

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

Microsoft Research, Cambridge, 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

14th International Workshop, JSSPP 2009
Rome, Italy, May 29, 2009
Revised Papers

Volume Editors

Eitan Frachtenberg
Microsoft
475 Brannan St., San Francisco, CA, 94107, USA
E-mail: etc_26@yahoo.com

Uwe Schwiegelshohn
TU Dortmund University
Robotics Research Institute
Section Information Technology
Otto-Hahn-Str. 8, 44227 Dortmund, Germany
E-mail: uwe.schwiegelshohn@udo.edu

Library of Congress Control Number: 2009936148

CR Subject Classification (1998): D.4.1, I.2.8, F.2.2, C.1.2, C.1.4

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

ISSN 0302-9743
ISBN-10 3-642-04632-0 Springer Berlin Heidelberg New York
ISBN-13 978-3-642-04632-2 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.com

© Springer-Verlag Berlin Heidelberg 2009
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12761739 06/3180 5 4 3 2 1 0

Preface

This volume contains the papers presented at the 14th workshop on Job Scheduling Strategies for Parallel Processing. The workshop was held in Rome, Italy, on May 29, 2009, in conjunction with the IEEE International Parallel Processing Symposium 2009.

This year 25 papers were submitted to the workshop. All submitted papers went through a complete review process, with the full version being read and evaluated by an average of four reviewers. We would like to especially thank the program committee members and additional referees for their willingness to participate in this effort and their excellent, detailed reviews: Su-Hui Chiang, Walfredo Cirne, Allen Downey, Dror Feitelson, Alexander Fölling, Allan Gottlieb, Christian Grimme, Andrew Grimshaw, Moe Jette, Joachim Lepping, Raquel Lopes, Reagan Moore, Jose Moreira, Bill Nitzberg, Alexander Papaspyrou, Lars Schley, Mark Squillante, John Towns, Dan Tsafir, Jon Weissman, and Philipp Wieder.

As a result of the review process 14 papers were accepted for oral presentation at the workshop. One additional paper is included in these proceedings after making substantial improvements based on the comments of the referees. The final versions of the papers in this volume have addressed the comments of the referees and partially reflect the discussions held during the workshop.

This year we observed an increasing trend towards heterogeneous and multi-core architectures. The paper by Gong, Pierces, and Fox proposes an improved heuristic approach to workflow scheduling and shows its efficiency with the help of simulations. This workshop series used to exclude task scheduling. However, precedence constraints are starting to play an important role in grid jobs. Therefore, DAGs and workflows are becoming more important in the context of job scheduling. Fölling, Grimme, Lepping, and Papaspyrou show in their paper that grids can produce win-win situations for independent sites if these sites are willing to collaborate by exchanging some jobs. However contrary to job scheduling on classical high performance architectures, there are hardly any workload traces available for grid computing, making it difficult to evaluate new scheduling approaches with respect to practical application. This problem is addressed by Lingrand, Montagnat, Martyniak, and Colling who analyze the workload of the presently largest production grid EGEE.

There are still a number of open issues in classical job scheduling on parallel architectures. For instance, Guim, Rodero, and Corbalan present an improvement of backfilling. Birkenheuer, Brinkmann, and Karl suggest applying overbooking, an approach well known in other scheduling and allocation areas, to job scheduling. Workload modeling also remains important in job scheduling for classical and new parallel architectures as workload traces remain rare and are not always applicable without modifications. To improve the present models,

Minh and Wolters present an approach to better include temporal locality. The problem of preventing faulty jobs from disrupting the schedule is addressed in the paper by Thebe, Bunde, and Leung who suggest using trial runs of restartable jobs.

Pascual, Navaridas, and Miguel-Alonso discuss allocation policies that better consider the topology of the parallel architecture. This issue may become important again in the context of very large parallel architectures. In situations in which performance is more important than efficiency, Sinnen, To, and Kaur propose to use task duplication in order to improve the speed-up of jobs. The paper by Wolf, Bansal, Hildrum, Parekh, Rajan, Wagle, and Wu presents approaches for scheduling and allocation of streaming applications. These applications are likely to become more important for parallel architectures as these architectures enter a broader market. Bobroff, Coppinger, Fong, Seelam, and Xu suggest an extension of the well-known LoadLeveler job scheduler to handle virtualization.

