Lecture Notes in Computer Science **10778**

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

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

Roman Wyrzykowski · Jack Dongarra
Ewa Deelman · Konrad Karczewski (Eds.)

Parallel Processing and Applied Mathematics

12th International Conference, PPAM 2017
Lublin, Poland, September 10–13, 2017
Revised Selected Papers, Part II

Springer

Editors
Roman Wyrzykowski
Czestochowa University of Technology
Czestochowa
Poland

Ewa Deelman
University of Southern California
Marina Del Rey, CA
USA

Jack Dongarra
University of Tennessee
Knoxville, TN
USA

Konrad Karczewski
Czestochowa University of Technology
Czestochowa
Poland

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-78053-5 ISBN 978-3-319-78054-2 (eBook)
https://doi.org/10.1007/978-3-319-78054-2

Library of Congress Control Number: 2018937375

LNCS Sublibrary: SL1 – Theoretical Computer Science and General Issues

Printed on acid-free paper

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

Preface

This volume comprises the proceedings of the 12th International Conference on Parallel Processing and Applied Mathematics – PPAM 2017, which was held in Lublin, Poland, September 10–13, 2017. It was organized by the Department of Computer and Information Science of the Czestochowa University of Technology together with Maria Curie-Skłodowska University in Lublin, under the patronage of the Committee of Informatics of the Polish Academy of Sciences, in technical cooperation with the IEEE Computer Society and ICT COST Action IC1305 "Network for Sustainable Ultrascale Computing (NESUS)". The main organizer was Roman Wyrzykowski.

PPAM is a biennial conference. Ten previous events have been held in different places in Poland since 1994. The proceedings of the last six conferences have been published by Springer in the *Lecture Notes in Computer Science* series (Nałęczów, 2001, vol. 2328; Częstochowa, 2003, vol. 3019; Poznań, 2005, vol. 3911; Gdańsk, 2007, vol. 4967; Wrocław, 2009, vols. 6067 and 6068; Toruń, 2011, vols. 7203 and 7204; Warsaw, 2013, vols. 8384 and 8385; Kraków, 2015, vols. 9573 and 9574).

The PPAM conferences have become an international forum for the exchange of ideas between researchers involved in parallel and distributed computing, including theory and applications, as well as applied and computational mathematics. The focus of PPAM 2017 was on models, algorithms, and software tools that facilitate efficient and convenient utilization of modern parallel and distributed computing architectures, as well as on large-scale applications, including big data and machine learning problems.

This meeting gathered more than 170 participants from 25 countries. A strict review process resulted in the acceptance of 100 contributed papers for publication in the conference proceedings, while approximately 42% of the submissions were rejected. For regular tracks of the conference, 49 papers were selected from 98 submissions, giving an acceptance rate of 50%.

The regular tracks covered such important fields of parallel/distributed/cloud computing and applied mathematics as:

- Numerical algorithms and parallel scientific computing, including parallel matrix factorizations and particle methods in simulations
- Task-based paradigm of parallel computing
- GPU computing
- Parallel non-numerical algorithms
- Performance evaluation of parallel algorithms and applications
- Environments and frameworks for parallel/distributed/cloud computing
- Applications of parallel computing
- Soft computing with applications

The invited talks were presented by:

- Rosa Badia from the Barcelona Supercomputing Center (Spain)
- Franck Cappello from the Argonne National Laboratory (USA)
- Cris Cecka from NVIDIA and Stanford University (USA)
- Jack Dongarra from the University of Tennessee and ORNL (USA)
- Thomas Fahringer from the University of Innsbruck (Austria)
- Dominik Göddeke from the University of Stuttgart (Germany)
- William Gropp from the University of Illinois Urbana-Champaign (USA)
- Georg Hager from the University of Erlangen-Nurnberg (Germany)
- Alexey Lastovetsky from the University College Dublin (Ireland)
- Satoshi Matsuoka from the Tokyo Institute of Technology (Japan)
- Karlheinz Meier from the University of Heidelberg (Germany)
- Manish Parashar from Rutgers University (USA)
- Jean-Marc Pierson from the University Paul Sabatier (France)
- Uwe Schwiegelshohn from TU Dortmund (Germany)
- Bronis R. de Supinski from the Lawrence Livermore National Laboratory (USA)
- Boleslaw K. Szymanski from the Rensselaer Polytechnic Institute (USA)
- Michela Taufer from the University of Delaware (USA)
- Andrei Tchernykh from the CICESE Research Center (Mexico)
- Jeffrey Vetter from the Oak Ridge National Laboratory and Georgia Institute of Technology (USA)

