

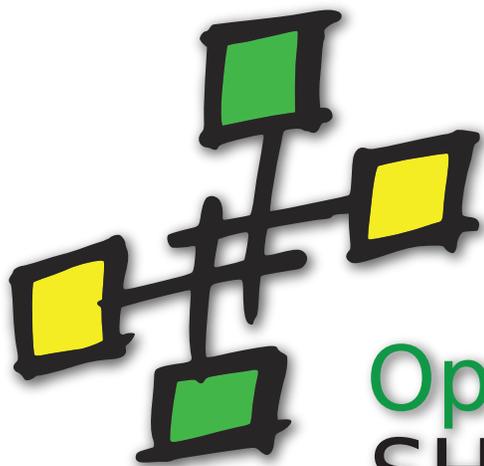
OpenSHMEM 2015

The Future of OpenSHMEM and Related Technologies

Workshop Program &
Abstracts

August 4-6, 2015

Annapolis, Maryland



Open
SHMEM

Register at

<http://www.csm.ornl.gov/workshops/openshmem2015>

OpenSHMEM

General Co-Chairs

Neena Imam

Oak Ridge National Laboratory

Pavel Shamis

Oak Ridge National Laboratory

Manjunath Gorentla Venkata

Oak Ridge National Laboratory

Steering Committee

Steve Poole

OSSS

Janice Elliott

DoD

Neena Imam

Oak Ridge National
Laboratory

Nicholas Park

DoD

William W. Carlson

IDA

Oscar Hernandez

Oak Ridge National
Laboratory

Barney Maccabe

Oak Ridge National
Laboratory

Lauren Smith

DoD

Tutorial Chair

Nicholas Park

DoD

Technical Program Chair

Graham Lopez

Oak Ridge National Laboratory

Technical Program Committee

George Bosilca

University of Tennessee

Gary Grider

Los Alamos National
Laboratory

Gregory Koenig

Oak Ridge National
Laboratory

Swaroop Pophale

Mellanox Technologies

Barbara Chapman

University of Houston

Jeff Hammond

Intel

Jeff Kuehn

Los Alamos National
Laboratory

Sreeram Potluri

NVIDIA

Tony Curtis

University of Houston

Chung-Hsing Hsu

Oak Ridge National
Laboratory

Barney Maccabe

Oak Ridge National
Laboratory

Michael Raymond

SGI

James Dinan

Intel

Dounia Khaldi

University of Houston

Nicholas Park

DoD

Pavel Shamis

Oak Ridge National
Laboratory

**Manjunath
Gorentla Venkata**

Oak Ridge National
Laboratory

David Knaak

Cray Inc.

Duncan Poole

NVIDIA

Sameer Shende

University of Oregon

Richard Graham

Mellanox Technologies

Andreas Knuepfer

TU-Dresden

Steve Poole

OSSS

Workshop Organizers

Corliss Thompson

Oak Ridge National Laboratory

Daniel Pack

Oak Ridge National Laboratory

Eric Mitchell

University of Tennessee

KEYNOTE

Presented By

Dr. Thomas Sterling

Center for Research in Extreme-Scale Technologies

School of Informatics and Computing

Indiana University

August 5, 8:00 AM - 9:00 AM



Dr. Thomas Sterling holds the position of Professor of Informatics and Computing at the Indiana University (IU) School of Informatics and Computing as well as serves as Chief Scientist and Executive Associate Director of the Center for Research in Extreme Scale Technologies (CREST). Since receiving his Ph.D from MIT in 1984 as a Hertz Fellow, Dr. Sterling has engaged in applied research in fields associated with parallel computing system structures, semantics, and operation in industry, government labs, and academia. Dr. Sterling is best known as the "father of Beowulf" for his pioneering research in commodity/Linux cluster computing. He was awarded the Gordon Bell Prize in 1997 with his collaborators for this work. He was the PI of the HTMT Project sponsored by NSF, DARPA, NSA, and NASA to explore advanced technologies and their implications for high-end system architectures. Other research projects included the DARPA DIVA PIM architecture project with USC-ISI, the Cray Cascade Petaflops architecture project sponsored by the DARPA HPCS Program, and the Gilgamesh high-density computing project at NASA JPL. Thomas Sterling is currently engaged in research associated with the innovative ParalleX execution model for extreme scale computing to establish the foundation principles to guide the co-design for the development of future generation Exascale computing systems. ParalleX is currently the conceptual centerpiece of the XPRESS project as part of the DOE X-stack program and has been demonstrated in proof-of-concept in the HPX-5 runtime system software. Dr. Sterling is the co-author of six books and holds six patents. He was the recipient of the 2013 Vanguard Award and is a Fellow of the AAAS.

