



# SATURDAY NOVEMBER 11

	Parallel Programming I Tower Court A	Parallel Programming II Tower Court B	Artificial Intelligence I Tower Court C	Artificial Intelligence II Tower Court D	Systems Silver	High Productivity Languages Gold	Hands On Tutorials I Denver	Hands On Tutorials II Spruce	Hands On Tutorials III Century	Visualization Development Colorado
7:00 - 9:00 am	<b>Registration and Breakfast</b>									
9:00 - 10:00 am	<b>Welcome and Opening Keynote: "The Intertwined Futures of HPC &amp; AI" - Grand Ballroom</b>									
10:00 - 10:15 am	<b>Break</b>									
10:15 - 10:45 am	Cataloging the Visible Universe through Bayesian Inference at Petascale in Julia <i>Julia Computing Inc., UC Berkely</i>	I/O Performance at Scale and Storage Class Memory <i>EPC, The University of Edinburgh</i>	Assessment of AVX512 to Speed up Crash Simulations <i>Altair</i>	Enabling the Future of Artificial Intelligence <i>Intel Corporation</i>	<i>Intel Corporation</i>	Stalking the Interactive Terabyte with R <i>Oak Ridge National Laboratory</i>	Introduction to Intel® 3D XPoint™ Large Capacity Persistent Memory and its Open Source API <i>Intel Corporation</i>	Parallelizing Python Applications with PyCOMPSs <i>Barcelona Supercomputing Center (BSC)</i>	Optimized Symbolic Finite Difference Computation with Devito <i>Imperial College London</i>	In Situ Software Defined Visualization Hands-On Lab Talk: "Using OSPRay API with Data Distributed Applications" <i>Argonne National Laboratory, Intel Corporation</i>
10:45 - 11:15 am	PSyclone - Portability & Scalability for a Large Intel® Xeon Phi™ (Code Named Knights Landing) System <i>Hartree Centre</i>	High-performance and Scalable MPI+X Library for Emerging HPC Clusters & Cloud Platforms <i>The Ohio State University</i>	Emerging AI Technologies on Intel® Client Platforms <i>Intel Corporation</i>	Optimizing TensorFlow for Intel® Xeon® and Xeon Phi™ Processors <i>HPE</i>	Containers for HPC Experience on Intel® Xeon Phi™ <i>Intel Corporation</i>	High Performance Computing with Python and Anaconda <i>Anaconda</i>				
11:15 - 11:45 am	ParaView and VTK w/OSPRay and OpenSWR <i>Kitware Inc.</i>	A Graph Model for Large Scale Simulation with Task Parallel Approach on HPC <i>Mississippi State University</i>	Low Precision Neural Networks with FPGAs <i>Intel Corporation</i>	Scalable Deep Learning with BigDL on Urika-XC <i>Cray Inc.</i>	Intel® Omni-Path™ SW Architecture Overview <i>Intel Corporation</i>	Harnessing the Power of High Performance Computing for R <i>Stanford University</i>				
11:45 - 12:15 pm	<b>Lunch Served</b>									
12:15 - 12:50 pm	<b>Industry Panel: "Taking Parallelism to the Next Level" - Grand Ballroom</b>									
12:50 - 1:00 pm	<b>Break</b>									
1:00 - 1:30 pm	SIMD Vectorization with High-Level Vectors <i>Zuse Institute Berlin</i>	OpenCL in Scientific High Performance Computing <i>Zuse Institute Berlin</i>	Optimization of Fluent on Latest Intel Processors <i>Ansys Inc.</i>	Deep Learning at 15PF: Supervised and Semi-supervised Classification for Scientific Data <i>Lawrence Berkeley National Laboratory, Intel Corporation</i>	Many-Cores for the Masses Part 2 - A Year With the Cori System at NERSC <i>NERSC</i>	Accelerating Scientific Python with Intel® Optimizations <i>Intel Corporation</i>	Leveraging System Diagnostics in an Application with Intel® Cluster Checker <i>Intel Corporation</i>	Tuning for the Intel® Xeon® Scalable Processor (Code Named Skylake) <i>NESC-ID/IST, University of Lisbon, Intel Corporation</i>	Accelerating Algorithms on FPGAs using OpenCL <i>Intel Corporation</i>	In Situ Software Defined Visualization Hands-On Lab Talk: "CSP Project featuring VTK-m" <i>Kitware Inc, Intel Corporation</i>

