August 2018 Newsletter

Highlights

 **Optimization Techniques**	**Scalable CPU-Based Visualization (SDVis) Enables Interactive Photoreal Visualization**: We are entering the era, based on the data size, where the scalability and constant runtime of SDVis wins over GPUs for visualization. **Solving Real-World Machine Learning Problems**: Selecting the correct algorithms based on two real-world machine learning problems that were taken from the well-known Kaggle platform used for predictive modeling and from analytics competitions where data miners compete to produce the best models. **Boosting Java* Performance in Big Data Applications**: BigDL**,** an open source distributed deep learning framework for Apache Spark*. With BigDL, users can write their DL applications as standard Spark programs, which can directly run on top of existing Spark or Hadoop* clusters. **Scaling read aligners to hundreds of thread on general-purpose processors**: Running exactly the same executables on all three systems(Broadwell, Skylake and KNL), achieved excellent thread scaling to hundreds of threads, and have officially released versions of Bowtie and Bowtie 2 tools that implement these ideas. **C++ Transactions for Persistent Memory Programming**: provides a detailed look at the use of transactions in persistent memory programming using PMDK C++ bindings.

 **Case Studies**	[In Situ Analysis and Visualization with SENSEI]: SENSEI enables connection of simulation data sources to visualization and analysis back ends through a data model and API. Simulations get run-time interchangeability of analysis/vis codes and analysis/vis codes can consume data from any simulation. [Optimised Data Decomposition For Reduced Communication Costs]: This new data layout, only active for part of the runtime, can produce a 5x improved performance for short simulations and 2x for more realistic simulations when using large numbers of processes and scaling up on large parallel computers. [OpenMP threading and vectorization of MPI Finite element code Elmer]:Utilizing mesh coloring, it is possible to have optimal global linear system assembly with OpenMP threading. Moreover, non-preconditioned Krylov iterative methods are marginally faster as threaded implementations. [Performance optimization for modern many core architecutres using PSyclone embedded-DSL]: by selecting the best transformation possible when architectural and system parameters are given, explore the performance characteristics of LFRic/PSylclone on many-core architectures, currently x2 speedup achieved SKL vs KNL with MCDRAM. [Devito - Automated High-Performance Finite-Difference for Geophysical Exploration]: the latest developments and results in Devito, an open-source framework for solving partial differential equations from symbolic problem definitions by the finite difference method. Performance results are shown on state-of-the-art Intel architectures (Skylake, KNL) for production-level seismic operators.
 **Science Breakthrough**	[Research team improves code to benefit industrial engineers]:The Ohio Supercomputer Center's(OSC) modernized WARP3D, an open-source code used by engineers to optimize the robotic welding process for heavy equipment. [Canada's Most Powerful Research Supercomputer Simulates Life of a Star]: Canada's fastest supercomputer, Niagara based upon Intel Xeon Gold processors and located at the University of Toronto, enables aster seismology research. [Deploy Anywhere, Anytime, Any-scale Visualization with Intel® Select Solutions]: The new Intel® Select Solution for Professional Visualization helps researchers with simulation and machine learning workloads at massive-scale while helping them make sense of increasingly large data sets and complex results. [A Little Mxing Goes a Long Way: Breakthroughs in Simulating Tides' Impact on Global Ocean Circulation]: Using the Intel Xeon Scalable processor-based Niagara supercomputer at

	the University of Toronto to link how small-scale forces affect the global circulation of our oceans. **Transfer Learning in Natural Language Processing**: Transfer Learning was limited to computer vision up till now, but recent work shows that the impact can be extended almost everywhere.

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.

Global Training Opportunities

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

Date	Location	Event
August 13-16, 2018	Eugene, Oregon	47th International Conf on Parallel Processing (ICPP) 2018
August 19-23, 2018	Boston, Massachusetts	ACS Chemistry for life
August 28-31, 2018	Nanjing, China	PRICAI 2018
August 26-30,2018	Las Vegas, NV	VM World 2018
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 12, 2018	Cheshire, UK	Accelerating codes on Intel processors
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, PRC	HPC China 2018
November 12-14, 2018	Milano, Italy	Introduction to Parallel Computing with MPI and OpenMP
November 14-16, 2018	Beijing, PRC	Asian Conference on Machine Learning
December 3-8, 2018	Montreal, Canada	Neural Information Processing Systems (NIPS) Conf. 2018
March 12-15, 2019	Warsaw, Poland	Supercomputing Frontiers Europe 2019
Anytime	Virtual	Introduction to TensorFlow* with Intel® Optimization
Anytime	Virtual	Getting Started with the Intel® AI DevCloud

August 10, 2018	Virtual	IXPUG Working Group: Machine Learning at Scale
Anytime	Virtual	How to identify causes of poor OpenMP parallel performance using the Intel® VTune Amplifier
Anytime	Virtual	Deep Learning and Natural Language Processing
Anytime	Virtual	Using Roofline Analysis to Analyze, Optimize, & Vectorize Iso3DFD with Intel® Advisor

More News...

Check out these latest news stories:

- HPC for Everyone Comes to Manufacturing
- NERSC Hosts First 'Big Data Summit'
- Growing HPC Adoption Among Manufacturers
- Three Mid-Year's Resolutions: Learn Some HPC, AI, and Cloud

© 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