

## **MARGO I. SELTZER**

*Harvard School of Engineering and Applied Sciences  
33 Oxford Street  
Cambridge, MA 02138  
Office Phone: (617)496-5663  
Fax: (617)496-5508  
Phone: (617)308-4872  
Email: [margo@eecs.harvard.edu](mailto:margo@eecs.harvard.edu)  
URL: <http://www.eecs.harvard.edu/~margo>*

### **Research Interests**

My research interests are in operating systems, file systems, provenance, and healthcare informatics (and anything else that strikes my fancy).

### **Education**

- 1988 - 1992 Ph.D., Computer Science, University of California at Berkeley.  
Dissertation: "File System Performance and Transaction Support." Advisor: Michael Stonebraker. Minors in Computer Theory and Japanese.
- 1979 - 1983 A.B., Applied Mathematics, Computer Science concentration, Magna cum Laude, Harvard/Radcliffe College.

### **Experience**

- 2015 - present Faculty Director, Center for Research on Computation and Society
- 2004 - present Herchel Smith Professor of Computer Science, Harvard University,  
Conducts research in many areas of computer systems, such as file systems, transaction systems, data provenance, graph analytic engines and databases, and healthcare informatics. Also conducting research in instructional pedagogy and the retention of women in computer science. Teaches undergraduate and graduate courses in operating systems and database systems. Also collaborates with Harvard Business School faculty in teaching courses on Innovating in Health Care, both on campus and as a combination MOOC/SPOC via EdX.
- 2006 - present Architect, Oracle Corporation, Redwood Shores, CA.  
One of two architects of the Berkeley DB product line.
- 2002 - 2006 Associate Dean for Computer Science and Engineering, Harvard University,  
Responsible for overall administration of Computer Science and Electrical Engineering including curricular planning, Jr. faculty recruiting and mentoring, departmental communication, and Industrial outreach.
- 2005 - 2010 Harvard College Professor, Harvard University,
- 2000 - 2004 Gordon McKay Professor of Computer Science, Harvard University,
- 1997 - 2000 Associate Professor of Computer Science, Harvard University,
- 1992 - 1997 Assistant Professor of Computer Science, Harvard University, Cambridge, MA.  
Responsible for teaching introductory computer science and graduate and undergraduate database and

operating systems. Conducting research in file systems, performance analysis, new operating system architectures, web application architectures and database systems.

- 1996 - 2006 Chief Technical Officer, Sleepycat Software, Lincoln, MA.  
Responsible for system architecture of a fully recoverable transactional database storage system. Personally responsible for hash access method, and the locking, transaction and recovery systems.
- 1988 - 1992 Research Assistant, University of California, Berkeley, CA.  
Investigated file system allocation strategies, performance, and transaction support. Using both simulation and implementation, evaluated the performance of read-optimized and write-optimized file systems across a variety of workloads representing transaction processing applications, supercomputer applications and time sharing environments. Designed, implemented, and evaluated a new log-structured file system for distribution with 4.4BSD.
- 1988 - 1992 Manager of Software Development, Quantum Consulting, Berkeley, CA.  
Managed a team of programmers responsible for the design, implementation and support of Microsoft Windows'-based graphical user interfaces for economic models.
- 1988 Teaching Assistant for Graduate Computer Architecture, U.C. Berkeley. Responsible for organizing course materials, preparing and grading assignments and exams, answering student questions, and advising students on final research projects.
- 1988 Senior Programmer, Silicon Compiler Systems, San Jose, CA.  
Surveyed existing database management systems and devised a strategy for future CAD database development. Implemented an object manager for next generation product.
- 1986 - 1987 Senior Engineer, Kendall Square Research Corp., Cambridge, MA.  
Responsible for VLSI logic design and simulation of full custom proprietary RISC CPU. Participated in design of instruction set architecture and interchip communication protocols. Also responsible for tape out of Kendall Square's first chip including DRC, ERC, layout versus schematic verification and MEBES generation. Designed and managed the implementation of automated CAD environment.
- 1985 - 1986 Senior Programmer, Stratus Computer Corp., Marlboro, MA.  
Member of the file system and transaction processing group. Designed and implemented new queuing mechanisms to provide increased concurrency in the operating system. Implemented performance enhancements to the transaction processing subsystem and performance metering facilities in the file system.
- 1983 - 1985 Senior Engineer, Sequoia Systems, Inc., Marlboro, MA.  
Member of the data management group. Responsible for design and implementation of Sequoia's indexed file system with transaction support. Participated in system design as well as design of buffer management support in the operating system. Ported the INGRES relational database system to the Sequoia Hardware.
- 1980-1983 Core Computer Requirement Instructor, Harvard University, Cambridge, MA.  
Taught two week introductory programming courses to enable students to pass the University Computer Requirement. Served as program administrator for one year. Responsible for writing and grading exams, giving introductory lectures each semester for 800 students, and administering tests to 1600 freshmen.
- 1980-1983 Teaching Fellow, Harvard University, Cambridge, MA.  
Teaching fellow for "Introduction to Computer Science," "Structure and Interpretation of Computer Programming," and "Introduction to Systems' Programming," at Harvard University and Harvard Summer School. Head Teaching Fellow for two semesters. Responsible for hiring other teaching fellows, establishing grading criteria, and preparing assignments and exams, as well as teaching weekly sections and grading assignments and exams.

**PhD Students Graduated**

- Christopher Small 1998
- Ellie Baker 1999
- Yasuhiro Endo 2000
- Keith Smith 2001
- Xiaolan Catherine Zhang 2002
- Kostas Maoutis 2003
- David Sullivan 2003
- Dan Ellard 2004
- Robert Fischer 2004
- Griffin Weber 2005
- Alexandra (Sasha) Fedorova 2006
- Christopher Alexander Stein 2007
- Jonathan Ledlie 2007
- Jeffrey Shneidman 2008
- Kiran-Kumar Muniswamy-Reddy 2010
- Chaki Ng 2011
- Jason Waterman 2012
- Uri Braun 2014
- Elaine Angelino 2014
- Peter Macko 2015

**Awards**

- ACM Fellow, 2011.
- Capers and Marion McDonald Award for Excellence in Mentoring and Advising, 2010.
- Best Paper, 2002 USENIX Technical Conference.
- Nominated to the National Academy of Engineering's Celebration of Women in Engineering, 1999.
- IBM Faculty Award, 2002, 2004, 2005

