



July 2019 Newsletter

Conference News

In September, CERN will host the Intel Extreme Performance Users Group (IXPUG) [Annual Conference](#) in Geneva, Switzerland. The [Call for Abstracts](#) is open until July 19. This is a great opportunity to interact with a dynamic community of HPC and AI experts and share your own research in areas of HPC, Data Analytics, and Artificial Intelligence (Machine and Deep Learning), Cloud Computing, and more.

Congratulations to University of California, Davis (project leader Adam Rupe) on Project DisCo, the first distributed HPC application of a physics-based, data-driven technique called “local causal states” (latent variables that decompose system into structurally relevant components), for being awarded the HPC Innovation Excellence award by the HPC User Forum, Hyperion Research at ISC 2019. The team at UC Davis, in collaboration with Intel and NERSC, used this technique on climate data from the CAM5.1 global atmospheric model, processing almost 90 terabytes of data in less than 7 minutes on the Cori supercomputer. Read about it [HERE](#) and [HERE](#).

The ISC 2019 (IXPUG) Workshop "Using FPGAs to Accelerate HPC & Data Analytics on Intel-Based Systems" featured the integration and use of FPGA devices in HPC, Data Analytics, and Artificial Intelligence (Machine and Deep Learning) workloads in Intel-based HPC systems. Presentations are available on the [IXPUG website](#).

Highlights





Optimization Techniques

[DAOS Revolutionizes High Performance Storage](#): learn how we're revolutionizing high performance storage with the Distributed Asynchronous Object Storage (DAOS), Intel open-source software stack, and Intel® Optane™ technology.

[Optimizing Deep Learning LSTM Topologies on Intel Xeon Architecture](#): a speed-up of 1.9× is achieved using LSTM cell over original TensorFlow implementation for 8-layer German-to-English translation model training.

[Effectively Train and Execute Machine Learning and Deep Learning Projects on CPUs](#): check out Intel's optimizations for frameworks like TensorFlow*, MXNet*, Pytorch*, and Caffe* and see how [Intel® MKL for DNN](#) helps speed training and inference execution on CPUs.

	<p><u>Understanding and Integrating Intel Deep Learning Boost</u>: an overview of Intel's DL Boost technology, featuring integer vector neural network instructions targeting future Intel® Xeon® scalable processors.</p>
 <p>Case Studies</p>	<p><u>Evolutionary Convolutional Neural Network for High Energy Physics Detector Simulation</u>: CERN has developed a 3DGAN model to simulate high granular calorimeter response, resulting in a 3000x speed-up while retaining 10% accuracy with respect to Monte Carlo.</p> <p><u>Supercomputer-Scale Training of Large AI Radiology Models</u>: in the case of ResNet-50, by using 256 nodes, the total time required to obtain a trained model is reduced by 187x.</p> <p><u>Petaflop Seismic Simulations in the Public Cloud</u>: demonstrate that it is possible to achieve petascale performance for the same scientific workload, when using an elastic petascale cluster in the public cloud.</p> <p><u>Quantum Natural Language Processing</u>: target corpuses with ~2000 most common words simulating up to 41 qubits. The implemented solution will be able to compute the meanings of two sentences and decide if their meanings match.</p> <p><u>Densifying Assumed-sparse Tensors</u>: using the Stampede2 supercomputer, CPU-only scaling tests achieved 91% weak scaling efficiency up to 1200 MPI processes (300 nodes), and up to 65% strong scaling efficiency up to 400 MPI processes (200 nodes).</p> <p><u>PartNet: A Large-scale Benchmark for Fine-grained and Hierarchical Part-level 3D Object Understanding</u>: present PartNet: a consistent, large-scale dataset of 3D objects annotated with fine-grained, instance-level, and hierarchical 3D part information.</p> <p><u>Twitter Gets Up to 50% Faster Runtimes for its Hadoop* Clusters</u>: result in approximately 30% lower TCO and over 50% faster runtimes compared to their legacy production cluster configuration.</p>
 <p>Scientific Breakthrough</p>	<p><u>Nation's First Exascale-Capable Supercomputer Advances Clean Fusion Research</u>: quick and accurate predictions regarding so-called 'disruptive events' which allow hot, thermonuclear plasma to escape quickly. Supervised machine-learning helps us predict it, then we can plan to control it."</p> <p><u>Aurora Supercomputer to Empower Advanced Chemistry Research</u>: using Exascale computers to develop a computational framework to accelerate the discovery of new catalysts, building capability to solve yesterday's insurmountable scientific problems.</p> <p><u>Advancing the Industrial Internet with Simulation-Based Digital Twins Powered by HPC</u>: digital twins are an excellent example of new, data-rich applications that are delivering continuous intelligence and radically improved decision-making across the enterprise.</p> <p><u>Interpretable Machine Learning for Generating Semantically Meaningful Formative Feedback</u>: a system is introduced for delivering</p>

	<p>Formative Feedback for Autistic Children Behavioral Therapy and it's built on an interpretable machine learning framework.</p> <p>Self-driving technology and the future autonomous depot-to-depot transport: PlusAI is developing a full stack self-driving technology to enable large-scale autonomous commercial fleets.</p>
--	--

Intel® PCC New Invited Talk Series

As simulations grow in scale, so does the need for *in situ* analysis methods, such as in-transit visualization, to handle the large amount of data produced. Get a demonstration on batch analysis and interactive visualization using libIS. -Please register and join us on July 25th 8:00-8:30am (Pacific Standard Time) to learn about Will Usher's update with [A Lightweight Library for Flexible In-Transit Visualization](#).

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
July 19, 2019	IXPUG Annual Conference – Call for Abstracts
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
July 11, 2019	Webinar	Scaling Distributed TensorFlow Training with Intel's nGraph Library on Xeon® Processor Based HPC Infrastructure
July 17-21, 2019	Cookeville, TN	Integrating Parallel and Distributed Computing in Introductory Programming Classes Workshops
July 22-31, 2019	TACC, Austin, TX	MoISSI 2019 Software Summer School
July 23-26, 2019	TACC, Austin, TX	Computational Science in the Cloud
July 28-August 1, 2019	Chicago, IL	PEARC 19
August 5-8, 2019	TACC, Austin, TX	Machine Learning Foundations
August 8-9, 2019	Arlington, VA	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, TX	HPC Leadership
September 24-27, 2019	Geneva, Switzerland	IXPUG Annual Conference 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

More News

Check out the latest Intel® news:

- [Intel's Data-Centric Portfolio Accelerates Convergence of High-Performance Computing and AI Workloads](#)
- [Intel's present and future AI chip business](#)
- [Intel's 'One API' Project Delivers Unified Programming Model Across Diverse Architectures](#)
- [Global Businesses Must Use Analytics and AI to Thrive](#)

© 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](#)