Leonel Sousa
Nuno Roma
Pedro Tomás (Eds.)

LNCS 12820 ARCoSS

Euro-Par 2021: Parallel Processing

27th International Conference
on Parallel and Distributed Computing
Lisbon, Portugal, September 1–3, 2021
Proceedings

Springer

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

Leonel Sousa · Nuno Roma ·
Pedro Tomás (Eds.)

Euro-Par 2021: Parallel Processing

27th International Conference on Parallel and Distributed Computing Lisbon, Portugal, September 1–3, 2021 Proceedings

Springer

Editors
Leonel Sousa
Universidade de Lisboa
Lisbon, Portugal

Nuno Roma
Universidade de Lisboa
Lisbon, Portugal

Pedro Tomás
Universidade de Lisboa
Lisbon, Portugal

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-85664-9 ISBN 978-3-030-85665-6 (eBook)
https://doi.org/10.1007/978-3-030-85665-6

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 2021, the 27th International European Conference on Parallel and Distributed Computing, held between September 1–3, 2021. Although Euro-Par 2021 had been planned to take place in Lisbon, Portugal, it was organized as a virtual conference, as a result of the COVID-19 pandemic.

For over 25 years, Euro-Par has brought together researchers in parallel and distributed computing. Founded by pioneers as a merger of the two thematically related European conference series PARLE and CONPAR-VAPP, Euro-Par started with the aim to create the main annual scientific event on parallel and distributed computing in Europe and to be the primary choice of professionals for the presentation of their latest results.

Since its inception, Euro-Par has covered all aspects of parallel and distributed computing, ranging from theory to practice; scaling from the smallest to the largest parallel and distributed systems; from fundamental computational problems and models to full-fledged applications; and from architecture and interface design and implementation to tools, infrastructures, and applications. Euro-Par's unique organization into topics provides an excellent forum for focused technical discussion as well as interaction with a large, broad, and diverse audience of researchers in academic institutions, public and private laboratories, and commercial stakeholders. Euro-Par's topics have always been oriented towards novel research issues and the current state of the art. Most topics have become constant entries, while new themes have emerged and been included in the conference. Euro-Par selects new organizers and chairs for every event, giving the opportunity to young researchers and leading to fresh ideas while ensuring tradition. Organizers and chairs of previous events support their successors. In this sense, the conference also promotes networking across borders, leading to the unique spirit of Euro-Par.

Previous 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, Santiago de Compostela, Turin, Göttingen, and Warsaw.

Thus, Euro-Par in Portugal followed the well-established format of its predecessors. The 27th Euro-Par conference was organized with the support of INESC-ID and Instituto Superior Técnico (Tecnico Lisboa) -the Faculty of Engineering of the University of Lisbon. Tecnico Lisboa is a renowned place for research in engineering, including in all fields related to electrical and computer engineering and computer science.

The topics of Euro-Par 2021 were organized into 10 tracks, namely:

- Compilers, Tools, and Environments
- Performance and Power Modeling, Prediction and Evaluation

- Scheduling and Load Balancing
- Data Management, Analytics, and Machine Learning
- Cluster, Cloud, and Edge Computing
- Theory and Algorithms for Parallel and Distributed Processing
- Parallel and Distributed Programming, Interfaces, and Languages
- Multicore and Manycore Parallelism
- Parallel Numerical Methods and Applications
- High-performance Architectures and Accelerators

Overall, 136 papers were submitted from 37 countries. The number of submitted papers, the wide topic coverage, and the aim of obtaining high-quality reviews resulted in a careful selection process involving a large number of experts. As the joint effort of the members of the Program Committee and of the 121 external reviewers, a total of 218 reviewers from 33 countries submitted 540 reviews: 10 papers received three reviews, 122 received four reviews, and 4 received 5 or more. On average, each paper received 4 reviews. The accepted papers were chosen after offline discussions in the reviewing system followed by a lively discussion during the paper selection meeting, which took place via video conference on April 19, 2021. As a result, 38 papers were selected to be presented at the conference and published in these proceedings, resulting in a 27.9% acceptance rate.

To increase reproducibility of the research, Euro-Par encourages authors to submit artifacts, such as source code, data sets, and reproducibility instructions. Along with the notification of acceptance, authors of accepted papers were encouraged to submit their artifacts. A total of 6 papers were complemented with their artifacts. These artifacts were then evaluated by the Artifact Evaluation Committee (AEC). The AEC managed to successfully reproduce the results of all of the 6 papers. These papers are marked in the proceedings by a special stamp; and the artifacts are available on-line in the Figshare repository.

In addition to the technical program, we had the pleasure of hosting three keynotes held by:

- Manish Parashar, University of Utah, USA,
- Alba Cristina Melo, University of Brasilia, Brazil,
- Keshav Pingali, University of Texas, USA.

Although Euro-Par 2021 was converted into a virtual event, it encouraged interaction and on-line discussions, in order to make it a successful and friendly meeting.

The conference program started with two days of workshops on specialized topics. Dora Blanco Heras and Ricardo Chaves ensured coordination and organization of this pre-conference event as workshop co-chairs. After the conference, a selection of the papers presented at the workshops will be published in a separate proceedings volume.

