March 2019 Newsletter

Highlights

 **Optimization Techniques**	**Improving Tensorflow* Inference Performance on Intel® Xeon® Processors**: recent testing by Intel shows up to a 60% performance improvement in latency at batch size one on a variety of popular models. **Expressing Parallelism in C++ with Threading Building Blocks**: with Intel TBB, you can parallelize computationally intensive work—including heterogeneous computing—in a fast, portable, and scalable way without having to delve in to the low-level details of threading. **Intel® Rendering Framework Using Software-Defined Visualization**: new multi- and many-core CPUs are your true ally for large rendering tasks, benchmarks show a100x increase in rendering performance. **OPENMP*: The Once and Future API**: as an open standard, allows developers to create portable code that withstands the continuous evolution of hardware revisions. **Parallelism in Python***: the base Python language remains single-threaded, find out how this language remain performant in the age of parallelism. **Optimizing High Performance Distributed Memory Parallel Hash Tables for DNA k-mer Counting**: this implementation completed index construction and query on an approximately 1 TB human genome dataset in just 11.8 seconds and 5.8 seconds, demonstrating speedups of 2.06x and 3.7x, respectively, over the previous state-of-the-art distributed memory k-mer counter.

 **Case Studies**	**Accelerating AIREBO: Navigating the Journey from Legacy to High-Performance Code**: present the journey from the C++ port of a previous Fortran code to performance-portable, KNC-hybrid, vectorized, scalable, and optimized code supporting full and reduced precision. **Running Tensorflow at Petascale and Beyond**: CosmoFlow, a cosmological research collaboration between Intel, NERSC, and Cray is the first project using Tensorflow* for a big science application to train a neural network on a supercomputer. Using 8,192 nodes on Cori, sustained performance during the training phase was 3.5 petaflops. **Lessons Learned from Optimizing Kernels for Adaptive Aggregation Multi-grid Solvers in Lattice QCD**: describe the implementation strategies for the restrictor including vectorization approach and the approaches they have investigated for parallelization with OpenMP. **S-Store: A Streaming NewSQL System for Big Velocity Applications**: a single, scalable data management system that can support both stream and transaction processing at the same time. **Extracting Semantic Relations using External Knowledge Resource with NLP Architect**: The Intel AI lab recently released a set of new features for NLP Architect, a Python* library for exploring the state-of-the-art deep learning topologies and techniques for NLP. **Deep Learning at Scale Using Cray Distributed Training**: for scientific applications like CosmoFlow and HEP-CNN where datasets can be significantly larger than ImageNet, scaling on the Cori Supercomputer is an especially relevant capability. How the Cray PE ML Plugin for distributed data-parallel training of DNNs can be used across a range of science domains is discussed on Github. **CHAOS: a parallelization scheme for training convolutional neural networks on Intel Xeon Phi**: design and implementation of CHAOS parallelization scheme for training CNNs on the Intel® Xeon Phi™ processors show speedups of up to 103× compared to the execution on one thread of the Xeon Phi. **Parallel K-means Clustering of Geospatial Data Sets Using Manycore CPU Architetures**: enable Parallel K to effectively use parallel machines based on Intel® architectures and employing significant single instruction, multiple data (SIMD) parallelism.
 **Scientific Breakthrough**	**How Intel Powers Dell EMC Ready Solutions for AI**: describes how Dell EMC is making enterprise AI implementation fast and easy with its Dell EMC Ready Solutions for AI. **JD Speeds up Image Analysis by Replacing GPUs with CPUs**: learn how Chinese cloud services provider JD processed 300,000 images in 162 seconds instead of

	2,800 seconds with Intel® C++ Compiler and Intel® Integrated Performance Primitives. **AI-Driven Medical Imaging Powered by Intel and Philips**: two healthcare use cases for deep learning inference models showed improved results using the OpenVINO™ toolkit and other optimizations, along with efficient multi-core processing from Intel Xeon Scalable processors, a speed improvement of 188.1x for the bone-age-prediction model, and a 37.7x speed improvement for the lung-segmentation model over the baseline measurements. **Innovative Use of HPC in the Cloud for AI, CFD, & LifeScience**: Ebru Taylak from the UberCloud provide recent examples of innovative use cases of HPC in the Cloud. **HPE & Intel Omni-Path Architecture Help Researchers Clear Path Toward Safe, Next-Gen Nuclear Fusion**: with Intel OPA, having this significant amount of injection bandwidth will allow deep learning training and inference workloads to scale well.

Intel® Parallel Computing Centers Invited Talk Series

The Intel® PCC Invited Talk Series are presented by global partners sharing optimization techniques, best practices and results. Please join register and join us March 28th 8:00-8:30am (Pacific Standard Time) to learn about Carnegie Mellon University's new Stochastic Local Search optimization method **Efficient Machine Learning Model Optimization by Stochastic Local Search and Thompson Sampling**.

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
April 14, 2019	ISC19 IXPUG Workshop: Using FPGAs to Accelerate HPC & Data Analytics on Intel-Based Systems
April 16, 2019	O'REILLY + Intel Artificial Intelligence Conference, London Call for Speakers

Global Event & Training Opportunities

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

Date	Location	Event
March 12 - 15, 2019	Warsaw, Poland	Supercomputing Frontiers Europe 2019
March 14, 2019	Webinar	Scalable and Flexible Distributed Rendering with OSPRay's Distributed API and FrameBuffer
March 31 – April 4, 2019	Orlando, FL	American Chemical Society Spring 2019 National Meeting & Exposition
April 11, 2019	Webinar	A Study of SIMD Vectorization for Matrix-Free Finite Element Method
April 15-18, 2019	New York, US	O'REILLY + Intel Artificial Intelligence Conference
April 30 – May 3, 2019	TACC Institute	Applied Parallel Programming
May 7-8, 2019	TACC Institute	Advanced Computing Foundations
May 6-9, 2019	New Orleans, US	ICLR 2019
June 10-13, 2019	TACC Institute	Scientific Visualization
June 10-15, 2019	Long Beach, CA	ICML 2019
June 15-21, 2019	Long Beach, CA	CVPR 2019
June 16-20, 2019	Frankfurt, Germany	ISC 2019
June 18-19, 2019	Beijing, PRC	O'REILLY + Intel Artificial Intelligence Conference
June 25-28, 2019	TACC Institute	Designing and Administering Large-scale Systems
July 9-12, 2019	TACC Institute	Workflows and Reproducibility in Scientific Computing
July 23-26, 2019	TACC Institute	Computational Science in the Cloud
July 28-August 1, 2019	Chicago, IL	PEARC 19
August 5-8, 2019	TACC Institute	Machine Learning Foundations
August 26-30, 2019	Gottingen, Germany	EURO-PAR 2019
September 9 -12, 2019	San Jose, CA	O'REILLY +AI Artificial Intelligence Conference
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 Institute	HPC Leadership

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:

- **Architected for HPC, AI, IaaS Leadership Performance**
- **Intel's MESO transistor promises vast leap in AI processing power**
- **Intel at MWC 2019: 5G & Edge Computing Remain Focal Points**