Summary of
Hybrid Computing – Coprocessors & Accelerators – Power-aware Computing & Performance of Application Kernels (hyPACK – 2013)
(*Initiatives on Measurement of Power Consumption and Performance***)**

Centre for Development of Advanced Computing (C-DAC) Pune and Centre for Modeling Simulation and Design (CMSD), University of Hyderabad (UoH), Hyderabad jointly conducted **four** days technology workshop titled "**Hybrid Computing – Coprocessors & Accelerators – Power-aware Computing & Performance of Application Kernels (HyPACK – 2013) (Initiatives on Measurement of Power Consumption & Performance)** at CMSD, UoH during **October 15-18, 2013**.

One of the objectives of **hyPACK-2013** technology workshop was to develop expertise on tuning and performance of application kernels involving numerical and non-numerical computations on Hybrid Computing Platforms i.e., Multi-Core processor systems with different devices such as Intel Xeon-Phi Coprocessors and NVIDIA & AMD GPU Accelerators as well as ARM multi-core processor systems. The *second* objective is write programs based on shared address programming such as OpenMP 4.0, Pthreads, Intel TBB, Cilk Plus, OpenCL, CUDA with explicit message passing MPI on multi-core systems with different devices as well understand various tools to get insight into behavior of codes. The *third* objective is to write programs on ARM based low power consumption multi-core systems and measure power consumption as well as performance of numerical linear algebra (NLA) algorithms and application kernels.

The aim of **hyPACK-2013** was to understand how to design and re-write application kernels for numerical linear algebra (NLA) algorithms and application kernels which can use compiler and vectorization features and different programming paradigms that are supported on hybrid parallel processing platforms with different devices such as Intel Xeon-Phi Processors and GPU accelerators. Codes which include to measure power consumption and performance of NLA kernels are on systems with different accelerators are included in **hyPACK-2013**.

The **four days** workshop provided an opportunity for interaction among the various participants from different academic institutes and research organizations in the country to understand more about performance aspects on multi-to-many core platforms. An opportunity was provided for software professionals to get performance of applications on cluster of many-core device accelerators on emerging heterogeneous parallel processing platforms with different devices (Intel Xeon-Phi Coprocessors, NVIDIA & AMD GPU accelerators).

By understanding the **hyPACK-2013** hard-copy and softcopy CD as building blocks, software professionals could piece together more complicated software tools that are tailored specifically for their needs, on Multi-Core Processors with Intel Xeon-Phi Coprocessors and GPGPU accelerators, HPC GPU Clusters and ARM Multi-core processors.

The workshop is organized in **six** sections as six **"Modes"** focusing on various programming paradigms of hybrid computing platform systems with coprocessors and GPUs. The **four** days workshop is aimed to cover classroom lectures in morning/forenoon session and four hours hands-on in afternoon session on distributed Shared memory platforms and message passing cluster with Intel Xeon-Phi coprocessors, and GPPUs on each day. The rich set of codes is provided on various computing platforms to understand and address performance issues of different codes that are written for this workshop. Participants will get an opportunity to walk-through and execute some of the programs designed for application kernels and understand various profiling tools.

The **Mode-1** and **Mode-2** section gives insights into performance aspects of software threading using different programming paradigms on Multi-Core processors and **ARM** processor based systems. The **Mode-3** discusses an overview of programming on Intel Xeon Multi-Core systems with Xeon-Phi Coprocessors. The focus is to discuss codes for NLA & application kernels in hands-on session which are based on compilers-vectorization, "Offload" & "Native" mode options provided for different programming paradigms such as MPI, OpenMP 3.0/ 4.0, Cilk Plus, Intel TBB, & OpenCL.

The **Mode-4** will cover an overview of GPU Computing - CUDA enabled NVIDIA Programming Software toolkit, GPGPUs - AMD-APP (SDK) with Hands-on Session for numerical and non-numerical computations. The programming on GPUs based on CUDA - OpenACC and OpenCL frame work is covered in order to solve prototype applications. Special programs on measurement of power consumption and performance of application kernels on Multi-Core processor systems with CUDA enabled NVIDIA GPUs, ARM Processor Systems are discussed. Industry experts will demonstrate the software for scientific kernels on Systems with GPGPUs / GPU accelerator devices.

