannot easily share data.

Modern multicore systems combine many different technologies. A single design can aggregate different processor architectures; closed and open source OSs and application stacks; hardware acceleration; and so on. Adequately tracing such a system requires coalescing many different trace data formats, and no tool for doing so exists. Even with a capable trace analyzer that displays events over time and in context, adapting multiple formats is difficult.

ENABLING MULTICORE TOOL INTEROPERABILITY

The Multicore Association's member companies saw the benefits of assembling a Tools Infrastruc-

ture Working Group (TIWG) to develop standards to advance multicore tool interoperability. At the inaugural meeting, various tool and OS vendors discussed pain points with representatives of the semiconductor industry to determine how the group might work together. It quickly became apparent that the increasing parallelism and complexity of multicore systems presented daunting problems. Although each participant had solutions for portions of the problem, none had a solution for everything. A farsighted goal of the participants was to provide tools for developing and optimizing multicore systems.

The working group agreed that the first step must be to visualize, understand, and benchmark the system's behavior. It would then be clear what to focus on, and the group could validate changes against previous benchmarks. The idea of an open standard for interchanging trace data seemed like a great first step that would serve as a foundation for future work.

TIWG was not the only group that recognized the need for a trace data standard. In the Linux world, in particular, there are several different tracing technologies with incompatible data formats. Similarly, there was a strong desire to share analysis tooling, and Ericsson began driving a unification effort for Linux in 2008. At the Embedded Linux Conference in April 2010, the TIWG and Linux communities decided to collaborate to create a standard that could meet the needs of Linux and embedded systems. Financed by Ericsson and the Embedded Linux Forum, Mathieu Desnoyers, LTTng maintainer, began the initial work to create the common trace format (CTF) specification, with input from TIWG.

Synchronizing traces from widely different sources

CTF offers a general trace data format that is application, architecture, and programming language agnostic. It is meant as a carrier for data that is temporal—ordered events that occur over time. The goal at its highest level is to easily analyze trace streams from various collection mechanisms on different cores or systems, each with different clock domains. OSs, hardware probes, bus analyzers, simulators, or instrumentation in any arbitrary application can create the traces.

A common use case might include traces from the following:

- a reduced-instruction-set computer (RISC) processor running Linux instrumented with LTTng to collect kernel-level trace data (high-level scheduling information for the kernel),
- a hardware trace probe collecting low-level instruction and data traces on the same RISC chip (to provide function-call-level detail),

- a digital signal processor running a proprietary realtime OS (RTOS) that also has a kernel instrumentation trace, and
- a hardware trace from a network accelerator hardware block in the same system.

Each trace source records different types of data that use different clocks. The clocks often relate in some way, and it is typically possible to correlate trace streams in other ways—for example, via interprocess communication among the cores that cause events in multiple contexts. Each trace collector is independent, resulting in decoupled traces. The goal is that single or multiple analyzers can correlate those traces and perform specific analyses on them, understanding their trace-collectorspecific event schema.

At its lowest level, CTF is a format to express arbitrary events versus time. A CTF-compatible analysis tool would, by default, display these events in a general way, showing the events' various attributes as plaintext, for example. The analyzer would also provide custom, event-specific, and potentially more meaningful rendering—displaying a video frame as an image, for example.

CTF's design anticipates many different collection mechanisms in the future, although the initial proof of concept used the following trace collectors:

- OS instrumentation, such as Ftrace and LTTng;
- application tracing, such as LTTng user space tracing (LTTng-UST); and
- hardware trace, such as instruction and data traces like Nexus.

Requirements from the trace collector (the trace data producer), the trace analyzer (the trace data consumer), and storage and transmission mediums shape the format's details.

The guiding principle of the trace collector's requirements is to minimize intrusion: recording the data must be optimally efficient for the tracer. This means that events are arbitrary binary data, encodable with natural byte order, alignment, and packing. The trace collector also must explicitly record or infer context information for an event—for example, core, thread, time stamp, and event type. Additionally, the trace collector should describe the trace data's layout with simple-to-generate human-readable metadata. The trace collector can leverage CTF facilities to describe the layout of any extra context information that it prepends to each event payload.

Trace analyzer requirements include functions such as handling extremely large traces (more than 10 Gbytes) as well as aggregating traces from different components, processes, cores, or systems. It must also be able to bias and offset time stamps on traces for event correlation.

How CTF models events

Anticipated collection and consumption use cases shape the details of CTF event encoding. CTF assumes that the collector stores raw events in an optimized buffer specific to the trace collector. At some point after the collector stores those events, it either inserts the events into CTF trace files or streams them onto a communication link (network). The terminology for the format aligns with those use cases. Preliminary definitions include

- *event trace:* a container of one or more event streams and metadata that describes them;
- event stream: an ordered sequence of events, containing a subset of the trace event types;
- *event packet:* a sequence of physically contiguous events within an event stream (event packets are variable size, with optional padding at the end); and
- *event*: the basic entry in a trace—a variable-size container with event-specific attributes that include type, context, and payload.

The guiding principle of the trace collector's requirements is to minimize intrusion: recording the data must be optimally efficient for the tracer.

File representation. After its generation and optional transport over the network, the final trace output is stored either permanently or temporarily in a virtual file system. Because the trace collector appends information to each event stream while recording a trace, it associates each stream with a separate file for output. Therefore, as Figure 2 shows, it is possible to represent a stored trace as a directory containing one file per stream.

Network streaming. A key use case for CTF is collecting events on a remote system and transmitting them over a network. It is also possible to stream events as they occur. The participants in the CTF group chose the event packets terminology to reflect the common use case of a group of events transmitted together. A tool or component generating the CTF format can choose the size of an event packet to fit within one user datagram protocol (UDP) packet to minimize the use of network and CPU bandwidth. The TIWG is developing complementary protocols for control of trace collection, as well as for streaming trace data.

CURRENT CTF STATUS

The CTF specification is currently at draft 1.8.2 and is under review by TIWG and other members of the open source community (www.efficios.com/ctf). Several open

COMPUTING PRACTICES



Figure 2. Block diagram of common trace format (CTF) data in a file system.

source projects have added CTF support in their most recent versions, as described in the "Open Source Tools and Frameworks" sidebar.

The Multicore Association TIWG is currently expanding the CTF specification to allow systems to produce information about the relationship between the clocks used in different traces. Doing so will support automatic correlation of trace events collected from different cores. Texas Instruments, a member of the working group, is spearheading this effort.

Extensions to the specification are being added that allow interpretation of event data with respect to common concepts. Take, for example, the tracing of an OS scheduler versus time. OSs often have vastly different execution models and OS-specific terminology for common concepts—processes, threads, tasks, blocks, and so on. Regardless of the implementation's details, it is universally desirable to present a Gantt chart of execution state versus time. CTF will be able to express a general concept of state that is easily describable in metadata for the events. The

OPEN SOURCE TOOLS AND FRAMEWORKS

everal open source tools either already use CTF-based formats or are the focus of plans to incorporate CTF as follows:

- BabelTrace is an MIT-licensed framework for reading and converting CTF and legacy trace formats. The purpose is to enable seamless transition to CTF for both trace collectors and analyzers. Sources are available for download at www. efficios.com/babeltrace.
- LTTng 2.x kernel tracer produces CTF format natively.
- LTTng-UST 2.x user space tracer produces CTF format natively.
- Eclipse Linux Tools LTTng viewer is an open source analyzer that reads CTF format natively. There is growing interest in extending this tool to support additional OSs and trace environments.
- LTTV trace analyzer support for CTF is under development.

result will be rich analysis capabilities at minimal cost for the tracer.

TIWG is also working on developing an MIT-licensed open source C library for creating CTF files. The vision is to facilitate an easy migration path to CTF for proprietary and open source OSs, tool suppliers, and others. With this library, adding CTF support inside various products—a hardware analyzer's firmware, RTOS source code, and so on—will be straightforward.

e believe that CTF will enable analyzing the behavior of complex multicore systems in a way

that was not previously possible due to the diversity of technologies. An open standard and the existence of easily accessible open source components for creating and analyzing traces will enable developers to quickly add tracing to applications and application frameworks. The result will be a better understanding of multicore systems' behavior. TIWG welcomes involvement and participation from interested companies and individuals.

Reference

 P. Chuong, D. Alexiev, and M. Khouzam, "Staying Ahead of the Multi-core Revolution with CDT Debug," 2011; www.slideshare.net/marckhouzam/ multi-coredebugeclipsecon11.

Aaron Spear is a staff engineer at VMware, where he works on tools and infrastructure for cloud computing. He is cochair of the Multicore Association's Tool Infrastructure Working Group, and works on resolving the parallel computing problems that occur in every domain. Spear received a BS in electrical and computer engineering from the University of Colorado, Boulder. Contact him at aspear@ vmware.com.

Markus Levy, president of the Multicore Association and chair of the Multicore Developers Conference, is also the founder and president of the Embedded Microprocessor Benchmark Consortium. Levy received a BS in electrical engineering from San Francisco State University. Contact him at markus.levy@eembc.org.

Mathieu Desnoyers is the CEO and a senior software architect at EfficiOS. He is maintainer of the LTTng project, Babeltrace, and Userspace RCU library, and he participated in the creation of the common trace format specification. Desnoyers works on tracing and efficient synchronization of scalable, real-time, and energy-constrained systems. He received a PhD in computer engineering from Ecole Polytechnique de Montréal. Contact him at mathieu.desnoyers@ efficios.com.

JaaS Cloud Architecture: From Virtualized Datacenters to Federated Cloud Infrastructures



Rafael Moreno-Vozmediano, Rubén S. Montero, and Ignacio M. Llorente Complutense University of Madrid

As a key component in a modern datacenter, the cloud operating system is responsible for managing the physical and virtual infrastructure, orchestrating and commanding service provisioning and deployment, and providing federation capabilities for accessing and deploying virtual resources in remote cloud infrastructures.

atacenters have evolved from expensive, rigid, mainframe-based architectures to agile distributed architectures based on commodity hardware that developers can dynamically shape, partition, and adapt to different business processes and variable service workloads.

Virtualization plays an important role as an enabling technology for datacenter implementation by abstracting compute, network, and storage service platforms from the underlying physical hardware.¹ Virtualized infrastructures support server consolidation and ondemand provisioning capabilities, which results in high server utilization rates and significant cost and energy savings.

Because secure, efficient, and scalable management of these virtualized infrastructures is essential to guarantee optimal datacenter operation, the virtual infrastructure manager is a key component of the datacenter architecture. The main role of this component is to deliver infrastructure as a service (IaaS), transforming the traditional data center into a cloud-like architecture.²

Future datacenters will look like private IaaS clouds, supporting the flexible and agile execution of virtualized services. In this context, the virtual infrastructure manager, also called the cloud operating system (cloud OS), orchestrates the deployment of virtual resources and manages the physical and virtual infrastructures to command-andcontrol service provisioning. In addition, management of the datacenter as a cloud makes it possible to complement the local infrastructure with remote resources from other federated datacenters or commercial clouds.

This cloud vision of the datacenter represents not only a new provisioning model but also a way to simplify and optimize infrastructure operation because applications are not tied to a specific physical server and data is not attached to a single storage device. This provides several advantages from the infrastructure management perspective, including server consolidation to reduce hardware

RESEARCH FEATURE





and power requirements, on-the-fly resizing of the physical infrastructure, service workload balance among physical resources to improve efficiency and utilization, server replication to support fault tolerance and high availability capabilities, and dynamic partitioning of physical infrastructure to execute and isolate different services and workloads.

Turning this vision into reality requires developing an open and flexible cloud architecture reference model that addresses the requirements of business use cases from IT companies and across multiple industries. This model should also provide some basic features such as adaptability to manage any hardware and software combination, interoperability and portability to prevent vendor lock-in, scalability to support large-scale infrastructures, and standardization by leveraging and implementing standards.

An open challenge in cloud computing is cloud federation,³ which involves different architectures⁴ and levels of coupling among federated cloud instances. The cloud architecture reference model must include support for the deployment of different federation scenarios so that cloud providers and IT companies can use external resources as well as share their internal resources.

CLOUD INFRASTRUCTURE ANATOMY

As Figure 1 shows, the key component of an IaaS cloud architecture is the cloud OS, which manages the physical and virtual infrastructures and controls the provisioning of virtual resources according to the needs of the user services.

There are many similarities between a typical computer system's threaded OS and a cloud OS. A computer OS's main role is to manage the computer resources—the CPU, memory, disks, and I/O devices—and provide a secure and isolated multithreaded execution environment for user applications. This environment enables resource sharing between different users, abstracting the user from the specifics of the underlying hardware and offering different interfaces for interacting with the computer.

Similarly, a cloud OS's role is to efficiently manage datacenter resources to deliver a flexible, secure, and isolated multitenant execution environment for user services that abstracts the underlying physical in-

frastructure and offers different interfaces and APIs for interacting with the cloud. While local users and administrators can interact with the cloud using local interfaces and administrative tools that offer rich functionality for managing, controlling, and monitoring the virtual and physical infrastructure, remote cloud users employ public cloud interfaces that usually provide more limited functionality.

OpenNebula (http://opennebula.org) is an example of an open cloud OS platform focused on datacenter virtualization that fits with the architecture proposed in Figure 1. Other open cloud managers, such as OpenStack (http:// openstack.org) and Eucalyptus (www.eucalyptus.com), primarily focus on public cloud features. They are not fully oriented to the management of virtualized datacenters and do not include an integrated module for managing different federation scenarios.

Infrastructure and cloud drivers

To provide an abstraction of the underlying infrastructure technology, the cloud OS can use adapters or drivers to interact with a variety of virtualization technologies. These include hypervisor, network, storage, and information drivers. The core cloud OS components, including the virtual machine (VM) manager, network manager, storage manager, and information manager, rely on these infrastructure drivers to deploy, manage, and monitor the virtualized infrastructures. In addition to the infrastructure drivers, the cloud OS can include different cloud drivers to enable access to remote providers.

Virtual machine manager

In the same way that a multithreaded OS defines the thread as the unit of execution and the multithreaded application as the management entity supporting several communication and synchronization instruments, a cloud OS defines the VM as the basic execution unit and the virtualized services (group of VMs for executing a multitier service) as the basic management entity supporting different communication instruments and their autoconfiguration at boot time. This concept helps create scalable applications because the user can either add VMs as needed (horizontal scaling) or resize a VM (if supported by the underlying hypervisor technology) to satisfy a VM workload increase (vertical scaling). Individual multitier applications are isolated from each other, but individual VMs in the same applications are not, as they all can share a communication network and services when needed.

A VM consists of a set of parameters and attributes, including the OS kernel, VM image, memory and CPU capacity, network interfaces, and so on. The VM manager is responsible for managing a VM's entire life cycle and performing different VM actions—deploy, migrate, suspend, resume, shut down—according to user commands or scheduling strategies. To perform these actions, the VM manager relies on the hypervisor drivers, which expose the basic functionality of underlying hypervisors' such as Xen, KVM, and VMware to avoid limiting the cloud OS to a specific virtualization technology.

The VM manager is also responsible for preserving the service-level agreements contracted with the users, which are usually expressed in terms of VM availability in infrastructure clouds. To guarantee this availability, the VM manager should include different mechanisms for detecting VM crashes and automatically restarting the VM in case of failure.

Network manager

The deployment of services in a cloud involves not only the provision of VMs but also the instantiation of communication networks to interconnect the different service components and to make the service reachable for external users, if needed. The network manager should be able to manage private networks to interconnect both the service's internal components and public IP address pools and connect the front-end service components to the Internet. The network manager uses the network drivers to provision virtual networks over the physical network infrastructure. As different virtual networks can share a common physical link, the network manager should provide an automated procedure for MAC and IP address assignment to avoid address overlap problems. It should also offer additional mechanisms to guarantee traffic isolation between different virtual networks. Traffic isolation can be achieved either by filtering the Ethernet traffic in the device the VM is attached to or by configuring a VLAN (IEEE 802.1Q) each time a new virtual network is instantiated.

Storage manager

The storage manager's main function is to provide storage services and final-user virtual storage systems as a commodity. Thus, the storage system must be scalable so that it can grow dynamically according to service needs;

Image management is an important challenge in current virtualized datacenters, since they must handle a huge amount of VM images belonging to different users.

highly available and reliable, to avoid data access disruption in data access in case of failure; high-performance, to support strong demands of data-intensive workloads; and easy to manage, abstracting users from the underlying physical storage's complexity.

To reach these goals, the storage manager relies on the existing storage drivers, which introduce a layer of abstraction between users or services and physical storage and enable the creation of a storage resource pool where storage devices appear as one, allowing data to be moved freely among devices.

Image manager

Image management is an important challenge in current virtualized datacenters, since they must handle a huge amount of VM images belonging to different users, with different operating systems and software configurations. Thus, the cloud OS must have the appropriate tools to manage these images efficiently and securely, as well as having additional functionality for administering image repositories.

A set of attributes defines the VM image, including the image's name, a description of its contents, the type of image—public, private, or shared—the image owner, and the image's location within the repository. Basic image functionality should include tools for creating a new image in a repository, deleting an image, cloning an image from an existing one, adding or changing an image

RESEARCH FEATURE

attribute, sharing an image with other users, publishing an image for public use, or listing the images available in the repository.

Information manager

The information manager is responsible for monitoring and gathering information about the state of VMs, physical servers, and other components of virtual and physical infrastructures such as network devices and storage systems. This monitoring function is essential to ensure that all these components are performing optimally.

The information manager uses the information drivers to collect monitoring information from virtual and physical resources. At the physical server level, the administrator can install different specialized tools for monitoring purposes⁵ such as Nagios (www.nagios.com) and Ganglia

Authorization policies control and manage user privileges and permissions to access different cloud resources, such as VMs, networks, or storage systems.

(http://ganglia.sourceforge.net). In contrast, monitoring at the VM level relies on the information provided by hypervisors, which can be very limited and might differ from one hypervisor to another.

Another way to obtain detailed VM-monitoring information is to install monitoring tools in the VM that can interact with the information manager. However, this solution is intrusive and requires the VM owner's consent.

The information manager can provide various predefined sensors, each one responsible for a different aspect of the system to be monitored such as CPU load, memory usage, running processes, disk usage, power consumption, and bandwidth consumption. It is also possible to design custom sensors to use in monitoring new metrics adapted to the characteristics of the deployed service or to collect information about various VMs performing a service.

Authentication and authorization

As in any kind of shared environment, clouds must incorporate mechanisms to authenticate users and administrators and to provide them with access only to authorized resources.

User authentication verifies and confirms the identity of users who try to access cloud resources. This function can be implemented using different methods, such as simple password verification mechanisms via LDAP or another kind of active directory; trusted authentication mechanisms based on public keys, X.509 certificates, or Kerberos; or Internet-based authentication mechanisms such as SAML or openID.

Authorization policies control and manage user privileges and permissions to access different cloud resources, such as VMs, networks, or storage systems. Access control can be implemented using role-based mechanisms, where a role defines a group of permissions to perform certain operations over specific cloud resources and users can be assigned particular roles. In addition, quota mechanisms can be used to limit the amount of resources—CPU, memory, network bandwidth, or disk—a specific user can access.

Accounting and auditing

Accounting's objective is to obtain and record resource usage information of deployed services. This function relies on the information manager to monitor resources and collect usage information from metric measurements. Accounting is essential to implementing the mechanisms that produce billing information.

Auditing provides information about activity in cloud resources, indicating who accessed cloud resources, when they gained access, and what operations they performed. This information is useful to improve cloud security and protect it from threats such as unauthorized access, abusive use of resources, and other forms of intrusion.

Federation manager

The federation manager enables access to remote cloud infrastructures, which can be either partner infrastructures governed by a similar cloud OS entity or public cloud providers. The federation manager should provide basic mechanisms for deployment, runtime management, and termination of virtual resources in remote clouds; remote resource monitoring; user authentication in remote cloud instances; access control management and remote resource permission; and tools for image building on different clouds with different image formats. The support for other advanced features such as creation of cross-site networks and virtual storage systems or cross-site VM migration will depend on the federation capabilities that remote clouds offer, as well as the level of coupling and interoperability supported by the different clouds involved.

The federation manager's design could differ depending on the supported types of federation—for example, cloud aggregation, bursting, or brokering—and levels of coupling and interoperability. The cloud OS must implement the federation manager as an internal component to support federation architectures at the infrastructure level. However, user-level federation scenarios could be implemented with a third-party stand-alone service, such as Aeolus (http://aeolusproject.org), which offers brokering services to access different cloud providers.

Table 1. Examples of optimization criteria for allocation and reallocation policies.							
Optimization criteria	Target	Allocation policy	Reallocation policy				
Server consolidation	Reduce the number of servers in use to minimize energy consumption.	VMs should be placed using a minimum number of servers.	VMs can be dynamically reallocated to reduce the number of servers in use.				
Workload balance	Balance the workload of all servers to avoid server saturation and performance slowdown.	VMs should be evenly distributed among the available servers.	A VM can be dynamically reallocated to balance VM distribution among servers.				
CPU balance	Balance the use of CPUs to avoid server saturation and performance slowdown.	A new VM should be located in the server with the highest amount of available CPUs.	In case of server saturation (CPU overloading), a VM can be dynamically reallocated to a less-loaded server.				
Thermal balance	Balance the temperature of all servers to avoid overheating and reduce cooling requirements.	A new VM should be located in the server exhibiting the lowest temperature.	In case of server overheating, a VM can be dynamically reallocated to a cooler server.				

Scheduler

There are two levels of scheduling within a cloud infrastructure: at the physical host level, managed by the hypervisor scheduler, which is responsible for deciding when VMs can obtain system resources—such as physical CPU or memory—and which physical CPUs are assigned to each VM; and at the cloud level, managed by the cloud OS scheduler, which is responsible for deciding the particular physical server where each VM is deployed.

The cloud OS scheduler's main function is to decide the initial placement of each VM following specific criteria. In a federated environment, the scheduler could decide to deploy the VM in a remote cloud when insufficient resources are available in the local infrastructure. In addition, the scheduler could also provide dynamic optimization capabilities, enabling the dynamic reallocation (migration) of VMs from one physical resource to another to meet specific optimization criteria.

Table 1 lists different scheduling policies, based on varying optimization criteria, to guide both initial placement and dynamic reallocation actions. In addition, the user can also specify the constraints that can restrict scheduler decisions such as, for example, hardware (amount of CPU, memory, and so on), platform (type of hypervisor, OS, and so on), affinity (two or more VMs that need to be deployed in the same physical server or the same physical cluster), location (geographical restrictions), or service-level agreement constraints (guaranteed CPU capacity or high operational reliability), among others.

The cloud OS invokes the scheduler every time a new VM is waiting to be deployed as well as periodically to optimize the entire virtual infrastructure, reallocating VMs if necessary to meet the established optimization criteria. The scheduler interacts with the VM manager to deploy or reallocate the VM in the selected server or with the federation manager to deploy VMs in remote clouds.

