



OPERATING SYSTEMS REVIEW

A Quarterly Publication of the
Special Interest Group on Operating Systems

Volume 21, Number 5

Special Issue

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

**8-11 November 1987
Stouffer Austin Hotel
Austin, Texas**

ACM ORDER NO. 534870

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

Secretary-Treasurer

Henry M. Levy
Dept. of Computer Science
University of Washington
Seattle, WA 98195

V. Chairperson

Sape J. Mullender
Digital Equipment Corporation
Systems Research Center
130 Lytton Avenue
Palo Alto, C 94301

Newsletter Editor

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

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 6¾ 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 14th of March, June, September, and December.

Proceedings of the
Eleventh ACM Symposium
on
Operating Systems Principles

8 - 11 November 1987

Stouffer Austin Hotel

Austin, Texas

Sponsored by
Special Interest Group on Operating Systems (SIGOPS)
Association for Computing Machinery (ACM)
Microelectronics and Computer Technology Corporation (MCC)

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

Copyright ©1987 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

0-89791-242-X

Additional copies may be ordered prepaid from:

ACM Order Department	Price
P.O. Box 64145	Members \$12.00
Baltimore, MD 21264	Non-Members \$16.00
ACM Order Number: 534870	

Eleventh ACM Symposium on Operating Systems

General Chairman

Les Belady, *MCC Software Technology Program*

Program Chairman

Alfred Spector, *Carnegie Mellon University*

Symposium Program Committee

Ozalp Babaoglu, *Cornell University*

Eric Cooper, *Carnegie Mellon University*

David Gifford, *Massachusetts Institute of Technology*

Edward Lazowska, *University of Washington*

Bruce Lindsay, *IBM Almaden Research Center*

John Ousterhout, *University of California, Berkeley*

Michael Schroeder, *DEC Systems Research Center*

Guy Steele, *Thinking Machines*

Norihisa Suzuki, *IBM Tokyo Research Laboratory*

Treasurer

John Rosenberg, *Carnegie Mellon University*

Local Arrangements Chairman

Brian Fugate, *MCC Software Technology Program*

Conference Organization

Claudia Raun, *MCC Software Technology Program*

Barbara Smith, *MCC Software Technology Program*

Forward to SOSP-11 Proceedings

Papers at the Eleventh Symposium on Operating Systems Principles (SOSP) range over a wide technical area and cover topics in communication systems, database systems, distributed systems, file systems, and multiprocessor operating systems. As is usual at SOSP, many of the papers report on experience with complete systems, though there are some that discuss algorithms, designs, or specification techniques. This year, many papers were on the subject of reliability in distributed systems, including a few focusing on transaction processing techniques. The approximately 80 papers not accepted for publication addressed a number of topics not represented in this symposium, including security, verification, and hardware architecture.

The diversity of paper topics is not unique to this year's SOSP. For a long while, both the community at large and the Program Committee have been taking the view that operating systems research covers more than such traditional topics of processor scheduling and virtual memory. The study of operating systems seems to include the study of algorithms, techniques, and architectures that are expected to be useful in a large number of application areas. While there has been concern that operating systems research will soon mature and grow dull, this will not happen if operating systems researchers and practitioners continue to investigate an increasing set of powerful building blocks. Based on the submissions, I believe the field is vibrant, and that there is plenty left to do as we develop ever more powerful virtual machines on top of an ever more powerful hardware base.

The diversity of topics and the quantity of submissions to SOSP makes the Program Committee's job difficult, but certainly educational. For SOSP-11, the committee considered each of the 100 papers in a first round of reviewing. The goal in this round was to generate two reports to authors for each paper and to choose the best 50 papers for continued consideration. This year, an outside referee wrote one of these two reports. In a second round, the Program Committee more carefully reviewed the top papers and wrote at least one more report on each paper. While refereeing is nominally blind, Program Committee members cannot help but guess the authors of many papers.

Final discussions ensued at a Program Committee meeting at Carnegie Mellon in late April. In most cases, there was continued interaction between Program Committee members and authors even after papers were selected. The result of this long process are the nineteen papers in the final program. The Program Committee recommended 6 outstanding papers for publication in *ACM Transactions on Computer Systems*, and we are happy to report that all of these were accepted. Preprints of these papers have been provided to all SOSP-11 attendees, but they are not included in the proceedings in accordance with ACM policy.

