

Founding Editors

Gerhard Goos

Karlsruhe Institute of Technology, Karlsruhe, Germany

Juris Hartmanis

Cornell University, Ithaca, NY, USA

Editorial Board Members

Elisa Bertino

Purdue University, West Lafayette, IN, USA

Wen Gao

Peking University, Beijing, China

Bernhard Steffen 

TU Dortmund University, Dortmund, Germany

Gerhard Woeginger 

RWTH Aachen, Aachen, Germany

Moti Yung

Columbia University, New York, NY, USA

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

Dalibor Klusáček · Walfredo Cirne ·
Narayan Desai (Eds.)

Job Scheduling Strategies for Parallel Processing

23rd International Workshop, JSSPP 2020
New Orleans, LA, USA, May 22, 2020
Revised Selected Papers

Editors
Dalibor Klusáček
CESNET
Prague, Czech Republic

Narayan Desai
Google
Seattle, WA, USA

Walfredo Cirne
Google
Mountain View, CA, USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-63170-3 ISBN 978-3-030-63171-0 (eBook)
<https://doi.org/10.1007/978-3-030-63171-0>

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© Springer Nature Switzerland AG 2020

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, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

This volume contains the papers presented at the 23rd workshop on Job Scheduling Strategies for Parallel Processing (JSSPP 2020) that was held on May 22, 2020, in conjunction with the 34th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2020). The proceedings of previous workshops are also available from Springer as LNCS volumes 949, 1162, 1291, 1459, 1659, 1911, 2221, 2537, 2862, 3277, 3834, 4376, 4942, 5798, 6253, 7698, 8429, 8828, 10353, 10773, and 11332.

This year eight papers were submitted to the workshop, of which we accepted six. All submitted papers went through a complete review process, with the full version being read and evaluated by an average of four reviewers. Additionally, one invited paper and one keynote were included in the workshop. We would like to especially thank our Program Committee members and additional reviewers for their willingness to participate in this effort and their excellent, detailed, and thoughtful reviews.

For the first time in its history the JSSPP workshop was held fully online due to the worldwide COVID-19 pandemic. Despite the obvious logistic problems, all talks were presented live, allowing for the participants to interact with the authors of the papers. We are very thankful to the presenters of accepted papers for their participation in the live workshop session. Recordings from all talks at the 2020 edition can be found at the JSSPP's YouTube channel: <https://bit.ly/3mXyT8F>.

This year, the workshop opened with a keynote delivered by César De Rose from the PUCRS, School of Technology, Brazil. De Rose discussed interference-aware scheduling in virtualized environments, where multiple applications contending for shared resources are susceptible to cross-application interference, thus leading to possible significant performance degradation and consequently an increase in the number of broken SLAs. Therefore, interference-aware scheduling has gained traction, with the investigation of ways to classify applications regarding their interference levels and the proposal of static cost models and policies for scheduling co-hosted applications. The keynote was concluded with a demonstration of how interference-aware scheduling can improve resource usage while reducing SLA violations, with further opportunities for improvement in the areas of application classification and pro-active dynamic scheduling strategies.

Papers accepted for this year's JSSPP focused on several interesting problems within the resource management and scheduling domains. The first two papers focused on the problem of resource contention and workload interference. Yoonsung et al. discussed the performance degradation due to the contention for shared resources, such as cache and memory bandwidth. In this paper, the trade-offs between software and hardware isolation techniques were illustrated. Also, authors showed the benefit of coordinated enforcement of multiple isolation techniques.

Thiyyakat et al. presented a new scheduling policy that improves the performance of critical workload that is co-located with less important batch workloads. The authors

showed that their policy decreases the slowdown for critical workloads compared to a solution that used standard Control Groups (cgroups).

In their invited paper, Jaroš et al. discussed the problems related to scheduling ultrasound simulation workflows. They described how therapeutic ultrasound plays an increasing role in modern medicine. To optimize its benefits, the treatment procedures must be adapted carefully to patients needs by computing various DAG-like workflows that refine the parameters needed for the actual ultrasound machine. In their paper, authors discussed several scheduling problems that must be solved in order to execute these workflows efficiently.

Cavicchioli et al. approached the problem of under-utilizing available memory bandwidth when avoiding memory interference in systems that feature high-performance multi-core CPUs tightly integrated with data-parallel accelerators. They performed a set of experiments where they showed that the standard conservative approach that relies on exclusive use of shared main memory can be extended by injecting controlled amounts of memory requests coming from other tasks than the one currently granted exclusive DRAM access, thus using the available bandwidth more efficiently.

