

SIGOPS

33-5



OPERATING SYSTEMS REVIEW

A Publication of the
Association for Computing Machinery
Special Interest Group on Operating Systems

Volume 33, Number 5

December, 1999

Proceedings of the

17th ACM Symposium on Operating Systems Principles (SOSP'99)

December 12-15, 1999

Kiawah Island Resort, near Charleston, South Carolina

OPERATING SYSTEMS REVIEW

A Publication of the ACM Special Interest Group on Operating Systems

Chairperson

William E. Weihl
Compaq Systems Research Center
130 Lytton Avenue
Palo Alto, CA 94301 USA
+1-650-853-2197
(weihl@pa.dec.com)

Newsletter Editor

William M. Waite
Dept. of Electrical and
Computer Engineering
University of Colorado
Boulder, CO 80309-0425 USA
+1-303-492-7204
(William.Waite@Colorado.edu)

V. Chairperson

Valerie Issarny
INRIA - Rennes
IRISA
Campus de Beaulieu
35042 Rennes Cedex France
+33-2-99-84-7462
(issarny@irisa.fr)

Secretary-Treasurer

David Kotz
Dept. of Computer Science
Dartmouth College
6211 Sudikoff Laboratory
Hanover, NH 03755 USA
+1-603-646-1439
(dfk@cs.dartmouth.edu)

Information Director

Mary Baker
Computer Science Department
Gates Building 4A
Stanford University
Stanford, CA 94305-9040 USA
+1-415-725-3711
(mgbaker@cs.stanford.edu)

OPERATING SYSTEMS REVIEW is an informal publication of the ACM Special Interest Group on Operating Systems (SIGOPS), whose scope of interest includes: Computer operating systems and architecture for multiprogramming, multiprocessing, and time sharing; resource management; evaluation and simulation; reliability, integrity, and security of data; communications among computing processors; and computer system modeling analysis.

Membership in SIGOPS (at \$15 per annum) is open to ACM members and associate members, and student members (at \$8 per annum); non-ACM members may join also (at \$42 per annum). Non-members of SIGOPS may subscribe to OPERATING SYSTEMS REVIEW at \$30 per year. All SIGOPS members receive OPERATING SYSTEMS REVIEW.

SIGOPS membership application forms are available from ACM Headquarters (or see back cover of this issue), 1515 Broadway, 17th Floor, New York, New York 10036, telephone (212) 869-7440. Changes of address and other matters pertaining to the SIGOPS mailing list should be directed to ACM Headquarters, not to any of the SIGOPS officers.

Contributions to OSR are sent to the editor and should be single spaced and printed on one side of the paper only. Text and diagrams must be contained within a 6½ inch wide by 8½ inch high (165mm by 216mm) area with at least one inch (25mm) top and side margins, 1½ inch (38mm) bottom margin, and be camera-ready (i.e., be an original having sufficient contrast and density for reproduction). All letters to the editor will be considered for publication unless accompanied by a request to the contrary. Except for editorial items, all sources of material appearing in OPERATING SYSTEMS REVIEW will be clearly identified. Items attributed to individuals will ordinarily be interpreted as personal rather than organizational opinions. The technical material in this newsletter is not refereed. **Deadlines for submitting material are the 15th of February, May, August, and November. The SIGOPS conference proceedings will be mailed as the December issue in odd numbered calendar years, and the ASPLOS conference proceedings will be mailed in even numbered calendar years.**

OPERATING SYSTEMS REVIEW (ISSN 0163-5980) is published five times a year (January, April, July, October and December) by the Association for Computing Machinery Inc., 1515 Broadway, New York, NY 10036. Periodicals postage paid at New York, NY 10001, and at additional mailing offices. POSTMASTER: Send address changes to OPERATING SYSTEMS REVIEW, ACM, 1515 Broadway, New York, NY 10036.

Subscriptions: Annual subscription cost of \$12.53 is included in the member dues of \$15.00 (for students cost is included in \$8.00); the non-member annual subscription is \$42.00.

RECEIVED JAN 10 2000

Proceedings of the


17th ACM Symposium on Operating Systems Principles (SOSP'99)

December 12–15, 1999

Kiawah Island Resort, near Charleston, South Carolina

Sponsored by ACM SIGOPS

(Association for Computing Machinery special interest group in operating systems).