The two work-in-progress sessions are an important part of SOSP this year. There will be approximately 20 short talks. Information on these sessions does not, unfortunately, appear in the proceedings.

It was my pleasure to work with the Program Committee, who was responsive to deadlines and diligent in following the complex refereeing plans. I'd like to thank the committee and outside referees for their hard work. I would like to particularly thank my secretary, Debra Lynn, who handled the large quantity of mail and maintained the database of submitted papers.

Alfred Spector
Program Chairman
August 1987

Eleventh ACM Symposium on Operating Systems

Table of Contents

Monday 9 November 1987

Session 1: Distributed File Systems
(E. Lazowska, Session Chairman)

Scale and Performance in a Distributed File System (extended abstract) 1
J. Howard, M. Kazar, S. Menees, D. Nichols, M. Satyanarayanan,
R. Sidebotham, M. West

Caching in the Sprite Network File System (extended abstract) 3
M. Nelson, B. Welch, J. Ousterhout

Session 2: Processing on Multiple Computers
(M. Schroeder, Session Chairman)

Using Idle Workstations in a Shared Computing Environment 5
D. Nichols

Attacking the Process Migration Bottleneck 13
E. Zayas

Session 3: Useful Techniques
(E. Cooper, Session Chairman)

*Hashed and Hierarchical Timing Wheels: Data Structures for the Efficient
Implementation of a Timer Facility* 25
G. Varghese, T. Lauck

The Packet Filter: An Efficient Mechanism for User-level Network Code 39
J. Mogul, R. Rashid, M. Accetta

A Name Service for Evolving Heterogeneous Systems 52
M. Schwartz, J. Zahorjan, D. Notkin

Session 4: Reports on Work in Progress I
(B. Lindsay, Session Chairman)

Seven minute presentations outlining current research projects.

Tuesday 10 November 1987

Session 5: Operating Systems (D. Gifford, Session Chairman)

The Duality of Memory and Communication in the Implementation of a Multiprocessor Operating System 63
M. Young, A. Tevanian, R. Rashid, D. Golub, J. Eppinger, J. Chew,
W. Bolosky, D. Black, R. Baron

Time Warp Operating System 77
D. Jefferson, B. Beckman, F. Wieland, L. Blume, M. DiLoreto,
P. Hontalas, P. Laroche, K. Sturdevant, J. Tupman, Van Warren,
J. Wedel, H. Younger, S. Bellenot

Synchronization Primitives for a Multiprocessor: A Formal Specification 94
A. Birrell, J. Guttag, J. Horning, R. Levin

Session 6: Systems (N. Suzuki, Session Chairman)

Managing Stored Voice in the Etherphone System (extended abstract) 103
D. Terry, D. Swinehart

Fine-Grained Mobility in the Emerald System (extended abstract) 105
E. Jul, H. Levy, N. Hutchinson, A. Black

Session 7: Transaction Facilities (A. Spector, Session Chairman)

Recovery Management in QuickSilver (extended abstract) 107
R. Haskin, Y. Malachi, W. Sawdon, G. Chan

801 Storage: Architecture and Programming (extended abstract) 109
A. Chang, M. Mergen

Implementation of Argus 111
B. Liskov, D. Curtis, P. Johnson, R. Scheifler

Session 8: Reports on Work in Progress II (J. Ousterhout, Session Chairman)

Seven minute presentations outlining current research projects.

Wednesday 11 November 1987

Session 9: Reliability Methods

(O. Babaoglu, Session Chairman)

Exploiting Virtual Synchrony in Distributed Systems 123
K. Birman, T. Joseph

Log Files: An Extended File Service Exploiting Write-Once Storage 139
R. Finlayson, D. Cheriton

Session 10: Efficient Recovery Techniques

(G. Steele, Session Chairman)

A Simple and Efficient Implementation of a Small Database 149
A. Birrell, M. Jones, E. Wobber

Reimplementing the Cedar File System Using Logging and Group Commit 155
R. Hagmann

The articles for which there are only extended abstracts will be published in *ACM Transactions on Computer Systems*. Pre-prints of these papers are available to symposium attendees.