

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, Lancaster, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Zurich, Switzerland

John C. Mitchell

Stanford University, Stanford, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Dortmund, Germany

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbrücken, Germany

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

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

OpenSHMEM and Related Technologies

Enhancing OpenSHMEM for Hybrid Environments

Third Workshop, OpenSHMEM 2016
Baltimore, MD, USA, August 2–4, 2016
Revised Selected Papers

Editors

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

Neena Imam
Oak Ridge National Laboratory
Oak Ridge, TN
USA

Swaroop Pophale
Oak Ridge National Laboratory
Oak Ridge, TN
USA

Tiffany M. Mintz
Oak Ridge National Laboratory
Oak Ridge, TN
USA

ISSN 0302-9743

ISSN 1611-3349 (electronic)

Lecture Notes in Computer Science

ISBN 978-3-319-50994-5

ISBN 978-3-319-50995-2 (eBook)

DOI 10.1007/978-3-319-50995-2

Library of Congress Control Number: 2016960195

LNCS Sublibrary: SL2 – Programming and Software Engineering

© Springer International Publishing AG 2016

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

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

The OpenSHMEM Workshop is the premier venue for presenting Partitioned Global Address Space (PGAS) research, particularly as it relates to OpenSHMEM. OpenSHMEM 2016 was the third event in the OpenSHMEM and Related Technologies workshop series. The workshop was organized by Oak Ridge National Laboratory and was held in Baltimore, Maryland, USA, and it was sponsored by ORNL, DoD, Intel, Mellanox, Cray, and SGI. The workshop was attended by participants from across academia, industry, and private and federal research organizations.

This year, the workshop focused on the role of OpenSHMEM in heterogeneous and hybrid environments. The two keynotes of the workshop included Steve Oberlin's (NVIDIA CTO) talk on the role of OpenSHMEM in future GPU-based extreme scale systems, and James Sexton's (IBM Fellow) talk on the usability of OpenSHMEM in data-centric architectures. Besides the keynote, the workshop included paper and vendor sessions as well as the OpenSHMEM committee meeting. The vendor session included talks from Intel, Cray, Mellanox, Allinea, and Paratools.

The paper session discussed a variety of concepts, including extending the OpenSHMEM API for future architectures, optimizing OpenSHMEM for current architectures, and enhancements to OpenSHMEM for the heterogeneous environments. All papers submitted to the workshop were peer-reviewed by the Technical Program Committee, which included members from universities, industry, and research labs. The Technical Program Committee members reviewed the papers with a very short turnaround time. Despite the short turnaround, each paper was reviewed by more than three reviewers, and in the end 14 full papers and 3 short 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 included active messages, non-blocking APIs, fault-tolerance capabilities, exploring implementation of OpenSHMEM using communication layers such as OFI and UCX, and implementing OpenSHMEM for heterogeneous architectures. The OpenSHMEM library is being explored as a high-performing communication layer for PGAS languages and Big Data frameworks, and those experiences from the developers were discussed at the OpenSHMEM workshop this year.

The third day of the OpenSHMEM workshop was focused on developing the OpenSHMEM specification. This year, like the year before, has been a very exciting year for the OpenSHMEM committee. The committee released OpenSHMEM version 1.3 in February 2016, and also built a very active community that participates in the development of the specification. The OpenSHMEM meeting at the workshop is an annual and only face-to-face OpenSHMEM committee meeting. This was one of the most important meetings, as a set of rules and procedures were defined to be adopted by the OpenSHMEM committee. This included operation procedures for the

OpenSHMEM committee and participants, and the formalization of the process of development and ratification of the specification.

The general and program chairs would like to thank everyone who contributed to the organization of the workshop. In particular, 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.

November 2016

Neena Imam
Manjunath Gorentla Venkata
Swaroop Pophale
Tiffany Mintz

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

Tiffany M. Mintz	Oak Ridge National Laboratory, USA
Swaroop Pophale	Oak Ridge National Laboratory, USA