Supported by:

COMPAQ

COMPAQ Research

Lucent Technologies
Bell Labs Innovations



Lucent Bell Laboratories

Computer Systems, Inc.
MERCURY

Mercury Computer Systems

hp HEWLETT®
PACKARD
Expanding Possibilities

Hewlett-Packard
Laboratories

IBM

IBM Research

Microsoft
Research


Microsoft Research

The Association for Computing Machinery, Inc.
1515 Broadway
New York, New York 10036

Copyright © 1999 by the Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted.

To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from: Publications Dept. ACM, Inc. Fax +1-212-869-0481 or E-mail permissions@acm.org.

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



ACM ISBN: 1-58113-140-2

Additional copies may be ordered prepaid from:

ACM Order Department
P.O. BOX 11405
Church Street Station
New York, NY 10286-1405

Phone: +1-800-342-6626 (U.S.A. and Canada)
+1-212-626-0500 (all other countries)
Fax: +1-212-944-1318
E-mail: acmhelp@acm.org

ACM Order Number: 534990

Printed in the U.S.A.

Seventeenth ACM Symposium on Operating Systems Principles

Message from the General Chair

This year, SOSOP is being held at the Kiawah Island Resort, just south of Charleston, South Carolina. The island offers a beautiful setting for our meeting, with plenty of diversions such as a stroll on the beach or a bike ride on the trails that lace the island.

As General Chair it has been my privilege to work with several outstanding people, who have helped to pull this conference together. John Wilkes and his program committee have done an outstanding job, and worked incredibly hard, to produce a first-rate program. John also produced the paper proceedings, and Chris Small arranged the CD-ROM version. Ann Chervenak has worked tirelessly to arrange all the little details necessary to run the show at Kiawah, with help from Ruthie Simmons of the Kiawah conference staff. Amin Vahdat arranged for a smooth registration process, with help from ACM's Haritini Kanthou who built the registration web site. Jeff Chase rounded up student volunteers to help with the registration desk, and selected the wines for our drinking pleasure. Eric Jul and his students Mads Dydensborg and Michael Svendsen arranged all the publicity and managed the web site. Cary Gray managed the finances. Guerney Hunt and his committee had the tough job of choosing which of many worthy students to receive travel scholarships. Jay Lepreau put many of them to work as session scribes, to provide input for a post-conference summary to be published in *Operating Systems Review*. Finally, thanks to all of the authors for their hard work in preparing their submitted papers.

I am also particularly thankful to the industrial sponsors this year: Compaq Research, Hewlett-Packard Laboratories, IBM Research, the Information Sciences Research Center at Lucent Bell Laboratories, Mercury Computer Systems, and Microsoft Research.

These proceedings include a CD-ROM that contains an electronic version of all of the papers, and code or data for some of the papers. An archive of SOSOP 17 information, including items like poster abstracts and synopses of the sessions, which are not included in either the paper or CD-ROM proceedings, will be available on the WWW at <http://www.acm.org/sigops/sosp99/>.

Many thanks to everyone!

David Kotz
Dartmouth College

Seventeenth ACM Symposium on Operating Systems Principles

Student scholarship awards

The student scholarship program for SOSOP 17 was provided courtesy of Compaq, Hewlett-Packard, IBM, Microsoft Research and SIGOPS. SIGOPS contributed \$20K and each industrial sponsor contributed \$10K. The committee was able to grant scholarships to the top 60 of the 97 students who applied.

Scholarships were awarded by a five member committee which was made up of representatives from SIGOPS, industry, and academia. I thank Nina Bhatti (HP Labs), Ken Birman (Cornell University), Garth Gibson (Carnegie Mellon University) and Marc Shapiro (SIGOPS) for serving as members of the scholarship committee.

The goals of the program were to increase student participation, recognize student authors, and support meritorious students. Each committee member evaluated the applicants individually. Applicants were evaluated based on a number of factors including: scholarship, how much the conference would benefit the student, how much the student would contribute to the conference, and need. Scholarship was primarily determined by the student position papers. Student position papers were judged based on relevance to the OS community and relevance to the conference. Students who were primary authors of accepted papers received a high scholarship score. The benefit of the conference to the student was based on the advisor's recommendation and the student's paper. The ability of the student to contribute to the conference was based on their contribution to the research community, progress in their program, and the advisor's recommendation. The assessment of each committee member was weighted equally in ranking the students. The consequences of this ranking were fairly diverse.

Students from thirty-five of the forty universities represented in the applicant pool were supported. Although the majority of the applicants were from North America, European students represented 19.5% (19) of the applicants and received 23.3% (14) of the scholarships. The university with the most supported students was Cambridge with four followed by CMU, DIKU, MIT, Rice, Kentucky, and Washington with three each. All students who were primary authors of accepted papers were funded. Even though the scholarship committee did not have an obligation to spend all of its budget, this year there were more meritorious students who applied than could be funded. Each supported student was granted their estimated airfare, an extra night's stay, and the conference registration fee.

