Lecture Notes in Computer Science **11014**

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Advanced Research in Computing and Software Science
Subline of Lecture Notes in Computer Science

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

Marco Aldinucci · Luca Padovani
Massimo Torquati (Eds.)

Euro-Par 2018: Parallel Processing

24th International Conference on Parallel and Distributed Computing Turin, Italy, August 27–31, 2018 Proceedings

Springer

Editors
Marco Aldinucci
Department of Computer Science
University of Torino
Torino
Italy

Massimo Torquati
Department of Computer Science
University of Pisa
Pisa
Italy

Luca Padovani
Department of Computer Science
University of Torino
Torino
Italy

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-96982-4 ISBN 978-3-319-96983-1 (eBook)
https://doi.org/10.1007/978-3-319-96983-1

Library of Congress Control Number: 2018949144

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

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 Euro-Par 2018, the 24th International European Conference on Parallel and Distributed Computing, held during August 27–31, 2018, in Turin, Italy. The whole computer hardware industry has embraced parallel computing. Today it is clear that in the long term, writing efficient, portable, and correct parallel programs must be no more challenging than writing the same programs for sequential computers. Euro-Par envisioned this challenge 24 years ago. It gracefully evolved from pioneering efforts to the mainstream of parallel computing maintaining its broad-spectrum coverage on parallel computing topics, from software to hardware to their co-design. Its adaptability and capability to frame emerging topics in an independent, consolidated structure is still the key to its success.

The main audience of Euro-Par comprises researchers in academia, 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, Klagenfurt, Pisa, Lisbon, Dresden, Rennes, Las Palmas, Delft, Ischia, Bordeaux, Rhodes, Aachen, Porto, Vienna, Grenoble, and Santiago de Compostela. The 24th edition of Euro-Par was organized by the Computer Science Department of the University of Turin. The topics were organized into 12 tracks, namely: Support Tools and Environments; Performance and Power Modeling, 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 Methods and Tools; Theory and Algorithms for Parallel Computation and Networking; Parallel Numerical Methods and Applications; and Accelerator Computing for Advanced Applications. Overall, 194 papers were submitted from 39 countries. The number of submitted papers, the wide topic coverage and the aim of obtaining high-quality reviews resulted in a difficult selection process involving a large number of experts. The join effort of the members of the Scientific Committee and of the 306 external reviewers resulted in 787 reviews: five papers received three reviews, 173 received four reviews and 16 received five, that is, on average, 4.06 reviews per paper. The accepted papers were chosen after lengthy discussions and finalized during the physical selection meeting, which took place on April 27, 2018 in Turin. All local chairs and three members of the Steering Committee participated in the meeting. In the end, 57 papers were selected to be presented at the conference and published in the proceedings, resulting in a 29.4% acceptance rate. The following three papers were nominated as "distinguished" and presented in a plenary session: "Resource-Efficient Execution of Conditional Parallel Real-Time Tasks," "VIoLET: A Large-Scale Virtual Environment for Internet of Things," "Design Principles for Sparse Matrix Multiplication on the GPU."

Apart the presentation sessions of accepted papers, we were honored to host three keynote talks given by esteemed colleagues, namely: "ALGORAND: A Better Distributed Ledger" by Silvio Micali, "Algorithmic Adaptations to Extreme Scale Computing" by David E. Keyes, and "Datacenters for the Post-Moore Era" by Babak Falsafi. The conference program was complemented by two days of workshops and tutorials on specialized topics. Dora B. Heras and Gabriele Mencagli deserve recognition for managing them efficiently and effectively. A selection of the papers presented at the workshops will be published in separated proceedings volumes after the conference.

With respect to previous editions of Euro-Par, the 2018 edition introduced two novelties both aimed at improving the relevance and impact of the scientific works presented at the conference. For the first time in the history of Euro-Par, authors of accepted papers were encouraged to submit an *artifact* (e.g., source code, tools, benchmarks, datasets, models) to assess the reproducibility of the experimental results presented in the paper. Overall, 13 artifacts were submitted and all of them were positively evaluated by a separate Artifact Evaluation Committee. Papers with an associated artifact received a seal of approval in the proceedings and Springer kindly agreed to permanently host all artifacts on their servers. Although the practice of evaluating artifacts is becoming commonplace in other computer science conferences, it should be mentioned that it poses substantial challenges in a conference like Euro-Par, in which artifacts may require large and dedicated hardware infrastructures and may involve the processing of gigabytes of data for long periods of time. The resulting additional effort required of the organizers and the members of the Artifact Evaluation Committee was largely compensated by the enthusiasm of the authors who decided to submit an artifact. In the end, nearly one quarter of the papers had an associated artifact. The second novelty experimented with in Euro-Par 2018 was a session of *chess-timer talks*, in which the audience was encouraged to interact with the speakers, and the session chair balanced solo presentation and discussions using a chess timer. We got the idea of proposing chess-timer talks from the successful CurryOn conference as an "unusual solution to making tech conferences a more interactive, more fun, and better place for learning and discussions."

