
Preface

Sunita Chandrasekaran*

Department of Computer and Information Sciences,
University of Delaware,
Newark 19702, USA
Email: schandra@udel.edu
*Corresponding author

M. Graham Lopez

Computer Science and Mathematics Division,
Oak Ridge National Laboratory,
Oak Ridge, TN 37830, USA
Email: lopezmg@ornl.gov

Biographical notes: Sunita Chandrasekaran is an Assistant Professor at the Computer and Information Sciences Department, University of Delaware. Her research interests include exploring the suitability of high-level programming models and runtime systems for HPC, and migrating scientific applications to heterogeneous computing systems. She was a Postdoctoral Fellow at the UH and holds a PhD from the NTU, Singapore. She received the 2016 IEEE-CS TCHPC Award for Excellence for Early Career Researchers in HPC. She has served on the program committees of conferences including SC, ISC, and ICPP, and co-chaired parallel programming workshops co-located with SC, ISC, IPDPS and SIAM.

M. Graham Lopez is a researcher in the Computer Science and Mathematics Division from the Oak Ridge National Laboratory. He works on programming environments preparation with the application readiness teams for the DOE CORAL and exascale computing projects. He has published research in the areas of computational materials science, application acceleration and benchmarking on heterogeneous systems, low-level communication APIs, and programming models. He earned his MS in Computer Science and PhD in Physics from the Wake Forest University. Prior to joining ORNL, he was a Research Scientist at the Georgia Institute of Technology where he worked on application and numerical algorithm optimisations for accelerators.

This special issue has five invited manuscripts from authors who presented their work at the Second International Workshop on Performance Portable Programming Models for Accelerators (P³MA) co-located with ISC 2017, held at Frankfurt, Germany on June 22, 2017.

The scope of the special issue includes defining language extensions for programming models, their implementations, and experiences with their deployment in HPC applications on multiple architectures, performance modelling and evaluation tools, asynchronous task and event-driven execution/scheduling. The workshop received eight submissions in total and all submitted manuscripts were peer reviewed. We invited five top papers to the special issue. The authors were asked to extend their paper by at the least 30% for the special issue.

These five manuscripts were peer-reviewed and the review process for the special issue was not double blind, i.e., authors were known to reviewers. Submissions were judged on correctness, originality, technical strength, and significance, quality of presentation, and interest and relevance to the conference scope. All five manuscripts

were accepted to the special issue after two rounds of review.

Special issue program committee

- Samuel Thibault, INRIA, University of Bordeaux, France
- James Beyer, NVIDIA, USA
- Wei Ding, AMD, USA
- Saber Feki, King Abdullah University, Saudi Arabia
- Robert Henschel, Indiana University, USA
- Michael Klemm, Intel, USA
- Eric Stotzer, Texas Instruments, USA
- Amit Amritkar, University of Houston, USA
- Guido Juckeland, HZDR, Germany
- Will Sawyer, ETH, Zurich

- Sameer Shende, University of Oregon, USA
- Costas Bekas, IBM, Zurich
- Toni Collis, University of Edinburgh, Scotland
- Adrian Jackson, University of Edinburgh, Scotland
- Henri Jin, NASA, USA
- Andreas Knuepfer, TU Dresden, Germany
- Steven Olivier, Sandia National Laboratory, USA
- Suraj Prabhakaran, TU Darmstadt, Germany
- Bora Ucar, ENS De Lyon, France
- Manisha Gajbe, Intel, USA
- Daniel Tian, PGI, USA