Dr. Thomas Sterling will present:

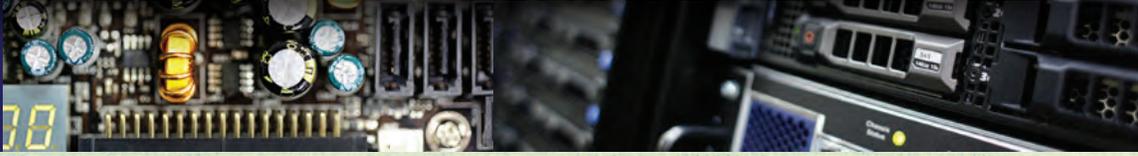
Eight Unquestioned Assumptions Blocking SHMEM Exascale Computing

The extraordinary momentum of Moore's Law has advanced HPC to the Petaflops performance regime even as this exponential progress of the enabling technologies is asymptotically flat-lining near the nanometer threshold. However HPC system architecture and programming models have struggled to keep up with increasingly complicated heterogeneous structures and multi-layered programming methods imposed on the user community. While SHMEM models convey unifying principles to return to efficient and scalable computing, a number of underlying assumptions that go unquestioned continue to force likely first generation exascale HPC to an increasingly limited form. In part these include restrictions of the commercial market, legacy codes, irrelevant benchmarking, and a culture of evolutionary incrementalism. Among the consequences of this lemming like community-wide approach are reduced generality, lack of performance portability, poorer efficiency, and degraded user productivity. This presentation will make explicit eight unquestioned assumptions permeating the conventional HPC trajectory and describe alternative advances based on the ParalleX execution model that will make possible the opportunities of future SHMEM exascale computing.

Day 1 Tutorials

Tuesday, August 4

OpenSHMEM 2015



7:00 AM Breakfast
(Provided)

8:00 AM **Tutorials Track 1:**

Introduction

Presented by *Pavel Shamis*

Latest in OpenSHMEM: Specification, API, and Programming

Presented by *Pavel Shamis, Manjunath Gorentla Venkata, Swaroop Pophale, Dounia Khaldi*

OpenSHMEM is a PGAS library interface specification. The model combines a global view of memory accessible to the entire program with local affinity of data storage. Various vendors have developed implementations of "SHMEM" over the last 30 years since it was first introduced by Cray. However, these diverged from each other as new features were added. OpenSHMEM is a standardization of SHMEM, designed to evolve the library and tool ecosystem to accommodate the needs of 21st century application development on increasingly larger computer systems.

This tutorial will cover both introductory and some intermediate level concepts. The introductory part will cover the core features of the OpenSHMEM API with focus on new and modified features for 1.2, including the interfaces, semantics, and programming model. The intermediate part of the tutorial will provide a hands-on experience with writing, compiling, and running a simple OpenSHMEM program, and then discusses the OpenSHMEM semantics and what it lacks for emerging architectures. Using an open source OpenSHMEM implementation, the last part of the tutorial will delve into the details of an OpenSHMEM implementation, with a brief overview of each layer.

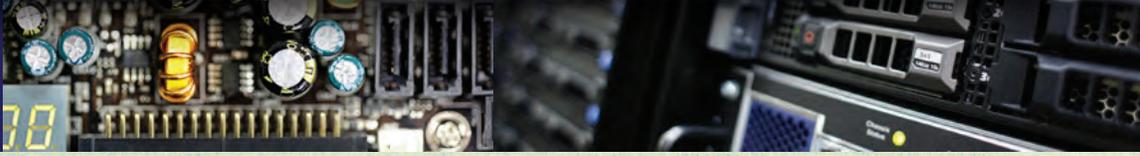
Prerequisites: The tutorial will be structured so as to be attractive to both programmers who are unfamiliar with the details of OpenSHMEM and PGAS environments and more seasoned programmers and scientists who want to acquaint themselves with the state of the art of such a technology. We expect that the attendees have familiarity with C, C++ or Fortran as those are the languages that OpenSHMEM targets, but note that we use C examples in the tutorial.



