






December 2018 Newsletter

Intel SC18 Actives & Event

What an amazing week of sessions, demos, booth interviews, and workshops! Thanks to all who participated in the [Intel activities and events during SC18](#). The sharing of your experiences and best known methods in optimizing applications using industry open standards truly helps the broader developer community to reach new discoveries fast. One example is the paper [CosmoFlow: Using Deep Learning to Learn the Universe at Scale](#). We look forward to sharing replays of the Intel® HPC Forum, including speakers from TACC, LRZ, NERSC, and CERN, soon in the New Year.

Highlights

 <p>Optimization Techniques</p>	<p>Advanced OpenMP: Host Performance and 5.0 Features: exploring the effects of different OpenMP parallelization strategies, their correctness, and performance.</p> <p>Anatomy of High-Performance Deep Learning Convolutions on SIMD Architectures: introduces direct convolution kernels for x86 architectures, in particular for Xeon and Xeon Phi systems, implemented via a dynamic compilation approach.</p> <p>24x Deep Learning Algorithm Performance Boost : see how video surveillance solution provider GeoVision turbo-charged its deep learning facial recognition solution using Intel® System Studio and OpenVINO™ Toolkit.</p> <p>Framework for Scalable Intra-Node Collective Operations Using Shared Memory: focusses on the performance benefit of optimizing intra-node collectives using shared memory.</p>
 <p>Case Studies</p>	<p>CosmoFlow: Using Deep Learning to explore the Universe at Scale: demonstrated fully synchronous data-parallel training on 8192 nodes of Cori with 77% parallel efficiency, achieving 3.5 Pflop/s sustained performance.</p>

	<p><u>Many-Core Graph Workload Analysis</u>: provided new insights on executing graph applications on many-core processors.</p> <p><u>Performance Tuning of Scientific Codes with the Roofline Model</u>: exposing the fundamental aspects behind different Roofline modeling principles and providing several practical use-case scenarios to highlight their efficacy for application optimization.</p> <p><u>Continuous Performance Benchmarking Framework for ROOT</u>: presenting a framework, built from industry best practices and tools, to help to understand ROOT code performance and monitor the efficiency of the code for a several processor architectures</p> <p><u>Apple* Machine Learning on Intel® Processor Graphics</u>: talks about taking advantage of Apple* Machine Learning Frameworks <u>Core ML*</u> and <u>Metal*</u> <u>Performance Shaders (MPS)</u> on Intel® Processor Graphics.</p>
 <p>Scientific Breakthrough</p>	<p><u>MoISSI announces the open force field consortium</u>: An academic-industry collaboration designed to improve classical molecular models to guide pharmaceutical drug discovery. All software, force fields, and development tools will be fully open source and available to everyone.</p> <p><u>How Supercomputers Can Save Lives in Airborne Chemical Attacks</u>: researchers developed a high-resolution simulation model of the Syrian deadly chemical attack air turbulence, reducing simulation time from 5 days to 4 hours the</p> <p><u>The Intersection of AI, HPC and HPDA: How Next-Generation Workflows Will Drive Tomorrow's Breakthroughs</u>: planning for the future isn't always easy. Intel's Trish Damkroger offers some insights on how Intel is working with their customers to help ensure they are prepared to take on the breadth of emerging HPC and AI workloads.</p> <p><u>Accelerating the Shift to Software Defined Visualization</u>: taking advantage of modern, general-purpose processors, SDVis offers a scalable, highly performant approach to visualizing massive data sets without GPU hardware. Intel PCCs demonstrate the superior capabilities of CPUs for visualization.</p>

Testing Your Code on Intel® Architecture

Intel® Xeon® Scalable Platform Access: We encourage testing applications using various configurations of Intel® architecture (Intel® Xeon Scalable processors, Intel® Omni-Path, etc. Click [HERE](#) to test your optimized application at scale using TACC, Stampede II system. Upon requesting access, create a new account (do not click on PI-eligible) and follow the email instructions. Then email the ipcc.program.office@intel.com account and include your username in the communication.