We would like to thank the authors, chairs, PC members, and reviewers for contributing to the success of Euro-Par 2021. Similarly, we would like to extend our appreciation to the Euro-Par Steering Committee for its support. Our mentor, Fernando Silva, devoted countless hours to this year's event, making sure we were on time and

on track with all the (many) key elements of the conference. Our virtual task force led by Tiago Dias took care of various aspects of moving a physical conference to cyberspace.

August 2021

Leonel Sousa
Nuno Roma
Pedro Tomás

Organization

The Euro-Par Steering Committee

Full Members

Luc Bougé (SC Chair)	ENS Rennes, France
Fernando Silva (SC Vice-Chair)	University of Porto, Portugal
Dora Blanco Heras (Workshops Chair)	CiTIUS, University of Santiago de Compostela, Spain
Marco Aldinucci	University of Turin, Italy
Emmanuel Jeannot	Inria, Bordeaux, France
Christos Kaklamanis	Computer Technology Institute, Patras, Greece
Paul Kelly	Imperial College London, UK
Thomas Ludwig	University of Hamburg, Germany
Tomàs Margalef	Autonomous University of Barcelona, Spain
Wolfgang Nagel	Dresden University of Technology, Germany
Francisco Fernàndez Rivera	CiTIUS, University of Santiago de Compostela, Spain
Krzysztof Rzadca	University of Warsaw, Poland
Rizos Sakellariou	University of Manchester, UK
Henk Sips (Finance Chair)	Delft University of Technology, The Netherlands
Leonel Sousa	Universidade de Lisboa, Portugal
Domenico Talia	University of Calabria, Italy
Massimo Torquati (Artifacts Chair)	University of Pisa, Italy
Phil Trinder	University of Glasgow, UK
Denis Trystram	Grenoble Institute of Technology, France
Felix Wolf	Technical University of Darmstadt, Germany
Ramin Yahyapour	GWDG, Göttingen, Germany

Honorary Members

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

General Chair

Leonel Sousa	INESC-ID, IST, Universidade de Lisboa, Portugal

Workshop Chairs

Ricardo Chaves	INESC-ID, IST, Universidade de Lisboa, Portugal
Dora Blanco Heras	University of Santiago de Compostela, Spain

PhD Symposium Chairs

Aleksandar Ilic	INESC-ID, IST, Universidade de Lisboa, Portugal
Didem Unat	Koç University, Turkey

Submissions Chair

Nuno Roma	INESC-ID, IST, Universidade de Lisboa, Portugal

Publicity Chairs

Gabriel Falcão	IT, Universidade de Coimbra, Portugal
Maurício Breternitz	ISCTE, Instituto Universitário de Lisboa, Portugal

Web Chairs

Pedro Tomás	INESC-ID, IST, Universidade de Lisboa, Portugal
Helena Aidos	LASIGE, FCUL, Universidade de Lisboa, Portugal

Local Chairs

Tiago Dias	INESC-ID, ISEL, Instituto Politécnico de Lisboa, Portugal
Ricardo Nobre	INESC-ID, Portugal

Artifact Evaluation Committee

Nuno Neves	INESC-ID, Universidade de Lisboa, Portugal
Massimo Torquati	University of Pisa, Italy

Scientific Organization

Topic 1: Compilers, Tools, and Environments

Global Chair

Frank Hannig	Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Local Chair

Gabriel Falcão	Instituto de Telecomunicações, Universidade de Coimbra, Portugal

Members

Izzat El Hajj	American University of Beirut, Lebanon
Simon Garcia Gonzalo	Barcelona Supercomputing Center, Spain
Hugh Leather	Facebook AI Research and The University of Edinburgh, UK
Josef Weidendorfer	TU München, Germany

Topic 2: Performance and Power Modeling, Prediction and Evaluation

Global Chair

Didem Unat	Koç University, Turkey

Local Chair

Aleksandar Ilic	INESC-ID, IST, Universidade de Lisboa, Portugal

Members

Laura Carrington	EP Analytics, USA
Milind Chabbi	Uber Technologies, USA
Bilel Hadri	KAUST Supercomputing Lab, Saudi Arabia
Jawad Haj-Yahya	ETH Zurich, Switzerland
Andreas Knüpfer	TU Dresden, Germany
Alexey Lastovetsky	University College Dublin, Ireland
Matthias Mueller	RWTH Aachen University, Germany
Emmanuelle Saillard	CRCN, Inria Bordeaux, France
Martin Schulz	Technical University Munich, Germany

Topic 3: Scheduling and Load Balancing

Global Chair

Oliver Sinnen	University of Auckland, New Zealand

Local Chair

Jorge Barbosa	Universidade do Porto, Portugal

Members

Giovanni Agosta	Politecnico di Milano, Italy
Bertrand Simon	University of Bremen, Germany
Luiz F. Bittencourt	University of Campinas, Brazil
Lúcia Drummond	Universidade Federal Fluminense, Brazil
Emmanuel Casseau	University of Rennes, Inria, CNRS, France
Fanny Dufossé	Inria Grenoble, France
Henri Casanova	University of Hawaii at Manoa, USA

