



# News Fact Sheet

## **Intel® Many Integrated Core (Intel® MIC) Architecture ISC'11 Demos and Performance Description**

June 20, 2011 — The Intel® Many Integrated Core (Intel® MIC) Architecture is another key addition to the company's existing products, including Intel® Xeon® processors, to lead the industry into the era of exascale computing solutions. Relative to the multicore Intel Xeon processors, Intel MIC Architecture has smaller cores and hardware threads and wider vector units. This is ideal for achieving higher aggregate performance for highly parallel applications. As developers embrace high degrees of parallelism (instruction, data, task, vector, thread, cluster, etc.), important and popular programming models for Intel® architecture processors extend to Intel MIC Architecture without need of changing the programming tools. The same techniques that deliver optimal performance on Intel processors -- scaling applications to cores and threads, blocking data for hierarchical memory and caches -- also apply to maximizing performance on Intel MIC Architecture. With greater reuse of parallel processor code, software companies and IT departments benefit from creating and maintaining a single code base binary and not having to re-train developers on proprietary programming models associated with accelerators.

The first Intel MIC product, codenamed "Knights Corner," is planned for production on Intel's 22-nanometer technology featuring the world's first 3-D Tri-Gate transistors<sup>1</sup>.

### **ISC 2011 Demonstrations**

Early results from the Intel MIC Architecture program are on display at the International Supercomputing Conference (ISC), in Hamburg Germany, June 20-23, 2011.

These results are shown in seven demonstrations. Four of these were developed by leading supercomputing centers with access to Intel's software development platform (SDP), codenamed "Knights Ferry." The remaining three demonstrations were developed by Intel engineers. The SDP uses an Intel co-processor, codenamed "Aubrey Isle" built on 45nm technology.

Software developers for these SDP systems have access to special versions of Intel's tools for multicore processors which also support Intel MIC co-processors. This includes the Intel® Fortran Compiler, the Intel® C++ Compiler and the Intel® Math Kernel Library.

---

<sup>1</sup> [http://newsroom.intel.com/community/intel\\_newsroom/blog/2011/05/04/intel-reinvents-transistors-using-new-3-d-structure](http://newsroom.intel.com/community/intel_newsroom/blog/2011/05/04/intel-reinvents-transistors-using-new-3-d-structure)

Copyright © 2011 Intel Corporation. All rights reserved.

Intel, Xeon and the Intel logo are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.

"MIC" is pronounced "Mike" / mahyk / mik

Together these demonstrations help to illustrate how the Knights Ferry SDP fulfills the promise of both performance and programmability, and shows how the MIC architecture will help the HPC industry address key challenges ahead.

### **FZ Jülich: SMMP Protein Folding (fix above page numbers)**

Simulating the protein folding process. SMMP (Simple Molecular Mechanics for Proteins) is an application to simulate and research the folding process of proteins to reach their final shape after they get produced by a cell. The demo visually shows the protein folding process with the Knights Ferry SDP doing computation in the background. The basic idea of SMMP is to find the appropriate folding by minimizing the total energy of the folding configuration. For the optimization Monte-Carlo and simulated annealing are used. Using the familiar programming models and tools such as OpenMP\*, Jülich was able to get this application running with Knights Ferry in less than 3 days and then moved on to further optimizations for performance.

SMMP is an open-source FORTRAN package for molecular simulation of proteins using Molecular Mechanics. It was augmented to use an Intel MIC co-processor through offloading directives and was tuned to take advantage of the 512-bit wide vector operations.

The ability of an Intel MIC co-processor to do protein folding using the computationally demanding Molecular Mechanics method is an example of the flexibility that full support of established programming models can offer. Protein folding can be done by two computational approaches: Molecular Dynamics, which is an approximation to reality using statistical mechanics and Molecular Mechanics, which is significantly more computationally demanding method that models the molecular system using Newtonian mechanics. Molecular Dynamics equations can be mathematically ill-conditioned which leads to problems in long simulations that require non-trivial effort to carefully select systems to simulate. Molecular Dynamics have been cited for protein folding on a GPU. Intel MIC co-processors can support either method of protein folding. This demo shows the more computationally complex Molecular Mechanics model for protein folding, yielding great results using the SMMP package on an Intel MIC co-processor. The KISTI demo shows Molecular Dynamics being used very effectively on an Intel MIC co-processor.

**System Config:** Intel Shady Cove Software Development Platform with 2 Intel® Xeon® X5690 processors (each with 6 cores, 12MB L3 Cache, 3.46GHz), 24GB DDR3 1333MHz Memory, single Intel® 5520 IOH, RHEL 6.0 OS, 1 Knights Ferry co-processor card (C0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.0GT/s), Alpha 5 Intel® MIC software stack and development tools.

### **KISTI Molecular Dynamics**

This application uses Empirical-potential molecular dynamics to model the interaction of molecules, simulating nano-materials such as carbon nanotube, graphene, fullerene, and silicon surfaces. It started on a multicore system using OpenMP and was easily extended to utilize an Intel MIC co-processor within mere hours because of the strong support for standard languages and programming models.