Note: breaks will be taken at the discretion of the instructor

Agenda

Tutorials



11:00 AM UCX - Communication Framework for Next Generation Programming Models

Presented by *Yossi Itigin, Mellanox*

Unified Communication X (UCX) is a set of network APIs and their implementations for high throughput computing. UCX comes from the combined effort of national laboratories, industry, and academia to design and implement a high-performing and highly-scalable network stack for next generation applications and systems. UCX design provides the ability to tailor its APIs and network functionality to suit a wide variety of application domains. We envision these APIs to satisfy the networking needs of many programming models such as Message Passing Interface (MPI), OpenSHMEM, Partitioned Global Address Space (PGAS) languages, task-based paradigms and I/O bound applications. UCX is an open source, BSD licensed software hosted on GitHub.

Prerequisites - none.

8:00 AM Tutorials Track 2:

Introduction

Presented by *Manjunath Gorentla Venkata*

Developing, Debugging and Profiling OpenSHMEM Applications

Presented by *David Lecomber, Alinea Software*

This tutorial session will present an overview of the Alinea tools for debugging and profiling. We begin with an introduction to Alinea DDT for OpenSHMEM debugging, in which we will cover how to use the debugger for efficiently solving OpenSHMEM application bugs, from small to large scale parallelism. Finally we consider the performance of OpenSHMEM applications and how to use Alinea MAP to profile performance and identify potential optimizations in areas such as vectorization, memory bandwidth, I/O and synchronization.

Prerequisites - none.

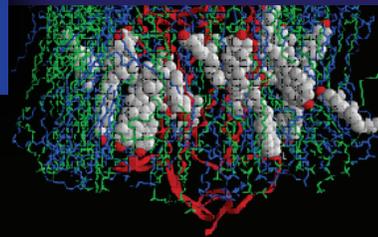
12:00 PM Lunch



Day 1 Tutorials

Tuesday, August 4

OpenSHMEM 2015



1:00 PM

Tutorials Track 1:

GP-GPU Programming with CUDA

Presented by *Larry Brown, NVIDIA*

This workshop will provide a tour of general purpose GPU (GP-GPU) programming with CUDA. The material will start at an introductory level and progress into intermediate to advanced topics. We will cover the fundamental concepts of heterogeneous computing using GPU accelerators. Topics will include an Introduction to CUDA, NVIDIA Developer Tools and GPU Computing Libraries, Optimization, Streams and Multi-GPU Programming.

Prerequisites - The tutorial will assume familiarity with traditional single threaded programming (as in standard C/Fortran/Java/Python).

Tutorials Track 2:

Performance Evaluation of OpenSHMEM Applications Using TAU

Presented by *Sameer Shende, University of Oregon*

The TAU Performance System is a powerful and highly versatile profiling and tracing tool ecosystem for performance analysis of parallel programs at all scales. Developed over the last two decades, TAU has evolved with each new generation of HPC systems and presently scales efficiently to hundreds of thousands of cores on the largest machines in the world. This tutorial will focus on performance data collection, analysis, and performance optimization of OpenSHMEM applications. The tutorial will introduce profiling and debugging support in TAU, cover memory usage, POSIX I/O, and support for various runtime systems (including OpenSHMEM, MPI, pthread, CUDA, OpenACC, and OpenCL). TAU's support for hardware performance counters and recent support for power and energy profiling will be demonstrated. TAU can generate traces in the OTF2 format for Vampir using the Score-P measurement library. The tutorial will also demonstrate TAU's 3D profile browser, ParaProf, and TAUdb, a data management framework used by TAU's PerfExplorer tool for cross-experiment analysis such as scalability studies. The participants are encouraged to download the OVA lite image from the HPCLinux distribution [<http://www.hpclinux.com>] on their laptops prior to the tutorial.

Prerequisites - Familiarity with building and running parallel programs on Linux.



