



September 2018 Newsletter

Highlights



Optimization Techniques


[TensorFlow at NERSC](#): The versions of TensorFlow that have been compiled to include specific optimizations contributed by Intel to improve compute performance on KNL are recommended.

[Hiding Communication Latency Using MPI-3* Non-Blocking Collectives](#): Linear scalability for applications is worthy pursuit. NBC in the Intel® MPI library 2018 update 2 provides the tools to make it happen in the MPI-3 standard.

[Advancing Artificial Intelligence on Apache Spark* with BigDL](#): A look at BigDL's implementation, real-world use-cases, and a glimpse into the new end-to-end analytics plus AI pipelines which being built on top of Apache Spark* and BigDL.

[Faster Gradient-Boosting Decision Trees](#): Explore the gradient boosted trees algorithm in Intel® Data Analytics Acceleration Library (Intel® DAAL) and why it's so widely used by ML practitioners.

[Code Modernization in Action: Thread, Memory, and Vector Optimizations](#): Imagine an expert telling you"

	<p>Your application underutilizes the Floating Point Unit,” then providing the steps to correct the issues. Intel® Parallel Studio XE is that expert and so much more.</p>
 <p>Case Studies</p>	<p><u>CERN Project Sees Orders-of-Magnitude Speedup with AI Approach</u>: Demonstrated that AI-based models have the potential to act as orders-of-magnitude-faster replacements for computationally expensive tasks in simulation, while maintaining a remarkable level of accuracy.</p> <p><u>Enabling Detailed, Biophysics-Based Skeletal Muscle Models on HPC Systems</u>: MegaMol provides the necessary basic infrastructure for enabling distributed rendering on an HPC system. And it is already capable of rendering discretized muscle fibers as continuous geometry.</p> <p><u>Visit-OSPRay: Toward an Exascale Volume Visualization System</u>: 7x-34x performance improvements brought to VisIt by OSPRay on modern Intel architectures. Also up to 100x reduction in memory consumption when using OSPRay on KNL.</p> <p><u>Local Causal States and Discrete Coherent Structures</u>: University of California at Davis paper laying out the foundations of their novel theory of coherent structures and its application to cellular automation is now published.</p> <p><u>CPU Isosurface Ray Tracing of Adaptive Mesh Refinement Data</u>: Utah presents “Octant” reconstruction algorithm to compute crack free implicit surfaces from BS-AMR data to providing better quality and higher performance isosurface rendering.</p> <p><u>Modernizing ‘Cranky’ Rendering Code for HPC-Based Animation</u>: Using the new highly parallel hardware and software to improve “the time to first decision” in new projects. It’s the reason that Pixar teamed up with Intel.</p>



Science Breakthrough

[Argonne Leverages HPC and Machine Learning to Accelerate Science](#): Planning to deploy Aurora A21, a new Intel-Cray system, slated to be the first Exascale supercomputer in the United States.

[Intel Hosts NASA Frontier Development Lab Demo Day for 2018 Research Presentations](#): Intel is addressing critical knowledge gaps by using AI to further space exploration and solve problems that can affect life on Earth.

[Intel Artificial Intelligence Helps Bring 'The Meg' Mega Shark to the Big Screen](#): "The Meg" shows a simulated 75-foot-long shark, the Megalodon. Intel artificial intelligence technology helped in the creation of the prehistoric shark for the 2018 film.

[Slashing HPC Energy costs with Automated, Dynamic Optimization](#): Static approach has demonstrated improvements in energy efficiency up to about 10 percent.

[Drive AI Innovation with Your Evolved HPC Infrastructure](#): AI is forming a symbiotic relationship with HPC, where the two are driving each other. AI bring massive data to table, while HPC delivers the high compute performance needed to manipulate it.

Testing Your Code on Intel® Architecture

We encourage testing applications using various configurations of Intel® architecture (i.e. Intel® Xeon processor, Intel® Xeon Phi™ processor, Intel® Omni-Path, etc. Click [HERE](#) to test your optimized application 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.

Global Training Opportunities

Please check the links below for details on upcoming global training opportunities.

Date	Location	Event
Anytime	Virtual	NASA FDL Event Horizon 2018
Anytime	Virtual	Deep Learning and Natural Language Processing
Anytime	Virtual	NERSC Data Day (presentations available to download)
Anytime	Virtual	NERSC Big Data Summit (presentations available to download)
Anytime	Virtual	Get Your Code Future-Ready with Free Webinars
Anytime	Virtual	IXPUG Working Group: Machine Learning at Scale
Anytime	Virtual	Introduction to TensorFlow* with Intel® Optimization
Anytime	Virtual	Getting Started with the Intel® AI DevCloud
Anytime	Virtual	How to Identify Causes of Poor OpenMP Parallel Performance Using the Intel® VTune Amplifier
Anytime	Virtual	Using Roofline Analysis to Analyze, Optimize, & Vectorize Iso3DFD with Intel® Advisor
September 24-28, 2018	Barcelona, Spain	OpenMP Conf. 2018 & International Workshop OpenMP 2018
September 25-28, 2018	Hillsboro, Oregon	IXPUG Annual Fall Conf. 2018
September 26-27, 2018	Berlin, Germany	Intel® Software Developer Workshop for Technical Computing and Artificial Intelligence
October 5, 2018	TACC, US	C++ for C Programmers
October 8-11, 2018	London, United Kingdom	O'REILLY Artificial Intelligence Conf.
October 14-16, 2018	Qingdao, China	2018 National Annual Conf. on HPC

October 14-19, 2018	Anaheim, California	Society of Exploration Geophysicists Annual Meeting 2018
October 18-20, 2018	Qingdao, PRC	HPC China 2018
October 21-26, 2018	Berlin, German	VIS 2018
November 2-4, 2018	Nanjing, PRC	MLA'18- The 16th China Symposium on Machine Learning and Applications
November 12-14, 2018	Milano, Italy	Introduction to Parallel Computing with MPI and OpenMP
November 14-16, 2018	Beijing, PRC	Asian Conference on Machine Learning
December 3-8, 2018	Montreal, Canada	Neural Information Processing Systems (NIPS) Conf. 2018
March 12-15, 2019	Warsaw, Poland	Supercomputing Frontiers Europe 2019

More News...

Check out these latest news stories:

- [NERSC, Intel, Cray Harness the Power of Deep Learning to Better Understand the Universe](#)
- [The Convergence of Big Data and Extreme-Scale HPC](#)
- [Intel Shows off its Xeon Roadmap Through 2020](#)
- [New Texas Supercomputer to Push the Frontiers of Science](#)
- [Technology Trends Driving HPC](#)
- [E-computing: The Next Steps](#)

© 2018, Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. 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 address: Intel Corporation, 2200 Mission College Blvd., M/S SC3-37, Attn: Unsubscribe/Privacy, Santa Clara, CA 95051. We do not have the practice of sharing information about individual subscribers or sharing it with third parties. [Intel Privacy Policy](#)