



February 2019 Newsletter

Highlights



Optimization Techniques

[Implementing an OpenCL compiler for CPU in LLVM](#): how an OpenCL compiler front-end can easily handle various target devices and how it all neatly revolves around LLVM/clang & tools.

[Floating- Point Reproducibility in Intel® Software Tools](#): explore the latest tools to help you get reproducible results from binary floating-point calculations.

[AI Practitioners Guide for Beginners](#): detailed instructions for installing and running popular AI examples three different ways: on bare metal, via containers, and on the cloud.

[Accelerate Machine Learning with Intel® Math Kernel Library \(Intel® MKL\)](#): Intel® MKL can speed up machine learning training by a factor of up to 7.9 compared to default BLAS libraries.

[Energy Flow Networks: Deep Sets for Particle Jets](#): Adaption of the "Deep Sets" framework to particle physics, including infrared and collinear safety factors essential for representing and learning from high-energy physics collider data.

[Rocketing the Development of Hypersonic Vehicles](#): Moscow Institute of Physics and Technology creates faster and more accurate computational fluid dynamics software with help from Intel® Math Kernel Library and Intel® C++ Compiler.




Case Studies

[Large Minibatch Training on Supercomputers with Improved Accuracy and Reduced Time to Train](#): demonstration of >90% scaling efficiency and a training time of 28 minutes using up to 104K x86 cores.

[Performance Characterization of Multi-threaded Graph Processing Applications on Many-Integrated-Core Architecture](#): Intel® Xeon® Phi™ demonstrates encouraging performance and power efficiency when running multi-threaded graph applications.

[Fast and Accurate Training of an AI Radiologist](#): explore ways of developing accurate models with fast parallel training on Zenith, the Intel Xeon-based supercomputer and improve training time from days to hours.

	<p><u>Evaluation of Intel Memory Drive Technology Performance for Scientific Applications:</u> found that for some applications DDR4/Optane hybrid configuration outperforms DDR4 setup by up to 20%.</p> <p><u>Massively scalable computing method to tackle large eigenvalue problems for nanoelectronics modeling:</u> learn how a single Lanczos iteration can be extended with a multi-level parallelism to solve highly degenerated systems.</p> <p><u>Arctic Ocean-Sea Ice Interactions:</u> the visualization project aimed at understanding and quantifying the Arctic ocean-sea ice mean state and its changes in response to the Earth's recent warming.</p>
 <p>Scientific Breakthrough</p>	<p><u>A Global HPC-Powered Hub for Natural Hazard Research:</u> using TACC's supercomputer to create machine learning algorithms to auto-tag disaster images so they don't have to be tagged manually after future disasters.</p> <p><u>Visualizing Outbursts of Massive Stars:</u> to better understand outbursts and variability on evolved massive stars, three dimensional simulations of radiation dominated massive stars were calculated on Argonne Leadership Computing Facility's supercomputer, Mira.</p> <p><u>A Big Data Journey While Seeking to Catalog our Universe:</u> Julia can offer high performance "at scale" (using hundreds of thousands of processor cores for compute), and their success certainly indicates that we will see more "at scale" big data work.</p> <p><u>Massachusetts General Hospital (MGH)* Achieves 20X Faster Colonoscopy Screening Processing Time:</u> MGH 3D Imaging Research*, working with Intel, Microsoft*, and Vectorform*, used Intel® Parallel Studio 2011 to optimize key image processing libraries, reducing the compute-intensive colon-screening processing time from 60 min to 3 min.</p> <p><u>Optimizing Application performance with Roofline Analysis:</u> NERSC boosts the performance of its scientific applications on Intel® Xeon® Scalable processors up to 35% using Intel® Parallel Studio and Intel® Advisor.</p> <p><u>AI and HPC Are Converging with Support from Intel® Technology:</u> Dr. Pradeep Dubey, Intel Fellow and director of its Parallel Computing Lab, explains why it makes sense for the HPC and AI to come together and how Intel is supporting this convergence.</p>

Intel® Parallel Computing Centers Invited Talk Series

The Intel® PCC Invited Talk Series are presented by global partners as they share experiences in using Intel architecture for scientific breakthrough. The presenters will share optimization techniques, best practices and results. This series is intended for students, educators, developers/programmers, scientist, data analyst, system administrators, etc. working to maximizing software efficiency using Intel technology. Please come to join February 2019 Intel® PCC Invited talk "**Hybrid Quantum-Classical Computing Architectures**".

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 TAC, 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.

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
February 13, 2019	ISC19 PhD Forum & Workshops(regular)
February 20, 2019	ISC19 BoF Session
February 26, 2019	O'REILLY + Intel Artificial Intelligence Conference, San Jose Call for Speakers
February 27, 2019	ISC19 Research Posters
April, 2019	O'REILLY + Intel Artificial Intelligence Conference, London Call for Speakers

Global Event & Training Opportunities

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

Date	Location	Event
February 19 & 26, 2019	Princeton University	Introduction to Programming Using Python
February 20, 2019	Webinar	Computer Vision and Deep Learning: A Synergistic Approach to Success
February 27 - March 2, 2019	Minneapolis, USA	SIGCSE 2019
March 6, 2019	Webinar	OpenMP* 5.0: A Story about Threads and Tasks
March 12 - 15, 2019	Warsaw, Poland	Supercomputing Frontiers Europe 2019
March 13, 2019	Webinar	Hybrid Cloud: Best Practices for HPC Performance On Premise and On Cloud
March 20, 2019	Webinar	Efficient HPC Communications: Profiling and Tuning MPI Applications
March 27, 2019	Webinar	Get Deep Learning Framework Performance on Intel® Architecture

March 31 – April 4, 2019	Orlando, FL	American Chemical Society Spring 2019 National Meeting & Exposition
April 15-18, 2019	New York, US	O'REILLY + Intel Artificial Intelligence Conference
May 6-9, 2019	New Orleans, US	ICLR 2019
June 10-15, 2019	Long Beach, CA	ICML 2019
June 15-21, 2019	Long Beach, CA	CVPR 2019
June 16-20, 2019	Frankfurt, Germany	ISC 2019
June 18-19, 2019	Beijing, PRC	O'REILLY + Intel Artificial Intelligence Conference
July 28-August 1, 2019	Chicago, IL	PEARC 19
August 26-30, 2019	Gottingen, Germany	EURO-PAR 2019
September 9 -12, 2019	San Jose, CA	O'REILLY +AI Artificial Intelligence Conference
September 15 – 20, 2019	San Antonio, TX	SEG 19
September 17-18, 2019	Mountain View, CA	AI Hardware Summit 2019
October 14-17, 2019	London, UK	O'REILLY + Intel Artificial Intelligence Conference
November 17 -20, 2019	Brisbane, Australia	SIGGRAPH Asia 2019
December 2-8 , 2019	Vancouver, CA	NeurIPS 2019
Anytime	Webinar	Hybrid Cloud, Containers & Developer Services to Transform the Enterprise

More News

Check out the latest Intel® news:

- [Raj Hazra from Intel presents: Our Journey...Accelerating](#)
- [Intel Xeon Scalable and Optane: Transforming the data centre](#)
- [Hybrid Cloud—Know What's Under the Hood](#)

© 2019, Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries. *Other names 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 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](#)