The Euro-Par conference in Turin 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. We are indebted to the members of the Euro-Par Steering Committee, especially Christian Lengauer, Luc Bougé, and Fernando Silva, for their trust, constant guidance, and support. We are grateful to the staff of Springer, particularly Anna Kramer, Alfred Hofmann, and Graham Smith, for their support in the preparation of the proceedings and the management of artifacts. Finally, a number of institutional and industrial sponsors

contributed to the organization of the conference. Their names appear on the Euro-Par 2018 website.

It was a pleasure and an honor to organize and host Euro-Par 2018 in Turin.

June 2018

Marco Aldinucci
Luca Padovani
Massimo Torquati

Organization

Steering Committee

Full Members

Luc Bougé (Chair)	ENS Rennes, France
Fernando Silva (Vice-Chair)	University of Porto, Portugal
Dora B. Heras (Workshops Chair)	CiTIUS, Santiago de Compostela, Spain
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
Tomàs Margalef	University Autonoma of Barcelona, Spain
Wolfgang Nagel	Dresden University of Technology, Germany
Francisco Fernández Rivera	CiTIUS, Santiago de Compostela, Spain
Rizos Sakellariou	University of Manchester, UK
Henk Sips	Delft University of Technology, The Netherlands
Domenico Talia	University of Calabria, Italy
Jesper Larsson Träff	TU Wien, Austria
Denis Trystram	Grenoble Institute of Technology, France
Felix Wolf	TU Darmstadt, Germany

Honorary Members

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

Observers

Marco Aldinucci	University of Turin, Italy
Ramin Yahyapour	GWDG/University of Göttingen, Germany

Euro-Par 2018 Organization

Co-chairs

Marco Aldinucci	University of Turin, Italy
Luca Padovani	University of Turin, Italy
Massimo Torquati	University of Pisa, Italy

Workshops

Dora B. Heras	University of Santiago de Compostela, Spain
Gabriele Mencagli	University of Pisa, Italy

Tutorials and Industry

Peter Kilpatrick	Queen's University Belfast, UK
Claudia Misale	IBM T.J. Watson Research Center, USA

Publicity

Javier Garcia Blas	Universidad Carlos III de Madrid, Spain
Dalvan Griebler	Pontifical Catholic University of Rio Grande do Sul, Brazil
Ivan Merelli	CNR-ITB, Italy
Fabio Tordini	University of Turin, Italy

Submission and Reviewing Process

Daniele D'Agostino	CNR-IMATI, Italy
Massimo Torquati	University of Pisa, Italy

Artifact Evaluation Process

Gabriele Mencagli	University of Pisa, Italy
Luca Padovani	University of Turin, Italy
Massimo Torquati	University of Pisa, Italy

Web and Social Media

Claudio Mattutino	University of Turin, Italy
Sergio Rabellino	University of Turin, Italy

Logistics

Katia Lupo	University of Turin, Italy
Claudio Mattutino	University of Turin, Italy
Sergio Rabellino	University of Turin, Italy

Program Committee

Topic 1: Support Tools and Environments

Global Chair

Siegfried Benkner	Universität Wien, Austria

Local Chair

Massimo Coppola	CNR-ISTI, Italy

Chairs

Franz Franchetti	Carnegie Mellon University, USA
Michael Gerndt	Technische Universität München, Germany
Erwin Laure	Royal Institute of Technology, Sweden
Nikos Parlavantzas	INSA Rennes and IRISA, France

Topic 2: Performance and Power Modeling, Prediction and Evaluation

Global Chair

Leonel Sousa	Universidade de Lisboa, Portugal

Local Chair

Daniele De Sensi	University of Pisa, Italy

Chairs

Giorgis Georgakoudis	Queen's University Belfast, UK
Aleksandar Ilic	Universidade de Lisboa, Portugal
Piotr Luszczek	University of Tennessee, USA
Federico Silla	Universitat Politecnica de Valencia, Spain
Guangming Tan	Chinese Academy of Sciences, China
Pedro Trancoso	Chalmers University of Technology, Sweden

Topic 3: Scheduling and Load Balancing

Global Chair

Anne Benoitć	ENS Lyon, France

Local Chair

Enrico Bini	University of Turin, Italy

Chairs

Maciej Drozdowski	Poznan University of Technology, Poland
Lionel Eyraud-Dubois	LaBRI-inria, Bordeaux, France
José Gracia	HLRS, Stuttgart, Germany
Nan Guan	Hong Kong Polytechnic University, Hong Kong, SAR China
Sascha Hunold	TU Wien, Austria
Krzysztof Rzadca	University of Warsaw, Poland

Topic 4: High-Performance Architectures and Compilers

Global Chair

Florian Brandner	Télécom ParisTech, Université Paris-Saclay, France

Local Chair

Fabio Luporini	Imperial College London, UK

Chairs

Alexandra Jimborean	Uppsala University, Sweden
Frank Hannig	University of Erlangen-Nürnberg, Germany
Gihan Mudalige	University of Warwick, UK

