

## hyPACK-2013

### Four-Days Technology Workshop on

#### Hybrid Computing - Coprocessors & Accelerators - Power-Aware Computing & Performance of Application Kernels

Jointly Organized by

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

**Venue** : CMSD, University of Hyderabad, Hyderabad

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

#### Tentative Technical Programme

##### Day 1: October 15, 2013 (Tuesday)

<b>8:30 AM ~ 9:00 AM: Registration</b>		<b>10:30 AM ~10:45 AM Coffee &amp; Tea Break</b>	
9:00 AM ~ 9:30 AM	Welcome & Inauguration		
9:30 AM ~ 10:00 AM	An Overview of HyPACK-2013: Technical Prog. & Hands-on Session		
<b>Coffee &amp; Tea Break: 10:00 AM - 10:15 AM</b>			
10:30 AM ~ 11:15 AM	<b>Classroom Lecture &amp; Lab.:</b> An Overview of Intel Xeon-Phi Prog. - X86 SMP Compiler & Vectorization Performance Issues - Lab Session		
11:15 AM ~ 11:45 AM	<b>Classroom Lecture &amp; Lab.</b> An Overview of Intel Xeon-Phi Prog. – X86 SMP – OpenMP Prog. & Intel MKL Performance Issues & Lab Session		
11:45 AM ~ 12:45 PM	<b>Keynote Talk (Industry/Academic) :</b> Intel Xeon-Phi Architecture – Prog. Models - Compilation Features & Demonstration Benchmarks		
12:45 PM ~ 1:00 PM	<b>Lab. Session:</b> Hands-on session on Intel Xeon-Phi Coprocessor		
<b>Lunch: 1:00 PM ~1:45 PM</b>		<b>Coffee &amp; Tea Break: 4:00 PM - 4:15 PM</b>	
1:45 PM ~ 2:00 PM	<b>Lab. Session:</b> Hands-on session on Intel Xeon-Phi Coprocessor		
2:00 PM ~4:00 PM	<b>Lab. Session:</b> Mixed Programming (OpenMP, Intel TBB, Pthreads, MPI); Numerical Computations; Multi-threaded I/O; Prog. on Multi-Core Processors with Xeon Phi Coprocessors; Performance of Benchmarks on Intel Xeon-Phi Coprocessors; MPI Example programs and performance on Intel Xeon-Phi Coprocessors		
4:15 PM ~ 5:00 PM	<b>Keynote Talk (Industry/Academic) :</b> Numerical Linear Algebra (NLA); Intel MKL - Performance on Intel Xeon-Phi Coprocessors		
5:00 PM ~ 6:00 PM	<b>Lab. Session:</b> Intel Xeon-Phi Prog : Mixed Prog. (OpenMP, MPI-TBB, MPI-Pthreads); Performance of Benchmarks: Compiler Optimizations & Vectorization; Math Kernel s; Libraries Performance on Intel Xeon-Phi Coprocessor Application Kernels on Multi-Cores		
6:00 PM ~ 6:30 PM	<b>Special Invited Talk :</b> An Overview of MPI programming - Intel Xeon-Phi Coprocessors & Performance Issues		

##### Day 2: October 16, 2013 (Wednesday)

9:00 AM ~ 9:45 AM	<b>Classroom Lecture &amp; Lab:</b> Prog. Intel Xeon Phi -Compiler Options; Compiler tips, Compiler Vectorization reports; Compiler Directives – Memory alignment; Performance Results Bandwidth – NLA Kernels
9:45 AM ~ 10:30 AM	<b>Classroom Lecture &amp; Lab (Part-I):</b> Prog. Intel Xeon Phi – Prog. Affinity Concepts – OpenMP, Intel TBB, Pthreads, Cilk Plus, Co-Processor prog. Models; Perf. of Numerical Computations, NLA & Intel MKL Kernels

10:30AM ~10:45 AM Coffee & Tea Break	
10:45 AM ~11:30 AM	<b>Classroom Lecture &amp; Lab (part-II):</b> Prog. Intel Xeon Phi – Prog. Affinity Concepts – OpenMP, Intel TBB, Pthreads, Cilk Plus, Co-Processor prog. Models; Perf. of Numerical Comps., NLA & Intel MKL Kernels
11:30 AM ~12:15 PM	<b>Keynote Talk (Industry/Academic):</b> Intel Xeon-Phi Co-processor Architecture; Off load Models – Overview of Co-Processor System Software Pragma/directive offload; Memory Map – Huge Page Enabling ; I/O files on Co-processor
12:15 AM ~1:00 PM	<b>Keynote Talk (Industry/Academic):</b> Topic: Performance of Application Kernels on Intel Xeon Phi Coprocessor – Profiling & Tuning
Lunch: 1:00 PM ~ 1:45 PM	Coffee & Tea Break: 4:00 PM - 4:15 PM
1:45 PM ~ 2:00 PM	<b>Lab. Session:</b> Hands-on session on Intel Xeon-Phi Coprocessor
2:00 PM ~ 4:00 PM	<b>Lab. Session:</b> Mixed Prog. (MPI-OpenMP, MPI-TBB, OpenCL, Pthreads); Basic Programs based on Intel Xeon-Phi Co-processors; Bandwidth Calculation Matrix Computations, Compiler Optimizations & Math Kernel Libraries on Intel Xeon Phi Coprocessors; Example programs based on Pthreads, Intel TBB, Cilk Plus, MPI, & OpenMP on Intel Xeon Phi Coprocessors; Application Kernels on Intel Xeon-Phi Coprocessors
4:15 PM ~ 5:00 PM	<b>Classroom Lecture &amp; Lab :</b> An overview of Introduction to GPU CUDA NVIDIA GPUs; Example programs on CUDA – Matrix Comps.
5:00 PM ~ 6:30 PM	<b>Lab. Session:</b> 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; CUDA Prog. for Numerical Computations; 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; Application kernels based on Mixed Prog. (MPI-CUDA, Pthreads-CUDA & OpenMP-CUDA); CUDA SDK ToolKit Demonstration

