



# Intel Labs SC'12 Papers Overview

## Follow Intel Labs here:

[www.twitter.com/intellabs](http://www.twitter.com/intellabs) #intellabs

<http://blogs.intel.com/research>

<http://www.facebook.com/IntelLabs>

## Intel Labs Supercomputing 2012 Papers

Supercomputing 2012 (SC12), Salt Lake City – Intel Corporation is delivering a host of papers, presentations, panel discussions and demonstrations at SC12, the international conference for high performance computing, networking, storage and analysis. A highlight paper discloses new details about algorithms which accelerate the correlation computation on massive amounts of telescopic data to help answer questions about the origin and fate of the universe. Other Intel papers describe innovative and efficient new algorithms delivering significant performance speedups on Intel® Xeon and Intel® Xeon Phi™ based systems for compute-intensive applications such as, matchmaking within big data to find meaningful relationships; Synthetic Aperture Radar computation to reveal the Earth; a new framework to double “Fast Fourier Transform” performance; and multigrid-based iterative solvers for very large linear systems. Below are highlights covering Intel’s scheduled presence at the event.

### **Big Data Astronomy (Gordon Bell Nominee, updated with results using Intel® Xeon Phi™)**

Intel is helping scientists use big data to answer big questions. Scientists are using supercomputers to correlate massive data on billions of stars and galaxies to find dark matter and better understand the origin and fate of the universe. In collaboration with Lawrence Berkeley National Laboratory, algorithms from Intel Labs were used to correlate 1.7 billion objects in just over five hours on Lawrence Livermore National Lab’s Zin supercomputer, a Petascale-class machine with 1600-nodes each containing two Intel® Xeon® processors E5-2670. This research improved speed by 35x and cost efficiency 11x by scaling these calculations better across thousands of cores. It also shows a path to real-time processing of telescopic data and for Exascale-class computation on even larger datasets. More recently, on a 1.8 PetaFlop section of the Texas Advanced Computing Center’s Stampede cluster, using Intel® Xeon Phi™ Coprocessors SE10, Intel Labs achieved a further speedup in run-time of 3.2 X on each node in comparison to the results above.

Paper and contributors details:

*Billion-Particle SIMD-Friendly Two-Point Correlation on Large-Scale HPC Cluster System.*

*Authors: Jatin Chhugani, Changkyu Kim, Hemant Shukla, Jongsoo Park, Pradeep Dubey, John Shalf, Horst D. Simon*

*Session: Tuesday, November 13, 1:30 p.m.-2:00 p.m. Room 115E*

### **Matchmaking within Big Data**

The ability to find meaningful relationships among millions or billions of data will lead to entirely new business insights, scientific trends, and social connections. The web-like connections among data are represented by structures called ‘graphs’ and for big data they span in clusters. Cluster graph searches are limited by the myriad node-node data transfers required in traversing the complex web. Intel Labs has demonstrated algorithms to mitigate the data transfers, increasing performance by 6.6x and energy efficiency by 8x using the Graph 500 Benchmark on an Intel® Xeon cluster.

Paper and contributors details:

*Large-Scale Energy-Efficient Graph Traversal - A Path to Efficient Data-Intensive Supercomputing.*

*Authors: Nadathur Satish, Changkyu Kim, Jatin Chhugani, Pradeep Dubey*

*Session: Tuesday, November 13, 11:30 a.m.-12:00 p.m. Room 255-EF*

### **Revealing the Earth (Intel® Xeon Phi™, Best Paper Nominee)**

Earth maps based on satellite images have led to many apps such as Google Earth, but require clear, well lit conditions. Synthetic Aperture Radar (SAR) can image the earth at night, through clouds and trees and provides information on surface materials. This has military as well as civilian applications – including mineral exploration, environmental monitoring, crop monitoring and navigation. SAR works by having a plane circle an area collecting radar data that is turned into an image through intense calculation. ‘Backprojection’ is one SAR calculation method that helps to ease data collection by enabling flexible flight paths and target area shapes. Using Intel® Xeon and Intel® Xeon Phi™, Intel and Georgia Tech demonstrate over 35 billion backprojections per second throughput per compute node on an Intel® Xeon E5-2670 based cluster, where each node is equipped with two cards of Intel® Xeon Phi™ coprocessors. Addition of these two co-processor cards delivered a node-level performance speedup of 4.8x. Overall, performance of our system corresponds to processing a 3000 x 3000 pixel image within a second on each node as well as over 500 billion backprojection/sec on 16-node system.

