Lecture Notes in Computer Science **11381**

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/7408

Sunita Chandrasekaran
Guido Juckeland · Sandra Wienke (Eds.)

Accelerator Programming Using Directives

5th International Workshop, WACCPD 2018 Dallas, TX, USA, November 11–17, 2018 Proceedings

Springer

Editors
Sunita Chandrasekaran
Department of Computer and Information Science
University of Delaware
Newark, DE, USA

Guido Juckeland
Helmholtz-Zentrum Dresden-Rossendorf
Dresden, Sachsen, Germany

Sandra Wienke
RWTH Aachen University
Aachen, Nordrhein-Westfalen, Germany

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-12273-7 ISBN 978-3-030-12274-4 (eBook)
https://doi.org/10.1007/978-3-030-12274-4

Library of Congress Control Number: 2018968279

LNCS Sublibrary: SL2 – Programming and Software Engineering

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

2018: 5th Workshop on Accelerator Programming Using Directives (WACCPD)

http://waccpd.org/

Current hardware trends lead to ever more complex compute node architectures offering multiple, heterogeneous levels of massive parallelism. As a result the 'X' in MPI +X demands more focus. A node in a future exascale system is expected to consist of GPU-like accelerators combined with processor architectures of other types. In order to exploit the maximum available parallelism out of such systems, we are in dire need of sophisticated programming approaches that can provide scalable as well as portable solutions without compromising on performance. The expectation from the scientific community is that such solutions should allow programmers to maintain a single code base whenever possible and to avoid requiring maintaining and debugging multiple versions of the same code.

Raising the abstraction of the code is one of the effective methodologies to reduce the burden on the programmer. At the same time such a methodology will require a heroic compiler to be designed. Software abstraction-based programming models such as OpenMP and OpenACC have been serving this purpose over the past several years. These programming models address the 'X' component by providing programmers with high-level directive-based approaches to accelerate and port scientific applications to heterogeneous platforms. Such programming paradigms played a decisive role in establishing heterogeneous node architectures as a valid choice for a multitude of HPC workloads.

The focus of this workshop is to explore this 'X' component in a hybrid MPI +X programming approach. We are looking forward to technical papers discussing innovative high-level language features and their (early prototype) implementations needed to address hierarchical heterogeneous systems, stories and lessons learnt while using directives to migrate scientific legacy code to parallel processors, state-of-the-art compilation and runtime scheduling techniques, techniques to optimize performance, as well as mechanisms to keep communication and synchronization efficient.

WACCPD has been one of the major forums for bringing together users, developers, and the software and tools community to share knowledge and experiences to program emerging complex parallel computing systems.

The WACCPD 2018 workshop received 12 submissions out of which six were accepted to be presented at the workshop and published in the proceedings. The Program Committee of the workshop comprised 26 members spanning university, national laboratories, and industries. Each paper received a maximum of four reviews. Four papers were accepted straight away whereas two papers went through a shepherding phase where the authors were asked to revisit and redo the paper based on feedback obtained from reviewers. The authors were given a 7-day window to revise the paper and resubmit for the shepherd to re-review and decide on an acceptance or a rejection for the workshop.

All the 14 authors were also strongly encouraged to add source files for reproducibility purposes, following SC18 guidelines, upon request from reviewers. Three out of six accepted papers were able to add these source files, which the reviewers greatly appreciated.

The program co-chairs invited Jack Wells from ORNL to give a keynote address on "Experiences in Using Directive-Based Programming for Accelerated Computing Architectures." Dr. Jack Wells is the Director of Science for the Oak Ridge Leadership Computing Facility (OLCF), a DOE Office of Science national user facility, and the Titan supercomputer, located at Oak Ridge National Laboratory (ORNL). Wells is responsible for the scientific outcomes of the OLCF's user programs.

Based on rigorous reviews and ranking scores of all papers reviewed, the following paper won the best paper award. The authors of the best paper award also included reproducibility results to their paper, which the WACCPD workshop organizers had indicated as a criterion to be eligible to compete for the best paper award.

- Anmol Padel and Satish Puri (Marquette University, USA)
- "OpenACC-Based GPU Parallelization of Plane Sweep Algorithm for Geometric Intersection"

Emphasizing the importance of using directives for legacy scientific applications, each presenter was given two recently released textbooks on programming models, one on *Using OpenMP – The Next Step* and the other on *OpenACC for Programmers: Concepts & Strategies*.

January 2019

Sunita Chandrasekaran
Guido Juckeland
Sandra Wienke

Organization

Steering Committee

Barbara Chapman	Stony Brook, USA
Oscar Hernandez	ORNL, USA
Michael Klemm	Intel, Germany
Kuan-Ching Li	Providence University, Taiwan
Satoshi Matsuoka	Titech, Japan
Duncan Poole	OpenACC, USA
Thomas Schulthess	CSCS, Switzerland
Jeffrey Vetter	ORNL, USA

Program Co-chairs

Sunita Chandrasekaran	University of Delaware, USA
Guido Juckeland	HZDR, Germany
Sandra Wienke	RWTH Aachen University, Germany

Program Committee

David Bernholdt	ORNL, USA
James Beyer	NVIDIA, USA
Sridutt Bhalachandra	Argonne National Laboratory, USA
Makus Eisenbach	ORNL, USA
Manisha Gajbe	Intel, USA
Jeff Hammond	Intel Labs, USA
Si Hammond	SNL, USA
Christian Iwainsky	TU Darmstadt, Germany
Adrian Jackson	EPCC, UK
Arpith Jacob	Google, USA
Henri Jin	NASA-Ames, USA
Gabriele Jost	NASA-Ames, USA
Jeff Larkin	NVIDIA, USA
Seyong Lee	ORNL, USA
Kelvin Li	IBM, Canada
C. J. Newburn	NVIDIA, USA
Antonio J. Pena	BSC, Spain
Thomas Schwinge	Mentor Graphics, Germany
Sameer Shende	University of Oregon, USA
Ray Sheppard	Indiana University, USA
Peter Steinbach	Scionics, Germany
Christian Terboven	RWTH Aachen University, Germany

Xiaonan Tian	NVIDIA/PGI, USA
Veronica Vergarra Larrea	ORNL, USA
Cheng Wang	Microsoft, USA
Michael Wolfe	NVIDIA/PGI, USA

Held in conjunction with

SC18: The International Conference for High-Performance Computing, Networking, Storage and Analysis
Dallas, TX
November 11–17, 2018

Contents