# SATURDAY NOVEMBER 11

	Parallel Programming I Tower Court A	Parallel Programming II Tower Court B	Artificial Intelligence I Tower Court C	Artificial Intelligence II Tower Court D	Systems Silver	High Productivity Languages Gold	Hands On Tutorials I Denver	Hands On Tutorials II Spruce	Hands On Tutorials III Century	Visualization Development Colorado
1:30 - 2:00 pm	Optimizations of Implicit Unstructured Aerodynamics Computations  <i>King Abdullah University of Science and Technology</i>	Accelerating Cryo-EM Reconstruction in RELION with x86 SIMD Instructions  <i>Stockholm University, Intel Corporation</i>	Software Engineering Challenges of Bringing an Aging Code Base Into the Modern Era  <i>Convergent Science</i>	Predicting Disease with Deep Learning on TensorFlow from Genetic Variant Data  <i>Scripps Translational Science Institute, Intel Corporation</i>	Administration of QoS for Lustre  <i>Johannes Gutenberg University Mainz</i>	Parallel Computing with Python and the Numba Compiler  <i>Anaconda</i>	(continued) Leveraging System Diagnostics in an Application with Intel® Cluster Checker  <i>Intel Corporation</i>	(continued) Tuning for the Intel® Xeon® Scalable Processor (Code Named Skylake)  <i>NESC-ID/IST, University of Lisbon, Intel Corporation</i>	(continued) Accelerating Algorithms on FPGAs using OpenCL  <i>Intel Corporation</i>	(continued) In Situ Software Defined Visualization Hands-On Lab Talk: "CSP Project featuring VTK-m"  <i>Kitware Inc, Intel Corporation</i>
2:00 - 2:30 pm	RckT: Scalable Physically Accurate Spectral Rendering in OSPRay  <i>SURVICE Engineering</i>	Practical Affinity for Threaded MPI Applications  <i>Intel Corporation</i>	Fast Deep Learning with Caffe and Tensorflow on Multi-core CPUs  <i>OnSpecta, Inc.</i>	Intel Nervana Graph: A Universal Deep Learning Compiler  <i>Intel Corporation</i>	Designing an Optimal Deep Learning Solution at Dell  <i>Dell</i>	Managing Data Science at Scale  <i>Domino Data Lab</i>				
2:30 - 2:45 pm	<b>Break</b>									
2:45 - 3:15 pm	Energy-efficient, Scalable Computing of Extremely Large Electronic Structures with Intel® Xeon Phi™ (Code Named Knights Landing) Processors <i>Intel Parallel Computing Center, Korea Institute of Science and Technology</i>	SWIFT: Task-based Calculation + Task-based MPI + Task-based I/O = Maximal Performance <i>ICC, Durham University</i>	Machine Learning with HPC: Optimizing for Big Data, Accuracy, and Response Time <i>Carnegie Mellon University</i>	Analysis of Effective Connectivity Using Python <i>Florida International University</i>	Hyper-acceleration of Apache Spark with Intel® Arria® 10 FPGAs <i>Bigstream</i>	Accelerate Big Data Processing (Hadoop, Spark, Memcached & TensorFlow) with HPC Technologies <i>The Ohio University</i>	Julia: Petascale Made Simple <i>Julia Computing Inc.</i>	Parallel Performance Evaluation Using the TAU Performance System <i>University of Oregon</i>	Accelerating Algorithms on FPGAs Using OpenCL <i>Intel Corporation</i>	In Situ Software Defined Visualization Hands-On Lab Talk: "Modernizing VisIT with VTK and SDVis" <i>Intel Corporation</i>
3:15 - 3:45 pm	MoonRay: Faster Path Tracing through Vectorization  <i>DreamWorks Animation</i>	TBB Flow Graph: Expressing and Analyzing Dependencies in Your C++ Application  <i>Intel Corporation</i>	From 'Hello World' to Scalable Deep and Machine Learning to 13 PF/s Average Sustained Performance  <i>TechEnablement</i>	Accelerated Characterization of Neural Circuits of the Brain  <i>Boston Children's Hospital &amp; Harvard Medical School</i>	Benefits of Intel® Memory Drive Technology for Scientific Applications  <i>Lomonosov Moscow State University, Argonne National Laboratory, Institute of Computational Mathematics and Mathematical Geophysics SB RAS &amp; Intel Corporation</i>	Mixed Language Debugging (Python and C/C++ with TotalView)  <i>Rogue Wave Software</i>				
3:45 - 4:15 pm	Interactive In Situ Visualization of LAMMPS Simulations with OSPRay  <i>Argonne National Laboratory</i>	High-performance Threaded Intel MPI Library on Intel® Omni-Path Architecture: Under the Hood  <i>Intel Corporation</i>	An AI-enabled Data Refinery for Satellite Imagery at Global Scale  <i>Descartes Labs</i>	Experiences of Scaling Tensorflow on up to 512 Nodes on CORI Supercomputer  <i>Intel Corporation</i>	Hybrid HPC and HPC Iaas Unbercloud  <i>HPE Intel Corporation</i>	Python Applications in the NERSC Exascale Science Applications Program for Data  <i>National Energy Research Scientific Computing Center (NERSC)</i>				
4:15 - 4:30 pm	<b>Break</b>									
4:30 - 5:30 pm	<b>Plenary Session "Gravitational Waves: The Role of Computing in Opening a New Field of Astronomy" – Grand Ballroom</b>									
5:30 - 7:30 pm	<b>Networking: Poster Sessions, Meet the Experts, and More</b>									