Note: breaks will be taken at the discretion of the instructor

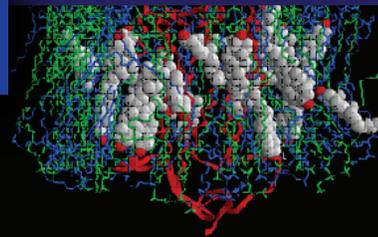
Agenda

Tutorials

Day 1 Tutorials

Tuesday, August 4

OpenSHMEM 2015



3:00 PM Break

3:30 PM Tutorials Track 1:

GP-GPU Programming with CUDA (continued)

Presented by *Larry Brown, NVIDIA*

This workshop will provide a tour of general purpose GPU (GP-GPU) programming with CUDA. The material will start at an introductory level and progress into intermediate to advanced topics. We will cover the fundamental concepts of heterogeneous computing using GPU accelerators. Topics will include an Introduction to CUDA, NVIDIA Developer Tools and GPU Computing Libraries, Optimization, Streams and Multi-GPU Programming.

Prerequisites - The tutorial will assume familiarity with traditional single threaded programming (as in standard C/Fortran/Java/Python).

Tutorials Track 2:

Parallel Performance Analysis for OpenSHMEM with Score-P and Vampir

Presented by *Joseph Schuchart, TU-Dresden*

The tutorial will present an introduction to the performance analysis tools Score-P and Vampir. Score-P is for instrumenting the parallel application code and for recording event traces during a measurement run of the target application on the target HPC machine. Vampir is for the visual and interactive analysis of the generated event traces after the measurement run. The tutorial shows the first steps with the tools, the most important parameters and switches, and some first-hand experiments how to identify typical performance flaws.

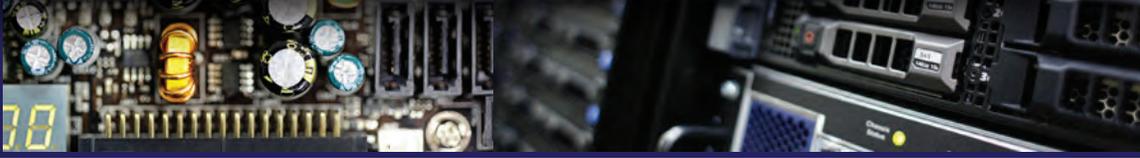
Prerequisites - A basic understanding of parallel processing and OpenSHMEM programming would be required to just follow the presentations. Attendees should have a laptop and access to Titan to follow the step by step examples.



Day 2 Technical Presentations

Wednesday, August 5

OpenSHMEM 2015



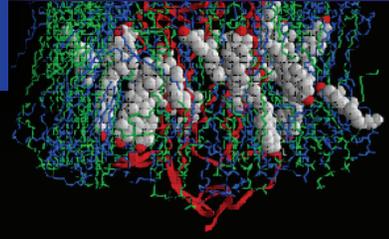
- 7:00 AM** Breakfast
(Provided)
- 8:00 AM** **Introduction**
Presented by *Neena Imam and Manjunath Gorentla Venkata*
- 8:15 AM** **Keynote:** Eight Unquestioned Assumptions Blocking SHMEM Exascale Computing
Presented by *Dr. Thomas Sterling, Indiana University*
- 9:05 AM** Dynamic Analysis to Support Program Development with the Textually Aligned Property for OpenSHMEM Collectives
Authored by *Andreas Knüpfer, Tobias Hilbrich, Joachim Protze and Joseph Schuchart*
- 9:30AM** Check-pointing Approach for Fault Tolerance in OpenSHMEM
Authored by *Pengfei Hao, Swaroop Pophale, Pavel Shamis, Anthony Curtis and Barbara Chapman*
- 9:55 AM** From MPI to OpenSHMEM: Porting LAMMPS
Authored by *Chunyan Tang, Aurelien Bouteiller, Thomas Herault, George Bosilca and Manjunath Gorentla Venkata*
- 10:20 AM** Break



Day 2 Technical Presentations

Wednesday, August 5

OpenSHMEM 2015



10:30 AM **Invited Talk:** Intel's Multifaceted PGAS Activities and Community Engagements

Presented by *Ulf Hanebutte, Intel*

11:20 AM Extending the Strided Communication Interface in OpenSHMEM

Authored by *Naveen Namashivayam, Dounia Khaldi, Deepak Eachempati and Barbara Chapman*

11:45 AM Graph 500 in OpenSHMEM

Authored by *Eduardo D'Azevedo and Neena Imam*

12:05 PM Lunch

1:00 PM **Invited Talk:** OpenSHMEM - The InfiniBand Advantage

Presented by *Rich Graham, Mellanox*

1:50 PM Scalable Out-of-core OpenSHMEM Library for HPC

Authored by *Antonio Gómez-Iglesias, Jerome Vienne, Khaled Hamidouche, William Barth and Dhableswar Panda*