Administrative tools

The cloud OS must provide different tools and interfaces (command-line or GUI) for users and administrators to perform various tasks. For example, for privileged administration, the system should include both user administration tools (to create, modify, or delete users and manage user authorization and access control policies) and physical infrastructure management tools (to boot or shut down physical servers, monitor physical infrastructure, and so on). In addition, unprivileged users should also be provided with tools to manage their own infrastructure-for example, VM management tools (to deploy, shut down, suspend, restore, or monitor a VM), virtual network management tools (to create or delete virtual networks), virtual storage management tools (to create, delete, or attach a virtual disk), and image management tools (to create, clone, or delete images).

Service manager

The cloud OS should be able to manage and support virtualized multitier services. A multitier service can comprise several component/tiers with some intrinsic dependencies among them. These services can be deployed as a group of interconnected VMs in the cloud with specific deployment dependencies and, optionally, some location, affinity, and elasticity requirements. In addition, multitier service deployment can involve some communication and storage elements such as virtual networks and virtual disks.

The service manager's admission control function entails deciding whether to accept or reject a service, depending on the service requirements and resource availability in the cloud. Once it accepts a service, the service manager is responsible for managing its life cycle, which can involve several actions, including deploying, suspending, resuming, or canceling the service. To deploy a new service, the service manager interacts with the scheduler to decide the best placement for the various VMs that comprise the

RESEARCH FEATURE

service, according to the selected optimization criteria and service constraints.

Another service manager function is the management of service elasticity. Elastic services can experience fluctuating workloads, and the service manager should adapt the capacity to this variable demand. The service manager can incorporate different mechanisms for service autoscaling based on elasticity rules, which trigger the deployment of new instances (horizontal scaling) or by resizing existing instances (vertical scaling) when user-specified service metrics exceed certain thresholds.

Independent of the service manager, users are always allowed to employ the interfaces provided by the administrative tools or the cloud interface to deploy, resize, migrate, or shut down their individual VMs.

Cloud interfaces

Cloud OS functionality can be exposed to consumers using some well-known public cloud interfaces. In the current cloud ecosystem, most cloud products and providers

Cloud developers and researchers have proposed or implemented numerous federation architectures, including cloud bursting, brokering, and aggregation.

offer their own APIs, such as Amazon EC2 or VMware's vCloud. Although some of these APIs are becoming de facto standards, this heterogeneity makes it difficult to achieve interoperability and portability across clouds.

Standards specifications and cloud adapters are the two main high-level approaches for interoperability. Open source projects such as Apache Deltacloud (http:// deltacloud.apache.org) or Libcloud (http://libcloud.apache.org) offer adapters for a range of clouds. Additionally, several standards bodies are addressing interoperability and portability issues surrounding cloud infrastructures.⁶ Open standards such as OGF OCCI (http://occi-wg.org), DMTF CIMI and OVF (http://dmtf.org/standards/cloud), and SNIA CDMI (www.snia.org/cdmi) are gaining increasing attention from cloud providers and consumers.

CLOUD FEDERATION

Cloud federation,^{3.7} which enables cloud providers and IT companies to collaborate and share their resources, is associated with many portability and interoperability issues.⁸ Cloud developers and researchers have proposed or implemented numerous federation architectures, including cloud bursting, brokering, and aggregation.⁴ These architectures can be classified according to the level of coupling or interoperation among the cloud instances involved, ranging from loosely coupled (with no or little interoperability among cloud instances) to tightly coupled (with full interoperability among cloud instances).

Coupling levels

In a federated scenario, cloud providers can exhibit different degrees of coupling pertaining to the level of cooperation among cloud instances, to the level of control and monitoring allowed over remote resources, to the possibility of deploying cross-site networks or even of migrating VMs between cloud instances.

Loosely coupled federation. This scenario is formed by independent cloud instances—for example, a private cloud complementing its infrastructure with resources from an external commercial cloud—with limited interoperation between them. Usually, a cloud instance has little or no control over remote resources (for example, decisions about VM placement are not allowed), monitoring information is limited (for example, only CPU, memory, or disk consumption of each VM is reported), and there is no support for advanced features such as cross-site networks or VM migration.

Partially coupled federation. This scenario typically consists of various partner clouds that establish a contract or framework agreement stating the terms and conditions under which one partner cloud can use resources from another. This contract can enable a certain level of control over remote resources (for example, allowing the definition of affinity rules to force two or more remote VMs to be placed in the same physical cluster); can agree to the interchange of more detailed monitoring information (for example, providing information about the host where the VM is located, energy consumption, and so on); and can enable some advanced networking features among partner clouds (for example, the creation of virtual networks across site boundaries).

Tightly coupled federation. This scenario usually includes clouds belonging to the same organization and is normally governed by the same cloud OS type. In this scenario, a cloud instance can have advanced control over remote resources—for example, allowing decisions about the exact placement of a remote VM—and can access all the monitoring information available about remote resources. In addition, it can allow other advanced features, including the creation of cross-site networks, cross-site migration of VMs, implementation of highavailability techniques among remote cloud instances, and creation of virtual storage systems across site boundaries.

Federation architectures

In practice, various federation architectures implement these coupling scenarios. Although there is no general


agreement on the classification of these architectures, Figure 2 shows the four main types of federation architectures: bursting (hybrid), broker, aggregation, and multitier. While the loosely coupled hybrid and broker architectures have been widely studied and implemented, there is still much work to be done regarding the development and implementation of more coupled architectures, especially in the case of multitier architectures.

Bursting (hybrid) architecture. As Figure 2a shows, the cloud bursting or hybrid architecture combines the existing on-premise infrastructure (usually a private cloud) with remote resources from one or more public clouds to provide extra capacity to satisfy peak demand periods. Because the local cloud OS has no advanced control over the virtual resources deployed in external clouds beyond the basic operations the providers allow, this architecture is loosely coupled. Most existing open cloud managers support the hybrid cloud architecture, which has been explored in various research efforts^{9,10} and is used in infrastructures such as StratusLab (http://stratuslab.eu).

Broker architecture. The central component of this architecture, shown in Figure 2b, is a broker that serves various users and has access to several public cloud infrastructures. A simple broker should be able to deploy virtual resources in the cloud as selected by the user. An advanced broker offering service management

capabilities could make scheduling decisions based on optimization criteria such as cost, performance, or energy consumption to automatically deploy virtual user services in the most suitable cloud, or it could even distribute the service components across multiple clouds. This architecture is also loosely coupled since public clouds typically do not allow advanced control over the deployed virtual resources.

Brokering is the most common federation scenario, with many commercial and open source brokering services operating in the cloud market. In addition to research projects,¹¹ experimental multisite cloud infrastructures based on the broker architecture include BonFIRE (www. bonfire-project.eu), Open Cirrus,¹² and FutureGrid (http:// futuregrid.org).

Aggregated architecture. As Figure 2c shows, cloud aggregation consists of two or more partner clouds that interoperate to aggregate their resources and provide users with a larger virtual infrastructure. This architecture is usually partially coupled, since partners could be provided with some kind of advanced control over remote resources, depending on the terms and conditions of contracts with other partners. These partner clouds usually have a higher coupling level when they belong to the same corporation than when they are owned by different companies that agree to cooperate and aggregate

RESEARCH FEATURE

their resources. The Reservoir federated infrastructure is an example of an aggregated cloud architecture.³

Multitier architecture. The multitier architecture, shown in Figure 2d, consists of two or more cloud sites, each running its own cloud OS and usually belonging to the same corporation, that are managed by a third cloud OS instance following a hierarchical arrangement. This upper cloud OS instance has full control over resources in different cloud sites—a tightly coupled scenario—and it exposes the resources available in the different cloud sites as if they were located in a single cloud. This architecture is beneficial for corporations with geographically distributed cloud infrastructures because it provides uniform access. It is also useful for implementing advanced management features such as high availability, load balancing, and fault tolerance.

irtualization and cloud computing technologies are gaining increasing acceptance from IT companies and institutions for use in the deployment of efficient, flexible, and scalable datacenters. The definition of an architecture reference model for IaaS clouds is essential for the widespread adoption of these technologies. The core component of this architecture—the cloud OS—is responsible for managing and monitoring the physical and virtual infrastructures, providing abstraction of the underlying infrastructure, and offering different tools and advanced functionality for cloud users. In addition, the cloud OS must offer federation capabilities to allow IT companies to scale out their local datacenters with external resources or to share and aggregate resources with partner infrastructures to increase computing capacity and reduce costs.

Acknowledgments

This research was partially supported by la Consejería de Educación of Comunidad de Madrid, el Fondo Europeo de Desarrollo Regional (FEDER), Fondo Social Europeo (FSE) through MediaNet Research Program S2009/ TIC-1468, el Ministerio de Economía y Competitividad through research grant TIN2012-31518, and by the European Union through the 4CaaSt (EU grant agreement 258862) and BonFIRE (EU grant agreement 257386) research projects.



References

- 1. M. Rosenblum and T. Garfinkel, "Virtual Machine Monitors: Current Technology and Future Trends," *Computer*, Apr. 2005, pp. 39-47.
- 2. R. Buyya et al., "Cloud Computing and Emerging IT Platforms: Vision, Hype, and Reality for Delivering Computing as the 5th Utility," *Future Generation Computer Systems*, vol. 25, no. 6, 2009, pp. 599-616.
- 3. B. Rochwerger et al., "Reservoir—When One Cloud Is Not Enough," *Computer*, Mar. 2011, pp. 44-51.
- 4. A.J. Ferrer et al., "OPTIMIS: A Holistic Approach to Cloud Service Provisioning," *Future Generation Computer Systems*, vol. 28, no. 1, 2012, pp. 66-77.
- 5. S.A. Chaves et al., "Toward an Architecture for Monitoring Private Clouds," *IEEE Comm.*, Dec. 2011, pp. 130-137.
- S. Ortiz, "The Problem with Cloud-Computing Standardization," *Computer*, Mar. 2011, pp. 13-16.
- 7. K. Keahey et al., "Sky Computing," *IEEE Internet Computing*, vol. 13, no. 5, 2009, pp. 43-51.
- D. Petcu, "Portability and Interoperability between Clouds: Challenges and Case Study," LNCS 6994, Springer, 2011, pp. 62-74.
- 9. B. Sotomayor et al., "Virtual Infrastructure Management in Private and Hybrid Clouds," *IEEE Internet Computing*, vol. 13, no. 5, 2009, pp. 14-22.
- R. Moreno-Vozmediano et al., "Multicloud Deployment of Computing Clusters for Loosely Coupled MTC Applications," *IEEE Trans. Parallel and Distributed Systems*, vol. 22, no. 6, 2011, pp. 924-930.
- J. Tordsson et al., "Cloud Brokering Mechanisms for Optimized Placement of Virtual Machines across Multiple Providers," *Future Generation Computer Systems*, vol. 28, no. 2, 2012, pp. 358-367.
- 12. A.I. Avetisyan et al., "Open Cirrus: A Global Cloud Computing Testbed," *Computer*, Apr. 2010, pp. 35-43.

Rafael Moreno-Vozmediano is an associate professor at Complutense University of Madrid (UCM). His research interests include high-performance and distributed computing, virtualization, and cloud computing. Moreno-Vozmediano received a PhD in computer architecture from UCM. Contact him at rmoreno@dacya.ucm.es.

Rubén S. Montero is the chief architect of the OpenNebula Project, a cofounder of C12G Labs, and an associate professor at Complutense University of Madrid. His research interests include resource provisioning models for distributed systems and cloud computing. Montero received a PhD in computer architectures from UCM. Contact him at rsmontero@opennebula.org.

Ignacio M. Llorente is the director of the OpenNebula Project, a cofounder of C12G Labs, and a full professor at Complutense University of Madrid. His research interests include highperformance computing, virtualization, cloud computing, and grid technology. Llorente received a PhD in computer architecture from UCM and an executive MBA from the Instituto de Empresa. He is a member of IEEE. Contact him at imllorente@ opennebula.org.

CN Selected CS articles and columns are available for free at http://ComputingNow.computer.org.

Factorizing Event Sequences

Naren Sundaravaradan and Naren Ramakrishnan, Virginia Tech David A. Hanauer, University of Michigan Medical School



Factorizing interleaved event sequences, such as those found in electronic medical records, into simpler processes can yield new insights from large datasets.

arge datasets often include sequences of events. For example, electronic medical records (EMRs) can be viewed as sequences of ICD-9 and CPT-4 codes. ICD-9 (International Classification of Diseases, ninth revision) is a coding system for injuries, diseases, and other health-related conditions, while CPT-4 (Current Procedural Terminology, fourth edition) is a categorization system for medical procedures such as surgeries and lab tests.

In EMRs, it's common for such coding sequences to be interleaved—thus, a patient's heart history is

intermingled with x-ray reports for a sports-related injury, or a plastic surgery procedure is mixed in with a record of kidney disease.

While one subsystem of the body can have complicating side effects on another subsystem, there's a need to factorize event sequences across a patient population into nonredundant processes to discover clinically relevant patterns.

EVENT SEQUENCE FACTORIZATION

Event sequence factorization draws on both process and sequence



mining. Process mining uses temporal data to reconstruct a process represented, say, by a Kripke structure or Petri net—that could account for it. Sequence mining aims to identify patterns in sequences that recur frequently or that optimize a user-defined objective.

There is an inherent tradeoff between mining local patterns, which are more efficient to mine and can be quite detailed, and global patterns, which yield more succinct representations. Our approach, which focuses on separating interleaved event sequences, finds the "sweet spot" between these two options. While mining local data, event sequence factorization doesn't generate an excessive number of patterns, but at the same time it can define a global model of the underlying processes.

Figure 1 shows two simple examples of event sequence factorization.

In Figure 1a, the factor sequences RABCD and EWXYZD generate the sequence AWXBYCDZ. Every element in the generated sequence must come from one of the given factors. Note that an element, such as D, can occur in multiple processes. Also, there must be an order-preserving mapping from a subset of factor elements to the generated sequence. However, factorization restricts these subsets

Kidneys

(D584.9: renal failure, acute nitric oxide synthase [NOS]) -> (DV42.0: transplant, kidney) -> (D585.6: renal disease, end stage) -> (D414.00: coronary atherosclerosis of unspecified vessel) -> (D577.1: pancreatitis, chronic)

Leading to a psychotherapy session

(D256.9: dysfunction, ovarian NOS) -> (D628.9: infertility, female NOS) -> (P90806: psychotherapy, office, 45-50 min.) -> (D309.9: reaction, adjustment NOS)

(D611.9: disorder, breast NOS) -> (P76091: mammogram, both breasts) -> (D611.72: lump or mass in breast) -> (P76090: mammogram, one breast) -> (D311: disorder, depressive, not elsewhere classified [NEC]) -> (P90844: psychotherapy, individual, 45-50 min.)

(D310.1: personality change due to clubbing, cyanosis, and edema [CCE]) -> (P90812: interactive psychotherapy, office, 45-50 min.) -> (D294.9: disorder, persistent mental, due to CCE/NOS) -> (D300.3: disorder, obsessive-compulsive) -> (P90806: psychotherapy, office, 45-50 min.)

Brain

(D345.01: epilepsy, generalized nonconvulsive, with intractable epilepsy) -> (P95819: EEG, awake and asleep) -> (D345.10: epilepsy, generalized convulsive) -> (P70553: MRI brain with and without dye) -> (D191.9: neoplasm, malignant, brain NOS)

(P70553: MRI brain with and without dye) -> (DV67.2: chemotherapy follow-up) -> (D191.3: neoplasm, malignant, brain, parietal lobe) -> (P96413: chemotherapy, IV infusion, 1 hr.)

Psoriasis

(D696.0: psoriatic arthritis) -> (D696.1: psoriasis NEC) -> (P99222: initial hospital care) -> (D244.9: hypothyroidism NOS) -> (D250.00: diabetes mellitus)

(D696.1: psoriasis NEC) -> (D250.02: type 2 diabetes) -> (D457.1: lymphedema NEC) -> (D250.00: diabetes mellitus) -> (P99222: initial hospital care)

Pelvis

(P73510: x-ray exam of hip) -> (D715.95: osteoarthrosis NOS, pelvis/ thigh) -> (DV43.64: hip joint replacement status) -> (D733.90) -> (P81000: urinalysis, nonautomated, with scope) -> (DV42.0: transplant, kidney) -> (P72170: x-ray exam of pelvis)

(D617.3: endometriosis, pelvic peritoneum) -> (D628.9: infertility, female NOS) -> (P76857: ultrasound exam, pelvic, limited) -> (D256.9: dysfunction, ovarian NOS)

Liver

(D996.82: complication, liver transplant) -> (P99141: sedation, conscious, IV, intramuscular, Isoniazid) -> (P75984: x-ray control catheter change) -> (P74305: x-ray bile ducts/pancreas) -> (P47525: change bile duct catheter) -> (P47505: injection for liver x-rays)

Lungs

(DV42.6: transplant, lung) -> (D996.84: complication, transplanted lung) -> (D792.9: abnormal finding, body substance NEC) -> (D512.8: pneumothorax, spontaneous NEC) -> (D212.3: neoplasm, benign, bronchus/lung) -> (P88312: special stains)

(D496: obstruction, chronic airway NEC) -> (P94720: monoxide diffusing capacity) -> (P94360: measure airflow resistance) -> (P94240: residual lung capacity) -> (P94060: evaluation of wheezing)

(P99283: emergency dept. visit) -> (D493.90: asthma) -> (D530.1: esophagitis NOS) -> (D518.82: insufficiency, pulmonary NEC) -> (D518.3: eosinophilia, pulmonary) -> (D493.91: asthma NOS with status asthmaticus) -> (P94010: breathing capacity test) -> (P94720: monoxide diffusing capacity) -> (P94240: residual lung capacity)

Others

(D427.41: fibrillation, ventricular) -> (P93737: analyze cardioverterdefibrillator without reprogramming) -> (D185: neoplasm, malignant, prostate) -> (D715.91: osteoarthrosis NOS, shoulder) -> (P20610: drain/ inject, joint/bursa)

(D250.01: diabetes mellitus, uncomplicated, type 1) -> (D240.9: goiter NOS) -> (P76536: ultrasound exam of head and neck) -> (D784.2: swelling in head/neck)

(D696.1: psoriasis NEC) -> (D571.2: cirrhosis, alcoholic, liver) -> (D696.0: psoriatic arthropathy)

Figure 2. Processes discovered by factorizing event sequences in an electronic medical record dataset.

to form a contiguous subsequence in this case, ABCD in the first factor sequence and WXYZ in the second. Note that some factor elements, such as R, need not appear in the generated sequence.

In Figure 1b, one of the factor sequences, BYDAWX, is repeated. Again, note the order-preserving mapping and the contiguity of subsequences AWX and BYD.

Given a particular sequence, it's easy to derive many factorizations from it. But given a large database of sequences, the goal is to derive a small set of processes that can generate all of the sequences in the database.

SEQUENCE FACTORIZATION ALGORITHM

Details of the sequence factorization algorithm we developed are beyond the scope of this article, but it's essentially incremental. We construct and maintain a model by adding one process at a time. As the algorithm encounters a dataset sequence, it maintains a working set of processes, each in one of four states: *waiting*—the process hasn't yet been factorized; *converging*—the process has been factorized but hasn't yielded sufficient evidence to make a decision; *converged*—the process is chosen to be included in the model (over its factorization); and *invalid* the process is no longer needed but can't be removed because another process depends on it. As the descriptions indicate, a process traverses the states in order: waiting, converging, converged, and possibly invalid.

Because the algorithm incrementally computes a factorization, the resulting model isn't optimal. Nevertheless, it significantly compresses EMR data. Furthermore, it has revealed several patterns about medical diagnoses and procedures.

EXPERIMENTAL RESULTS

With approval from the University of Michigan Health System, we organized a dataset of de-identified information from about 1.6 million patients who received care there. The actual medical records contained about 100 million time-stamped ICD-9 and CPT-4 codes.

We ran our algorithm on three sets of 150,000 patients with an alphabet size of 10,000 for each set. To further condense the representation of patient records, we collapsed a sequence of contiguous identical events into one event—thus, AABBCCBD became ABCBD.

As with most large-scale studies involving the discovery of clinical associations, we manually reviewed a subset of the data to determine significant and interesting patterns. Figure 2 shows a sampling of results from our analysis.

Many of the processes we discovered were consistent with known medical information.

For example, ventricular fibrillation (Dx 427.41) -> automatic implantable cardioverter-defibrillator check (Px 93737) -> malignant prostate cancer (Dx 185) -> shoulder osteoarthrosis (Dx 715.91) -> shoulder joint injection (Px 20610) is a process that might be found in an elderly man. Similarly, the medical history of a patient with an autoimmune disorder might include type 1 diabetes (Dx 250.01) -> goiter (Dx 240.9) -> neck ultrasound (Px 76536) -> neck swelling (Dx 784.2). Goiter is often associated with autoimmunethyroiditis, and type 1 diabetes is also an autoimmune disorder.

Other processes revealed unfortunate stories about patients.

For example, intractable seizures (Dx 345.01) -> EEG, brain scan (Px 96819) -> brain MRI, to look for pathology (Px 70553) -> malignant brain neoplasm (Dx191.9) indicates that a patient had severe seizures and, after a workup to determine the cause, was found to have a brain tumor. Another example is ovarian dysfunction (Dx 256.9) -> female infertility (Dx 628.9) -> psychotherapy (Px 90806) -> adjustment reaction (Dx309.9), which is a psychiatric diagnosis defining a significant emotional response to a specific stressor—in this case, being unable to bear children.

Some patterns were clinically interesting but are less well known in the medical domain. For example, we discovered an association between psoriasis and hypothyroidism that has been documented but is rare: psoriatic arthritis (Dx 696.0) -> psoriasis (Dx 696.1) -> hospitalization (Px 99222) -> hypothyroidism (Dx 244.9) -> diabetes mellitus (Dx 250.00).

Another interesting pattern with psoriasis involves lymphedema, which is swelling caused by blockage of the lymphatic system: psoriasis (Dx 696.1) -> type 2 diabetes (Dx 250.02) -> lymphedema (Dx 457.1) -> diabetes mellitus (Dx 250.00) -> hospitalization (Px 99222). This condition has also been reported in the medical literature.

Both patterns include diabetes, which might be expected given that

psoriasis and diabetes are both associated with elevated body mass index and obesity.

