

Preface

OpenSHMEM 2015 was the second event in the OpenSHMEM and Related Technologies workshop series. The workshop was organized by Oak Ridge National Laboratory and held in Annapolis, Maryland, USA, and it was and sponsored by ORNL, DoD, Mellanox, NVIDIA, Intel, and SGI.

The OpenSHMEM Workshop is the premier venue for exhibiting and presenting research on partitioned global address space (PGAS), particularly as it relates to OpenSHMEM. The workshop was attended by participants from across academia, industry, and private and federal research organizations. The topics for the workshop included extensions to the OpenSHMEM API, implementation of the API for current and emerging architectures, tools to debug and profile OpenSHMEM programs, experience porting applications to the OpenSHMEM programming model, and changes to the OpenSHMEM specification to address the needs of programming exascale systems.

The response to the OpenSHMEM Workshop call for papers was very encouraging. The Program Committee members reviewed the papers with a very short turnaround. Despite the short turnaround, each paper was reviewed by more than three reviewers, and 12 papers were selected to be presented at the workshop.

This proceedings volume is a collection of papers presented at the second workshop, during August 4–6, 2015. The technical papers provided a multitude of ideas for extending the OpenSHMEM specification and making it efficient for current and next-generation systems. This included non-blocking APIs, teams (similar to “communicators” in the message-passing paradigm), extended capabilities for collective operations, and considerations for additional memory architectures such as accelerators. OpenSHMEM is now being used to explore new parallel algorithms for applications, and those experiences from the developers were also an integral part of the technical program at the OpenSHMEM Workshop this year.

Besides contributed papers, the technical program consisted of tutorials, invited talks, and specification discussion. The tutorials were presented by speakers from NVIDIA, Mellanox, Oak Ridge National Laboratory, Allinea, ParaTools, University of Houston, and TU Dresden. The invited talk from NVIDIA discussed the role of OpenSHMEM in programming GPUs, the talk from Mellanox discussed the effectiveness of using InfiniBand hardware for OpenSHMEM, and the talk from Intel provided an overview of OpenSHMEM activities in Intel.

The third day of the workshop was focused on developing the OpenSHMEM specification. This was a very exciting year in the OpenSHMEM community since the first OpenSHMEM Workshop in March 2014. Two updates to the OpenSHMEM specification have been introduced in the meantime: version 1.1 was released during summer 2014 and version 1.2 in spring 2015. The discussion in this workshop focused on features for the upcoming 1.3 version and future versions.

The general and program chairs would like to thank everyone who contributed to the organization of the workshop. Particularly, we would like to thank the authors, Program Committee members, reviewers, session chairs, participants, and sponsors. We are grateful for the excellent support we received from our ORNL administrative staff and Daniel Pack, who maintained our workshop website.

August 2015

Manjunath Gorenthla Venkata
Pavel Shamis
Neena Imam
M. Graham Lopez

OpenSHMEM and Related Technologies. Experiences,
Implementations, and Technologies
Second Workshop, OpenSHMEM 2015, Annapolis, MD,
USA, August 4-6, 2015. Revised Selected Papers
Gorentla Venkata, M.; Shamis, P.; Imam, N.; Lopez, M.G.
(Eds.)
2015, X, 199 p. 84 illus. in color., Softcover
ISBN: 978-3-319-26427-1