

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison, UK

Josef Kittler, UK

Friedemann Mattern, Switzerland

Moni Naor, Israel

Bernhard Steffen, Germany

Doug Tygar, USA

Takeo Kanade, USA

Jon M. Kleinberg, USA

John C. Mitchell, USA

C. Pandu Rangan, India

Demetri Terzopoulos, USA

Gerhard Weikum, Germany

Advanced Research in Computing and Software Science

Subline of Lecture Notes in Computer Science

Subline Series Editors

Giorgio Ausiello, *University of Rome 'La Sapienza', Italy*

Vladimiro Sassone, *University of Southampton, UK*

Subline Advisory Board

Susanne Albers, *TU Munich, Germany*

Benjamin C. Pierce, *University of Pennsylvania, USA*

Bernhard Steffen, *University of Dortmund, Germany*

Deng Xiaotie, *City University of Hong Kong*

Jeannette M. Wing, *Microsoft Research, Redmond, WA, USA*


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


Francisco F. Rivera · Tomás F. Pena
José C. Cabaleiro (Eds.)

Euro-Par 2017: Parallel Processing

23rd International Conference
on Parallel and Distributed Computing
Santiago de Compostela, Spain, August 28 – September 1, 2017
Proceedings

Editors

Francisco F. Rivera 
University of Santiago de Compostela
Santiago de Compostela
Spain

José C. Cabaleiro 
University of Santiago de Compostela
Santiago de Compostela
Spain

Tomás F. Pena 
University of Santiago de Compostela
Santiago de Compostela
Spain

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-64202-4 ISBN 978-3-319-64203-1 (eBook)
DOI 10.1007/978-3-319-64203-1

Library of Congress Control Number: 2017947501

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

© Springer International Publishing AG 2017

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. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Printed on acid-free paper

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

Preface

This volume contains the papers presented at Euro-Par 2017: the 23rd International Conference on Parallel and Distributed Computing, held from 28 August to 1 September 2017 in Santiago de Compostela (Spain).

Euro-Par is a prestigious annual series of international conferences dedicated to parallel and distributed computing. The topics covered by the conference include aspects related to both software and hardware technologies and, in particular, applications in different hardware platforms, ranging from small embedded systems to cloud computing and supercomputers. The specific topics on which the conference focuses have been renewed along the years extending the state of the art in the field. Nowadays, the challenges of building exascale performance computing systems and their programming are among the main motivations in the parallel and distributed computing community. This challenge opens opportunities to deal with issues related to health, climate, security, and many more. Various topics are deeply impacted by this scenario like energy optimization, scalability, heterogeneous computing, fault-tolerance, etc.

The main audience of Euro-Par are researchers in academic institutions, public and private laboratories, and industrial organizations. Euro-Par's main objective is to be the primary choice of such professionals for the presentation of new results in the field.

Previous Euro-Par conferences took place in Stockholm, Lyon, Passau, Southampton, Toulouse, Munich, Manchester, Paderborn, Klagensfurt, Pisa, Lisbon, Dresden, Rennes, Las Palmas, Delft, Ischia, Bordeaux, Rhodes, Aachen, Porto, Vienna, and Grenoble. This year Euro-Par 2017 was the 23rd conference and was organized in Santiago de Compostela, Spain, by the IT Research Centre of the University of Santiago de Compostela, called CiTIUS – Centro de Investigación en Tecnoloxías da Información. The topics were organized into 12 tracks, namely: Support Tools and Environments; Performance and Power Modelling, Prediction and Evaluation; Scheduling and Load Balancing; High-Performance Architectures and Compilers; Parallel and Distributed Data Management and Analytics; Cluster and Cloud Computing; Distributed Systems and Algorithms; Parallel and Distributed Programming, Interfaces, and Languages; Multicore and Manycore Parallelism; Theory and Algorithms for Parallel Computation and Networking; Parallel Numerical Methods and Applications; and Accelerator Computing. In all, 176 papers were submitted from 39 countries from all continents. Finally, only 50 papers were accepted in a selection meeting in which all the global or local chairs, as well as three members of the Steering Committee, participated. A selective rate of acceptance of 28.4% resulted: 691 reviews were performed by 317 experts; 151 papers received four reviews, 19 papers were reviewed by three experts, and 6 papers by five. The huge work of bringing many innovative ideas by the Scientific Committee made that the evaluation and selection processes proceed smoothly.

Apart from the parallel sessions to present the accepted papers, we were pleased to present two keynote talks of well-recognized colleagues, namely, David Padua “High-Level Abstractions and Automatic Optimization Techniques for the Programming of Irregular Algorithms,” and Jürgen Döllner “Software Analytics – Effectively Managing Complex Software Systems,” as well as an invited paper by Ian Foster et al. entitled “Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales.” The program was complemented by two days of dedicated workshops and tutorials on specialized topics. The huge task of managing them was efficiently conducted by Dr. Dora B. Heras. The selected papers will be published in separated proceedings volumes after the conference.

The Euro-Par conference in Santiago de Compostela would not have been possible without the support of many individuals and organizations. We owe special thanks to the authors of all the submitted papers, the members of the topic committees, in particular the global and local chairs, as well as the reviewers for their contributions to the success of the conference. We would also like to express our gratitude to the members of the Organizing Committee and the local staff who helped us. Moreover, we are indebted to the members of the Euro-Par Steering Committee, especially Christian Lengauer, Luc Bougé, and Fernando Silva, for their trust, guidance, and support. Finally, a number of institutional and industrial sponsors contributed to the organization of the conference. Their names appear on the Euro-Par 2017 website.

It was a pleasure and an honour to organize and host Euro-Par 2017 in Santiago de Compostela.