Finally, some of the temporal patterns we discovered are quite complex, such as acute renal failure (Dx 584.9) -> kidney transplant (Dx V42.0) -> end stage renal disease (Dx 585.6) -> coronary atherosclerosis (Dx 414.00) -> chronic pancreatitis (Dx 577.1). Acute (as opposed to chronic) pancreatitis is a known cause of acute renal failure, but this process suggests that chronic renal failure—akin to end stage renal disease—might cause pancreatitis.

e've developed a novel approach to factorizing events sequences into a small set of processes, and have demonstrated its effectiveness in deriving insights from EMR data. A major advantage of our approach is that it can be used in a distributed data mining setting, making it ideal for mining remote databases as well as for when privacy preservation is important. Another benefit is that this approach combines local and global considerations of pattern mining.

Naren Sundaravaradan is a PhD student in the Department of Computer Science at Virginia Tech. Contact him at narens@vt.edu.

Naren Ramakrishnan, Discovery Analytics column editor, is the Thomas L. Phillips Professor of Engineering in the Department of Computer Science and director of the Discovery Analytics Center at Virginia Tech. Contact him at naren@vt.edu.

David A. Hanauer, MD, is assistant director, Comprehensive Cancer Center Bioinformatics Core, University of Michigan Medical School. Contact him at hanauer@umich.edu.

Editor: Naren Ramakrishnan, Dept. of Computer Science, Virginia Tech, Blacksburg, VA; naren@cs.vt.edu

Digital Fabrication

Manfred Lau, Lancaster University, UK Jun Mitani, University of Tsukuba, Japan Takeo Igarashi, University of Tokyo, Japan



For the first time in history, laypeople can participate in the product design and manufacturing process by directly interacting with the underlying hardware and software.

any of the items that we use on a daily basis are created by professional designers, mass-produced at factories, and then transported, through a complex distribution network, to regional warehouses and local retail sales outlets.

In the past, consumers often had to physically go to a neighborhood store to browse through existing products and make a purchase. Although mailorder catalogues have been around since the mid-eighteenth century, the number of goods that could be shipped directly to homes was extremely limited, and it was difficult to evaluate offerings without seeing them in person.

The Internet has made it possible for consumers to easily compare and buy competing products, with websites providing up-todate textual descriptions, photos, videos, and customer reviews. Items purchased online can be shipped to the consumer from the nearest available distribution center rather than to local stores, reducing costs and increasing convenience. Online shopping has also made it easier to customize products—there is a dizzying array of choices for everything from T-shirt logos to sofa fabric patterns—although for most goods, consumers are still restricted to established designs.

A PARADIGM SHIFT

The next stage in the evolution of consumer product design and manufacturing is *digital fabrication*, in which individuals design products to meet their unique needs and preferences. These products are then manufactured and delivered to them on demand (N. Gershenfeld, *Fab: The Coming Revolution on Your Desktop— From Personal Computers to Personal Fabrication*, Basic Books, 2005; J.A. Landay, "Technical Perspective: Design Tools for the Rest of Us," *Comm. ACM*, Dec. 2009, p. 80).

Emerging digital fabrication technologies, such as desktop 3D printing, are increasingly affordable and might soon make what some observers have called the next industrial revolution a reality (C. Anderson, "Atoms Are the New Bits—The New Industrial Revolution," *Wired.co.uk*, 1 Feb. 2010; "The Third Industrial Revolution," *The Economist*, 21 April 2012). For the first time in history, laypeople can participate in the product design and manufacturing process by directly interacting with the underlying hardware and software. The notion that average people not just professionals—can leverage their unique creative skills and ideas to create real, physical objects is fun and exciting, and will open up a world of innovation.

Because digital content is weightless and can be moved instantly on a global scale at little cost, digital fabrication will lead to a paradigm shift not just in product design and manufacturing, but also in the storage, transportation, and energy sectors.

INTERFACES AND MODELING TOOLS

To fully realize this transformation, users need to be able to create digital 3D shapes before they can be fabricated into physical objects. Thus, a critical research challenge is developing 3D modeling tools for use by people with little or no experience in modeling and design. We have explored two types of user interfaces and modeling tools for this purpose.

Modeling-in-context

We have developed an easy-to-use interface that lets users sketch a new object using a photo of an existing object as a reference (M. Lau et al., "Modeling-in-Context: User Design of Complementary Objects with a Single Photo," Proc. 7th Symp. Sketch-Based Interfaces and Modeling [SBIM 10], Eurographics Assoc., 2010, pp. 17-24). The photo provides a background context for the user to draw a 2D sketch of the new object and annotate its 3D geometric properties. The system then creates a 3D digital model of the 2D sketch that matches the real object's dimensions.

Figure 1 illustrates the use of modeling-in-context to design a replacement lid for a teapot. The user first takes a photo of the teapot and uploads the photo to the system. The user then draws 2D lines and curves on the photo to represent the new lid, and annotates the drawing with the lid's geometric properties. An algorithm expands the 2D sketch into a 3D shape. A 3D printer fabricates the resulting digital shape into a real teapot lid that fits well with the original teapot.

Situated modeling

We have also developed an augmented-reality-based interface that lets users create models for digital fabrication by manipulating small real-world shapes, such as wooden cylinders, spheres, and prisms; the user "stamps" digital 3D copies of these shapes and combines these into a single, life-size model (M. Lau et al., "Situated Modeling:

WHERE THE PHYSICAL AND DIGITAL WORLDS COLLIDE

We surround ourselves with manmade objects that range in scale from tiny pieces of jewelry to the huge buildings where we live and work. Historically, the separation of the physical and digital worlds has been very clear, but digital fabrication is blurring that dividing line.

Thus far, digital printing technology can only generate relatively simple objects. However, artifacts in our modern world can in many cases be seen as an instantiation of a plan, which can be data or software, and fabricating programmable devices that have complex physical characteristics and functional parts is the next logical step in the technology's evolution.

Initially, digitally fabricating such smart objects will rely on combining 3D printing with functional and programmable components (A. Schmidt, T. Döring, and A. Sylvester, "Changing How We Make and Deliver Smart Devices: When Can I Print Out My New Phone?," *IEEE Pervasive Computing*, Oct.-Dec. 2011, pp. 6-9). Using computing component platforms such as Microsoft .NET Gadgeteer, it's possible to build a fully functional digital camera or game console within hours (S. Hodges et al., "A New Era for Ubicomp Development," *IEEE Pervasive Computing*, Jan.-Mar. 2012, pp. 5-9). In the long term, digital fabrication will become more sophisticated and enable direct printing of circuits as well as displays into objects.

In the not-so-distant future, the distribution and payment processes established in software—for example, app stores and in-app billing—might be applicable to physical objects as well. Concepts well researched in software engineering, such as versioning and product lines, could play a major role in all sorts of products.

Albrecht Schmidt, University of Stuttgart

A Shape-Stamping Interface with Tangible Primitives," *Proc. 6th Int'l Conf. Tangible, Embedded, and Embodied Interaction* [TEI 12], ACM, 2012, pp. 275-282).

The user wears a head-mounted display to visualize the creation of the digital model and its placement in virtual space. Instead of stamping in empty space, the interaction between the small primitive shapes and the real-world environment allows for tactile feedback, and leads to a more precise overall digital shape.

Figure 2 illustrates the use of situated modeling to design a table that will fit in the empty corner of a living room. Wearing a head-mounted display, the user manipulates a set of wooden primitive shapes that are outfitted with markers for identification purposes (Figures 2a and 2b). The user stamps digital copies of these physical shapes in the virtual 3D environment, which the user sees on the display. The corner of the room, including the physical ground and walls, serves as a background reference to create the digital model, which is later used to assemble a real table (Figure 2c).

FABRICATING PHYSICAL OBJECTS FROM DIGITAL MODELS

Fabricating physical objects from digital 3D models is another major research challenge. Although there



Figure 1. Using modeling-in-context to design a replacement lid for a teapot: (a) original photo of the teapot and 2D user sketch and annotations of the new lid; (b) 3D digital model of the new lid; (c) new physical lid from 3D printer that fits well with the original teapot.

INVISIBLE COMPUTING



Figure 2. Using situated modeling to design a table to fit in the empty corner of a living room: (a) tangible primitive wooden shapes with markers for identification purposes; (b) a user wearing the head-mounted display and holding one of the shapes; (c) digital table model created by the system, and corresponding physical table later assembled from wood pieces.

are many 3D model datasets, such as the Princeton Shape Benchmark (P. Shilane et al., "The Princeton Shape Benchmark," *Proc. Int'l Conf. Shape Modeling and Applications* [SMI 04], IEEE CS, 2004, pp. 167-178) and Google 3D Warehouse (www. sketchup.com/3dwh), as well as various software tools for creating and editing digital 3D shapes, it is not obvious how a digital 3D model can be fabricated into a physical object.

As one possible approach to this problem, we introduced a method that converts 3D furniture models into usable real-world furniture, specifically cabinets and tables (M. Lau et al., "Converting 3D Furniture Models to Fabricatable Parts and Connectors," *Proc. SIGGRAPH* 2011, ACM, 2011, article no. 85). The principal challenge is that the input models are triangular mesh models that only contain information about a shape's surface, as these models are typically used for display and visualization; they contain no information about the shape's semantics.

Our solution calls for defining formal grammars that describe how different types of cabinets and tables are constructed from individual parts and connectors (such as screws and nails). We then use these grammars to analyze a given model to extract its parts. After extracting the parts, we use data obtained beforehand from real furniture to generate the connectors. Figure 3 shows an example of the process applied to a 3D cabinet model, along with the equivalent realworld cabinet based on the structure and dimensions of the digital parts and connectors.

espite the progress we and other researchers have made in advancing digital fabrication technology, many open challenges remain.

Current 3D digital modeling interfaces only allow for designing static objects, not objects with dynamic functional parts such as a closet door or a foldable chair. Researchers must develop more general user interfaces for handling these cases.

In addition, such tools do not consider the integrity of the resulting physical objects. Constructing a digital model doesn't guarantee that the object fabricated from it will be structurally stable, which severely limits the technology's usefulness. More research is required in this area.

New business and manufacturing models are also needed. If products



Figure 3. Converting a 3D cabinet model into usable real-world furniture: (a) digital model of IKEA ALVE cabinet from Google 3D Warehouse; (b) fabricatable parts and connectors generated by our algorithm; and (c) equivalent real-world cabinet based on the structure and dimensions of the digital parts and connectors.

can be created locally on demand, less storage, transportation, and energy will be required. The related industries must accordingly adapt to changes in demand.

Intellectual property issues must be addressed. Emerging commercial services such as Ponoko and Shapeways offer the ability to design, buy, and sell user-created 3D digital models. As more of these models are shared online, there must be adequate mechanisms in place to protect IP rights and prevent copyright infringement.

Manfred Lau is an assistant professor (lecturer) in the School of Computing and Communications at Lancaster University, UK. Contact him at m.lau@ lancaster.ac.uk.

Jun Mitani is an associate professor in the Department of Computer Science

at the University of Tsukuba, Japan. Contact him at mitani@cs.tsukuba. ac.jp.

Takeo Igarashi is a professor in the Department of Computer Science at the University of Tokyo, Japan. Contact him at takeo@acm.org.

Cn Selected CS articles and columns are available for free at http://ComputingNow.computer.org.

computing now

NEW+EXPANDED

Learn industry solutions you can use from practical articles, case studies, blogs, and interviews that address high-interest, focused areas of technology.

Mobile Computing • Cloud Computing • Security • Software • High-Performance Computing • Networking



IEEE (Computer society

Visit http://computingnow.computer.org

SECURITY

Atomic-Level Security for Web Applications in a Cloud Environment

Arnold Brown, US Navy Benjamin Apple, US Department of the Navy James Bret Michael, Naval Postgraduate School Michael Schumann, US Navy



Adopting an atomic-separation paradigm for provisioning the level of access control needed in a cloud-computing environment requires taking a holistic approach to security.

s organizations rush to gain the operational efficiencies that collaborative Web applications such as social networking tools promise, decision makers often ignore effective information security or are simply unaware of the weaknesses and possible consequences associated with these tools' use in a highly distributed, data-centric cloudcomputing environment (K. Johnson and E. Savitz, "Cloud Computing Hides Big Issues in Corporate Data Sharing," Forbes.com, 17 Feb. 2012).

Many cloud-computing-based applications aren't designed to support the stringent information security controls that corporate and government markets demand. Consequently, information security professionals responsible for assessing the compliance of these cloud-based applications often find that they were designed for use in a single security domain.

When designers fail to anticipate the additional security requirements of enterprise cross-domain computing, the security controls necessary for a complex multiclassification environment often are either absent or weak. The potential for increased security risks alone isn't sufficient to dissuade senior decision makers from seeking to take advantage of the operational efficiencies associated with using the cloud (J. Jaeger, "Shop Talk: Cloud Computing Poses New Risks, Opportunities," *Compliance Week*, 15 Feb. 2011).

MEETING EXPECTATIONS

Information security professionals are challenged to meet senior decision makers' expectations by ensuring the confidentiality, integrity, and availability of high-value information assets for their organization while facilitating the sharing of sensitive information such as intellectual property, intelligence, or financial data in a dynamic, collaborative environment (S.F. Gale, "Who's Guarding the Cloud?" *PM Network*, Mar. 2012, pp. 30-37).

This high level of information sharing requires information security professionals to adopt an object- or atomic-separation paradigm that provides strong access controls (that is, kernel space) using fine-grained access decisions, as opposed to the traditional network-centric paradigm that grants broad user-access levels based on network-security services (J.W. DeWitt, "Cyber Risk in 2012: Get Your Head in the Cloud," *National Underwriter/P&C*, Mar. 2012, pp. 18-19).

Strong access control is particularly important for government organizations because they can't purchase indemnity against loss or rebrand themselves when a major information security incident occurs. Adopting an atomic-separation paradigm for provisioning the level of control needed in a cloud-computing environment requires taking a holistic approach to security.

REAL-TIME COLLABORATIVE ENVIRONMENT

Our research project is a broadly scoped, holistic engineering effort focused on developing a secure cloudcomputing environment that supports highly collaborative enterprise Web applications. This real-time collaborative environment reaches beyond database-cell-level security, abstracting the end points from the network infrastructure for confidentiality and integrity risks as well as changing existing information management and usage paradigms to provide an end-to-end solution that addresses root inhibitors to secure sharing.

We aim to reduce the barrier to entry for developers as well as decrease the time and cost to deliver secure applications by providing an open source licensed development platform in an environment that includes a robust set of security controls. This will enable developers to contribute software-as-a-service applications to a centrally managed environment that meets the high standard of information system and organization controls required in the current cyberthreat environment. The atomic-level security (ALS) access control model is a core element of this approach.

THE PROBLEM

Historically, fine-grained access control was constrained to select applications used to protect operationally critical information. However, today's cyberthreat environment demands stronger information access controls in all applications.

In addition, today's constrained fiscal environment demands distributed, scalable, and costeffective solutions for all phases of the system development life cycle. This downward fiscal pressure has also stimulated efforts to reduce infrastructure and support costs, with a focus on virtualizing datacenters and desktops. However, this only leads to virtualization of current paradigms with applications designed using controls intended for operation in a single security domain.

Although these efforts might lead to cost savings, they don't address multilevel cross-securitydomain issues that inhibit realtime collaboration in an enterprise environment.

DIFFERING SECURITY POLICIES

Enterprises are often divided into networks managed under different security policies, with *cross-domain access* and *transfer devices* filling most secure-information-sharing needs. Cross-domain access devices allow secure simultaneous access to desktops in different network security domains. Transfer devices enable secure and reliable movement of files between network security domains. However, these devices enforce a container-separation paradigm in which networks, applications, and files are used to manage information

The network-centric and container-separation paradigms make it difficult to meet user expectations for interactive collaboration.

as a composite versus an aggregate. This creates physical choke points and introduces transaction latency and, for transfer devices, leads to multiple unsynchronized copies of files across network security domains.

A third type of cross-domain device, *multilevel security* (MLS), applies fine-grained access control to more granular information. MLS is defined as the concept of "processing information with different classifications and categories that simultaneously permits access by users with different security clearances and denies access to users who lack authorization" (R. Kissel, *Glossary of Key Information Security Terms*, NIST IR 7298, rev. 1, Nat'l Inst. Standards and Technology, Feb. 2011, pp. 122-123).

Often overlooked is the key phrase "processing information with different classifications," which has led to confusing MLS with *multiple security levels* (MSLs). With MSLs, an information system "is trusted to contain, and maintain separation between, resources (particularly stored data) of different security domains."

Over the years, MLS has been perceived as complex and inflexible, and its goals have proven both elusive and costly, leading to the adoption of network-centric and containerseparation security paradigmsfor example, MSLs via multiple networks and access solutions. However, implementing these paradigms has also proven costly in terms of infrastructure, licensing, and maintenance, and they don't scale well enough to meet users' and decision makers' expectations for data transparency (B. Ames and F. Brown, "Auditing the Cloud," Internal Auditor, vol. 68, no. 4, 2011, pp. 35-39).

The assumption of an *air gap*, or physical separation, between networks of different security domains is flawed in reality, particularly for cross-domain collaboration.

FINDING SOLUTIONS

As users and decision makers increase their expectations for collaboration while applying pressure to reduce IT costs, internal and external IT service providers are responding by rapidly adopting virtualization and other cloud-enabling technologies. However, moving to the cloud might not resolve the issues resulting from managing information as a composite as opposed to an aggregate of portioned data.

In addition, the network-centric and container-separation paradigms make it inordinately difficult to meet user expectations for interactive collaboration because validation, malware scanning, and malware removal are pushed to an intermediary who must operate on a composition of information—for example, a large brief with mixed content—at network speed.

In contrast, the atomic-separation paradigm addresses these issues

SECURITY



Figure 1. The atomic-level security model architecture factors the policy decision point into components, allowing decisions to occur at different times within a session and at different locations in a distributed architecture.

closer to the point of origin, where they are inherently distributed and human or data ingest speed is the reference. To implement the atomicseparation paradigm, information security professionals need an endto-end access control model that operates at the atomic level of information assets, that is, in portions.

ATOMIC-LEVEL SECURITY

The ALS access control model is based on the atomic-separation paradigm, in which information is created and managed at the portion level from end to end—that is, client, application, database, and cyber ecosystem. This differs from the container-separation paradigm in which information is collected and managed in networks, applications, and files as a composite.

When managed as a composite, the associated metadata characterizes the container rather than the individual portions of information it contains. This makes it more difficult to differentiate or validate the classification prior to transferring the contents.

DIFFERENTIATING BETWEEN ALS AND MLS

A fundamental difference between the ALS and MLS models is that ALS is designed to offer the dynamism and flexibility required to scale on the order of n labels as opposed to mnetworks, where n >> m. Thus, the ALS access control model supports atomic-level collaboration within dynamic groups in which a single network infrastructure can host all users and devices.

The design strategy is to be dynamic, flexible, and granular while establishing session subject labels centrally, and less flexible and granular as the model scales toward the client end points. ALS is intended to provide dynamic, distributed, fine-grained access control for database-centric enterprise Web applications, particularly those in which the database is part of the access control mechanisms providing its protection as well as a trusted source of persistent information.

ALS is characterized as atomic because the smallest portion of information that must be classified for example, a paragraph or imagerather than the nature of its access control—for example, a role or policy—determines the access control resolution.

The basis for characterization is different because the model applies logic and attributes normally associated with role-based access control, attribute-based access control, and policy-based access control to arrive at intermediate decisions that inform mandatory access control (MAC) decisions.

ALS also includes a spillcontainment strategy involving collaboration group (CG), CG manager (CGM), application, project, and multitag subject and object labels. Access requests are always in context, meaning the request is from a specific subject representing a user in an assigned role in a CG that the responsible CGM associates with a specific project.

In an implementation of the ALS access control model, the system authenticates both the user and client and uses the clearance and accesses of each to produce a subject label representing the pair. This allows users to have more than one persona, each correlated to a specific certificate, while constraining distribution to applicable CGs on a need-to-know basis.

The policy-driven decision process has three components:

- an attribute-based subject label composition,
- a role-based preliminary access decision, and
- a label-based final access decision.

As Figure 1 shows, this enables distribution of the components to achieve the effect of policy-driven attribute-, role-, and label-based decisions at different points in a distributed architecture informing MAC decisions. Currently, attribute-based decisions require storing user information in many application and authorization

servers, leading to data maintenance, trusted source, and personally identifiable information security issues. In contrast, the ALS model limits persistence of user and device attributes to their authoritative source.

Policies defined a priori and applied during subject-label formulation prior to each session govern membership in CGs. Each tag within the subject label is formed by applying user and device attributes from their authoritative source, along with a set of rules, as input to a policy engine. The policies establish or validate a subject's current membership status—whether it's implicit or explicit—before that CG's tag is included in the multitag subject label.

Once included, each tag serves as a token representing the result of a potentially complex application-space decision involving user and client affiliation attributes. However, unlike an access decision, an individual CG membership decision leads only to its inclusion in a subject label for a session. This allows creation of dynamic subject labels that update when triggered by a change of user or client attributes persistent in their authoritative source.

aking a holistic approach scoped to a cloud Web application environment has led to the development of new user interface, information management, and access control paradigms. Together they support a complete atomic-level secure enterprise Web application environment.

The ALS access control model is a core element that reaches beyond database-cell-level security to apply strong kernel-space access controls throughout the environment and its supporting cyber ecosystem. By focusing information management and access control at the atomic level, end-to-end solutions that address root inhibitors make it possible to ensure security in a dynamic, real-time, collaborative environment.

We are developing an open source reference implementation to demonstrate opportunities presented by approaching cloud deployment from a different perspective to achieve significantly enhanced security.

Acknowledgments

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements of the US government.

Arnold Brown is a commander in the US Navy. Contact him at karma.als. project@gmail.com.

Benjamin Apple is a computer security expert with the US Department of the Navy. Contact him at bengapple@gmail.com.

James Bret Michael is a professor in the Naval Postgraduate School's Computer Science and Electrical and Computer Engineering departments. Contact him at bmichael@nps.edu.

Michael Schumann is a commander in the US Navy. Contact him at maschuma@nps.edu.

Editor: Jeffrey Voas, National Institute of Standards and Technology; jeffrey.m.voas@gmail.com



Awareness in Software-Intensive Systems

Emil Vassev and Mike Hinchey Lero—the Irish Software Engineering Research Centre



Closely related to artificial Intelligence, awareness depends on the knowledge transferred to software-intensive systems so they can use it to exhibit intelligence.

onceptually, awareness is a product of knowledge and monitoring. A large class of software-intensive systems-including those for industrial automation, consumer electronics, airplanes, automobiles, medical devices, and civic infrastructure-must interact with the physical world. More advanced systems, such as unmanned autonomous systems, don't just interact but also perceive important structural and dynamic aspects of their operational environment. To become interactive, an autonomous software system must be aware of its physical environment and whereabouts, as well as its current internal status. This ability helps intelligent software systems sense, draw inferences, and react.