Chair: Guerney Hunt, IBM Research

The scholarship program was made possible by financial support from:



ACM SIGOPS



COMPAQ Research



Hewlett-Packard
Laboratories



IBM Research



Microsoft Research

Seventeenth ACM Symposium on Operating Systems Principles

Message from the Program Chair

It is with great pleasure that I introduce to you this collection of 19 papers chosen by the SOSP program committee. We received 90 full submissions this time; compared to the highly popular 16th SOSP (held in Saint-Malo, France), this is a drop of 17%, but it is roughly the same as the historical average for the prior decade's submissions. 80 of the submissions came from the USA: 60 from the edu domain and 17 from the com domain, although at least one author submitted a paper from both domains. Of the others, there were 4 submissions from each of Korea and Germany. 68 sites submitted just one paper; 11 submitted 3 or more.

The reviewing process was a combination of approaches that had proved successful in the past: in the first pass, every paper was assigned to three program committee members to review, one of whom solicited 2–3 external reviews. Using the first-round reviews as input, we focussed our attention on the most promising half of the papers: all of these were read by the entire committee, including all the papers that ended up in the proceedings. We also appointed a program committee member for each paper to lead the discussion in the program committee meeting, which generated a spirited and highly informed debate. Each accepted paper was then shepherded by a program committee member, to help ensure that the feedback provided by the reviewers was put to the best possible use. All the authors put in great efforts to do so, under very tight deadlines, to achieve the high quality results you see before you.

The reviewing process also represents a huge contribution of time, expertise and effort by the outside reviewers and the program committee: between them, they wrote over 800 reviews containing over 280 thousand words. I cannot think of a better endorsement for our field than the amount of effort that people were willing to put into it.

As always, our choice of what to accept was made difficult by the number, variety and quality of the papers submitted. The inventiveness and range of the submissions remains a strong point of this conference. Several submissions were a joy to read, and the program committee asked me to express their hope that more future submitters would pay greater attention to the guidance offered by Roy Levin and David D. Redell¹, as well as find ways to add more statistical rigor in their submissions (especially repeatability data, and error- or confidence-estimates). Papers that did so were perceived as significantly stronger.

For the first time, we used an electronic submission process (David Nicol's WIMPE). My thanks to David for being so helpful in responding to a stream of requests for help and bug fixes. Dejan Milojicic and Mike Rodriguez helped greatly in bringing up WIMPE in the HP Labs environment. My other HP Labs colleagues deserve a great deal of credit for being so tolerant of my idiosyncrasies throughout the entire process!

To complement the more formal papers, the conference includes the traditional informal work-in-progress session, as well as a poster session, to encourage discussion of less fully-formed results. Data from these will be made available at the conference web site.

The high quality of the SOSP conference is entirely due to the hard work of the authors, the program committee and the reviewers. It has been a pleasure to work with such a dedicated and enthusiastic group of people. I hope you like the result.

John Wilkes
October 1999

1. Roy Levin and David D. Redell. An evaluation of the ninth SOSP submissions; or how (and how not) to write a good systems paper. *Operating Systems Review* 17(3):35-40 (July 1983).
Available at: <http://www.usenix.org/events/samples/submit/advice.html> as of Oct. 1999.

