

# INTEL® PARALLEL COMPUTING CENTERS



## May 2018 Newsletter

### Highlights



#### Optimization Techniques

**[Intel TensorFlow\\* Installation Guide](#)**: Performance library, using Intel® MKL-DNN, predominately used for machine learning framework in deep learning arena, with different versions compiled to support specific instruction sets.

**[Compiler Prefetching](#)**: Tips on tuning for performance, using the recent changes in the compiler prefetching options, pragmas and prefetching intrinsics.

**[New Approximation Algorithms for Minimum Weighted Edge Cover](#)**: Solving problems faster in serial, taking advantage of parallelism and developing new algorithms with high degrees of concurrency is being addressed by Purdue University.

**[Merging MPI Intercommunicators](#)**: Using the MPI\_Comm\_accept/MPI\_Comm\_connect technique creates an intercommunicator that connects completely separate instances for an MPI job.



#### Case Studies

**[High-throughput Object Detection on Edge Platforms with FPGA](#)**: Software and Deep Neural Network (DNN) architecture level optimizations and tweaks to achieve high throughput with deep learning based object detection applications and FPGAs on edge platforms.

**[Illustrating Success of CPU-based Visualization](#)**: Running simulation and eliminating data movement by many is a requirement for Exascale computing. The CPU's ability to access large amounts of memory is key to realizing trillion triangle per second rendering capability.

**[Intel® DAAL Python\\* API: Gentle Introduction to PyDAAL](#)**: Mpi4py can be easily interfaced with PyDAAL (Intel DAAL's Python API), serialize and de-serialize class of data exchanged between nodes during parallel computation.



**Science  
Breakthrough**

**[University of Stuttgart Advances MegaMol Cross-Platform Visualization Framework](#)**: Enabling image and exploration of tiny-scale complexities like molecular dynamics (Big Data) to handle increasingly-challenging image workloads. Handling terabytes of information with high fidelity rendering in unprecedented three-dimensional details.

**[Google I/O 2018 "Advances in Machine Learning & TensorFlow"](#)**: NERSC scaling training with accuracy on the network for deep learning using TensorFlow on a supercomputer.

**[Deep Learning Foundations to Enable Natural Language Processing \(NLP\)](#)**: The advancement of hardware and software capabilities for deep learning and the availability of foundational elements in an open and flexible stack is particularly suitable for solving a variety of NLP problems.

**[New Solution for a True Hybrid Cloud](#)**: Red Hat and Intel announce new robust, high-performance solution for customers that want to accelerate their path to enterprise containers and hybrid cloud—the Intel® Select Solution for Red Hat OpenShift\* Containers.

**[Durham University Seeks to Unlock Mysteries of the Cosmos](#)**: EAGLE-XL reduces time required for simulation post-processing and ultimately reducing time-to-science for the entire simulation process. Astrophysicists can test many simulations of cosmos based on different theories, ratios of matter and energy.

**[Microsoft\\* Turbocharges AI with Intel® FPGAs. You Can, Too](#)**: Microsoft\* announced a public preview of Azure Machine Learning Hardware Accelerated Models powered by Project Brainwave\*. New AI interface service, uses Intel® Arria® 10 FPGAs, configured as “soft DNN processing units” highly-tuned to the ResNet-50 image recognition and Microsoft calls it “real time AI.”

**[Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling](#)**: Common association between sequence modeling & recurrent networks should be reconsidered & convolutional networks should be regarded as natural starting point for sequence modeling tasks.

## 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](mailto:ipcc.program.office@intel.com) account and include your username in the communication.

## Call for Abstracts –Submit NOW!

Sharing learnings, best practices and techniques around the benefits you've received in leveraging Intel® architecture, By participating in abstract opportunities by the submission deadline.

