

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, Lancaster, 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, Zürich, Switzerland

John C. Mitchell

Stanford University, Stanford, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbrücken, Germany

More information about this series at <http://www.springer.com/series/7407>

Walfredo Cirne · Narayan Desai (Eds.)

Job Scheduling Strategies for Parallel Processing

18th International Workshop, JSSPP 2014
Phoenix, AZ, USA, May 23, 2014
Revised Selected Papers

Editors
Walfredo Cirne
Google
Mountain View, CA
USA

Narayan Desai
Ericsson
San Jose, CA
USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-15788-7 ISBN 978-3-319-15789-4 (eBook)
DOI 10.1007/978-3-319-15789-4

Library of Congress Control Number: 2015931917

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

Springer Cham Heidelberg New York Dordrecht London
© Springer International Publishing Switzerland 2015

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media
(www.springer.com)

Preface

This volume contains the papers presented at the 18th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP 2014), which was held in Phoenix, USA, on May 23, 2014 in conjunction with the IEEE International Parallel Processing Symposium 2014 (IPDPS 2014).

The proceedings of the previous workshops are also available from Springer-Verlag as LNCS volumes 949, 1162, 1291, 1459, 1659, 1911, 2221, 2537, 2862, 3277, 3834, 4376, 4942, 5798, 6253, 7698, and 8429. These volumes are available as printed books and online.

This year, we had 24 papers submitted to the workshop, of which we accepted nine. 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 the additional referees for their willingness to participate in this effort and their detailed, constructive reviews.

As a primary venue of the parallel scheduling community, the Job Scheduling Strategies for Parallel Processors workshop offers a good vantage point to witness its evolution. During these two decades, we have seen parallel scheduling grow in scope and importance, following the popularization of parallel systems. Fundamental issues in the area remain relevant today (e.g., scheduling goal and evaluation, workload modeling, and performance prediction). Meanwhile, a new set of issues have emerged, due to new workloads, increased scale, and the differing priorities of cloud systems. Together, the traditional and new issues make for a lively and discussion-rich workshop, where academic researchers and participants from industry meet and exchange ideas and experiences.

The workshop began with a keynote talk by Liana Fong, from IBM. She discussed how cognitive computing places new challenges for parallel job scheduling. These computations must not merely achieve performance levels that were unthinkable a decade ago, but do so in a very flexible manner, interactively steering the computation to support human cognitive processes.

At the paper presentations, the blend of old and new challenges in parallel job scheduling set the tone of the discussions. Even the most classical scenario of single node parallelism remained a very active area. This trend is no surprise, both as the number of cores sharing the same memory increases and special purpose parallel processors like GPUs have become prevalent.

We had three papers tackling single-core parallelism. Seo et al. introduced Bubble Task, which scheduled memory access among tasks running in a multi-core machine by throttling the tasks. Looking beyond memory contention, Herz and Pinkau addressed the problem of scheduling of task graphs on shared memory machine, with special focus on task graphs generated automatically. Finally, for the second time in the workshop history, we visited the intersection of parallel job scheduling and real-time systems. Qamhieh and Midonnet investigated the effects of parallelism in real-time

systems, showing that some reservations of the real-time community regarding parallelism were not fully justified.

Moving to distributed-memory, larger scale systems, scheduling fairness was a particularly hot topic this year. Klusáček and Rudová presented a new approach that supported multiresource aware user prioritization mechanism to ensure fairness. Importantly, this approach is capable of dealing with the heterogeneity of both jobs and resources. Rodrigo et al. investigated which basic prioritization primitives would make it easier to achieve fair scheduling in large, decentralized distributed. Tóth and Klusáček had a new take on the basic question of how to evaluate different scheduling algorithms. They proposed a user-centric approach that tries to measure the deviation of job end time from what the user would expect. The expectation of the user is based on fairness, which in itself generated a very interesting discussion on whether this model favors schedulers that strive for fairness.

