Four-Days Technology Workshop
on
Hybrid Computing - Coprocessors & Accelerators - Power-Aware Computing & Performance of Application Kernels (HyPACK-2013)

Jointly Organized by

Centre for Development of Advanced Computing (C-DAC), Pune University Campus, Pune
Centre for Modelling & Simulation (CMSD), HPC Facility, University of Hyderabad,

Venue : CMSD, University of Hyderabad, Hyderabad

Dates : October 15 (Tuesday) – October 18 (Friday)

Technical Programme

Day 1: October 15, 2013 (Tuesday)

8:30 AM ~ 9:00 AM: Registration	**10:30 AM ~10:45 AM Coffee & Tea Break**
9:00 AM ~ 9:30 AM	Welcome & Inauguration
9:30 AM ~ 10:00 AM	An Overview of hyPACK-2013: Technical Prog. &Hands-on Session
Coffee & Tea Break: 10:00 AM - 10:30 AM	
10:30 AM ~ 11:15 AM	**Classroom Lecture & Lab. :** An Overview of Intel Xeon-Phi Arch. & Programming Models; OpenMP Prog. & Performance Issues - Numerical Computations – Lab Sessions
11:15 AM ~ 11:45 AM	**Classroom Lecture & Lab.:** An Overview of Intel Xeon-Phi – System Software; X86 SMP Compiler & Vectorization; Native Compilation & Compiler Offload Pragmas; Performance Issues - Numerical Computations – Laboratory Sessions
11:45 AM ~ 12:45 PM	**Classroom Lecture & Demonstration:** Intel Xeon-Phi Architecture – Prog. Models - Compilation features; Compiler tips, Compiler Vectorization reports; Compiler Directives & Demonstration Benchmarks - OpenMP on Xeon Host and Xeon-Phi Coprocessor
12:45 PM ~ 1:00 PM	**Lab. Session:** Hands-on session on Intel Xeon-Phi Coprocessor – OpenMP framework on Intel Xeon-Phi – Tuning & Performance
Lunch: 1:00 PM ~1:45 PM	**Coffee & Tea Break: 4:00 PM - 4:15 PM**

Day 1: October 15, 2013 (Tuesday)

Lunch: 1:00 PM ~1:45 PM	**Coffee & Tea Break: 4:00 PM - 4:15 PM**
1:45 PM ~ 2:30 PM	**Lab. Session:** Programming – Intel Xeon-Phi Compiler Offload Pragmas; Compiler Technology & Vectorization – Numerical Computations
2:30 PM ~4:00 PM	**Demonstration & Lab. Session:** Mixed Programming **(**OpenMP, Intel TBB, Pthreads); Compilation & Vectorization - Numerical Computations based on Intel MKL on Xeon Phi Coprocessors; Tuning and Performance of Benchmarks on Intel Xeon-Phi; Compilation – Intel Offload pragmas
4:15 PM ~ 5:00 PM	**Keynote Talk (Academic) :** Numerical Linear Algebra (NLA); Intel MKL - Performance on Intel Xeon-Phi Coprocessors
5:00 PM ~ 5:45 PM	**Lab. Session:** Intel Xeon-Phi Prog **(**OpenMP, TBB, Pthreads); Performance of Benchmarks – Numerical Computations: Compiler Optimizations & Vectorization; Intel MKL Math Kernels on Xeon Phi; Thread Affinity – OpenMP & Pthreads on Intel Xeon Phi; Performance of Application Kernels on Intel Xeon-Phi
5:45 PM ~ 6:30 PM	**Invited Talk** : An Overview of MPI programming - Intel Xeon-Phi Coprocessors & Performance Issues

Day 2: October 16, 2013 (Wednesday)

