

INTEL® PARALLEL COMPUTING CENTERS



July 2018 Newsletter

Highlights



Optimization Techniques

[Using Roofline Analysis to Analyze, Optimize and Vectorize with Intel® Advisor:](#)

Using Iso3DFD as a case study for detecting wrong memory access patterns or loop carried dependency in your application.

[Using Intel® Xeon® for Multi-node Scaling of TensorFlow* with Horovod*:](#) Time to train a Deep Learning network can be accelerated by as much as 57x (Resnet 50) and 58x (Inception V3) by using Intel-optimized TensorFlow and Horovod MPI for multi-node training.

[AWS Optimizing CNN Inference on CPUs:](#) Gains performance through graph and Tensor co-optimization with separation of hardware-specific schedules.

[Improving VASP* Materials Simulation Performance:](#) Learn how to generate a concise summary of application performance using Intel® VTune™ Amplifier's application performance snapshot, and see an interesting new preview feature in Intel® MPI Library.

[Harp-DAAL for High-Performance Big Data Computing:](#) Explore a new framework that combines Hadoop* and HPC techniques to simultaneously boost productivity and performance by taking advantage of the highly optimized Intel® DAAL.



Case Studies

[Deep Learning Tool for Fast Simulation: \(ISC'18 Research Poster Awardee\)](#) Faster approach to treat Monte Carlo simulation as a replacement for deep learning algorithm training, resulting in achievement of a speedup of at least x100.

[Deep Learning HW Accelerates Fused Discontinuous Galerkin Simulations: \(ISC'18 Research Poster Awardee\)](#) Extended solver to support Intel® Xeon processor with 14 TFLOPS, compared to KNL with speed-ups of up to 1.6x (scenario dependent) and matching dual socket top-bin Xeon performance in case of FP32 execution.

[How-to Guide for Intel® SDK for OpenCL™ Applications:](#) Comprehensive environment for developing and optimizing OpenCL applications on Intel® platforms, and part of an increasingly rich portfolio of Intel tools for heterogeneous programming.

[Deep Learning In Resource-Constrained Environments:](#) Easily deploy CNN models with Python APIs and derive fast inference performance through hardware acceleration, all on an efficient and cost effective solution for embedded applications.

[Advancing AI Performance with Intel® Xeon® Scalable Systems:](#) Harness the power of systems readily available through demonstrations of new methodologies that advance deep learning inference and training throughput on CPUs.

[Powerful New Tools for Exploring and Healing the Human Brain:](#) Diffusion Compartment Imaging optimization work at Harvard Medical School is providing a remarkable 75x performance improvement when running on Intel® Xeon processor.

[Jülich Supercomputer Due to Begin User Operation in the Coming Days:](#) Initial test runs with the LINPACK benchmark, shows cluster module achieved a peak performance of 6.2 petaflop/s.

[Deep Learning and Transformation of Science Exploration:](#) Unique characteristics of Deep Learning and how they constitute a 'game changer' for science exploration.

[Visit Application Speeds Visualization Workloads and Empowers Global Research:](#) Scalability of the newest VisIt codebase with SWR rendering to visualize a terabyte-scale CM1 tornado dataset. Also used by many fields of exploration (astrophysics & climate-related phenomenon).

[Automatic 3D MRI KNEE Damage Classification with 3D CNN Using BigDL on Spark:](#) Center for Digital Health Innovation at UCSF, leveraging new digital health technologies to transform healthcare. Top model accuracy for training and test data showed >90% for no lesions, limited data sets for mild/sever lesions.



**Science
Breakthrough**

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.

Call for Abstracts – Submit NOW!

Share learnings, best practices, and techniques around the benefits you've received in leveraging Intel® architecture—participate in upcoming abstract opportunities by the submission deadline.

Submission Deadline	Events
July 13, 2018	IXPUG Annual Fall Conf. Technical Lecture, Lightning Talk, Tutorial, Poster
July 31, 2018	SC18 Poster
July 31, 2018	SC18 BoF
July 31, 2018	SIAM/ACM Prize in Computational Science & Engineering
Will be announced soon	Supercomputing Frontiers Europe 2019

Global Training Opportunities

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

Date	Location	Event
July 9, 2018	Roma, Italy	Summer School on Parallel Computing
July 10-12, 2018	Lemont, Illinois	IXPUG Software-Defined Visualization Workshop
July 11, 2018	Virtual	Deploy Real –Time Object Detectors on Intel® Movidius™ Neural Compute Stick Webinar
July 13, 2018	Warrington, UK	IPCC Parallel Software Workshop
July 18, 2018	Berkeley, California	Big Data Summit 2018
July 22-28, 2018	East Lansing, Michigan	The 36th Annual International Symposium on Lattice Field Theory
July 26-28, 2018	Shanghai, China	China Visualization 2018
August 6-8, 2018	NYU, New York	2018 New York Scientific Data Summit (NYSDS)
August 6-9, 2018	Minneapolis, Minnesota	Society for Industrial & Applied Mathematics on Life Science
August 10, 2018	Virtual	IXPUG Working Group: Machine Learning at Scale
August 13-16, 2018	Eugene, Oregon	47th International Conf on Parallel Processing (ICPP) 2018
August 19-23, 2018	Boston, Massachusetts	ACS Chemistry for life
September 4-7, 2018	San Francisco, California	O'REILLY Artificial Intelligence Conf.
September 10, 2018	Munich, Germany	Compact Course: Iterative Linear Solvers and Parallelization
September 24-28, 2018	Barcelona, Spain	OpenMP Conf. 2018 & International Workshop OpenMP 2018
September 25-28, 2018	Hillsboro, Oregon	IXPUG Annual Fall Conf. 2018
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, China	HPC China 2018
November 12-14, 2018	Milano, Italy	Introduction to Parallel Computing with MPI and OpenMP
December 3-8, 2018	Montreal, Canada	Neural Information Processing Systems (NIPS) Conf. 2018
Anytime	Virtual	How to identify causes of poor OpenMP parallel performance using the Intel® VTune Amplifier
Anytime	Virtual	Deep Learning and Natural Language Processing Webinar
Anytime	Virtual	Using Roofline Analysis to Analyze, Optimize, & Vectorize Iso3DFD with Intel® Advisor

More News...

Check out these latest news stories:

- [Next-Gen HPC Symposium Spotlights Exascale, AI and Personalized Medicine](#)
- [See How the Convergence of HPC and AI is Changing Supercomputing at ISC 2018](#)
- [How Intel is Powering the Convergence of HPC & AI](#)
- [Bridging the Internal Divides Between IT, HPC, AI That Slow Computing Innovation](#)
- [New eGuide: Practical Considerations for Embracing AI in Your HPC Environment](#)

© 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](#)