Topic 5: Parallel and Distributed Data Management and Analytics

Global Chair

K. Selçuk Candan	Arizona State University, USA

Local Chair

Ruggero Pensa	University of Turin, Italy

Chairs

Lei Chen	Hong Kong University of Science and Technology, SAR China
Gianmarco De Francisci Morales	Qatar Computing Research Institute, Qatar
Ming Zhao	Arizona State University, USA

Topic 6: Cluster and Cloud Computing

Global Chair

Ivona Brandić	Vienna University of Technology, Austria

Local Chair

Domenico Talia	Università della Calabria, Italy

Chairs

Toni Mastelic	Vienna University of Technology, Austria
Raffaele Montella	Università degli Studi di Napoli Parthenope, Italy
Anne-Cécile Orgerie	CNRS, Rennes, France
Thomas Renner	Technische Universität Berlin, Germany
Rafael Brundo Uriarte	IMT School for Advanced Studies, Lucca, Italy

Topic 7: Distributed Systems and Algorithms

Global Chair

Sonia Ben-Mokhtar	LIRIS, INSA de Lyon, France

Local Chair

Alberto Montresor	University of Trento, Italy

Chairs

Christof Fetzer	TU Dresden, Germany
Indranil Gupta	University of Illinois, Urbana-Champaign, USA

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

Global Chair

J. Daniel García	University Carlos III of Madrid, Spain

Local Chair

Patrizio Dazzi	CNR-ISTI, Italy

Chairs

Bryce Adelstein-Lelbach	NVIDIA, USA
Marcelo Pasin	Université de Neuchâtel, Switzerland
Mitsuhisa Sato	Riken and University of Tsukuba, Japan
Paolo Trunfio	DIMES, University of Calabria, Italy
Chan-Hyun Youn	KAIST, South Korea

Topic 9: Multicore and Manycore Methods and Tools

Global Chair

Christoph Kessler	Linköping University, Sweden

Local Chair

Marco Danelutto	University of Pisa, Italy

Chairs

Rudolf Eigenmann	University of Delaware, USA
Arturo González Escribano	Universidad de Valladolid, Spain
Kevin Hammond	University of St. Andrews, UK
Jesper Larsson Träff	TU Wien, Austria
Ana Lucia Varbanescu	University of Amsterdam, The Netherlands

Topic 10: Theory and Algorithms for Parallel Computation and Networking

Global Chair

Christos Zaroliagis	University of Patras, Greece

Local Chair

Tiziano De Matteis	University of Pisa, Italy

Chairs

Leszek Gąsieniec	University of Liverpool, UK
Ulrich Meyer	Goethe-Universität Frankfurt am Main, Germany
Henning Meyerhenke	Universität zu Köln, Germany

Topic 11: Parallel Numerical Methods and Applications

Global Chair

Elisabeth Larsson	Uppsala University, Sweden

Local Chair

Pasqua D'Ambra	CNR-IAC, Naples, Italy

Chairs

Aneta Karaivanova	Institute of Information and Communication Technology (IICT-BAS), Bulgaria
Ángeles Martínez Calomardo	University of Padova, Italy
Ulrike Meier Yang	CASC-LLNL, USA

Topic 12: Accelerator Computing for Advanced Applications

Global Chair

Angeles Navarro	Universidad de Málaga, Spain

Local Chair

Maurizio Drocco	University of Turin, Italy

Chairs

Raphael de Camargo	Federal University of ABC, Brazil
Jaejin Lee	Seoul National University, South Korea
Jose Luis Nunez-Yanez	University of Bristol, UK

Artifact Evaluation Committee

Javier Garcia Blas	University Carlos III of Madrid, Spain
Massimo Coppola	CNR-ISTI Pisa, Italy
Tiziano De Matteis	University of Pisa, Italy
Daniele De Sensi	University of Pisa, Italy
Manuel F. Dolz	University Carlos III of Madrid, Spain
Maurizio Drocco	University of Turin, Italy
Javier Fernandez Munoz	University Carlos III of Madrid, Spain
Salvatore Filippone	Cranfield University, UK
Dalvan Griebler	Pontifical Catholic University of Rio Grande do Sul, Brazil
Lu Li	University of Edinburgh, UK
Fabrizio Marozzo	DIMES, University of Calabria, Italy
Matteo Nardelli	University of Rome Tor Vergata, Italy
Mauricio Pilla	Federal University of Pelotas, Brazil
Paolo Viviani	University of Turin, Italy