9:00 AM ~ 9:45 AM	**Classroom Lecture & Lab:** Prog. Intel Xeon Phi – Prog. Affinity Concepts Intel Xeon-Phi Programming Models – Intel TBB, Overview of Co-Processor Prog. Models; Tuning & Performance of NLA Kernels
9:45 AM ~ 10:30 AM	**Classroom Lecture & Lab:** Prog. Intel Xeon Phi -Compiler Options; Cilk Plus Compiler tips, Compiler Vectorization reports; Compiler Directives – Memory alignment; Math Kernel Library - Performance Results
10:30AM ~10:45 AM Coffee & Tea Break	
10:45 AM ~11:15 AM	**Classroom Lecture & Lab :** MPI execution models on Intel Xeon and Intel Xeon Phi coprocessors, including pure MPI or hybrid MPI applications
11:15 AM ~12:15 PM	**Keynote Talk (Industry) : Speaker : Intel** **Topic :** An Overview of Intel Xeon-Phi Co-processor Archi.; An Overview of Co-Processor System Software; Compiler Offload Pragmas & Tuning for performance on Xeon Phi(tm) coprocessor using VTune Amplifier XE
12:15 AM ~1:00 PM	**Keynote Talk (Industry) : Speaker : Intel** Intel Xeon Phi / Demos: Compilers, VTune Amplifier Demonstrate Compilers and VTune AXE on Intel ; Tuning methodologies; Intel(r) Trace Analyzer and Collector (ITAC) on MIC
Lunch: 1:00 PM ~ 1:45 PM	**Coffee & Tea Break: 4:00 PM - 4:15 PM**

Day 2: October 16, 2013 (Wednesday)

1:45 PM ~ 3:00 PM	**Demonstration & Lab. Session:** MPI execution models on Intel Xeon processors and Intel Xeon Phi coprocessors, including pure MPI or hybrid MPI applications.; Mixed Prog. (MPI-OpenMP, MPI-TBB, OpenCL, Pthreads); Basic Programs based on Intel Xeon-Phi Co-processors; Bandwidth Calculation Matrix Computations; Using Intel Xeon-Phi Tools for Numerical Computations
3:00 PM ~ 4:00 PM	**Keynote Talk (Industry) : Speaker : Intel** Performance of Application Kernel – based on OpenMP /MPI - Intel Xeon-Programming Framework
4:15 PM ~ 5:15 PM	**Lab. Session:** Mixed Prog. (MPI-OpenMP, MPI-TBB, MPI-Pthreads) on Intel Xeon Co-Processors; I/O files on Intel Xeon Co-processors; Memory Map (mmap) Examples on Intel Xeon-Phi Co-processors; MPI Prog. for Numerical Computations on Intel Xeon Phi; Intel Tool Kit Demonstration
5:15 PM ~ 6:00 PM	**Classroom Lecture & Lab (part-II):** Prog. Intel Xeon Phi – Memory Map (mmap) & Huge Page Enabling

Day 3: October 17, 2013 (Thursday)

9:00 AM ~ 9:45 AM	**Classroom Lecture & Lab :** Introduction to GPU Computing : Memory Optimization; Tuning & Performance on CUDA enabled NVIDIA GPUs;
9:45 AM ~ 10:30 AM	**Classroom Lecture & Lab:** CUDA – enabled NVIDIA GPUs – Use of CUDA Toolkit Math Libraries – OpenACC Pragmas Framework – CUDA Enabled NVIDIA Multi-GPUs
10:30AM ~10:45 AM Coffee & Tea Break	
10:45 AM ~ 11:45 AM	**Keynote Talk (Industry) : Speaker : NVIDIA** An Overview of CUDA enabled NVIDIA GPUs – Programming & Performance Issues
11:45 AM ~ 12:45 PM	**Keynote Talk (Industry) : Speaker : NVIDIA** An Overview of CUDA enabled NVIDIA GPUs – Programming & Performance Issues.
Lunch: 1:00 PM ~ 2:00 PM	**Coffee & Tea Break: 4:00 PM - 4:15 PM**

Day 3 : October 17, 2013 (Thursday)

