DOE Centers of Excellence Performance Portability Meeting 2017

August 22 - 24, 2017

Crowne Plaza Denver Downtown, Denver, CO

AGENDA (updated 8.16.17)

http://www.lanl.gov/asc/doe-coe-mtg-2017.php

General Sessions: First Floor - Urban Ballroom - Breakout Sessions: Park / Office / Library / Museum Rooms

coepp-meeting@lanl.gov

Tuesday, August 22, Technical Sessions

Start	End	Duration			
7:00 AM	8:30 AM	1:30	Breakfast (paid by Registration) + Pick Up Badges - First floor foyer outside Urban Ballroom		
8:30 AM	8:40 AM	0:10	Welcome/Kickoff (Hai Ah Nam, LANL)		
			Session 1 - Performance Portability, Best Practices, Words of Wisdom - Chair: Rob Neely		
8:40 AM	9:00 AM	0:20	Levesque, John	Cray	Setting expectations for Performance Portability between Companion Accelerator and Manycore Systems
9:00 AM	9:20 AM	0:20	O'Connor, Mark	ARM	When Performance Portability is less than Perfect: Matching Applications to Architectures
9:20 AM	9:40 AM	0:20	Rupp, Karl	ANL	Exascale Computing Without Templates
9:40 AM	10:00 AM	0:20	John Pennycook, Jason Sewall, Vic	Intel	Implications of a Metric for Performance Portability: Necessity of Specialization and Application-Specific Abstractions
10:00 AM	10:30 AM	0:30	Performance Portability Panel Discussion - Moderator: Rob Neely		
10:30 AM	10:50 AM	0:20	Morning Break		
			Session 2 - Memory Hierarchy (on-node) - Chair: Doug Doerfler		
10:50 AM	11:05 AM	0:15	Olivier, Stephen	Sandia	Memory Management Extensions for OpenMP 5.0
11:05 AM	11:20 AM	0:15	Scogland, Tom	LLNL	Deep Copy and Unified Memory in OpenMP
11:20 AM	11:40 AM	0:20	Newburn, CJ	NVIDIA	A declarative approach to managing memory properties
11:40 AM	12:00 PM	0:20	Williams, Sean; Lang, Mike; Latche	LANL	Simplified Interface to Complex Memory (SICM)
12:00 PM	12:20 PM	0:20	Beckingsale, David	LLNL	Umpire: Resource Management for Heterogeneous Memory Hierarchies
12:20 PM	1:20 PM	1:00	Lunch (paid by Registration) - Elevation+Altitude Rooms		
1:20 PM	2:50 PM	1:30	**Breakout 1**		
			Topic 1 - Multi-Level Memory Current Experiences; Abstractions for managing memory for systems with diverse memory resources		
			--> Session 1: Lead: Ian Karlin (LLNL) ** PARK ROOM **		
			--> Session 2: Lead: CJ Newburn (NVIDIA) ** OFFICE ROOM **		
			Topic 2 - Performance, Portability and Productivity: Definitions & Metrics; Session Lead: John Pennycook (Intel) ** MUSEUM ROOM **		
			*** Scribe: Tina Macaluso		
			Topic 3 - Productivity Issues in Preparing Large Codes for Performance Portability; Session Lead: Anshu Dubey (ANL) ** LIBRARY ROOM **		
			*** Scribe: Emily Simpson		
2:50 PM	3:10 PM	0:20	Afternoon Break		
			Session 3 - Applications Experience 1 - Chair: John Pennycook		
3:10 PM	3:30 PM	0:20	Gunter, David	LANL	Kokkos port of CoMD mini-app for Trinity-class systems and NVIDIA Pascal nodes
3:30 PM	3:45 PM	0:15	Pankajakshan, Ramesh	LLNL	SW4Lite: Performance Portability using RAJA
3:45 PM	4:05 PM	0:20	Vaquero, Alejandro	Utah	Performance Portability Experiments with the Grid C++ Lattice QCD Library
4:05 PM	4:25 PM	0:20	Ryujin, Brian	LLNL	Experiences Porting a Multiphysics Code to GPUs
4:25 PM	5:05 PM	0:40	Breakout 1 - Recap: 4 tracks @ 10 minutes each		
5:05 PM	5:30 PM	0:25	Poster Session - 1 slide, 2 minutes for poster intro + Break		
5:30 PM	7:00 PM	1:30	**Poster Session (light appetizers) - Elevation+Altitude Rooms**		
			Arlie Capps, Peter Robinson, Josep	LLNL	Progress Porting ALE3D to the GPU
			Appelhans, David	IBM	Overlapping Data Movement and Compute with XLF and OpenMP 4: Experiences in UMT
			Gorentla Venkata, Manjunath; Ferro	ORNL	SharP: Towards Programming Hierarchical-Heterogeneous Memory based Extreme-Scale Systems
			Peles, Slaven	LLNL	Performance portability of numerical time integrators in SUNDIALS library
			North, Geraint	ARM	Experiences with performance portability across ARM microarchitectures
			Sadayappan, Saday	OSU	High-performance GPU code generation for high-order stencils: Alleviating register pressure
			Aravind Sukumaran-Rajam	OSU	TTLG: Tensor Transpose Library for GPUs
			Gonsiorowski, Elsa	LLNL	MACSio Development and Proxy Application Validation
			Luo, Lixiang	IBM	OpenACC/OpenMP4.5 Interoperability
			Liao, Chunhua "Leo"	LLNL	AutoPar: Semantics-Aware Automatic Insertion of OpenMP directives
			Dinner on your own		

