Scaling the future with Intel Xeon Workshop

November 18-20, 2014

Taipei, Taiwan

Over the past few years, the use of heterogeneous architectures in HPC at the large scale has become increasingly common. One exciting new technology for HPC is the Intel Xeon Phi co-processor also known as the MIC. The Xeon Phi is x86 based, hosts its own Linux OS, and is capable of running most codes with little porting effort. However, the MIC architecture has significant features that are different from that of current x86 CPUs, and attaining optimal performance requires an understanding of possible execution models and the architecture. Multiple lectures and hands-on exercises will be used to acquaint attendees with the MIC platform and to explore the different execution modes as well as parallelization and optimization through example testing and reports.

This 3-day event explores multi-core and many-core computing with Intel® Xeon® processors and Intel® Xeon Phi™ coprocessors. Technical experts at Intel will discuss development tools, programming models, vectorization, and execution models that will get your development efforts powered up to get the best out of your applications and platforms.

Target: This tutorial is intended for application developers (or application modifiers) who wish to port their applications to supercomputers consisting of both multi-core processors and many-core co-processors on a single node.

Program Overview & Agenda:

** Limited seating

Day 1	Day 2	Day 3
Date: Tue., Nov. 18, 2014 **Time: 13:00- 15:00 (2hr)** **In conjunction with** **Intel Asia Innovation Summit**	**Date: Wed., Nov 19, 2014** **Time: 13:30-15:30(2hr)** **In conjunction with** **Intel Asia Innovation Summit**	**Date: Thu., Nov.20, 2014** **Time: 09:00-16:30(6.5hr)**
Regent Hotel Taipei 晶華酒店	**Regent Hotel Taipei** 晶華酒店	**National Taiwan University (Computer Center)** 台灣大學計算中心

High Performance Application Development for Big Data	Architecture & Tools	Hands-on Practice & Implementation
■ Future of high-performance, energy-efficient, reliable computing platform ■ Accelerators for computational intelligence	■ Scaling the future with Intel Xeon and Xeon Phi processors ■ Parallel Programming with Intel® Parallel Studio XE 2015	■ Intel® Xeon Phi™ Coprocessor Overview ■ Intel Xeon Phi™ Programming Models ■ Optimizing for Intel Xeon Phi ■ Performance analysis on Intel Xeon Phi ■ Intel-MKL on the Intel Xeon Phi ■ Stepwise optimization (hands-on lab)
Yen-Kuan Chen (Intel Labs, US) Debbie Marr (Intel Labs, US)	Feilong Huang (Intel China) Ying Hu (Intel China)	Feilong Huang (Intel China) Ying Hu (Intel China)
11/18-19 活動報名請上: http://www.intelasiainnovationday.com		11/20 活動報名請洽: (02)3366-5059 黃小姐<臺大> course@ntu.edu.tw

- Agenda subject to change

Speaker & Lecturer Profile

Yen-Kuang Chen
Principal Engineer at Intel Corporation
Associate Director of Intel-NTU Connected Context Computing Center.

Yen-Kuang Chen (also known as, Y.-K.) is a Principal Engineer at Intel Corporation and Associate Director of Intel-NTU Connected Context Computing Center. His research interests include

Internet of Things, machine-to-machine, multimedia signal processing, video compression, parallel processing, and computer architecture. He is a key contributor to SSSE3 and AVX instruction sets on Intel® Core™ processor family. He has 25+ US patents, 25+ pending patent applications, and 85+ technical publications. He has served as an editorial board member of 5+ journals and a program committee member of 40+ international conferences, on multimedia, video compression/communication, image/signal processing, VLSI circuits and systems, parallel processing, and software optimization. He is an invited participant to 2002 Frontiers of Engineering Symposium (National Academy of Engineering) and to 2003 German-American Frontiers of Engineering Symposium (Alexander von Humboldt Foundation). He received his Ph.D. from Princeton University, and is an IEEE Fellow.