August 2017

Francisco F. Rivera
Tomás F. Pena
José C. Cabaleiro

Organization

Steering Committee

Full Members

Christian Lengauer (Chair)	University of Passau, Germany
Luc Bougé (Vice-Chair)	ENS Rennes, France
Emmanuel Jeannot	LaBRI-Inria, Bordeaux, France
Christos Kaklamanis	Computer Technology Institute, Patras, Greece
Paul Kelly	Imperial College, London, UK
Thomas Ludwig	University of Hamburg, Germany
Emilio Luque	University Autònoma of Barcelona, Spain
Tomàs Margalef	University Autònoma of Barcelona, Spain
Wolfgang Nagel	Dresden University of Technology, Germany
Rizos Sakellariou	University of Manchester, UK
Fernando Silva	University of Porto, Portugal
Henk Sips	Delft University of Technology, The Netherlands
Domenico Talia	University of Calabria, Italy
Jesper Larsson Träff	TU Vienna, Austria
Denis Trystram	Grenoble Institute of Technology, France
Felix Wolf	TU Darmstadt, Germany

Honorary Members

Ron Perrott	Oxford e-Research Centre, UK
Karl Dieter Reinartz	University of Erlangen-Nürnberg, Germany

Observers

Marco Aldinucci	University of Turin, Italy
Francisco F. Rivera	University of Santiago de Compostela, Spain

Euro-Par 2017 Organization

Chair

Francisco F. Rivera

Co-chairs

Tomás F. Pena
José C. Cabaleiro
Dora B. Heras

Proceedings

Tomás F. Pena
José C. Cabaleiro

Workshops

Dora B. Heras
Luc Bougé

Local Organization

Elisardo Antelo
Francisco Argüello
Antonio G. Loureiro
Juan C. Pichel
Natalia Seoane
David L. Vilarinho

Web and Publicity

Tomás F. Pena

Program Committee

Topic 1: Support Tools and Environments

Chair

Matthias Müller RWTH Aachen University, Germany

Local Chair

Andrés Gómez CESGA, Spain

Members

Martin Schulz	LLNL, Livermore, USA
Olivier Richard	LIG/Inria, France
João M.P. Cardoso	University of Porto, Portugal
Tomàs Margalef	Universitat Autònoma de Barcelona, Spain
Michael Gerndt	Technische Universität München, Germany

Topic 2: Performance and Power Modelling, Prediction, and Evaluation

Chair

Petr Tůma Charles University, Czech Republic

Local Chair

Basilio Fraguela University of A Coruña, Spain

Members

Ana Lucia Varbanescu	University of Amsterdam, The Netherlands
Denis Barthou	Inria, France
Lizy Kurian John	University of Texas, USA
Marc González Tallada	Universitat Politècnica de Catalunya, Spain
Andreas Knüpfer	T.U. Dresden, Germany
Diwakar Krishnamurthy	University of Calgary, Canada

Topic 3: Scheduling and Load Balancing**Chair**

Florina Ciorba University of Basel, Switzerland

Local Chair

Ester Garzón University of Almería, Spain

Members

José Luis Bosque Orero	University of Cantabria, Spain
Radu Prodan	University of Innsbruck, Austria
José Gracia	High-Performance Computing Center Stuttgart, Germany
Ioana Banicescu	Mississippi State University, USA
Julius Zilinskas	Vilnius University, Lithuania
Bora Uçar	CNRS and LIP ENS Lyon, France

Topic 4: High-Performance Architectures and Compilers**Chair**

Christophe Dubach University of Edinburgh, UK

Local Chair

Juan Touriño University of A Coruña, Spain

Members

Aaron Smith	Microsoft Research, USA
Louis-Nôel Pouchet	Colorado State University, USA
Laura Pozzi	University of Lugano, Switzerland
Jerónimo Castrillón	TU Dresden, Germany
Thomas Fahringer	University of Innsbruck, Austria
Chris Adeniyi-Jones	ARM, UK

Topic 5: Parallel and Distributed Data Management and Analytics

Chair

Bruno Raffin Inria, France

Local Chair

David E. Singh Carlos III University of Madrid, Spain

Members

Julian Kunkel	German Climate Computing Center, Germany
Lars Nagel	Johannes Gutenberg University of Mainz, Germany
Toni Cortés	Barcelona Supercomputing Center, Spain
Matthieu Dorier	Argonne National Laboratory, USA
Wolfgang Frings	Jülich Supercomputing Centre, Germany

Topic 6: Cluster and Cloud Computing

Chair

Alfredo Goldman University of São Paulo, Brazil

Local Chair

Patricia González University of A Coruña, Spain

Members

Laura Ricci	University of Pisa, Italy
Luiz Bittencourt	University of Campinas, Brazil
Ian Foster	Argonne National Laboratory, USA
Frédéric Desprez	Inria, France
Ivona Brandic	Technische Universität Wien, Austria
Giorgio Lucarelli	Inria, France
Rizos Sakellariou	University of Manchester, UK
Ramón Doallo	University of A Coruña, Spain

Topic 7: Distributed Systems and Algorithms

Chair

Luís Veiga INESC-ID, Portugal

Local Chair

Rafael Asenjo University of Málaga, Spain

Members

Sonia Ben Mokhtar	LIRIS CNRS, France
Óscar Plata González	University of Málaga, Spain
Gheorghe Almasi	IBM, USA
Rui Oliveira	Universidade do Minho, Portugal
Javier Navaridas Palma	The University of Manchester, UK
Fabio Kon	University of São Paulo, Brazil

Topic 8: Parallel and Distributed Programming, Interfaces, Languages**Chair**

María Jesús Garzarán	University of Illinois, USA
----------------------	-----------------------------

Local Chair

Vicente Blanco	University of La Laguna, Spain
----------------	--------------------------------

Members

M ^a Ángeles González Navarro	University of Málaga, Spain
Evelyn Duesterwald	T. J. Watson Research Center, IBM, USA
Didem Unat	Koç Universitesi, Turkey
Francisco Almeida	University of La Laguna, Spain
Georges da Costa	Irit, France
Marco Danelutto	University of Pisa, Italy
Mary Hall	University of Utah, USA

Topic 9: Multicore and Manycore Parallelism**Chair**

Hans Vandierendonck	Queen's University, UK
---------------------	------------------------

Local Chair

Juan Carlos Pichel	University of Santiago de Compostela, Spain
--------------------	---

Members

Bingsheng He	National University of Singapore, Singapore
Paul Harvey	Queen's University Belfast, UK
Michele Weiland	EPCC at University of Edinburgh, UK
Yiannis Nikolakopoulos	TU Chalmers, Sweden
Polyvios Pratikakis	FORTH, Greece
Martin Burtcher	Texas State University, USA

Georgios Goumas	National Technical University of Athens, Greece
Rutger Hofman	Vrije Universiteit Amsterdam, The Netherlands
Vania	Grenoble University, France
Marangozova-Martin	

Topic 10: Theory and Algorithms for Parallel Computation and Networking

Chair

Geppino Pucci	University of Padua, Italy
---------------	----------------------------

