# Lecture Notes in Computer Science 10679

*Commenced Publication in 1973*
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

## Editorial Board

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

Manjunath Gorentla Venkata · Neena Imam
Swaroop Pophale (Eds.)

# OpenSHMEM and Related Technologies

## Big Compute and Big Data Convergence

4th Workshop, OpenSHMEM 2017
Annapolis, MD, USA, August 7–9, 2017
Revised Selected Papers

Springer

*Editors*
Manjunath Gorentla Venkata
Oak Ridge National Laboratory
Oak Ridge, TN
USA

Swaroop Pophale
Oak Ridge National Laboratory
Oak Ridge, TN
USA

Neena Imam
Oak Ridge National Laboratory
Oak Ridge, TN
USA

ISSN 0302-9743 ISSN 1611-3349 (electronic)
Lecture Notes in Computer Science
ISBN 978-3-319-73813-0 ISBN 978-3-319-73814-7 (eBook)
https://doi.org/10.1007/978-3-319-73814-7

Library of Congress Control Number: 2017963768

LNCS Sublibrary: SL2 – Programming and Software Engineering



Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

# Preface

OpenSHMEM is a Partitioned Global Address Space (PGAS) library specification. The main abstractions of the programming model are execution contexts called processing elements (PEs) and symmetric memory objects. The programming model also provides explicit mechanisms to access and transfer data between the symmetric memory objects of different PEs. The key factors that make OpenSHMEM an excellent choice for parallel, communicating HPC applications is its simple application programming interface (API), support for remote direct memory access (RDMA), and constant innovation in the library API to keep abreast with the current scientific and hardware changes. Over the past few years there is a growing momentum behind the development and usage of the OpenSHMEM programming model.

The OpenSHMEM Workshop is the premier venue for presenting new and innovative PGAS research in the context of OpenSHMEM. OpenSHMEM 2017, held in Annapolis, Maryland, was the fourth event in the OpenSHMEM and Related Technologies workshop series. The workshop was organized by Oak Ridge National Laboratory and sponsored by ORNL, DoD, Cray, Nvidia, Mellanox, ARM, and HPE. The workshop was attended by participants from across academia, industry, and private and federal research organizations.

This year, the workshop focused on "OpenSHMEM and the Big Compute and Big Data Convergence." The workshop included two days of technical presentations followed by one day dedicated to the OpenSHMEM Specification discussions and development. The technical segment commenced with a keynote from Dr. William Carlson. Apart from being a member of the research staff at the IDA Center for Computing Sciences since 1990, he also leads the UPC language effort. The title of his talk was "Shared Memory HPC Programming: Past, Present, and Future?".

The paper session discussed a variety of concepts, including extending the OpenSHMEM API for future architectures, new applications using OpenSHMEM, evaluation and implementation of OpenSHMEM for new architectures, novel use of OpenSHMEM for the heterogeneous environments, and new development in the tools eco-system for OpenSHMEM. All papers submitted to the workshop were peer-reviewed by the Program Committee (PC) which included members from universities, industry, and research labs. Despite the short turnaround, each paper was reviewed by at least three reviewers. In all, 11 full papers were selected to be presented at the workshop.

This proceedings volume is a collection of papers presented at the workshop. The technical papers provided a variety of ideas for extending the OpenSHMEM specification and making it efficient for current and next-generation systems. This includes new research for communication contexts in OpenSHMEM, different optimizations for OpenSHMEM on shared memory machines, exploring the implementation of OpenSHMEM and its memory model on Intel's KNL architecture, and implementing new applications and benchmarks with OpenSHMEM.

The third day of the workshop was focused on developing the OpenSHMEM specification. This year, 2017, like the year before, has been a very exciting year for the OpenSHMEM committee. Thanks to the active participation at the workshop, the committee is in the process of ratifying OpenSHMEM Specification 1.4. The OpenSHMEM meeting at the workshop is an annual face-to-face OpenSHMEM committee meeting making it an important and impactful venue.

The general and technical program Chairs would like to thank everyone who contributed to the organization of the workshop. Particularly, we would to thank all authors, PC 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 helped maintain and update our workshop website.

December 2017

Neena Imam
Manjunath Gorentla Venkata
Swaroop Pophale

# Organization

## General Co-chairs

| | |
|---|---|
| Neena Imam | Oak Ridge National Laboratory, USA |
| Manjunath Gorentla Venkata | Oak Ridge National Laboratory, USA |
| Nick Park | Department of Defense, USA |

## Technical Program Co-chairs

| | |
|---|---|
| Manjunath Gorentla Venkata | Oak Ridge National Laboratory, USA |
| Swaroop Pophale | Oak Ridge National Laboratory, USA |

## Technical Program Committee

| | |
|---|---|
| Ferrol Aderholdt | Oak Ridge National Laboratory, USA |
| Matthew Baker | Oak Ridge National Laboratory, USA |
| Pavan Balaji | Argonne National Laboratory, USA |
| Swen Boehm | Oak Ridge National Laboratory, USA |
| Bob Cernohous | Cray Inc., USA |
| Zheng Cui | VMWare, USA |
| Tony Curtis | Stony Brook University, USA |
| James Dinan | Intel Corporation, USA |
| Jeff Hammond | Intel Labs, USA |
| Bryant Lam | Department of Defense, USA |
| Arthur Maccabe | Oak Ridge National Laboratory, USA |
| Dhabaleswar (DK) Panda | Ohio State University, USA |
| Nick Park | Department of Defense, USA |
| Stephen Poole | OSSS, USA |
| Sreeram Potluri | NVIDIA, USA |
| Michael Raymond | SGI, USA |
| Gilad Shainer | Mellanox Technologies, USA |
| Pavel Shamis | ARM, USA |
| Sameer Shende | University of Oregon, USA |
| Min Si | Argonne National Laboratory, USA |
| Weikuan Yu | Florida State University, USA |

## Sponsors

### Diamond Sponsors


### Silver Sponsors



### Bronze Sponsors

# Contents