

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Eitan Frachtenberg Uwe Schwiegelshohn (Eds.)

Job Scheduling Strategies for Parallel Processing

13th International Workshop, JSSPP 2007
Seattle, WA, USA, June 17, 2007
Revised Papers

Volume Editors

Eitan Frachtenberg
Powerset, Inc.
475 Brannan St., San Francisco, CA 94107, USA
E-mail: eitan@powerset.com

Uwe Schwiegelshohn
Technische Universität Dortmund, Institut für Roboterforschung
Otto-Hahn-Str. 8, 44227 Dortmund, Germany
E-mail: uwe.schwiegelshohn@udo.edu

Library of Congress Control Number: 2008924626

CR Subject Classification (1998): D.4, D.1.3, F.2.2, C.1.4, B.2.1, B.6, F.1.2

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

ISSN 0302-9743
ISBN-10 3-540-78698-8 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-78698-6 Springer 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. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2008
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12243488 06/3180 5 4 3 2 1 0

Preface

This volume contains the papers presented at the 13th workshop on Job Scheduling Strategies for Parallel Processing. The workshop was held in Seattle, WA, USA, on June 17, 2007, in conjunction with ICS 2007.

All submitted papers went through a complete review process, with the full version being read and evaluated by an average of five reviewers. We would like to thank the Program Committee members and additional referees for their willingness to participate in this effort and their excellent, detailed reviews: Nazareno Andrade, Su-Hui Chiang, Walfredo Cirne, Alvaro Coelho, Lauro Costa, Dror Feitelson, Allan Gottlieb, Andrew Grimshaw, Moe Jette, Richard Lagerstrom, Virginia Lo, Reagan Moore, Bill Nitzberg, Mark Squillante, John Towns, Jon Weissman, and Ramin Yahyapour.

The accepted workshop papers in recent years show a departure from the supercomputer-centric viewpoint of parallel job scheduling. On the one hand, the field of supercomputer scheduling is showing some signs of maturity, exhibited in many widely accepted practices for job scheduling. On the other hand, many nontraditional high-performance computing and parallel environments are emerging as viable solutions to many users and uses that cannot or need not access a traditional supercomputer, such as Grids, Web services, and commodity parallel computers. With the growing ubiquity of these technologies, the requirement to schedule parallel jobs well on these various architectures also grows.

Some of the papers in this and recent JSSPP proceedings address these contemporary architectures, such as Balle and Palermo's paper on resource management with awareness of multi-core and multi-threaded parallel processors, and whole sessions devoted to Grid scheduling. We plan to follow through with these shifts in parallel architectures and related scheduling in JSSPP. Opening this collection, we present a detailed exposition of recent technological trends in parallel computing, and how they affect, and are affected by, parallel job scheduling. To remain a relevant field of study in the face of constant progress in the underlying technology, we are currently expanding the scope of JSSPP to include these trends, as described in the introductory paper.

We call for the researchers in the JSSPP community and related fields to contribute from their experience with traditional parallel job scheduling, and to participate in the description, generalization, and resolving of the contemporary scheduling challenges. We hope to continue the long tradition of this workshop, which has now reached its 13th consecutive year. We would also love to hear from authors with any questions or suggestions regarding the future direction of JSSPP.

December 2007

Eitan Frachtenberg
Uwe Schwiegelshohn

Organization

Workshop Organizers

Eitan Frachtenberg	Powerset, Inc.
Uwe Schwiegelshohn	Dortmund University

Program Committee

Su-Hui Chiang	Portland State University
Walfredo Cirne	Google
Dror Feitelson	The Hebrew University
Allan Gottlieb	New York University
Andrew Grimshaw	University of Virginia
Moe Jette	Lawrence Livermore National Lab
Richard Lagerstrom	Cray Inc.
Virginia Lo	University of Oregon
Reagan Moore	San Diego Supercomputer Center
Bill Nitzberg	Altair Grid Technologies
Mark Squillante	IBM T. J. Watson Research Center
John Towns	NCSA
Jon Weissman	University of Minnesota
Ramin Yahyapour	Dortmund University

Table of Contents

New Challenges of Parallel Job Scheduling	1
<i>Eitan Frachtenberg and Uwe Schwiegelshohn</i>	
Group-Wise Performance Evaluation of Processor Co-allocation in Multi-cluster Systems	24
<i>John Ngubiri and Mario van Vliet</i>	
Enhancing an Open Source Resource Manager with Multi-core/Multi-threaded Support	37
<i>Susanne M. Balle and Daniel J. Palermo</i>	
A Job Self-scheduling Policy for HPC Infrastructures	51
<i>Francesc Guim and Julita Corbalan</i>	
QBETS: Queue Bounds Estimation from Time Series	76
<i>Daniel Nurmi, John Brevik, and Rich Wolski</i>	
Probabilistic Backfilling	102
<i>Avi Nissimov and Dror G. Feitelson</i>	
Impact of Reservations on Production Job Scheduling	116
<i>Martin W. Margo, Kenneth Yoshimoto, Patricia Kovatch, and Phil Andrews</i>	
Prospects of Collaboration between Compute Providers by Means of Job Interchange	132
<i>Christian Grimme, Joachim Lepping, and Alexander Papaspyrou</i>	
GridARS: An Advance Reservation-Based Grid Co-allocation Framework for Distributed Computing and Network Resources	152
<i>Atsuko Takefusa, Hidemoto Nakada, Tomohiro Kudoh, Yoshio Tanaka, and Satoshi Sekiguchi</i>	
A Self-optimized Job Scheduler for Heterogeneous Server Clusters	169
<i>Elad Yom-Tov and Yariv Aridor</i>	
Author Index	189