

Founding Editors

Gerhard Goos

Karlsruhe Institute of Technology, Karlsruhe, Germany

Juris Hartmanis

Cornell University, Ithaca, NY, USA

Editorial Board Members

Elisa Bertino

Purdue University, West Lafayette, IN, USA

Wen Gao

Peking University, Beijing, China

Bernhard Steffen 

TU Dortmund University, Dortmund, Germany

Gerhard Woeginger 

RWTH Aachen, Aachen, Germany

Moti Yung

Columbia University, New York, NY, USA

More information about this series at <http://www.springer.com/series/7408>

Sandra Wienke · Sridutt Bhalachandra (Eds.)

Accelerator Programming Using Directives

6th International Workshop, WACCPD 2019
Denver, CO, USA, November 18, 2019
Revised Selected Papers

Editors

Sandra Wienke 
RWTH Aachen University
Aachen, Germany

Sridutt Bhalachandra 
Lawrence Berkeley National Laboratory
Berkeley, CA, USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-49942-6 ISBN 978-3-030-49943-3 (eBook)
<https://doi.org/10.1007/978-3-030-49943-3>

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer Nature Switzerland AG 2020

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

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

Preface

The ever-increasing heterogeneity in supercomputing applications has given rise to complex compute node architectures offering multiple, heterogeneous levels of massive parallelism. As a result, the ‘X’ in MPI+X demands more focus. Exploiting the maximum available parallelism out of such systems necessitates sophisticated programming approaches that can provide scalable as well as portable solutions without compromising on performance. A programmer’s expectation from the scientific community is to deliver solutions that would allow maintenance of a single code base whenever possible avoiding duplicate effort.

Raising the abstraction of the code is one of the effective methodologies to reduce the burden on the programmer while improving productivity. Software abstraction-based programming models, such as OpenMP and OpenACC, have been serving this purpose over the past several years as the compiler technology steadily improves. 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.

These proceedings contain the papers accepted for presentation at the 6th Workshop on Accelerator Programming using Directives (WACCPD 2019) – <http://waccpd.org/>. WACCPD is one of the major forums for bringing together users, developers, and the software and tools community to share knowledge and experiences when programming emerging complex parallel computing systems.

Recent architectural trends indicate a heavy reliance of future exascale machines on accelerators for performance. Toward this end, the workshop highlighted improvements to the state of the art through the accepted papers and prompted discussion through keynotes/panels that drew the community’s attention to key areas that will facilitate the transition to accelerator-based high-performance computing (HPC). The workshop aimed to showcase all aspects of heterogeneous systems discussing innovative high-level language features, lessons learned while using directives to migrate scientific legacy code to parallel processors, compilation and runtime scheduling techniques, among others.

The WACCPD 2019 workshop received 13 submissions out of which 7 were accepted to be presented at the workshop and published in these proceedings. The Program Committee of the workshop comprised 24 members spanning universities, national laboratories, and industries. Each paper received an average of five reviews.

For 2019, we encouraged all authors to add the Artifact Description (AD) to their submissions. Two additional pages were made available to authors (however without obligations) to make their code and data publicly available (e.g. on GitHub, Zenodo, Code Ocean, etc.) in support of the reproducibility initiative. As a further push, only papers with AD were considered for the Best Paper Award.

Of the 7 accepted papers, 86% had reproducibility information and these manuscripts are highlighted with an ‘artifacts available’ logo in this book.

The program co-chairs invited Dr. Nicholas James Wright from Lawrence Berkeley National Laboratory (LBL) to give a keynote address on “Perlmutter – A 2020 Pre-Exascale GPU-accelerated System for NERSC: Architecture and Application Performance Optimization.” Dr. Nicholas J. Wright is the Perlmutter chief architect and the Advanced Technologies Group lead in the National Energy Research Scientific Computing (NERSC) center at LBL. He led the effort to optimize the architecture of the Perlmutter machine, the first NERSC platform designed to meet the needs of both large-scale simulation and data analysis from experimental facilities. Nicholas has a PhD from the University of Durham in computational chemistry and has been with NERSC since 2009.

Robert Henschel from Indiana University gave an invited talk titled “The SPEC ACCEL Benchmark – Results and Lessons Learned.” Robert Henschel is the director of Research Software and Solutions at Indiana University. He is responsible for providing advanced scientific applications to researchers at Indiana University and national partners as well as providing support for computational research to the Indiana University School of Medicine. Henschel serves as the chair of the Standard Performance Evaluation Corporation (SPEC) High-Performance Group and in this role leads the development of production quality benchmarks for HPC systems. He also serves as the treasurer of the OpenACC organization. Henschel has a deep background in HPC and his research interests focus on performance analysis of parallel applications.