***New Software: [Intel® Parallel Studio XE 2019](#):** new features that include high performance Python* innovations, intuitive Intel® VTune™ Amplifier interface, visualization capabilities with Intel® Advisor's Flow Graph Analyzer, next-gen Intel® MPI Library for preparing HPC solutions for Exascale, and more.

***Tools: advance vision solutions for AI at the Edge:** the latest [Intel® Distribution of OpenVINO™ toolkit \(2018 R4\)](#) supports [new vision accelerators](#): Intel® Vision Accelerator Design with an [Intel® Arria® 10 FPGA](#) PCIe add-in card and [Intel® Neural Compute Stick 2](#). This hardware/software combo unleashes [deep learning performance](#) and speeds real-time image analysis. Other new toolkit features expand topologies support and add three new pre-trained models. [Download it free](#).

Speaker & Publication Opportunities

There are several opportunities for you to share your learnings, best practices and techniques around the benefits you've received in leveraging Intel® architecture. We would like bring to your attention some key abstract submission deadlines for 2019 conferences and workshops. Feel free to submit abstracts to all that interest you.

Submission Deadline	Event
December 12, 2018	ISC19 Research Papers
February 6, 2019	ISC19 Tutorials & Project Posters
February 13, 2019	ISC19 PhD Forum & Workshops(regular)
February 20, 2019	ISC19 BoF Session
February 27, 2019	ISC19 Research Posters

Global Event & Training Opportunities

We encourage you to participate in any of the upcoming global training and free webinar opportunities.

Date	Location	Event
December 12, 2018	Webinar	Under What Conditions will My Application give Reproducible Results?
December 14, 2018	Webinar	AI on PC: Introduction to Deep Learning Inference on the PC
December 17 – 20, 2018	Bengaluru, India	25th IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC 2018)
December 13, 2018	CERN openlab, Switzerland	Workshop: Build Deep Learning Applications for Big Data using Intel Analytics Zoo
January 9, 2019	Webinar	Accelerating the Heterogeneous World with OpenCL*
January 14 -16, 2019	Guangzhou, China	HPC Asia 2019
January 7, 2019	Guangzhou, China	IXPUG Workshop at HPC Asia 2019

February 27 - March 2, 2019	Minneapolis, USA	SIGCSE 2019
March 12 - 15, 2019	Warsaw, Poland	Supercomputing Frontiers Europe 2019
March 31 – April 4, 2019	Orlando, FL	American Chemical Society Spring 2019 National Meeting & Exposition
June 16-20, 2019	Frankfurt, Germany	ISC 2019
July 28-August 1, 2019	Chicago, IL	PEARC 19
September 15 – 20, 2019	San Antonio, TX	SEG 19
November 17 -20, 2019	Brisbane, Australia	SIGGRAPH Asia 2019
Anytime	Webinar	What Intel® Processor Graphics GEN9 Unlocks in OpenCL*
Anytime	Webinar	All About Persistent Memory Flushing
Anytime	Webinar	Configure and Manage Persistent Memory DIMMs
Anytime	Webinar	Persistent Memory Programming Today and in the Future

More News...

Check out the latest Intel® news:

- [Turn Unsolvable into Breakthrough](#)
- [Could Intel's upcoming event reveal more about its new discrete GPU?](#)
- [Intel Shows Breadth of Data-Centric Platform with Cascade Lake Advanced Performance and Xeon E-2100](#)

© 2018, Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

To subscribe to the Intel PCC mailing list, please register [HERE](#). To [unsubscribe](#) from other Intel communications, please reply to those directly. Or contact us at this address: Intel Corporation, 2200 Mission College Blvd., M/S SC3-37, Attn: Unsubscribe/Privacy, Santa Clara, CA 95054. Intel Corporation has never engaged in the practice of sharing information about individual subscribers or sharing it with third parties. [Intel Privacy Policy](#)