Wednesday, August 23, Technical Sessions

Start	End	Duration			
7:00 AM	8:30 AM	1:30	Breakfast (paid by Registration) - First floor foyer outside Urban Ballroom		
8:30 AM	8:35 AM	0:05	Welcome/Kickoff/Announcements		
			Session 4 - Abstractions, DSL - Chair: David Richards (LLNL)		
8:35 AM	8:55 AM	0:20	Bergen, Ben	LANL	The Flexible Computational Science Infrastructure (FleCSI)
8:55 AM	9:15 AM	0:20	Newburn, CJ	NVIDIA	HiHAT, a way forward to perf portability with retargetable infrastructure
9:15 AM	9:35 AM	0:20	Kunen, Adam	LLNL	Recent experiences with RAJA nested loop abstractions and CHAI
9:35 AM	9:55 AM	0:20	Edwards, Carter	Sandia	Kokkos' Task-DAG Parallel Capabilities & Evolution of Kokkos' Back-ends
9:55 AM	10:15 AM	0:20	Abstractions, DSL Panel - Moderator: David Richards (LLNL)		

10:15 AM	10:35 AM	0:20	Morning Break		
			Session 5 - Applications Experience 2 - Chair: Charles Ferenbaugh		
10:35 AM	10:55 AM	0:20	Howard, Micah	Sandia	Performance Portability in SPARC – Sandia's Hypersonic CFD Code for Next-Generation Platforms
10:55 AM	11:15 AM	0:20	Scott Parker, Ron Rahaman, Sudhe	ANL	Portability and Performance of the Nekbone Mini-App
11:15 AM	11:35 AM	0:20	Friesen, Brian	LBL	Performance Portability Experiences at NERSC
11:35 AM	11:55 AM	0:20	Tharrington, Arnold	ORNL	Portability Initiatives for Scientific Computing and Simulation: Molecular Dynamics as a Case Study
			Session 6 - Hierarchical Memory, off-node, I/O: Chair: Charles Ferenbaugh		
11:55 AM	12:15 PM	0:20	Gonsiorowski, Elsa	LLNL	SCR and Preparing for Burst Buffers
12:15 PM	12:35 PM	0:20	Tessier, Francois	ANL	Toward portable I/O performance by leveraging system abstractions of deep memory and interconnect hierarchies
12:35 PM	1:35 PM	1:00	Lunch (paid by Registration) - Elevation+Altitude Rooms		
			Session 7 - Languages, Compilers, Frameworks, Tools - Chair: CJ Newburn		
1:35 PM	1:55 PM	0:20	Bertolli, Carlo	IBM	Performance Analysis and Optimizations for Lambda-based Applications in OpenMP 4.5
1:55 PM	2:15 PM	0:20	Larkin, Jeff	NVIDIA	Early Results of OpenMP 4.5 Portability on NVIDIA GPUs
2:15 PM	2:35 PM	0:20	Katz, Max	NVIDIA	Adaptive Mesh Refinement for Exascale
2:35 PM	2:55 PM	0:20	Jin, Xiaoyong	ANL	Performance portability via Nim metaprogramming
2:55 PM	3:15 PM	0:20	Liakh, Dmitry	ORNL	Portable Heterogeneous High Performance Computing via Domain-Specific Virtualization
3:15 PM	3:35 PM	0:20	Edwards, Carter; Trott, Christian; H	Sandia	High Performance Atomics are Critical for Thread Scalability
3:35 PM	3:55 PM	0:20	Afternoon break		
3:55 PM	4:15 PM	0:20	Bercea, Gheorghe-Teodor	IBM	Towards a portable OpenMP data sharing implementation for NVIDIA accelerators in the CLANG/LLVM toolchain
4:15 PM	4:35 PM	0:20	Languages, Compilers, Tools Panel - Moderator: CJ Newburn		
4:35 PM	4:45 PM	0:10	Breakout setup / Break		
4:45 PM	6:15 PM	1:30	**Breakout 2**		
			Topic 1 - OpenMP 5.0 features		
			--> Session Leads: Tom Scogland (LLNL) & Kevin O'Brien (IBM) ** PARK & OFFICE **		
			*** Scribe: Tina Macaluso		
			Topic 2 - Jack Deslippe (LBNL); SC Facilities Performance Portability Best Practices Website ** LIBRARY ROOM **		
			*** Scribe: Emily Simpson		
			Topic 3 - Carter Edwards (SNL); ISO/C++17 and Beyond - Parallelism and Concurrency ** MUSEUM ROOM **		
			*** Scribe:		