The workshop concluded with a panel “Convergence, Divergence, or New Approaches? – The Future of Software-Based Abstractions for Heterogeneous Supercomputing” moderated by Fernanda Foertter from NVIDIA. The panelists included:

- Christian Trott, Sandia National Laboratories, USA
- Michael Wolfe, Nvidia, USA
- Jack Deslippe, Lawrence Berkeley National Laboratory, USA
- Jeff Hammond, Intel, USA
- Johannes Doerfert, Argonne National Laboratory, USA

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 criteria to be eligible to compete for the Best Paper Award.

- Hongzhang Shan and Zhengji Zhao from Lawrence Berkeley National Laboratory, and Marcus Wagner from Cray: “Accelerating the Performance of Modal Aerosol Module of E3SM Using OpenACC”

Emphasizing the importance of using directives for legacy scientific applications, each keynote/invited speakers, panelists, and Best Paper Award winners were given a book on “OpenACC for Programmers: Concepts & Strategies.”

Organization

Steering Committee

Barbara Chapman	Stony Brook, USA
Duncan Poole	OpenACC, USA
Kuan-Ching Li	Providence University, Taiwan
Oscar Hernandez	ORNL, USA
Jeffrey Vetter	ORNL, USA

Program Co-chairs

Sandra Wienke	RWTH Aachen University, Germany
Sridutt Bhalachandra	Lawrence Berkeley National Laboratory, USA

Publicity Chair

Neelima Bayyapu	NMAM Institute of Technology, Karnataka, India
-----------------	--

Web Chair


Shu-Mei Tseng	University of California, Irvine, USA
---------------	---------------------------------------

Program Committee

Adrian Jackson	Edinburgh Parallel Computing Centre, University of Edinburgh, UK
Andreas Herten	Forschungszentrum Jülich, Germany
Arpith Jacob	Google, USA
Cheng Wang	Microsoft, USA
Christian Iwainsky	Technische Universität Darmstadt, Germany
Christian Terboven	RWTH Aachen University, Germany
Christopher Daley	Lawrence Berkeley National Laboratory, USA
C. J. Newburn	NVIDIA, USA
David Bernholdt	Oak Ridge National Laboratory, USA
Giuseppe Congiu	Argonne National Laboratory, USA
Haoqiang Jin	NASA Ames Research Center, USA
Jeff Larkin	NVIDIA, USA
Kelvin Li	IBM, USA
Manisha Gajbe	Intel, USA
Michael Wolfe	NVIDIA/PGI, USA
Ray Sheppard	Indiana University, USA
Ron Lieberman	AMD, USA





Ronan Keryell	Xilinx, USA
Seyong Lee	Oak Ridge National Laboratory, USA
Simon Hammond	Sandia National Laboratories, USA
Sameer Shende	University of Oregon, USA
Thomas Schwinge	Mentor Graphics, Germany
Tom Scogland	Lawrence Livermore National Laboratory, USA
William Sawyer	Swiss National Supercomputing Centre, Switzerland

Held in conjunction with the International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2019), Denver, USA:




Contents


Porting Scientific Applications to Heterogeneous Architectures Using Directives

	GPU Implementation of a Sophisticated Implicit Low-Order Finite Element Solver with FP21-32-64 Computation Using OpenACC	3
	<i>Takuma Yamaguchi, Kohei Fujita, Tsuyoshi Ichimura, Akira Naruse, Maddegadara Lalith, and Muneo Hori</i>	
	Acceleration in Acoustic Wave Propagation Modelling Using OpenACC/OpenMP and Its Hybrid for the Global Monitoring System	25
	<i>Noriyuki Kushida, Ying-Tsong Lin, Peter Nielsen, and Ronan Le Bras</i>	
	Accelerating the Performance of Modal Aerosol Module of E3SM Using OpenACC	47
	<i>Hongzhang Shan, Zhengji Zhao, and Marcus Wagner</i>	
	Evaluation of Directive-Based GPU Programming Models on a Block Eigensolver with Consideration of Large Sparse Matrices.	66
	<i>Fazlay Rabbi, Christopher S. Daley, Hasan Metin Aktulga, and Nicholas J. Wright</i>	

Directive-Based Programming for Math Libraries

	Performance of the RI-MP2 Fortran Kernel of GAMESS on GPUs via Directive-Based Offloading with Math Libraries	91
	<i>JaeHyuk Kwack, Colleen Bertoni, Buu Pham, and Jeff Larkin</i>	

Performance Portability for Heterogeneous Architectures

	Performance Portable Implementation of a Kinetic Plasma Simulation Mini-App.	117
	<i>Yuuichi Asahi, Guillaume Latu, Virginie Grandgirard, and Julien Bigot</i>	
	A Portable SIMD Primitive Using Kokkos for Heterogeneous Architectures . . .	140
	<i>Damodar Sahasrabudhe, Eric T. Phipps, Sivasankaran Rajamanickam, and Martin Berzins</i>	

	Author Index	165
--	-------------------------------	------------