2:15 PM Proposing OpenSHMEM Extensions Towards a Future for Hybrid Programming and Heterogeneous Computing

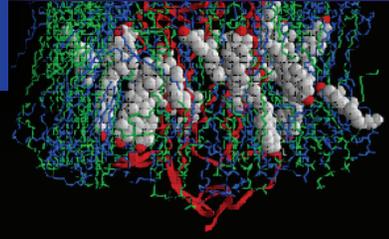
Authored by *David Knaak and Naveen Namashivayam*



Day 2 Technical Presentations

Wednesday, August 5

OpenSHMEM 2015



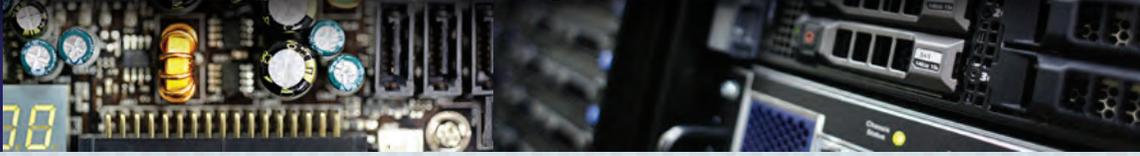
- 2:40 PM** A Case for Non-Blocking Collectives in OpenSHMEM: Design, Implementation, and Performance Evaluation using MVAPICH2-X
Authored by *Ammar Ahmad Awan, Khaled Hamidouche, Ching-Hsiang Chu and Dhableswar Panda*
- 3:00 PM** Break
- 3:30 PM** **Invited Talk:** Improving Application Scaling using OpenSHMEM for GPU-Initiated Communication
Presented by *Sreeram Potluri, NVIDIA*
- 4:20 PM** An Evaluation of OpenSHMEM Interfaces for Variable-length Collective Operations
Authored by *M. Graham Lopez, Pavel Shamis and Manjunath Gorentla Venkata*
- 4:45 PM** Accelerating k-NN Algorithm with Hybrid MPI and OpenSHMEM
Authored by *Jian Lin, Khaled Hamidouche, Jie Zhang, Xiaoyi Lu, Abhinav Vishnu and Dhableswar Panda*
- 5:10 PM** Parallelizing the Smith-Waterman Algorithm using OpenSHMEM and MPI-3 One-Sided Interfaces
Authored by *Matthew Baker, Aaron Welch and Manjunath Gorentla Venkata*



Day 3 OpenSHMEM Specifications

Thursday, August 6

OpenSHMEM 2015



7:00 AM Breakfast
(Provided)

8:00 AM **OpenSHMEM 1.3 -Overview**

Presented by *Steve Poole*

8:30 AM **OpenSHMEM 1.3 Specification Discussion**

This session will introduce and discuss features that are ready for OpenSHMEM 1.3, and take input from the wider community on the features and priorities for the OpenSHMEM specification that will follow 1.3. The OpenSHMEM 1.3 topics include non-blocking remote memory access (RMA) operations, Alltoall collective interface, Group communication (Teams), fine-grained completions, ordering for relaxed memory architectures, and C11 Generic selection based RMA and Atomic operation interfaces. We encourage participation from all attendees, irrespective of you regularly participating in the specification discussion or not.



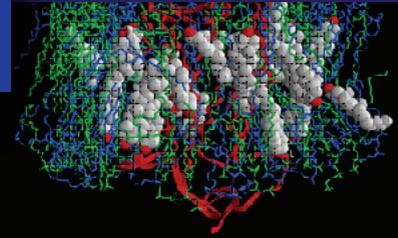
<http://www.csm.ornl.gov/workshops/openshmem2015>

Agenda

OpenSHMEM Specifications

INVITED SPEAKER BIOGRAPHIES

OpenSHMEM 2015



Larry Brown

NVIDIA



Larry is a Solution Architect with NVIDIA, where he assists customers and partners with their questions about Tesla server products and CUDA. Larry has over 15 years of experience designing, implementing and supporting a variety of advanced hardware and software systems for defense system integrators and major research universities. He has designed electro-optical systems for head-mounted displays, and adapted computer vision code for lane following in UGVs for the GPU. His current areas of focus are Data Analytics and Deep Learning. Larry has a Ph.D. from Johns Hopkins University and was with Booz Allen Hamilton prior to joining NVIDIA.