- Roslyn Abramson Teaching Award, 1999.
- Best Student Paper, 1999 USENIX NT Symposium.
- Phi Beta Kappa Prize for Excellence in Teaching, 1996.
- Radcliffe Junior Faculty Fellowship, 1996 - 1997.
- Best Student Paper, 1996 USENIX Technical Conference.
- Award Paper, 1995 Symposium on Operating Systems.
- Best Student Paper, 1995 USENIX Technical Conference.
- Sloan Foundation Fellowship in Computer Science, 1994 - 1996.
- 1993 USENIX Lifetime Achievement Award (Contributor).
- Best Presentation, 1993 Winter USENIX Conference.
- Best Student Paper, 1991 Winter USENIX Conference.
- Best Student Paper, 1990 Winter USENIX Conference.
- University of California Microelectronics Scholarship, 1988.
- Elizabeth Cary Agassiz Scholarship (for grade point average).
- John Harvard Scholarship (for grade point average).

#### **University Service**

- HarvardX Research Committee (2013-present)
- SEAS Resource Advisory Committee (2013-present)
- IACS Advisory/Program Committee (2011-present)
- SEAS Entrepreneurship Committee (2013-2015)
- SEAS Graduate Admissions Diversity Committee (2012-present; chair 2014-present)
- President's Challenge Judge Co-Chair (2014-2015)
- FAS Committee on the Status of Women (2000-2014)
- Student Information Systems Advisory Committee (2012-2013)
- Computer Science Curriculum Sub Committee (chair 2009-2010)
- Speaker for the Women's Leadership Project (3/2009)
- FAS IT Committee (2009-2010)

- PRISE Guest Lecturer (2008)
- Faculty Chair for Reaccreditation (2008-2009)
- SEAS Committee on Higher Degrees (2007-2012, chair 2010-2011)
- FAS Committee on Pedagogical Improvement (2007-2010, chair 2008-2009)
- FAS Committee on Writing and Speaking (2007-2008)
- FAS Committee on Activity-Based Learning (Chair: 2007-2008)
- FAS Educational Policy Committee (2003-2006, 2007-2009)
- Provost's Committee on Technology Transfer (2003-2006)
- Provost's Committee on Science (2002-2003)
- Provost's Committee on the Life Sciences (2002-2003)
- Provost's Committee on Computational Biology (2002-2003)
- Faculty Advisor to the Student Entrepreneurship Counsel (2003-2006)
- Committee on Higher Degrees in the Joint Information Technology PhD Program (2002-2006)
- DEAS Dean's Advisory Group (2002-present)
- FAS Standing Committee on the Library (2000 - 2001).
- Board of Freshman Advisors (2000-2001, 2002-2003, 2006-present).
- Reunion Speaker for the 55th Reunion Class, 1999.
- Faculty Council (1998-1999).
- Committee on Undergraduate Education (1998-1999).
- Harvard Hoopes Selection Committee (1993, 1995, 1996).
- Herchel Smith Harvard Scholarship Selection Committee (1996).
- GSAS Alumni Day Presentation (1993, 1996).
- DEAS Building Committee (1996).
- DEAS Committee on Teaching Practices (1995 - 1996).
- DEAS Admissions Committee 1994-1996, 1998-2001.
- DEAS Graduate Orientation 1997-1999.
- University Committee on Rights and Responsibilities (1995 - 2001).

- FAS Standing Committee on Information Technology (1993 - 2002).
- Faculty advisor for the Harvard University Band (1994-present).
- Faculty advisor to the Harvard chapter of the Society of Women Engineers.
- Frequent participant in development office panels on the role of information technology in undergraduate education.
- Radcliffe Science Alliance Program (1993, 1994). Freshman week panel on Women in Science (1996-1998).
- Coordinated gatherings for undergraduate Computer Science Concentrators.
- Women in Science at Harvard/Radcliffe speaker.
- Panelist on the Academic Panel for the Women's Leadership Conference (1999, 2000, 2002).
- Panelist for orientation of the Radcliffe Bunting Institute (2000).

#### **Invited Talks**

- Automatically Scalable Computation, RICON East Keynote (2013).
- Provenance in Digital Collections, Library of Congress (2012).
- Provenance Everywhere. Distinguished Lecturer, Brown University (2011), University of Massachusetts, Amherst (2011), Simon Fraser University (2011), Distinguished Lecturer University of Chicago (2010), Grace Murray Hopper Speaker, University of Pennsylvania (2010).
- Provenance Everywhere, MIT (2011).
- A Day without Provenance is a like a day without Sunshine (HPTS, 2011).
- Speaker in the E-Science Institute Symposium on Provenance in Secure and Advanced Systems (5/09).
- Invited speaker to the eScience Institute Provenance Workshop (5/09).
- Panelist for Tufts Engineering Panel on Engineering Leadership (4/09).
- Mentoring and Management: SEAS Professional Development Seminar (4/09).
- Panelist for Berkman Center's Gender and Technology Workshop (3/09).
- Invited Speaker for the First Workshop on the Theory and Practice of Provenance (2/09).
- Riding the Tiger versus Petting the Cat: Different models for software startups, Invited Speaker for the Harvard Computer Society Technology Career Conference (12/08).
- Provenance-Aware Storage Systems, Florida International University (1/08), Northwestern University (10/07), State University of New York, Stonybrook (4/06).
- The Cat's Eye View: Building a Business on Open Source, Sun Microsystems First Friday Talk (4/2006).

- Doing It All, National Symposium on the Advancement of Women In Sciences (4/05).
- Princeton Long-range planning workshop, "Recruiting and Retaining Junior Faculty," Princeton, NJ, November 2003.
- Distinguished Lecture Series, University of Wisconsin, Madison Hourglass: An Architecture for Pervasive Applications (3/03).
- CRAW Distinguished Lecture Series, University of Toronto, Operating System Extensibility to Improve Application Performance and Functionality (12/00).
- Invited Talk on Embedded Database Systems, 1999 USENIX Technical Conference (6/99).
- Invited Panelist on Embedded Database Systems, 1999 SIGMOD Conference (6/99).
- Invited Talk on Measuring Computer Systems, 1997 USENIX Technical Conference (1/97).
- Invited Panelist on The Source Dilemma, 1997 USENIX Windows NT workshop.
- Invited Speaker for the CSTB Workshop on Information Survivability, "Why Safe Languages and Java Don't Matter," (2/97)
- PCTV UNIX Series, UNIX Architecture and Philosophy, (aired 9/96).
- Invited Panelist for CRA Workshop on Academic Careers for Women (11/94).
- Invited speaker to the President's Council of Advisors on Science and Technology (6/92).
- Distinguished Lecturer for the 1994, 1995, 1996, 1997 Research Summer Institute Program at MIT.
- Invited Speaker on Information Technology, Beneficial Board of Directors (8/95).