Closely related to artificial intelligence, awareness depends on the knowledge we transfer to software systems so they can use it to exhibit intelligence. In addition to knowledge, artificial awareness also requires a means of sensing changes so that the system can perceive both external and internal worlds through raw events and data. Thus, self and environmental monitoring are crucial to awareness: to exhibit awareness, software-intensive systems must sense and analyze their internal components and the environment in which they operate. Such systems should be able to notice a change and understand its implications. Moreover, an aware system should apply both pattern analysis and pattern recognition to determine normal and abnormal states.

Ideally, awareness should be part of the cognitive process that underlies learning. An efficient awareness mechanism should also rely on both past experience and new knowledge. Awareness via learning is the basic mechanism for introducing new facts into the cognitive system—other possible ways are related to interaction with a human operator who manually introduces new facts into the knowledge base. Here, we clarify the nature of artificial awareness and its impact on contemporary software-intensive systems.

CLASSES OF AWARENESS

Awareness generally is classified into two major areas: *self-awareness*,

pertaining to the internal world; and *context-awareness*, pertaining to the external world. Autonomic computing research defines these two classes ("Autonomic Computing: IBM's Perspective on the State of Information Technology," *IBM Autonomic Computing Manifesto*, 2001; www.research. ibm.com/autonomic/manifesto/ autonomic_computing.pdf):

- A self-aware system has detailed knowledge about its own entities, current states, capacity and capabilities, physical connections, and ownership relations with other systems in its environment.
- A context-aware system knows how to sense, negotiate, communicate, and interact with environmental systems and how to anticipate environmental system states, situations, and changes.

Perhaps a third class could be *situational awareness*, which is self-explanatory; other classes could draw attention to specific problems, such as operational conditions and per-

formance (operational awareness), control processes (control awareness), interaction processes (interaction awareness), and navigation processes (navigation awareness). Although classes of awareness can differ by subject, they all require a subjective perception of events and data "within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (M.R. Endsley, "Toward a Theory of Situation Awareness in Dynamic Systems," *Human Factors*, vol. 37, no. 1, 1995, pp. 32-64).

To better understand the idea of awareness in software-intensive systems, consider an exploration robot. Its navigation awareness mechanism could build a map on the fly, with landmarks represented as part of the environment knowledge, so that navigation becomes simply a matter of reading sensor data from cameras and plotting the robot's position at the time of observation. Via repeated position plots, the robot's course and land-reference speed can be established.

STRUCTURING AWARENESS

Recent research efforts have focused on awareness implementations in software-intensive systems. For example, commercially available server-monitoring platforms, such as Nimbus (www.nimbusproject. org) and Cittio's Watch Tower (www. networkcomputing.com/dataprotection/cittios-watchtower-30/229611534), offer robust, lightweight sensing and reporting capabilities across large server farms. Such solutions are oriented toward massive data collection and performance reporting, so they leave much of the final analysis and decision making to a human administrator. Other approaches achieve awareness through model-based detection and response based on offline training and models constructed to represent different scenarios that the system can recognize at runtime.



Figure 1. The awareness pyramid. The first three levels include monitoring tasks; the fourth, recognition tasks; the fifth and sixth, assessment tasks; and the last, learning tasks.

To function, the mechanism implementing the awareness must be structured to take into consideration different stages—for example, it might be built over a complex chain of functions such as raw data gathering, data passing, filtering, conversion, assessment, projection, and learning. As Figure 1 shows, ideally, all the awareness functions could be structured as an awareness pyramid, forming the mechanism that converts raw data into conclusions, problem prediction, and eventually learning.

The pyramid levels in Figure 1 represent awareness functions that can be grouped into four specific tasks:

- monitoring collects, aggregates, filters, manages, and reports internal and external details such as metrics and topologies gathered from the system's internal entities and its context;
- recognition uses knowledge structures and data patterns to aggregate and convert raw data into knowledge symbols;
- assessment tracks changes and determines points of interest, generates hypotheses about situations involving these points, and recognizes situational patterns; and
- *learning* generates new situational patterns and maintains a history of property changes.

Aggregation can be included as a subtask at any function level; it's intended to improve overall awareness performance. For example, it can pull together large amounts of sensory data during the filtering stage or recognition tasks can apply it to improve classification.

The awareness process isn't as straightforward as it might seem rather, it's cyclic, with several iterations over the various awareness functions. Closing the chain of awareness functions can form an awareness control loop in which different awareness classes can emerge (E. Vassev and M. Hinchey, "The Challenge of Developing Autonomic Systems," *Computer*, Dec. 2010, pp. 93-96).

The process's cyclic nature is why awareness itself is so complex, with several levels of exhibition and degrees of perception. The levels can be related to data readability and reliability—that is, they might include noisy data that must be cleaned up and eventually interpreted with some degree of probability.

Other levels might include *early awareness*, which is a product of one or two passes of the awareness control loop, and *late awareness*, which should be more mature in terms of conclusions and projections. Similar to humans who react to their first impression but then find that a later and better realization of the

SOFTWARE TECHNOLOGIES



Figure 2. Implementing awareness. KnowLang provides the constructs and mechanisms for specifying knowledge models at the ontology and logic foundation levels.

situation shifts their reaction, an aware software system should rely on early awareness to react quickly to situations when fast reaction is needed and on late awareness when more precise thinking is required.

IMPLEMENTING AWARENESS

The four awareness functions require a comprehensive and wellstructured knowledge base (KB) to hold knowledge representation (KR) symbols that can express the system itself with its proper internal structures and functionality as well as the environment (E. Vassev and M. Hinchey, "Knowledge Representation and Reasoning for Intelligent Software Systems," *Computer*, Apr. 2011, pp. 96-99).

Building an efficient awareness mechanism requires properly integrating the awareness pyramid within the implemented software system. The goal is to provide a means of monitoring and KR with a reasoner supporting awareness reasoning. KR adds a new open-world KR context to the program, and the reasoner operates in this context, taking into account the monitoring activities that drive the awareness control loop and deliver awareness results to the system itself.

Figure 2 depicts an approach the Autonomic Service-Component Ensembles (ASCENS) FP7 European Project is using to tackle the problem (www.ascens-ist.eu).

In this approach, KnowLang, a special KR language, provides the constructs and mechanisms for specifying knowledge models at the ontology and logic foundation levels (E. Vassev and M. Hinchey, "Knowledge Representation for Cognitive Robotic Systems," Proc. 15th IEEE Int'l Symp. Object/Component/ Service-oriented Real-time Distributed Computing Workshops, IEEE CS, 2012, pp. 156-163). To specify knowledge with KnowLang, we need to think about domain concepts and their properties and functionalities, important states of major concepts, objects as realizations of concepts, relations that show how concepts and objects connect to each other, self-adapting scenarios for the system in question, remarkable behavior in terms of policies driving the system out of specific situations, and other important specifics that can be classified as concepts.

As Figure 2 shows, the KB comprises KR structures such as concept trees, object trees, and concept maps. The system talks to the KnowLang reasoner via a predefined set of TELL and ASK operators, forming a communication interface that connects both the system and the KB. TELL operators feed the KB with important information driven by errors, executed actions, new sensory data, and so forth, thus helping the KnowLang reasoner update the KR context with recent changes in both the system and execution environment. The system uses ASK operators to receive recommended behavior, where knowledge is used against perception to generate appropriate actions in compliance with specific goals and beliefs. In addition, ASK operators can provide awareness-based conclusions about the system or environment's current state and ideally with behavior models for reaction.

In addition to the awareness abilities initiated via ASK and TELL operators, we can envision additional awareness capability based on selfinitiation, where the reasoner initiates actions of its own based on state changes. For example, the system might decide to switch to an energysaving mode if the current state is related to insufficient energy supply.

One of the biggest challenges in this approach is how to map sensory raw data to KR symbols. An aware software-intensive system has sensors that connect it to the world and eventually help it listen to its internal components. These sensors generate raw data that represent the world's physical characteristics. The problem is that these low-level data streams must be converted to programming variables or more complex data structures that represent collections of sensory data, and those programing data structures must be labeled with KR symbols.

Another considerable challenge is how to express states and reason about them. KnowLang introduces an explicit STATES attribute that helps us specify concepts with a set of important states in which the concepts instances can be. Thus, we explicitly specify a variety of states for important concepts, for example, "operational" and "nonoperational" for a robot's motion system. Furthermore, a state in KnowLang is specified as a Boolean expression where we can use event activation, action execution, or property changes to build a state's Boolean expression. To facilitate complex state evaluation, the reasoner can use special predicates in which complex system states are evaluated as the product of other states.

he long-term impact of awareness-related research and development is a road map leading to artificial intelligence. Machine intelligence depends on the ability to perceive the environment and react to changes in it. The awareness mechanism uses raw data gathered via system's sensors to recognize objects, project situations, track changes, and learn new facts. A successful awareness mechanism can exhibit awareness at different levels of maturity and relevance. Noisy data can affect awareness relevance, which can lead to awareness results gradually changing over time and data input.

Ideally, the awareness mechanism should help intelligent systems behave like humans, realizing situations and reacting progressively: the first impression triggers a reaction that can change based on progressive realization of the current situation. *Emil Vassev* is a research fellow at Lero—the Irish Software Engineering Research Centre at the University of Limerick, Ireland. Contact him at emil. vassev@lero.ie.

Mike Hinchey is the director of Lero—the Irish Software Engineering Research Centre and a professor of software engineering at the University of Limerick, Ireland. Contact him at mike.hinchey@lero.ie.

CN Selected CS articles and columns are available for free at http://ComputingNow.computer.org.



IEEE@computer society Corporate Affiliate Program

Increases technical training while cutting costs.

Provides company-wide, employee access to 4,300 technical courses, 600 technical and business books, dozens of Brainbench Exams and free or discounted training webinars and software development certifications.

For more information, call 1-855-727-3632 or email us at cap@computer.org



Subscribe today for the latest in computational science and engineering research, news and analysis, CSE in education, and emerging technologies in the hard sciences.

www.computer.org/cise

IEEE computer society



IEEE @ computer society

IEEE Computer Society Membership— Focused on Your Future

Now when you join or renew your IEEE Computer Society membership, you can choose the membership package focused specifically on advancing your career:

- Software and Systems—includes IEEE Software Digital Edition
- Information and Communication Technologies (ICT)—includes IT Professional Digital Edition
- Security and Privacy—includes IEEE Security & Privacy Digital Edition
- Computer Engineering—includes IEEE Micro Digital Edition

In addition to your standard benefits, each package gives outstanding new benefits never offered before:

- A digital edition of the most requested leading publication specific to your interest
- A monthly digital newsletter developed EXCLUSIVELY for your focus area
- Your choice of three FREE webinars from the extensive IEEE Computer Society selection
- Downloads of 12 free articles of your choice from the IEEE Computer Society Digital Library (CSDL)
- Discounts on training courses specific to your focus area

Grow with IEEE Computer Society

IEEE Computer Society is expanding its services every day and we encourage you to take full advantage of all the opportunities that membership affords.

Check out the new "Welcome Webinar," designed to orient members to the many benefits of Computer Society membership. Here, you may discover new opportunities for participation, continuing education, or mentoring that appeal to you. You can watch the Welcome Webinar at www.computer.org/ membership.

In addition to your new benefits, Computer Society members still enjoy:

- Computer magazine in print and digital editions
- Our extensive eLearning Library made up of technical and management courses, videos, and books
- Six technical and community newsletters

- Discounts on publications, certifications, and conferences
- A computer.org email alias
- The opportunity to join technical committees
- Membership in local chapters
- Mentoring and networking

To choose one of the new, focused membership packages, go to www.computer.org/membership when you join or renew your membership. For just US\$123, you receive thousands of dollars of value and much more.

JOIN OR RENEW TODAY!

Select your membership package at www.computer.org/membership or call us at 1-800-272-6657.

Building a Virtual World: The Pipeline and Process

Brad Hallisey, ArenaNet



Game creation is a highly iterative and challenging process that involves three major disciplines—art, design, and engineering.

n 2007, ArenaNet began working on the sequel to its hit game, *Guild Wars*. More than five years later, the efforts of approximately 50 engineers, 75 artists, and 50 designers were realized when *Guild Wars 2* was released in August 2012.

This visually stunning game offers players the epic grandeur of a massive multiplayer online (MMO) playing environment combined with innovative combat mechanics, dynamic events, and customized personal storytelling. With a concurrent player base in the hundreds of thousands and a land mass of more than 60 square miles, it's a truly massive game.

Following its release, *Guild Wars 2* was met with critical acclaim, and more than 2 million copies were sold within two months.

THE PIPELINE

Game creation is a highly iterative process that involves three major disciplines—art, design, and engineering. Figure 1 shows the basic workflow and demonstrates that any stage can have feedback from the previous stages. The iterative process can be challenging, as it can result in invalidating existing work or changing technical requirements.

Genesis

Although Guild Wars 2 was created from the lore of the original Guild Wars, the design team still needed to clearly define the new content. For example, humans were the only playable species (referred to as "race") in the original game, but Guild Wars 2 supports five different playable races, and a team of writers and conceptual artists had to flesh them out. The design team creates the lore that encompasses game-play concepts and defines a realistic universe in which the game will take place. In game design, the lore and game-play concepts are the foundation upon which the game is built. However, it isn't uncommon for interesting art to feed back into and alter that process.

The decision to include five races required designing and building five different cities and starting areas. Concept art had to be created for each race, including the architecture and set pieces for each distinct area. For example, the Char is a warlike race that uses machinery, while the Sylvari is a more organic plant-based race. These lore choices drove not only the art but also the technology, as the techniques used to build a giant war machine are very different from those for building a tree city. Figure 2 shows a piece of concept art for the



Figure 1. Game creation workflow. Any stage in the iterative process can provide feedback to the other stages.

Char Sphere that exemplifies the Char race's strong characteristics.

Art assets

As the design team formalizes lore ideas and lays out concepts, it hands them off to level designers and artists. Level designers use MapEdit, the graphical editing program shown in Figure 3, to lay out a map's general physical characteristics. This gray-boxing process, often uses simple, solid-color (gray) models to fill in an area. Artists then begin creating textures and models that are representative of the concept art they're working from.

Next, the artists use Mapedit to add custom models to the map. This map editor lets artists populate the world rapidly and easily place, rotate, and scale models to build cities, caves, or other explorable areas. They can use a vegetation placement system to add vast forests or a simple farm. They can also design environmental effects such as lighting, haze, and weather. All this must come together to create the world within the game's larger fabric.

Content

When the world begins to mature and take shape, designers use a content editing tool like the one shown in Figure 4 to add game content atop the art. This includes enemy monsters, nonplayer characters that the player can interact with, and special bosstype monsters. There are also objects that the player can interact with, such as switches that open gates or mining nodes that offer crafting resources.

All of this dynamic content can change based upon predefined events and players' actions. In the persistent game worlds that MMOs usually contain, players' actions can impact other players, which can rapidly expand the possible scenarios. From an engineering standpoint, this can make the game developer's job challenging because the game engine must support these nearly infinite scenarios.



Figure 2. Char Sphere concept art, which exemplifies the Char race's strong characteristics.



Figure 3. MapEdit graphical editing program. Designers use solid-color models to fill in an area and lay out a map's general physical characteristics prior to a complete art pass.

Performance

Although performance is the last stage in the game creation workflow, it's an important consideration throughout the process because it's a significant limiting factor for the game's scope. The term *performance* covers many areas of technical concern, including software stability, memory usage, disk usage, frame rate, and patch sizes. Like most software development, good engineering balances performance and features. Additionally, for videogames, it's important to maintain a real-time simulation of 30 frames per second across a complex combination of hardware and workload.

This complex combination of processing power and demand is

ENTERTAINMENT COMPUTING

			Guil: Script Mirre = 1.0.0.4337 (Release)		
2 0	Go 2 Hord Tyre Quest Human C2 Q2 Human Cr	rous Map 1 Carriva.Mma.Script Mma			0
harrer Laurch ColTest X	Go2: Splat Hine X Go2: Locator Rover D	a X te			
		The start foliable of Societar Validate His	· D Additionalian D Additionalian		
Typer					
lare		E Script rune			
Cear Filter	Apply Filter	Description			
		(B) (T) Parameters			
80-0	1 Q1 Human Circus Hap 1 Logan				
±0	1 Q1 Human Circus Hap 2 Tent	Has Been Optimized			
	1 Q1 Human Circus Hap 3 Logan				
	Of Human Family Han 2 Seamon	SorptType	* Creature		
	Of Human Sister Han 1 Lonan	- Hinney	And the second se		
80.0	Q1 Human Sister Map 2 Tervela	C Interact	Bucches Creative Char Interact		
8000	Q1 Human Sister Hap 3 Centaur	10 Tripper Here Start Battle	Situation Tripper Set	Thipper Have Start Battle) get.	
8740	Q2 Human Circus Map 1 Carniva	E Hanual Hune Cover	Meruel		
800	Audio	E - Hanual Hune Pender	Manual		
805	Cinematics	Hanual Hirse Salute	Manual		
800	Landmarks	Hannal Hanne Shrug	Manual		
80-	Sectors	- Hanval Here Wave	Manual		
000	Anire Cambral Texasion Canada	- Quest Step Complete Byvestigate	Situation Quest Step Complete	[\$ Investigate Workshop] completes [Step].	
	Clean	- Quest step complete verent comes	promon Quest sith condition	to never caused contrains tough-	
BZA	Eogen Stave Counter Nime Progress Counter Nime Success Locaton Nime				
BZ	Logan Have Counter Here Progress Counter Here Progress Counter Here Success Counter Here Success Contact Here Here Codopt Here Codopt Here Here				
67.4	Login Hine Progress Counter Mine Progress Counter Mine Success Codest Mine Codest Mine Codest Mine Codest Mine Codest Mine Hisp Stroft Codest Mine	🕑 🍘 Mant Editable 🗢 Scriptber 🖌 Validate 🕥 Ad	lAction 📡 Add Branch 🔘 Add Loop		
67.4	Login Kine Propess	E 💓 Not Estable 🗢 Sciptfel 🖌 Veldere 🔘 Ad	lAtion ∀AddBanch ©Addloop		
87	Logan Flowe Counter Hine Propess Counter Hine Propess Loadation Hine Counter Loadation Hine Cadget Hine Sort Cadget Hine	E ● Not Situble → Solgibil → Veldee ○ Ad	LAction 🐨 Add Branch 💿 Add Leop Sects Trigger (Trigger Himse Listes Court)		
87.4	Logan titike Counter Hine Progress Counter Hine Success Locabat Hine Gadget Hine Hig Sopol Gadget Hine Sopol Gadget Hine Sopol Site Sopol Hine Counter Hine Counter Hine Counter Hine Counter Hine Counter Hine Counter Hine Sopol Hin	Phot Editable ← Solgitied → Velable → Ad Photod New Case Photod New Case Photod States Common Traper Net B) ← Addine Common Traper Net	l Action ∵ Add Banch © Add Leop Sech Trigger (Trigger Hone Lecter Cource) Sech Trigger (Trigger Hone Teal Reset)		
82.4	Logan Eliver Coulter Rine Propess Counter Rine Success Counter Rine Counter Rine Counter Rine Counter Rine Counter Rine Counter Rine Soper Counter Rine Soper Rine Soper Rine Counter Rine	PostSitabis + Sciptor Valdes Ad PostSitabis + Sciptor Valdes Ad PostSitabis Consens Tager Set D + Adas Consens Tager Set D + Adas Consens Tager Set D + Adas Consens Tager Set	Letters T Add Barech O Add Leep Sets Tetgeer (Tragger Hans Lates Cover) Sets Tetgeer Tragger Hans Lates (Cover) (Stills) ensuine (Cover).		
	Logan line Couter None Suppose Couter None Suppose Couter None Suppose Couter None Couter None Couter None Couter None Couter None Couter None Soper Couter None Soper Couter None Soper None Couter None Couter None Soper	Portalizada • Solgette Value Ao Portalizada Coure D	Lation 및 Add Isson ③ Add Loop Sets Trigger Floigger Hore Lation (Loop) Sets Trigger Floigger Hore Take Keyel (Filted constar (Court) () : 0		
BZ A	Logan Biner — Contre Response — Contre New Excess — Collection New — Colle	2: ■ tot totals ● toget ● toget ● toget 3:: ● total total Context ■ total Context	Letters T Add Baseds D Add Leop Sets Trapper Transport Hows Letters (Leoner) Sets Trapper Transport Hows Takes Kensell (Triffee) exercises Courses.		
	Logan Logan Security Inter Regard Control Nee Sutters Control Nee	Prestatusis * Sogether * Valdare @ Add Prestatusis * Sogether * Valdare @ Add Prestatusis Comments Prestatusis Comments Prestatusis Comments Prestatusis Comments Prestatusis Comments Prestatusis Comments Prestatusis	Hotton 'T Add Band O Add Leop Sets Topper Triager Hote Lates Cours? Sats Topper Topper Hote Lates Cours? (Stat Topper Topper Hote Lates Cours?) : 0 Topper Tragger Hote Lates Cours? In Set		_
	Logan Counter Mine Register Septimis Septimis Niger Mine Register	□ □ + trajector - ' vacator □ □ > trade there there □ - trade the there □ > trade there there □ - trade the there □ > trade the there □ - trade the there □ > trade the there □ - trade the there □ □ - trade the there □ □ □ - trade there □ □ □ - trade there □	Tables T Add David D Add Long Set Topper Player How takes Concel Set Topper Player How takes Add Set Topper Player How takes Add Topper Player How takes Concel is Set Topper Topper How takes Concel		
	Legal Courter from Program - Courter from Crosses - Courter from Crosses - Courter for Crosses - Cognetifications - Cogn	Plant States & Stoppher Vinitian () All Plant Alter Game Plant Alter Game Plant Alter Games Theorem Int Plant Alter Games Theorem Int Plant Alter Games Theorem Int Plant Pla	Falson T Additional O Additional facto Texper Friend Sector Sector Sector Sector facto Texper Friend Sector factor Texper Friend Sector Friend Sector Sector Friend Sector Friend Sector In Sec. Texper Friend Sector In Sec.		
	Leas Leas Control from Pages () Control from Pages () Control from Pages () Control from Science () Co	Bread and a single of	Hadnes "Q" Add Based "O Add Leop Sets Topological Files takes from the former Sets Topological Former Sets (1984) construction from the Sets (1984) construction from the Sets (1984) respectively for the Sets (1984) respectively for the Sets (1984) respectively for the Sets (1984) respectively for the Sets (1984) (1984) respe		
9 2 2 2 2 2 2 2 2 3 2 3 2 3 2 3 2 3 2 3	Lease Lease Country free Pages III - Country free States - Country free - Cou	Bertanse + broyder - valuer bertarier bertarier bertarier bertarier bertarier bertarier	Tables T Add Baseb D Add Leagt Tables Topop Floops Have tables Council tables Topop Floops Have tables Council Topop Have tables Council Topop Have tables Council table Topop Floops Have Council Topop Have tables Council table Topop Have tables Council table Topop Have tables Council table Topop Have tables Council tables Topop Have tables Council tables Tables		
	Local Control From Pages 1 — Counter from Pages 2 — Stage From Counter from Pages 2 — Stage	Bitchanis + Sought - Vename () Ad Bitchanis + Sought - Vename	Technin T Add Tained O Add Leop Technic Technic O Add Leop Technic Technic O Technic O Add Leop Technic O Technic		
S S S S S S S S S S S S S S S S S S S	Lease Lease Country free Pages and Country free States - Country free - Count	Statistica - Surgitar Vision (2) Statistica Care Stat	Antons ™ Ark Bonds © Anticipy Not Topoge Theore Theore Inter Annual Third Inter Topoge Theore Inter Annual Third Inter Theore Inter Annual Neuron Theore Endon Constants Inter Neuron Theore Endon Inter Annual Neuron Theorem Inter Annual Neuron Theorem Inter Annual Neuron Inter Annual Neuro	plestada	
	Course free houses — Course free houses — Course free forms — Course free forms — Course free forms — Course forms — Course forms — Course forms — Course forms — Course forms — Course forms — State forms — State forms — State forms — Course forms — Course — Course forms — Course forms	Statistics + Supplier Supplier Supplier	Hallow T And Branch D And Large Sets Topper Transport Network And Council In Strategy Transport Network And Council In Strategy Transport And Council In Set Topper Trans Lation Later Stores and State State And And And And And State State St	softer Janfier.	
	Course from suppose Course fr	Supersonal & Standard & Venderal () Add Supersonal Addression () Addression (Taken TARbons O Adding the type Phase Here is a down to type Phase Here is a down that operating Phase Phase Phase Here is a down the phase Phase Here is a down to the type Phase Here is a down to	alayinda	
	Constructions of the second se	Plant Standau or Sanghar Vindaum (2014) Plant Standau Career Plant Standau Career Plant Standau Career Theorem Theore	NAME: "P Add Books" () Add Lang be Stream Folgon: Heat Joint County Into Toury Folgon: Heat Joint County () Tours Tour Folgon: Heat Joint County () Tours Tours Heat Joint County Tours Heat Joint County () Tours Heat Join	nden kanden	
	Control free space Control free space Control free space Control free space Sp	Statistics + Stayford - Values () Ad Nord Way Gar Should be Gar Shoul	Note: The Additional To Addition Not Report Phases Home Lake Court The Additional Additional To Additional Theory of Phases Additional To Additional Theory of Phases Additional To Additional Report Phases Additional To Additional Report Phases Additional To Additional Report Phases Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To Additional To	phylada	
	Concerning and an	Statistics + Surght - Vision () All Statistics - Surght - Vision () - All Statistics - Surght - Vision - Surght - Su	Taken Takethow Database Takethow Take	adar kada.	
	Concerning and an analysis of the second sec	Part Estatis & Supplet IV Mark (2) Al Source Strategies (3) Source Strategies (NAMES 'P Add Books' D Add Lang be Transe Tologo: How Tologo How Tologo How Tologo Tologo How Tologo How Tologo (***********************************	adas Ladas	
	Control Technologies and a control Technologies	Bitchenie + Sopher I viewen () Ad Bitchenie Versene Tearre III Bitchenie Versene Tearre IIII Bitchenie Versene Tearre IIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIIII	Hates ▼Address © Address Intel Imper Frage Heat Links Court Network Progen Trans Links Court Network Provide Courts Network Provide Courts Netwo	ader la de	