Local Chair

Pedro Ribeiro	University of Porto, Portugal
---------------	-------------------------------

Members

Kieran T. Herley	University College Cork, Ireland
Christos Zaroliagis	University of Patras, Greece
Mauro Bianco	Swiss National Supercomputing Centre, Switzerland
Henning Meyerhenke	Karlsruher Institut für Technologie, Germany
Michele Scquizzato	University of Houston, USA

Topic 11: Parallel Numerical Methods and Applications

Chair

Maya Neytcheva	Uppsala University, Sweden
----------------	----------------------------

Local Chair

María Martín	University of A Coruña, Spain
--------------	-------------------------------

Members

Yvan Notay	Université Libre de Bruxelles, Belgium
Peter Arbenz	ETH Zürich, Switzerland
Enrique S. Quintana	Jaime I University, Spain
Fred Wubs	University of Groningen, The Netherlands
Osni Marques	Lawrence Berkeley National Laboratory, USA

Topic 12: Accelerator Computing

Chair

Bertil Schmidt	Johannes Gutenberg University of Mainz, Germany
----------------	---

Local Chair

Arturo González	University of Valladolid, Spain
-----------------	---------------------------------

Members

Tobias Grosser	ETH Zürich, Switzerland
Josef Weidendorfer	Technische Universität München, Germany
Rob Van Nieuwpoort	Netherlands eScience Center, The Netherlands
Seyong Lee	Oak Ridge National Laboratory, USA
Jorge González-Domínguez	University of A Coruña, Spain
Deming Chen	University of Illinois, USA

Euro-Par 2017 Reviewers

Euro-Par is grateful to all the reviewers for their willingness and effort in providing good feedback to authors and topic committees. All external reviewers are listed and hereby thanked.

Abuín, José M.	Bonifaci, Vincenzo	Dichev, Kiril
Acosta, Alejandro	Braghetto, Kelly Rosa	Dolgov, Sergey
Aliaga, José Ignacio	Brown, Jed	Dorostkar, Ali
Alonso, Pedro	Brunie, Hugo	Dreher, Matthieu
Alonso-Monsalve, Saul	Buenabad-Chávez, Jorge	Dufossé, Fanny
Amannejad, Yasaman	Calotoiu, Alexandru	Durillo, Juan J.
Amor, Margarita	Carlini, Emanuele	Elafrou, Athena
Ananthanarayanan, Ganesh	Caron, Eddy	Eleliemy, Ahmed
Andión, José M.	Cała, Jacek	Erpen De Bona, Luis Carlos
Andrade, Diego	Cesar, Eduardo	Espinosa, Toni
Anzt, Hartwig	Chard, Kyle	Expósito, Roberto R.
Aparício, David	Chard, Ryan	Farsarakis, Emmanouil
Aral, Atakan	Chau, Vincent	Fernández, Javier
Arantes, Luciana	Chen, Lizhong	Flegar, Goran
Argüello, Francisco	Cheng, Xuntao	Flehmig, Martin
Arif, Mahwish	Cheptsov, Alexey	Foerster, Klaus-Tycho
Arlandini, Claudio	Chowdhury, Anamika	Frangoudis, Pantelis
Atalar, Aras	Chowdhury, Mosharaf	Fujita, Hajime
Azimi, Sahar	Chuvelev, Michael	Gabriel, Edgar
Baars, Sven	Coimbra, Miguel E.	Galante, Guilherme
Baiardi, Fabrizio	Cojean, Terry	García Blas, Javier
Bamha, Mostafa	Collange, Sylvain	Garcia, Islene
Barros Lourenço, Ricardo	Comprés Ureña, Isaías	Genez, Thiago
Belviranlı, Mehmet	Coplin, Jared	Georgakoudis, Giorgis
Benedict, Shajulin	Coppola, Massimo	Ghimire, Amrita
Benson, Austin	Cordeiro, Daniel	Giannoula, Christina
Bleuse, Raphaël	Costa, Fabio	Giménez, Domingo
Boisvert, Sébastien	Dai, Dong	Glantz, Roland
Boito, Francieli Zanon	De Maio, Vincenzo	Gonçalves, Rui
	De Sande, Francisco	

Gorman, Gerard	Lachaize, Renaud	Mulder, Thomas
Grelck, Clemens	Lago, Daniel	Mäsker, Markus
Greve, Fabiola	Lan, Haidong	Nachtmann, Mathias
Grimley Evans, Edmund	Lančinskas, Algirdas	Nadjaran Toosi, Adel
Gschwandtner, Philipp	Latham, Robert	Nagarajan, Arthi
Guidi, Barbara	Lebeane, Michael	Nasre, Rupesh
Gulur, Nagendra	Lee, Wooseok	Netto, Marco
Gupta, Abhishek	Letsios, Dimitrios	Nickolay, Sam
Gupta, Amit	Lirkov, Ivan	Niethammer, Christoph
Hager, Georg	Liu, Yongchao	Nou, Ramon
Haine, Christopher	Liu, Zhengchun	Nussbaum, Lucas
Hashemian Harandi, Raoufehsadat	Llanos, Diego	O'Neil, Molly
Hernández, Francisco	Llopis Sanmillán, Pablo	Oeste, Sebastian
Herold, Christian	Lobeiras Blanco, Jacobo	Ortega, Gloria
Herrera, Juan F.R.	Lopez Redondo, Juani	Padrón, Emilio
Hijma, Pieter	Lorenzo del Castillo, Juan Ángel	Pai, Sreepathi
Hirsch, Alex	Lujic, Ivan	Palka, Michal
Horký, Vojtěch	Lulli, Alessandro	Panda, Reena
Huchant, Pierre	Luque, Emilio	Papadopoulou, Nikela
Huedo, Eduardo	Lèbre, Adrien	Papagiannis, Anastasios
Hugo, Andra-Ecaterina	Madduri, Ravi	Papakonstantinou, Nikolaos
Hundt, Christian	Madeira, Edmundo	Pardo, Xoán C.
Hupp, Daniel	Magni, Alberto	Parsons, Mark
Hünich, Denis	Maia, Francisco	Pascual, Fanny
Igumenov, Aleksandr	Maleki, Saeed	Pascual, Jose A.
Ilic, Aleksandar	Mandli, Kyle	Passarella, Andrea
Iliev, Hristo	Mantas Ruiz, J. Miguel	Paul, Johns
Ilsche, Thomas	Marathe, Yashwant	Pérez, Borja
Iosup, Alexandru	Marinescu, María Cristina	Pérez Diéguez, Adrián
Jaiganesh, Jayadharini	Mathà, Roland	Pérez, Christian
Janetschek, Matthias	Mercier, Michael	Petcu, Dana
Jorba, Josep	Meyer, Marcel	Phan, Tien-Dat
Kalbasi, Amir	Michelogiannakis, George	Pietri, Ilia
Kamienski, Carlos	Mijaković, Robert	Pilla, Laércio L.
Kanellou, Eleni	Milenkovic, Aleksandar	Polato, Ivanilton
Karakostas, Vasileios	Miranda, Alberto	Poquet, Millian
Katsogridakis, Pavlos	Mohammed, Ali	Prieto-Matias, Manuel
Kavoussanakis, Konstantinos	Mommessin, Clement	Prokosch, Thomas
Kecskemeti, Gabor	Moreno-Vozmediano, Rafael	Protze, Joachim
Kimovski, Dragi	Moti, Nafiseh	Prountzos, Dimitrios
Klinkenberg, Jannis	Mounié, Grégory	Pruyne, Jim
Kotselidis, Christos	Moure, Juan Carlos	Queralt, Anna
Kruliš, Martin	Mouriño Gallego, José Carlos	Reid, Fiona
Kumaraswamy, Madhura	Mukherjee, Joydeep	Rexachs, Dolores
Küstner, Tilman		Reyes, Ruymán
		Richard, Jerome