Seventeenth ACM Symposium on Operating Systems Principles

Outside reviewers

Martín Abadi	S. G. Hild	Thomas L. Rodeheffer
Tarek Abdelzaher	Wilson Hsieh	Drew Roselli
Atul Adya	Marty Humphrey	Norm Rubin
Lorenzo Alvisi	Yury Izrailevsky	M. Satyanarayanan
David G. Andersen	Kevin Jeffay	James B. Saxe
Thomas Anderson	Michael Jones	Dan Scales
Godmar Back	Eddie Kohler	Elizabeth Seamans
Gaurav Banga	P. Krishnan	Margo Seltzer
Frank Bellosa	Paul Leach	Lui Sha
Steven M. Bellovin	Dennis Lee	Marc Shapiro
Brian Bershad	Ed Lee	Douglas Shepherd
Guy Blelloch	Jorg Liebeherr	Kang G. Shin
Bill Bolosky	Beng-Hong Lim	Olin Shivers
Edouard Bugnion	Steven B. Lipner	Liuba Shrira
Mike Burrows	Jacob Lorch	Christopher Small
John Carter	Chenyang Lu	Stephen Smalley
Miguel Castro	Mark Manasse	Jonathan M. Smith
Jeff Chase	Timothy Mann	Keith A. Smith
Alan Demers	Keith Marzullo	Sang H. Son
Fred Douglass	Matthew Mathis	Sudhir Srinivasan
Carla Ellis	Jeanna Neeffe Matthews	David Steere
Dawson Engler	Roland McGrath	Dave Stewart
Mike Feeley	Greg Morrisett	Ion Stoica
Anja Feldmann	Sape Mullender	Sai Susarla
Paulo Ferreira	Sue Nagy	Patrick Tullmann
Matthew Flatt	George Necula	Uresh Vahalia
Armando Fox	Scott Nettles	Amin Vahdat
Eran Gabber	Wee Teck Ng	Robbert van Renesse
Greg Ganger	Jason Nieh	Ben Verghese
Sanjay Ghemawat	Brian Noble	Radek Vingralek
Ed Gould	Banu Ozden	Werner Vogels
Kinshuk Govil	Vivek Pai	David Wagner
Steven Gribble	Craig Partridge	Randolph Wang
David Halls	Radia Perlman	William E. Weihl
Steven Hand	Karin Petersen	Matt Welsh
John H. Hartman	Shankar Ponnkanti	Richard Wheeler
Chris Hawblitzel	John Regehr	Jeannette Wing
Mark Hayden	Mike Reiter	Ted Wobber
John Heidemann	Martin Rinard	Bennet Yee
Mike Hibler	Ohad Rodeh	John Zahorjan

Seventeenth ACM Symposium on Operating Systems Principles

Conference officers

- General chair: David Kotz, *Dartmouth College*
- Program committee: John Wilkes, *HP Laboratories (chair)*
Jean Bacon, *Cambridge University*
David Black, *EMC Corporation*
Eric Brewer, *University of California, Berkeley*
Peter Chen, *University of Michigan*
Peter Lee, *Carnegie Mellon University*
Jay Lepreau, *University of Utah*
Barbara Liskov, *Massachusetts Institute of Technology*
Jeff Mogul, *Compaq Western Research Laboratory*
Mendel Rosenblum, *Stanford University*
Fred Schneider, *Cornell University*
Jack Stankovic, *University of Virginia*
Chandu Thekkath, *Compaq Systems Research Center*
- Treasurer: Cary Gray, *Wheaton College*
- Local arrangements: Ann Chervenak, *Georgia Tech*
- Volunteers: Jeff Chase, *Duke University*
- Registration: Amin Vahdat, *Duke University*
- Publicity: Eric Jul, *DIKU, University of Copenhagen (chair)*
Mads Dydensborg, *DIKU, University of Copenhagen*
Michael Svendsen, *DIKU, University of Copenhagen*
- Session scribes: Jay Lepreau, *University of Utah*
- CD-ROM: Christopher Small, *Bell Labs*
- Scholarship committee: Guerney Hunt, *IBM Research (chair)*
Nina Bhatti, *HP Laboratories*
Ken Birman, *Cornell University*
Garth Gibson, *Carnegie Mellon University*
Marc Shapiro, *INRIA and SIGOPS*
- Web sites: <http://www.acm.org/sigops/sosp99/> (*primary*)
<http://www.diku.dk/sosp99/>
<http://www.cs.dartmouth.edu/SOSP99/>
<http://www-sor.inria.fr/mirrors/sosp99/>

Seventeenth ACM Symposium on Operating Systems Principles

Table of contents

Distributed systems (1)

- Manageability, availability and performance in Porcupine: a highly scalable, cluster-based mail service*
Yasushi Saito, Brian N. Bershad, Henry M. Levy
(University of Washington) 1
- On the scale and performance of cooperative web proxy caching*
Alec Wolman, Geoffrey M. Voelker, Nitin Sharma, Neal Cardwell, Anna Karlin, Henry M. Levy
(University of Washington) 16

Client systems

- The interactive performance of SLIM: a stateless, thin-client architecture*
Brian K. Schmidt, Monica S. Lam, J. Duane Northcutt[†]
(Stanford University and [†]Sun Microsystems Laboratories) 32
- Energy-aware adaptation for mobile applications*
Jason Flinn, M. Satyanarayanan
(Carnegie Mellon University) 48

Networking (1)

- Active network vision and reality: lessons from a capsule-based system*
David Wetherall
(University of Washington) 64
- Building reliable, high-performance communication systems from components*
Xiaoming Liu, Christoph Kreitz, Robbert van Renesse, Jason Hickey, Mark Hayden, Kenneth Birman,
Robert Constable
(Cornell University) 80

