



A. Karshmer J. Nehmer (Eds.)

Operating Systems of the 90s and Beyond

International Workshop

Dagstuhl Castle, Germany, July 8-12, 1991

Proceedings

Springer-Verlag

Berlin Heidelberg New York

London Paris Tokyo

Hong Kong Barcelona

Budapest

Series Editors

Gerhard Goos
Universität Karlsruhe
Postfach 69 80
Vincenz-Priessnitz-Straße 1
W-7500 Karlsruhe, FRG

Juris Hartmanis
Department of Computer Science
Cornell University
5148 Upson Hall
Ithaca, NY 14853, USA

Editors

Arthur Karshmer
Computer Science Department, New Mexico State University
Las Cruces, NM 88003, USA

Jürgen Nehmer
Computer Science Department, University of Kaiserslautern
W-6750 Kaiserslautern, FRG

CR Subject Classification (1991): D.4

ISBN 3-540-54987-0 Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-54987-0 Springer-Verlag New York Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1991
Printed in Germany

Typesetting: Camera ready by author
Printing and binding: Druckhaus Beltz, Hemsbach/Bergstr.
45/3140-543210 - Printed on acid-free paper

Preface

Operating Systems of the 1990s And Beyond: Where Do We Go From Here?

Architectural and hardware advances in computing systems design are occurring at an ever quickening rate, but it is not clear that the operating systems that make these new systems useful are keeping pace. Indeed, it is the operating system that masters the complexity of the ever more complex computing devices being built to make them useful tools. In the past, the existence of a variety of operating systems has made the difference between an interesting architecture and a useful computing environment.

Today, as more and more complex computational structures are emerging, and new and more powerful communications technologies are becoming available, we are faced with the need to develop new generations of operating systems to harness their power. A few of the new challenges to face the operating system implementor include, but are not restricted to:

- Unreliable communications,
- Fault tolerance,
- Issues of size and scalability,
- Integration of heterogeneous systems,
- Supporting advanced applications, such as multimedia,
- Protection and security issues in faulty/untrustworthy distributed systems,
- Coping with existing systems in networked environments.

What form should future operating systems take to address these and numerous other complex problems? In which area is further research indicated? Is the current generation of operating systems a valid platform for the operating systems of the next decade and beyond, or should we be designing a whole new generation of operating systems from the bottom up? What type of architectural support for operating systems and communications hardware should be built into the next generation of computers?

These are the general questions that our workshop attempted to address. It is our belief that now is the time to address these questions in a manner that will have some chance of producing useful results in the form of guidelines for future operating systems design and development. Three compelling circumstances lead us to believe that this is the time to act in terms of plotting a rational course for operating systems of the 1990s and beyond:

- a) The commercially available operating systems of today, including UNIX, have been designed, or based on designs, that are twenty years old, with timesharing on a single node machine being the foundation of the design. Networking and distributed computing facilities were “add- ons” and not part of the basic system design. Today’s new computing architectures are moving in different directions, unfortunately, with operating systems that have been pieced together to do the job.
- b) UNIX has formed the foundation for a de facto international standard for operating systems. Is UNIX the proper foundation for future efforts? Is its basic structure appropriate for computing equipment currently coming onto the market, or more importantly for new machines that will be unveiled in the near future?
- c) We are now at a crossroad characterized by a change in technology from the traditional single- node systems to networks of computers, distributed machines and massively parallel systems. It is not clear that current operating systems technology is poised to take advantage of these new machines when they do become available.

For all of the above reasons, an international workshop entitled *Operating Systems of the 1990s And Beyond: Where Do We Go From Here* was organized and held at the International Conference and Research Center for Computer Science (IBFI) at Dagstuhl Castle in the Federal Republic of Germany. The overriding motivation behind the workshop was to provide an opportunity for a relatively small number of leading researchers in operating systems from both universities and industry to meet and discuss current problems and future directions. An initial organizing committee was formed in April 1990 to help identify the key topics and key people to be involved in the workshop. Six specific topics for discussion were chosen along with people to present a *white paper* on each of the topics. A seventh session was planned as a summary of the week’s work.

After the selection of topics and discussion leaders was made, approximately 30 other participants were selected and invited to present position papers on one or several of the selected topics. The participants came from Europe, the United States and Japan. During each session, a white paper was presented followed by several position papers, all of which were mingled with open discussion. The materials contained in the remainder of this volume include the text of all white and position papers as well as a condensation of the discussion that took place during each session. Each session is represented by a separate section in the book which includes the papers presented and the discussion that took place.

Note on the Session Organization

The papers in the book were grouped with respect to the major contributions they made to the session themes. The actual presentation of position statements in the sessions differed substantially from this grouping since the discussions were given priority over the presentation of papers. This is reflected in the summary parts of each session by the list of speakers who actually gave position papers.