Euro-Par 2018 Additional Reviewers

Aguilar, Xavier
Ahmed, Laeeq
Ahvar, Ehsan
Alguwaifli, Yasir
Amarís González, Marcos
Ancourt, Corinne
Angriman, Eugenio
Aral, Atakan
Areias, Miguel
Atalar, Aras
Atanassov, Emanouil
Aublin, Pierre-Louis
Aupy, Guillaume
Aurangzeb
Bacciu, Davide
Bachstein, Matthew
Bajic, Fani
Bajrovic, Enes
Balouek-Thomert, Daniel
Barker, Andrew
Baroli, Davide
Barwell, Adam
Basumallik, Ayon
Beaumont, Olivier
Behrens, Hans
Belcastro, Loris
Beltran, Vicenç
Benedict, Shajulin
Bergamaschi, Luca
Berlinska, Joanna
Bernaschi, Massimo
Bigot, Julien
Bolten, Matthias
Bourgoin, Mathias
Brandes, Thomas
Brown, Christopher
Brown, Dominic
Buesing, Henrik
Buttari, Alfredo
Cardellini, Valeria
Carlini, Emanuele
Caron, Eddy
Carpentieri, Bruno
Carretero, Jesus
Carribault, Patrick
Catena, Matteo
Cavalheiro, Gerson
Caíno-Lores, Silvina
Cecilia, José M.
Cesario, Eugenio
Chen, Yitao
Chester, Dean
Chester, Dean G.
Collange, Sylvain
Comito, Carmela
Corbera, Francisco
Coullon, Helene
Danelutto, Marco
Dao, Than Tuan
De Maio, Vincenzo
De Simone, Valentina
de Souza, Diego F.
Decouchant, Jérémie
Del Río Astorga, David
Denoyelle, Nicolas
Di Girolamo, Salvatore
Dickson, James
Do, Youngdong
Dobrev, Veselin
Dokulil, Jiri
Dolz, Manuel
Dreuning, Henk
Drocco, Maurizio
Du, He
Durastante, Fabio
Dutot, Pierre-François
Dutta, Hritam
Engblom, Stefan
Falcao, Gabriel
Faqeh, Rasha
Fernandez Fabeiro, Jorge
Fernandez, Javier
Fey, Dietmar
Fey, Florian
Filippone, Salvatore
Flores, Paulo
Folino, Francesco
Folino, Gianluigi
Forestiero, Agostino
Gadkari, Ashish
Gante, João
Garcia Blas, Javier
Gardner, David
Garg, Yash
Gaspar, Francisco
Gendron, Bernard
Geronimo, Guilherme
Gorlatch, Sergei
Gracia, José
Grahn, Håkan
Grelck, Clemens
Grubel, Patricia
Gu, Chuancai
Guermouche, Abdou
Guerreiro, João
Gurov, Todor
Guyon, David
Ha, Phuong
Herlihy, Maurice
Hiraishi, Tasuku
Hoefer, Martin
Hugo, Andra
Hückelheim, Jan
Ienco, Dino
Ivanovska, Sofiya
Janna, Carlo
Jeong, Chang-Sung
Jiang, Xu
Jo, Gangwon
Jung, Wookeun
Kallimanis, Nikolaos
Kang, Dong-Ki
Karlsson, Sven
Keller, Jörg
Kim, Heehoon
Kim, Hyngmo
Kim, Jungwook
Kim, Seong-Hwan
Kim, Woojoong

Kirk, Richard
Kontogiannis, Spyros
Koslovski, Guilherme
Kosta, Sokol
Krall, Andreas
Kritikos, Kyriakos
Kronbichler, Martin
Kukreja, Navjot
Kumar, Rakesh
Kumaraswamy, Madhura
Kunji, Khalid
Lambert, Thomas
Larsson, Elisabeth
Le Quoc, Do
Lee, Jinpil
Lee, Yongjun
Legrand, Arnaud
Li, Ruipeng
Li, Xueqi
Lin, Yuhan
Lirkov, Ivan
Liu, Jun
Liu, Sicong
Liu, Songran
Liu, Weifeng
Low, Tze Meng
Lujic, Ivan
Lulli, Alessandro
Madi-Wamba, Gilles
Maier, Tobias
Mao, Bo
Maouche, Mohamed
Marozzo, Fabrizio
Marques, Diogo
Martin, André
Martins, Paulo
Matsuda, Motohiko
Mazzia, Francesca
McKenney, Paul
Mehofer, Eduard
Mele, Valeria
Melot, Nicolas
Mendes, Rafael
Meng, Ke
Mijaković, Robert
Misale, Claudia
Morais, Mayuri
Mordacchini, Matteo
Mounié, Grégory
Nagarkar, Parth
Nakao, Masahiro
Netzer, Gilbert
Neves, Nuno
Nidito, Francesco
Oh, Pyeongseok
Ohshima, Satoshi
Oleksenko, Oleksii
Osborn, Sarah
Owen, Herbert
Owenson, Andrew
Padoin, Edson Luiz
Palkowski, Marek
Park, Sangdon
Pascual, Fanny
Penschuck, Manuel
Perarnau, Swann
Pereira, Fernando
Petri, Matthias
Pham, Linh Manh
Pllana, Sabri
Poccia, Silvestro
Poke, Marius
Ponce, Colin
Popov, Mihail
Popovici, Doru Thom
Prokopec, Aleksandar
Qiao, Bo
Qiu, Sheng
Raca, Valon
Rais, Issam
Raoofy, Amir
Rauchwerger, Lawrence
Rebonatto, M. Trindade
Reguly, Istvan
Rehn-Sonigo, Veronika
Reiche, Oliver
Reis, Valentin
Rezaei, Mohamad
Ricketson, Lee
Rivas-Gomez, Sergio
Romano, Diego
Ros, Alberto
Rupp, Karl
Sabne, Amit
Sadi, Fazle
Sakdhnagool, Putt
Sampathirao, Ajay Kumar
Sandrieser, Martin
Santana, Luis
Santander-Jiménez, Sergio
Santos, Danilo
Sarkar, Subhadeep
Sarkar, Susmit
Sato, Yukinori
Saveski, Martin
Schickedanz, Alexander
Schmitt, Christian
Schroder, Jacob
Schuchart, Joseph
Serafini, Marco
Serrano, Estefania
Shin, Dong-Jae
Shun, Julian
Sjogreen, Bjorn
Skowron, Piotr
Soliman, Amira
Spampinato, Daniele
Spataro, William
Spiga, Filippo
Stramondo, Giulio
Sun, Jinghao
Sun, Tianjiao
Sutra, Pierre
Tairum, Miguel
Tanase, Alexandru
Tang, Yue
Teabe, Boris
Terboven, Christian
Thamsen, Lauritz
Tomas, Andres
Tomov, Vladimir
Torres de La Sierra, Yuri
Trach, Bohdan
Trinitis, Carsten
Trunfio, Paolo
Träff, Jesper Larsson
Tsai, Yaohung

