

Proceedings of the 23rd International Workshop on High-Level Parallel Programming Models and Supportive Environments

May 21, 2018, Vancouver, British Columbia, Canada

The 23rd HIPS workshop, held as a full-day meeting on Monday, May 21, 2018 at the IEEE International Parallel & Distributed Processing Symposium (IPDPS) 2018 conference in Vancouver, focuses on high-level programming of multicore/manycore systems, accelerators, compute clusters, and massively parallel machines. Like previous workshops in the series, which was established in 1996, this event serves as a forum for research in the areas of parallel applications, language design, compilers, runtime systems, and programming tools.

In our call for papers, we especially encouraged innovative approaches in the areas of emerging programming models for large-scale parallel systems, data-intensive applications, and memory-centric architectures. The program committee used a rigorous review process in which our they provided at least three reviews for each submission. We selected six high-quality submissions to include in these proceedings. The papers are presented in two sessions: 1) Tool Support for Parallel Programming Environments; and 2) Distributed Memory and Task-based Programming.

Our workshop starts with a keynote address by Dr. Christian Trott of Sandia National Laboratories. In his talk, Dr. Trott describes the challenges of introducing a new programming model into an established applications portfolio, drawing lessons from the Kokkos software that provides performance portable multi-threading and acceleration support to parallel applications.

The first session focuses on tool support for parallel application development and optimization. Mix, Herold, and Weber propose novel approaches to visualizing I/O performance events observed at multiple layers in the I/O stack. Next, Atzeni and Gopalakrishnan describe their approach for automatically

identifying data races in OpenMP programs and its implementation in the *SWORD* tool. Mehrabi, Giacaman, and Sinnen conclude the session by describing their annotation-based approach for specifying asynchronous graphical user interface operations and its implementation in the *@PT* parallel programming environment.

The second workshop session focuses on distributed memory and task-based programming models. Li, Chen, Gupta, and Xie begin by proposing an approach to avoid scaling problems when entering or exiting collective operations in applications that use the Message Passing Interface (MPI) for communication and synchronization. Van Veen and Jongmans then describe a domain-specific language for specifying synchronization and communication in a task-based programming model. To conclude the workshop, Whitlock, Kolla, Treichler, Pebay, and Bennett propose an approach for providing scalable collective operations for asynchronous many-task programming models such as *Legion*.

On behalf of the HIPS 2018 workshop committees, we thank the authors of all submitted papers for sharing their recent results, and we especially thank the authors of accepted submissions for the additional work required to prepare their papers for these proceedings. We also thank the program committee and all other reviewers for volunteering their time, talents, and expertise during the review process, and we thank our steering committee for their helpful advice. Finally, we thank the organizers of the 32nd IPDPS conference, and especially its workshops chair and vice chair, Erik Saule and Jaroslaw Zola for hosting our workshop.

Karl Fuerlinger and Philip C. Roth
HIPS 2018 Workshop Chairs

HIPS 2018 Committees

Workshop Chairs

Karl Fuerlinger, Ludwig-Maximilians-Universität (LMU) München, Munich, Germany
Philip C. Roth, Oak Ridge National Laboratory, Oak Ridge, TN, USA

Steering Committee

Rudolf Eigenmann, Purdue University, West Lafayette, IN, USA
Michael Gerndt, Technische Universität München, Munich, Germany
Frank Mueller, North Carolina State University, Raleigh, NC, USA
Craig Rasmussen, University of Oregon, Eugene, OR, USA
Martin Schulz, Technische Universität München, Munich, Germany

Program Committee

Richard Barrett, Sandia National Laboratories, USA
Guoyang Chen, Qualcomm Inc., USA
Lydia Duncan, Cray Inc., USA
Alessandro Fanfarillo, National Center for Atmospheric Research, USA
Karl Fuerlinger, LMU München, Germany
Markus Geimer, Jülich Supercomputing Center, Germany
Jeff Hammond, Intel, USA
Oscar Hernandez, Oak Ridge National Laboratory, USA
Kevin Huck, University of Oregon, USA
Hartmut Kaiser, Louisiana State University, USA
Andreas Knüpfer, Technische Universität Dresden, Germany
JP Lehr, Technische Universität Darmstadt, Germany
Graham Lopez, Oak Ridge National Laboratory, USA
Stefano Markidis, KTH Stockholm, Sweden
Hitoshi Murai, RIKEN, Japan
CJ Newburn, NVIDIA, USA
Philip C. Roth, Oak Ridge National Laboratory, USA
Massimo Torquati, University of Pisa, Italy
Josef Weidendorfer, Technische Universität München, Germany