### **Talks**

- Data Modeling for Non-Relational Data, Dr. Dobbs Architecture World (7/08).
- Erlang + EDTK + BDB: Disrupting the Web, HPTS (10/07).
- Matching Database Needs with Database Solutions, eBay (3/05), Juniper Systems (3/05).
- An Architecture A Day Keeps the Hacker Away, IBM (3/05).
- Provenance Aware Storage Systems, IBM (3/06), EMC (1/06), Network Appliance(4/05), Sun Microsystems (1/05).
- Berkeley DB, Veritas Software (5/04), Yahoo (3/04).
- Non-Relational Data Management, IBM (5/04).
- Hourglass: An Architecture for Pervasive Applications, Princeton University (11/02), Tufts University (2/04), Amherst College (9/04).

- Issues and Challenges for the WWW, Lucent Technologies (5/98) Lotus Corporation (3/97).
- Issues and Challenges in Extensible Operating Systems, Cornell University (3/98), Intel Corporation (6/97), University of Wisconsin, Madison (4/97), Hewlett-Packard Laboratories (4/97).
- The VINO Project, HP Research Laboratories, IBM Almaden Research Center (1/97).
- Application Performance in the VINO Operating System, Bunting Institute (12/96).
- The VINO Operating System, Dartmouth College (11/96).
- Clustering, Logging, and Journaling: Approaches to High Performance File System Design, Depauw University (7/96).
- Panelist on Research and Education at the 1996 NSF Infrastructure PI Workshop (7/96).
- World Wide Web Developments, Lotus Corporation (3/97), IBM Almaden (7/96), World Wide Web Consortium Workshop (4/96).
- The VINO Project, Sun Eastern Research Lab (9/95), Novell Corporation (1/96).
- The Berkeley Fast File System, Clam Associates (3/95).
- Approaches to High Performance File Systems, MIT (2/95), Sun Eastern Research Lab (3/95).
- The Harvard Ubiquitous Information Project, Xerox PARC (5/93).
- File systems and transaction support, Digital Equipment Corporation (11/93), Tandem Computers (5/93), MIT (5/93), Boston University (4/93), IBM Almaden Research Center (1/93), Auspex Computer (1/93), Harvard University (11/91), Hewlett Packard (6/91), Sequent Computer Corp. (2/91), Portland State University (2/91).

### **Professional Activities**

- NSF PROBE Steering Committee (2012 - present)
- Sloan Selection Committee (2011-2017).
- Acting Executive Director, USENIX Association (2011-2012)
- Nominating Committee of the CCC (2010-2012, Chair 2012).
- Member of the Computing Community Consortium of the CRA (2010-2013, Executive Committee 2012-2013).
- President of the USENIX Association, 2012-2014; Acting Executive Director 2011-2012; Vice-President, 2008-2012; Director at Large, 2006-2008, 1996 - 1998. Chair of the scholastic committee, responsible for variety of student programs including the student research projects program.
- Program Co-Chair for the 2010 Workshop on Theory and Practice of Provenance (TaPP 2010).
- Program Chair for the 2009 High Performance Transaction Processing Workshop (HPTS 2009).
- Steering Committee for the USENIX Workshop on Theory and Practice of Provenance (2009-present).



- Program Co-Chair for the 2009 File and Storage Technologies Symposium (FAST'09).
- CMU Visiting Committee (2006).
- Program Chair for the 2005 Workshop on Hot Topics in Operating Systems (HotOS'05).
- Program Co-Chair for the 2004 ACM Conference on Electronic Commerce.
- National Academy Panel on Digital Archiving (2002-2004).
- Steering Committee for the USENIX Symposium on File and Storage Technologies (2002-present).
- Steering Committee for the USENIX Symposium on Network Design and Implementation (2006-2009).
- Program Co-Chair for the 1998 Symposium on Operating System Design and Implementation (OSDI 1998).
- Member of the HTTP-NG Working Group on Web Characterization, 1998.
- Steering Committee for the USENIX Workshop on NT.
- Awards Chair for the 1997 International Conference on Data Engineering.
- Program Co-Chair for 1994 Summer USENIX Conference.
- Program committee member for the 1992 Summer USENIX, 1993 Workshop on Workstation Operating Systems, 1994 Winter USENIX, 1994 ASPLOS conference, 1994 SIGMOD conference, 1995 Symposium on Operating System Principles, 1996 International Workshop on Object Orientation in Operating Systems, 1997 Workshop in Internet Systems and Technologies, 2001 USENIX Technical Conference, 2001 Symposium on Operating System Principles, 2003 USENIX File And Storage Technologies Symposium, 2003 Workshop on Hot Topics in Operating Systems, 2004 ACM Conference on Electronic Commerce, 2004 Symposium on Operating System Design and Implementation, 2004 ACM Conference on High Performance Distributed Computing, 2008 USENIX File and Storage Technologies Symposium, 2011 Symposium on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2011 Workshop on the Theory and Practice of Provenance (TaPP), 2011 High Performance Transaction Processing (HPTS), 2011 Workshop on Hot Topics in Security (HotSec), 2011 Symposium on Operating System Principles (SOSP), 2012 Symposium on File and Storage Technology (FAST), 2012 Workshop on Hot Topics in Storage Systems (HotStorage), 2013 Symposium on File and Storage Technology(FAST), 2013 ACM International Systems and Storage Conference (SYSTOR), 2013 High Performance Transaction Processing Workshop (HPTS), 2014 Symposium on File and Storage Technology (FAST), 2015 Hot Topics in Operating Systems (HotOS).
- Referee for IEEE Computer, IEEE Software Practice and Experience, IEEE Transactions on Computer Systems, IEEE Parallel and Distributed Technology: Systems and Applications, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Knowledge and Data Engineering, VLDB Journal, Digital Systems Journal, Communications of the ACM, and Algorithmica.
- Reviewer for 1992, 1996 ASPLOS Conference, 1993 SOSP Conference, 1994 ISCA. 1994 Workshop on Operating System Design and Implementation, 1995 USENIX Technical Conference, 1996 Symposium on Operating System Design and Implementation, 1997 International Symposium on Computer Architecture, 1997 Symposium on Internet Systems and Technologies, 1999 International Conference on Distributed Computing Systems, 1999 Symposium on Operating System Principles, 2001 Hot Topics in Operating Systems Workshop.
- Book reviewer for Addison-Wesley, Prentice Hall, Kluwer Associates and John Wiley & Sons, Inc.