Important and integral parts of the PPAM 2017 conference were the workshops:

- Workshop on Models, Algorithms, and Methodologies for Hierarchical Parallelism in New HPC Systems organized by Giulliano Laccetti and Marco Lapegna from the University of Naples Federico II (Italy), and Raffaele Montella from the University of Naples Parthenope (Italy)
- Workshop on Power and Energy Aspects of Computation — PEAC 2017 organized by Ariel Oleksiak from the Poznan Supercomputing and Networking Center (Poland) and Laurent Lefevre from Inria (France)
- Workshop on Scheduling for Parallel Computing — SPC 2017 organized by Maciej Drozdowski from the Poznań University of Technology (Poland)
- The 7th Workshop on Language-Based Parallel Programming Models — WLPP 2017 organized by Ami Marowka from Bar-Ilan University (Israel)
- Workshop on PGAS Programming organized by Piotr Bała from Warsaw University (Poland)
- Special Session on Parallel Matrix Factorizations organized by Marian Vajtersic from the University of Salzburg (Austria) and Slovak Academy of Sciences
- Minisymposium on HPC Applications in Physical Sciences organized by Grzegorz Kamieniarz and Wojciech Florek from the A. Mickiewicz University in Poznań (Poland)
- Minisymposium on High-Performance Computing Interval Methods organized by Bartłomiej J. Kubica from Warsaw University of Technology (Poland)
- Workshop on Complex Collective Systems organized by Paweł Topa and Jarosław Wąs from the AGH University of Science and Technology in Kraków (Poland)

The PPAM 2017 meeting began with three tutorials:

- Scientific Computing with GPUs, by Dominik Göddeke from the University of Stuttgart (Germany) and Robert Strzodka from Heidelberg University (Germany)
- Advanced OpenMP Tutorial, by Dirk Schmidl from RWTH Aachen University (Germany)
- Parallel Computing in Java, by Piotr Bała from Warsaw University (Poland), and Marek Nowicki from the Nicolaus Copernicus University in Toruń (Poland)

A new topic at PPAM 2017 was "Particle Methods in Simulations." Particle-based and Lagrangian formulations are all-time classics in supercomputing and have been wrestling with classic mesh-based approaches such as finite elements for quite a while now, in terms of computational expressiveness and efficiency. Computationally, particle formalisms benefit from very costly inter-particle interactions. These interactions with high arithmetic intensity make them reasonably "low-hanging" fruits in supercomputing with its notoriously limited bandwidth and high concurrency.

Surprisingly, PPAM 2017 was shaped by articles that give up on expensive particle–particle interactions: discrete element methods (DEM) study rigid bodies which interact only rarely once they are in contact, while particle-in-cell (PIC) methods use the physical expressiveness of Lagrangian descriptions but make the particles interact solely locally with a surrounding grid. It is obvious that the lack of direct long-range particle–particle interaction increases the concurrency of the algorithms. Yet, it comes at a price. With low arithmetic intensity, all data structures have to be extremely fine-tuned to perform on modern hardware, and load-balancing has to be lightweight. Codes cannot afford to resort data inefficiently all the time, move around too much data, or work with data structures that are ill-suited for vector processing, while notably the algorithmic parts with limited vectorization potential have to be revisited and maybe rewritten for emerging processors tailored toward stream processing.