Loris Marchal	CNRS, University of Lyon, LIP, France
Maciej Malawski	AGH University of Science and Technology, Poland
Malin Rau	Universität Hamburg, Germany
Pierre-Francois Dutot	Université Grenoble Alpes, France
Rizos Sakellariou	The University of Manchester, UK
Lucas Mello Schnorr	Federal University of Rio Grande do Sul, Brazil
Gang Zeng	Nagoya University, Japan

Topic 4: Data Management, Analytics, and Machine Learning

Global Chair

Alex Delis	University of Athens, Greece

Local Chair

Helena Aidos	LASIGE, Universidade de Lisboa, Portugal

Members

Angela Bonifati	Lyon 1 University, France
Pedro M. Ferreira	LASIGE, Universidade de Lisboa, Portugal
Panagiotis Liakos	University of Athens, Greece
Jorge Lobo	Universitat Pompeu Fabra, Spain
Paolo Romano	INESC-ID, Universidade de Lisboa, Portugal
Jacek Sroka	University of Warsaw, Poland
Zahir Tari	Royal Melbourne Institute University, Australia
Goce Trajcevski	Iowa State University, USA
Vassilios Verykios	Hellenic Open University, Greece
Demetris Zeinalipour	University of Cyprus, Cyprus

Topic 5: Cluster, Cloud, and Edge Computing

Global Chair

Radu Prodan	University of Klagenfurt, Austria

Local Chair

Luís Veiga	INESC-ID, Universidade de Lisboa, Portugal

Members

Juan J. Durillo	Leibniz Supercomputing Centre (LRZ), Germany
Felix Freitag	Polytechnic University of Catalonia, Spain
Aleksandar Karadimche	University for Information Science and Technology "St. Paul the Apostle", Macedonia
Attila Kertesz	University of Szeged, Hungary
Vincenzo de Maio	Technical University of Vienna, Austria
Joan-Manuel Marquès	Universitat Oberta de Catalunya, Spain

Rolando Martins	University of Porto, Portugal
Hervé Paulino	Universidade Nova de Lisboa, Portugal
Sasko Ristov	University of Innsbruck, Austria
José Simão	INESC-ID, ISEL, Portugal
Spyros Voulgaris	Vrije Universiteit Amsterdam, The Netherlands

Topic 6: Theory and Algorithms for Parallel and Distributed Processing

Global Chair

Andrea Pietracaprina	Università di Padova, Italy

Local Chair

João Lourenço	Universidade Nova de Lisboa, Portugal

Members

Adrian Francalanza	University of Malta, Malta
Ganesh Gopalakrishnan	University of Utah, USA
Klaus Havelund	NASA's Jet Propulsion Labratory, USA
Ulrich Meyer	Goethe-Universität Frankfurt, Germany
Emanuele Natale	Université Côte d'Azur, Inria-CNRS, France
Yihan Sun	University of California, Riverside, USA

Topic 7: Parallel and Distributed Programming, Interfaces, and Languages

Global Chair

Alfredo Goldman	University of São Paulo, Brazil

Local Chair

Ricardo Chaves	INESC-ID, IST, Universidade de Lisboa, Portugal

Members

Ivona Brandic	Vienna University of Technology, Austria
Dilma Da Silva	Texas A&M University, USA
Laurent Lefevre	Inria, France
Dejan Milojicic	HP Labs, USA
Nuno Preguiça	Universidade Nova de Lisboa, Portugal
Omer Rana	Cardiff University, UK
Lawrence Rauchwerger	University of Illinois at Urbana Champaign, USA
Noemi Rodriguez	PUC-Rio, Brazil
Stéphane Zuckerman	ETIS Laboratory, France

Topic 8: Multicore and Manycore Parallelism

Global Chair

Enrique S. Quintana Ortí	Universitat Politècnica de València, Spain

Local Chair

Nuno Roma	INESC-ID, IST, Universidade de Lisboa, Portugal

Members

Emmanuel Jeannot	Inria, France
Philippe A. Navaux	Federal University of Rio Grande do Sul, Brazil
Robert Schöne	TU Dresden, Germany

Topic 9: Parallel Numerical Methods and Applications

Global Chair

Stanimire Tomov	University of Tennessee, USA

Local Chair

Sergio Jimenez	Universidad de Extremadura, Spain

Members

Hartwig Anzt	Karlsruhe Institute of Technology, Germany
Alfredo Buttari	CNRS/IRIT Toulouse, France
José M. Granado-Criado	University of Extremadura, Spain
Hatem Ltaief	King Abdullah University of Science and Technology, Saudi Arabia
Álvaro Rubio-Largo	University of Extremadura, Spain

Topic 10: High-performance Architectures and Accelerators

Global Chair

Samuel Thibault	Université de Bordeaux, France

Local Chair

Pedro Tomás	INESC-ID, IST, Universidade de Lisboa, Portugal

Members

Michaela Blott	Xilinx, Ireland
Toshihiro Hanawa	University of Tokyo, Japan
Sascha Hunold	Vienna University of Technology, Austria
Naoya Maruyama	NVIDIA, USA
Miquel Moreto	Barcelona Supercomputing Center, Spain

PhD Symposium

Chairs

Didem Unat	Koç University, Turkey
Aleksandar Ilic	INESC-ID, IST, Universidade de Lisboa, Portugal