Paper and contributors details:

*Efficient Backprojection-Based Synthetic Aperture Radar Computation with Many-Core Processors*

*Authors: Jongsoo Park, Ping Tak Peter Tang, Mikhail Smelyanskiy, Daehyun Kim, Thomas Benson*

*Session: Tuesday, November 13, 2:30 p.m.-3:00 p.m. Room 255-BC*

### **Big Data Signal Processing (Best Paper Nominee)**

Numerous scientific and technical applications rely on the “Fast Fourier Transform” (FFT) in working with waves including signal processing, communications and multi-media. For big FFT problems running on large clusters, 50-90 percent of the time can be spent waiting on node-node data transfers. Intel’s Software and Services Group and Intel Labs devised a new framework for distributed 1-D FFT problems which traditionally require three all-to-all inter-node data exchanges (also called global transposes). Intel’s new framework delivers multiple 1D FFT algorithms requiring single all-to-all inter-node data exchange. For large-scale problems, this can

double FFT performance. Another key feature is that users can opt to further increase FFT performance by accepting reduced-accuracy results.

Paper and contributors details:

*A Framework for Low-Communication 1-D FFT*

*Authors: Ping Tak Peter Tang, Jongsoo Park, Daehyun Kim, Vladimir Petrov*

*Session: Wednesday, November 14, 10:30 a.m.-11:00 a.m. Room 255-BC*

### **Solving Differential Equations Fast (Xeon Phi)**

“Multigrid Methods” are vital to solving differential equations for a wide variety of applications, for example, special distributions of heat, fluids or other mechanical systems.

Intel Labs, in partner with University of California, Berkeley and Lawrence Berkeley National Laboratory, present optimizations that reduce data movement through communication aggregation and communication avoidance techniques. These were used to accelerate a multi-grid solver on a variety of multi/many-core platforms including Xeon and Intel® Xeon Phi™. Intel® Xeon Phi™ achieved the highest speedup of 3.5x over a Cray XE6 based reference implementation.

Paper and contributors details:

*Optimization of Geometric Multigrid for Emerging Multi- and Manycore Processors*

*Authors: Samuel W. Williams, Dhiraj D. Kalamkar, Amik Singh, Anand M. Deshpande, Brian Van Straalen, Mikhail Smelyanskiy, Ann Almgren, Pradeep Dubey, John Shalf, Leonid Oliker*

*Session: Thursday, November 15, 2:30 p.m.-3:00 p.m. Room 255-BC*

### **Creating Exascale Benchmarks**

The NAS Parallel Benchmarks (NPB) proxy scientific computing apps are a key tool for HPC system design. However, even the hardest NPB problems do not challenge a Petascale machine, let alone future Exascale ones. Intel illustrates the steps needed to extend an example benchmark, Block Tridiagonal to Exascale. The paper describes the impact to hardware, key bottlenecks, solutions to overcome them and methods to simulate the performance of future Exascale systems running the app. This will aid in the design of future Exascale systems.

Paper and contributors details:

*Extending the BT NAS Parallel Benchmark to Exascale Computing*

*Authors: Rob F. Van Der Wijngaart, Srinivas Sridharan, Victor W. Lee*

*Session: Thursday, November 15, 1:30 p.m.-2:00 p.m. Room 255-BC*

Intel, the Intel logo and Atom are trademarks of Intel Corporation in the United States and other countries.

\*Other names and brands may be claimed as the property of others.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804

CONTACT: Connie Brown  
503-791-2367  
[connie.m.brown@intel.com](mailto:connie.m.brown@intel.com)