The new session "Particle Methods in Simulations" provided a platform for some presentations with interesting and significant contributions addressing these challenges:

- Contact problems are rephrased as continuous minimization problems coupled with a posteriori validity checks, which allows codes to vectorize at least the first step aggressively (by K. Krestenitis, T. Weinzierl, and T. Koziara)
- Classic PIC is recasted into a single-touch algorithm with only few synchronization points, which releases pressure from the memory subsystem (by Y. Barsamian, A. Chargueraud, and A. Ketterlin)
- Cell-based shared memory parallelization of PIC is revised from a scheduling point of view and tailored parallelization schemes are developed, which anticipate the enormous per-cell load imbalances resulting from clustered particles (by A. Larin et al.)
- Particle sorting algorithms are revisited that make the particles be stored in memory in the way they are later accessed by the algorithm even though the particles tend to move through the domain quickly (by A. Dorobisz et al).

Another new topic at PPAM 2017 was "Task-Based Paradigm of Parallel Computing." Task-based parallel programming models have appeared in the recent years as an alternative to traditional parallel programming models, both for fine-grain and

coarse-grain parallelism. In this paradigm, the task is the unit of execution and traditionally a data-dependency graph of the application tasks represents the application. From this graph, the potential parallelism of the application is exploited, enabling an asynchronous execution of the tasks that do not require explicit fork-join structures.

Research topics in the area are multiple, from the specification of the syntax or programming interfaces, the definition of new scheduling and resource management algorithms that take into account different metrics, the design of the interfaces with the actual infrastructure, or new algorithms specified in this parallel paradigm. As an example of the success of this paradigm, the OpenMP standard has adopted this paradigm in its latest releases.

This topic was presented at PPAM 2017 in the form of a session that consisted of several presentations from various topics:

- "A Proposal for a Unified Interface for Task-Based Programming Models That Enables the Execution of Applications in Multiple Parallel Environments" (by A. Zafari)
- "A Comparison of Time and Energy Oriented Scheduling for Task-Based Programs, Which Is Based on Real Measured Data for the Tasks Leading to Diverse Effects Concerning Time, Energy, and Power Consumption" (by T. Rauber and G. Rünger)
- "A Study of a Set of Experiments with the Sparse Cholesky Decomposition on Multicore Platforms, Using a Parametrized Task Graph Implementation" (by I. Duff and F. Lopez)
- "A Task-Based Algorithm for Reordering the Eigenvalues of a Matrix in Real Schur Form, Which Is Realized on Top of the StarPU Runtime System" (by M. Myllykoski)

A new topic at PPAM 2017 was the "Special Session on Parallel Matrix Factorizations." Nowadays, in order to meet demands of high-performance computing, it is necessary to pay serious attention to the development of fast, reliable, and communication-efficient algorithms for solving kernel linear algebra problems. Tasks that lead to matrix decomposition computations are undoubtedly some of the most frequent problems encountered in this field. Therefore, the aim of the special session was to present new results from parallel linear algebra with an emphasis on methods and algorithms for factorizations and decompositions of large sparse and dense matrices. Both theoretical aspects and software issues related to this problem area were considered for submission.

The topics of the special session focused on: (a) efficient algorithms for the EVD/SVD/NMF decompositions of large matrices, their design and analysis; (b) implementation of parallel matrix factorization algorithms on parallel CPU and GPU systems; (c) usage of parallel matrix factorizations for solving problems arising in scientific and technical applications. Seven papers were accepted for presentation, which covered the session topics. Geographically, the authors were dispersed among two continents and five countries. The individual themes of the contributions included:

- "New Preconditioning for the One-Sided Block-Jacobi Singular Value Decomposition Algorithm" (by M. Bečka, G. Okša, and E. Vidličková)