Richard Graham

Mellanox Technologies, Inc.



Dr. Richard Graham is a Staff Architect at Mellanox Technologies, Inc.

His primary focus is on the High Performance Computing market, working on OFED and communication middleware architecture issues, as they relate to extreme-scale computing. Prior to moving to Mellanox, Rich spent thirteen years at Los Alamos National Laboratory and Oak Ridge National Laboratory, in computer science technical and administrative roles, with a technical focus on communication libraries and application analysis tools. He is cofounder of the Open MPI collaboration, and he was chairman of the MPI 3.0 standardization efforts.

Manjunath Gorentla Venkata

Oak Ridge National Laboratory

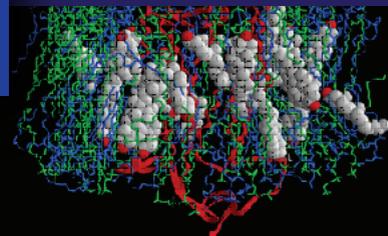


Manjunath Gorentla Venkata is a research scientist in Oak Ridge National Laboratory's Computer Science and Mathematics Division pursuing research and development efforts focused on abstractions and mechanisms that enables non-computer scientists to use the supercomputers and clusters in an efficient way. He is primarily responsible for conceiving, designing and leading the development of scalable communication interfaces, protocols and implementations for extreme-scale systems. Dr. Gorentla has published several peer-reviewed research articles in this area, contributed to various international standards, and his research has influenced commercially available network interfaces. He is a senior member of the Institute of Electrical and Electronics Engineers, and holds an affiliate assistant professor position in Auburn University's Department of Computer Science and Software Engineering.



INVITED SPEAKER BIOGRAPHIES

OpenSHMEM 2015



Ulf Hanebutte

Intel, Inc.



Dr. Ulf Hanebutte is a researcher and architect in the Extreme Scale Software Pathfinding team at Intel Corporation. A computational scientist at heart, he concentrates his expertise on a holistic system co-design for extreme scale computing with an emphasis on PGAS programming models. His career is marked by staying at the forefront of computer science, having designed, developed and analyzed complex software stacks ranging from large-scale scientific applications on the world's largest supercomputers to Cyber Physical Systems and the Internet of Things (IoT). At Intel Labs he advanced the field of platform energy efficiency with his research on ultra-mobile devices, server virtualization and memory power management. Ulf holds an advanced degree in Aerospace Engineering from the University of Stuttgart, Germany and a Ph.D. in Mechanical Engineering from Northwestern University, Evanston, IL. Prior to joining Intel, he was a research scientist at ICASE at NASA Langley, Argonne National Laboratory and CASC at Lawrence Livermore National Laboratory. Ulf is a Senior member of IEEE, and holds currently 8 US and international patents with more pending.

Yossi Itigin

Mellanox Technologies



Yossi Itigin is a Staff HPC Developer at Mellanox Technologies focusing on high-performance communication libraries. Mr. Itigin is the lead developer and maintainer of the MXM messaging library, and was the lead developer of the FCA collectives library

Dounia Khaldi

University of Houston

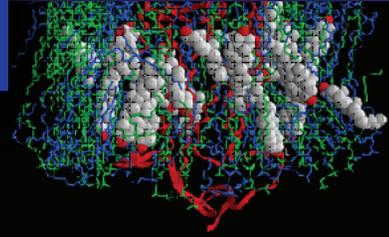


Dounia Khaldi is a post-doctoral researcher with the HPC Tools group at the University of Houston. In her current role at UH, she is involved in developing parallel intermediate representations of compilers for optimizing PGAS programs; specifically, she is the leader of the OpenSHMEM analyzer tool project in the LLVM compiler. Before joining the University of Houston, Dounia Khaldi received a PhD degree at MINES ParisTech school in France in 2013. Her research work was focused on the automatic task parallelization of sequential programs. She was granted a Master Recherche, majoring in HPC (High Performance Computing), by the Université de Versailles Saint-Quentin-en-Yvelines, France in 2010. During her studies, she specialized in compilation, HPC and parallelism issues. She obtained her engineering diploma in computer science at the Ecole supérieure d'informatique in Algiers, Algeria, in 2009.