Members

Mehmet Esat Belviranli	Colorado School of Mines, USA
Xing Cai	Simula Research Laboratory, Norway
Anshu Dubey	Argonne National Lab, USA
Francisco D. Igual	Universidad Complutense de Madrid, Spain
Sergio Santander-Jiménez	University of Extremadura, Spain
Lena Oden	FernUniversität Hagen, Germany

Artifacts Eavaluation

Chairs

Nuno Neves	INESC-ID, IST, Universidade de Lisboa, Portugal
Massimo Torquati	University of Pisa, Italy

Members

Diogo Marques	INESC-ID, IST, Universidade de Lisboa, Portugal
Alberto Martinelli	Università degli Studi di Torino, Italy
João Bispo	INESCTEC, Faculty of Engineering, University of Porto, Portugal
João Vieira	INESC-ID, IST, Universidade de Lisboa, Portugal
Iacopo Colonnelli	Università di Torino, Italy
Ricardo Nobre	INESC-ID, IST, Universidade de Lisboa, Portugal
Nuno Paulino	INESCTEC, Faculty of Engineering, University of Porto, Portugal
Luís Fiolhais	INESC-ID, IST, Universidade de Lisboa, Portugal

Additional Reviewers

Ahmad Abdelfattah	Hamza Baniata	Adrien Cassagne
Prateek Agrawal	Raul Barbosa	Sebastien Cayrols
Kadir Akbudak	Christian Bartolo Burlò	Daniel Chillet
Mohammed Al Farhan	Naama Ben-David	Terry Cojean
Maicon Melo Alves	Jean Luca Bez	Soteris Constantinou
Alexandra Amarie	Mauricio Breternitz	João Costa Seco
Filipe Araujo	Rafaela Brum	Nuno Cruz Garcia
Cédric Augonnet	Rocío Carratalá-Sáez	Arthur Carvalho Cunha
Alan Ayala	Caroline Caruana	Alexandre Denis
Bartosz Balis	Pablo Carvalho	Nicolas Denoyelle

Vladimir Dimić
Petr Dobias
Panagiotis Drakatos
Anthony Dugois
Lionel Eyraud-Dubois
Pablo Ezzatti
Carlo Fantozzi
Johannes Fischer
Sebti Foufou
Emilio Francesquini
Gabriel Freytag
Giulio Gambardella
João Garcia
Afton Geil
Anja Gerbes
Fritz Goebel
David Gradinariu
Yan Gu
Amina Guermouche
Harsh Vardhan
Changwan Hong
Sergio Iserte
Kazuaki Ishizaki
Marc Jorda
Nikolaos Kallimanis
Ioanna Karageorgou
Alexandros Karakasidis
Dimitrios Karapiperis
Raffi Khatchadourian
Cédric Killian
Andreas Konstantinides
Orestis Korakitis
Matthias Korch
Sotiris Kotsiantis
Thomas Lambert
Matthew Alan Le Brun
Lucas Leandro Nesi
Ivan Lujic
Piotr Luszczek
Manolis Maragoudakis
Andras Markus
Roland Mathà
Miguel Matos
Kazuaki Matsumura
Narges Mehran
Aleksandr Mikhalev
Lei Mo
Lucas Morais
Grégory Mounié
Paschalis Mpeis
Pratik Nayak
Alexandros Ntoulas
Ken O'Brien
Cristobal Ortega
Yiannis Papadopoulos
Marina Papatriantafilou
Miguel Pardal
Maciej Pawlik
Lucian Petrica
Geppino Pucci
Long Qu
Vinod Rebello
Tobias Ribizel
Rafael Rodríguez-Sánchez
Evangelos Sakkopoulos
Matheus S. Serpa
João A. Silva
João Soares
João Sobrinho
Nehir Sonmez
Philippe Swartvagher
Antonio Teofilo
Luan Teylo
Yu-Hsiang Tsai
Volker Turau
Jerzy Tyszkiewicz
Manolis Tzagarakis
Yaman Umuroglu
Alejandro Valero
Huijun Wang
Jasmine Xuereb

Euro-Par 2021 Invited Talks

Big Data and Extreme-Scales: Computational Science in the 21st Century

Manish Parashar

University of Utah, USA
manish.parashar@utah.edu

Extreme scales and big data are essential to computational and data-enabled science and engineering in the 21st, promising dramatic new insights into natural and engineered systems. However, data-related challenges are limiting the potential impact of application workflows enabled by current and emerging extreme scale, high-performance, distributed computing environments. These data-intensive application workflows involve dynamic coordination, interactions and data coupling between multiple application processes that run at scale on different resources, and with services for monitoring, analysis and visualization and archiving, and present challenges due to increasing data volumes and complex data-coupling patterns, system energy constraints, increasing failure rates, etc. In this talk I will explore some of these challenges and investigate how solutions based on data sharing abstractions, managed data pipelines, data-staging service, and in-situ/in-transit data placement and processing can be used to help address them. This research is part of the DataSpaces project at the Scientific Computing and Imaging (SCI) Institute, University of Utah.