- "Using the Cholesky QR Method in the Full-Blocked One-Sided Jacobi Algorithm" (by S. Kudo and Y. Yamamoto)
- "Parallel Divide-and-Conquer Algorithm for Solving Tridiagonal Eigenvalue Problems on Manycore Systems" (by Y. Hirota and I. Toshiyuki)
- "Structure-Preserving Technique in the Block SS-Hankel Method for Solving Hermitian Generalized Eigenvalue Problems" (by A. Imakura, Y. Futamura, and T. Sakurai)
- "Parallel Inverse of Non-Hermitian Block Tridiagonal Matrices" (by L. Spellacy and D. Golden)
- "Tunability of a New Hessenberg Reduction Algorithm Using Parallel Cache Assignment" (by M. Eljammaly, L. Karlsson, and B. Kågström)
- "Convergence and Parallelization of Nonnegative Matrix Factorization (NMF) with Newton Iteration" (by R. Kutil, M. Flatz, and M. Vajtersic).

The organizers are indebted to the PPAM 2017 sponsors, whose support was vital for the success of the conference. The main sponsor was the Intel Corporation. Another important sponsor was Lenovo. We thank all the members of the international Program Committee and additional reviewers for their diligent work in refereeing the submitted papers. Finally, we thank all the local organizers from the Częstochowa University of Technology, and Maria Curie-Skłodowska University in Lublin, who helped us run the event very smoothly. We are especially indebted to Grażyna Kołakowska, Urszula Kroczewska, Łukasz Kuczyński, Adam Tomaś, and Marcin Woźniak from the Częstochowa University of Technology; and to Przemysław Stpiczyński and Beata Bylina from Maria Curie-Skłodowska University. Also, Paweł Gepner from Intel offered great help in organizing social events for PPAM 2017, including the excursion to the Zamoyski Palace in Kozłówka and the concert of the youth accordion orchestra "Arti Sentemo" at the Royal Castle in Lublin.

We hope that this volume will be useful to you. We would like everyone who reads it to feel invited to the next conference, PPAM 2019, which will be held during September 8–11, 2019, in Białystok, the largest city in northeastern Poland, located close to the world-famous Białowieża Forest.

January 2018

Roman Wyrzykowski
Jack Dongarra
Ewa Deelman
Konrad Karczewski

Organization

Program Committee

Jan Węglarz (Honorary Chair)	Poznań University of Technology, Poland
Roman Wyrzykowski (Program Chair)	Częstochowa University of Technology, Poland
Ewa Deelman (Program Co-chair)	University of Southern California, USA
Pedro Alonso	Universidad Politecnica de Valencia, Spain
Hartwig Anzt	University of Tennessee, USA
Peter Arbenz	ETH, Zurich, Switzerland
Cevdet Aykanat	Bilkent University, Ankara, Turkey
Marc Baboulin	University of Paris-Sud, France
David A. Bader	Georgia Institute of Technology, USA
Michael Bader	TU München, Germany
Piotr Bała	Warsaw University, Poland
Krzysztof Banaś	AGH University of Science and Technology, Poland
Olivier Beaumont	Inria Bordeaux, France
Włodzimierz Bielecki	West Pomeranian University of Technology, Poland
Paolo Bientinesi	RWTH Aachen, Germany
Radim Blaheta	Czech Academy of Sciences, Czech Republic
Jacek Błażewicz	Poznań University of Technology, Poland
Pascal Bouvry	University of Luxembourg
Jerzy Brzeziński	Poznań University of Technology, Poland
Marian Bubak	AGH Kraków, Poland and University of Amsterdam, The Netherlands
Tadeusz Burczyński	Polish Academy of Sciences, Warsaw, Poland
Christopher Carothers	Rensselaer Polytechnic Institute, USA
Jesus Carretero	Universidad Carlos III de Madrid, Spain
Raimondas Čiegis	Vilnius Gediminas Technical University, Lithuania
Andrea Clematis	IMATI-CNR, Italy
Zbigniew Czech	Silesia University of Technology, Poland
Pawel Czarnul	Gdańsk University of Technology, Poland
Jack Dongarra	University of Tennessee and ORNL, USA
Maciej Drozdowski	Poznań University of Technology, Poland
Mariusz Flasiński	Jagiellonian University, Poland
Tomas Fryza	Brno University of Technology, Czech Republic
Jose Daniel Garcia	Universidad Carlos III de Madrid, Spain
Pawel Gepner	Intel Corporation, Poland
Shamsollah Ghanbari	Universiti Putra, Malaysia