Figure 4. Duo content editing tool showing its data editing mode. The tool has a graphical editing mode similar to Mapeditor for defining locations, regions, and paths that game scripts use.



Figure 5. *Guild Wars 2* options panel. A player can use the options to improve performance or alter visual details.

unique to MMO development. PC videogames must support the vast array of hardware that customers use, including CPUs and GPUs from several hardware generations. In addition, an MMO game has many scenarios with vastly different workloads. For example, open areas with long sight lines require the ability to hide details of objects based on their distance from the player. Conversely, smaller environments demand a high level of detail for pleasing visual appearance. Chaotic player combat requires constant streaming of different art assets as combatants enter and leave the fray.

Guild Wars 2 includes many subsystems that support this scaling, such as models that reduce geometric complexity, reduce the density of grass and other vegetation, drop objects such as floating dandelions, or eliminate the rendering of other players' models. Players can use the option panel shown in Figure 5 to make adjustments to the tradeoff between performance and visual detail.

BRINGING IT ALL TOGETHER

All games comprise many moving parts that often don't come together until a project's final stages. MMOs are among the largest and most complex games, with many parts that must come together at the end, a process that requires constant testing, tweaking, and more testing.

From the in-house quality assurance department to external invited testers, constant testing is always a focus of development. Reaching that final bar of creating a fun product is often the biggest challenge of all.

G ame programming is challenging, but the programmers really love their jobs. The end product is a visible piece of software that shows itself off. Having customers express enthusiasm for that final product is the ultimate reward.

This combination of challenge, visible product, and enthusiastic customers makes game programming an extremely satisfying line of work that is anchored in a multibillion-dollar industry. Given the project's size and scope, the challenges and rewards of participating in the development of *Guild Wars* 2 might have been a bit more extreme, but it was well worth the investment of time and effort over the course of five years.

Brad Hallisey is a senior software engineer at ArenaNet, where he works on Guild Wars game tools and technology. Contact him at brad@arena.net.

Editor: Kelvin Sung, Computing and Software Systems, University of Washington, Bothell; ksung@u.washington.edu

REPORT TO MEMBERS

Milojičić Voted 2013 Computer **Society** President-Elect



Members also choose new vice presidents and Board of Governors members.

EEE Computer Society members recently selected Dejan S. Milojičić of HP Labs to serve as president-elect for 2013. Milojičić, is currently the Society's first vice president, has served on its Board of Governors, and participated in developing the IEEE-CS 2011 Strategic Plan. He is the founding editor of the Society's Computing Now portal and an IEEE Internet Computing editorial board member. He was appointed as the first Special Technical Communities chair, and has served as chair of the IEEE-CS Technical Committee on Operating Systems and on many program committees, including ICDCS, CLOUD, and EDOC. An IEEE Fellow, Milojičić is also an IEEE-CS Golden Core member and an ACM Distinguished Engineer. He has published two books and more than 120 papers and has been awarded 11 patents and 22 patent applications.

Candidates elected to the Computer Society presidency serve in leadership roles for a three-year term. After serving a year as presidentelect under 2013 president David Alan Grier, Milojičić will assume the duties of Society president in 2014. Following his 2014 term as president, he will continue to be an active Society leader in 2015 as past president.



2013 IEEE Computer Society president-elect Dejan S. Milojičić will focus on aligning the CS with the technology transformation reflected in cloud computing and social networking.



2013 IEEE Computer Society president David Alan Grier is exploring ways to strengthen support for conferences, intellectual property rights management, technical education, and plans for the society's future.

LEADERS SERVE MEMBERS

Each year, Society members vote for the next year's president-elect, first and second vice presidents, and seven members of the IEEE Computer Society Board of Governors. The Society president and vice presidents each serve a one-year active term, while the 21 Board of Governors members serve three-year terms, rotating in three groups of seven.

The three presidents-presidentelect, current president, and past president—work together in setting Society policies and making operational decisions. The current Society president is responsible for heading three annual Board of Governors meetings and for addressing major issues that affect the Computer Society during the year.

NEW VICE PRESIDENTS ELECTED

Thomas M. Conte of the Georgia Institute of Technology was elected 2013 first vice president, while Elizabeth (Liz) Burd of Durham University in the UK topped the balloting for 2013 second vice president. Conte will serve as chair of the Publications Board; Burd will chair the Member and Geographic Activities Board.

The sitting president also appoints vice presidents to complement the two elected VPs as leaders of individual Society activities boards: Publications, Educational Activities, Standards Activities, Technical and Conference Activities, and Professional Activities.

REPORT TO MEMBERS

Table 1. New officers who will begin serving the IEEE Computer Society on 1 January 2013.

Office	Officer	Number of votes	Percent
2013 president-elect/2014 president	Dejan S. Milojičić	3,831	57.3
2013 first vice president	Thomas M. Conte	4,597	69.5
2013 second vice president	Elizabeth (Liz) Burd	5,089	76.9
2013-2015 Board of Governors	Ann DeMarle	4,623	69.9
	Cecilia Metra	4,315	65.2
	Nita Patel	3,377	51.0
	Diomidis Spinellis	3,250	49.1
	Phillip Laplante	3,156	47.7
	Jean-Luc Gaudiot	3,093	46.7
	Stefano Zanero	3,074	46.4

Table 2. The full ballot for the 2012 Computer Society election also included the following candidates.

Office	Officer	Number of votes	Percent
2013 president-elect/2014 president	Don Shafer	2,830	42.3
2013 first vice president	Paul R. Croll	1,997	30.2
2013 second vice president	Paul K. Joannou	1,517	22.9
2013-2015 Board of Governors	Xun Luo	2,767	41.8
	Harold Javid	2,752	41.6
	William (Bill) Pitts	2,564	38.7
	Ronald Jensen	2,444	36.8
	Simon Y. Liu	2,344	35.4
	Sattupathu V. Sankaran	2,037	30.8

COMPUTING Then

Learn about computing history and the people who shaped it.

http://computingnow. computer.org/ct

The appointed Society vice presidents also serve as nonvoting members of the Board of Governors. Holding voting positions on the Board are the president, past president, president-elect, and the first and second vice presidents. Additional nonvoting members include the Society's staff executive director and the IEEE directors for Divisions V and VIII—the Computer Society's elected representatives on the IEEE Board of Directors.

BOARD OF GOVERNORS ADDS SEVEN NEW MEMBERS

In the 2012 Society election, which closed in early October, voters also cast ballots to fill seven openings on the IEEE Computer Society Board of Governors. The seven members chosen for 2013-2015 terms are Ann DeMarle, Jean-Luc Gaudiot, Phillip Laplante, Cecilia Metra, Nita Patel, Diomidis Spinellis, and Stefano Zanero. Many of the successful candidates have had recent Computer Society Board of Governors and leadership experience.

Elected officers volunteer their time and talents to further the Society's goals and to elevate the profile of the computing profession in general. Society officers take a lead role in promoting new publications, educational efforts, technical focus groups, and international standards that help Society members attain career goals.

Table 1 shows the breakdown of votes cast for each office. The full ballot for the 2012 election also included the candidates listed in Table 2.

NOMINATE A CANDIDATE

Any member can nominate candidates for Society offices. Most members are also eligible to run for a seat on the Board of Governors. Candidates for other offices must be full members of IEEE and must have been Society members for at least the preceding three years.

Visit www.computer.org/portal/ web/election for more details on the 2012 IEEE Computer Society election.

J. ROBERTO B. DE MARCA CHOSEN AS IEEE PRESIDENT-ELECT FOR 2013

EEE members recently selected J. Roberto B. de Marca as their 2013 president-elect. De Marca was a Fulbright Scholar at the University of Southern California where he received a PhD in electrical engineering. Since 1978, he has been on the faculty of the Catholic University, Rio de Janeiro, where he also held several administrative positions, including associate academic vice president. As National Research Council Scientific Director, de Marco authorized the start-up money for the national research network, which paved the way to the widespread Internet use in Brazil. He was the Brazilian Telecommunications Society founding president and is a member of the Brazilian National Academy of Sciences and the National Academy of Engineering.

As president-elect, De Marco will succeed Peter W. Staecker, AMP M/A-COM (retired), Massachusetts, who will become IEEE president in 2013. De Marco will serve one year as IEEE president-elect, participating in Board of Directors activities. He will then assume the roles of president in 2014 and past president in 2015.

In the same election, IEEE members

chose Computer Society Board of Governors member Susan K. (Kathy) Land, CSDP, as Division V director-elect for 2013.

Division directors serve on the IEEE Board of Directors and Technical Activities Board. Division directors V and VIII are elected to represent the Computer Society membership. Land will act as director-elect in 2013 and as division director for 2014-2015. The division directors also serve as ex officio members of the Computer Society's Board of Governors and Executive Committee.

Showcase Your Multimedia Content on Computing Now!

IEEE Computer Graphics and Applications seeks computer graphics-related multimedia content (videos, animations, simulations, podcasts, and so on) to feature on its Computing Now page, www.computer.org/portal/web/ computingnow/cga.

If you're interested, contact us at **cga@computer.org**. All content will be reviewed for relevance and quality.





CS Announces 2012 Winners of Kennedy and Rau Awards



he IEEE Computer Society recently named the 2012 recipients of two technical awards. Mary Lou Soffa, Owen R. Cheatham Professor at the University of Virginia, received the ACM-IEEE Computer Society Ken Kennedy Award for outstanding contributions in highperformance computing together with significant community service or mentoring. Joseph A. (Josh) Fisher, Hewlett-Packard Senior Fellow (Emeritus), received the B. Ramakrishnan Rau Award for significant accomplishments in microarchitecture and compiler code generation.

MARY LOU SOFFA



Mary Lou Soffa, recipient of the 2012 Ken Kennedy Award.

Mary Lou Soffa received the Kennedy Award at SC12, held in Salt Lake City in November 2012, for contributions to compiler technology and software engineering, exemplary service to the profession, and lifelong dedication to mentoring and improving diversity in computing. A leading researcher in programming languages, Soffa has developed software tools for debugging and testing programs as well as modelbased strategies for optimizing compilers to produce higher-quality code.

Soffa was elected an ACM Fellow in 1999, and received the Presidential Award for Excellence in Science, Mathematics, and Engineering Mentoring from the White House the same year. In 2006, she received the Computing Research Association (CRA) Nico Habermann Award for contributions toward increasing the numbers and successes of underrepresented members in the computing research community. She has held leadership roles in prominent national and international organizations, among them CRA and its Committee on the Status of Women in Computer Science and Engineering.

ACM and the Computer Society cosponsor the Kennedy Award, which was established in 2009 to recognize substantial contributions to programmability and productivity in high-performance computing (HPC) and significant community service or mentoring contributions. It was named for the late Ken Kennedy, founder of Rice University's computer science program and a world expert on HPC. The Kennedy Award carries a \$5,000 honorarium.

JOSEPH A. (JOSH) FISHER



Josh Fisher, recipient of the 2012 B. Ramakrishnan Ray Award.

Josh Fisher received the Rau Award for the development of trace scheduling compilation and pioneering work in very long instruction word (VLIW) architectures. He devised the trace scheduling compiler algorithm, an optimization technique for compilers, and coined the term instruction-level parallelism (ILP) during his undergraduate and graduate studies at the Courant Institute of Mathematical Sciences at New York University.

Fisher created and named VLIW architectures and invented many fundamental ILP technologies while a professor at Yale University. In 1984, he started Multiflow Computer with two members of his Yale team. In 1984, he received the ACM/IEEE Computer Society Eckert-Mauchly Award. Fisher joined HP Labs in 1990, where he worked with ILP and custom-embedded VLIW processors

and their compilers before retiring in 2006. At HP, Fisher started and managed the HP Labs Cambridge research facility.

The IEEE Computer Society established the Rau Award in 2010 in memory of the late Bob Rau, also an HP Senior Fellow. Rau, who passed away in 2002, managed HP Labs' Compiler and Architecture Research group. He started HP Labs' research program in VLIW and ILP processing when he joined the facility in 1989, resulting in the development of the Explicitly Parallel Instruction Computing (EPIC) style of architecture that is the basis for the IA-64. The award comes with a \$2,000 honorarium.

Fisher received the Rau Award at the 45th Annual ACM/IEEE International Symposium on Microarchitecture (MICRO-45) in December in Vancouver.

CHANGES TO SOCIETY BYLAWS AVAILABLE ONLINE

A t its November 2012 meeting, the IEEE Computer Society Board of Governors approved the first reading of amendments to two sections of the Society's bylaws.

The amendment to Article IX, Section 5, eliminates the editor in chief position for the Computer Society Press, which will cease to exist after 31 December 2012. The amendment to Article XII, Section 11, describes changes to the role and responsibilities of the Electronic Products and Services Committee.

Changes to existing Society bylaws that receive approval by the Board of Governors are listed by title in *Computer*, with links to website locations hosting the actual documents. The documents will be accessible at the website location until such time as the changes receive final approval.

The documents for the currently proposed amendment are posted at the following URLs:

Bylaws Article IX—Publications; Section 5—Editor in Chief Appointments and Terms http://bit.ly/TYjcM3

Bylaws Article XII—Standing Committees; Section 11— Electronic Products and Services Committee http://bit.ly/SZD2po

Deletions are marked in strikeout text. Only relevant segments of the bylaws in question are reproduced.

Members can send comments to Anne Marie Kelly, amkelly@ computer.org, by close of business, 4 January 2013.

MARIO R. BARBACCI (1945–2012)

Mario R Barbacci

1996 Computer Society president, helped establish

the Software Engineering

Institute at Carnegie

Mellon University.

Mario R. Barbacci, 1996 president of the IEEE Computer Society, died at his home on 23 October 2012.

Barbacci left Peru in 1969 and became Carnegie Mellon University's first PhD student in computer science from his country in 1974. By the time he left CMU 30 years later, he had helped found the school's prestigious Software Engineering Institute.

Barbacci was a research scientist in CMU's Department of Computer Science from 1969 to 1985. In 1984, he was part of a small team that wrote a proposal for a new federal contract to study advances in software engineering and how

to use them to benefit the US Department of Defense and private industry. The team's proposal resulted in the creation of SEI, which

now has more than 500 employees operating under a federal contract worth more than half a billion dollars. Before his retirement in 2004, Barbacci held a variety of SEL positions, including associate director

variety of SEI positions, including associate director, project leader, program director, and senior member of technical staff. He was known as an innovative analyst of how computers interact with one another.

"He helped create a language that described computers and how they were built," and he did it with a quiet, sociable manner that provided many humorous insights over the years, said Charles Weinstock, a longtime friend and senior member of SEI's technical staff.

Barbacci is survived by his partner Josephine Olsen, one brother in Lima, and two sisters in the US.



SELECTED CS ARTICLES AND COLUMNS ARE ALSO AVAILABLE FOR FREE AT COMPUTING NOW

http://ComputingNow.computer.org

Thanks to Volunteers

he IEEE Computer Society thanks the following associate editors and editorial board members who are retiring at the end of this year for giving their valuable time and support to our publications.

- Computer Tom Conte Robert France Steven Reinhardt
- Computer Architecture Letters Angelos Bilas Bruce Jacob James Larus Ravi Nair Kevin Skadron Yuanyuan Zhou
- Computing in Science & Engineering Klaus-Jurgen Bathe Rachel Kuske David Winch
- IEEE Intelligent Systems Hsinchun Chen

IEEE Micro Charles Moore, deceased Kevin Skadron

IEEE Software Art Sedighi

IT Professional Chien-Yeh Hsu George Kraft Bruce Potter

IEEE Transactions on Affective Computing Kristina Hook Ana Paiva Brian Parkinson Helmut Prendinger Marilyn Walker

IEEE Transactions on Computational Biology and Bioinformatics Dannie Durand Tao Jiang Bertram Ludäscher Joerg Stelling

IEEE Transactions on Computers Elisardo Antelo Cristiana Bolchini George Constantinides Kanad Ghose Dimitris Gizopoulos Victor C.M. Leung John Chi-Shing Lui Radu Marculescu Patrick McDaniel Paolo Montuschi Walid Najjar Sotiris Nikoletseas Stephan Olariu Sanjay Ranka Sandeep Shukla Sang Hyuk Son Spyros Tragoudas Yuanyuan Yang

IEEE Transactions on Dependable and Secure Computing Shiuhpyng Shieh Frank Stajano Neeraj Suri

IEEE Transactions on Haptics Federico Barbagli Cagatay Basdogan Brent Gillespie Danny Grant Matthias Harders Vincent Hayward Hiroyuki Kajimoto Astrid (A.M.L.) Kappers Karon Maclean Jeha Ryu Hong Tan

IEEE Transactions on Knowledge and Data Engineering Yixin Chen Juergen Dix Sergio Greco Eamonn Keogh Vipin Kumar Jianzhong Li Qing Li Yannis Papakonstantinou V.S. Subrahmanian Zhi-Hua Zhou Xingquan (Hill) Zhu

IEEE Transactions on Learning Technologies Pierre Dillenbourg

IEEE Transactions on Mobile Computing Amotz Bar-Noy Sajal Das Richard Han Ekram Hossain Mary Ann Ingram Yunhao Liu Prasant Mohapatra Chiara Petrioli Tajana Simunic Rosing Paolo Santi Terry Todd Wade Trappe Nalini Venkatasubramanian

IEEE Transactions on Parallel and Distributed Systems Ananth Grama Rabi Mahapatra Albert Zomaya

IEEE Transactions on Pattern Analysis and Machine Intelligence Gregory Hager Marina Meila Nikos Paragios Vladimir Pavlovic Bernt Schiele Kaloom Siddiqi Antonio Torralba Daphna Weinshall

IEEE Transactions on Software Engineering Elisabetta Di Nitto Paul Strooper Tetsu Tamai

IEEE Transactions on Visualization and Computer Graphics Kavita Bala Sheelagh Carpendale Gerik Scheuermann Wenping Wang

CALL AND CALENDAR

CALLS FOR ARTICLES FOR COMPUTER

Computer seeks submissions for a June 2013 special issue on big data, exploring aspects of discovery, productivity, and policy, with a focus on their socioethical implications.

This special issue will address how deeply personal data will likely be scrutinized in the context of datadriven decision making, such as how often and with whom we communicate using our cell phones or other digital devices; our location; what products and services we buy; where we eat, sleep, and work every day; and our photos and videos. In a way, we will be trading our privacy for a new kind of "surveillance" based on the premise of customer care.

The guest editors seek varied perspectives regarding the challenges, possibilities, and benefits of big data. They particularly welcome articles that feature qualitative assessments, case studies from government agencies, perspectives from Internet search companies and other hardware and software vendors, predictive studies demonstrating paradigm shifts, and social impact research.

Articles should appeal to a broad interdisciplinary audience or policy professionals in the information and communications technology sector. The writing should be original, avoiding long discussions about theories, theorems, algorithms, or mathematical notations.

All manuscripts are subject to peer review on both technical merit and

SUBMISSION INSTRUCTIONS

The Call and Calendar section lists conferences, symposia, and workshops that the IEEE Computer Society sponsors or cooperates in presenting.

Visit www.computer.org/conferences for instructions on how to submit conference or call listings as well as a more complete listing of upcoming computingrelated conferences.



relevance to *Computer*'s readership. Accepted papers will be professionally edited for content and style.

The guest editors for this special issue are Katina Michael (katina@ uow.edu.au), an associate professor on the Faculty of Informatics at the University of Wollongong, Australia, and Keith Miller (miller.keith@uis. edu), a professor in the Department of Computer Science at the University of Illinois Springfield.

