

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

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

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Eitan Frachtenberg Uwe Schwiegelshohn (Eds.)

Job Scheduling Strategies for Parallel Processing

12th International Workshop, JSSPP 2006
Saint-Malo, France, June 26, 2006
Revised Selected Papers

Volume Editors

Eitan Frachtenberg

Los Alamos National Laboratory

Computer and Computational Sciences Division, Los Alamos, NM 87545, USA

E-mail: eitanf@lanl.gov

Uwe Schwiegelshohn

University of Dortmund

Robotics Research Institute (IRF-IT), 44221 Dortmund, Germany

E-mail: uwe.schwiegelshohn@udo.edu

Library of Congress Control Number: 2007920905

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

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

ISSN 0302-9743

ISBN-10 3-540-71034-5 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-71034-9 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 2007

Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper SPIN: 11980940 06/3142 5 4 3 2 1 0

Preface

This volume contains the papers presented at the 12th workshop on Job Scheduling Strategies for Parallel Processing. The workshop was held in Saint-Malo, France, on June 16, 2006, in conjunction with SIGMETRICS 2006.

This year, the presented papers covered a large variety of topics. The first three papers address workflow problems. “Provably efficient two-level adaptive scheduling” by Yuxiong He et al. provides a theoretical analysis of a scheduling approach for independent jobs consisting of threads, that are represented by a DAG. Job and thread scheduling are separately addressed with different algorithms. The task graph is not known a priori in the paper “Scheduling dynamically spawned processes in MPI-2” by Márcia Cera et al., but processes are spawned dynamically. This paper is based on the features of MPI-2 and evaluates its scheduler with the help of an experiment. The DAG of a Grid job is known at submission time in the problem discussed in the paper “Advance reservation policies for workflows” by Henan Zhao and Rizos Sakellariou. Here, the tasks of this job are automatically scheduled on heterogeneous machines using advance reservation such that the overall execution time frame of the user is obeyed. The proposed approach is again experimentally evaluated.

The next three papers describe classical job scheduling problems that arise when parallel jobs are submitted to parallel systems with little or no node heterogeneity. The paper “On advantages of scheduling using Genetic Fuzzy systems” by Carsten Franke et al. presents scheduling algorithms that support arbitrary scheduling criteria. The algorithms are trained with recorded workloads using Fuzzy concepts. Their performances are evaluated by simulations with those workloads. In their paper “Moldable parallel job scheduling using job efficiency: An iterative approach,” Gerald Sabin et al. show that scalability information of a job can help to improve the efficiency of this job. As in the previous paper, they use real workload traces for evaluation. The missing scalability information is provided with the help of a well-established speedup model. This model is also used in the paper “Adaptive job scheduling via predictive job resource allocation” by Lawrence Barsanti and Angela Sodan. Similar to the previous paper, the scalability of jobs improves the schedule performance. In addition, the resource allocation considers future job submissions based on a suitable prediction.

Many scientific applications are data intensive. For those applications, it is important to consider the network latency to transfer data from the storage facility to the parallel processing system. It is possible to improve schedule performance by scheduling those jobs on compute resources that are local to the storage resources. This is the subject of the paper “A data locality-aware online scheduling approach for I/O-intensive jobs with file sharing” by Gaurav Kanna et al. The next two papers address job migration issues. “Volunteer computing on clusters” by Deepti Vyas and Jaspal Subhlok demonstrates that nodes of

a compute cluster are often underutilized while executing parallel applications. Exploiting this observation by a cycle stealing approach will lead only to a small slowdown of the parallel host application while system throughput increases significantly. Idleness of processors is also the subject of the paper “Load balancing: Toward the Infinite Network and Beyond” by Javier Bustos-Jiménez. There, active objects are sent to underutilized processors that are determined with the help of a peer-to-peer approach. The performance of the approach is evaluated by an experiment with a real application and also by simulations. Jonathan Weinberg and Allan Snaveley observed in their paper “Symbiotic space-sharing on SDSC’s DataStar system” that the hierarchical architecture of modern parallel processing systems leads to a significant amount of resource sharing among independent jobs and thus to performance degradation. They propose to generate better schedules by considering combinations of jobs with minimum interference between them. Again the performance is evaluated with the help of experiments with real applications.