Rico, Juan Antonio
Riedel, Morris
Ristov, Sasko
Rocki, Kamil
Rodríguez Gutiez,
 Eduardo
Rodríguez Martínez,
 Diego
Rodríguez, Gabriel
Ropars, Thomas
Ryoo, Jee Ho
Santana, Eduardo
Saurabh, Nishant
Schoene, Robert
Schuchart, Joseph
Schulz, Christian
Sclocco, Alessio
Sedaghati, Naser
Senger, Hermes
Seoane, Natalia
Shamakina, Anastasia
Shontz, Suzanne
Siakavaras, Dimitrios
Sikora, Anna

Silla, Federico
Silva, Miguel
Silva, Pedro Paulo
Singh, Shikhar
Sinnen, Oliver
Siqueira, Rodrigo
Sitchinava, Nodari
Skruzacek, Tyler
Solsona, Francesc
Song, Shuang
Srivastav, Abhinav
Srivastava, Srishti
Stafford, Esteban
Starikovicius, Vadimas
Strout, Michelle
Stylianopoulos,
 Charalampos
Symeonidou, Christi
Tchoua, Roselyne
Terboven, Christian
Termier, Alexandre
Tomás, Andrés
Tonello, Nicola
Torquati, Massimo

Toss, Julio
Tschueter, Ronny
Tzovas, Charilaos
van der Plas, P
van Werkhoven, Ben
Vázquez, Álvaro
Verdi, Fabio
Walulya, Ivan
Wang, Feiyi
Wang, Jiajun
Weber, Matthias
Wellein, Gerhard
Werner, Matthias
Winkler, Frank
Wozniak, Justin
Yang, Chih-Chieh
Yang, Dai
Zafari, Afshin
Zangerl, Peter
Zhani, Mohamed Faten
Zhou, Huan
Zois, Georgios

Euro-Par 2017 Invited Talks

High Level Abstractions and Automatic Optimization Techniques for the Programming of Irregular Algorithms

David Padua, University of Illinois at Urbana-Champaign, USA

High-performing irregular algorithms are typically implemented using simple operations and conventional control structures. In addition, due to today's compilers inability to manipulate these implementations, program tuning must usually be done by hand. Better notations and automatic optimization would help improve programmer productivity, portability, and maintainability. This talk will review high level notation proposals for the description of irregular algorithms, as well as compiler and autotuning techniques for the optimization of these algorithms. A short discussion of open research problems and necessary conditions for adoption of these more advanced notations and strategies will conclude the presentation.

Software Analytics – Effectively Managing Complex Software Systems

Jürgen Döllner, Hasso-Plattner-Institute for IT Systems Engineering, Germany

Digital transformation and industry 4.0 are among the key terms that reference a fundamental change in almost all branches of industry and society: Information technologies become essential building blocks of systems, applications, and processes. Access to and analytics for big data, along with machine learning, become key and competitive factors for transforming businesses in the next decade.

In this talk, we focus on and reflect how methods and techniques of big data analytics can be adapted and applied to the context of software engineering and IT industry. Here, the so called “software crisis” yet persists regardless of the manifold progress in programming concepts, languages, software modelling, software development methodologies, etc. Software analytics aims at boosting effectiveness of software development by providing new means of transparency within their corresponding ecosystems.

Euro-Par 2017 Topics Overview

Topic 1: Support Tools and Environments

*Matthias Müller, Andrés Gómez, Martin Schulz, Olivier Richard,
João M. P. Cardoso, Tomàs Margalef, Michael Gerndt*

Hardware and software of high performance computing (HPC) platforms are evolving every day. This evolution is very fast and is contributing to a very complex ecosystem. Applications must cope with large systems, with thousands of cores (even with millions in the largest HPC environments), several levels of memory hierarchy, hardware accelerators, heterogeneity, etc. Even more, it is becoming of paramount importance to extract the best performance with a strong control of the power consumption. Thus, HPC designers and programmers must have the tools to manage this complex scenario.

The Euro-Par Support Tools and Environments is a privileged forum to show new techniques and tools that allow all the stakeholders in the development and execution of HPC applications to manage the complexity involved, focusing on main challenges regarding programmability, resilience, performance and energy efficiency, monitoring, correctness, etc.

This track received 10 papers. After a reviewing process involving all the track TPC members and 21 external reviews, overall resulting in at least 4 reviews per paper, we decided to accept 3 of the submitted papers.

We acknowledge here the work of the reviewers who provided important feedback to the authors and helped us to select the best papers. Finally, we thank all authors who submitted papers. They really make this conference a key world event for presenting new Support Tools and Environments.

