Lecture Notes in Computer Science 12655

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

Sridutt Bhalachandra · Sandra Wienke ·
Sunita Chandrasekaran ·
Guido Juckeland (Eds.)

Accelerator Programming Using Directives

7th International Workshop, WACCPD 2020
Virtual Event, November 20, 2020
Proceedings

Springer

Editors
Sridutt Bhalachandra
Lawrence Berkeley National Laboratory
Berkeley, CA, USA

Sunita Chandrasekaran
University of Delaware
Newark, DE, USA

Sandra Wienke
RWTH Aachen University
Aachen, Germany

Guido Juckeland
Helmholtz-Zentrum Dresden-Rossendorf
Dresden, Germany

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-030-74223-2 ISBN 978-3-030-74224-9 (eBook)
https://doi.org/10.1007/978-3-030-74224-9

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

Preface

The course of high-performance computing (HPC) system architecture is at crossroads. Heterogeneous systems like the current Summit and Sierra, and the up-and-coming systems such as Perlmutter, Frontier, and El Capitan, all have graphics processing units (GPUs). On the other hand, ARM-based systems such as GW4 Isambard and the future Isambard 2 and Fugaku are relatively homogeneous. Nonetheless, system performance relies heavily on acceleration through existing ever-improving vector units in the homogeneous systems or dedicated acceleration units in heterogeneous systems.

With increasing complexity to exploit the maximum available parallelism, the importance of sophisticated programming approaches that can handle performance, scalability, and portability is increasing. Programmers, especially, prefer to keep a single code base to help ease maintenance and avoid the need to debug multiple versions of the same code. In the literature, it has been shown that the abstraction can be raised at different levels - at the high level using directives and frameworks or at a relatively lower level by language modifications.

Software abstraction-based programming models such as OpenMP and OpenACC have been serving this purpose over the past several years and are likely to represent one path forward. These programming models address the 'X' component in a hybrid MPI+X programming approach by providing programmers high-level directives and delegating some burden to the compiler. Such programming paradigms have played a decisive role in establishing heterogeneous node architectures as a valid choice for a multitude of HPC workloads. In addition, frameworks like Kokkos and Raja, along with modifications to the language, are trying to help improve the performance as well as portability.

These proceedings contain the papers accepted for presentation at the 7th Workshop on Accelerator Programming using Directives (WACCPD 2020) held on November 13, 2020. 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 (https://www.waccpd.org).

Like in the previous years, the workshop highlighted improvements to the state of the art through the accepted papers and prompted discussion through keynotes 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 other subjects.

The WACCPD 2020 workshop received seven submissions out of which five were accepted to be presented at the workshop and published in these proceedings. The Program Committee of the workshop comprised 27 members spanning universities, national laboratories, and several industries. Each paper received a minimum of six

single-blind reviews. Similar to WACCPD 2019, we encouraged all authors to add the Artifact Description (AD) to their submissions and make their code and data publicly available (e.g. on GitHub, Zenodo, Code Ocean, etc.) in support of the reproducibility initiative. Of the five accepted papers, 40% had reproducibility information and these manuscripts are highlighted with an 'artifacts available' logo in this book.

The program co-chairs invited Prof. Mary Hall from the University of Utah to give a keynote address on "Achieving Performance Portability for Extreme Heterogeneity." Mary Hall is the Director of the School of Computing at the University of Utah. Her research focus brings together compiler optimizations and performance tuning targeting current and future high-performance architectures on real-world applications. Professor Hall is an IEEE Fellow, an ACM Distinguished Scientist, and a member of the Computing Research Association Board of Directors. She actively participates in mentoring and outreach programs to encourage the participation of groups underrepresented in computer science.

Nicholas Malaya from AMD gave an invited talk titled "Enabling Portable Directive-Based Programming at Exascale." Nicholas Malaya is a computational scientist at AMD Research, and is AMD's technical lead for the Frontier and El Capitan Centers of Excellence (COEs). These COEs are focused on close collaborations between AMD, DOE, and HPE to ensure application readiness, so that key workloads can run on the computers from Day-1 of machine deployment. Nick's research interests include Exascale Computing, CFD, Bayesian Inference, and Machine Learning.

Usually, the text of the preface focuses on factual content only. However, 2020 was (unfortunately) different – we feel that we cannot leave this unmentioned: the COVID-19 pandemic hit the world really hard in 2020. It has affected the daily lives of all of us and cost way too many lives. To limit the spread of the virus, the most important action has become social distancing. With that, big events were cancelled all over the world. However, in the HPC community, we have been rather lucky to be able to "easily" switch to solely digital and virtual conference formats. To this end, Supercomputing 2020 was held online, and so was WACCPD 2020 for the first time in its seven-year history. Thanks to all of you that contributed to this success! Hopefully, we will be able to meet in-person again next time. Stay tuned!

February 2021

Sridutt Bhalachandra
Sandra Wienke
Sunita Chandrasekaran
Guido Juckeland

Organization

Steering Committee

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

General Chairs

Sunita Chandrasekaran	University of Delaware, USA
Guido Juckeland	Helmholtz-Zentrum Dresden-Rossendorf, Germany

Program Chairs

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

Publicity Chair

Neelima Bayyapu	National Institute of Technology Karnataka, India

Web Chair

Shumpei Shiina	University of Tokyo, Japan

Program Committee

Daniel Abdi	National Oceanic and Atmospheric Administration, USA
James Beyer	Nvidia Corporation, USA
Maciej Cytowski	Pawsey Supercomputing Center, Australia
Christopher Daley	Lawrence Berkeley National Laboratory, USA
Joel Denny	Oak Ridge National Laboratory, USA
Johannes Doerfert	Argonne National Laboratory, USA
Millad Ghane	Samsung Semiconductor Inc., USA
Priyanka Ghosh	Washington State University, USA
Haoqiang Jin	NASA Ames Research Center, USA
Ronan Keryell	Xilinx, USA
Jeongnim Kim	Intel, USA
John Leidel	Tactical Computing Laboratories LLC, USA

Ron Lieberman	AMD, USA
Kelvin Li	IBM Corporation, Canada
Meifeng Lin	Brookhaven National Laboratory, USA
Chun-Yu Lin	National Center for High-Performance Computing, Taiwan
Zhao Liu	Tsinghua University, China; National Supercomputing Center, Wuxi, China
Stephen Lecler Olivier	Sandia National Laboratory, USA
Arpith Jacob	Google, USA
Thomas Schwinge	Mentor Graphics, Germany
Bharatkumar Sharma	Nvidia Corporation, Bengaluru, India
Ray Sheppard	Indiana University, USA
Gregory Stoner	Intel, USA
Cheng Wang	Microsoft Corporation, USA
Michael Wolfe	Nvidia Corporation, USA
Rengan Xu	Dell EMC, USA
Charlene Yang	Lawrence Berkeley National Laboratory, USA

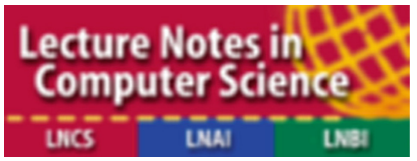

Contents