File systems

- File system usage in Windows NT 4.0*
Werner Vogels
(Cornell University) 93
- Deciding when to forget in the Elephant file system*
Douglas S. Santry, Michael J. Feeley, Norman C. Hutchinson, Alistair C. Veitch[†],
Ross W. Carton, Jacob Ofir
(University of British Columbia and [†]HP Laboratories) 110
- Separating key management from file system security*
David Mazières, Michael Kaminsky, M. Frans Kaashoek, Emmett Witchel
(MIT Laboratory for Computer Science) 124

OS kernels

<i>Integrating segmentation and paging protection for safe, efficient and transparent software extensions</i> Tzi-cker Chiueh, Ganesh Venkitachalam, Prashant Pradhan (State University of New York at Stony Brook)	140
<i>Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors</i> Kinshuk Govil, Dan Teodosiu [†] , Yongqiang Huang, Mendel Rosenblum (Stanford University and [†] HP Laboratories)	154
<i>EROS: a fast capability system</i> Jonathan S. Shapiro [†] , Jonathan M. Smith, David J. Farber (University of Pennsylvania and [†] IBM T.J. Watson Research Center)	170
Distributed systems (2)	
<i>The design and implementation of an intentional naming system</i> William Adjie-Winoto, Elliot Schwartz, Hari Balakrishnan, Jeremy Lilley (MIT Laboratory for Computer Science)	186
<i>Design and implementation of a distributed virtual machine for networked computers</i> Emin Gün Sirer, Robert Grimm, Arthur J. Gregory, Brian N. Bershad (University of Washington)	202
Networking (2)	
<i>The Click modular router</i> Robert Morris, Eddie Kohler, John Jannotti, M. Frans Kaashoek (MIT Laboratory for Computer Science)	217
<i>Soft timers: efficient microsecond software timer support for network processing</i> Mohit Aron, Peter Druschel (Rice University)	232
Real time	
<i>Progress-based regulation of low-importance processes</i> John R. Douceur, William J. Bolosky (Microsoft Research)	247
<i>Borrowed-Virtual-Time (BVT) scheduling: supporting latency-sensitive threads in a general-purpose scheduler</i> Kenneth J. Duda, David R. Cheriton (Stanford University)	261
<i>EMERALDS: a small-memory real-time microkernel</i> Khawar M. Zuberi, Padmanabhan Pillai, Kang G. Shin (University of Michigan)	277

Seventeenth ACM Symposium on Operating Systems Principles

Author index

Adjie-Winoto, William	186	Levy, Henry M.	1, 16
Aron, Mohit	232	Lilley, Jeremy	186
Balakrishnan, Hari	186	Liu, Xiaoming	80
Bershad, Brian N.	1, 202	Mazières, David	124
Birman, Kenneth	80	Morris, Robert	217
Bolosky, William J.	247	Northcutt, J. Duane	32
Cardwell, Neal	16	Ofir, Jacob	110
Carton, Ross W.	110	Pillai, Padmanabhan	277
Cheriton, David R.	261	Pradhan, Prashant	140
Chiueh, Tzi-cker	140	Rosenblum, Mendel	154
Constable, Robert	80	Saito, Yasushi	1
Douceur, John R.	247	Santry, Douglas S.	110
Druschel, Peter	232	Satyanarayanan, M.	48
Duda, Kenneth J.	261	Schmidt, Brian K.	32
Farber, David J.	170	Schwartz, Elliot	186
Feeley, Michael J.	110	Shapiro, Jonathan S.	170
Flinn, Jason	48	Sharma, Nitin	16
Govil, Kinshuk	154	Shin, Kang G.	277
Gregory, Arthur J.	202	Sirer, Emin Gün	202
Grimm, Robert	202	Smith, Jonathan M.	170
Hayden, Mark	80	Teodosiu, Dan	154
Hickey, Jason	80	van Renesse, Robbert	80
Huang, Yongqiang	154	Veitch, Alistair C.	110
Hutchinson, Norman C.	110	Venkitachalam, Ganesh	140
Jannotti, John	217	Voelker, Geoffrey M.	16
Kaashoek, M. Frans	124, 217	Vogels, Werner	93
Kaminsky, Michael	124	Wetherall, David	64
Karlin, Anna	16	Witchel, Emmett	124
Kohler, Eddie	217	Wolman, Alec	16
Kreitz, Christoph	80	Zuberi, Khawar M.	277
Lam, Monica S.	32		