Topic 2: Performance and Power Modelling, Prediction, and Evaluation

*Petr Tůma, Basilio Fragueta, Ana Lucia Varbanescu, Denis Barthou,
Lizy Kurian John, Marc González Tallada, Andreas Knüpfer, Diwakar Krishnamurthy*

In recent years, a range of novel methods and tools have been developed for the evaluation, design, and modelling of parallel and distributed systems and applications. At the same time, the term ‘performance’ has broadened to also include scalability and energy efficiency, and touching reliability and robustness in addition to the classic resource-oriented notions. The aim of the ‘performance’ topic is to gather researchers working on different aspects of performance modelling, evaluation, and prediction, be it for systems or for applications running on the whole range of parallel and distributed systems (multi-core and heterogeneous architectures, HPC systems, grid and cloud contexts etc.)

This year, the track proved very popular, receiving a large number of submissions. Out of them, six papers were selected for presentation following a rigorous review process in which each manuscript received four independent reviews, either from the committee members or their subreviewers. We would like to thank all the authors who submitted papers to this topic as well as the external reviewers, for their contribution to the success of the conference.

Topic 3: Scheduling and Load Balancing

Florina M. Ciorba, Ester Garzón, José Luis Bosque Orero, Radu Prodan, José Gracia, Ioana Banicescu, Julius Zilinskas, Bora Uçar

New computer systems offer an opportunity to improve the performance and the energy consumption of the applications by the exploitation of several parallelism levels. Heterogeneity and complexity are the main characteristics of modern computer architectures. Thereby, the optimal exploitation of modern computing platforms becomes a challenge. The scheduling and load balancing techniques are relevant topics for the optimal exploitation of modern computers in terms of performance, energy consumption, cost of using resources, and so on.

This topic covered all aspects related to scheduling and load balancing on parallel and distributed systems, ranging from theoretical foundations for modelling and designing efficient and robust strategies, to experimental studies, applications, and practical tools and solutions. The main interest was focussed on modern multi/many-core processors, distributed/cloud platforms and data centres. The proposals to improve the performance were centred on the simulation of dynamic load balancing; scheduling based on genetic algorithms, approximations, and pinning; resource co-allocation; communications optimization; and graph partitioning.

A total of seventeen full-length submissions were received in this track, each of which received at least four reviews, from the eight program committee members and/or from the thirty-five additional sub-reviewers. Following the thorough discussion of the reviews, seven submissions have been accepted (42% acceptance rate), including one that was nominated as distinguished paper.

The chair and local chair sincerely thank all the authors for their submissions, the Euro-Par 2017 Organizing Committee for all their valuable help, and the reviewers and sub-reviewers for their excellent review work. Each has contributed to making this topic and Euro-Par an excellent forum to discuss Scheduling and Load Balancing challenges.

Topic 4: High Performance Architectures and Compilers

Christophe Dubach, Juan Touriño, Chris Adeniyi-Jones, Jerónimo Castrillón, Thomas Fahringer, Louis-Nöel Pouchet, Laura Pozzi, Aaron Smith

This topic deals with architecture design, languages, and compilation for parallel high performance systems. The areas of interest range from microprocessors to large-scale parallel machines (including multi-/many-core, possibly heterogeneous, architectures);

from general-purpose to specialized hardware platforms (e.g., graphic coprocessors, low-power embedded systems); and from architecture design to compiler technology and language design.

On the compilation side, topics of interest include programmer productivity issues, concurrent and/or sequential language aspects, vectorization, program analysis, program transformation, automatic discovery and/or management of parallelism at all levels, autotuning and feedback directed compilation, and the interaction between the compiler and the system at large. On the architecture side, the scope spans system architectures, processor micro-architecture, memory hierarchy, and multi-threading, architectural support for parallelism, and the impact of emerging hardware technologies.

The track received 13 submissions, all of which received, in a first stage, at least 3 reviews. In a second stage, all the papers and reviews were thoroughly discussed by all PC members. As a result, three papers were finally accepted for the conference (23% acceptance rate) covering both architecture and compiler topics.

Topic 5: Parallel and Distributed Data Management and Analytics

Bruno Raffin, David E. Singh, Julian Kunkel, Lars Nagel, Toni Cortés, Matthieu Dorier, Wolfgang Frings

Many areas of science, industry, and commerce are producing extreme-scale data that must be processed—stored, managed, analysed—in order to extract useful knowledge. This topic seeks papers in all aspects of distributed and parallel data management and data analysis. For example, HPC in situ data analytics, cloud and grid data-intensive processing, parallel storage systems, and scalable data processing workflows are all in the scope of this topic. More in detail, aspects in which this conference topic is interested are:

- Parallel, replicated, and highly-available distributed databases
- Cloud and HPC storage architectures and systems
- Scientific data analytics (Big Data or HPC based approaches)
- Middleware for processing large-scale data
- Programming models for parallel and distributed data analytics
- Workflow management for data analytics
- Coupling HPC simulations with in situ data analysis
- Parallel data visualization
- Distributed and parallel transaction, query processing and information retrieval
- Internet-scale data-intensive applications
- Sensor network data management
- Data-intensive clouds and grids
- Parallel data streaming and data stream mining
- New storage hierarchies in distributed data systems
- Parallel and distributed knowledge discovery and data mining

Thirteen full-length papers were submitted to this topic, and each paper received four reviews. After discussion with the reviewers and track chairs, two papers were selected for publication, one related to distributed database design, the second one to workload partitioning and scheduling algorithms for Apache Spark.

Topic 6: Cluster and Cloud Computing

Alfredo Goldman, Patricia González, Laura Ricci, Luiz Bittencourt, Ian Foster, Frédéric Desprez, Ivona Brandic, Giorgio Lucarelli, Rizos Sakellariou, Ramón Doallo

Cloud Computing is not a concept anymore, but a reality with many providers around the world. The use of massive storage and computing resources accessible remotely in a seamless way has become essential for many applications in various areas, including High Performance Computing. While significant progresses have been achieved in the past decade, the complete adoption of the Utility Computing paradigm is still facing important challenges. There are still unsolved challenges related to performance, reliability and energy efficiency of the infrastructures that should be addressed by research. Moreover, up to this time fundamental capabilities and services are required to achieve the goals of user-friendliness, security, privacy and service guarantees in such environments. Finally, there are important trends as going from large centralized infrastructures to smaller ones massively distributed at the edge of the network, and also to execute more efficiently High Performance Computing applications on Clouds.