The **Mode-5** section will cover an overview of Message Passing Cluster with different devices such as Intel Xeon-Phi Coprocessors, AMD and NVIDIA GPU accelerators. The hands-on session is focused on writing programs using explicit message passing (MPI) and programming paradigms such as CUDA, OpenCL, Intel Xeon-Phi Offload /native pragmas, OpenMP, Cilk Plus, Intel TBB, and Intel MKL.

The **Mode-6** section covers performance of application kernels on hybrid heterogeneous HPC cluster with different devices (Intel Xeon Phi Coprocessors, CUDA, OpenCL enabled NVIDIA GPUs, AMD GPUs. It is aimed to write programs on HPC GPU Cluster to solve compute intensive applications. Experts from private sector demonstrate software and hardware components based on Multi-Core processor systems with Intel Xeon-Phi Coprocessors and GPU accelerators.

The programing frameworks such as Intel Xeon-Phi coprocessor Offload and native mode pragmas, Intel Xeon-Phi Compiler & vectorization features, CUDA enabled NVIDIA GPUs, GPGPUs & GPU Computing - CUDA enabled NVIDIA GPUs, NVIDIA-OpenACC- Compiler Directives, GPGPUs AMD-APP SDK, OpenMP 4.0 frame work are discussed and explained several codes for numerical and non-numerical computations in Hands-on Session. The programming on HPC systems with GPUs with

different programming on host-CPU and CUDA / OpenCL Programming on device-GPUs are covered. Exhaustive hands-on session is made available to address tuning and performance issues on Intel Xeon-Phi Coprocessors, CUDA, OpenACC and OpenCL programming on heterogeneous multi-to-many core computing platforms.

Several Programming tools focused on activities of devices such as Intel Xeon Phi coprocessors, AMD and NVIDIA GPUs are also covered. The profiling tools supports event monitoring registers, open source tools such PAPI, MPI visualization tools, NVIDIA CUDA 5.5 development tools (Visualization Profiler, CUDA MEMCHK, Debugger), AMD OpenCL Vidual Studio and profiler are discussed.

The **hyPACK-2013** laboratory session softcopy document offers the application users a great opportunity to learn about the fundamentals of writing multi-threaded programs on multi-to-many core computing platforms. Also, different programming paradigms are used to write codes for NLA kernels and application kernels, emphasizing on optimization techniques to extract the performance on cluster of Multi-Core Processor Platforms with different devices. An overview of different accelerators such as FPGA, GPGUs, Intel Xeon-Phi coprocessors are also discussed.

The **hyPACK-2013** laboratory session is provided foundation for application user to implement parallel algorithms for heterogeneous computing platforms such as HPC Cluster computing systems with Intel Xeon-Phi Coprocessors and GPU Accelerators. The workshop proceedings provide a balanced coverage of theoretical and practical aspects on emerging trends in multi-to-may core parallel processing platforms focusing on performance enhancement through software multi-threading.

The **hyPACK-2013** technology workshop CD soft-copy proceedings offer the application users a great opportunity to learn about the fundamentals of writing parallel programs using different programming paradigms on different devices. The participants can easily port the data-parallel and task-parallel computations of their application on many-core platforms. However, tuning and performance of application kernels require extra effort on many-core platforms and techniques are discussed to achieve the performance of application kernels. Participants can understand measurement of power consumption of their applications and extract optimum performance for application kernels on many-core platforms.

C-DAC views the **hyPACK-2013 technology** workshop Proceedings (soft-copy and hardcopy) and the Hands-on session softcopy as a continuously evolving resource on emerging multi-to-many core processing platforms. The soft-copy provides a strong foundation to port and enable many application kernels on message passing clusters with coprocessors and GPU accelerators. Most of the articles of the proceedings include broad coverage of practical aspects of programming on *multi-to-many* core parallel processing platforms and these have been selected from several important research articles, books and web sites. The material is prepared from various references that are included in the web-pages.