Paper submissions are due by **15 January 2013**. Please email the guest editors a brief description of the article you plan to submit by 15 December 2012. Visit www.computer. org/computer/cfp6 to view the complete call for papers.

Computer seeks submissions for an October 2013 special issue on multi-core memory coherence.

As we enter an era of large multicore systems, the question of efficiently supporting a shared memory model has become more important. Massively parallel architectures lacking coherent shared memory have enjoyed great success in niche applications such as 3D rendering, but general programming developers still demand the convenience of a shared memory abstraction.

Efficiently using the Message Passing Interface requires that individual computation tasks must be relatively large to overcome communication latencies. It becomes difficult to use the MPI at the fine-grained level when fast on-chip communication is available. Higher-level mechanisms like MapReduce or shard-based databases are popular in particular application domains, though researchers have not yet efficiently applied them at the chip/node level.

This special issue will focus on approaches to providing scalable, shared on-chip memory, paramount in a future where individual nodes will have on the order of 1,000 cores each.

Suggested topics include but are not limited to private and shared cache hierarchies; scalable memory coherence protocols, directorybased and otherwise; data layout and placement techniques; on-chip interconnects to support shared-memory abstractions; and hardware, software, and hybrid approaches.

Articles are due by **1 March 2013**. Visit www.computer.org/computer/ cfp10 to view the complete call for papers..

CALLS FOR ARTICLES FOR OTHER IEEE CS PUBLICATIONS

IEEE IEEE Internet Computing plans a September/October 2013 special issue on dynamic collective work.

As the Internet has changed the way in which data circulates, we have shifted from a world of paper documents to one of online documents, databases, and provenance

CALL AND CALENDAR

EVENTS IN 2012 & 2013

January 2013

15-18.	•	• •	•	•	• •	•	• •	•	•	•	• •	•	•	•	•	•	WACV 2013
28-30	•			•	•••	•			•	•		•	•	•	•	•	ICOIN 2013
28-31.																	. ICNC 2013

February 2013

23-27 HPCA 201	3
----------------	---

March 2013

6-8	ISADS 2013
16-23	VR 2013
25-28	AINA 2013
25-28	SOSE 2013

April 2013

1-5 LAD	C 2013
8-11 ST	C 2013
8-12 ICD	E 2013
15-18SYSCOI	N 2013
29 Apr-1 May NSV	V 2013

systems. This has also increased the size and complexity of systems that support today's globally distributed, rapidly changing, and agile collaborative enterprises. Such systems are becoming increasingly federated and are generating a huge amount of data at different granularity levels that include tweets, blog posts, instant messages, Facebook updates, and other social media content. These systems and data are fueling explosive growth in dynamic collective work in the healthcare, insurance, banking, and other industries.

Dynamic and collective activities are characterized by their flexibility and people-driven nature. Automobile insurance claims handling, order processing of prescription drugs, hospital patient case management, and recovery and response assistance during natural disasters are just a few examples. In these and other tasks, various factors determine the set of actions that must be performed and the order in which they're executed, including human judgment and document contents. The guest editors seek original articles describing research efforts and experiences concerning Internetsupported dynamic collective work.

Articles are due **4 January 2013**. Visit www.computer.org/portal/web/ computingnow/iccfp5 to view the complete call for papers.

IEEE Micro plans a July/August 2013 special issue on reliability.

Over the past decade, designers have sought better ways to balance power, performance, and cost. Of these, power has emerged as a firstorder design challenge. In the coming era, this challenge may be subsumed by that of building robust and reliable systems. As technology advances, systems are becoming increasingly susceptible to transient errors such as timing violations, parameter variations, and aging. Without innovations in the areas of microprocessor and software reliability, future systems may face continuous failure. Thus, we need new computing paradigms that incorporate adaptive techniques at both the hardware and software layers to ensure resilient execution. The system, as a whole, must dynamically detect and recover from errors to meet historically established high reliability standards without exceeding power budgets and cost constraints or violating performance targets.

This guest editors seek original articles on all topics related to reliability that span the layers in the system stack, from device, circuit, and architecture design to the role of software in enabling robust and reliable computing.

Articles are due **8 January 2013**. Visit http://www.computer.org/portal/ web/computingnow/micfp4 to view the full call for papers.

IEEE Internet Computing plans a November/December 2013 special issue on smart cities.

Smart cities are currently the focus of a broad research community as well as of many government

ICNC 2013

he International Conference on Computing, Networking and Communication is technically cosponsored by the IEEE Computer Society and the IEEE Communications Society. ICNC 2013 will address areas such as cloud computing and networking; cognitive computing and networking; multimedia computing and communications; green computing, communications, and networking; mobile computing and vehicle communications; Internet services and applications; optical and grid networking; wireless ad hoc and sensor networks; wireless communications and networks; and communications and information security.

ICNC 2013 will take place 28-31 January 2013 in San Diego, California. Visit www.conf-icnc.org/2013 for complete conference information.

and industry innovation agendas. The Internet plays a fundamental role in communication, information sharing and processing, data transfer and analysis, and distributed computing in many of today's cities. The rise of the Internet of Things and the large-scale adoption of Web technologies in urban environments have proved that Internet-based solutions can successfully address smart cities' multifaceted, cross-domain challenges.

The guest editors seek submissions describing recent or ongoing research efforts and experiences in applying Internet technologies to realize the smart city vision.

Email the guest editors (ic6-2013@ computer.org) a brief description of the article you plan to submit by **15 February 2013**. Articles are due **1 March 2013**. Visit www.computer. org/portal/web/computingnow/iccfp6 to view the complete call for papers.

IEEE Intelligent Systems plans a January/February 2014 special issue on the Web of Things.

The new concept of the Web of Things encompasses the ability of

mobile devices and sensors to observe and monitor their environments, increasing the coordination between things in the real world and their counterparts on the Web. This Web of Things is expected to produce large volumes of data related to the physical world, and intelligent solutions are required to enable connectivity and internetworking and ensure relevance between the physical world and the corresponding digital world resources.

The guest editors seek innovative contributions to intelligent system and interaction design, information processing and knowledge engineering, and adaptive solutions to assist in efficient utilization of the Web of Things. Submissions that include audio, video, and community content are encouraged.

Articles are due **1 March 2013**. Visit http://www.computer.org/portal/ web/computingnow/iscfp1 to view the complete call for papers.

IEEE Software plans a March/April 2014 special issue on next-generation mobile computing.

We use mobile computing not only when we interact with our smartphones to connect with friends and family, but also when we use ticketing systems on a bus or train, purchase food from a mobile vendor at a park, and watch videos or listen to music on our phones and portable devices.

Any computation system expected to move and interact with end users or other computational systems despite potential network changes such as loss of connectivity represents an aspect of mobile computing. The number of such systems is expected to grow significantly each year over the coming decades. Mobile technology also is expected to evolve, creating new challenges.

IEEE Software seeks articles that explore the next generation of mobile computing within the contexts of

mission-critical scenarios, quality-ofservice differentiation, and resource constraints.

Articles are due **30 June 2013**. Visit http://www.computer.org/portal/ web/computingnow/swcfp2 to view the complete call for papers.

CALENDAR

JANUARY 2013

15-18 Jan: WACV 2013, IEEE Workshop on Applications of Computer Vision, Clearwater Beach, Florida; http://cvl. cse.sc.edu/wacv2013/index.html

28-30 Jan: ICOIN 2013, 2013 Int'l Conf. on Information Networking, Bangkok, Thailand; www.icoin.org

28-31 Jan: ICNC 2013, Int'l Conf. on Computing, Networking and Communications, San Diego; www. conf-icnc.org/2013

FEBRUARY 2013

23-27 Feb: HPCA 2013, IEEE Int'l Symp. on High-Performance Computer Architecture, Shenzhen, China; www.cs.utah.edu/~lizhang/HPCA19/ index.html

MARCH 2013

6-8 Mar: ISADS 2013, IEEE Int'l Symp. on High-Performance Computer Architecture, Mexico City, Mexico; www.isadsmexico2013.mx/isads

16-23 Mar: VR 2013, IEEE Virtual

Reality 2013, Orlando, Florida; http:// ieeevr.org/2013

25-28 Mar: AINA 2013, 27th IEEE Int'l Conf. on Advanced Information Networking and Applications, Barcelona, Spain; http://voyager.ce.fit.ac.jp/conf/ aina/2013//index.html

25-28 Mar: SOSE 2013: 7th Int'l Symp. on Service Oriented System Engineering, San Francisco; http:// sei.pku.edu.cn/conference/sose2013/ index.htm

APRIL 2013

1-5 Apr: LADC 2013, 6th Latin-American Symp. on Dependable Computing, Rio de Janeiro, Brazil; www.ft.unicamp.br/ladc2013/

8-11 Apr: STC 2013, 25th Annual Software Technology Conf., Salt Lake City, Utah; http://sstc-online.org/home. cfm?pg=hm

8-12 Apr: ICDE 2013, 29th IEEE Int'l Conf. on Data Engineering, Brisbane, Australia; www.icde2013.org/index. html

15-18 Apr: SYSCON 2013, IEEE Int'l Systems Conf., Orlando, Florida; http://ieeesyscon.org/

29 Apr-1 May: NSW 2013, IEEE 2nd Intl Workshop on Network Science, West Point, New York; http://ieee-nsw. org

Calls for Papers



IEEE Micro seeks general-interest submissions for publication in upcoming issues. These works should discuss the design, performance, or application of microcomputer and microprocessor systems. Of special interest are articles on performance evaluation and workload characterization. Summaries of work in progress and descriptions of recently completed works are most welcome, as are tutorials. *IEEE Micro* does not accept previously published material.

www.computer.org/micro/cfp



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.

THE UNIVERSITY OF KENTUCKY, Engineer Associate III/Research SM534701, Center for Visualization. Engineer Associate III/Research needed full time for development of hardware system and software modules related to computer graphics and image processing. Requires Masters and 12 months experience in programming in C++C/Java, image processing and computer graphics (Open GL/Direct X programming), autostereoscopic display design and implementation, multi-projector displays and its automatic calibration, optic design and web page authoring and management. To apply for job # SM534701submit a UK Online Application at www.uky. edu/ukjobs. If you have any questions,

contact HR/Employment, phone (859) 257-9555 press 2, or email ukjobs@email. uky.edu. Application deadline is January, 1 2012. Any candidate offered this position may be required to pass preemployment screenings as mandated by University of Kentucky Human Resources. These screenings may include a national background check and/or drug screen. The University of Kentucky is an equal opportunity employer and encourages applications from minorities and women.

UNIVERSITY OF NORTHERN IOWA, Assistant Professor of Computer Sci-

ence. The Department of Computer Science at the University of Northern Iowa invites applications for a tenure-track assistant professor position to begin August 2013. Applicants must hold a Ph.D. in Computer Science or a closely-related discipline. The department seeks candidates able to participate widely in the CS curriculum and conduct a research program involving undergraduates. Detailed information about the position and the department are available at http://www.cs.uni.edu/ To apply, visit http://jobs.uni.edu/. Applications received by January 15, 2013, will be given

UNIVERSITY OF MARYLAND



A. JAMES CLARK SCHOOL OF ENGINEERING DEPARTMENT OF ELECTRICAL & COMPUTER ENGINEERING

TENURE-TRACK AND TENURED FACULTY POSITIONS IN CYBERSECURITY

The Department of Electrical and Computer Engineering at the University of Maryland, in collaboration with the Maryland Cybersecurity Center (cyber. umd.edu), seeks exceptionally qualified candidates for tenure-track and tenured faculty positions to begin in August 2013 in the area of: Cyber Security. This includes all aspects of security-oriented research in computer engineering, communications, signal processing, and networking, at the hardware, software, protocol, algorithm, system, and physical layer levels.

Appointments at all ranks will be considered. Applicants should have received or expect to receive their PhD in Electrical Engineering or a related discipline prior to August 2013. Candidates for the rank of Assistant Professor should be creative and adaptable, and should have a high potential for both teaching and research. Candidates for the ranks of Associate and Full Professor should have distinguished records in research and a strong interest in educational programs.

For best consideration, applications should be submitted by January 21, 2013 to https://jobs.umd.edu (position number 105043). Applications should include a cover letter, curriculum vitae with list of publications, research and teaching statements, and the names and contact information of at least four references.

The University of Maryland is an equal opportunity, affirmative action employer with a strong commitment to the principle of diversity. Applications from minority groups and women are especially invited. full consideration. EOE/AA. Pre-employment background checks are required. UNI is a smoke-free campus.

UNIVERSITY OF MICHIGAN-DEAR-**BORN, Assistant/Associate Professor** in Software Engineering. The Department of Computer and Information Science (CIS) at the University of Michigan-Dearborn invites applications for a tenure-track faculty position in software engineering. Rank and salary will be commensurate with gualifications and experience. We offer competitive salaries and start-up packages. Qualified candidates must have, or expect to have, a Ph.D. in computer science or a closely related discipline by the time of appointment and will be expected to do scholarly and sponsored research, as well as teaching at both the undergraduate and graduate levels. Candidates at the associate professor rank should already have an established funded research program. The CIS Department offers several BS and MS degrees, and participates in several interdisciplinary degree programs, including an MS program in software engineering and a Ph.D. program in information systems engineering. The current research areas in the department include computer graphics and geometric modeling, database systems, multimedia systems and gaming, networking, computer and network security, and software engineering. These areas of research are supported by several established labs and many of these areas are currently funded. The University of Michigan-Dearborn is located in the southeastern Michigan area and offers excellent opportunities for faculty collaboration with many industries. We are one of three campuses forming the University of Michigan system and are a comprehensive university with over 8900 students. One of the university's strategic visions is to advance the future of manufacturing in a global environment. The University of Michigan-Dearborn is dedicated to the goal of building a culturally-diverse and pluralistic faculty committed to teaching and working in a multicultural environment, and strongly encourages applications from minorities and women. A cover letter, curriculum vitae including e-mail address, teaching statement, research statement, and three letters of reference should be sent to, Dr. William Grosky, Chair Department of Computer and Information Science University of Michigan-Dearborn 4901 Evergreen Road Dearborn, MI 48128-1491 Email: wgrosky@umich.edu, Internet: http://www.cis.umd.umich.edu Phone: 313.583.6424, Fax: 313.593.4256 The University of Michigan-Dearborn is an equal opportunity/affirmative action employer.

Mountain View, CA; Palo Alto, CA; and Sunnyvale, CA

Design Verification/Validation Engineers: Responsible for ensuring the quality of Microsoft hardware products. http://bit.ly/MSJobs-Hardware

Hardware Dev. or Design Engineers, Hardware Engineers, and Design Engineers: Design, implement, and test computer hardware. http://bit.ly/MSJobs-Hardware

Researchers/Scientists: Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms, and prototypes. http://bit.ly/ MSJobs-Research

Service Engineers, Service Operations Engineers, and Systems/Operations Engineers: Plan, architect, deploy and/or support complex client/server or database software systems. http://bit.ly/MSJobs-SysOps

User Experience Researchers/Designers and User Research Engineers: Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. http://bit.ly/MSJobs-UX

San Francisco, CA

User Experience Researchers/Designers and User Research Engineers: Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. http://bit.ly/MSJobs-UX

Ft. Lauderdale, FL

Evangelist: Secure future growth of the Microsoft platform by engaging a community of customers, partners, and academics to embrace and adopt Microsoft technology. https://bitly.com/MSJobs-OtherTech

Operations Program Managers: Responsible for the design, implementation, and release of programs or projects. http://bit.ly/MSJobs-ProgMgr

Solutions Sales Professional/Specialist: Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. http://bit.ly/MSJobs-SalesEng

Alpharetta, GA

Technology Solutions Professionals / CATM Specialists: Drive product win rates by proving the value of product(s) to customers and partners. http://bit.ly/ MSJobs-SalesEng

Chicago, IL and Downers Grove, IL

Premier Field Engineer: Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Requires from 50-75% travel throughout the U.S. http://www.jobs-microsoft.com/job/go/2249316/

Technical Account Managers: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. http://bit.ly/MSJobs-Support

Charlotte, NC

Premier Field Engineers: Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. http://bit.ly/MSJobs-Support

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs-Support

Consultant: Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Requires from 50-75% travel throughout the U.S. http://www.jobs-microsoft.com/job/go/2249313/

Fargo, ND

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs-Support

Technical Account Manager: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires up to 25% travel throughout the U.S. http://www.jobs-microsoft.com/job/go/2249314/

<u>Reno, NV</u>

Operations Program Managers: Responsible for the design, implementation, and release of programs or projects. http://bit.ly/MSJobs-ProgMgr

New York, NY

Account Managers and Directors: Develop business opportunities for sales of software or services. http://bit.ly/ MSJobs-NonTech

Global Account Technology Strategist: Provide presales technical and architectural support for sales of software, solutions, and related products. Requires travel to various unanticipated locations up to 35% of the time. http://www.jobs-microsoft.com/job/go/2229102/

Solution Sales Specialist: Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. Requires travel to various unanticipated locations up to 25% of the time. http://www.jobs-microsoft.com/ job/go/2215933/

Account Technology Strategist: Provide pre-sales technical/architectural support. http://www.jobs-microsoft. com/job/go/2219038/

Pittsburgh, PA

Account Managers and Directors: Develop business opportunities for sales of software or services. http://bit.ly/ MSJobs-NonTech

Irving, TX

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs-Support

Technical Account Managers: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. http://bit.ly/MSJobs-Support

Technical Account Manager: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires travel throughout South Central U.S. up to 25%. http://www.jobs-microsoft.com/job/go/2249315/

Redmond, WA

Account Managers and Directors: Develop business opportunities for sales of software or services. http://bit.ly/ MSJobs-NonTech

Architects: Manage the sales, discovery, and design phases of computer software deployments. http://bit.ly/ MSJobs-Architects

Artists, Art Leads and Animators: Responsible for designing and creating art assets that meet or exceed industry standards for quality while supporting Microsoft Game Studio (MGS) business goals. https://bitly.com/ MSJobs-OtherTech

Build Engineers/Managers: Responsible for developing, managing, and ensuring effective and efficient builds of Microsoft products. http://bit.ly/MSJobs-SDE

Consultants: Deliver design, planning, and implementation services that provide IT solutions to customers and partners. http://bit.ly/MSJobs-SysOps

Consultants: Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. http://bit.ly/MSJobs-SysOps

Design Verification/Validation Engineers: Responsible for ensuring the quality of Microsoft hardware products. http://bit.ly/MSJobs-Hardware

Evangelist: Secure future growth of the Microsoft plat-

Microsoft®

form by engaging a community of customers, partners, and academics to embrace and adopt Microsoft technology. https://bitly.com/MSJobs-OtherTech

Hardware Dev. or Design Engineers, Hardware Engineers, and Design Engineers: Design, implement, and test computer hardware. http://bit.ly/MSJobs-Hardware

International Project or Localization Engineers/Managers: Ensure the successful localization of software components for foreign markets. http://bit.ly/MSJobs -Localization

Operations Program Managers: Responsible for the design, implementation, and release of programs or projects. http://bit.ly/MSJobs-ProgMgr

Premier Field Engineers: Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. http://bit.ly/MSJobs-Support

Researchers/Scientists: Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms, and prototypes. http://bit.ly/ MSJobs-Research

Service Engineers, Service Operations Engineers, and Systems/Operations Engineers: Plan, architect, deploy and/or support complex client/server or database software systems. http://bit.ly/MSJobs-SysOps

Solution Managers: Identify and analyze internal client and partner business needs, and translate needs into business requirements and value-added solutions and solution roadmaps. http://bit.ly/MSJobs-SalesEng

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs-Support

Technical Account Managers: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. http://bit.ly/MSJobs-Support

Technical Account Managers: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Roving Employee requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. http://bit. ly/MSJobs-Support

User Experience Researchers/Designers and User Research Engineers: Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. http://bit.ly/MSJobs-UX

Director, New Product Introduction: Manage the endto-end lifecycle of Microsoft's hardware development programs. Position requires domestic and international travel up to 15% with work to be performed at various unknown worksites throughout the U.S. http://www.jobs-microsoft. com/job/qo/2239157/

Product Intelligence Manager II: Responsible for building a platform to monitor the daily trends of key metrics across major dimensions. http://www.jobs-microsoft. com/job/go/2190005/

Senior Privacy and Safety Strategist: Analyze engineering requirements and business processes to identify privacy and safety issues. http://www.jobs-microsoft.com/ job/go/2231328/

Senior Program Management Lead: Coordinate program development of computer software applications, systems or services, working with development and product planning teams. Requires travel up to 20% with work to be performed at various unknown worksites throughout the U.S. http://www.jobs-microsoft.com/job/ go/2236246/

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

CAREER OPPORTUNITIES

THE UNIVERSITY OF ALABAMA IN HUNTSVILLE. The Computer Science Department of the University of Alabama in Huntsville (UAHuntsville) invites applicants for a tenure-track faculty position at the Assistant Professor level. A Ph.D. in Computer Science or closely related area is required. Our areas of interest are Data Science and Cyber Security & Information Assurance. We have a strong commitment to excellence in teaching, research and service. Successful applicants should have good communication and teaching abilities, and will be expected to develop and maintain a strong externally funded research program. UA-Huntsville is strategically located along the Tennessee River in scenic North Alabama in a rapidly expanding high technology area. Huntsville has one of the largest Research Parks in the nation and is home to NASA's Marshall Space Flight Center, the Army's Redstone Arsenal, and many high-tech industries that present many opportunities to pursue collaborative research. The surrounding population of approximately 375,000 well-educated and highly technically skilled people has access to excellent public schools and inexpensive housing. The University has an enrollment of approximately 7700 students with more than ten research centers in diverse areas including information technology, modeling & simulation, Earth Science, and space science that provide further opportunities for collaborative projects. The Computer Science Department has 12 full-time faculty and offers B.S., M.S., and Ph.D. degrees in Computer Science and M.S. degree in Software Engineering. There are approximately 254 undergraduate CS majors, 92 MS, and 34 PhD. candidates enrolled in our program. Current faculty research interests include Software Engineering, Visualization, Pattern Recognition and Image Processing, Distributed Systems, Graph Theory, Data Mining, Information Technology, Artificial Intelligence, Modeling and Simulation, Multimedia Systems, Information Assurance and Networking. According to recent NSF figures the CS department ranks 30th in the nation in federal research funding.

SR. SYSTEMS ANALYST (SAP), Memphis, TN. Travel to multiple client locations nationwide to analyze, design, develop, manage, administer sophisticated applications using SAP/ABAP in web based multiplatform environment. Work

with database related development, SAP smart forms, Reports, ALE, BAPI. Test, troubleshoot, maintain existing systems. Reply to: Comspark International Inc., 3265 W. Sarazen Circle # 201, Players Crossing Office Park, Memphis, TN 38125.