HPC for Bioinformatics: The Genetic Sequence Comparison Quest for Performance

Alba Cristina Melo

University of Brasilia, Brazil
alves@unb.br

Genetic Sequence Comparison is an important operation in Bioinformatics, executed routinely worldwide. Two relevant algorithms that compare genetic sequences are the Smith-Waterman (SW) algorithm and Sankoff's algorithm. The Smith-Waterman algorithm is widely used for pairwise comparisons and it obtains the optimal result in quadratic time - $O(n^2)$, where n is the length of the sequences. The Sankoff algorithm is used to structurally align two sequences and it computes the optimal result in $O(n^4)$ time. In order to accelerate these algorithms, many parallel strategies were proposed in the literature. However, the alignment of whole chromosomes with hundreds of millions of characters with the SW algorithm is still a very challenging task, which requires extraordinary computing power. Likewise, obtaining the structural alignment of two sequences with the Sankoff algorithm requires parallel approaches. In this talk, we first present our MASA-CUDAlign tool, which was used to pairwise align real DNA sequences with up to 249 millions of characters in a cluster with 512 GPUs, achieving the best performance in the literature in 2021. We will present and discuss the innovative features of the most recent version of MASA-CUDAlign: parallelogram execution, incremental speculation, block pruning and score-share balancing strategies. We will also show performance and energy results in homogeneous and heterogeneous GPU clusters. Then, we will discuss the design of our CUDA-Sankoff tool and its innovative strategy to exploit multi-level wavefront parallelism. At the end, we will show a COVID-19 case study, where we use the tools discussed in this talk to compare the SARS-CoV-2 genetic sequences, considering the reference sequence and its variants.

Knowledge Graphs, Graph AI, and the Need for High-performance Graph Computing

Keshav Pingali

Katana Graph & The University of Texas at Austin, USA
pingali@cs.utexas.edu

Knowledge Graphs now power many applications across diverse industries such as FinTech, Pharma and Manufacturing. Data volumes are growing at a staggering rate, and graphs with hundreds of billions edges are not uncommon. Computations on such data sets include querying, analytics, and pattern mining, and there is growing interest in using machine learning to perform inference on large graphs. In many applications, it is necessary to combine these operations seamlessly to extract actionable intelligence as quickly as possible. Katana Graph is a start-up based in Austin and the Bay Area that is building a scale-out platform for seamless, high-performance computing on such graph data sets. I will describe the key features of the Katana Graph Engine that enable high performance, some important use cases for this technology from Katana's customers, and the main lessons I have learned from doing a startup after a career in academia.

Euro-Par 2021 Topics Overview

Topic 1: Compilers, Tools and Environments

Frank Hannig and Gabriel Falcao

This topic addresses programming tools and system software for all kinds of parallel computer architectures, ranging from low-power embedded high-performance systems, multi- and many-core processors, accelerators to large-scale computers and cloud computing. Focus areas include compilation and software testing to design well-defined components and verify their necessary structural, behavioral, and parallel interaction properties. It deals with tools, analysis software, and runtime environments to address the challenges of programming and executing the parallel architectures mentioned above. Moreover, the topic deals with methods and tools for optimizing non-functional properties such as performance, programming productivity, robustness, energy efficiency, and scalability.

The topic received eight submissions across the subjects mentioned before. The papers were thoroughly reviewed by the six program committee members of the topic and external reviewers. Each submission was subjected to rigorous review from four peers. After intensely scrutinizing the reviews, we were pleased to select two high-quality papers for the technical program, corresponding to a per-topic acceptance ratio of 25%.

The first accepted paper "ALONA: Automatic Loop Nest Approximation with Reconstruction and Space Pruning" by Daniel Maier et al. deals with approximate computing. More specifically, the contribution advances the state-of-the-art in pattern-based multi-dimensional loop perforation. The second paper "Automatic low-overhead load-imbalance detection in MPI applications" by Peter Arzt et al. presents a method for detecting load imbalances in MPI applications. The approach has been implemented into the Performance Instrumentation Refinement Automation (PIRA) framework.

We would like to thank the authors who responded to our call for papers, the members of the Program Committee and the additional external reviewers who, with their opinion and expertise, ensured a program of the highest quality. Many thanks to Leonel Sousa for the tremendous overall organization of Euro-Par 2021 and his engaging interaction.

Topic 2: Performance and Power Modeling, Prediction and Evaluation

Didem Unat and Aleksandar Ilic

In recent years, a range of novel methods and tools have been developed for the evaluation, design, and modeling 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 papers submitted to this topic represent researchers working on different aspects of performance modeling, 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.). The accepted papers present novel research in all areas of performance modeling, prediction and evaluation, and help bringing together current theory and practice.

The topic received 14 submissions, which were thoroughly reviewed by the 11 members of the topic program committee and external reviewers. Out of all submissions and after a careful and detailed discussion among committee members, we finally decided to accept 3 papers, resulting in a per-topic acceptance ratio of 21%.

We would like to thank the authors for their submissions, the Euro-Par 2021 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

Oliver Sinnen and Jorge Barbosa

New computing systems offer the opportunity to reduce the response times and the energy consumption of the applications by exploiting the levels of parallelism. Modern computer architectures are often composed of heterogeneous compute resources and exploiting them efficiently is a complex and challenging task. Scheduling and load balancing techniques are key instruments to achieve higher performance, lower energy consumption, reduced resource usage, and real-time properties of applications.