Nobre et al. proposed a highly optimized GPU+CPU based approach for epistasis detection. Epistasis (multiple interacting variations in DNA) detection is an important research topic in the field of DNA analysis as it allows to better understand various DNA variations that may cause, e.g., Alzheimer's disease, breast cancer, or Crohn's disease. As such, epistasis detection represents a computationally intensive optimization problem.

The sixth paper focused on walltime prediction and its impact on job scheduling performance and predictability. Job walltimes estimates, usually specified by users, are known to be very imprecise which causes problems both to the users and to the scheduling policies. Klusáček et al. presented an experimental analysis that demonstrated how the use of walltime predictors impacts the actual performance of a job scheduler as well its ability to provide accurate predictions concerning future job execution.

Last but not least, Geng et al. presented PDAWL, a novel dynamic approach for scheduling tasks that are capable of running simultaneously on both CPUs and general-purpose accelerators. It uses machine learning to build communication and computation performance estimation model of the workload with respect to the actual CPU and GPU performance. The online scheduler then adaptively adjusts the workload allocation based on the runtime situation.

We hope you can join us at the next JSPP workshop, this time in Portland, Oregon, USA, on May 21, 2021. Enjoy your reading!

September 2020

Dalibor Klusáček
Walfredo Cirne
Narayan Desai

Organization

Workshop Organizers

Dalibor Klusáček	CESNET, Czech Republic
Walfredo Cirne	Google, USA
Narayan Desai	Google, USA

Program Committee

Ashvin Agrawal	Microsoft, USA
Amaya Booker	Facebook, USA
Julita Corbalan	Barcelona Supercomputing Center, Spain
Stratos Dimopoulos	Apple, USA
Hyeonsang Eom	Seoul National University, South Korea
Dror Feitelson	Hebrew University, Israel
Liana Fong	IBM T. J. Watson Research Center, USA
Eitan Frachtenberg	Facebook, USA
Alfredo Goldman	University of São Paulo, Brazil
Allan Gottlieb	New York University, USA
Zhiling Lan	Illinois Institute of Technology, USA
Bill Nitzberg	Altair, USA
P-O. Östberg	Umeå University, Sweden
Gonzalo P. Rodrigo	Apple, USA
Larry Rudolph	Two Sigma, USA
Uwe Schwiegelshohn	TU Dortmund, Germany
Yingchong Situ	Google, USA
Leonel Sousa	Universidade de Lisboa, Portugal
Ramin Yahyapour	University of Göttingen, Germany

Additional Reviewers

Diogo Marques
Ricardo Nobre

Contents

Towards Interference-Aware Dynamic Scheduling in Virtualized Environments	1
<i>Vinícius Meyer, Uillian L. Ludwig, Miguel G. Xavier, Dionatrã F. Kirchoff, and Cesar A. F. De Rose</i>	
Towards Hybrid Isolation for Shared Multicore Systems	25
<i>Yoonsung Nam, Byeonghun Yoo, Yongjun Choi, Yongseok Son, and Hyeonsang Eom</i>	
Improving Resource Isolation of Critical Tasks in a Workload	45
<i>Meghana Thiyyakat, Subramaniam Kalambur, and Dinkar Sitaram</i>	
Optimizing Biomedical Ultrasound Workflow Scheduling Using Cluster Simulations	68
<i>Marta Jaros, Dalibor Klusáček, and Jiri Jaros</i>	
Evaluating Controlled Memory Request Injection to Counter PREM Memory Underutilization	85
<i>Roberto Cavicchioli, Nicola Capodieci, Marco Solieri, Marko Bertogna, Paolo Valente, and Andrea Marongiu</i>	
Accelerating 3-Way Epistasis Detection with CPU+GPU Processing	106
<i>Ricardo Nobre, Sergio Santander-Jiménez, Leonel Sousa, and Aleksandar Ilic</i>	
Walltime Prediction and Its Impact on Job Scheduling Performance and Predictability	127
<i>Dalibor Klusáček and Mehmet Soysal</i>	
PDAWL: Profile-Based Iterative Dynamic Adaptive WorkLoad Balance on Heterogeneous Architectures	145
<i>Tongsheng Geng, Marcos Amaris, Stéphane Zuckerman, Alfredo Goldman, Guang R. Gao, and Jean-Luc Gaudiot</i>	
Author Index	163