- Board of Directors, Blue Leaf (2009-2011).
- Chairman of the Sleepycat Software Board of Directors (1996-2006).
- Member of the Traxit Board of Directors (2000-2001)
- Member of the Marble Technical Advisory Board and Board of Directors (1996-1997).
- Member IEEE, ACM, USENIX Association.
- Contributor to 4.4BSD Project.
- Nominating Committee for the USENIX Board of Directors, 1994, 1999.

### Refereed Publications

Margo, D., Seltzer, M., “A Scalable Distributed Graph Partitioner,” Proceedings of the 41st International Conference on Very Large Databases (VLDB), Kohola Coast, HA, September 2015.

Eldridge, S., Appavoo, J., Joshi, A., Waterland, A., Seltzer, M., “Towards General-Purpose Neural Network Computing,” Proceedings of the 13th International Conference on Parallel Architectures and Compilation Techniques (PACT 2015), Petrozavodsk Russia, September 2015.

Balakrishnan, N., Bytheway, T., Carata, L., Chick, O., Snee, J., Akoush, S., Sohan, R., Seltzer, M., Hopper, A., “Recent Advances in Computer Architecture: The Opportunities and Challenges for Provenance,” Proceedings of the 7th Workshop on the Theory and Practice of Provenance (TaPP 2015), Edinburgh Scotland, July 2015.

Macko, P., Margo, D., Marathe, V., Seltzer, M., “LLAMA: Efficient Graph Analytics Using Large Multiversioned Arrays,” Proceedings of the 31st IEEE International Conference on Data Engineering (ICDE 2015), Seoul Korea, April 2015.

Appavoo, J., Waterland, A., Eldridge, S., Zhao, K., Joshi, A., Homer, S., Seltzer, M., “Programmable Smart Machines: A Hybrid Neuromorphic approach to General Purpose Computation,” Proceedings of the first workshop on Neuro-morphic Architectures (NeuroArch) collocated with ISCA 2014.

Angelino, E., Kohler, E., Waterland, A., Seltzer, M., Adams, M., “Accelerating MCMC via parallel predictive prefetching,” Proceedings of the 2014 Conference on Uncertainty in Artificial Intelligence (UAI '14), Quebec City, Quebec Canada, July 2014.

Rao, M., Parkes, D., Seltzer, M., and Bacon, D., “A Framework for Incentivizing Deep Fixes,” Proceedings of the AAAI 2014 Workshop on Incentives and Trust in E-Communities (WIT-EC'14), Quebec City, Quebec Canada, July 2014.

Carata, L., Akoush, S., Balakrishnan, N., Bytheway, T., Sohan, R., Seltzer, M., and Hopper, A., 2014. A primer on provenance. *Commun. ACM* 57, 5 (May 2014), 52-60. DOI=10.1145/2596628 <http://doi.acm.org/10.1145/2596628>.

Waterland, A., Angelino, E., Adams, R., Appavoo, J., Seltzer, M., “ASC: Automatically Scalable Computation” Proceedings of the 2014 Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '14), Salt Lake City Utah, March 2014.

Herzlinger, R., Seltzer, M., Gaynor, M., “Applying KISS to Healthcare Information Technology,” *Computer* 46(11), November 2013, 72--74.

Borkin, M., Yeh, C., Boyd, M., Macko, P., Gajos, K., Seltzer, M., Pfister, H., “Evaluation of Filesystem Provenance Visualization Tools,” Proceedings of the 2013 Conference on Information Visualization (Vis 2013), Atlanta Georgia, October 2013.

Macko, P., Margo, D., Seltzer, M., “Local Clustering in Provenance Graphs,” Proceedings of the 2013 International Conference on Information and Knowledge Management (CIKM 2013), Burlingame CA, October 2013.