Job scheduling problems that are relevant in the context of multi-core architectures were the topic of the last session of the workshop. Sun, Cao, and Hsu suggest using resource augmentation to handle non-clairvoyant and malleable jobs. They evaluate their approach with both simulation and theoretical analysis. Zeng and Sodan show that resource utilization on multi-core architectures can be improved with the help of forming appropriate groups of jobs. This holds for time and space sharing. Then Sodan reports on first experiences with adaptive scheduling that adjusts the size of jobs according to the actual load situation. This is also done in the context of virtual machines. The final paper by Vrba, Espeland, Halvorsen, and Griwodz discusses the benefit of workload stealing for complex applications with respect to utilization and load balancing in multi-core architectures.

The proceedings of previous workshops are available from Springer as LNCS volumes 949, 1162, 1291, 1459, 1659, 1911, 2221, 2537, 2862, 3277, 3834, 4376, and 4942. Since 1998 these volumes have also been available online.

Finally, we would like to explicitly thank Joachim Lepping for his support in organizing the publication of this volume.

July 2009

Eitan Frachtenberg
Uwe Schwiegelshohn

Organization

Workshop Organizers

Eitan Frachtenberg	Microsoft
Uwe Schwiegelshohn	TU Dortmund University

Program Committee

Su-Hui Chiang	Portland State University
Walfredo Cirne	Google
Allen Downey	Olin College
Dror Feitelson	The Hebrew University
Allan Gottlieb	New York University
Andrew Grimshaw	University of Virginia
Moe Jette	Lawrence Livermore National Lab
Virginia Lo	University of Oregon
Jose Moreira	IBM T.J. Watson Research Center
Reagan Moore	San Diego Supercomputer Center
Bill Nitzberg	Altair Engineering
Mark Squillante	IBM T.J. Watson Research Center
Dan Tsafir	IBM T.J. Watson Research Center
John Towns	NCSA
Jon Weissman	University of Minnesota
Ramin Yahyapour	TU Dortmund University

Table of Contents

Dynamic Resource-Critical Workflow Scheduling in Heterogeneous Environments	1
<i>Yili Gong, Marlon E. Pierce, and Geoffrey C. Fox</i>	
Decentralized Grid Scheduling with Evolutionary Fuzzy Systems	16
<i>Alexander Fölling, Christian Grimme, Joachim Lepping, and Alexander Papaspyrou</i>	
Analyzing the EGEE Production Grid Workload: Application to Jobs Submission Optimization	37
<i>Diane Lingrand, Johan Montagnat, Janusz Martyniak, and David Colling</i>	
The Resource Usage Aware Backfilling	59
<i>Francesc Guim, Ivan Rodero, and Julita Corbalan</i>	
The Gain of Overbooking	80
<i>Georg Birkenheuer, André Brinkmann, and Holger Karl</i>	
Modeling Parallel System Workloads with Temporal Locality	101
<i>Tran Ngoc Minh and Lex Wolters</i>	
Scheduling Restartable Jobs with Short Test Runs	116
<i>Ojaswirajanya Thebe, David P. Bunde, and Vitus J. Leung</i>	
Effects of Topology-Aware Allocation Policies on Scheduling Performance	138
<i>Jose Antonio Pascual, Javier Navaridas, and Jose Miguel-Alonso</i>	
Contention-Aware Scheduling with Task Duplication	157
<i>Oliver Sinnen, Andrea To, and Manpreet Kaur</i>	
Job Admission and Resource Allocation in Distributed Streaming Systems	169
<i>Joel Wolf, Nikhil Bansal, Kirsten Hildrum, Sujay Parekh, Deepak Rajan, Rohit Wagle, and Kun-Lung Wu</i>	
Scalability Analysis of Job Scheduling Using Virtual Nodes	190
<i>Norman Bobroff, Richard Coppinger, Liana Fong, Seetharami Seelam, and Jing Xu</i>	
Competitive Two-Level Adaptive Scheduling Using Resource Augmentation	207
<i>Hongyang Sun, Yangjie Cao, and Wen-Jing Hsu</i>	

Job Scheduling with Lookahead Group Matchmaking for Time/Space Sharing on Multi-core Parallel Machines	232
<i>Xijie Zeng and Angela C. Sodan</i>	
Adaptive Scheduling for QoS Virtual Machines under Different Resource Allocation – Performance Effects and Predictability	259
<i>Angela C. Sodan</i>	
Limits of Work-Stealing Scheduling.	280
<i>Željko Vrba, Håvard Espeland, Pål Halvorsen, and Carsten Griwodz</i>	
Author Index	301