Debbie Marr
Director of Accelerator Architecture Lab
Intel, Corporation

Debbie Marr is Director of the Accelerator Architecture Lab (AAL) in Intel Labs. She is responsible for research in accelerator architecture which are key to the future of computing efficiency. Marr has played leading roles on many of Intel's leading CPU products for over 20 years, from the 386SL (Intel's first mobile computing processor), to Intel's famous P6 processor, to Intel's leading-edge 2017 products. She was Intel's first Xeon server architect, launching Intel into being the predominant provide of server computing platforms. On the Pentium 4 Processor, she was the Hyperthreading Architect. She was the lead architect for the Haswell Core, and was the lead architect for advanced development for Intel's 2017/2018 Core and Xeon CPUs. Debbie has over 20 patents in many aspects of CPU architecture/microarchitecture, performance analysis, multiprocessing, workload analysis, and simulator development. Marr holds a Ph.D. in electrical and computer engineering from University of Michigan, a M.S. in electrical engineering and computer science from Cornell University, and a B.S. in electrical engineering and computer science from the University of California, Berkeley.

Hu Ying
Sr. Technical Consulting Engineer.
Software and Services Group
Intel Corporation

Ying Hu (胡英), joined Intel since 2003 and is a senior technical consulting engineer with a focus on enabling global developers, engineers & researchers to use Intel High Performance Libraries (include Math Kernel Library and Integrated Performance Primitive) on Intel architectures. As the only technical support team located in APAC, Ying worked as key interface for Intel software tools enabling for OEMs, channels & developer ecosystem from first-generation core archticture to many integrated core (MIC). She had rich experience in supporting strategic accounts in various HPC and parallel computing areas including Digital Data Processing, Industry and Academic segments. Her current interest is to explore the advantage of math library and Intel® Xeon Phi in machine learning and big data analysis domain. Ying earned her Ph.D. degree from Shanghai Jiao Tong University.

Topic of Talk

Parallel Programming with Intel® Parallel Studio XE 2015

Abstract

Intel Xeon Phi coprocessor brings a new era in high-performance computing. With the combination of Intel software, Intel Parallel Studio XE, it achieved unprecedented performance for innovative breakthroughs in manufacturing, life sciences, energy and other areas. Intel® Parallel Studio XE is the software tool suite provided by Intel SSG. The ability to quick design, development, debug, and tuning of code that utilizes parallel processing extremely help scientists and engineers optimize their code to take full advantage of Intel Xeon Phi coprocessors. This session will provide an up-to-date Parallel Studio XE overview from programming practice and introduce its components: Intel C/C++, Fortran compiler, high performance library (MKL/IPP), Threading Build block, Inspector & VTune Amplifier etc. Being familiar with Parallel Studio XE and parallel programming will increase the productivity and efficiency

of high- performance computing and make high-performance computing easier to access for developers and scientists on Intel Xeon processor and Intel Xeon phi.

Huang Feilong
Sr. Technical Consulting Engineer, PRC
Software and Services Group (SSG)
Intel Corporation

Huang Feilong is a senior technical consulting engineer in SSG. He mainly focuses on Intel compiler, code optimization and compiler compatibility analysis. His responsibility includes enabling global ISVs' applications for latest Intel processors, performance tuning on Intel architecture, product definition of Intel compiler, and delivering trainings on Intel software development tools. His current interest is to extract performance from Xeon Phi coprocessor with Intel compiler as well as C/C++ new language extensions.

Feilong joined Intel in 2002 as a technical consulting engineer. Before joining Intel, he worked in National Instruments as a software developer working on Linux device driver development. He graduated from Fudan University in Shanghai and got his master degree of engineering.

Topic of Talk

Scaling the future with Intel Xeon and Xeon Phi processors

Abstract

This talk covers roadmap of Intel® Xeon processor and Intel® Xeon Phi™ coprocessor, as well as their architectures. It explains the differences between these two architectures. It also shows Xeon Phi performance metrics from a number of ISVs in a few vertical segments. There are various workloads in the real world. Some of them run better on Xeon, while the others run better on Xeon Phi. After this talk, audience would be expected to tell whether a workload is suitable to run on Xeon or on Xeon Phi.

The Intel® Xeon Phi™ coprocessor, first generation product code name Knights Corner, provides high performance, and performance per watt for highly parallel workloads, while not requiring a

new programming model, API, language or restrictive memory model.