Lunch: 1:00 PM ~ 2:00 PM	Coffee & Tea Break: 4:00 PM - 4:15 PM
2:00 PM ~ 3:00 PM	**Lab. Session & Demonstration :** Lab. Session: An overview of CUDA enabled NVIDIA GPUs / OpenCL – GPGPUs/Example Programs.
3:00 PM ~ 4:00 PM	**Lab. Session & Demonstration**: Tuning & Performance on CUDA enabled NVIDIA-GPUs; Matrix-matrix multiplication - tiled techniques for partitioning of a matrix, shared memory optimization, Warp level parallelism; Tuning & Performance on Multi-CUDA enabled NVIDIA-GPUs;
4:15 PM ~ 5:00 PM	**Keynote Talk (Academic):** Tuning & Performance on CUDA NVIDIA GPUs
5:00 PM ~ 5:30 PM	**Lab. Session & Demonstration:** Application kernels based on Mixed Prog. (MPI-CUDA, Pthreads-CUDA & OpenMP-CUDA); CUDA SDK ToolKit Demonstration; Prog. On Heterogeneous Comp. Platforms – AMD-APP; Basic OpenCL Programs based on Single /Multiple GPUs on AMD-APP GPUs; Use of Work Groups & Work-items – Memory Optimizations; Prog. on ARM Multi-Cores with CUDA NVIDIA carma – Using NVML APIs
5:30 PM ~ 6:00 PM	**Class-room Lecture & Demonstration :** Measurement of Power Consumption–Performance of Application Kernels – Using NVML Lib Calls

Day 4: October 18, 2013 (Friday)

9:00 AM ~ 9:45 AM	**Classroom Lecture & Lab** : Heterogeneous Programming – CUDA enabled NVIDIA GPUs /AMD APP – OpenCL; Tuning & Performance – Matrix Computations; AMD APP Tech. – SDK & Prog. Env /Libraries
9:45 AM ~ 10:30 AM	**Classroom Lecture & Lab :** An Overview of HPC GPU Cluster – OpenCL Performance Issues – Numerical Linear Algebra; AMD APP Tech –Tuning & Performance OpenCL; Demonstration of Application Kernels
Coffee & Tea Break: 10:30 AM - 10:45 AM	
10:45 AM ~11:45 PM	**Keynote Talk (Industry):** Measurement of Power Consumption – Systems with Coprocessors/Accelerators; Case Studies - Partial differential Eqs – Solution of Matrix Systems on Cluster with Single /Multiple GPUs
11.45 AM ~12:45 PM	**Keynote Talk (Academic):** Performance of Application Kernels on Parallel Processing Platforms with GPU Accelerators
Lunch: 1:00 PM ~ 1:45 PM	**Coffee & Tea Break: 4:00 PM - 4:15 PM**
2:00 PM ~ 3:00 PM	**Lab. Session & Demonstration:** Hands-on session on NVIDIA GPUs /AMD OpenCL;Prog. on ARM Multi-Core system with CUDA NVIDIA carma – Using NVML APIs
3:00 PM ~ 4:00 PM	**Lab. Session - Demonstration:** Programming based on OpenCL, Tuning & Performance of OpenCL on GPGPUs; matrix-matrix multiplication – algorithms based on OpenCL; Shared memory optimization, Wavefront level parallelism; – Memory Optimizations;
4:15 PM ~ 5:00 PM	**Keynote Talk (Academic):** An Overview of Application Kernels on Parallel Processing Systems with Multi-GPU – Power aware Performance Issues – NVML Library calls & external Power-Off Meter
5:00 PM ~ 5:45 PM	**Lab. Session:** Example programs on host-cpu (Pthreads, MPI, OpenMP) and OpenCL on Multiple GPUs; Tuning & Performance of Matrix Computations on AMD-APPs; Memory Optimization on AMD APP–OpenCL; Application kernels based on Mixed Prog. (MPI,Pthreads, OpenMP- with OpenCL); OpenCL programs for Numerical Linear Algebra on HPC GPU Cluster (OpenCL on NVIDIA/AMD-APP GPUs) Benchmarks
5:45 PM ~ 6:00 PM	**Closing & Feedback Session**