This topic attracts papers on all aspects related to scheduling and load balancing on parallel and distributed machines, from theoretical foundations for modelling 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 submissions were received in this track, each of which received at least three reviews, the large majority four. Following a thorough and lively discussion of the reviews among the 16 PC members, 7 submissions were accepted, giving an acceptance rate of 30%.

The chairs would like to sincerely thank all the authors for their high quality submissions, the Euro-Par 2021 Organizing Committee for all their valuable help, and the PC members 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: Data Management, Analytics and Machine Learning

Alex Delis and Helena Aidos

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, cloud and grid data-intensive processing, parallel and distributed machine learning, HPC in situ data analytics, parallel storage systems, scalable data processing workflows, and distributed stream processing were all in the scope of this topic.

This year, the topic received 20 submissions, which were thoroughly reviewed by the 12 members of the track program committee and external reviewers. Out of all the submissions and after a careful and detailed discussion among committee members, we finally decided to accept 4 papers, resulting in a per-topic acceptance ratio of 20%.

We would like to express our thanks to the authors for their submissions, the Euro-Par 2021 Organizing Committee for their help throughout all the process, the PC members and the external reviewers for providing timely and detailed reviews, and for participating in the discussion we carried on after the reviews were received.

Topic 5: Cluster, Cloud and Edge Computing

Radu Prodan and Luís Veiga

While the term Cluster Computing is hardware-oriented and determines the organization of large computer systems at one location, the term Cloud Computing usually focuses on using these large systems at a distributed scale, typically combined with virtualization technology and a business model on top. Despite their differences, there exist many complementary interdependencies between both fields that need further exploration. This topic of the Euro-Par conference is open to new research, which spans across these two related areas. In both Cluster and Cloud Computing, much relevant research works focus on performance, reliability, energy efficiency, and the impact of novel processor architectures. Since Cloud Computing tries to hide the hardware and system software details from the users, research issues include various forms of virtualization and their impact on performance, resource management, and business models that address system owner and user interests.

In the last years, and primarily due to the increasing number of IoT applications, the combination of local resources and Cloud Computing, also referred to as Fog and Edge Computing, has received growing interest. This concept has led to many research questions, like an appropriate distribution of subtasks to the available systems under various constraints.

This year, 14 papers were submitted to this track and received at least four reviews from the 13 program committee members. Following the thorough discussion of the reviews, the conference chairs decided to include five submissions to the track in the overall conference program. The accepted papers address relevant and timely topics such as cloud computing, fault-tolerance, edge computing, energy efficiency.

The track chairs thank all the authors for their submissions, the Euro-Par 2021 Organizing Committee for their help throughout the process, the PC members, and the reviewers to provide timely and detailed reviews and participate in the discussions.

Topic 6: Theory and Algorithms for Parallel and Distributed Processing

Andrea Pietracaprina and João M. Lourenço

Nowadays parallel and distributed processing is ubiquitous. Multicore processors are available on smartphones, laptops, servers and supercomputing nodes. Also, many devices cooperate in fully distributed and heterogeneous systems to provide a wide array of services. Despite recent years have witnessed astonishing progress in this field, many research challenges remain open concerning fundamental issues as well as the design and analysis of efficient, scalable, and robust algorithmic solutions with provable performance and quality guarantees.

This year, a total of 12 submissions were received in this track. Each submission received four reviews from the eight PC members. Following the thorough discussion of the reviews, four original and high quality papers have been accepted to this general topic of the theory of parallel and distributed algorithms, with the acceptance rate of 33%.

We would like to thank the authors for their excellent submissions, the Euro-Par 2021 Organizing Committee for their help throughout all the process, and the PC members and the external reviewers for providing timely and detailed reviews, and for participating in the discussions that helped reach the decisions.

Topic 7: Parallel and Distributed Programming, Interfaces, and Languages

Alfredo Goldman and Ricardo Chaves

Parallel and distributed applications require appropriate programming abstractions and models, efficient design tools, parallelization techniques and practices. This topic attracted papers presenting 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 accepted papers emphasize research on high-performance, resilient, portable, and scalable parallel programs via appropriate parallel and distributed programming model, interface and language support. Contributions that assess programming abstractions and automation for usability, performance, task-based parallelism or scalability were valued.

This year, the topic received 14 submissions, which were thoroughly reviewed by the 11 members of the track program committee and external reviewers. After careful and detailed discussion among committee members, we decided to accept 5 of the submissions, giving a per-topic acceptance ratio of 36%.

The topic chairs would like to thank all the authors who submitted papers for their contribution to the success of this track, the Euro-Par 2021 Committee for their support, and the external reviewers for their high-quality reviews and their valuable feedback.

Topic 8: Multicore and Manycore Parallelism

Enrique S. Quintana Ortí and Nuno Roma

Modern homogeneous and heterogeneous multicore and manycore 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 address 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 multicore and manycore systems, as well as stand-alone systems with large numbers of cores and various types of accelerators; this can also include hybrid and heterogeneous systems with different types of multicore processors. It focuses on novel research and solutions in the form of programming models, frameworks and languages; compiler optimizations and techniques; lock-free algorithms and data structures; transactional memory advances; performance and power trade-offs and scalability; libraries and runtime systems; innovative applications and case studies; techniques and tools for discovering, analysing, understanding, and managing multicore parallelism; and in general, tools and techniques to increase the programmability of multicore, manycore, and heterogeneous systems, in the context of general-purpose, high-performance, and embedded parallel computing.