The last two papers address job modeling issues in Grid computing. “Modeling job arrivals in a data-intensive Grid” by Hui Li et al. analyzes job arrival processes in workloads from high-energy physics and uses a special Markov process to model them. Virtual organizations determine the granularity of the model. The paper “On Grid performance evaluation using synthetic workloads” by Alexandru Iosup et al. discusses various aspects of performance analysis. The authors review different performance metrics and show important properties of existing workloads. Then, they present workload modeling requirements that are specific for Grid computing.

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 for their willingness to participate in this effort and their excellent, detailed reviews: Su-Hui Chiang, Walfredo Cirne, Allen Downey, Dror Feitelson, Allan Gottlieb, Andrew Grimshaw, Moe Jette, Richard Lagerstrom, Virginia Lo, Jose Moreira, Bill Nitzberg, Mark Squillante, John Towns, Jon Weissman, and Ramin Yahyapour.

The continued interest in this area is reflected by the longevity of this workshop, which has now reached its 12th consecutive year. The proceedings of previous workshops are available from Springer as LNCS volumes 949, 1162, 1291, 1459, 1659, 1911, 2221, 2537, 2862, 3277, and 3834 (and since 1998 they have also been available online).

Finally, we would like to give our warmest thanks to Dror Feitelson and Larry Rudolph, the founding co-organizers of the workshop. Their efforts to promote this field are evidenced by the continuing success of this workshop.

Table of Contents

Provably Efficient Two-Level Adaptive Scheduling	1
<i>Yuzhong He, Wen-Jing Hsu, and Charles E. Leiserson</i>	
Scheduling Dynamically Spawned Processes in MPI-2.....	33
<i>Márcia C. Cera, Guilherme P. Pezzi, Maurício L. Pilla, Nicolas Maillard, and Philippe O.A. Navaux</i>	
Advance Reservation Policies for Workflows	47
<i>Henan Zhao and Rizos Sakellariou</i>	
On Advantages of Scheduling Using Genetic Fuzzy Systems	68
<i>Carsten Franke, Joachim Lepping, and Uwe Schwiegelshohn</i>	
Moldable Parallel Job Scheduling Using Job Efficiency: An Iterative Approach	94
<i>Gerald Sabin, Matthew Lang, and P. Sadayappan</i>	
Adaptive Job Scheduling Via Predictive Job Resource Allocation.....	115
<i>Lawrence Barsanti and Angela C. Sodan</i>	
A Data Locality Aware Online Scheduling Approach for I/O-Intensive Jobs with File Sharing.....	141
<i>Gaurav Khanna, Umit Catalyurek, Tahsin Kurc, P. Sadayappan, and Joel Saltz</i>	
Volunteer Computing on Clusters	161
<i>Deepti Vyas and Jaspal Subhlok</i>	
Load Balancing: Toward the Infinite Network and Beyond.....	176
<i>Javier Bustos-Jiménez, Denis Caromel, and José M. Piquer</i>	
Symbiotic Space-Sharing on SDSC's DataStar System	192
<i>Jonathan Weinberg and Allan Snively</i>	
Modeling Job Arrivals in a Data-Intensive Grid	210
<i>Hui Li, Michael Muskulus, and Lex Wolters</i>	
On Grid Performance Evaluation Using Synthetic Workloads	232
<i>Alexandru Iosup, Dick H.J. Epema, Carsten Franke, Alexander Papaspyrou, Lars Schley, Baiyi Song, and Ramin Yahyapour</i>	
Author Index	257