January 2019 Newsletter

Highlights

 **Optimization Techniques**	[Precision and Recall for Time Series]: research at the most fundamental levels to change the way that scientists and engineers perform anomaly detection evaluation and training. [Scalable Methods for 8-bit Training of Neural Networks]: by reducing precision requirements from 32 floating point to 8-bit precision, get a reduction in memory and power due to the increased efficiency of computing at 8-bit precision. [Deep Defense: Training DNNs with Improved Adversarial Robustness]: perturbations for fooling a DNN model can be 1000x smaller in magnitude when compared with real images, making these perturbations imperceptible to the naked eye. [Intel memory hashing update/apps affected]: Intel processors use a proprietary hashing algorithm to distribute physical addresses pseudo-randomly across the enabled L3 “slices” on the chip. Understanding these mappings is critical to modeling throughput on the chip, and for optimizing tree-based shared-memory synchronization algorithms. [Norm matters: efficient and accurate normalization schemes in deep networks]: present a novel view on the purpose and function of normalization methods and weight-decay, as tools to decouple weights' norm from the underlying optimized objective.

 **Case Studies**	**Graph Neural Networks for IceCube Signal Classification**: the application of GNNs for the IceCube signal classification task, quantitatively, this architecture results in a 3x improvement in signal-to-noise ratio, and is capable of detecting on the order of 6.3x more signal events than traditional physics-based baselines. ("Best Paper" award at IEEE ICMLA) **OpenMP in VASP: Threading and SIMD**: the high-level vector coding scheme applied to VASP's general gradient approximation routine gives up 9× performance gain on AVX512 platforms with the Intel compiler. **Topology Can Help Us Find Patterns in Weather**: many applications are being recast to use topology, For instance, looking for weather and climate patterns using Intel Data Analytics Acceleration Library for the SVM(Support Vector Machine), good scaling was achieved. **Supercomputers Help Uncover the Mysteries of Black Hole Tidal Disruption Events**: optimize Cosmos++ code to take advantage of the Intel Xeon Scalable and Xeon Phi processors on the Stampede2, have measured about a 50 percent improvement in overall performance. **Speeding Up a Big Data Platform**:after optimizing the algorithms with Intel® Data Analytics Acceleration Library (Intel® DAAL) and Intel Math Kernel Library (Intel® MKL) performance improved by an average of 3x, with peak performance gains of 14x.
 **Scientific Breakthrough**	**Intel® Software Development Tools Optimize Deep Learning Performance for Healthcare Imaging**: optimizing the trained GE solution with Intel's Deep Learning Deployment Toolkit and Intel MKL-DNN improved throughput an average of 14 times over a baseline version of the solution and exceeded GE's throughput goals by almost six times. **Outbursts of luminous blue variable stars from variations in the helium opacity**: using ParaView knows the cause of the explosions of these luminous blue variables gives scientists a more complete picture of the life and death of the biggest stars in the universe.

	[How AI is helping Scientists with the large Hadron Collider]: describes how AI is being used in design of experiments for the Large Hadron Collider. [Accelerate Machine Learning on Google Cloud with Intel® Xeon® Scalable processors]: shows how machine learning workflows can benefit from running on optimized Intel® Architectures. [Exploring the Frontiers of Chemistry with HPC]: access to TACC Stampede 2, the chemistry researchers can enlist the help of supercomputers to run simulations of the components and see how they react.

Intel® Parallel Computing Centers Invited Talk Series

The Intel® PCC Invited Talk Series are presented by global partners as they share experiences in using Intel architecture for scientific breakthrough. The presenters will share optimization techniques, best practices and results. This series is intended for students, educators, developers/programmers, scientist, data analyst, system administrators, etc. working to maximizing software efficiency using Intel technology. Please come to join January 2019 Invited talk "A Hybrid MPI+Threads Approach to Group Finding Using Union-Find".

Testing Your Code on Intel® Architecture

Intel® Xeon® Scalable Platform Access: We encourage testing applications using various configurations of Intel® architecture (Intel® Xeon Scalable processors, Intel® Omni-Path, etc. Click HERE to test your optimized application at scale using TAC, 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.

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
February 1, 2019	2019 Molecular Science Software Workshops Call for Proposals
February 6, 2019	ISC19 Tutorials & Project Posters
February 13, 2019	ISC19 PhD Forum & Workshops(regular)
February 20, 2019	ISC19 BoF Session
February 27, 2019	ISC19 Research Posters

Global Event & Training Opportunities

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

Date	Location	Event
January 14 -16, 2019	Guangzhou, China	HPC Asia 2019
January 23, 2019	Munich, German	EMEA Intel® AI DevCon 2019
February 27 - March 2, 2019	Minneapolis, USA	SIGCSE 2019
March 12 - 15, 2019	Warsaw, Poland	Supercomputing Frontiers Europe 2019
March 31 – April 4, 2019	Orlando, FL	American Chemical Society Spring 2019 National Meeting & Exposition
April 15-18, 2019	New York, US	The Artifical Intelligence Conference in New York
June 16-20, 2019	Frankfurt, Germany	ISC 2019
July 28-August 1, 2019	Chicago, IL	PEARC 19
August 26-30, 2019	Gottingen, Germany	EURO-PAR 2019
September 15 – 20, 2019	San Antonio, TX	SEG 19
November 17 -20, 2019	Brisbane, Australia	SIGGRAPH Asia 2019
Anytime	Webinar	What Intel® Processor Graphics GEN9 Unlocks in OpenCL*
Anytime	Webinar	All About Persistent Memory Flushing

Anytime	Webinar	Configure and Manage Persistent Memory DIMMs
Anytime	Webinar	Persistent Memory Programming Today and in the Future

More News...

Check out the latest Intel® news:

- **Intel's Framework for Automated Vehicle Safety Standard Gaining Global Acceptance**
- **Intel Showcases New Technology for Next Era of Computing**
- **Mobileye, Beijing Public Transport Corp. and Beijing Beytai Collaborate to Bring Autonomy to China's Public Transportation**