# SUNDAY NOVEMBER 12

	Parallel Programming I	Parallel Programming II	Artificial Intelligence I	Artificial Intelligence II	Systems	Enterprise	Hands On Tutorials I	Hands On Tutorials II	Hands On Tutorials III	Visualization Development
	Tower Court A	Tower Court B	Tower Court C	Tower Court D	Silver	Gold	Denver	Spruce	Century	Colorado
7:00 - 8:30 am	<b>Registration and Breakfast</b>									
8:30 - 9:30 am	<b>Plenary Session "Deep Learning for Science" – Grand Ballroom</b>									
9:30 - 9:45 am	<b>Break</b>									
9:45 - 10:15 am	High-fidelity Rendering of Bioscience Data Using OspaRay and VR <i>Texas Advanced Computing Center</i>	Parallel Programming with Distributed Datastructures in a Global Address Space <i>Ludwig-Maximilian-University (LMU) Munich</i>	Enabling DL Frameworks with MPI <i>Intel Corporation</i>	Achieving Deep Learning Training in Less than 40 Minutes on ImageNet-1K with Intel® Xeon® and Xeon Phi™ Architectures <i>SURFsara B.V., Intel Corporation</i>	Experiences and Evolution of the Stampede Supercomputer <i>Texas Advanced Computing Center, University of Texas</i>	HyperLoom: A Platform for Defining and Executing Scientific Pipelines in Distributed Environments <i>IT4Innovations, National Supercomputing Center</i>	Harp-DAAL: A High Performance Machine Learning Framework for HPC-Cloud <i>Indiana University</i>	Discover All Benefits of Intel® Xeon® Scalable Processor (Skylake) Inside Google Cloud Platform (GCP) <i>Google Cloud</i>	Singularity, Containers for Science <i>SingularityWare, LLC.</i>	In Situ Software Defined Visualization Hands-On Lab Talk: "In Situ Open Lab Recap and Insight to the Future of Visualization" <i>Intel Corporation</i>
10:15 - 10:45 am	Challenges and Opportunities in using Software-defined Visualization in MegaMol <i>VISUS University of Stuttgart</i>	Capability Models for Many-core Memory Systems: A Case-Study with Intel® Xeon Phi™ <i>ETH Zurich</i>	Performance Evaluation of FPGA-based Bit-Scaled CNN Architecture <i>San Jose State University &amp; Mando Corporation</i>	Machine Learning for Boosted Jet Classification in High Energy Physics <i>Center for Scientific Computing (NCC/UNESP)</i>	DAOS: An Architecture for Extreme Scale Storage <i>Intel Corporation</i>	Java for HPC - A Story of Performance <i>Intel Corporation</i>		Deep Learning for Fast Particle Simulation Using Neon <i>CERN</i>		
10:45 - 11:00 am	<b>Break</b>									
11:00 - 11:30 am	Volume Visualization of Massive Seismic and Rock Physics Data Using OSPRay in HPC <i>BP</i>	Comparison and Analysis of Parallel Tasking Performance for an Irregular Application <i>University of Bristol</i>	Machine Learning on the Latest Intel® Architecture - Three Case Studies <i>Bayncore</i>	High-performance Machine Learning for Weather Prediction Applications <i>KAUST</i>	Simplify System Software Stack Development and Maintenance <i>Intel Corporation</i>	Accelerating Workloads Using the Acceleration Stack for Intel® Xeon® CPU with FPGAs <i>Intel Corporation</i>	(continued) Harp-DAAL: A High Performance Machine Learning Framework for HPC-Cloud <i>Indiana University</i>	(continued) Deep Learning for Fast Particle Simulation Using Neon <i>CERN</i>	(continued) Singularity, Containers for Science <i>SingularityWare, LLC.</i>	(continued) In Situ Software Defined Visualization Hands-On Lab Talk: "In Situ Open Lab Recap and Insight to the Future of Visualization" <i>Intel Corporation</i>
11:30 - 12:00 pm	Rendering in Blender Cycles Using Intel® Xeon Phi™ (Code Named Knights Landing) <i>IT4Innovations, VSB – Technical University of Ostrava</i>	Towards Exascale Computing with Fortran 2015 on Intel® Many-core Processors <i>Sourcery Institute, National Center for Atmospheric Research</i>	Taming the Many-core Beast: View from Industry on Impact of New Architectures on Algorithm Design <i>Schlumberger</i>	Best Practices for On-demand HPC in Enterprises <i>SVSG</i>	A Guided Tour of PBS Professional's Basic and Advanced Configurations with OpenHPC <i>Altair</i>	Reconfigurable Accelerator Platform with FPGAs <i>Intel Corporation</i>				
12:00 - 12:15 pm	<b>Break</b>									
12:15 - 1:00 pm	<b>Plenary Session and Awards Ceremony "Leading the Evolution of Compute: Neuromorphic and Quantum Computing" – Grand Ballroom</b>									
1:00 - 3:00 pm	<b>Networking and Light Snacks</b>									