October 2017 Newsletter

Highlights

Intel® HPC Developer Conference: Checkout the agenda, technical sessions, etc. and register now, for the November 11-12, 2017 in Denver, Colorado, prior to SC17. This is a free and open to the public conference where industry luminaries will share best practices and techniques through technical lectures, hands-on tutorials, poster sessions, networking and more.

A

RSVP for Intel® PCC Members only meeting: at Venice Ristorante Downtown on November 14, 2017 from 6:00PM-9:30PM (MST). Each member should have received a direct invite by Dayle Smith. If not, please send an email to IPCC.Program.Office@intel.com.

Intel® Xeon Phi™ Discount: Intel is offering a promotion to accelerate code modernization efforts for Intel® Xeon Phi™ and Intel® Xeon® Scalable family processors with a low cost solution. Promotional savings of up to 60% off and is available through all authorized Intel distributors and resellers.

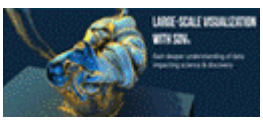

SDVis Appliance: Advancements to open source libraries such as OpenSWR, Embree and OSPRay have been proven to perform and out-perform using Intel hardware. The SDVis appliance has all the necessary software for running and managing a render farm and supports all types of professional visualization workloads and data sets up to 1.5 TB's.

Case Studies

University Of California-Davis presented a complementary physics-based, data-driven approach that exploits the causal nature of spatiotemporal data sets generated by local dynamics, illustrate how novel patterns and coherent structures can be discovered in cellular automata and outline the path from them to climate data.

National Energy Research Scientific Computing Center solved one of cosmology's toughest challenges by calculating the 3-point correlation function on the Outer Rim dataset of two billion galaxies using the NERSC Cori supercomputer.

Colfax optimized Hamerly's K-Means Clustering Algorithm on CFXKMeans Library with a

speedup of 85.6x by applying SIMD reduction with OpenMD, re-using registers with unroll and jam, analyzing various parallel reduction algorithms, and detecting workload imbalance and resolving it with scheduling.

ETH Zurich shows how the capability models of the memory subsystem, derived by systematic measurements, can be used to automatically develop new close-to-optimal algorithms for various communication functions with improvements of 5x-24x over tuned OpenMP and MPI implementations.

George Washington University saw performance improvement of 1.4x-2x by implementing techniques within Chapel, an emerging scalable, productive parallel programming language and their relation to the OpenMP implementations of the parallel research kernels.

Ohio Supercomputing Center increased vector lengths and benchmarked a 3.7x speedup that enables more convenient and in-depth exploration of loop-level performance on **WARP3D**, a 3D nonlinear finite element analysis of solids for fracture and fatigue simulation.

Free Training Opportunities: Join. Learn. Excel.

Join us at any of these upcoming educational workshops and conferences and learn about new Parallel Programming concepts, Intel® Libraries, Software Development tools and Artificial Intelligence frameworks. They are open to the public and free to attend.

Date	Location	Event
Oct 2-3, 2017	Espoo, Finland	Geocomputing Using CSC Resources
Oct 3, 2017	London, UK	Intel Software Developers Conference
Oct 4, 2017	Virtual	Memory Access Profiling: Find and Fix Common Performance Bottlenecks
Oct 11, 2017	Virtual	Is Python* Almost as Fast as Native Code? Believe It!
Oct 18, 2017	Virtual	Speed Up Small-Matrix Multiplication with Intel® Math Kernel Library
Oct 20-21, 2017	Virtual	Intel® Nervana™ AI Student Webcast
Oct 25, 2017	Virtual	Accelerating Lossless Data Compression Code for Cloud
Nov 1, 2017	Virtual	Parallel Programming Standards Update: MPI*, OpenMP* and Intel® TBB
Nov 8, 2017	Virtual	Better, Faster and More Scalable: The March To Exascale
Nov 11-12, 2017	Denver, CO	Intel® HPC Developer Conference 2017
Nov 14, 2017	Denver, CO	IXPUG BoF: Usability, Scalability & Productivity on Many-Core Processors - Intel Xeon Phi & Beyond
Jan 28-31, 2018	Tokyo, Japan	IXPUG Workshop at HPC Asia

Access to Intel® Xeon Phi™ Processor

We encourage all Intel PCC members to leverage the TACC cluster to testing your optimized application for multi-node. To request access, please click **HERE** and create a new account (do not click on PI-eligible) and follow the email instructions. Please email the **ipcc.program.office@intel.com** account and include your username in the communication.

More News…

Check out these latest HPC news stories:

- **Intel HPC Orchestrator Solves HPC's Software Challenges**
- **AI: Deeper Learning with Intel® Omni-Path Architecture**
- **The AI Revolution: Unleashing Broad and Deep Innovation**
- **Joint HPE/Intel® Xeon Phi™ Technical Brief Now Available**
- **Intel® Xeon Phi™ Processors Help Map the Brain at Argonne National Laboratory**
- **Driving Code Performance with Intel Advisor's Flow Graph Analyzer**
- **Embed Blazing Performance with Intel® System Studio 2018 Beta**
- **Lab7 Systems Helps Manage an Ocean of Information**
- **Capabilities of Intel® AVX-512 in Intel® Xeon® Scalable Processors (Skylake)**