Topic 6 sought papers covering many aspects of Cluster and Cloud Computing dealing with infrastructure layer challenges, such as performance/energy optimizations, and security enhancements, as well as cloud-enabled applications, workflow management and High Performance Computing on Clouds. This year, 26 papers have been submitted to Topic 6. There were authors from 18 different countries from all the continents. Four expert reviewers analysed each submission. Overall, more than 70 specialists were involved into the reviewing process. Finally, despite the high quality of the submitted papers, only 7 papers were accepted for publication.

We would like to thank all the authors for their submissions, the PC members and the reviewers for providing us with constructive and informative reviews, and the Euro-Par 2017 Organizing committee for all the help that allows us to smoothly take over the whole process.

Topic 7: Distributed Systems and Algorithms

Luís Veiga, Rafael Asenjo, Gheorghe Almasi, Sonia Ben Mokhtar, Fabio Kon, Javier Navaridas, Rui Oliveira, Oscar Plata

Parallel computing today is increasingly related to and dependent on developments and challenges of distributed systems. Problems including load balancing, asynchrony, failures, malicious and selfish behaviour, long latencies, network partitions, disconnected operations, distributed computing models and concurrent data structures, and

heterogeneity are representative of typical distributed issues that often appear along the design of parallel applications.

This track of Euro-Par provides a forum for both theoretical and practical research, of interest to both academia and industry, on distributed computing, distributed algorithms, distributed systems, distributed data structures, and parallel processing on distributed systems, in particular in relation to efficient high performance computing. This year, 8 complete papers have been submitted to this track. After a bidding phase, each paper has been evaluated by 4 or 5 reviewers with high expertise. Overall, 14 experts have been involved into the review process. Finally, from this set of high quality submitted papers, only three papers have been selected for publications.

The PC chairs, Luís Veiga (INESC-ID/IST, University of Lisbon, Portugal) and Rafael Asenjo (Universidad de Málaga, Spain), are very grateful to all the authors, and all researchers that have participated to the review process and permitted to select three high-quality papers.

Topic 8: Parallel and Distributed Programming, Interfaces, Languages

María Jesús Garzarán, Vicente Blanco, Didem Unat, Angeles Navarro, Mary Hall, Evelyn Duesterwald, Marco Danelutto, Francisco Almeida, George Da Costa

Parallel and distributed applications require adequate programming abstractions and models, efficient design tools, parallelization techniques and practices. This topic was open for the submission of new results and practical experience in this domain: efficient and effective parallel languages, interfaces, libraries and frameworks, as well as solid practical and experimental validation.

It provides a forum for research on high-performance, correct, portable, and scalable parallel programs via adequate parallel and distributed programming model, interface and language support. Contributions that assess programming abstractions, models and methods of usability, performance prediction, scalability, self-adaptations, rapid prototyping and fault-tolerance, as is needed, for instance, in dynamic heterogeneous parallel and distributed infrastructures, were accepted.

All twelve papers on this topic received four reviews that were further discussed among all nine PC members. As a result, four strong papers were accepted for the conference, covering important topics. One of them was proposed for the best paper award.

Topic 9: Multicore and Manycore Parallelism

Hans Vandierendonck, Juan Carlos Pichel, Bingsheng He, Paul Harvey, Michele Weiland, Yiannis Nikolakopoulos, Polyvios Pratikakis, Martin Burtscher, Georgios Goumas, Rutger Hofman, Vania Marangozova-Martin

Over the last ten years the trend in processor design has been towards an ever-increasing number of cores. The complexity of emerging many- and multi-core architectures makes it increasingly hard to program these devices efficiently. Efficient

algorithms must scale to large degrees of parallelism, use optimized data formats, minimize runtime system overhead and must use efficient synchronization mechanisms. Moreover, it is important to tune algorithms to the specific organization and dimensions of the target processor. The breadth of approaches that are investigated to achieve high-performance on multi- and many-core processors is a reflection of the complexity of these processors and the difficulty of designing algorithms that match the architecture of the processor.

This topic presents novel research contributions on a wide range of performance optimization techniques that are indispensable to programming multi- and many-cores, including efficient sparse matrix formats, optimization of linear algebra operations through batching, optimization of the fast multipole method on Intel many-cores, parallelization of remeshing algorithms, parallelization of model checking algorithms, thread-level speculation using transactional memories, non-blocking algorithms for radix trees, and concurrency-optimal search trees.

Eight papers out of 26 submissions were selected for publication in this track. All papers received at least 3 reviews.

We thank the authors who submitted papers, the PC members and referees who rigorously reviewed the submissions and provided constructive and informative feedback. We also thank the organizing committee for creating a smooth process and we look forward to an exciting edition of Euro-Par.

Topic 10: Theory and Algorithms for Parallel Computation and Networking

Geppino Pucci, Pedro Ribeiro, Mauro Bianco, Kieran T. Herley, Henning Meyerhenke, Michele Scquizzato, and Christos Zaroliagis

Parallel computing is everywhere, on smartphones, laptops; at online shopping sites, universities, computing centres; behind the search engines. Efficiency and productivity at these scales and contexts are only possible by scalable parallel algorithms using efficient communication schemes, routing and networks. Theoretical tools enabling scalability, modelling and understanding parallel algorithms, and data structures for exploiting parallelism are more important than ever. Topic 10 solicits high quality, original papers on the general topic of theory and algorithms for parallel computation including communication and network algorithms.

Topic 10 received 10 submissions, all of which received 4 reviews, either from the 7 PC members or from their subreviewers. The papers and their reviews were discussed extensively, and 2 submissions were eventually accepted.

Topic 11: Parallel Numerical Methods and Applications

Maya Neytcheva, María Martín, Yvan Notay, Peter Arbenz, Enrique S. Quintana, Fred Wubs, Osni Marques

The demand for high performance computations is driven by the need for large-scale computer simulations in nearly all activity areas - science and engineering, finance, life sciences etc. In turn, high performance computing goes hand in hand with the necessity to develop highly scalable numerical methods and algorithms that are able to efficiently exploit modern computer architectures and to fully utilize their computing power. The scalability of these algorithms and methods and their suitability to efficiently utilize the available high performance, but in general heterogeneous, computer resources, is a key point to improve the performance of the target applications and enable fast and reliable computer simulations.

This conference topic aims at presenting and discussing recent developments in parallel numerical algorithms and their implementation on current parallel architectures, including many-core and hybrid architectures.

This year the topic received 9 contributions. Each submission was reviewed by at least four reviewers. Overall, 27 experts have been involved into the review process. Finally, three papers were accepted for presentation. We thank all authors for their valuable contributions, as well as the Program Committee members and the external reviewers for investing their time, sharing their expertise and keeping the high scientific level of the Euro-Par conference.