Domingo Gimenez	University of Murcia, Spain
Jacek Gondzio	University of Edinburgh, Scotland, UK
Andrzej Gościński	Deakin University, Australia
Laura Grigori	Inria, France
Inge Gutheil	Forschungszentrum Juelich, Germany
Georg Hager	University of Erlangen-Nuremberg, Germany
José R. Herrero	Universitat Politecnica de Catalunya, Barcelona, Spain
Ladislav Hluchy	Slovak Academy of Sciences, Bratislava, Slovakia
Sasha Hunold	Vienna University of Technology, Austria
Aleksandar Ilic	Technical University of Lisbon, Portugal
Florin Isaila	Universidad Carlos III de Madrid, Spain
Ondrej Jakl	Institute of Geonics, Czech Academy of Sciences, Czech Republic
Emmanuel Jeannot	Inria, France
Bo Kagstrom	Umea University, Sweden
Grzegorz Kamieniarz	A. Mickiewicz University in Poznań, Poland
Eleni Karatza	Aristotle University of Thessaloniki, Greece
Ayse Kiper	Middle East Technical University, Turkey
Jacek Kitowski	Institute of Computer Science, AGH, Poland
Joanna Kołodziej	Cracow University of Technology, Poland
Jozef Korbicz	University of Zielona Góra, Poland
Stanislaw Kozielski	Silesia University of Technology, Poland
Tomas Kozubek	Technical University of Ostrava, Czech Republic
Dieter Kranzlmueller	Ludwig-Maximillian University, Munich and Leibniz Supercomputing Centre, Germany
Henryk Krawczyk	Gdańsk University of Technology, Poland
Piotr Krzyżanowski	University of Warsaw, Poland
Krzysztof Kurowski	PSNC, Poznań, Poland
Jan Kwiatkowski	Wrocław University of Technology, Poland
Giulliano Laccetti	University of Naples Federico II, Italy
Marco Lapegna	University of Naples Federico II, Italy
Alexey Lastovetsky	University College Dublin, Ireland
Laurent Lefevre	Inria and University of Lyon, France
Joao Lourenco	University Nova of Lisbon, Portugal
Tze Meng Low	Carnegie Mellon University, USA
Hatem Ltaief	KAUST, Saudi Arabia
Emilio Luque	Universitat Autonoma de Barcelona, Spain
Piotr Luszczek	University of Tennessee, USA
Victor E. Malyshkin	Siberian Branch, Russian Academy of Sciences, Russian Federation
Pierre Manneback	University of Mons, Belgium
Tomas Margalef	Universitat Autonoma de Barcelona, Spain
Svetozar Margenov	Bulgarian Academy of Sciences, Sofia
Ami Marowka	Bar-Ilan University, Israel
Norbert Meyer	PSNC, Poznań, Poland