Technical Program Committee

Ferrol Aderholdt	Oak Ridge National Laboratory, USA
Aurelien Bouteiller	University of Tennessee - Knoxville, USA
Tony Curtis	Stony Brook University, USA
James Dinan	Intel Corporation, USA
Richard Graham	Mellanox Technologies, USA
Khaled Hamidouche	Ohio State University, USA
Jeff Hammond	Intel Labs, USA
Dounia Khaldi	Stony Brook University, USA
David Knaak	Cray Inc., USA
Andreas Knuepfer	ZIH, TU Dresden, Germany
Gregory Koenig	KPMG, USA
Bryant Lam	Department of Defense, USA
Arthur Maccabe	Oak Ridge National Laboratory, USA
Dhabaleswar (DK) Panda	Ohio State University, USA
Stephen Poole	OSSS, USA
Sreeram Potluri	NVIDIA, USA
Sarah Powers	Oak Ridge National Laboratory, USA
Michael Raymond	SGI, USA
Gilad Shainer	Mellanox Technologies, USA
Pavel Shamis	ARM, USA
Sameer Shende	University of Oregon, USA
Tom St. John	Intel Corporation, USA
Weikuan Yu	Florida State University, USA

Sponsors

Diamond Sponsors



Gold Sponsors



Silver Sponsors



Contents

OpenSHMEM Extensions

Integrating Asynchronous Task Parallelism with OpenSHMEM	3
<i>Max Grossman, Vivek Kumar, Zoran Budimlić, and Vivek Sarkar</i>	
Evaluating OpenSHMEM Explicit Remote Memory Access Operations and Merged Requests	18
<i>Swen Boehm, Swaroop Pophale, and Manjunath Gorentla Venkata</i>	
Increasing Computational Asynchrony in OpenSHMEM with Active Messages	35
<i>Siddhartha Jana, Tony Curtis, Dounia Khaldi, and Barbara Chapman</i>	
System-Level Transparent Checkpointing for OpenSHMEM	52
<i>Rohan Garg, Jérôme Vienne, and Gene Cooperman</i>	
Surviving Errors with OpenSHMEM	66
<i>Aurelien Bouteiller, George Bosilca, and Manjunath Gorentla Venkata</i>	
On Synchronisation and Memory Reuse in OpenSHMEM	82
<i>Aaron Welch and Manjunath Gorentla Venkata</i>	

OpenSHMEM Implementation and Use Cases

Design and Implementation of OpenSHMEM Using OFI on the Aries Interconnect	97
<i>Kayla Seager, Sung-Eun Choi, James Dinan, Howard Pritchard, and Sayantan Sur</i>	
OpenSHMEM-UCX: Evaluation of UCX for Implementing OpenSHMEM Programming Model	114
<i>Matthew Baker, Ferrol Aderholdt, Manjunath Gorentla Venkata, and Pavel Shamis</i>	
SHMemCache: Enabling Memcached on the OpenSHMEM Global Address Model	131
<i>Huansong Fu, Kunal SinghaRoy, Manjunath Gorentla Venkata, Yue Zhu, and Weikuan Yu</i>	
An OpenSHMEM Implementation for the Adapteva Epiphany Coprocessor . . .	146
<i>James Ross and David Richie</i>	

Hybrid Programming and Benchmarking with OpenSHMEM

An Evaluation of Thread-Safe and Contexts-Domains Features in Cray SHMEM	163
<i>Naveen Namashivayam, David Knaak, Bob Cernohous, Nick Radcliffe, and Mark Pagel</i>	
OpenCL + OpenSHMEM Hybrid Programming Model for the Adapteva Epiphany Architecture	181
<i>David A. Richie and James A. Ross</i>	
OpenSHMEM Implementation of HPCG Benchmark	193
<i>Eduardo D'Azevedo, Sarah Powers, and Neena Imam</i>	
Using Hybrid Model OpenSHMEM + CUDA to Implement the SHOC Benchmark Suite	204
<i>Megan Grodowitz, Eduardo D'Azevedo, Sarah Powers, and Neena Imam</i>	

OpenSHMEM Tools

Profiling Production OpenSHMEM Applications.	219
<i>John C. Linford, Samuel Khuvis, Sameer Shende, Allen Malony, Neena Imam, and Manjunath Gorentla Venkata</i>	

Short Papers

SHMEM-MT: A Benchmark Suite for Assessing Multi-threaded SHMEM Performance	227
<i>Hans Weeks, Matthew G.F. Dosanjh, Patrick G. Bridges, and Ryan E. Grant</i>	
Investigating Data Motion Power Trends to Enable Power-Efficient OpenSHMEM Implementations.	232
<i>Tiffany M. Mintz, Eduardo D'Azevedo, Manjunath Gorentla Venkata, and Chung-Hsing Hsu</i>	

Author Index	239
-------------------------------	-----