Topic 12: Accelerator Computing

Bertil Schmidt, Arturo González, Tobias Grosser, Josef Weidendorfer, Rob Van Nieuwpoort, Seyong Lee, Jorge González-Domínguez, Deming Chen

The need for high-performance computing is constantly growing in all kind of scenarios, from high-end scientific applications, to consumer electronics software. Hardware manufactures are involved in a race to develop specialized hardware to cover these critical demands.

Currently, hardware accelerators of various kinds offer a potential for achieving massive performance in applications that can leverage their high degree of parallelism and customization. Examples include graphics processors (GPUs), manycore co-processors, as well as more customizable devices, such as FPGA-based systems or streaming data-flow architectures.

The research challenge for this topic is to explore new directions for actually realizing this potential. Significant advances in all areas related to accelerators are considered, with special focus in architectures, algorithms, languages, compilers, libraries, runtime systems, coordination of accelerators and CPU, debugging and profiling tools, and application-related contributions that provide new insights into fundamental problems or solution approaches in this domain.

The program committee of this topic was formed by seven members of different backgrounds and specializations in the accelerators field, with the collaboration of several other subreviewers. We received 13 contributions from researchers in many different countries. After the review process and the general PC meeting, two high-quality papers were selected for presentation in Euro-Par 2017 at Santiago de Compostela. They are focused on important hot-topics: exploiting the GPUs potential on sparse linear algebra, and the question of load balancing for performance or energy.

The committee members want to thank all the authors that submitted their work to this track, the reviewers for their timely and constructive comments, and the organization committee for the efforts to ease our task, and to provide a nice conference environment in Santiago de Compostela for a high-quality discussion of research results in this interesting topic.

Contents

Invited Paper

Computing Just What You Need: Online Data Analysis and Reduction at Extreme Scales	3
<i>Ian Foster, Mark Ainsworth, Bryce Allen, Julie Bessac, Franck Cappello, Jong Youl Choi, Emil Constantinescu, Philip E. Davis, Sheng Di, Wendy Di, Hanqi Guo, Scott Klasky, Kerstin Kleese Van Dam, Tahsin Kurc, Qing Liu, Abid Malik, Kshitij Mehta, Klaus Mueller, Todd Munson, George Ostouchov, Manish Parashar, Tom Peterka, Line Pouchard, Dingwen Tao, Ozan Tugluk, Stefan Wild, Matthew Wolf, Justin M. Wozniak, Wei Xu, and Shinjae Yoo</i>	

Support Tools and Environments

Scaling Energy Adaptive Applications for Sustainable Profitability	23
<i>Fabien Hermenier, Giuliani Giovanni, Andre Milani, and Sophie Demassey</i>	
Off-Road Performance Modeling – How to Deal with Segmented Data	36
<i>M. Kashif Ilyas, Alexandru Calotoiu, and Felix Wolf</i>	
Online Dynamic Monitoring of MPI Communications	49
<i>George Bosilca, Clément Foyer, Emmanuel Jeannot, Guillaume Mercier, and Guillaume Papauré</i>	

Performance and Power Modeling, Prediction and Evaluation

Micro-benchmarking MPI Neighborhood Collective Operations	65
<i>Felix Donatus Lübke</i>	
Performance Characterization of De Novo Genome Assembly on Leading Parallel Systems	79
<i>Marquita Ellis, Evangelos Georganas, Rob Egan, Steven Hofmeyr, Aydin Buluç, Brandon Cook, Leonid Oliker, and Katherine Yelick</i>	
NVIDIA Jetson Platform Characterization	92
<i>Hassan Halawa, Hazem A. Abdelhafez, Andrew Boktor, and Matei Ripeanu</i>	

Following the Blind Seer – Creating Better Performance Models Using Less Information	106
<i>Patrick Reisert, Alexandru Calotoiu, Sergei Shudler, and Felix Wolf</i>	
An Accurate Simulator of Cache-Line Conflicts to Exploit the Underlying Cache Performance	119
<i>Yukinori Sato and Toshio Endo</i>	
Shutdown Policies with Power Capping for Large Scale Computing Systems.	134
<i>Anne Benoit, Laurent Lefèvre, Anne-Cécile Orgerie, and Issam Raïs</i>	
Scheduling and Load Balancing	
Partitioning Strategy Selection for In-Memory Graph Pattern Matching on Multiprocessor Systems.	149
<i>Alexander Krause, Thomas Kissinger, Dirk Habich, Hannes Voigt, and Wolfgang Lehner</i>	
Efficient Dynamic Pinning of Parallelized Applications by Reinforcement Learning with Applications.	164
<i>Georgios C. Chasparis, Michael Rossbory, and Vladimir Janjic</i>	
Accelerating by Idling: How Speculative Delays Improve Performance of Message-Oriented Systems	177
<i>Aleksandar Prokopec</i>	
Using Simulation to Evaluate and Tune the Performance of Dynamic Load Balancing of an Over-Decomposed Geophysics Application	192
<i>Rafael Keller Tesser, Lucas Mello Schnorr, Arnaud Legrand, Fabrice Dupros, and Philippe Olivier Alexandre Navaux</i>	
Optimizing Egalitarian Performance in the Side-Effects Model of Colocation for Data Center Resource Management	206
<i>Fanny Pascual and Krzysztof Rzađca</i>	
Generic Algorithms for Scheduling Applications on Hybrid Multi-core Machines	220
<i>Marcos Amaris, Giorgio Lucarelli, Clément Mommessin, and Denis Trystram</i>	
Low-Cost Approximation Algorithms for Scheduling Independent Tasks on Hybrid Platforms.	232
<i>Louis-Claude Canon, Loris Marchal, and Frédéric Vivien</i>	

High Performance Architectures and Compilers

Runtime-Assisted Shared Cache Insertion Policies Based on Re-reference Intervals	247
<i>Vladimir Dimić, Miquel Moretó, Marc Casas, and Mateo Valero</i>	
Rewriting System for Profile-Guided Data Layout Transformations on Binaries	260
<i>Christopher Haine, Olivier Aumage, and Denis Barthou</i>	
Hardware Support for Scratchpad Memory Transactions on GPU Architectures	273
<i>Alejandro Villegas, Rafael Asenjo, Angeles Navarro, Oscar Plata, Rafael Ubal, and David Kaeli</i>	