Tsigas, Philippas
Tsuji, Miwako
Ujaldon, Manuel
Utrera, Gladys
Vaddina, Kameswar
Vasiloudis, Theodore
Vaumourin, Gregory
Veiga, Luís
Veith, David
Venetis, Ioannis
Verbitskiy, Ilya
Verdoolaege, Sven
Vespa, Emanuele
Villegas, Alejandro
Vinci, Andrea
Viviani, Paolo
von Looz, Moritz
Voronin, Kirill
Wacrenier, Pierre-André
Wagner, Martin
Wallschläger, Marcel
Walulya, Ivan
Wang, Zhan
Weidendorfer, Josef
Wilhelm, Andreas
Wong, Michael
Wu, Meng-Ju
Xie, Zhen
Xu, Guanglin
Yang, Dai
Yang, Eunju
Yang, Qirui
Yao, Erlin
Yarkhan, Asim
Yasugi, Masahiro
Yeh, Tsung Tai
Zamani, Ali Reza
Zang, Dawei
Zhang, Fa
Zhang, Jiyuan
Zilic, Josip
Zlatev, Zahari
Zounon, Mawussi
Żuk, Paweł

Euro-Par 2018 Invited Talks

ALGORAND: A Better Distributed Ledger

Silvio Micali

CSAIL, MIT, USA

A distributed ledger is a tamperproof sequence of data that can be read and augmented by everyone. Distributed ledgers stand to revolutionize the way a democratic society operates. They secure all kinds of traditional transactions – such as payments, asset transfers, titling – in the exact order in which they occur; and enable totally new transactions – such as cryptocurrencies and smart contracts. They can remove intermediaries and usher in a new paradigm for trust. As currently implemented, however, distributed ledgers cannot achieve their enormous potential. Algorand is an alternative, democratic, and efficient distributed ledger. Unlike prior ledgers based on "proof of work", it dispenses with "miners". Indeed, Algorand requires only a negligible amount of computation. Moreover, its transaction history does not "fork" with overwhelming probability: i.e., Algorand guarantees the finality of all transactions.

Algorithmic Adaptations to Extreme Scale Computing

David E. Keyes

King Abdullah University, Saudi Arabia

Algorithmic adaptations to use next-generation computers close to their potential are underway. Instead of squeezing out flops – the traditional goal of algorithmic optimality, which once served as a reasonable proxy for all associated costs – algorithms must now squeeze synchronizations, memory, and data transfers, while extra flops on locally cached data represent only small costs in time and energy. After decades of programming model stability with bulk synchronous processing, new programming models and new algorithmic capabilities (to make forays into, e.g., data assimilation, inverse problems, and uncertainty quantification) must be co-designed with the hardware. We briefly recap the architectural constraints and application opportunities. We then concentrate on two types of tasks each of occupies a large portion of all scientific computing cycles: large dense symmetric/Hermitian linear systems (covariances, Hamiltonians, Hessians, Schur complements) and large sparse Poisson/Helmholtz systems (solids, fluids, electromagnetism, radiation diffusion, gravitation). We examine progress in porting "exact" and hierarchically rank-reduced solvers for these tasks to the hybrid distributed-shared programming environment, including the GPU and the MIC architectures that make up the cores of the top scientific computers "on the floor" and "on the books."