- Waterland, A., Angelino, E., Cubuk, E., Kaziras, E., Adams, R., Appavoo, J., Seltzer, M., "Computational Caches," Proceedings of the 2013 SYSTOR Conference, Haifa Israel, July 2013.
- Macko, P., Margo, D., Seltzer, M., "Performance Introspection of Graph Databases," Proceedings of the 2013 SYSTOR Conference, Haifa Israel, July 2013.
- Holland, D., Angelino, E., Wald, G., Seltzer, M., "Flash Caching on the Storage Client," Proceedings of the 2013 USENIX Annual Technical Conference, San Jose, CA, June 2013.
- Guo, P., Seltzer, M., "Burrito: Wrapping Your Lab Notebook in Computational Infrastructure," Proceedings of the Fourth Workshop on the Theory and Practice of Provenance (TaPP 2012), Boston MA, June 2012.
- Macko, P., Seltzer, M., "A General-Purpose Provenance Library," Proceedings of the Fourth Workshop on the Theory and Practice of Provenance (TaPP 2012), Boston MA, June 2012.
- Waterland, A., Appavoo, J., Seltzer, M., "Parallelization by Simulated Tunneling," Proceedings of the Fourth USENIX Conference on Hot Topics in Parallelism (HotPar 2012), Berkely, CA, June 2012.
- Macko, P., Seltzer, M., "Provenance Map Orbiter: Interactive Exploration of Large Provenance Graphs," Proceedings of the Third Workshop on the Theory and Practice of Provenance (TaPP 2011), Heraklion Greece, June 2011.
- Angelino, E., Braun, U., Holland, D., Macko, P., Margo, D., Seltzer, M., "Provenance Integration Requires Reconciliation," Proceedings of the Third Workshop on the Theory and Practice of Provenance (TaPP 2011), Heraklion Greece, June 2011.
- Macko, P., Chiarini, M., Seltzer, M., "Collecting Provenance via the Xen Hypervisor," Proceedings of the Third Workshop on the Theory and Practice of Provenance (TaPP 2011), Heraklion Greece, June 2011.
- Holland, D., Seltzer, M., "Multicore OSES: Looking Forward from 1991, er, 2011," Proceedings of the 2011 Workshop on Hot Topics in Operating Systems (HOTOS XIII), May 2011, Napa, CA.
- Tarasov, V., Bhanage, S., Zadok, E., Seltzer, M., "Benchmarking File System Benchmarking: It \*IS\* Rocket Science," Proceedings of the 2011 Workshop on Hot Topics in Operating Systems (HOTOS XIII), May 2011, Napa, CA.
- Bailis, P., Reddi, V., Gandhi, S., Brooks, D., Seltzer, M., "Dimetrodon: Processor-level Preventive Thermal Management via Idle Cycle Injection," Proceeding of the 2011 Design Automation Conference (DAC 2011), San Diego, CA, June 2011.
- Maclean, D., Seltzer, M., "Mining the Web for Medical Hypotheses: A Proof-of-Concept System," Proceedings of the International Conference on Health Informatics, Rome Italy, January 2011.
- Macko, P., Seltzer, M., Smith, K.A., "Tracking Back References in a Write-Anywhere File System," Proceedings of the 8th Conference on File and Storage Technologies (FAST'10), San Jose, CA, February 2010.
- Muniswamy-Reddy, K., Macko, P., Seltzer, M., "Provenance for the Cloud," Proceedings of the 8th Conference on File and Storage Technologies (FAST'10), San Jose, CA, February 2010.
- "Towards Query Interoperability: PASSing PLUS," Braun, U., Seltzer, M., Chapman, A., Blaustein, B., Allen, M.D., Seligman, L., Proceedings of the 2nd Workshop on the Theory and Practice of Provenance (TaPP'10) San Jose, CA, February 2010.
- Muniswamy-Reddy, K. and Seltzer, M. 2010. Muniswamy-Reddy, K. and Seltzer, M. "Provenance as first class cloud data." ACM SIGOPS Operating Systems Review, 43, 4 (Jan. 2010), 11-16.
- Muniswamy-Reddy, K., Seltzer, M., "Provenance as First-Class Cloud Data," 3rd ACM SIGOPS International Workshop on Large Scale Distributed Systems and Middleware (LADIS'09), October 2009.
- Cheney, J., Chong, S., Foster, N., Seltzer, M., and Vansummeren, S. 2009. "Provenance: a future history." In Proceeding of the 24th ACM SIGPLAN Conference Companion on Object Oriented Programming Systems Languages and Applications (Orlando, Florida, USA, October 25 - 29, 2009). OOPSLA '09. ACM, New York, NY, 957-964.

Muniswamy-Reddy, K., Braun, U., Holland, D., Macko, P., Maclean, D., Margo, D., Seltzer, M., Smogor, R., "Layering in Provenance Systems," In proceedings of the 2009 USENIX Annual Technical Conference, San Diego, CA, June 2009.

Seltzer, M., Murphy, N., "Hierarchical File Systems are Dead," Proceedings of the 12th Workshop on Hot Topics in Operating Systems (HotOS'09), Monte Verita, Switzerland, May 2009.

Margo, D., Seltzer, M., "The Case for Browser Provenance," 1st Workshop on the Theory and Practice of Provenance (TaPP'09), February 2009.

Muniswamy-Reddy, K., Seltzer, M., "Making a Cloud Provenance-Aware," 1st Workshop on the Theory and Practice of Provenance (TaPP'09), February 2009.

Braun, U., Shinnar, A., Seltzer, M., "Securing Provenance," In Proceedings of the 3rd USENIX Workshop on Hot Topics in Security (HotSec), San Jose, CA, July 2008.

Holland, D., Braun, U., Maclean, D., Muniswamy-Reddy, K., Seltzer, M., "Choosing a Data Model and Query Language for Provenance," In proceedings of the 2nd International Provenance and Annotation Workshop, Salt Lake City, UT, Jun 2008.

Fedorova, A., Seltzer, M., Smith, M.D., "Improving Performance Isolation on Chip Multiprocessors via an Operating System Scheduler," Proceedings of the Sixteenth International Conference on Parallel Architectures and Compilation Techniques (PACT), Brasov, Romania, September 2007.

Stein, L., Holland, D., Seltzer, M., Zhang, Z., Proceedings of the First EuroSys Workshop on Virtualization for HPC, Lisbon, Portugal, March 2007.

Ledlie, J., Gardner, P., Seltzer, M., "Network Coordinates in the Wild," Proceedings of the Conference on Network System Design and Implementation, Boston MA, April 2007.

Holland, D., Seltzer, M., Braun, U., Muniswamy-Reddy, K., "PASS-ing the Provenance Challenge," Special Edition of Concurrency and Computation: Practice and Experience Wiley and Sons, 2007.

Muniswamy-Reddy, K., Holland, D., Braun, U., Seltzer, M., "Provenance Aware Storage Systems," Proceedings of the 2006 USENIX Annual Technical Conference, Boston, MA, June 2006.

Fedorova, A., Seltzer, M., Smith, M., "A Non-Work-Conserving Operating System Scheduler for SMT Processors," Proceedings of the Workshop on the Interaction between Operating Systems and Computer Architecture (WIOSCA), in conjunction with ISCA-33, June 2006.

Braun, U., Garfinkel, S., Holland, D., Muniswamy-Reddy, K., Seltzer, M., "Issues in Automatic Provenance Collection," Proceedings of the International Provenance and Annotation Workshop (IPAW'06), Chicago, IL, May 2006.

Brunelle, J., Hurst, P., Huth, J., Kang, L., Ng, C., Parkes, D., Seltzer, M., Shank, J., Youssef, S., "EGG: An Extensible and Economics-inspired Open Grid Computing Platform," Proceedings of the 2006 Grid Asia, Singapore, May 2006.

Pietzuch, P., Ledlie, J., Shneidman, J., Roussopoulos, M., Welsh, M., Seltzer, M., "Network-Aware Operator Placement for Stream-Processing Systems" Proceedings of the 22nd International Conference on Data Engineering (ICDE '06), Atlanta, GA, April 2006.