The standards-based parallel programming models, common between Intel Xeon processors and the Knights Ferry SDP, and the simple but powerful offload programming model turned out to be much more productive when compared to working with GPUs. The simulation is highly parallel with minimum data dependencies, and is well parallelized with standard multi-threading

programming models such as OpenMP and Intel® Cilk™ Plus. The application shows near-theoretical linear performance scaling with the number of cores both on Intel Xeon processors and Knights Ferry. The application is fully utilizing the 512-bit wide Vector Processing Unit in Knights Ferry via intrinsics in C.

Dr. Hongsuk Yi, the Heterogeneous Computing team leader, KISTI Supercomputing Center, reported that “by just utilizing standard programming on both Intel Xeon processor and Intel MIC architecture based platforms, the performance met multi-threading scalability expectations and we observed near-theoretical linear performance scaling with the number of threads.”

**System Config:** Dell Precision\* Workstation with one (1) Intel Xeon X5620 processor (4 cores, 12MB L3 Cache, 2.40GHz), 24GB DDR3 1333MHz Memory, single Intel® 5520 IOH, RHEL 6.0 OS, 1 Knights Ferry co-processor card (C0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.0GT/s).

### **LRZ: TifaMMMy Matrix Multiplication**

This shows a recursive scheme to partition input data for computation and parallelization using a cache-oblivious implementation of matrix-matrix multiply.

TifaMMMy is a cache-oblivious implementation of a matrix-matrix multiply using the space filling Peano curve and offering an intuitive C++ API. It uses a recursive scheme to partition the input data for the computation and parallelization. While this worked well for CPUs it could not be ported to GPUs due to the use of recursion. With the CPU-like programming model on the Intel MIC co-processor, LRZ could get this application running on Knights Ferry within hours and after some optimizations saw performance in excess of 650 GFLOP.

This program is written in C++ using OpenMP and taking advantage of the 512-bit wide vector capabilities via intrinsics in C++. The 650 GFLOPs demonstrated is close to optimal for this particular algorithm (a little more than half the theoretical peak of the card) even though there remains some small additional opportunity for gain as tools mature.

Most importantly, the use of a C++ application using a sophisticated self-adapting recursive algorithm has tremendous implications in terms of “performance portability” of codes. Such an algorithm does not require manual tuning as cache sizes and memory hierarchy changes even when moving to the Intel MIC architecture. By being able to easily utilize Intel MIC co-processors via standard programming methods, these self-adapting methods can be available for other important algorithms, including PDE solvers, in addition to matrix multiply, and are applicable to dense and sparse matrices. This helps illustrate the enormous potential of Intel’s MIC architecture and its ability to support programming methods, such as recursion, already in use.

**System Config:** Intel Shady Cove Software Development Platform with 2 Intel® Xeon® X5690 processors (each with 6 cores, 12MB L3 Cache, 3.46GHz), 24GB DDR3 1333MHz Memory, single Intel® 5520 IOH, RHEL 6.0 OS, 1 Knights Ferry co-processor card (C0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.0GT/s), Alpha 5 Intel® MIC software stack and development tools.

### **CERN openlab: Core Scaling of Intel® MIC Architecture**

CERN is showing a kernel extracted from the CBM/ALICE HLT software development for collider experiments. It estimates real trajectories based upon imprecise measurements.

This (kernel) benchmark used for collider experiments estimates real trajectories from imprecise measurements. It is important in performing full event reconstruction online, to select the most interesting events that will be sent from the detectors to the computer centers. CERN experimented with Intel MIC architecture on this kernel to test scalability and found that when code is well parallelized and vectorized, Intel MIC delivers linear scaling of performance with added cores. The demo shows a plot of measured performance versus number of cores employed for the computation. CERN observed linear scaling of performance through the 32 cores on Knights Ferry which is great news looking forward as Intel plans to add more cores in Knights Corner.

This application started on a multicore system and was easily extended to utilize an Intel MIC co-processor because of the strong support for standard languages and programming models.

**System Config:** SGI H4002 platform with 2 Intel® Xeon® X5690 processors (each with 6 cores, 12MB L3 Cache, 3.46GHz), 24GB DDR3 1333MHz Memory, RHEL 6.0 OS, 1 Knights Ferry co-processor card (C0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.0GT/s).

#### **7.4 TFLOP SGEMM in a node**

In this demo we have SGEMM (Single Precision General Matrix Multiplication) running simultaneously on eight Knights Ferry cards installed in a system from Colfax. On adding up the individual SGEMM performance on each card in the system we see around 7.4 TFLOP of delivered performance. This is very interesting to customers looking for high compute-density. It gives an early measure of the efficiency of the FLOPs in the Intel MIC architecture. Without data transfer overheads on the PCIe bus, we are seeing 925 GFLOP of measured sustained SGEMM performance per Knights Ferry card, each with 1200 GFLOP of peak theoretical performance.

**System Config:** Colfax CXT8000 Platform: With 2 Intel® Xeon® X5690 processors (each with 6 cores, 12MB L3 Cache, 3.46GHz), 24GB DDR3 1333MHz Memory, dual Intel® 5520 IOH, 4 PLX PeX8647 Gen 2 PCIe switches, RHEL 6.0 OS, with 8 Knights Ferry co-processor cards (D0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.6GT/s), Alpha 5 Intel® MIC software stack and development tools.