Datacenters for the Post-Moore Era

Babak Falsafi

EPFL, Switzerland

Datacenters are growing at unprecedented speeds fueled by the demand on global IT services, investments in massive data analytics and economies of scale. Worldwide data by some accounts (e.g., IDC) grows at much higher rates than server capability and capacity. Conventional silicon technologies laying the foundation for server platforms, however, have dramatically slowed down in efficiency and density scaling in recent years. The latter, now referred to as the post-Moore era, has given rise to a plethora of emerging logic and memory technologies presenting exciting new challenges and abundant opportunities from algorithms to platforms for server designers. In this talk, I will first motivate the post-Moore era for server architecture and present avenues to pave the path forward for server design.

Euro-Par 2018 Topics Overview

Topic 1: Support Tools and Environments

Siegfried Benkner, Massimo Coppola, Franz Franchetti,
Michael Gerndt, Erwin Laure, and Nikos Parlavantzas

Despite an impressive body of research over the last decades, parallel and distributed programming remains a complex task, a process that is prone to subtle software issues that can affect both the correctness and the performance of an application. The amount of implementation details and their hidden connections are getting harder and harder to manage as multilevel parallel hierarchical and hybrid architectures become more and more commonplace along the path to Exascale computing systems.

The Euro-Par Support Tools and Environments track focuses on tools, techniques and environments that help tackling that complexity by addressing the many challenges related to programmability, portability, correctness, reliability, scalability, efficiency, performance and energy consumption.

The papers submitted and accepted for this track do well represent the community that this topic brings together, gathering tool designers, developers, and users to share their ideas, solutions, products, and concerns for a wide range of parallel platforms. Key points of the evaluation were solid theoretical foundations and strong experimental validations on production-level parallel and distributed systems, as well as the novelty of program development tools and environments that tackle the daunting complexity of current and future parallel systems.

The track received 14 submissions, which were thoroughly reviewed by the members of the track program committee with the help of 27 external reviewers delivering in total 55 distinct reviews. Out of all the submissions and after a careful and detailed discussion among committee members, we finally decided to accept 5 papers, resulting in a per-topic acceptance ratio of 36%.

We would like to thank all the authors who submitted papers for their contribution to the success of this track, as well as all the external reviewers for their high-quality reviews and their valuable feedback.

Topic 2: Performance and Power Modeling, Prediction and Evaluation

Leonel Sousa, Daniele De Sensi, Giorgis Georgakoudis, Aleksandar Ilic, Piotr Luszczek, Federico Silla, Guangming Tan, and Pedro Trancoso

Power consumption is becoming a major factor to consider when designing hardware and applications. Due to the tight correlation to performance, these two goals need to be addressed together in a synergistic way. This topic covers different aspects of performance and power consumption modeling, prediction and evaluation on different types of computing architectures and for a wide variety of applications.

This year we received 24 submissions and each paper received 4 reviews, either from the nine program committee members and/or from external reviewers. After discussion, we accepted 6 papers (25% acceptance rate). The papers cover different aspects of performance optimization, power and energy efficiency, addressing the problem at different levels, from low-level optimizations to visualization tools and considering different platforms, from mobile devices to large-scale HPC systems.

We would like to thank the authors for their submissions, the Euro-Par 2018 Organizing Committee for their help throughout all the process, and the PC members and the reviewers for providing timely and detailed reviews, and for participating in the discussion we carried on after the reviews were received.

Topic 3: Scheduling and Load Balancing

Anne Benoit, Enrico Bini, Maciej Drozdowski,
Lionel Eyraud-Dubois, José Gracia, Nan Guan, Sascha Hunold,
and Krzysztof Rzadca

New computing systems offer the opportunity to reduce the response times and the energy consumption of the applications by exploiting the levels of parallelism. Heterogeneity and complexity are the distinguishing characteristics of modern architectures. Thereby, the optimal exploitation of modern platforms is challenging. Scheduling and load balancing techniques are key instruments to achieve higher performance, lower energy consumption, reduced resource usage, and predictability of applications.

This topic invites papers on all aspects related to scheduling and load balancing on parallel and distributed machines, from theoretical foundations for modeling and designing efficient and robust scheduling policies to experimental studies, applications and practical tools and solutions. It applies to multi-/manycore processors, embedded systems, servers, heterogeneous and accelerated systems, HPC clusters as well as distributed systems such as clouds and global computing platforms.

A total of 23 full-length submissions were received in this track, each of which received at least four reviews, from the eight chairs and/or from 21 additional experts. Following the thorough discussion of the reviews, five submissions have been accepted, including one that was nominated as distinguished paper.

The chair and local chair sincerely thank all the authors for their submissions, the Euro-Par 2018 Organizing Committee for all their valuable help, and the reviewers for their excellent work. They all have 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

Florian Brandner, Fabio Luporini, Alexandra Jimborean,
Frank Hannig, and Gihan Mudalige

This topic deals with architecture design, programming 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 programming 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, multi-threading, architectural support for parallelism, and the impact of emerging hardware technologies.