This year 7 papers covering some of these issues were submitted. Each of them was reviewed by four reviewers. Finally, one regular paper was selected. It proposes an extension to the oneAPI programming model in order to provide the ability to exploit multiple heterogeneous devices when executing the same kernel following a co-execution strategy. It also discusses the implementation of common load-balancing schemes, including static and dynamic allocation.

We would like to express our gratitude to all the authors for submitting their work. We also thank the reviewers for their great job and useful comments. Finally, we would like to thank the Euro-Par organization and steering committees for their continuous help, and for producing a nice working environment to smooth the process.

Topic 9: Parallel Numerical Methods and Applications

Stanimire Tomov and Sergio Jimenez

The need for high-performance computing is driven by the need for large-scale simulation and data analysis in science and engineering, finance, life sciences, etc. This requires the design of highly scalable numerical methods and algorithms that are able to efficiently exploit modern computer architectures. The scalability of these algorithms and methods and their ability to effectively use high-performance heterogeneous resources is critical to improving the performance of computational and data science applications.

This conference topic provides 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. The submitted papers address algorithmic design, implementation details, performance analysis, as well as integration of parallel numerical methods in large-scale applications.

This year, the topic received 14 submissions, which were thoroughly reviewed by the 7 members of the track program committee and a number of external reviewers. Each submission received four reviews. After careful and constructive discussions among committee members, we decided to accept five papers, resulting in a per-topic acceptance ratio of 36%.

We would like to sincerely thank all the authors for their submissions, the Euro-Par 2021 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 parallel numerical methods and applications.

Topic 10: High-performance Architectures and Accelerators

Samuel Thibault and Pedro Tomás

Various computing platforms provide distinct potentials for achieving massive performance levels. Beyond general-purpose multi-processors, examples of computing platforms include graphics processing units (GPUs), multi-/many-core co-processors, as well as customizable devices, such as FPGA-based systems, streaming data-flow architectures or low-power embedded systems. However, fully exploiting the computational performance of such devices and solutions often requires tuning the target applications to identify and exploit the parallel opportunities.

Hence, this topic explores new directions across this variety of architecture possibilities, with contributions along the whole design stack. On the hardware side, the scope of received papers spans from general-purpose to specialized computing architectures, including methodologies to efficiently exploit the memory hierarchy and to manage network communications. On the software side, the scope of received papers included libraries, runtime tools and benchmarks. Additionally, we also received application-specific submissions that contributed with new insights and/or solutions to fully exploit different high-performance computing platforms for signal processing, big data and machine learning application domains.

The topic received 10 submissions, which were thoroughly reviewed by the 7 members of the track program committee and external reviewers. Out of all the submissions and after a careful and detailed discussion among committee members, we finally decided to accept 2 papers, resulting in a per-topic acceptance ratio of 20%.

The topic chairs would like to thank all the authors who submitted papers for their contribution to the success of this topic, the Euro-Par 2021 Committee for their support, as well as all the external reviewers for their high-quality reviews and their valuable feedback.

Contents

Compilers, Tools and Environments

ALONA: Automatic Loop Nest Approximation with Reconstruction and Space Pruning 3
Daniel Maier, Biagio Cosenza , and Ben Juurlink

Automatic Low-Overhead Load-Imbalance Detection in MPI Applications . . . 19
Peter Arzt, Yannic Fischler, Jan-Patrick Lehr, and Christian Bischof

Performance and Power Modeling, Prediction and Evaluation

Trace-Based Workload Generation and Execution 37
Yannis Sfakianakis, Eleni Kanellou, Manolis Marazakis, and Angelos Bilas

Update on the Asymptotic Optimality of LPT. 55
Anne Benoit , Louis-Claude Canon, Redouane Elghazi, and Pierre-Cyrille Héam

E2EWatch: An End-to-End Anomaly Diagnosis Framework for Production HPC Systems 70
Burak Aksar, Benjamin Schwaller, Omar Aaziz, Vitus J. Leung, Jim Brandt, Manuel Egele, and Ayse K. Coskun

Scheduling and Load Balancing

Collaborative GPU Preemption via Spatial Multitasking for Efficient GPU Sharing 89
Zhuoran Ji and Cho-Li Wang

A Fixed-Parameter Algorithm for Scheduling Unit Dependent Tasks with Unit Communication Delays 105
Ning Tang and Alix Munier Kordon

Plan-Based Job Scheduling for Supercomputers with Shared Burst Buffers. . . 120
Jan Kopanski and Krzysztof Rzadca

Taming Tail Latency in Key-Value Stores: A Scheduling Perspective 136
Sonia Ben Mokhtar, Louis-Claude Canon, Anthony Dugois, Loris Marchal, and Etienne Rivière

A Log-Linear $(2+5/6)$-Approximation Algorithm for Parallel Machine Scheduling with a Single Orthogonal Resource 151
Adrian Naruszko, Bartłomiej Przybylski, and Krzysztof Rządca

An MPI-based Algorithm for Mapping Complex Networks onto Hierarchical Architectures 167
Maria Predari, Charilaos Tzovas, Christian Schulz, and Henning Meyerhenke