Pietzuch, P., Ledlie, J., Seltzer, M., "Supporting Network Coordinates on PlanetLab," Proceedings of the Second Workshop on Real, Large Distributed Systems (WORLDS'05), San Francisco, CA, December 2005.

Seltzer, M., "Beyond Relational Databases," ACM Queue, Vol 3. No 3, April 2005.

Fedorova, A., Seltzer, M., Small, C., Nussbaum, D., "Performance of Multithreaded Chip Multiprocessors and Implications for Operating System Design," Proceedings of the 2005 Annual Technical Conference, Anaheim, CA April 2005.

Shneidman, J., Pietzuch, P., Welsh, M., Seltzer, M., Roussopoulos, M., "A Cost-Space Approach to Distributed Query Optimization in Stream-Based Overlays," Proceedings of the 1st IEEE International Workshop on Networking Meets Databases (NetDB'05), Tokyo, Japan, April 2005.

- Ledlie, J., Ng, C., Holland, D., Muniswamy-Reddy, K., Braun, U., "Provenance Award Sensor Data Storage," Proceedings of the 1st IEEE International Workshop on Networking Meets Databases (NetDB'05), Tokyo, Japan, April 2005.
- Ledlie, J., Seltzer, M., "Distributed, Secure Load Balancing with Skew, Heterogeneity, and Churn," Proceedings of the 2005 INFOCOM, March 2005.
- Pietzuch, P., Shneidman, J., Ledlie, J., Welsh, M., Seltzer, M., Roussopoulos, M., "Evaluating DHT Service Placement in Stream-Based Overlays," Proceedings of the International Workshop on Peer to Peer Systems (IPTPS'05), Ithaca, NY, February 2005.
- Holland, D., Lim, A., Seltzer, M., "An Architecture A Day Keeps the Hacker Away," Proceedings of the 2004 Workshop on Architectural Support for Security and Anti-Virus," Boston, MA, April 2004.
- Ledlie, J., Shneidman, J., Welsh, M., Roussopoulos, M., Seltzer, M., "Open Problems in Data Collection Networks," *Proceedings of the 2004 SIGOPS European Workshop, September 2004, Leuven, Belgium.*
- Fedorova, A., Small, C., Nussbaum, D., Seltzer, M., "Chip Multithreading Systems Need a New Operating System Scheduler," *Proceedings of the 2004 SIGOPS European Workshop, September 2004, Leuven, Belgium.*
- Mesnier, M., Thereska, E., Ellard, D., Ganger, G., Seltzer, M., "File classification in self-\* storage systems," Proceedings of the International Conference on Autonomic Computing (ICAC-04), May 2004, New York, NY.
- Magoutis, K., Seltzer, M., Gabber, E., "The Case Against User-level Networking," *Proceedings of the Third Annual Workshop on System-Area Networks (SAN-3)*, February 2004, Madrid Spain.
- Fedorova, A., Seltzer, M., Magoutis, K., Addetia, S., "Application Performance on the Direct Access File System," *Proceedings of Workshop on Software and Performance 2004 (WOSP'04)*, January 2004, Redwood City, CA.
- Ellard, D., Seltzer, M., "New NFS Tracing Tools and Techniques for System Analysis," *Proceedings of the 17th Annual Large Installation System Administration Conference (LISA'03)*, October 2003, San Diego, CA, pp 73-85.
- Ellard, D., Seltzer, M. "NFS Tricks and Benchmarking Traps," *Proceedings of the 2003 FREENIX Technical Conference*, June 2003, San Antonio, TX, pp 101-114.
- Ellard, D., Ledlie, J., Malkani, P., Seltzer, M., "Passive NFS Tracing of Email and Research Workloads," *Proceedings of the Second Symposium on File and Storage Technologies*, March 2003, San Francisco, CA, pp 203-216.
- Magoutis, K., Fedorova, A., Addetia, S., Seltzer, M., "Making the Most Out of Direct-Access Network Attached Storage," *Proceedings of the Second Symposium on File and Storage Technologies*, March 2003, San Francisco, CA.
- Ledlie, J., Shneidman, J., Seltzer, M., Huth, J., "Scooped Again," *Proceedings of the Second International Workshop on Peer-to-Peer Systems (IPTPS-03)*, Berkeley, CA, February 2003.
- Stein, C., Tucker, M., Seltzer, M., "Building a Reliable Mutable File System on Peer-to-peer Storage," *Proceedings of the International Workshop on Reliable Peer-to-peer Distributed Systems*, Osaka, Japan, October 2002.
- Ledlie, J., Taylor, J., Serban, L., Seltzer, M., "Self-Organization in Peer-to-Peer Systems," *Proceedings of the 10th ACM SIGOPS European Workshop*, Saint-Emilion, France, September 2002.
- Magoutis, K., Addetia, S., Fedorova, A., Seltzer, M., Chase, J., Gallatin, D., Kisley, R., Wickremesinghe, R., Gabber, E., "Structure and Performance of the Direct Access File System," *Proceedings of the 2002 USENIX Technical Conference*, June 2002, Monterey, CA, 1-14.
- Ellard, D., Holland, D., Murphy, N., Seltzer, M., "On the Design of a New CPU Architecture for Pedagogical Purposes" *Proceedings of the Workshop on Computer Architecture Education*, May, 2002, Anchorage, Alaska.
- Holland, D., Lim, A., and Seltzer, M., "A New Instructional Operating System", Proceedings of the 2002 SIGCSE Conference, February 2002, Covington, KY, 111-115.
- Stein, C., Howard, J., Seltzer, M., "Unifying File System Protection," *Proceedings of the 2001 USENIX Technical Conference*, June 2001, Boston, MA, 79-90.
- Holland, D., Josephson, W., Magoutis, K., Seltzer, M., Stein, C., "Research Issues in No-Futz Computing," Proceedings of the 2001 Workshop on Hot Topics in Operating Systems (HotOS VII), Elmau Germany, May 2001.

