






May 2019 Newsletter

ISC19 Intel Activities

ISC High Performance 2019 (ISC19) June 16 - 20, 2019 in Frankfurt, Germany will bring together the international supercomputing community—a gathering of scientists, engineers, researchers, educators, programmers, system administrators, and developers that is unequaled in the world. Join [Intel & partner presentations](#) across the community, including the [ISC19 IXPUG Workshop: "Using FPGAs to Accelerate HPC & Data Analytics on Intel-Based Systems"](#). Additional events, activities, and the Intel mobile agenda will be announced soon.

Highlights

 <p>Optimization Techniques</p>	<p><u>Containers in HPC</u>: a brief introduction on how to build images, run containers on HPC systems and best practices to ensure containers can take full advantage of HPC systems.</p> <p><u>Introducing the Model Zoo for Intel® Architecture</u>: more than 20 pre-trained models, benchmarking scripts, best practice documents, and step-by-step tutorials for running deep learning models optimized for Intel® Xeon® Scalable processors.</p> <p><u>Future Directions for Natural Language Processing in Commercial Environments</u>: new NLP practices are addressing transfer learning by learning language structures from domains and tasks rich with labelled examples and applying the learned model with some adaptation to a different domain/task.</p> <p><u>nGraph-HE: A Graph Compiler for Deep Learning on Homomorphically Encrypted Data</u>: modifying existing TensorFlow code to use nGraph-HE requires adding only a single line of code. This makes nGraph-HE extremely easy to use by data scientists.</p>
---	--

	<p><u>Quantitatively Assessing Performance Portability with Roofline</u>: discussion around 'performance portability', why it is important and desirable, and how to quantitatively measure it.</p>
 <p>Case Studies</p>	<p><u>Impact of mixed precision and storage layout on additive Schwarz smoothers</u>: acceleration of the smoother application depends mainly on the reduction of the smoother storage size, and to some extent also on the access patterns.</p> <p><u>A particle-filter framework for robust CryoEM 3D reconstruction</u>: a particle-filter algorithm to achieve higher resolution and observe atomic details in biochemical systems.</p> <p><u>Performance and Energy Impact of OpenMP Runtime Configurations on Power Constrained System</u>: careful configuration selection can improve execution time and energy consumption of an OpenMP parallel region up to 67% and 72% respectively compared to system default.</p> <p><u>Optimization Practice of Deep Learning Inference Deployment on Intel® Processors</u>: tests completed by iQIYI on their own application development environment show that the performance of different FP32 DL neural networks can be improved from 4 to 9 times.</p> <p><u>Real-World AI at Enterprise Scale</u>: get more from the Xeon foundation you already know, accelerate with purpose, and use software to simplify the environment.</p>
 <p>Scientific Breakthrough</p>	<p><u>TACC Supercomputers Play Pivotal Role in Event Horizon Telescope's First-Ever Black Hole Image</u>: researchers were able to turn theoretical models of black holes into specific prediction, which ultimately gave scientists confidence that the image produced by the EHT collaboration was accurate.</p> <p><u>40 Powers of 10 – Simulating the Universe with the DiRAC HPC Facility</u>: a revolutionary new cosmological hydrodynamics code developed at Durham which is 20x faster than the state-of-the-art.</p> <p><u>A New Approach to Multiplication Opens the Door to Better Quantum Computers</u>: tail call optimization allows the quantum algorithm to avoid creating intermediate information that a quantum computer can never discard.</p> <p><u>BSC Researchers Shrink Floating Point Formats to Accelerate Deep Neural Network Training</u>: evaluation of the applications of many formats and codes, including Intel DL Boost (such as VNNI and others), floating point, Flexpoint, and integer formats, at various phases of training neural networks and inferencing.</p> <p><u>Argonne's Aurora System and the Road to Exascale Computing</u>: the new Aurora Exascale system will provide unprecedented capability to integrate data analytics and simulation in order to advance fundamental science research.</p>

Intel® Parallel Computing Centers Invited Talk Series

Working with researchers in various domains, Argonne Leadership Computing Facility (ALCF) scientists are optimizing popular machine learning frameworks (such as TensorFlow* and PyTorch*) to understand scaling performance and reduce the time to a solution. This presentation describes this research program and tools that support the program. Please register and join us on May 30th 8:00-8:30am (Pacific Standard Time) to learn about ALCF's new update with [Scale Data-Intensive & Deep Learning Applications](#).

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
May 31, 2019	HPC CHINA 2019
July 31, 2019	SC19 Research Posters
July 31, 2019	SC19 Scientific Visualization & Data Analytics Showcase
July 31, 2019	SC19 Bird of A Feather(BoF)
July 31, 2019	SC19 Early Career Program Applications
July 31, 2019	SC19 Doctoral Showcase
July 31, 2019	SC19 Exhibitor Forum
July 31, 2019	SC19 HPC Impact Showcase

Global Event & Training Opportunities

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

Date	Location	Event
May 7-8, 2019	TACC, Austin	Advanced Computing Foundations
May 6-9, 2019	New Orleans, US	ICLR 2019
June 4-5, 2019	Edinburgh, UK	2nd UK OpenMP Users Conference
June 5, 2019	CERN, Switzerland	Speeding up Scientific Codes in HPC Architectures by Code Modernization: Lessons Learned
June 10-13, 2019	TACC, Austin	Scientific Visualization
June 10-14, 2019	Houston, Texas	Rice University Tapia Center Workshop in Computational Chemistry
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
June 25-28, 2019	TACC, Austin	Designing and Administering Large-scale Systems
July 9-12, 2019	TACC, Austin	Workflows and Reproducibility in Scientific Computing
July 17-21, 2019	Cookeville, TN	Integrating Parallel and Distributed Computing in Introductory Programming Classes Workshops
July 22-31, 2019	TACC, Austin	MoISSI 2019 Software Summer School
July 23-26, 2019	TACC, Austin	Computational Science in the Cloud
July 28-August 1, 2019	Chicago, IL	PEARC 19
August 5-8, 2019	TACC, Austin	Machine Learning Foundations
August 8-9, 2019	Arlington, Virginia	MoISSI "Teach the Teachers" Instructor Training
August 26-30, 2019	Gottingen, Germany	EURO-PAR 2019
September 9 -12, 2019	San Jose, CA	O'REILLY +AI Artificial Intelligence Conference
September 9-11, 2019	Auckland, New Zealand	OpenMPCon 2019
September 11-13, 2019	Auckland, New Zealand	IWOMP 2019
September 15 – 20, 2019	San Antonio, TX	SEG 19
September 17-18, 2019	Mountain View, CA	AI Hardware Summit 2019
September 17-19, 2019	TACC, Austin	HPC Leadership
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
May 9, 2019	Webinar	Deeply-Pipelined FPGA Clusters Make DNN Training Scalable
Anytime	Webinar	Parallelism in Python*

More News

Check out the latest Intel® news:

- [We're Stepping on the Gas Pedal for Hybrid Cloud](#)
- [The Long View on the Intel Xeon Architecture](#)
- [Intel to Deliver Exascale for the Advancement of Science](#)

© 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 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](#)