



OPERATING SYSTEMS REVIEW

A Quarterly Publication of the
Special Interest Group on Operating Systems

Volume 23, Number 5

Special Issue

**Proceedings of
the Twelfth ACM Symposium
on
Operating Systems Principles**

**December 3-6, 1989
The Wigwam
Litchfield Park, Arizona**

OPERATING SYSTEMS REVIEW

A Quarterly Publication of the ACM Special Interest Group on Operating Systems

Chairperson

Roy Levin
Digital Equipment Corporation
Systems Research Center
130 Lytton Avenue
Palo Alto, CA 94301
(415)853-2100
(levin@src.dec.com)

V. Chairperson

Sape J. Mullender
CWI
Kruislaan 413
1098 SJ Amsterdam
Netherlands
+31 20 5924139
(sape@cw.nl) (mcvax!sape)

Secretary-Treasurer

Henry M. Levy
Dept. of Computer Science
University of Washington
Seattle, WA 98195
(206) 543-9204
(levy@whistler.cs.washington.edu)

Newsletter Editor

William M. Waite
Dept. of Electrical and
Computer Engineering 425
University of Colorado
Boulder, CO 80309-0425
(waite@boulder.colorado.edu)

OPERATING SYSTEMS REVIEW is an informal quarterly 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 \$10 per annum) is open to ACM members and associate members, and student members (at \$6 per annum); non-ACM members may join also (at \$32 per annum). Non-members of SIGOPS may subscribe to OPERATING SYSTEMS REVIEW at \$16 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), 11 West 42nd Street, 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, should not ordinarily exceed 1000 words, and should be single spaced and printed on one side of the paper only. Text and diagrams must be printed on 8½ by 11 inch paper and 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.**

**Proceedings of the
Twelfth ACM Symposium
on
Operating Systems Principles**

**December 3-6, 1989
The Wigwam
Litchfield Park, Arizona**

**Sponsored by
Special Interest Group on Operating Systems (SIGOPS)
Association for Computing Machinery (ACM)**

The Association for Computing Machinery
11 West 42nd Street
New York, New York 10036

Copyright © 1989 by the Association for Computing Machinery, Inc. Copying without fee is permitted provided that the copies are not made or distributed for direct commercial advantage, and credit to the source is given. Abstracting with credit is permitted. For other copying of articles that carry a code at the bottom of the first page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 27 Congress Street, Salem, MA 01970. For permission to republish write to: Director of Publications, Association for Computing Machinery. To copy otherwise, or republish, requires a fee and/or specific permission.

ISBN 0-89791-338-8

Additional copies may be ordered prepaid from:

ACM Order Department	Price
P.O. Box 64145	Members \$14.00
Baltimore, MD 21264	All Others \$19.00

ACM Order Number: 534890

Twelfth ACM Symposium on Operating Systems Principles

General Chair:

Greg Andrews, *University of Arizona*

Program Chair:

John Ousterhout, *University of California at Berkeley*

Program Committee:

Andrew Black, *Digital Equipment Corporation*

Jim Gray, *Tandem Computers*

Hugh C. Lauer, *Kodak Boston Technology Center*

Sape Mullender, *Centrum voor Wiskunde en Informatica*

David Redell, *Digital Equipment Corporation*

M. Satyanarayanan, *Carnegie Mellon University*

Fred B. Schneider, *Cornell University*

Doug Terry, *Xerox Palo Alto Research Center*

Marvin Theimer, *Xerox Palo Alto Research Center*

Local Arrangements:

Rick Schlichting, *University of Arizona*

Treasurer:

Larry Peterson, *University of Arizona*

Publicity:

Norm Hutchinson, *University of Arizona*

Secretary:

Beth Stair, *University of Arizona*

Outside Reviewers

The following people assisted the program committee by writing reviews for one or more of the submitted papers:

Guy Almes	Kieren Hartley	Robert Sansom
Yeshayahu Artsy	John Hartman	Frank Schmuck
Ozalp Babaoglu	Roger Haskin	Mike Schroeder
Jean Bacon	Pat Helland	Peter Schwarz
Mary Baker	Norman Hutchinson	Santosh Shrivastava
Joel Bartlett	Michael Kazar	Ellen Siegel
Wendy Bartlett	Chris Kent	Karen Sollins
Tsahi Birk	James Kistler	Jim Stamos
Ken Birman	Manfried Krause	Chris Stephenson
David Black	Paul Leach	Patrick Stephenson
Toby Bloom	Roy Levin	Howard Sturgis
Anita Borg	Bruce Lindsay	Dave Tennenhouse
Bob Bruce	Bryan Lyles	Ted Wobber
Mike Burrows	David Maier	John Wroclawski
Luis-Felipe Cabrera	Tim Mann	Jim Wyllie
Richard Carr	Keith Marzullo	Michael Young
Eric Cooper	Paul McJones	
Robert Cooper	Ralph Merkle	
Flaviu Cristian	Nathaniel Mishkin	
Dean Daniels	Jeff Mogul	
Alan Demers	Eliot Moss	
David Detlefs	Brian Oki	
Anthony Discolo	Susan Owicki	
Graeme Dixon	Prakash Panangaden	
Fred Douglass	Radia Perlman	
Alessandro Forin	Richard Rashid	
Lois Frampton	Eric Roberts	
Michel Gien	Tom Rodeheffer	
Robert Hagmann	Mendel Rosenblum	
Bert Halstead	Kevin St. Pierre	

Foreword

The twenty papers of the 12th Symposium on Operating Systems Principles (SOSP) reflect much of the best original work in the field over the last two years. Three-fourths of the papers in the conference fall under one of four themes. The first theme is remote procedure call (RPC). This is at least the fourth consecutive SOSP where RPC has been an important topic; the emphasis on performance rather than semantics in this year's papers suggests that RPC is reaching maturity. The second theme this year is file caching in distributed systems (five papers in two sessions). This is the third consecutive SOSP where file caches have been discussed; the broad range of caching papers this year suggests that distributed file caches have evolved from a research curiosity into a central component of any distributed system.

