

Four Days C-DAC Internal Training Programme on

Hybrid Computing: Co-processors & Accelerators – Case Studies (Era of Multi-to-Many Core Processors)

at C-DAC ACTS PUNE (September 24-27, 2013) or CMSD, U of HYDERABAD (October 15-18, 2013)

(Co-located with hyPACK-2013 at CMSD, U of HYDERABAD during October 15-18, 2013)

(hyPACK –2014 is four days National Tech. Workshop on “Hybrid Computing-Coprocessors & Accelerators- Power-aware Computing and Performance of App. Kernels. The training programme is designed for Class-room lectures and laboratory sessions on Coprocessors & GPUs including case studies for applications. Topics of Interest are given below.

Day-1 & Day-2:

- Programming on - Intel Xeon Phi Coprocessors multicore processor HPC Systems
- Xeon Phi usage model: MPI *versus* Offload; Compiler and Prog. Model; Approaches to Vectorization. Compiler Directives; OpenMP-4.0, Intel TBB, Intel Cilk Plus, Intel MKL
- Intel Xeon-Phi Coprocessor Architecture; Coprocessor System software; Tuning Memory Allocation Perf., Huge Page Sizes; Profiling & Tuning Tools
- Tuning & Performance Issues- Measurement of Power Consumption for App Kernels; External Power-Off Meter; Application Kernels; Energy Efficiency & Performance Issues

Day-3 & Day-4:

- HPC Cluster-NVIDIA GPUs and AMD GPUs /HPC Cluster with Accelerators
- An Overview of CUDA enabled NVIDIA GPUs : CUDA SDK/APIs; CUDA - Optimization & Performance Issues; Efficient use of different memory types, (CUBLAS, CUFFT, CULA Tools, MAGMA, CUSPARSE, Thrust); CUDA-OpenACC APIs; NVIDIA- OpenCL CUDA NVIDIA GPU Cluster - Kepler GPUs; NVIDIA CUDA – for Application Kernels on GPGPUs (NVIDIA KEPLER Features- Directives-OpenACC,OpenCL)
- An Overview of AMD Accelerated Parallel Processing (APP) Capabilities; AMD APUs - OpenCL Prog. On Multi-Core CPUs & Multi-GPUs; AMD APP Math libraries- BLAS & FFTs; AMD APP SDK, AMD tools- Aparapi APIs; AMD OpenCL tuning & performance; HPC AMD GPU Cluster: Host CPU (Pthreads, OpenMP, MPI) with OpenCL on AMD GPUs; Health Monitoring of GPUs

An Overview of Hands-on:

- Memory Allocators, Software Threading, Mixed Programs (OpenMP 3.0/4.0, POSIX Threads, Intel TBB, MPI, Memory allocators) for Numerical (Matrix Comps.); & Non-Numerical Comps. Multi-Core Software tools; Application and System Benchmarks-Top-500 & HPC Challenge Benchmarks;
- Prog. on Intel Xeon Phi; MPI *versus* Offload; Compiler & Prog. Model; Prog. Paradigms - OpenMP, Intel Cilk Plus, Intel MKL; Tuning & Perf, Huge Page Sizes; Profiling
- Basic prog. (NVIDIA GPU Comp. CUDA 4.0 SDK & AMD GPGPUs - AMD-APP SDK); CUDA Toolkit; CUDA Matrix Comps. Lib.; OpenCL - NVIDIA - CUDA Prog. on Numerical Comps. CUDA Streams; Multi-GPU progres; NVIDIA CUDA OpenACC - Prog.
- AMD-APP- OpenCL; OpenCL/CUDA -Multi-GPU; NVIDIA /AMD-APP Profilers & Tools; NVIDIA NVML APIs - Prog.; Measure Power-aware Perf. for application kernels
- CUDA /OpenCL Programs on Numerical Comps. (Dense /Sparse Matrix Comps.); Partial Diff. Eqs.; String Search Alg.; FFTs; Image Processing Algorithms, Programming Environment on HPC GPU Cluster; Performance issues of Benchmarks & Application Kernels Mixed Prog. on Host CPU (MPI, OpenMP, Pthreads) and CUDA/OpenCL on devices, OpenACC, OpenMP 4.0.

Special lectures: Application Kernels & Case Studies

- Mixed Programming for Numerical /Non-Numerical Computations on multi-core Processors with Intel Xeon-Phi coprocessors, and NVIDIA /AMD GPUs accelerators; Application & System Benchmarks; Image Processing Applications; Bio-Informatics String Search Algorithms & Sequence Analysis; Matrix Computations on HPC Cluster with accelerators and Co-processors; Solution of Partial Diff. Eqs. (FDM & FEM); FFT Libraries; Information Science applications.

Experts from Industry:

- NVIDIA and Intel experts are invited to deliver talks on HPC Multi-core processors and GPUs