Acknowledgments

It would have been difficult to organize a workshop like this without help and financial support from a variety of people and institutions. First, we would like to thank all of the participants for their efforts, both written and oral. Next, the organizing committee deserves a large measure of gratitude for their efforts over a period of 16 months in putting together an excellent agenda and assembling the appropriate participants. Clearly an effort such as this could not have been successful without financial support. For this we would like to thank Siemens Corporation and IBM Germany for their contribution of general purpose funding. The IBFI supplied funding for the residence and meeting facilities. Travel support for all participants from the United States was generously provided by a grant from the National Science Foundation. Finally, our thanks go to Peter Buhler and Peter Sturm of the University of Kaiserslautern who took notes of all discussions for inclusion in this volume and to Mrs. Alt from the University who did a splendid job with the clerical work and arranged the social events.

Kaiserslautern
October 1991

Arthur Karshmer
Jürgen Nehmer

Contents

Session 1: Issue of Size, Scalability and Distribution in Future Operating Systems

An Agenda for Research in Large-Scale Distributed Data Repositories	2
<i>M. Satyanarayanan</i>	
File System Indexing and Backup	13
<i>J.H. Saltzer</i>	
Intelligent File Systems for Object Repositories	20
<i>D.W. Gifford, J.W. O'Toole</i>	
Using Domains for Specifying Management Policy	25
<i>M. Sloman</i>	
Discussion Summary	31

Session 2: Impact an Operating Systems by Future Trends in Hardware and Communication Technology

Future Operating Systems: An Environmental Scenario and Consequential Challenges	34
<i>B. Joy</i>	
Operating System Support for High-Performance Architectures	40
<i>E.D. Lazowska</i>	
Closely Coupled Systems	44
<i>P. Schlichtiger</i>	
Are Operating Systems at RISC?	48
<i>A.I. Karshmer, J.N. Thomas</i>	
Operating Systems of the 90s and Beyond (with Emphasis on Multiprocessing)	53
<i>J. Zahorjan</i>	
Munin: Distributed Shared Memory Using Multi-Protocol Release Consistency	56
<i>J.K. Bennett, J.B. Carter, W. Zwaenepoel</i>	
Discussion Summary	61

Session 3: Integrating Heterogeneous Operating Systems

Autonomous Heterogeneous Computing - Some Open Problems	63
<i>H. Schmutz</i>	
Coping with the Past	72
<i>B. Randell</i>	
The Immortality of Operating Systems, or: Is Research in Operating Systems still Justified?	77
<i>J. Nehmer</i>	
Discussion Summary	84

Session 4: Trends in Real-Time Operating Systems

Event-Triggered Versus Time-Triggered Real-Time Systems	87
<i>H. Kopetz</i>	
Pre-Scheduling for Synchronization in Hard Real-Time Systems	102
<i>S.K. Tripathi, V. Nirkhe</i>	
Notes on Future Operating Systems for Real-Time Dependable Distributed Computing	109
<i>G. Le Lann</i>	
Discussion Summary	116

Session 5: Fault Tolerance Support in Future Operating Systems

Basic Concepts and Issues in Fault-Tolerant Distributed Systems	119
<i>F. Cristian</i>	
Operating Systems and Fault-Tolerance	150
<i>F.D. Schlichting</i>	
Fault-Tolerance in Sixth-Generation Operating Systems	154
<i>K. Birman</i>	
Fault Tolerance Support in Future Operating Systems	162
<i>Ö. Babaoglu</i>	
Fast Stable Storage as a Basis for Fault-Tolerant Architectures	167
<i>J.P. Banatre, M. Banatre</i>	
Stable Memory - Another Look	171
<i>C. Horn, B. Coghlan, N. Harris, J. Jones</i>	
System Level Support for Dependable Distributed Applications	178
<i>R. Kröger, E. Nett</i>	
Discussion Summary	182

Session 6: Security and Protection Support in Future Operating Systems

Protection and Security Issues for Future Systems	184
<i>B.C. Neuman</i>	
Some Issues in Security Arising from Distribution	202
<i>R. Cole</i>	
Security: System Development in a New Context	207
<i>H.G. Stiegler</i>	
Issues in Security and Fault Tolerance	212
<i>H. Härtig, W. Kühnhauser, J. Liedtke</i>	
Discussion Summary	217

Session 7: Closing Session: The Next Generation of Operating Systems: What Is Needed and How We Get It

What Next? Some Speculations	220
<i>R. Needham</i>	
Some Thoughts on Systems Challenges for the 1990s	223
<i>A.Z. Spector</i>	
Next Generation Operating Systems Architecture	227
<i>M. Gien</i>	
Toward Computing Systems for the 2000s	233
<i>M. Tokoro</i>	
Encapsulation and Interaction in Future Operating Systems	241
<i>F. Hofmann, P. Schlenk, T. Eirich</i>	
The Explainable Operating System	247
<i>H.D. Wettstein</i>	
Soul: An Object-Oriented OS Framework for Object Support	251
<i>M. Shapiro</i>	
Operating System Support for Object-Oriented Distributed Systems	256
<i>S. Shrivastava, D.McCue</i>	
Discussion Summary	258
 Evening Session: Supporting Multimedia Applications in Distributed Systems	
Towards a Human Memory Prosthesis	262
<i>M.G. Lamming</i>	
Systems for the Nineties - Distributed Multimedia Systems	273
<i>S.J. Mullender</i>	
The Role of Performance, Scheduling and Resource Reservation in Multimedia Systems	279
<i>R.G. Herrtwich</i>	
List of Participants	285