Parallel and Distributed Data Management and Analytics

Execution of Recursive Queries in Apache Spark	289
<i>Pavlos Katsogridakis, Sofia Papagiannaki, and Polyvios Pratikakis</i>	
Replica-Aware Partitioning Design in Parallel Database Systems.	303
<i>Liming Dong, Weidong Liu, Renchuan Li, Tiejun Zhang, and Weiguo Zhao</i>	

Cluster and Cloud Computing

A Simplified Model for Simulating the Execution of a Workflow in Cloud.	319
<i>Roland Mathá, Sasko Ristov, and Radu Prodan</i>	
Dealing with Performance Unpredictability in an Asymmetric Multicore Processor Cloud	332
<i>Boris Teabe, Patrick Lavoisier Wapet, Alain Tchana, and Daniel Hagimont</i>	
Deadline-Aware Deployment for Time Critical Applications in Clouds	345
<i>Yang Hu, Junchao Wang, Huan Zhou, Paul Martin, Arie Taal, Cees de Laat, and Zhiming Zhao</i>	
More Sharing, More Benefits? A Study of Library Sharing in Container-Based Infrastructures.	358
<i>José Bravo Ferreira, Marco Cello, and Jesús Omana Iglesias</i>	
An Efficient Communication Aware Heuristic for Multiple Cloud Application Placement.	372
<i>Pedro Silva and Christian Perez</i>	

Energy-Driven Straggler Mitigation in MapReduce	385
<i>Tien-Dat Phan, Shadi Ibrahim, Amelie Chi Zhou, Guillaume Aupy, and Gabriel Antoniu</i>	

Leveraging Cloud Heterogeneity for Cost-Efficient Execution of Parallel Applications	399
<i>Eduardo Roloff, Matthias Diener, Emmanuell Diaz Carreño, Luciano Paschoal Gaspary, and Philippe O.A. Navaux</i>	

Distributed Systems and Algorithms

A Consensus-Based Fault-Tolerant Event Logger for High Performance Applications.	415
<i>Edson Tavares de Camargo, Elias P. Duarte Jr., and Fernando Pedone</i>	

Families of Graph Algorithms: SSSP Case Study	428
<i>Thejaka Amila Kanewala, Marcin Zalewski, and Andrew Lumsdaine</i>	

SEMem: Deployment of MPI-Based In-Memory Storage for Hadoop on Supercomputers	442
<i>Thanh-Chung Dao and Shigeru Chiba</i>	

Parallel and Distributed Programming, Interfaces, and Languages

Supporting the Xeon Phi Coprocessor in a Heterogeneous Programming Model	457
<i>Ana Moreton-Fernandez, Eduardo Rodriguez-Gutierrez, Arturo Gonzalez-Escribano, and Diego R. Llanos</i>	

GLT: A Unified API for Lightweight Thread Libraries	470
<i>Adrián Castelló, Sangmin Seo, Rafael Mayo, Pavan Balaji, Enrique S. Quintana-Ortí, and Antonio J. Peña</i>	

PASCAL: A Parallel Algorithmic SCALable Framework for N -body Problems	482
<i>Laleh Aghababaie Beni and Aparna Chandramowlishwaran</i>	

GASPI/GPI In-memory Checkpointing Library	497
<i>Valeria Bartsch, Rui Machado, Dirk Merten, Mirko Rahn, and Franz-Josef Pfreundt</i>	

Multicore and Manycore Parallelism

Optimized Batched Linear Algebra for Modern Architectures	511
<i>Jack Dongarra, Sven Hammarling, Nicholas J. Higham, Samuel D. Relton, and Mawussi Zounon</i>	

New Efficient General Sparse Matrix Formats for Parallel SpMV Operations	523
<i>Jan Philipp Ecker, Rudolf Berrendorf, and Florian Mannuss</i>	
Lazy Parallel Kronecker Algebra-Operations on Heterogeneous Multicores. . .	538
<i>Wasuwee Sodsong, Robert Mittermayr, Yoojin Park, Bernd Burgstaller, and Johann Blieberger</i>	
Performance Evaluation of Computation and Communication Kernels of the Fast Multipole Method on Intel Manycore Architecture	553
<i>Mustafa Abduljabbar, Mohammed Al Farhan, Rio Yokota, and David Keyes</i>	
Efficient Non-blocking Radix Trees.	565
<i>Varun Velamuri</i>	
A Concurrency-Optimal Binary Search Tree.	580
<i>Vitaly Aksenov, Vincent Gramoli, Petr Kuznetsov, Anna Malova, and Srivatsan Ravi</i>	
Scalable Fine-Grained Metric-Based Remeshing Algorithm for Manycore/NUMA Architectures.	594
<i>Hoby Rakotoarivelo, Franck Ledoux, Franck Pommereau, and Nicolas Le-Goff</i>	
Performance Evaluation of Thread-Level Speculation in Off-the-Shelf Hardware Transactional Memories.	607
<i>Juan Salamanca, José Nelson Amaral, and Guido Araujo</i>	
Theory and Algorithms for Parallel Computation and Networking	
Addressing Volume and Latency Overheads in 1D-parallel Sparse Matrix-Vector Multiplication	625
<i>Seher Acer, Oguz Selvitopi, and Cevdet Aykanat</i>	
Improving the Network of Search Engine Services Through Application-Driven Routing	638
<i>Joe Carrión, Daniel Franco, Veronica Gil-Costa, Mauricio Marin, and Emilio Luque</i>	
Parallel Numerical Methods and Applications	
Accelerating the Tucker Decomposition with Compressed Sparse Tensors . . .	653
<i>Shaden Smith and George Karypis</i>	
Shared Memory Pipelined Parareal	669
<i>Daniel Ruprecht</i>	

Nonintrusive AMR Asynchrony for Communication Optimization. 682
*Muhammad Nufail Farooqi, Didem Unat, Tan Nguyen, Weiqun Zhang,
Ann Almgren, and John Shalf*

Accelerator Computing

Balanced CSR Sparse Matrix-Vector Product on Graphics Processors 697
Goran Flegar and Enrique S. Quintana-Ortí

To Distribute or Not to Distribute: The Question of Load Balancing
for Performance or Energy 710
*Esteban Stafford, Borja Pérez, Jose Luis Bosque, Ramón Beivide,
and Mateo Valero*

Author Index 723