MICROSOFT CORPORATION currently has the following openings in Redmond, WA; Sunnyvale, CA; and Fargo, ND (all levels, e.g., Principal, Senior and Lead levels): Program Managers: Coordinate program development of computer software applications, systems or services. http://bit.ly/MSJobs-ProgMgr. Multiple job openings are available. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

COMPUTER CONSULTING MANAGER

(NY, NY). Conduct logical analyses of bus., scientific, engg, & other tech. problems, formulating mathematical models of problems for solution by computers for CIS & Billing S/ware Systems in Energy Utilities mkt space. Apply theoretical expertise & innovation to create or apply new tech., such as adapting principles for applying computing to new user in-



Department of Computer and Information Sciences Temple University Tenure Track Faculty

Applications are invited for tenuretrack, open rank, faculty positions in the Department of Computer and Information Sciences at Temple University. The junior position is in the software systems area, which includes

- Software Engineering and Applications,
- Database Systems, and
- Programming Languages.

The senior position for Associate or Full Professor is open to all areas of computer science/engineering. Applicants for the senior position are expected to have an outstanding track record.

Please submit applications with all requested information online at <u>http://academicjobsonline.org</u>.

For further information check <u>http://www.cis.temple.edu</u> or send email to search committee chair Dr. Eugene Kwatny at <u>gkwatny@temple.edu</u>. Review of candidates will begin on January 2, 2013 and will continue until the positions are filled. Temple University is an equal opportunity, equal access, affirmative action employer.

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.

terface reqmts. Trav'l to unanticipated worksites. Bachelor of Sci in Comp Engg & 5 yrs of exp reqd. If interested, send resumes to Peace Software Inc.: recruiting. us@hsntech.com.

MICROSOFT CARIBBEAN currently has the following opening in Guaynabo, Puerto Rico: Senior Consultant: Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Roving Employeerequires travel up to 100% with work to be performed at various unknown worksites throughout Puerto Rico, the Caribbean, and the Americas. Fluent in Spanish & English. To view detailed job descriptions and minimum requirements, and to apply, visit: http://www. jobs-microsoft.com/job/go/2236566/. EOE.

MICROSOFT CORPORATION currently has the following openings in Redmond, WA; Mountain View, CA; Palo Alto, CA; San Francisco, CA; Sunnyvale, CA; Cambridge, MA; Durham, NC; Charlotte, NC; Fargo, ND; St. Paul, MN; Salt Lake City, UT, and Madison, WI (all levels, e.g., Principal, Senior and Lead levels): Software Dev. Engineers, Software Dev. Engineers in Test, Dev. Leads, Test Engineers/ Leads, and Research Software Dev. Engineers: Responsible for developing or testing computer software applications, systems or services. http://bit.ly/ MSJobs-SDE. Multiple job openings are available. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.



Juniper Networks is recruiting for our Sunnyvale, CA office: **Technical Support Engineer** Senior Staff: Provide technical leadership to support teams by providing an example of excellent customer service & technical ability. Act as an escalation point for company's Advanced Technical Assistance Center focused on, but not limited to the North America region. Mail single-sided resume with job code #3007 to Attn: MS A1.2.1.435 Juniper Networks, 1194 N. Mathilda Ave., Sunnyvale, CA 94089.

FACULTY POSITIONS AT WRIGHT STATE UNIVERSITY

The Department of Computer Science and Engineering (CSE) at Wright State University seeks applicants for up to two positions at assistant, associate, or full rank with rank and tenure status appropriate to qualifications and experience. Candidates for these positions are expected to have an earned Ph.D. in computer science, computer engineering or a closely related field anticipated by the start date with outstanding academic credentials. Candidates applying for a position at the rank of assistant professor must clearly demonstrate the potential to develop a vibrant funded research program that engages graduate students and produces peer reviewed publications while candidates for the rank of associate/full professor must have an outstanding record of funded research and scholarly publications. In addition, candidates must possess excellent communication skills and a commitment to engage in both undergraduate and graduate education.

For one of the positions the departments seeks a faculty member specializing in cyber security including areas such as software security; trustworthy systems; information assurance, security and privacy; mobile and embedded security; computer forensics; and security tools and visualization. Outstanding applicants with a high potential for contributing to the department's newly developed Master program in Cyber Security are strongly encouraged to apply.

For the other position, the department seeks faculty specializing in Big Data research. Particular areas of interest include, but are not limited to, data management and lifecycle, data analytics, data visualization, data fusion and integration, semantics and ontologies, social and sensor Web, biomedical and health informatics. Outstanding applicants with a high potential for collaborations with existing strengths of the department and the Kno.e.sis Center (<u>http://knoesis.</u> wright.edu) are particularly welcome to apply.

Outstanding applicants specializing in other emerging research areas are also welcome to apply.

The Department has 26 faculty members, more than 500 undergraduate, 75 M.S. and 40 Ph.D. students and offers B.S., M.S. and Ph.D. degrees both in Computer Science and Computer Engineering, and an M.S. in Cyber Security. Information about the Department can be found at: http://www.cs.wright.edu/cse/. The Department is located in the Russ Engineering Center and Joshi Research Center, which includes the Kno.e.sis Center and the Appenzeller Visualization Laboratory. The Department is one of four departments in the College of Engineering and Computer Science, which houses 4 out of 7 of Wright State's University System of Ohio Centers of Excellence (http://webapp2.wright.edu/web1/coe/ category/centers-of-excellence/). Wright State University, an institution of nearly 19,000 students, is located on a spacious campus within a growing suburban community. A variety of affordable and pleasant living environments with schools and parks attractive to professionals are conveniently located close to campus. Wright State University is surrounded by industry leaders including Lexis-Nexis, Reynolds & Reynolds, CSC, Ball Aerospace, Northrop Grumman, Teradata, and SAIC. Wright State is also located adjacent to the Wright-Patterson Air Force Base, which houses the headquarters of the Air Force Research Laboratory. The university is committed to industrial and government partnerships for research and economic development ventures and has a strong institutional commitment to underrepresented groups, women, persons with disabilities, and veterans. Applicants should provide a brief statement of their research, teaching interests, and professional goals. The application should include a cover letter indicating the rank desired and a complete vita with the names, addresses, telephone numbers and e-mail addresses of at least four references. Applications and supporting information for the Big Data area is completed on-line at: https://jobs.wright. edu/postings/5916 and applications for Cyber Security area is completed on-line at: https://jobs.wright.edu/postings/5914.

Consideration of candidates begins January 1, 2013 and continues until the positions are closed or filled. Salaries and resources are competitive and based on rank. For details and additional information, you may contact Prof. Mateen Rizki, Chair, at <u>mateen.rizki@wright.edu</u>, or Prof. Thomas Wischgoll (for the cyber security search) at <u>thomas.wischgoll@wright.edu</u>, or Prof. Pascal Hitzler (for the Big Data search) at <u>pascal.hitzler@wright.edu</u>.

Wright State University is an equal opportunity/affirmative action employer.

CAREER OPPORTUNITIES

SR. PROGRAMMER ANALYST - New York, NY: Perform Analysis, design & devel of apps working in JAVA techs using J2EE, EJB, Servlets, JDBC on NT & UNIX platforms; Employ design patterns & best practices for high performance, scalable, extensible apps on weblogic/ websphere App Servers; Define interfaces using Object Oriented methodologies to be implemented by the programmers; Lead cross regional teams in Americas, EMEA & APAC to implement & integrate tech solutions; Design & devel Webservices interfaces w/SOA & JAX WS tech; Integrate business rules configured in ILog JRules with app; Performance tuning of app compenents & database queries for Oracle 10g; Will use Java 6, Java EE 5, Weblogic 10g, Oracle 11g, SOA & Webservices, Java/J2EE design patterns, JSF, ILog JRules, & Messaging. Reqd: Bachelor's deg in Comp Sci, Eng, Math or MIS & 5 yrs exp. Any suitable combin of edu, training & exp is acceptable. Send resume to: Archer IT, 469 7th Ave, Ste 248, New York, NY 10018.

THE UNIVERSITY OF HAWAII. The Department of Information and Computer Sciences (ICS) at the University of Hawai'i at Mānoa (UHM) is pleased to announce two tenure track faculty positions, pending availability of funds. The ICS department at UHM has a more than 40-year history of excellence and inno-

Peabody Investments Corp.

has an opening in St. Louis, MO for a

SAP BI/BO ARCHITECT

Responsible for implementation of BO Integration Kit for SAP Solutions. Bachelor's in Ind Eng, Comp Sci., or related field and 5 yrs exp in SAP Business Intelligence/Business Objects req'd. Must have SAP BI/BO certification. Must have legal auth. to work permanently in the US. EOE.

> Send resumes to: L. Peterson, Peabody Investments Corp. HR Dept., 701 Market Street, St. Louis, MO 63101 Must note Job ID #: SAPBIBOA

vation in information and computer science and offers 6 degrees: BA, BS, MS, MLISc and 2 PhDs. Applicants must have a PhD from an accredited college or university in computer science or a closely related discipline. The ICS department is particularly interested in applicants with a specialization in one of these areas: security/information assurance, big data analytics, computer graphics/visualization, or foundations of computing/ algorithms. Salary and hiring rank will be commensurate with qualifications and experience. Applications should include a cover letter, vita, statement of research and teaching interests, and contact information for three references. Please send your application or any questions to cssearch@hawaii.edu or Search Committee Chair, Univ of Hawaii Dept of ICS, 1680 East West Rd, POST 317, Honolulu, Hawaii 96822. Additional information about the department can be found at www.ics. hawaii.edu. Closing Date: 1/20/2013. The University of Hawai'i is an Equal Opportunity Affirmative Action Employer.

FIRST RESERVE CORPORATION seeks Sharepoint Architects in Greenwich, CT to interface w/ end users & stake holders, defining req'ts, designing, dvlpg & deploying SharePoint based bus. apps. Provide tech support, SharePoint admin, training & oversight for clients. Req's: Master's or equiv in IT, CS, Eng'g or rel. discipline & 2 yrs exp providing support, training & consulting srvcs to end users. Exp must incl prog'g in C#.net, VB.net, ASP.net & Java; designing pages & style sheets using Custom SharePoint Workflows; bldg custom web parts; req'ts gathering: SharePoint Indexing: backup & restore planning & ops for SharePoint in multi-homed SharePoint environs; dvlpg apps across broad spectrum of user req'ts, HR apps, & IT support apps; & documenting mgmt solutions. MCTS in SharePoint 2007 & ITIL Foundations certs are req'd. Pls submit resume to Human Resources, First Reserve Corporation, LLC, 1 Lafayette Place, Greenwich, CT, 06830. Please indicate job code VM11912IEEE.

SOFTWARE DEVELOPMENT SPE-CIALIST, HDNY, INC. (Lyme, CT) Install, modify & implmt new product or enhancements to current Digital mgmt Asset & content system for music & video, for desktop, laptop & mobile devices. Communicate w/ client to collect, interpret, doc, customized system reqmts. Apply knowl of music & video environments. Perform gap analysis, resolve issues. Test modules & integrate w/internal site modules. Microsoft .Net, SQL,

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.


YAHOO! HAS THE FOLLOWING OPENINGS AVAILABLE:

YAH00! INC. seeks Tech Yahoo, Database Administrator in Richardson, TX. Define, implement and maintain standard operating procedures for the operations team within the Corporate Applications group. TO APPLY: Submit your resume to the following URL: https:// tas-yahoo.taleo.net/careersection/ y a h o o _ p mt _ c s / j o b s e a r c h . ftl? lang=en&portal=34140220106. Must reference Reg ID #: 1249239. EOE.

YAHOO! INC. seeks Tech Yahoo, Software Apps Dev Eng, Sr. in Atlanta, GA to Design and develop software for mobile devices, and web that support. TO APPLY: Submit your resume to the following URL: https:// tas-yahoo.taleo.net/careersection/ y a h o o _ p m t _ c s / j o b s e a r c h . ftl?lang=en&portal=34140220106. Must reference Req ID #: 1249191. EOE.

UNIVERSITY OF NOTRE DAME

The Department of Computer Science and Engineering at the University of Notre Dame invites applications for Assistant or Associate Professor. Excellent candidates in all areas will be considered.

The Department offers the PhD degree and undergraduate Computer Science and Computer Engineering degrees. Faculty are expected to excel in classroom teaching and to build and lead highly-visible research projects that attract substantial external funding.

The University of Notre Dame is a private, Catholic university with a doctoral research extensive Carnegie classification, and consistently ranks in USNWR as a top-twenty national university. The South Bend area has a vibrant and diverse economy with affordable housing and excellent school systems, and is within easy driving distance of Chicago and Lake Michigan.

Applicants should send (pdf format preferred) a CV, statement of teaching and research interests, and contact information for three professional references to: cse2013@nd.edu.

The University of Notre Dame is an Equal Opportunity, Affirmative Action Employer.

FIORIDA INTERNATIONAL UNIVERSITY

Florida International University is a multi-campus public research university located in Miami, a vibrant, international city. FIU is recognized as a Carnegie engaged university. Its colleges and schools offer more than 180 bachelor's, master's and doctoral programs in fields such as computer science, engineering, international relations, architecture, law, and medicine. As one of South Florida's anchor institutions, FIU is worlds ahead in its local and global engagement, finding solutions to the most challenging problems of our time. FIU emphasizes research as a major component of its mission and enrolls 48,000 students in two campus and three centers including FIU Downtown on Brickell and the Miami Beach Urban Studios. More than 160,000 alumni live and work in South Florida. For more information about FIU, visit http://www.fiu.edu/.

The School of Computing and Information Sciences seeks exceptionally qualified candidates for multiple tenure-track and tenured faculty positions at all levels as well as non-tenure track faculty positions at the level of Instructor.

TENURE TRACK/TENURED POSITIONS (JOB ID# 505004)

We seek well-qualified candidates in all areas of Computer Science and researchers in the areas of programming languages, compilers, databases, information retrieval, computer architecture, scientific computing, big data, natural language processing, computational linguistics, health informatics, and robotics, are particularly encouraged to apply. Preference will be given to candidates who will enhance or complement our existing research strengths.

Ideal candidates for junior positions should have a record of exceptional research in their early careers. Candidates for senior positions must have an active and proven record of excellence in funded research, publications, and professional service, as well as a demonstrated ability to develop and lead collaborative research projects. In addition to developing or expanding a high-quality research program, all successful applicants must be committed to excellence in teaching at both graduate and undergraduate levels. An earned Ph.D. in Computer Science or related disciplines is required.

NON-TENURE TRACK INSTRUCTOR POSITIONS (JOB ID# 505000)

We seek well-qualified candidates in all areas of Computer Science and Information Technology. Ideal candidates must be committed to excellence in teaching a variety of courses at the undergraduate level. A graduate degree in Computer Science or related disciplines is required; significant prior teaching and industry experience and/or a Ph.D. in Computer Science is preferred.

Florida International University (FIU), the state university of Florida in Miami, is ranked by the Carnegie Foundation as a comprehensive doctoral research university with high research activity. The School of Computing and Information Sciences (SCIS) is a rapidly growing program of excellence at the University, with 36 faculty members and over 1,500 students, including 75 Ph.D. students. SCIS offers B.S., M.S., and Ph.D. degrees in Computer Science, an M.S. degree in Telecommunications and Networking, and B.S., B.A., and M.S. degrees in Information Technology. SCIS has received approximately \$17.5M in the last four years in external research funding, has six research centers/clusters with first-class computing infrastructure and support, and enjoys broad and dynamic industry and international partnerships.

HOW TO APPLY:

Applications, including a letter of interest, contact information, curriculum vitae, academic transcript, and the names of at least three references, should be submitted directly to the FIU Careers website at <u>https://jobsearch.fiu.edu</u>; refer to Job ID# 505004 for tenure-track or tenured positions and to Job ID# 505000 for instructor positions. The application review process will begin on January 7, 2013, and will continue until the position is filled. Further information can be obtained from the School website <u>http://www.cis.fiu.edu</u>, or by e-mail to recruit@cis.fiu.edu.

FIU is a member of the State University System of Florida and is an Equal Opportunity, Equal Access Affirmative Action Employer.

CAREER OPPORTUNITIES

2 yrs exp. 40 hrs/wk. Send res to jobs@ hdny.net

SENIOR PROGRAMMER ANALYST, Piscataway, NJ: Travel to multiple client locations nationwide to perform programming services including, requirement gathering, defining functional specifications, designing architecture, coding, integration, testing, deploying, troubleshooting using Informatica, SQL Server, ETL, OWB, Oracle/PLSQL in multiplatform environment. Support users. Reply to: Infotech Global, Inc., 371 Hoes Court, ste 300 A Piscataway, NJ 08854.

EMPHASIS SOFTWARE DEVELOP-MENT has an opening in Charlotte, NC for Sr. Application Developer to analyze requirements & design for TLM product & Oracle ODI ETL tool. Requires BS+7 yrs exp in systems or network programming. Email resumes to employment@ citco.com & refer to job code ESD72 in email subject line.

SENIOR SOFTWARE ENGINEER, Gainesville, FL – Design, implement and test software solutions, algorithms, and data structures to solve transportation and logistics problems. Master's degree in Computer Science, Engineering or a related field and 5 years progressive experience in the field required. Salary commensurate with exp. 40 hrs/ wk, 8 AM - 5 PM, M – F. Mail resume to: Innovative Scheduling, Inc., 2153 SE Hawthorne Road, Suite 128, Gainesville, FL 32641.

HOUSTON INTERNATIONAL INSUR-ANCE GROUP, LTD., Houston TX seeks a Senior Analyst Product Development

STATE UNIVERSITY OF NEW YORK AT BINGHAMTON Department of Computer Science

Applications are invited for two tenure-track Assistant Professor Positions beginning Fall 2013. Preferred specializations include embedded systems, energy-aware computing and security. The Department has about 800 majors, including 63 full-time PhD students. Junior faculty have a significantly reduced teaching load for at least the first three years. A new NSF supported industry-university collaborative research center on energy-smart electronic systems offers an added venue for research and funding. Apply online at:

http://binghamton.interviewexchange.com First consideration given to applications received by January 31, 2013.

We are an EE/AA employer.

to research, design, document, and modify software specifications for software utilized in the litigation and/or insurance industry throughout the production lifecycle. Bachelor's degree in Computer Science and 5 yrs. exp. in job offered. Domestic travel required 10% of the time. International travel to Pakistan required 15% of the time. Email resume to jmasterson@hiig.com. REF: SAPD.

SR. CONTROL SYSTEMS ENGINEER

(N. Reading, MA) Must hv PhD in controls, mech. engg or EE or foreign equiv & 2 yr exp (pre or post PhD) in job offd or in systems engg applying estimation & control theory, technical principles in electro-mechanical engg & embedded s/ware dsgn; or must hv MS in same & 4 yr exp (pre or post MS) in same. Send resumes to Mary Beth Christopher, Kiva Systems, Inc., 300 River Park Dr., N. Reading, MA 01864 or to mchristopher@kivasystems.com. Refer to job#002.

SR. QA AUTOMATION ENGINEER (N. Reading, MA) Must hv MS in CS, EE or s/ ware engg, industrial tech., math, info systems or foreign equiv & 2 yrs exp in job offd or in quality assurance engg w/ automation &/or prgmg skills; or must hv BS in same & 5 yrs exp in same. Send resumes to Mary Beth Christopher, Kiva Systems, Inc., 300 River Park Dr., N. Reading, MA 01864 or to mchristopher@kiva-systems.com. Refer to job#004.

PROGRAMMER ANALYST (South Plainfield, NJ) Dsgn data architecture in SQL Server 2008; perform data modeling w/use of Erwin tool; dvlp/understand/ improve operational & installation procedures for wide range of regmts incl. communication systems, h/ware, network, security & s/ware storage; data migration planning/dsgn; performance tuning (Dynamic Mgmt. Views, SQL Profiler); unit test planning/dsgn; dvlp T-SQL code, stored procedures, Triggers, Functions, Partitioning Functions/ Schemes & User Defined Types. Min. req. BSC in Mgmt. Info. Systems or Comp Sci (or foreign equiv. deg) + 5 yrs work exp in job off'd. Travel req. to unanticipated client sites in USA. Resume to Job Loc: InfoLogitech Inc., 50 Cragwood Rd, Ste 209, South Plainfield, NJ 07080.

CALIFORNIA STATE UNIVERSITY, FULLERTON (CSUF) seeks applications for a tenure-track Assistant Professor in Computer Engineering beginning Fall 2013. Please refer to http://diversity. fullerton.edu/jobs/ft/ECS_CompEng. asp for additional details. Review of applications will commence immediately and will continue until the position is filled. Applicants must submit electronic versions of curriculum vita and a list of three references to: snorell@ fullerton.edu.

THE UNIVERSITY OF NORTH TEXAS. **Department of Computer Science and** Engineering, Assistant Professor. The Department of Computer Science and Engineering at the University of North Texas (UNT) is seeking candidates for a tenure-track faculty position at the Assistant Professor level beginning August 15, 2013. The department plans to build on its existing strengths in Computer Security, including network security and intrusion detection, secure software systems, vulnerability analysis, and machine learning techniques applied to computer security. Candidates should have demonstrated the potential to excel in research in one or more of these areas and in teaching at all levels of instruction. A Ph.D. in Computer Science, Computer Engineering or closely related field is required at the time of appointment. An Applicant's record must include high quality publications. The Computer Science and Engineering department is home to 812 bachelor students, 140 masters students and 82 Ph.D. students. The UNT Center for Information and Computer Security, housed in the department, has been recognized by the National Security Agency as a National Center for Academic Excellence in Information Assurance Research and Education and offers several certificate programs in computer security. Additional information about the department and center are available at the websites: www.cse.unt.edu and www. cics.unt.edu, respectively. Application Procedure: All applicants must apply online to: https://facultyjobs.unt.edu/applicants/Central?quickFind=51736. Submit nominations and questions regarding the position to Dr. Philip Sweany (sweanv@cse.unt.edu). Application Deadline: The committee will begin its review of applications on December 1, 2012 and continue until the position is closed. The University: With about 36,000 students, UNT is the nation's 33rd largest university. As the largest, most comprehensive university in Dallas-Fort Worth, UNT drives the North Texas region. UNT offers 97 bachelor's, 82 master's and 35 doctoral degree programs, many nationally and internationally recognized. A studentfocused public research university, UNT is the flagship of the UNT System. UNT is strategically located in Denton, Texas, a vibrant city with a lively arts and music culture, at the northern end of the Dallas-Fort Worth metroplex. The DFW area has more than six million people, with significant economic growth, numerous industrial establishments, and excellent school districts. The University of North Texas is an AA/ADA/EOE committed to diversity in its educational programs.