The last three works presented this year touched different areas of parallel job scheduling. Kumar and Vadhiyar revisited the traditional question for batch systems of wait time prediction. They pushed the state of the art in prediction accuracy by using more sophisticated statistical models. They also evaluated how much meta-schedulers can improve performance by using such predictions.

Kuzmanovska et al. addressed how to schedule parallel jobs when these jobs are written using frameworks to deal with parallelism (e.g., MapReduce, Dryad, Pregel). They proposed a two-level approach in which each framework asks for the resources needed, and then distributes them to the jobs written in the framework.

Schwiegelshohn explored his long-time experience to distill lessons for scheduler algorithm designers. Starting with the provocative observation that most research papers in the area have had negligible impact on the state-of-practice, he identified constraints, objectives, and evaluation as the key reasons for this mismatch, and prescribed solutions based on his observations. The paper concluded by applying its own advice on designing a scheduling solution for an Infrastructure as a Service provider.

Enjoy the reading!

We hope you can join us in the next JSSPP workshop, in Hyderabad, India, in May 2015.

November 2014

Walfredo Cirne
Narayan Desai

Organization

Program Committee

Henri Casanova	University of Hawaii at Manoa, USA
Julita Corbalan	Technical University of Catalonia, Spain
Dick Epema	Delft University of Technology, The Netherlands
Gilles Fedak	Inria, France
Dror Feitelson	The Hebrew University, Israel
Liana Fong	IBM T.J. Watson Research Center, USA
Eitan Frachtenberg	Facebook, USA
Ali Ghodsi	University of California, Berkeley, USA
Alfredo Goldman	University of São Paulo, Brazil
Allan Gottlieb	New York University, USA
Alexandru Iosup	Delft University of Technology, The Netherlands
Morris Jette	SchedMD LLC, USA
Rajkumar Kettimuthu	Argonne National Laboratory, USA
Dalibor Klusáček	Masaryk University, Czech Republic
Zhiling Lan	Illinois Institute of Technology, USA
Bill Nitzberg	Altair Engineering, USA
Larry Rudolph	MIT, USA
Uwe Schwiegelshohn	Technical University of Dortmund, Germany
Mark Squillante	IBM T.J. Watson Research Center, USA
Murray Stokely	Google, USA
Wei Tang	Argonne National Laboratory, USA
Dan Tsafirir	Technion – Israel Institute of Technology, Israel
Ramin Yahyapour	GWDG – University of Göttingen, Germany

Contents

Bubble Task: A Dynamic Execution Throttling Method for Multi-core Resource Management	1
<i>Dongyou Seo, Myungsun Kim, Hyeonsang Eom, and Heon Y. Yeom</i>	
Real-World Clustering for Task Graphs on Shared Memory Systems	17
<i>Alexander Herz and Chris Pinkau</i>	
Experimental Analysis of the Tardiness of Parallel Tasks in Soft Real-Time Systems	36
<i>Manar Qamhieh and Serge Midonnet</i>	
Multi-resource Aware Fairsharing for Heterogeneous Systems	53
<i>Dalibor Klusáček and Hana Rudová</i>	
Priority Operators for Fairshare Scheduling	70
<i>Gonzalo P. Rodrigo, Per-Olov Östberg, and Erik Elmroth</i>	
User-Aware Metrics for Measuring Quality of Parallel Job Schedules	90
<i>Šimon Tóth and Dalibor Klusáček</i>	
Prediction of Queue Waiting Times for Metascheduling on Parallel Batch Systems	108
<i>Rajath Kumar and Sathish Vadhiyar</i>	
Dynamically Scheduling a Component-Based Framework in Clusters	129
<i>Aleksandra Kuzmanovska, Rudolf H. Mak, and Dick Epema</i>	
How to Design a Job Scheduling Algorithm	147
<i>Uwe Schwiegelshohn</i>	
Author Index	169