### Day 3: October 17, 2013 (Thursday)

9:00 AM ~ 9:45 AM	<b>Classroom Lecture &amp; Lab :</b> GPU Computing : Memory Optimization; Tuning & Performance on CUDA enabled NVIDIA GPUs;
9:45 AM ~ 10:30 AM	<b>Classroom Lecture &amp; Lab:</b> <i>An Overview of OpenCL – on GPGPUs &amp; Tuning /Performance – CUDA/OpenCL Enabled NVIDIA Multi-GPUs</i>
10:30AM ~10:45 AM Coffee & Tea Break	
10:45 AM ~ 11:45 AM	<b>Keynote Talk ( Industry/Academic):</b> An Overview of AMD-APP Tech-Heterogeneous Programming : OpenCL
11:45 AM ~ 12:45 PM	<b>Keynote Talk ( Industry/Academic):</b> An Overview of ARM Processor Multi-Core System with NVIDIA carma DevKit – Prog. Performance Issues.
12:45 AM – 1:00 PM	<b>Lab. Session:</b> An overview of CUDA enabled NVIDIA GPUs /AMD-APP OpenCL – GPGPUs/Example Programs.
Lunch: 12:45 PM ~ 1:45 PM	Coffee & Tea Break: 4:00 PM - 4:15 PM

### Day 3 : October 17, 2013 (Thursday)

Lunch: 12:45 PM ~ 1:45 PM		Coffee & Tea Break: 4:00 PM - 4:15 PM	
1:45 PM ~ 2:00 PM	<b>Lab. Session:</b> Hands-on session on NVIDIA GPUs /AMD OpenCL		
2:00 PM ~ 4:00 PM	<b>Lab. Session &amp; Demonstration:</b> CUDA Prog. for Numerical Computations; 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	<b>Keynote Talk ( Industry/Academic):</b> An Overview of Application Kernels & Benchmarks on Parallel Processing Systems with Multi-GPU - OpenCL		
5:00 PM ~ 6:30 PM	<b>Lab. Session &amp; Demonstration:</b> Application kernels based on Mixed Prog. (MPI-CUDA, Pthreads-CUDA & OpenMP-CUDA); CUDA SDK ToolKit Demonstration; Prog. On Heterogeneous Computing Platforms – AMD-APP; Basic OpenCL Programs based on Single /Multiple GPUs on AMD-APP GPUs; OpenCL-Matrix Computations on Multi-GPUs; Use of Work Groups & work-items – Memory Optimizations; Programming on ARM Multi-Core system with CUDA NVIDIA carma – Using NVML APIs		

### Day 4: October 18, 2013 (Friday)

9:00 AM ~ 9:45 AM	<b>Classroom Lecture &amp; Lab ::</b> 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	<b>Classroom Lecture &amp; Lab :</b> An Overview of HPC GPU Cluster – OpenCL Performance Issues – Numerical Linear Algebra; AMD APP Tech –Tuning & Performance OpenCL; Demonstration of Application Kernels		
<b>Coffee &amp; Tea Break: 10:30 AM - 10:45 AM</b>			
10:45 AM ~11:45 PM	<b>Keynote Talk (Industry):</b> Partial differential Eqs (Finite Element /Finite Difference Methods) – Solution of Sparse Matrix Systems on HPC GPU Cluster & Numerical Linear Algebra on Multi-Core Processor Systems with Single /Multiple GPUs - HPC GPU Cluster		
11.45 AM ~12:45 PM	<b>Keynote Talk (Industry):</b> Performance of Application Kernels on Parallel Processing Platforms - String Search Algorithms – OpenCL		
Lunch: 1:00 PM ~ 1:45 PM		Coffee & Tea Break: 4:00 PM - 4:15 PM	
1:45 PM ~ 2:00 PM	<b>Lab. Session:</b> Hands-on session on NVIDIA GPUs /AMD OpenCL		
2:00 PM ~ 4:00 PM	<b>Lab. Session - Demonstration :</b> Programming based on OpenCL, Tuning & Performance of OpenCL on GPGPUs; OpenCL for matrix-matrix multiplication – Optimization; , shared memory optimization, Wavefront level parallelism; – Memory Optimizations; Programming on ARM Multi-Core system with CUDA NVIDIA carma – Using NVML APIs		
4:15 PM ~ 5:00 PM	<b>Keynote Talk ( Industry/Academic):</b> 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 ~ 6:30 PM	<b>Lab. Session:</b> 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); Tuning & Performance for matrix computations based on AMD – ACML, OpenCL programs for Numerical Linear Algebra on HPC GPU Cluster (OpenCL on NVIDIA/AMD-APP GPUs) Benchmarks		