The third theme of this year's conference is multiprocessor operating systems. Multiprocessors have re-emerged as a popular research topic after several years of relative dormancy; five papers in two sessions deal with multiprocessor issues. In particular, large-scale multiprocessors, called "NUMA"s because of their non-uniform memory access times, pose several interesting new problems for operating system designers. The fourth theme is security. The importance of this topic dates back to the earliest SOSP's, and new challenges continue to appear as operating systems become more distributed in nature.

The program committee used a reviewing process similar to that of other recent SOSP's. Each of the ten members of the committee read all 84 of the submitted papers. During this process two written reviews were obtained for each paper: one from a program committee member and one from an outside reviewer. At that point the committee members ranked all of the papers and we eliminated the bottom half from further consideration. The program committee then re-examined the top half of the papers and produced at least one more written review for each paper in the top half.

The reviewing process culminated in a meeting of the program committee in Berkeley on May 22. Since every program committee member had read every paper, the discussions were both informed and lively. We accepted twenty papers and nominated three particularly outstanding ones for publication in *ACM Transactions on Computer Systems*. I am happy to report that all of the papers nominated for *Transactions on Computer Systems* were accepted and will appear in a special issue in February 1990. The reviewing process was nominally blind, in that none of the program committee members (including myself) was told the identity of any of the authors until after the program committee meeting. Of course, some of us couldn't help but guess the authors of some of the papers.

I would like to thank the members of the program committee for the enormous amount of time they invested in the reviewing process; their diligence in meeting the reviewing deadlines and their concern for producing the best possible conference made it a pleasure to work with them. The outside reviewers also deserve thanks for providing a "second opinion" in the reviewing process. Lastly, I'd like to thank my administrative assistant Bob Miller, who managed nearly all of the administrative details of the reviewing process and assembled these proceedings.

John Ousterhout
Program Chair
September 1989

Twelfth ACM Symposium on Operating Systems

Table of Contents

Monday, December 4, 1989

Session 1: Security

(F. Schneider, Session Chairman)

- A Logic of Authentication* 1
M. Burrows, M. Abadi, R. Needham
- Reducing Risks from Poorly Chosen Keys* 14
T. Lomas, L. Gong, J. Saltzer, R. Needham

Session 2: NUMA Memory Management

(D. Redell, Session Chairman)

- Simple But Effective Techniques for NUMA Memory Management* 19
W. Bolosky, R. Fitzgerald, M. Scott
- The Implementation of a Coherent Memory Abstraction on a NUMA Multiprocessor: Experiences with PLATINUM* 32
A. Cox, R. Fowler

Session 3: File Caching

(M. Satyanarayanan, Session Chairman)

- Spritely NFS: Experiments with Cache-Consistency Protocols* 45
V. Srinivasan, J. Mogul
- Exploiting Read-Mostly Workloads in the FileNet File System* 58
D. Edwards, M. McKendry
- Improving the efficiency of UNIX file buffer caches* 71
A. Braunstein, M. Riley, J. Wilkes

Session 4: Work in Progress and Outrageous Opinions

(J. Ousterhout, Session Chairman)

Short presentations on topics of current interest.

Tuesday, December 5, 1989

Session 5: Remote Procedure Call (D. Terry, Session Chairman)

<i>Performance of Firefly RPC</i>	83
M. Schroeder, M. Burrows	
<i>RPC in the x-Kernel: Evaluating New Design Techniques</i>	91
L. Peterson, N. Hutchinson, S. O'Malley, M. Abbott	
<i>Lightweight Remote Procedure Call</i>	102
B. Bershad, T. Anderson, E. Lazowska, H. Levy	

Session 6: Structural Issues (S. Mullender, Session Chairman)

<i>The Portable Common Runtime Approach to Interoperability</i>	114
M. Weiser, A. Demers, C. Hauser	
<i>Generic Virtual Memory Management for Operating System Kernels</i>	123
V. Abrossimov, M. Rozier, M. Shapiro	

Session 7: Multiprocessors (J. Ousterhout, Session Chairman)

<i>Low-Synchronization Translation Lookaside Buffer Consistency in Large-Scale Shared-Memory Multiprocessors</i>	137
B. Rosenberg	
<i>The Amber System: Parallel Programming on a Network of Multiprocessors</i>	147
J. Chase, F. Amador, E. Lazowska, H. Levy, R. Littlefield	
<i>Process Control and Scheduling Issues for Multiprogrammed Shared-Memory Multiprocessors</i>	159
A. Tucker, A. Gupta	

Session 8: Work In Progress and Outrageous Opinions (H. Lauer, Session Chairman)

Short presentations on topics of current interest.

Wednesday, December 7, 1989

Session 9: Performance

(A. Black, Session Chairman)

- A Lazy Buddy System Bounded by Two Coalescing Delays* 167
R. Barkley, T. Lee
- Analysis of Transaction Management Performance* 177
D. Duchamp
- Threads and Input/Output in the Synthesis Kernel* 191
H. Massalin, C. Pu

Session 10: Time-based Distributed Coherency

(M. Theimer, Session Chairman)

- Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency* 202
C. Gray, D. Cheriton
- Mirage: A Coherent Distributed Shared Memory Design* 211
B. Fleisch, G. Popek