- Zhang, X., Seltzer, M., "HBench:JGC - An Application-Specific Benchmark Suite for Evaluating JVM Garbage Collector Performance," *Proceedings of the 6th USENIX Conference on Object-Oriented Technology (COOTS 2001)*, February 2001, San Antonio, TX.
- Seltzer, M., Ganger, G., McKusick, M., Smith, K., Soules, C., Stein, C., "Logging Versus Soft Updates: Asynchronous Metadata Protection in File Systems," *Proceedings of the 2000 USENIX Technical Conference*, June 2000, San Diego, CA, 71–84.
- Sullivan, D., Seltzer, M., "Isolation with Flexibility: A Resource Management Framework for Central Servers," *Proceedings of the 2000 USENIX Technical Conference*, June 2000, San Diego, CA, 337–350.
- Wong, A., Seltzer, M., "Operating System Support for Multi-User, Remote Graphical Interaction," *Proceedings of the 2000 USENIX Technical Conference*, June 2000, San Diego, CA, 183–196..
- Endo, Y., Seltzer, M., "Improving Interactive System Performance Using TIPME," *Proceedings of the 2000 ACM Sigmetrics Conference*, June 2000, Santa Clara, CA.
- Zhang, X., Seltzer, M., "HBench:Java: An Application-Specific Benchmark for JVMs" *Proceedings of the ACM 2000 Java Grande Conference*, June 2000, San Francisco, CA.
- Wong, A., Seltzer, M., "Evaluating Windows NT Terminal Server Performance," *Proceedings of the 3rd USENIX Windows NT Symposium*, July, 1999, Seattle, WA, 145–154.
- Zhang, X., Barrientos, M., Chen, J., Seltzer, M., "HACC: An Architecture for Cluster-Based Web Server," *Proceedings of the 3rd USENIX Windows NT Symposium*, July, 1999, Seattle, WA, 155–164.
- Olson, M., Bostic, K., Seltzer, M., "Berkeley DB," *Proceedings of the 1999 FREENIX Conference*, June, 1999, Monterey, CA.
- Sullivan, D., Haas, R., Seltzer, M., "Tickets and Currencies Revisited: Extensions to Multi-Resource Lottery Scheduling," *Proceedings of the 1999 Workshop on Hot Topics in Operating Systems (HotOS VII)*, March, 1999, Rio Rico, AZ.
- Seltzer, M., Olson, M., "Challenges in Embedded Database System Administration," *Proceedings of the 1999 USENIX Workshop on Embedded Systems*, Cambridge, MA, March 1999.
- Seltzer, M., Krinsky, D., Smith, K., Zhang, X., "The Case for Application-Specific Benchmarking," *Proceedings of the 1999 Workshop on Hot Topics in Operating Systems (HotOS VII)*, March, 1999, Rio Rico, AZ.
- Ellard, D. J., Ellard, P. A., Megquier, J. M., Chen, J. B., Seltzer, M. I., "The ANT Architecture — An Architecture for CS1," *The IEEE Computer Society Technical Committee on Computer Architecture Newsletter*, February, 1999, 25–27.
- Small, C., Seltzer, M., "MiSFIT: Constructing Safe Extensible Systems," *IEEE Concurrency*, (6), 3, July–Sep 1998, 34–41.
- Baker, E., Seltzer, M., "The Mug Shot Search Problem," *Proceedings of the 1998 Vision Interfaces Conference*, Vancouver, Canada, June 1998, 421–430.
- Manley, S., Seltzer, M., "Web Facts and Fantasy," *Proceedings of the Symposium on Internet Technologies and Systems*, Monterey, CA, December 1997, 125-134.
- Endo, Y., Seltzer, M., "Measuring Windows NT--Possibilities and Limitations," *Proceedings of the First USENIX Workshop on NT*, Seattle, WA, August 1997.
- Brown, A., Seltzer, M., "Operating System Benchmarking in the Wake of *Lmbench* Case Study of the Performance of NetBSD on the Intel Architecture," *Proceedings of the 1997 Sigmetrics Conference*, Seattle, WA, June 1997, 214-224.
- Smith, K., Seltzer, M., "File System Aging," *Proceedings of the 1997 Sigmetrics Conference*, Seattle, WA, June 1997, 203-213.
- Seltzer, M., Small, C., "Self-Monitoring and Self-Adapting Operating Systems," *Proceedings of the 1997 Workshop on Hot Operating Systems*, Cape Cod, MA, May 5-6, 1997, 124-129.