Southern Methodist University Position #50049 Department of Computer Science and Engineering Faculty Position in Computer Science and Engineering

The Department of Computer Science and Engineering in the Lyle School of Engineering at Southern Methodist University invites applications for a faculty position in computer science and engineering beginning Fall 2013. Individuals with experience and research interests in all areas of computer science and engineering are encouraged to apply. Priority will be given to individuals with expertise and research interest in data mining, informatics, computer systems and networking, and related areas. The search is focused at the tenure-track assistant professor level. The successful candidates must have or expect to have a Ph.D. in computer science, computer engineering, or a closely related area by date of hire. Successful applicants will demonstrate a deep commitment to research activity in computer science and engineering and a strong record of excellence in teaching.

The Dallas/Fort Worth area, one of the top three high-tech industrial centers in the country, has the largest concentration of telecommunications corporations in the US, providing abundant opportunities for industrial research cooperation and consulting. Dallas/Fort Worth is a multifaceted business and high-tech community, offering exceptional museums, diverse cultural attractions, and a vibrant economy.

The CSE Department resides within the Bobby B. Lyle School of Engineering and offers BS, MS, and Ph.D. degrees in Computer Engineering and Computer Science, the Doctor of Engineering in software engineering, and the MS in Security Engineering and Software Engineering. The department currently has 15 faculty members with research concentrations in security engineering, software engineering, computer networks, telecommunications, data mining, database systems, VLSI and digital systems, and computer arithmetic. Additional information may be found at: <u>www.lyle.smu.edu/cse</u>.

To receive full consideration, interested individuals should send a complete resume and names of three references, including a one-page statement of research interests and accomplishments by December 21, 2012 to:

csesearch@lyle.smu.edu

or

CSE Faculty Search Position #50049 Department of Computer Science and Engineering SMU

Dallas, TX 75275-0122

Review of applicants will begin immediately and will continue until the positions are filled. Hiring is contingent upon the satisfactory completion of a background check.

SMU will not discriminate on the basis of race, color, religion, national origin, sex, age, disability, or veteran status. SMU is committed to nondiscrimination on the basis of sexual orientation.

Apple has the following job opportunities in Cupertino, CA:

Cellular SW Development Engineer (4 Openings) [Req. #21821226] iOS Software Engineer [Req. #21875705] Software Engineer [Req. #22024307] Software Development Engineer [Req. #22023585] Sr. Software Engineer [Req. #22880339] Software Engineer [Req. #23463645] Software Engineer Apps [Req. #23465376]

Some positions may require travel and may have direct reports. Mail resumes to 1 Infinite Loop M/S: 104-1GM, Attn: SA, Cupertino, CA 95014. Principals only. EOE.

Siri Speech Runtime Engineer [Req. #21738825] Optimize the environment by coming up with methods to fine tune systems, which also contributes to capacity planning. Requires ten (10) years professional experience in job offered or in a related occupation. Must have professional experience with: Linux systems (RHEL/CentOS/OEL 5.x); Shell/Bash, Perl, PHP, Python, Ruby and Java; IP networking and core services (NFS/DHCP/NTP/SSH/DNS/SNMP/HTTPD); hardware management (firmware/BIOS/driver updating); troubleshooting abilities and team environment to fix large production issues.

Instrumentation Engineer [Req. #21822368] Design innovative touch technology products with best-in-class technologies and user experience. Requires Master's degree, or foreign equivalent, in Electrical Engineering, or related field plus two (2) years of professional experience in the job offered or in a related occupation. Must have academic knowledge of or experience with: Application software development in Mac OS X environment; electronic circuit design and debugging, including analog and digital subsystems; reviewing mechanical CAD; scientific method, experimental process, root cause analysis; authoring engineering specification documents. Requires international travel approximately 25% of the time.

Senior Timeline UI Engineer [Req. #21821412] Design, implement and support architecture for user interactions and layout of timeline using Core Animation. Requires Bachelor's degree or foreign equivalent in Computer Science or related field plus five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate progressive in nature. Must have professional experience with: Cocoa and Objective-C; object oriented analysis and software architecture design skills; applied software methodology, process and quality skills; analyzing, debugging and improving UI interaction and drawing performance with Core Animation.

Software Development Engineer [Req. #21735982] Responsible for development of image processing algorithms for state of the art photo and computational software features in Apple's embedded cameras. Req.'s Master's degree, or foreign equivalent, in Electrical Engineering, Computer Science, or related. Six (6) years professional experience in job offered or in a related occupation. Must have academic background or professional experience with: Computer programming; Computational and photo imaging algorithm development; Code performance optimization.

iOS Device Driver Engineer [Req. #21736252] Design and develop state of the art Board Support Packages involving device drivers for Universal Serial Bus, Universal Asynchronous Receiver Transmitter, Serial Peripheral Interface Bus and Inter Integrated-Chip Sound on real-time/embedded systems development platforms such as iOS operating system and system boot-up and initialization. Requires Bachelor's degree in Computer Science, Computer Engineering or Electrical Engineering or related field plus five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate progressive in nature. Must have professional experience with: C Coding; Debugging with JTAG (Joint Test Action Group) tools; Universal Asynchronous Receiver Transmitter; Universal Serial Bus; Crash analysis; Throughput analysis; Embedded real time operating system.

Senior Cellular Performance Engineer [Req #21483719] Responsible for baseband/radio firmware development with C/C++. Req.'s Bachelor's degree, or foreign equivalent, in Computer Engineering, Computer Science or related plus five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have academic background or professional experience with: C/C++ programming; CPU Core dump analysis; CPU Memory; Baseband IPC.

Software Engineer [Req. #21738347] Work on software development of large scale graph processing algorithms. Req.'s PhD, or foreign equivalent, in Computer Science, or related plus six (6) months of professional experience in job offered or in a related occupation. Must have academic background or professional experience with: 'Big Data" programming and processing; building hosted services in high-volume environment; Linux, Solaris, BSD, OS X; C/C++ or Objective-C programming; Graph algorithms.

Camera Module EPM (2 Openings) [Req. #21873021] Coordinate the design, development and integration of camera modules into iPhone, iPad, Mac, and iPod products. Req's Bachelor's degree, or foreign equivalent, in Electrical Engineering, Mechanical Engineering, Industrial Engineering, Physics, or related field. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have academic background or professional experience with: cross functional communication; high volume consumer device production and test processes; communication and presentation of technical topics to senior staff; one or more of the following: OS sensors and SOCs, optics and lens design, camera module design, development management, project management, mobile phone system design including system integration and manufacturing. May require 25% of international travel.

Silicon Validation Engineer (2 Openings) [Req. #21876121] Provide silicon bring-up, functional validation and debug from prototype to production. Req's Master's degree, or foreign equivalent, in Electrical Engineering, Computer Engineering, or related field. Two (2) years professional experience in job offered or in a related occupation. Must have academic knowledge or professional experience with: C/C++ programming; debuggers (GDB/LLDB, JTAG) disassemblers; TCL/Expect/Perl Scripting, embedded programming; Technical understanding of system-on-chip architecture, logic blocks and IO devices; Lab tools, Oscilloscopes, Logic Analyzers.

Baseband Firmware Engineer [Req. #21875877] Provide software crash analysis with coredump registers. Req's Bachelor's degree, or foreign equivalent, in Computer Science, Computer Engineering, or related field. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have academic background or professional experience with: C/C++ programming; JTAG Debugging; Communication Transport including USB, UART, I2C, SPI; Battery/Power. May require 10% of international travel.

Retail Materials Planning Lead [Req. #21823282] Responsible for service parts planning for assigned commodities. Req's Master's degree, or foreign equivalent, in Business, Operations, Supply Chain or related field. Three (3) years professional experience in job offered or in a related occupation. Must have professional experience with: demand planning and supply chain management for Hi-Tech/consumer products industry; Forecasting and Demand Planning; S&OP; Business Process re-engineering; Presenting and Leading business meetings; Statistical modeling; Microsoft Excel including data models, macros and pivot tables.

System Design Engineer (2 Openings) [Req. #21816897] Work on radiated performance of wireless portable devices. Requires Master's degree, or foreign equivalent, in Electrical & Computer Engineering, Engineering Science or related. Must have academic background or professional experience with: Electromagnetic theory and antenna and wave propagation; radiation performance; RF instrumentation; GSM/GPRS/EDGE, 802.11, Bluetooth; Wireless consumer devices, including FCC, PTCRB, CTIA, ETSI, TIA//EIA. May require less than 10% of international travel.

Senior Camera Technology Specialist [Req. #20979827] Responsible for actuator technology assessment and selection. Req.'s Master's degree, or foreign equivalent, in Electrical Engineering, Mechanical Engineering, Engineering, Business Administration or related plus Ten (10) years professional experience in job offered or in a related occupation. Must have professional experience with: Actuator structural design; Actuator dynamic design; Actuator control system design; Actuation fundamental mechanism design; Supplier management in research and development; Innovation and IP generation in advanced actuator technology. May require 25% of both international and domestic travel time. Position includes direct reports.

Senior Software Engineer [Req. 20980687] Responsible for performance analysis and tuning of iPhone, iPod, and iPad products. Req.'s Bachelor's degree, or foreign equivalent, in Computer Science, or related. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional experience with: Performance tuning and analysis of system; Microprocessor Architecture; Kernel; Assembly language programming; C, C++ and Objective C, Mac OS X and Mac OS Performance tools.

Quality Assurance Engineer [Req. #22675032] Design and write on-the-fly automated stress and stability hardware-centric software tests for embedded systems. Req.'s Bachelor's degree, or foreign equivalent, in Electrical Engineering, Computer Science, Computer Engineering, or related. Must have academic background or professional experience with: Perl, Python, Ruby, or JavaScript; object oriented programming experience in C, C++ or Java; core engine design, game logic, animation techniques, and environment interaction; OpenGL experience, including knowledge of shader / GPU programming.

Search and Discover Engineer [Req. #22024659] Responsible for the iTunes Match, Genius, Popularity Charts, Recommendations, and many other personalized features of the iTunes Store. Requires Master's degree, or foreign equivalent, in Computer Science, Industrial Engineering or related. Must have academic background or professional experience with: Highly scalable applications and web services in Java; Hadoop; Berkely DB; Matrix computations; Machine learning; R language; collaborative filtering.

Engineering Manager [Req. #22024970] Responsible for design and process development of advanced flat panels from concept to product ramp. Requires Master's degree, or foreign equivalent, in Physics, Industrial Engineering, Chemistry, Electronic Engineering, Engineering, or related plus three (3) years professional experience in job offered or in a related occupation. Must have professional experience with: TFT/CF/Cell process; TFT/CF/Cell process DOE, SPC and failure analysis; TFT device engineering, TFT LCD circuit; TFT-LCD structures and display optics; panel design and array process; negotiation with panel suppliers. Will have direct reports. May require 5-10% international travel to organize technical discussion with panel suppliers in Japan and other Asian areas.

Physical Design Engineer (multiple openings) [Req. # 22882938] Responsible for physical design of high speed large blocks, about 1 million instances, floor planning, placement, timing optimization, clock tree synthesis, routing, and post route optimization. Requires Bachelor's degree, or foreign equivalent, in Electrical Engineering, Electronic Engineering, Computer Science, or related field plus five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional experience with: physical design of a block in the chip; closing timing at block level in a chip, deep sub-micron circuit phenomena including cross talk, temperature inversion, sub-threshold conductance and impact on leakage current and power; performing Engineering Change Order (ECO) on the physical design; performing Design Rule Checks (DRC) and Layout versus schematic (LVS) checks. May have direct reports. Requires 4% international travel.

Senior ABAP Developer [Req. #22881665] Create optimal and scalable technical solutions. Req's Bachelor's degree, or foreign equivalent, in Computer Engineering, Computer Science, Management Information Systems, or related field. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional development experience with: ABAP objects in SD/MM/FI Modules of SAP; Classical/ALV/Interactive Reports, SmartForms/SAPScripts, Dialog Programs and BDC Programs; ALE/IDoc's, RFC's, and BAPI's; BADI's, BTE's and Enhancement Framework, SPAU Resolution/OSS Note Applications; ABAP Web Dynpro, Workflows, and Web Services (SOA).

Test Engineer [Req. #23461170] Responsible for system test process development, including upgrades to product feature testing requirements. Requires Bachelor's degree, or foreign equivalent, in Computer Science, Electrical or Computer Engineering, or related field plus seven (7) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have academic or professional experience with: Diagnostic/test infrastructure for computer product; computer industry experience; programming languages including C, C++ and script program; and high-volume consumer product manufacturing skill. Domestic and international travel required approximately 35% of time.

Engineering Systems Administrator [Req. #23308784] Responsible for installing, configuring and maintaining all systems related to iCloud Mail. Req's Bachelor's degree, or foreign equivalent, in Computer Engineering, Information Technology, Computer Science, or related field. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional experience with: InfiniBand and FiberChannel Connectivity; Network Appliance products FAS6000 and FAS6200 series with Data ONTAP 7.3, 8.0 (snapvault, snapmirror and syncmirror); Oracle Solaris 10/11 including ZFS Filesystem; Oracle Communications Messaging Server (formerly Sun Internet Messaging and Directory Server); Anti-spam solutions including at least one of the following: Brightmail, Proofpoint, Ironport; Network Load Balancers, Citrix Netscalers; Worked for large scale environment (with at minimum of 50+ QA Unix servers and 50+ storage appliances).

IOS Engineer Lead [Req. #23639927] Build and manage an engineering team focused on building a high volume web service. Req's Bachelor's degree, or foreign equivalent, in Computer Engineering, Computer Science, or related field. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have academic background or professional experience with: leading teams of engineers in developing high volume web service applications; designing and architecting large scale hosted services in a mobile web environment; C++; data structures and algorithms; base technologies including: networking, TCP/IP, HTTP, Sockets, Threads, STL and templates; Computational Geometry or Graph Theory. May have direct reports.

Sr. Software Engineer [Req. #23489232] Responsible for designing/architecting large-scale enterprise applications and services using object-oriented languages. Req.'s Bachelor's degree or foreign equivalent in Engineering, Information Technology or related degree plus five (5) years of experience in the job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional experience with: object-oriented programming and design in Java; relational database systems such as Oracle and SQL/PL/SQL; RESTful web services, SOAP, XML-RPC, and JSON, XML; build tools Ant or Maven; version control systems Subversion or CVS.

Senior Software Quality Assurance Engineer [Req. #23486691] Responsible for Accessory Protocol verification for Apple devices. Requires Master's degree or foreign equivalent in Computer Science, Electrical Engineering, or related degree plus one year of experience in the job offered or in a related occupation. Must have professional experience or academic knowledge with: Shell scripting and high-level scripting languages such as Perl or Python; testing harness and writing system and user interface test automation; testing data transfer protocol; connect, configure, and verify operation status of a device interface; verify device configuration and network connectivity using ping, traceroute, telnet, SSH or other utilities; Linux/Unix/Mac OS; Database and Structured Query Language; Software Quality Assurance methodologies & practices, software development lifecycle; Mobile application development; development board bring-up; source control system; HTML, XML; operation of the protocols in the OSI and TCP models; understanding the components required for network and Internet communications.

APPLIED MATERIALS®

Applied Materials, Inc. is accepting resumes for the following positions in Santa Clara/Sunnyvale, CA:

Software Engineer (SCRKL): Develops requirements and design, completes programming and performs testing and debugging of applications.

Technical Support Manager (SCSBA): Responsible for managing the division's warranty responsibilities and ensuring that customer difficulties encountered with Applied Materials products are resolved in a timely and efficient manner. Position may be assigned to work at unanticipated worksites throughout the US as determined by headquarters.

Member of Technical Staff (SVTKU): Develops new or modified process formulations, defines process or handling equipment requirements and specifications, reviews process techniques and methods applied in the fabrication of integrated circuits.

Please mail resumes with reference number to Applied Materials, Inc., 3225 Oakmead Village Drive, M/S 1217, Santa Clara, CA 95054. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

www.appliedmaterials.com

Experimenting with your hiring process?

Finding the best computing job or hire shouldn't be left to chance. IEEE Computer Society Jobs is your ideal recruitment resource, targeting over 85,000 expert researchers and qualified top-level managers in software engineering, robotics, programming, artificial intelligence, networking and communications, consulting, modeling, data structures, and other computer science-related fields worldwide. Whether you're looking to hire or be hired, IEEE Computer Society Jobs provides real results by matching hundreds of relevant jobs with this hard-to-reach audience each month, in *Computer* magazine and/or online-only!

http://www.computer.org/jobs

The IEEE Computer Society is a partner in the AIP Career Network, a collection of online job sites for scientists, engineers, and computing professionals. Other partners include *Physics Today*, the American Association of Physicists in Medicine (AAPM), American Association of Physics Teachers (AAPT), American Physical Society (APS), AVS Science and Technology, and the Society of Physics Students (SPS) and Sigma Pi Sigma.

IEEE computer society

Work/Play

David Alan Grier and Erin Dian Dumbacher



Gaming has an odd relationship with software and software engineering.

oes the thrill of competition propel computer science forward?

//DAG// The two men had different reasons for playing a game of foosball that night, but neither indicated what they were. One was taking time out to compete in a casual game, where the stakes were low; the other was facing the only human competition of the day. They both stood over the table with the same hunched stance, sending the ball flying with equal force.

For the CEO of a midsized software services company, the game came at the end of the day. He had completed hours of meetings to raise capital, find customers, and keep the organization working harmoniously. He was involved in a competition with his board, having crossed swords with them earlier this year, temporarily losing his job. He was not about to lose it again.

For his opponent, one of his IT programmers, the game was an opportunity for competition with a human before settling into his nightly race against time. Before the sun rose, he needed to complete a given number of lines of code, assure their accuracy, and integrate them into the system.

//EDD// By day, many computer
science practitioners work in dull

office buildings with gray walls and long hallways. By night, they play for their own benefit, their own fun. Whether that means playing games or creating them, the thrill of competition is their prize for the working hours.

"Why buy when you can code?" asked one computer scientist friend as he demo-ed his new app. "I created this just for fun. It's in the store, but I don't expect many to buy it." That's what they all say, I thought.

Most wannabe technology tycoons have heard the stories of Silicon Valley garages. Yet, after spending time in the office or otherwise supporting the health of their bank accounts, the creative often find ways to compete against themselves. Can I make this, and will it work? Will people like it and find it useful? If the answer is "yes" to all of the above, then it's not just a personal success but also likely a commercial one.

Forecasts say the worth of the gaming industry will amount to more than \$80 billion in the next five years. Software for the workplace is widespread; there's a growing mar-

ket for software to fill free time. //**DAG**// Gaming has an odd relationship with software and software engineering. Many individuals entered the field because they had played videogames in their youth. Few probably realized that the work of designing and developing programming systems, even when those systems are videogames, has little in common with the act of playing a game. Nonetheless, games and gaming seem to hover over software engineering, and our technical staffs turn to competition even when their jobs fail to provide it.

B oth foosball players that night cared deeply about the game's result. They cheered at each point and cursed each misplay. When the CEO scored the final goal in the waning moments of play, the programmer demanded a rematch, which would certainly happen when time, tide, and the demands of the software business calendar allowed.

David Alan Grier is an IEEE Fellow and author of the forthcoming book The Company We Keep. Contact him at grier@gwu.edu or on Twitter @ dagrier.

Erin Dian Dumbacher is a research director and consultant in Washington, DC. Contact her at erin. dumbacher@fulbrightmail.org or follow her on Twitter @erin_dian.

CN Selected CS articles and columns are available for free at http://ComputingNow.computer.org.



IEEE 9th World Congress on Services (SERVICES 2013)

June 27—July 2, 2013, Santa Clara Marriott, CA, USA (center of Silicon Valley) Federation of 5 Theme Topic Conferences on Services from Different Angles (<u>http://www.servicescongress.org</u>)



IEEE 6th International Conference on Cloud Computing (CLOUD 2013)

Cloud Computing is becoming a scalable services delivery and consumption platform in the field of Services Computing. The technical foundations of Cloud Computing include Service-Oriented Architecture and Virtualizations. Major topics cover Infrastructure Cloud, Software Cloud, Application Cloud, Social Cloud, & Business Cloud. Visit <u>http://thecloudcomputing.org</u>.

IEEE 20th International Conference on Web Services (ICWS 2013)

ICWS 2013 will feature data-centric, web-based



services modeling, design, development, publishing, discovery, composition, testing, QoS assurance, adaptation, and delivery technologies and standards. Visit <u>http://icws.org</u>.

IEEE 2nd International Conference on Mobile Services (MS 2013)

MS 2013 will feature all aspects of mobile services including modeling, construction, deployment, mid-



dleware, and user experience with a special emphasis on contextawareness in social settings. Visit <u>http://themobileservices.org/2013</u>.

IEEE 10th International Conference on Services Computing (SCC 2013)

SCC 2013 will focus on services innovation lifecy-



cle e.g., enterprise modeling, business consulting, solution creation, services orchestration, optimization, management, business process integration and management. Visit http://conferences.computer.org/

IEEE 2nd International Conference on Services Economics (SE-BigData 2013) SE 2013 will focus on quantitative analysis of impact on service financial and enterprise oper-



ational outcome from BigData initiative, including outcome metrics identification, risk assessment and trade-off of IT solution selection and optimization of IT expense. Visit http://ieeese.org/2013.

Sponsored by IEEE Technical Committee on Services Computing (TC-SVC, http://tab.computer.org/tcsc)

Conference proceedings are El indexed. Extended versions of invited ICWS/SCC/CLOUD/MS/SE papers will be published in IEEE Transactions on Services (TSC, SCI & El indexed), International Journal of Web Services Research (JWSR, SCI & El indexed), International Journal of Business Process Integration and Management (IJBPIM), and IEEE IT Pro (SCI & El Indexed).



Submission Deadlines

ICWS 2013: 1/28/2013 CLOUD 2013: 1/28/2013 SCC 2013: 2/11/2013 MS 2013: 2/11/2013 SE 2013: 3/1/2013 SERVICES 2013: 3/1/2013

Contact: Liang-Jie Zhang (LJ) zhanglj@ieee.org (Steering Committee Chair)



Fuel your imagination

The IEEE Member Digital Library gives you the latest technology research—so you can connect ideas, hypothesize new theories and invent better solutions.

Get full-text access to the IEEE *Xplore*[®] digital library—at an exclusive price—with the only member subscription that includes any IEEE journal article or conference paper.

Choose from two great options designed to meet the needs of every IEEE member:

IEEE Member Digital Library

Designed for the power researcher who needs a more robust plan. Access all the IEEE content you need to explore ideas and develop better technology.

- 25 article downloads every month
- Only US\$35 per month

IEEE Member Digital Library Basic

Created for members who want to stay up-to-date with current research.

- 3 new article downloads every month
- Rollover unused downloads for 12 months
- Only US\$15 per month



Get the latest technology research.

Try the IEEE Member Digital Library —and get your FIRST MONTH FREE!

www.ieee.org/go/freemonth



IEEE Member Digital Library is an exclusive subscription available only to active IEEE members.