#### **Hybrid Computing – SGEMM with Intel® MKL**

This demo shows that users can already get over 1 TFLOP of SGEMM performance in just 18 lines of code using Intel® Math Kernel Library (MKL) with the early version of tools available to Intel MIC Partners. Intel® MKL utilizes the Intel Xeon® processors and Knights Ferry in concert to share the processing to achieve incredible platform level performance. It is critical to note that not a single line of those 18 lines of code needs to change when going from Intel Xeon processor-only platforms to the Xeon processor and MIC implementation. Intel MKL abstracts the work of checking for the availability of Knights Ferry in the system and executing SGEMM on Intel Xeon processor-only platforms or on Xeon processor and MIC implementations depending on availability of Knights Ferry.

**System Config:** Intel Shady Cove Software Development Platform with 2 Intel® Xeon® X5680 processors (each with 6 cores, 12MB L3 Cache, 3.33GHz), 24GB DDR3 1333MHz Memory, single Intel® 5520 IOH, RHEL 6.0 OS, 1 Knights Ferry co-processor card (D0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.6GT/s), Alpha 5 Intel® MIC software stack and development tools.

### **Hybrid Computing – LU Factorization**

This is a demonstration associated with Intel's ISC technical paper "**Designing and Dynamically Load Balancing Hybrid LU for Multi/Many-core.**" Last year we demonstrated the Knights Ferry SDP achieving more than 500 GFLOP on LU Factorization with a very early version of the software stack. In this demo, we take advantage of the common programming model between Intel Xeon processors and Intel MIC to use Knights Ferry and Xeon processors simultaneously to deliver high performance on LU Factorization – we measured up to 772 GFLOP of sustained performance. The common programming model allows us to dynamically process smaller sized matrices on Intel Xeon processors and larger sized matrices on Intel MIC resulting in optimal overall performance.

**System Config:** Intel Shady Cove Software Development Platform with 2 Intel® Xeon® X5680 processors (each with 6 cores, 12MB L3 Cache, 3.33GHz), 24GB DDR3 1333MHz Memory, single Intel® 5520 IOH, RHEL 6.0 OS, 1 Knights Ferry co-processor card (D0 Si, 1.2GHz, 2GB GDDR5 memory running at 3.6GT/s), Alpha 5 Intel® MIC software stack and development tools.

### **About CERN**

CERN, the European Organization for Nuclear Research, is the world's leading laboratory for particle physics. It has its headquarters in Geneva. At present, its Member States are Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, the Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom. Romania is a candidate for accession. India, Israel, Japan, Russia, , the United States, Turkey, the European Commission and UNESCO have Observer status. Additional information about CERN is available at <http://press.web.cern.ch/press>.

### **About FZ Jülich**

The Jülich Research Center (FZ Jülich) is one of the largest research centers in Europe. The Jülich Research Center pursues cutting-edge interdisciplinary research on solving the grand challenges facing society in the fields of health, energy and the environment and information technologies. The center revolves around key competencies in physics and supercomputing. Additional information about FZ Jülich is available at <http://www.fz-juelich.de>.

### **About KISTI**

The Korea Institute of Science & Technology Information is a leading national institute in building the nationwide infrastructure for knowledge and information by linking the high-performance research network with its supercomputers. The KISTI Supercomputing Center is the largest provider of supercomputing resources and high-performance networks in Korea. Additional information about KISTI is available at [www.ksc.re.kr/eng/](http://www.ksc.re.kr/eng/).

### **About LRZ**

The Leibniz-Rechenzentrum (LRZ) is a common computing center of the Ludwig-Maximilians University in München, the Technical University of München and the Bavarian Academy of Sciences. LRZ operates high performance computing systems for all Bavarian universities and federal high-performance computers, and is makes systems available for scientific research at all German universities. Additional information about LRZ is available at <http://www.lrz.de>.

## About Intel

Intel (NASDAQ: INTC) is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. Additional information about Intel is available at [newsroom.intel.com](http://newsroom.intel.com) and [blogs.intel.com](http://blogs.intel.com).

## Notices and Disclaimer

Information regarding the demonstrations including, but not limited to the precise system configuration, is subject to change without notice.

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.

Intel® compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel® and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the "Intel® Compiler User and Reference Guides" under "Compiler Options." Many library routines that are part of Intel® compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel® compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel® compilers, associated libraries and associated development tools 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 Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3) 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.

While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel® and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

The Intel products discussed herein may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel, Xeon, Cilk and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

-- 30 --

## About Intel

Intel (NASDAQ: INTC) is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing

devices. Additional information about Intel is available at [newsroom.intel.com](http://newsroom.intel.com) and [blogs.intel.com](http://blogs.intel.com).

CONTACT:   Radoslaw Walczyk  
              Intel Corporation  
              408-765-0012  
              [radoslaw.walczyk@intel.com](mailto:radoslaw.walczyk@intel.com)