








June 2018 Newsletter

Highlights

 	<p>Welcome our newest Intel® PCC members:</p> <p>Nanjing University (gcForest): focus on gcForest application for Deep Forest using Learning & Mining from Data.</p> <p>First Institute of Oceanography (MASNUM): enabling MASNUM application used in ocean numerical simulation community around the world.</p>
 <p>Optimization Techniques</p>	<p>Enhancing Outlier Detection with Intel® DAAL: Detecting outliers or anomalies using Intel® Data Analytics Acceleration Library when running systems equipped with Intel® Xeon® processors.</p> <p>Intel® Software Guard Extensions (Intel® SGX): Software Developer Kit is a collection of APIs, libraries, documentation, sample source code, and tools that allows software developers to create and debug applications enabled for Intel SGX in C and C++.</p> <p>Accelerate Computer Vision from Edge to Cloud: OpenVINO™ Toolkit: Computer vision & deep learning are the big thing driving AI. Streamline building powerful applications with OpenVINO™ (Open Visual Inference & Neural Network Optimization) toolkit.</p> <p>HPC Applications Running on Amazon EC2* Cluster: demonstrating how to scale-out your HPC applications compiled with Intel® Software Development Tools, leveraging Intel® Xeon® Scalable processors on Amazon Elastic Compute Cloud* (Amazon EC2*) environment.</p>

	<p><u>Intel® Parallel Studio XE 2019 Beta</u>: Early access to new features and analysis techniques to optimize HPC, Enterprise and Cloud applications, while contributing to Intel's next gen SW Tools.</p>
 <p>Case Studies</p>	<p><u>Parallel CFD with HiFUN* Solver on Intel® Xeon® Scalable Processor</u>: Increasing the memory channels, enhances single-node performance through large cache availability, higher core density per CPU, higher memory speed, and provides opportunity for larger memory bandwidth.</p> <p><u>Using TensorFlow* Object Detection API</u>: Traffic Light detections at 90% accurate using Microsoft Common Objects in Context pre-trained model called Single Shot Multibox Detector MobileNet for transferring learning.</p> <p><u>Exasol Accelerates In-Memory Data Analytics by up to 114%</u>: Using Intel® Xeon® Platinum processor and implementing optimizations that take full advantage of low level processor features resulted in customers being able to accelerate their data processing from hours to minutes.</p> <p><u>Extending High Accuracy Quantum Chemistry Calculations with HPC</u>: implemented a new method to represent the data contained in 950 GB and reduce the size by a factor of ~7 without losing accuracy.</p> <p><u>Analyzing and Understanding Visual Data</u>: Methods for acquiring, processing, analyzing and understanding images, for high-dimensional data generated by the real world in order to produce numerical or symbolic information.</p>
 <p>Science Breakthrough</p>	<p><u>LIBXSMM Brings Deep-learning "Lessons Learned" to Many HPC Applications</u>: Library targeting Intel processors for small, dense or sparse matrix multiplications, and small convolutions.</p> <p><u>Amazing Inference Performance with Intel® Xeon® Scalable Processors</u>: machine translation which uses recurrent neural networks, Intel Xeon Scalable processor outperforms NVidia* V100* by 4x on <u>AWS Sockeye</u> Neural Machine Translation model with <u>Apache* MXNet*</u> on Intel® Math Kernel Library.</p>

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
June 14, 2018	NIPS Conf. 2018 Tutorials
June 15, 2018	OpenMP Conf. 2018 Technical Papers
June 16, 2018	IEEE Visualization Conf. 2018 Short Paper , Posters
June 30, 2018	Amazon Web Services Grants
July 1, 2018	2018 National Annual Conf. on HPC in China Paper
July 13, 2018	IXPUG 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

Global Training Opportunities

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

Date	Location	Event
June 14, 2018	Virtual	IXPUG Working Group: Using Roofline Analysis
June 16, 2018	Munich, Germany	PRACE Workshop: Intel Manycore Programming Workshop
June 17, 2018	Salt Lake City, Utah	Intel® AI DevJam at CVPR
June 19-21, 2018	Las Vegas, Nevada	HP Discover 2018

June 21-22, 2018	Munich, Germany	PRACE Workshop: HPC Code Optimization
June 24-28, 2018	Frankfurt, Germany	ISC High Performance 2018
June 26-July 4	Espoo, Finland	CSC Summer School in HPC 2018
July 2-4, 2018	Basel, Switzerland	SIGHPC Platform for Adv. Scientific Computing (PASC) Conf.
July 3-21, 2018	Oregon	Oregon Programming Languages Summer School: Parallelism and Concurrency (OPLSS 2018)
July 9, 2018	Roma, Italy	Summer School on Parallel Computing
July 10-12, 2018	Lemont, Illinois	IXPUG Software-Defined Visualization Workshop
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 Kimmel Center, 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 Conference 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
November, 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:

- [Accelerating AI Transformation by Making HPC Easier](#)
- [Accelerating AI Transformation by Making HPC Easier\(Part 2\)](#)
- [Extending High Accuracy Quantum Chemistry Calculations with HPC](#)
- [Performance Insights Using the Intel Advisor Python API](#)
- [How Google One Actually Stacks Up to the Cloud Competition](#)
- [Intel launches Xeon with integrated FPGA](#)

© 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](#)