Iosif Meyerov	Lobachevsky State University of Nizhni Novgorod, Russian Federation
Marek Michalewicz	ICM, Warsaw University, Poland
Ricardo Morla	INESC Porto, Portugal
Jarek Nabrzyski	University of Notre Dame, USA
Raymond Namyst	University of Bordeaux and Inria, France
Edoardo Di Napoli	Forschungszentrum Juelich, Germany
Gabriel Oksa	Slovak Academy of Sciences, Bratislava, Slovakia
Tomasz Olas	Częstochowa University of Technology, Poland
Ariel Oleksiak	PSNC, Poland
Ozcan Ozturk	Bilkent University, Turkey
Marcin Paprzycki	IBS PAN and SWPS, Warsaw, Poland
Dana Petcu	West University of Timisoara, Romania
Jean-Marc Pierson	University Paul Sabatier, France
Radu Prodan	University of Innsbruck, Austria
Enrique S. Quintana-Ortí	Universidad Jaime I, Spain
Omer Rana	Cardiff University, UK
Thomas Rauber	University of Bayreuth, Germany
Krzysztof Rojek	Częstochowa University of Technology, Poland
Jacek Rokicki	Warsaw University of Technology, Poland
Leszek Rutkowski	Częstochowa University of Technology, Poland
Robert Schaefer	Institute of Computer Science, AGH, Poland
Stanislav Sedukhin	University of Aizu, Japan
Franciszek Seredyński	Cardinal Stefan Wyszyński University in Warsaw, Poland
Happy Sithole	Centre for High Performance Computing, South Africa
Jurij Silc	Jozef Stefan Institute, Slovenia
Karolj Skala	Ruder Boskovic Institute, Croatia
Renata Słota	Institute of Computer Science, AGH, Poland
Leonel Sousa	Technical University of Lisbon, Portugal
Vladimir Stegailov	Joint Institute for High Temperatures of RAS, Moscow, Russian Federation
Radek Stompor	Universite Paris Diderot and CNRS, France
Przemysław Stpiczyński	Maria Curie-Skłodowska University, Poland
Maciej Stroiński	PSNC, Poznań, Poland
Reiji Suda	University of Tokyo, Japan
Lukasz Szustak	Częstochowa University of Technology, Poland
Boleslaw Szymanski	Rensselaer Polytechnic Institute, USA
Domenico Talia	University of Calabria, Italy
Andrei Tchernykh	CICESE Research Center, Ensenada, Mexico
Christian Terboven	RWTH Aachen, Germany
Parimala Thulasiraman	University of Manitoba, Canada
Roman Trobec	Jozef Stefan Institute, Slovenia
Giuseppe Trunfio	University of Sassari, Italy
Denis Trystram	Grenoble Institute of Technology, France

Marek Tudruj	Polish Academy of Sciences and Polish-Japanese Academy of Information Technology, Warsaw, Poland
Pavel Tvrdik	Czech Technical University, Prague, Czech Republic
Bora Ucar	Ecole Normale Superieure de Lyon, France
Marian Vajtersic	Salzburg University, Austria, and Slovak Academy of Sciences, Slovakia
Vladimir Voevodin	Moscow State University, Russian Federation
Kazimierz Wiatr	Academic Computer Center CYFRONET AGH, Poland
Bogdan Wiszniewski	Gdańsk University of Technology, Poland
Roel Wuyts	IMEC, Belgium
Andrzej Wyszogrodzki	Institute of Meteorology and Water Management, Warsaw, Poland
Ramin Yahyapour	University of Göttingen/GWDG, Germany
Jiangtao Yin	University of Massachusetts Amherst, USA
Krzysztof Zielinski	Institute of Computer Science, AGH, Poland
Julius Žilinskas	Vilnius University, Lithuania
Jarosław Żola	University of Buffalo, USA

Steering Committee

Jack Dongarra	University of Tennessee and ORNL, USA
Leszek Rutkowski	Częstochowa University of Technology, Poland
Boleslaw Szymanski	Rensselaer Polytechnic Institute, USA

Contents – Part II

Workshop on Language-Based Parallel Programming Models (WLPP 2017)

Workshop on PGAS Programming

Contents – Part I

Numerical Algorithms and Parallel Scientific Computing

Particle Methods in Simulations

Task-Based Paradigm of Parallel Computing

GPU Computing

Parallel Non-numerical Algorithms

Performance Evaluation of Parallel Algorithms and Applications

Applications of Parallel Computing

Soft Computing with Applications

Special Session on Parallel Matrix Factorizations