This year the topic received 11 submissions, covering a wide range of topics ranging from hardware designs over compilation techniques to programming models. The five topic co-chairs solicited at least four experts in the respective fields to review each paper. A lively online discussion followed the reviewing phase, during which the various co-chairs frequently solicited additional input from the expert reviewers. Based on the online discussions, 3 papers were proposed for acceptance, which were ultimately confirmed during the final selection meeting held in Turin.

Topic 5: Parallel and Distributed Data Management and Analytics

K. Selçuk Candan, Ruggero Pensa, Lei Chen,
Gianmarco De Francisci Morales, and Ming Zhao

Many areas of science, industry, and commerce are producing extreme-scale data that must be processed – stored, managed, analyzed – 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, IoT data management and analytics, and scalable data processing workflows are all in the scope of this topic. Privacy and trust issues in parallel and distributed data management and analytics systems are also aspects of interest for this conference topic.

Seven full-length papers were submitted to this topic, and each paper received at least four reviews, mostly performed by track chairs. After discussion with the reviewers and track chairs, two papers were selected for publication, one related to the minimization of network traffic for distributed joins, the second one to privacy-preserving top-k query processing in distributed systems.

Topic 6: Cluster and Cloud Computing

Ivona Brandić, Domenico Talia, Toni Mastelic, Raffaele Montella,
Anne-Cécile Orgerie, Thomas Renner, and Rafael Brundo Uriarte

Cloud Computing evolved from Cluster Computing and Grid Computing as a new parallel and scalable architecture. Cloud Computing is a paradigm and a technology that today is largely used. Together with Grid and Cluster computing, Cloud Computing is 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, in all these cases Clusters, Grids and Clouds are useful tools.

Beyond the scene, most of Cloud Computing solutions rely on federations of large-scale clusters where well-known but still unsolved challenges related to performance, reliability and energy efficiency of the infrastructures should be addressed by research. Moreover, Cloud Computing emphasized the importance of fundamental capabilities and services that are required to achieve the goal of user-friendly, security and service guarantees. Our community should also investigate these aspects.

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 High Performance Computing applications on Clouds. The first referred as "fog/edge" computing, such a dawning paradigm is attracting growing interests as it brings computing resources closer to end-users, tackling the network overhead issues that prevent the use of the UC paradigm by latency-aware applications. The second still needs a large research effort, to allow the use of compute and network intensive applications without loss of performance 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, 24 papers have been submitted to Topic 6. There were authors from several countries from all the continents. Four expert reviewers analyzed each submission. Overall, many specialists were involved into the reviewing process and, despite the high quality of the submitted papers, only 8 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 2018 Organizing committee for all the help that allows us to smoothly take over the whole process.

Topic 7: Distributed Systems and Algorithms

Sonia Ben-Mokhtar, Alberto Montresor, Christof Fetzer,
and Indranil Gupta

Parallel computing is heavily dependent on and interacts with the developments and challenges concerning distributed systems, such as load balancing, asynchrony, failures, malicious and selfish behavior, high latencies, network partitions, disconnected operations and heterogeneity. 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 computing models, distributed data structures, and parallel processing on distributed systems, in particular in relation to efficient high performance computing.

This year the track received 10 submissions on various topics of the call and accepted 3 papers. Each paper had a minimum of four reviews and was discussed within the track PC meeting. A subset of papers was then proposed to the PC chairs for final discussions and decisions.

The track chairs would like to warmly thank the track members Indranil Gupta (University of Illinois Urbana-Champaign, USA) and Christof Fetzer (TU Dresden, Germany) for their work as well as the 12 external reviewers that greatly helped in the reviewing process.

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

J. Daniel García, Patrizio Dazzi, Bryce Adelstein-Lelbach, Marcelo Pasin, Mitsuhisa Sato, Paolo Trunfio, and Chan-Hyun Youn

Parallel and distributed applications requires adequate programming abstractions and models, efficient design tools, parallelization techniques and practices. This topic was open for submissions 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.

The topic emphasizes 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 for usability, performance prediction, scalability, self-adaptation, rapid prototyping and fault-tolerance, as needed, for instance, in dynamic heterogeneous parallel and distributed infrastructures, were welcome.

We received nineteen submissions on this topic that went through four independent reviews. Those reviews where further discussed among the PC members. As a result eight paper were accepted.

We would like to express our gratitude to all author for submitting their work. We received very good submissions and the selection of accepted paper was quite hard. We also would like to thank all the reviewers for their detailed reviews and their participation in discussions following the reviews. Finally we would like to also thank the organizing and steering committees for all their help, support and hard work.

