Lecture Notes in Computer Science **10732**

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 (Eds.)

Accelerator Programming Using Directives

4th International Workshop, WACCPD 2017
Held in Conjunction with the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2017
Denver, CO, USA, November 13, 2017
Proceedings

Springer

Editors
Sunita Chandrasekaran
University of Delaware
Newark, DE
USA

Guido Juckeland
Helmholtz-Zentrum
Dresden-Rossendorf e.V.
Dresden
Germany

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-74895-5 ISBN 978-3-319-74896-2 (eBook)
https://doi.org/10.1007/978-3-319-74896-2

Library of Congress Control Number: 2018931882

LNCS Sublibrary: SL2 – Programming and Software Engineering

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

Welcome to the proceedings of WACCPD 2017, the 4th Workshop on Accelerator Programming Using Directives (http://waccpd.org/).

In the current pre-exascale era, domain and computational scientists still struggle to adapt large applications or prototype new ideas on the plethora of novel hardware architecture with diverse memory subsystems or cores with different ISAs or accelerators of varied types. The HPC community is in constant need for sophisticated software tools and techniques to port legacy code to these emerging platforms.

Given the complexity in hardware, maintaining a single code base yet achieving performance portable solutions continues to pose a daunting task. Directive-based programming models such as OpenACC and OpenMP have been tackling this issue by offering scientists a high-level approach to accelerate scientific applications and develop solutions that are portable and yet do not compromise on performance or accuracy. Such programming paradigms have facilitated complex heterogeneous systems in order to be classified as first-class citizens for HPC.

This workshop aims to solicit papers that explore innovative language features and their implementations, stories and lessons learnt while using directives to migrate scientific legacy code to parallel processors, state-of-the-art compilation and runtime scheduling techniques, performance optimization and analysis on state-of-the-art hardware etc.

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

WACCPD 2017 received 14 submissions out of which nine were accepted for presentation at the workshop and for publication in these proceedings. The Program Committee of the workshop comprised 26 members spanning various university, national labs, and industries. Each paper received at most four reviews. Four papers were accepted directly, while five 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 15-day window to fix the paper and resubmit for the reviewer to make a decision.

All 14 authors were also strongly encouraged to add source files for reproducibility purposes upon request from reviewers. Ten out of 14 authors were able to add these source files, which the reviewers greatly appreciated.

The program co-chairs invited John E. Stone from UIUC to give a keynote address on “Using Accelerator Directives to Adapt Science Applications for State-of-the-Art HPC Architectures.” John is Senior Research Programmer at the Theoretical and Computational Biophysics Group and NIH Center for Macromolecular Modeling and Bioinformatics of the University of Illinois at Urbana-Champaign.

The invited talk was given by Randy Allen, Director of Advanced Research in the Embedded Systems Division of Mentor Graphics. His talk was titled "The Challenges Faced by OpenACC Compilers."

Based on rigorous reviews and ranking scores of all papers reviewed, we arrived at two best paper award recipients this year. They were:

- Takuma Yamaguchi, Kohei Fujita, Tsuyoshi Ichimura, Muneo Hori, Maddegedara Lalith, and Kengo Nakajima (University of Tokyo, Japan).
  "Implicit Low Order Unstructured Finite-Element Multiple Simulation Enhanced by Dense Computation Using OpenACC"
- Khalid Ahmad (University of Utah, USA) and Michael Wolfe (PGI/NVIDIA)
  "Automatic Testing of OpenACC Applications"

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." The attendees were given reference guides of both models.

January 2018

Sunita Chandrasekaran
Guido Juckeland

Organization

Steering Committee

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

Program Co-chairs

Sunita Chandrasekaran	University of Delaware, USA
Guido Juckeland	HZDR, Germany

Program Committee

James Beyer	NVIDIA, USA
Henri Callandra	TOTAL, USA
Robert Dietrich	TU Dresden, Germany
Mark Govette	NOAA, USA
Georg Hager	FAU, Germany
Jeff Hammond	Intel, USA
Christian Iwainsky	TU Darmstadt, Germany
Arpith J. Jacob	IBM, USA
Henri Jin	NASA-Ames, USA
Adrian Jackason	EPCC, UK
Wayne Joubert	ORNL, USA
Michael Klemm	Intel, Germany
Jeff Larkin	NVIDIA, USA
Seyong Lee	ORNL, USA
C. J. Newburn	NVIDIA, USA
Antonio J. Pena	BSC, Spain
William Sawyer	CSCS, Switzerland
Thomas Schwinge	MentorGraphics, Germany
Ray Sheppard	Indiana University, USA
Sameer Shende	University of Oregon, USA
Peter Steinbach	Scionics, Germany
Christian Terboven	RWTH Aachen University, Germany

Xiaonan Tian	NVIDIA/PGI, USA
Cheng Wang	Microsoft, USA
Michael Wolfe	PGI, USA

Publicity and Publication Chair

Sebastian Starke	HZDR, Germany

Held in conjunction with

SC17: The International Conference for High Performance Computing, Networking, Storage and Analysis
Denver, Colorado,
November 12–18, 2017

Contents