

Lecture Notes in Computer Science
Edited by G. Goos, J. Hartmanis and J. van Leeuwen

1459

Dror G. Feitelson Larry Rudolph (Eds.)

Job Scheduling Strategies for Parallel Processing

IPPS/SPDP'98 Workshop

Orlando, Florida, USA, March 30, 1998

Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany
Juris Hartmanis, Cornell University, NY, USA
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Dror G. Feitelson
The Hebrew University of Jerusalem
Institute of Computer Science
91904 Jerusalem, Israel
E-mail: feit@cs.huji.ac.il

Larry Rudolph
Laboratory for Computer Science, MIT
Cambridge, MA 02139, USA
E-mail: rudolph@lcs.mit.edu

Cataloging-in-Publication data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

**Job scheduling strategies for parallel processing : proceedings /
IPPS/SPDP '98 workshop, Orlando, Florida, USA, March 30, 1998.
Dror G. Feitelson ; Larry Rudolph (ed.). - Berlin ; Heidelberg ; New
York ; Barcelona ; Budapest ; Hong Kong ; London ; Milan ; Paris ;
Santa Clara ; Singapore ; Tokyo : Springer, 1998
(Lecture notes in computer science ; Vol. 1459)
ISBN 3-540-64825-9**

CR Subject Classification (1991): D.4, D.1.3, C.2, F.2.2, C.1.2, B.2.1, B.6.1,
F.1.2

ISSN 0302-9743

ISBN 3-540-64825-9 Springer-Verlag Berlin Heidelberg New York

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 1998
Printed in Germany

Typesetting: Camera-ready by author
SPIN 10638279 06/3142 - 5 4 3 2 1 0 Printed on acid-free paper

Preface

This volume contains the papers presented at the fourth workshop on Job Scheduling Strategies for Parallel Processing, which was held in conjunction with IPPS/SPDP '98 in Orlando, Florida, on March 30, 1998. All the papers have gone through a complete refereeing process, with the full version being read and evaluated by five or six members of the program committee in most cases. We would like to take this opportunity to thank the program committee, Stephen Booth, Allan Gottlieb, Atsushi Hori, Phil Krueger, Richard Lagerstrom, Miron Livny, Virginia Lo, Reagan Moore, Bill Nitzberg, Uwe Schwiegelshohn, Ken Sevcik, Mark Squillante, John Zahorjan, and Songnian Zhou, for an excellent job. Thanks are also due to the authors for their submissions, presentations, and final revisions for this volume. Finally, we would like to thank the MIT Laboratory for Computer Science and the Computer Science Institute at Hebrew University for the use of their facilities in preparation of these proceedings.

As multi-user parallel supercomputers become more widespread, job scheduling takes on a crucial role. The number of users of parallel supercomputers is growing at an even faster pace and so there is an increasing number of users who must share a parallel computer's resources. Job scheduling strategies must address this need.

There is a spectrum of groups that are interested in job scheduling strategies for parallel processors. At one end are the vendors of parallel supercomputers who supply the scheduling software for managing jobs on their machines. In the middle are researchers in academia, National Labs, and industrial research labs who propose new scheduling strategies and methods for evaluating and comparing them. At the other end of the spectrum are the users and system administrators of parallel processing facilities who have a set of demands and requirements.

The goal of the workshop was to bring together people from all three groups, in order to exchange ideas and discuss ongoing work. Indeed, many interesting discussions took place, and the workshop was quite lively. We were encouraged by this since we believe it is important to increase communication so that academics work on the right problems and vendors and computation centers make the best use of the novel solutions. We hope these proceedings help parallel supercomputing to achieve its fundamental goal of satisfying the needs of the user.

This was the fourth annual workshop in this series, which reflects the continued interest in this field. The previous three were held in conjunction with IPPS '95 through IPPS '97. Their proceedings are available from Springer-Verlag as volumes 949, 1162, and 1291 in the Lecture Notes in Computer Science series.

Jerusalem, May 1998

Dror Feitelson
Larry Rudolph

Table of Contents

Metrics and Benchmarking for Parallel Job Scheduling	1
<i>Dror G. Feitelson, Larry Rudolph</i>	
A Comparative Study of Real Workload Traces and Synthetic Workload Models for Parallel Job Scheduling	25
<i>Virginia Lo, Jens Mache, and Kurt Windisch</i>	
Lachesis: A Job Scheduler for the Cray T3E	47
<i>Allen B. Downey</i>	
A Resource Management Architecture for Metacomputing Systems	62
<i>Karl Czajkowski, Ian Foster, Nick Karonis, Carl Kesselman, Stuart Martin, Warren Smith, and Steven Tuecke</i>	
Implementing the Combination of Time Sharing and Space Sharing on AP/Linux	83
<i>Kuniyasu Suzaki and David Walsh</i>	
Job Scheduling Scheme for Pure Space Sharing Among Rigid Jobs	98
<i>Kento Aida, Hironori Kasahara, and Seinosuke Narita</i>	
Predicting Application Run Times Using Historical Information	122
<i>Warren Smith, Ian Foster, and Valerie Taylor</i>	
Job Scheduling Strategies for Networks of Workstations	143
<i>B. B. Zhou, R. P. Brent, D. Walsh, and K. Suzaki</i>	
Probabilistic Loop Scheduling Considering Communication Overhead	158
<i>Sissades Tongsima, Chantana Chantrapornchai, and Edwin H.-M. Sha</i>	
Improving First-Come-First-Serve Job Scheduling by Gang Scheduling	180
<i>Uwe Schwiegelshohn and Ramin Yahyapour</i>	
Expanding Symmetric Multiprocessor Capability Through Gang Scheduling	199
<i>Morris A. Jette</i>	
Overhead Analysis of Preemptive Gang Scheduling	217
<i>Atsushi Hori, Hiroshi Tezuka, and Yutaka Ishikawa</i>	
Dynamic Coscheduling on Workstation Clusters	231
<i>Patrick G. Sobalvarro, Scott Pakin, William E. Weihl, and Andrew A. Chien</i>	
Author Index	257