Pipelined Model Parallelism: Complexity Results and Memory Considerations 183
Olivier Beaumont, Lionel Eyraud-Dubois, and Alena Shilova

Data Management, Analytics and Machine Learning

Efficient and Systematic Partitioning of Large and Deep Neural Networks for Parallelization 201
Haoran Wang, Chong Li, Thibaut Tachon, Hongxing Wang, Sheng Yang, Sébastien Limet, and Sophie Robert

A GPU Architecture Aware Fine-Grain Pruning Technique for Deep Neural Networks 217
Kyusik Choi and Hoeseok Yang

Towards Flexible and Compiler-Friendly Layer Fusion for CNNs on Multicore CPUs 232
Zhongyi Lin, Evangelos Georganas, and John D. Owens

Smart Distributed DataSets for Stream Processing 249
Tiago Lopes, Miguel Coimbra, and Luís Veiga

Cluster, Cloud and Edge Computing

Colony: Parallel Functions as a Service on the Cloud-Edge Continuum 269
Francesc Lordan, Daniele Lezzi, and Rosa M. Badia

Horizontal Scaling in Cloud Using Contextual Bandits 285
David Delande, Patricia Stolf, Raphaël Feraud, Jean-Marc Pierson, and André Bottaro

Geo-distribute Cloud Applications at the Edge 301
Ronan-Alexandre Cherrueau, Marie Delavergne, and Adrien Lèbre

A Fault Tolerant and Deadline Constrained Sequence Alignment Application on Cloud-Based Spot GPU Instances 317
Rafaela C. Brum, Walisson P. Sousa, Alba C. M. A. Melo, Cristiana Bentes, Maria Clicia S. de Castro, and Lúcia Maria de A. Drummond

Sustaining Performance While Reducing Energy Consumption: A Control Theory Approach 334
Sophie Cerf, Raphaël Bleuse, Valentin Reis, Swann Perarnau, and Éric Rutten

Theory and Algorithms for Parallel and Distributed Processing

Algorithm Design for Tensor Units 353
Rezaul Chowdhury, Francesco Silvestri, and Flavio Vella

A Scalable Approximation Algorithm for Weighted Longest Common Subsequence 368
Jeremy Buhler, Thomas Lavastida, Kefu Lu, and Benjamin Moseley

TSLQueue: An Efficient Lock-Free Design for Priority Queues 385
Adones Rukundo and Philippas Tsigas

G-Morph: Induced Subgraph Isomorphism Search of Labeled Graphs on a GPU 402
Bryan Rowe and Rajiv Gupta

Parallel and Distributed Programming, Interfaces, and Languages

Accelerating Graph Applications Using Phased Transactional Memory 421
Catalina Munoz Morales, Rafael Murari, Joao P. L. de Carvalho, Bruno Chinelato Honorio, Alexandro Baldassin, and Guido Araujo

Efficient GPU Computation Using Task Graph Parallelism 435
Dian-Lun Lin and Tsung-Wei Huang

Towards High Performance Resilience Using Performance Portable Abstractions 451
Nicolas Morales, Keita Teranishi, Bogdan Nicolae, Christian Trott, and Franck Cappello

Enhancing Load-Balancing of MPI Applications with Workshare 466
Thomas Dionisi, Stephane Bouhrour, Julien Jaeger, Patrick Carribault, and Marc Pérache

Particle-In-Cell Simulation Using Asynchronous Tasking 482
Nicolas Guidotti, Pedro Ceyrat, João Barreto, José Monteiro, Rodrigo Rodrigues, Ricardo Fonseca, Xavier Martorell, and Antonio J. Peña

Multicore and Manycore Parallelism

Exploiting Co-execution with OneAPI: Heterogeneity from a Modern Perspective. 501
Raúl Nozal and Jose Luis Bosque

Parallel Numerical Methods and Applications

Designing a 3D Parallel Memory-Aware Lattice Boltzmann Algorithm on Manycore Systems 519
Yuankun Fu and Fengguang Song

Fault-Tolerant LU Factorization Is Low Cost 536
Camille Coti, Laure Petrucci, and Daniel Alberto Torres González

Mixed Precision Incomplete and Factorized Sparse Approximate Inverse Preconditioning on GPUs. 550
Fritz Göbel, Thomas Grützmacher, Tobias Ribizel, and Hartwig Anzt

Outsmarting the Atmospheric Turbulence for Ground-Based Telescopes Using the Stochastic Levenberg-Marquardt Method 565
Yuxi Hong, El Houcine Bergou, Nicolas Doucet, Hao Zhang, Jesse Cranney, Hatem Ltaief, Damien Gratadour, Francois Rigaut, and David Keyes

GPU-Accelerated Mahalanobis-Average Hierarchical Clustering Analysis. . . . 580
Adam Šmelko, Miroslav Kratochvíl, Martin Kruliš, and Tomáš Sieger

High Performance Architectures and Accelerators

PrioRAT: Criticality-Driven Prioritization Inside the On-Chip Memory Hierarchy 599
Vladimir Dimić, Miquel Moretó, Marc Casas, and Mateo Valero

Optimized Implementation of the HPCG Benchmark on Reconfigurable Hardware 616
Alberto Zeni, Kenneth O'Brien, Michaela Blott, and Marco D. Santambrogio

Author Index 631