Topic 9: Multicore and Manycore Methods and Tools

Christoph Kessler, Marco Danelutto, Rudolf Eigenmann,
Arturo González Escribano, Kevin Hammond, Jesper L. Träff,
and Ana L. Varbanescu

Modern homogeneous and heterogeneous multi-core and many-core architectures are now part of the high-end, embedded, and mainstream computing scene and can offer impressive performance for many applications. This architecture trend has been driven by the need to reduce power consumption, increase processor utilization, and deal with the memory-processor speed gap. However, the complexity of these new architectures has created several programming challenges, and achieving performance on these systems is often a difficult task. This topic seeks to explore productive programming of multi- and many-core systems, as well as stand-alone systems with large numbers of cores like GPUs and various types of accelerators. This can also include hybrid and heterogeneous systems with different types of multi-core processors. It focuses on novel research and solutions in the form of programming models, algorithms, languages, compilers, libraries, runtime and analysis tools to increase the programmability of multi-core, many-core, and heterogeneous systems, in the context of general-purpose, high-performance, and embedded parallel computing. It also covers issues such as lock-free algorithms and data structures, transactional memory, static and dynamic analysis and optimization techniques and tools, performance and power trade-offs, scalability aspects, and hardware support for programming models and runtime systems.

This year, 25 papers discussing some of these issues were submitted to this topic. Each paper was reviewed by four reviewers. Eventually, 6 regular papers were selected.

The accepted papers discuss the following issues: load balancing for parallel graph traversal algorithms on GPUs, energy-efficient stencil computations on clustered many-core processors, optimizing the thread placement for overlapping MPI-3 non-blocking collective communication operations on many-core processors, a lock-free cache-trie data structure, improving performance of multi-program workloads by cache-criticality aware last-level cache partitioning, and NUMA optimizations for algorithmic skeletons.

The topic chairs wish to thank all authors contributing their work to the topic, the PC members and the additional reviewers for their highly useful comments, as well as the Euro-Par Organizing Committee for creating a smooth process.

Topic 10: Theory and Algorithms for Parallel Computation and Networking

Christos Zaroliagis, Tiziano De Matteis, Leszek Gąsieniec,
Ulrich Meyer, and Henning Meyerhenke

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 9 submissions, all of which received 4 reviews. The papers and their reviews were discussed extensively, and 3 submissions were eventually accepted. We thank all authors for their valuable contributions, as well as the PC Committee members and external reviewers for investing their time in reviewing the papers, for providing constructive feedback and sharing their expertise, and for keeping the high scientific level of the Euro-Par conference.

Topic 11: Parallel Numerical Methods and Applications

Elisabeth Larsson, Pasqua D'Ambra, Aneta Karaivanova, Ángeles Martínez Calomardo, and Ulrike Meier Yang

The need for high performance computing is driven by the need for predictive simulations in science and engineering, as well as in areas such as finance, life sciences, and humanities, where computational needs have more recently been increasing. This requires the development of highly scalable numerical methods and algorithms that are able to efficiently exploit modern, and in general heterogeneous, computer architectures. Another need that is currently arising with the increasing size of computer systems is fault tolerance, which puts additional demands on algorithms, run-time systems, and tools such as MPI.

This conference topic aims at providing a forum for presenting and discussing recent developments in parallel numerical algorithms and their implementation on current parallel architectures, including many-core and hybrid architectures. We encouraged submissions addressing algorithmic design, implementation details, performance analysis, as well as integration of parallel numerical methods in large-scale science and engineering applications.

The program committee for this topic consisted of five women with different specializations in high-performance parallel computing for numerical applications. We received 17 submissions on a broad variety of topics. Fourty-one additional experts were involved in the review process. Each submission received at least four reviews. After the paper selection meeting at the University of Turin, four high quality papers were accepted for presentation at EuroPar 2018. The topics of the papers cover application areas in plasma physics, quantum physics, seismic wave propagation, and matrix factorizations. Algorithmic aspects of the Particle-in-cell method and Cholesky factorization for dense matrices with compressed blocks are considered. Implementations are performed using task based parallel programming models as well as explicit programming models using MPI+threads.

We thank all the authors for their contributions, the reviewers for their careful reading of the papers, and the organizing committee members for their smooth operation of the whole process.

Topic 12: Accelerator Computing for Advanced Applications

Angeles Navarro, Maurizio Drocco, Raphael de Camargo, Jaejin Lee, and Jose Luis Nunez-Yanez

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.

Nowadays, 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 coprocessors, as well as more customizable devices, such as FPGA-based systems, and 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 on architectures, algorithms, languages, compilers, libraries, runtime systems, coordination of accelerators and CPU, debugging and profiling tools, as well as 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 five members of different backgrounds and specializations in the accelerators field, with the collaboration of several other sub-reviewers. We received 11 contributions from researchers in many different countries. After the review process and the general PC meeting, three high-quality papers were selected for presentation in Euro-Par 2018 at Turin. They are focused on important hot-topics: exploiting the GPUs potential towards advanced hierarchical matrix computations in large-scale sparse applications, proposing runtime systems for dynamically adapting the state on heterogeneous systems to enhance its energy efficiency, or introducing stream processing frameworks that enable easily programmable and high-performance computations on hybrid CPU/Xeon Phi systems.

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 easy our task, and to provide a nice conference environment in Turin for a high-quality discussion of research results in this emerging topic.

Contents

Cluster and Cloud Computing

Distributed Systems and Algorithms