- Seltzer, M., Endo, Y., Small, C., Smith, K., "Dealing with Disaster: Surviving Misbehaving Kernel Extensions," *Proceedings of the Second Symposium on Operating System Design and Implementation*, Seattle, WA, October 1996, 213-227.
- Endo, Y., Wang, Z., Chen, B., Seltzer, M. "Understanding Latency in GUI-based Applications," *Proceedings of the Second Symposium on Operating System Design and Implementation*, Seattle, WA, October 1996, 195-199.
- Seltzer, M., Small, C., Smith, M., "Symbiotic System Software: Fast Operating Systems for Fast Applications," *Proceedings of the First Workshop on Compiler Support for System Software*, Tuscon, AZ, Feb 1996.
- Gwertzman, J., Seltzer, M., "World-Wide Web Cache Consistency," *Proceedings of the 1996 USENIX Technical Conference*, San Diego, CA January 1996, 141-152.
- Small, C., Seltzer, M., "A Comparison of OS Extension Technologies," *Proceedings of the 1996 USENIX Technical Conference*, San Diego, CA January 1996, 41-54.
- Smith, K., Seltzer, M., "A Comparison of FFS Disk Allocation Policies," *Proceedings of the 1996 USENIX Technical Conference*, San Diego, CA January 1996, 15-26.
- Chen, B., Endo, Y., Chan, K., Mazieres, D., Dias, A., Seltzer, M., Smith, M., "The Measured Performance of Personal Computer Operating Systems," *Proceedings of the 15th ACM Symposium on Operating Systems Principles*, Copper Mountain, Colorado, December 3-6, 1995. Also in *ACM Transactions on Computer Systems*, January 1996, 3-40.
- Small, C., Seltzer, M., "Structuring the Kernel as a Toolkit of Extensible, Reusable Components," *Proceedings of the International Workshop on Object Orientation in Operating Systems*, Lund, Sweden, August 1995, 134-137.
- Gwertzman, J., Seltzer, M., "The Case for Geographic Push-Caching," *Proceedings of the Fifth Annual Workshop on Hot Operating Systems*, Orcas Island, WA, May 1995, 51-55.
- Seltzer, M., Smith, K., Balakrishnan, H., Chang, J., McMains, S., Padmanabhan, V., "File System Logging versus Clustering: A Performance Comparison," *Proceedings of the 1995 Winter USENIX Technical Conference*, New Orleans, LA, January 1995, 249-264.
- Blackwell, T., Harris, J., Seltzer, M., "Heuristic Cleaning Algorithms for Log-Structured File Systems," *Proceedings of the 1995 Winter USENIX Technical Conference*, New Orleans, LA, January 1995, 277-288.
- Blackwell, T., Chan, K., Chang, K., Charuhas, T., Karp, B., Kung, H. T., Lin, D., Morris, R., Seltzer, M., Smith, M., Young, C., Bahgat, O., Chaar, M., Chapman, A., Depelteau, G., Grumble, K., Huang, S., Hung, P., Kemp, M., McLaughlin, I., Ng, T., Vincent, J., and Watchor, J., "An Experimental Flow Controlled Multicast ATM Switch," *Proceedings of the First Annual Conference on Telecommunications Research and Development in Massachusetts*, Vol. 6, October 1994, 33-38.
- Baker, E., Seltzer, M., "Evolving Line Drawings," *Proceedings of Graphics Interface '94*, Banff, Canada, May 1994, 91-100.
- Seltzer, M., "Transaction Support in a Log-Structured File System," *Proceedings of the Ninth International Conference on Data Engineering*, Vienna, Austria, April 1993, 503-510.
- Seltzer, M., Bostic, K., McKusick, M., Staelin, C., "An Implementation of a Log-Structured File System for UNIX," *Proceedings of the 1993 Winter USENIX Conference*, San Diego, CA January 1993, 307-326.
- Baker, M., Asami, S., Deprit, E., Ousterhout, J., Seltzer, M., "Non-Volatile Memory for Fast, Reliable File Systems," *Proceedings of ASPLOS-V*, Boston, MA, October 1992, 10-22.
- Seltzer, M., Olson, M., "LIBTP: Portable, Modular Transactions for UNIX," *Proceedings 1992 Winter USENIX Conference*, San Francisco, CA, January 1992, 9-26.
- Seltzer, M., Stonebraker, M., "Read Optimized File Systems: A Performance Evaluation," *Proceedings 7th Annual International Conference on Data Engineering*, Kobe, Japan, April 1991, 602-611.
- Seltzer, M., Yigit, O., "A New Hashing Package for UNIX," *Proceedings 1991 Winter USENIX Conference*, Dallas, TX, January 1991, 173-184.
- Seltzer, M., Stonebraker, M., "Transaction Support in Read Optimized and Write Optimized File Systems," *Proceedings 16th International Conference on Very Large Data Bases*, Brisbane, Australia, August 1990, 174-185.

Seltzer, M., Chen, P., Ousterhout, J., "Disk Scheduling Revisited," *Proceedings 1990 Winter USENIX Conference*, Washington, D.C., January 1990, 313-324.

### Invited Papers

- Gwertzman, J., Seltzer, M., "People, Places, and Things: The Next Generation Web," *Proceedings of the 1996 Computer Conference of the IEEE (CompCON)*, Santa Clara, CA, February 1996, 65-70.

### Refereed Poster Presentations

Manley, S., Courage, M., Seltzer, M., "A Self-Scaling and Self-Configuring Benchmark for Web Servers," Presented at the 1998 Sigmetrics Conference, Madison, WI, June 1998.

Gwertzman, J., Seltzer, M., "Geographical Push-Caching," Presented at the 1995 Symposium on Operating System Principles, Copper Mountain, CO, December 1995.

### Unrefereed Publications

Seltzer, M., "Building Highly-Available and Scalable Applications with Berkeley DB," Oracle Corporation White Paper, <http://www.oracle.com/technologies/embedded/docs/building-ha-scalable-applications-with-berkeley-db-whitepaper.pdf>, Oracle Corporation, August 2009.

Seltzer, M., "Berkeley DB: A Retrospective," *IEEE Data Engineering Bulletin*, 30(3): 21-28 (2007).

Seltzer, M., Ganger, G., McKusick, M. K., Smith, K., Soules, C., Stein, C., "Logging versus Soft Updates: Asynchronous Meta-data Protection in File Systems," Harvard University Computer Science Technical Report TR-07-99, December 1999.

Sullivan, D., Seltzer, M., "A Resource Management Framework for Central Servers," Harvard University Computer Science Technical Report TR-13-99, December 1999.

Wong, A., Seltzer, M., "Operating System Support for Multi-User, Remote, Graphical Interaction," Harvard University Computer Science Technical Report TR-14-99, December 1999.

Endo, Y., Seltzer, M., "Improving Interactive System Performance Using TIPME," Harvard University Computer Science Technical Report TR-09-99, December 1999.

Seltzer, M., Endo, Y., Small, C., Smith, K., "Issues in Extensible Operating Systems," Harvard University Computer Science Technical Report TR-18-97, November 1997.

Manley, S., Courage, M., Seltzer, M., "A Self-Scaling, Self-Configuring Web Benchmark," Harvard University Computer Science Technical Report TR-17-97, November 1997.

Small, C., Ghosh, N., Saleeb, H., Seltzer, M., Smith, K., "Does Systems Research Measure Up?," Harvard University Computer Science Technical Report TR-16-97, November 1997.

Smith, K., Seltzer, M., "File Layout and File System Performance," Harvard University Technical Report TR-35-94, December 1994.

Endo, Y., Gwertzman, J., Seltzer, M., Small, C., Smith, K., Tang, Diane "VINO: The 1994 Fall Harvest," Harvard University Technical Report TR-34-94.

Small, C., Seltzer, M., VINO: An Integrated Platform for Database and Operating System Research, Harvard University Technical Report TR-30-94.

Grosz, B., Kung, H. T., Seltzer, M., Shieber, S., Smith, M., "Infrastructure for Research towards Ubiquitous Information Systems," Harvard University Technical Report TR-28-94.

Seltzer, M., "File System Performance and Transaction Support," PhD Dissertation, University of California, Berkeley, College of Engineering, Electrical Engineering and Computer Science Department, Computer Science Division, December 1992.