August 2017 Newsletter

Highlights

2017 IXPUG Annual Meeting at TACC, September 26-28, 2017 in Austin, Texas: Register now to meet with experts from around the world to learn more about code modernization on Intel® Xeon Phi™ processors.

Intel® HPC Developer Conference November 11-12, 2017 in Denver, Colorado: We encourage all Intel® PCCs to share their best practices, techniques, etc. by submitting an abstract (August 20, 2017 deadline) and registering.

SAVE THE DATE for the Intel® PCC members only meeting during SC17: Mark your calendar and join us at Venice Ristorante Downtown on November 14, 2017 from 6:00PM-9:30PM MST. Invitations will be sent out soon.

Updated Tuning Guides for Intel® Architecture: Intel® VTune™ Amplifier XE Tuning Guides explains common software performance issues and suggested optimization techniques has now been updated.

Case Studies

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.

Georgia 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.

University of Utah found a new method for progressive volume rendering by accumulating object-space samples over successively rendered frames on OSPRay, a software-defined, high fidelity visualization code. It demonstrates that it is particularly useful for volumetric data with costly sampling functions.

Training Opportunities

Join any of the following upcoming parallel programing trainings with hands-on experience.

Date	Location	Event
Aug 7-11, 2017	Sydney, Australia	International Conference on Machine Learning
Aug 8, 2017	Virtual	Scalable, Distributed Deep Learning: Python* & Pachyderm*
Aug 11, 2017	Seoul, South Korea	Intel Xeon Phi™ Programming
Aug 17, 2017	Boulder, CO	KNL Training at RMACC
Aug 22, 2017	Virtual	Achieving High-Performance Computing with Intel® Python
Sept 6, 2017	Espoo, Finland	Intro to C Programming
Sept 9-13, 2017	Portland, USA	Parallel Architectures and Compilation Techniques
Sept 18-29, 2017	Virtual	Colfax: HOW Series Webinar
Sept 19-21, 2017	Virtual	Introducing Intel's Thread Building Blocks
Oct 2-3, 2017	Espoo, Finland	Geocomputing Using CSC Resources
Anytime	Virtual	Intel® High Performance Computing Webinar Archive
Anytime	Virtual	Machine Learning Part 1, Part 2, Part 3, Part 4

Access to Intel® Xeon Phi™ Processor

Optimize your applications for multi-node by testing on the following clusters:

Texas Advanced Computing Center (TACC) Stampede Cluster:

- Click HERE and create a new account (**do not click on PI-eligible**) and follow the email instructions.
- Register account by emailing ipcc.program.office@intel.com with username.

More News...

Catch up on the latest HPC news:

- The Parallel Universe
- Texas Advanced Supercomputing Center Taps Latest HPC Tech
- Technology Requirements for Deep and Machine Learning
- AI: Scaling Neural Networks Through Cost-Effective Memory Expansion