INVITED SPEAKER BIOGRAPHIES

OpenSHMEM 2015



David Lecomber

Allinea Software



Dr. David Lecomber is the CEO and a founder of Allinea Software. He has played a role in parallel and high performance computing for over two decades. He holds a DPhil from Oxford University, where his research interests included programming models for concurrency and general purpose parallel computing, and he subsequently held research and teaching positions within the university. In 2002 he helped found Allinea Software to create tools for the forthcoming parallel era. The Allinea team has a passion for making the development of multi-process or multi-threaded codes easier - from the desktop through to supercomputers with millions of cores. David led the development team at Allinea that created the first extreme-scale capable debugger and profiler - enabling developers and scientists in government, academia and industry to achieve ground-breaking science and simulations that exploit supercomputing to its full potential.

Swaroop Pophale

Mellanox Technologies



Dr. Swaroop Pophale is a HPC Developer working at Mellanox Technologies. She was previously a post-doctoral researcher with the HPC Tools group at the University of Houston. She has been involved in high performance computing research and the OpenSHMEM project for the last six years. Her main area of research is collective optimizations, OpenSHMEM benchmarks, and compiler-based analysis for collectives. Swaroop was one of the main drivers for the OpenSHMEM 1.1 specification and has vast experience in porting codes to OpenSHMEM.

Sreeram Potluri

NVIDIA

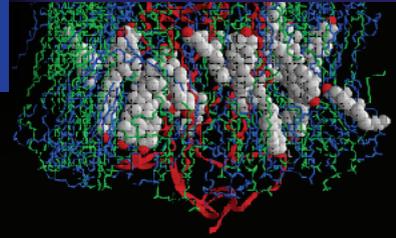


Sreeram Potluri received his PhD in Computer Science and Engineering from The Ohio State University and is currently a Senior Software Engineer at NVIDIA Corp. His research interests include high-performance interconnects, heterogeneous architectures, parallel programming models and high-end computing applications. He has published over 30 papers in major peer-reviewed journals and international conferences related to these research areas. His current focus is on designing runtime and network solutions that enable high performance and scalable communication on clusters with NVIDIA GPUs.



INVITED SPEAKER BIOGRAPHIES

OpenSHMEM 2015



Joseph Schuchart

TU-Dresden



Joseph Schuchart has been involved with scalable performance analysis tools for many years. His interests include performance analysis and tuning as well as energy-efficiency of parallel applications. Joseph has received his M.Sc. in computer science from Technische Universität Dresden, where he continued to work on scalable and efficient analysis tools as a research associate after a one year visit to work on performance tools at ORNL. He is also a passionate hobby photographer and traveler.

Pavel Shamis

Oak Ridge National Laboratory



Pavel Shamis is a staff researcher at Oak Ridge National Laboratory, who focuses on a high-performance communication middleware and collective communication. Prior to joining Oak Ridge National Laboratory, Mr. Shamis spent ten years at Mellanox Technologies in different technical roles, including Senior Software Developer and HPC Team Leader.

Sameer Shende

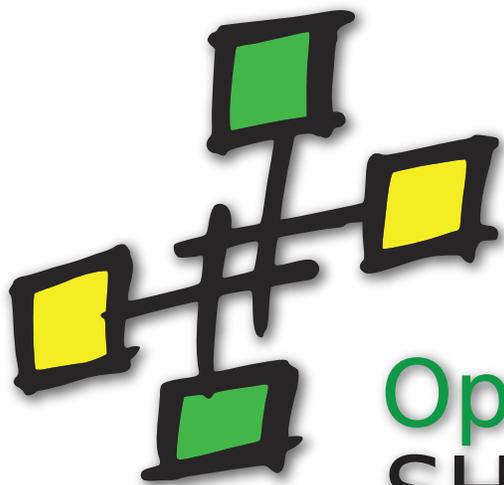
University of Oregon



Dr. Shende directs the Performance Research Laboratory at the University of Oregon, and contributes to the TAU Performance System and the Program Database Toolkit (PDT) projects. He also serves as the President of ParaTools, Inc. and ParaTools, SAS. He has worked extensively in the area of performance instrumentation and integration of performance evaluation tools with compilers over the past fifteen years. His research interests include performance instrumentation, measurement, and analysis tools, compiler optimizations, and runtime systems.



This workshop is funded by the DoD HPC program of the Computing and Computational Sciences Directorate of Oak Ridge National Laboratory.
The DOD HPC program is supported by the United States Department of Defense



Open
SHMEM