Submission Deadline	Events
May 18, 2018	NIPS Conf. 2018 <a href="#">Papers</a>
May 20, 2018 (Beijing time)	China Visualization 2018 <a href="#">Papers</a>
June 1, 2018	<a href="#">James H. Wilkinson Prize for Numerical Software</a>
June 5, 2018	SIGGRAPH ASIA 18 <a href="#">Papers</a>
June 11, 2018	NIPS Conf. 2018 <a href="#">Workshops</a>
June 14, 2018	NIPS Conf. 2018 <a href="#">Tutorials</a>
June 15, 2018	OpenMP Conf. 2018: <a href="#">Technical Papers</a>
June 16, 2018	IEEE Visualization Conf. 2018 <a href="#">Short Paper</a> , <a href="#">Posters</a>
June 30, 2018	Amazon Web Services <a href="#">Grants</a>
July 1, 2018	2018 National Annual Conf. on HPC in China <a href="#">Paper</a>
July 13, 2018	IXPUG Fall Conf. <a href="#">Technical Lecture</a> , <a href="#">Lightning Talk</a> , <a href="#">Tutorial</a> , <a href="#">Poster</a>
July 31, 2018	SC18 <a href="#">Poster</a>
July 31, 2018	SC18 <a href="#">BoF</a>
July 31, 2018	<a href="#">SIAM/ACM Prize in Computational Science &amp; Engineering</a>

## Global Training Opportunities

Join us at any of the upcoming events. Please check the links below for additional details.

Date	Location	Event
Anytime	Virtual	<a href="#">Topology and Cache Coherence in Intel Architecture</a>
Anytime	Virtual	<a href="#">Making Machine Learning Accessible</a>
Anytime	Virtual	<a href="#">What is FPGA?</a>
Anytime	Virtual	<a href="#">Acceleration Stack for Intel® Xeon® with FPGA</a>
May 10, 2018	Virtual	<a href="#">High Productivity Languages</a>
May 14-15, 2018	Espoo, Finland	<a href="#">Introduction to Julia</a>
May 21-22, 2018	Oxford, United Kingdom	<a href="#">UK OpenMP Users' Conf.</a>
May 21-25, 2018	Vancouver BC, Canada	<a href="#">32<sup>nd</sup> IEEE International Parallel &amp; Distributed Processing</a>
May 23-24, 2018	San Francisco, California	<a href="#">Intel® AI Developer Conf.</a> (promo code: CMGPCC)
May 31, 2018	Virtual	<a href="#">Intro to Persistent Memory Configuration &amp; Analysis Tools (Intel® 3D XPoint™)</a>
June 21-22, 2018	Munich, German	<a href="#">PRACE Workshop: HPC Code Optimization</a>
June 26-July 4	Espoo, Finland	<a href="#">CSC Summer School in HPC 2018</a>
July 2-4, 2018	Basel, Switzerland	<a href="#">SIGHPC Platform for Adv. Scientific Computing (PASC) Conf.</a>

July 9, 2018	Roma, Italy	<a href="#">Summer School on Parallel Computing</a>
July 26-28, 2018	Shanghai, China	<a href="#">China Visualization 2018</a>
August 6-9, 2018	Minneapolis, Minnesota	<a href="#">Society for Industrial &amp; Applied Mathematics on Life Science</a>
August 19-23, 2018	Boston, Massachusetts	<a href="#">ACS Chemistry for life</a>
September 4-7, 2018	San Francisco, California	<a href="#">O'REILLY Artificial Intelligence Conf.</a>
September 24-28, 2018	Barcelona, Spain	<a href="#">OpenMP Conf. 2018</a> & <a href="#">International Workshop OpenMP 2018</a>
September 25-28, 2018	Hillsboro, Oregon	<a href="#">IXPUG Annual Fall Conf. 2018</a> (open to public)
October 8-11, 2018	London, United Kingdom	<a href="#">O'REILLY Artificial Intelligence Conf.</a>
October 14-16, 2018	Qingdao, China	<a href="#">2018 National Annual Conf. on HPC</a>
October 14-19, 2018	Anaheim, California	<a href="#">Society of Exploration Geophysicists Annual Meeting 2018</a>
November, 2018	Milano, Italy	<a href="#">Introduction to Parallel Computing with MPI and OpenMP</a>
December 3-8, 2018	Montreal, Canada	<a href="#">Neural Information Processing Systems (NIPS) Conf. 2018</a>

## More News...

Check out these latest news stories:

- [The meaning of eleven-point-nine-billion dollars. \(It's not what you think.\)](#)
- [Datacenter Dynamics—Next Generation Hyperscale Hybrid Architectures](#)
- [Cray to build FPGA-Accelerated Supercomputer for Paderborn University](#)
- [Machine Learning in manufacturing: Collect, Contextualize, and Predict](#)
- [Generating Electricity with Bladeless Wind Turbines – One Household at a Time](#)
- [Data Management at NERSC in the Era of Petascale Deep Learning](#)

© 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](#)