Thursday, August 24, Technical Sessions

Start	End	Duration			
7:00 AM	8:30 AM	1:30	Breakfast (paid by Registration) - First floor foyer outside Urban Ballroom		
8:30 AM	8:35 AM	0:05	Welcome/Kickoff/Announcements		
8:35 AM	9:15 AM	0:40	Breakout 2 - Recap: 4 tracks @ 10 minutes each		
9:15 AM	10:15 AM	1:00	**Vendor Performance Portability Panel - Moderator: Dave Bernholdt (ORNL)**		
			CJ Newburn, Principal HPC Architect for NVIDIA Compute Software		
			John Levesque, Director of Supercomputing Center of Excellence at Cray		
			John Pennycook, HPC Application Engineer at Intel		
			Kathryn O'Brien, Principal Research Staff Member at IBM T.J. Watson Research Center		
			Jonathan Gallmeier, AMD		
			Geraint North, Distinguished Engineer of ARM HPC Tools		
10:15 AM	10:35 AM	0:20	Morning Break		
			Session 8 - Applications Experience 3 - Chair: Ian Karlin		
10:35 AM	10:55 AM	0:20	Pearce, Olga	LLNL	Experiences Utilizing CPUs and GPUs for Computation Simultaneously on a Heterogeneous Node
10:55 AM	11:15 AM	0:20	Lo, Li-Ta; Woodring, Jonathan; Pop	LANL, ANL	Performance Portable Halo and Halo Center Finding in HACC
11:15 AM	11:30 AM	0:15	Quadros, Roshan; Carnes, Brian	Sandia	Scaling Post-meshing Operations on Next Generation Platforms
11:30 AM	11:50 AM	0:20	Govett, Mark	NOAA	Parallelization and Performance of the NIM Weather Model for CPU, GPU and MIC Processors
11:50 AM	12:10 PM	0:20	Teranishi, Keita	Sandia	Portability and Scalability of Sparse Tensor Decompositions on CPU/MIC/GPU Architectures
12:10 PM	12:30 PM	0:20	Blake, Robert	LLNL	Melodee: Solving ODEs with platform-specific code generation.
12:30 PM	12:35 PM	0:05	**Closing/Survey**		
12:35 PM			**Adjourn**		

Steering Committee:

Hai Ah Nam, LANL hnam@lanl.gov	Jack Deslippe, NERSC jrdeslippe@lbl.gov
Charles Ferenbaugh, LANL cferenba@lanl.gov	Tim Williams, ANL tjwilliams@anl.gov
Rob Neely, LLNL neely4@llnl.gov	Kalyan Kumaran, ANL kumaran@alcf.anl.gov
Bert Still, LLNL still1@llnl.gov	Nick Romero, ANL naromero@anl.gov
Ian Karlin, LLNL karlin1@llnl.gov	Victor Lee, Intel victor.w.lee@intel.com
Mike Glass, SNL mwglass@sandia.gov	John Pennycook, Intel john.pennycook@intel.com
Rob Hoekstra, SNL rjhoeks@sandia.gov	John Levesque, Cray levesque@cray.com
Tjerk Straatsma, ORNL straatsmatp@ornl.gov	Kathryn O'Brien, IBM kmob@us.ibm.com
Wayne Joubert, ORNL joubert@ornl.gov	Cyril Zeller, NVIDIA czeller@nvidia.com
Rebecca Hartman-Baker, NERSC rjhartmanbaker@lbl.gov	CJ Newburn, NVIDIA cnewburn@nvidia.com