# PREFACE

The energy and power density consumption in modern Cluster of High Performance Computing Systems is growing for HPC applications and these efforts led to design power-aware computer architectures. With Power dissipation becoming an increasingly serious problem, the modern ARM processors, GPUs and many-core systems are used for calculation of power consumption and performance of application kernels. The hybrid computing systems with devices such as Intel Xeon Coprocessors, GPUs, ARM multi-core systems with devices are being used. To-day, a major challenge in High Performance Computing (HPC) is to get performance for applications which cans scale upto few hundred core processors and It has become the central theme. It is therefore imperative to set a goal to tackle the new challenges in "HPC-hybrid computing" using software multi-threaded programming on multi-core systems with Intel Xeon-Phi coprocessors and GPU accelerator devices.

Programming environment on Intel Xeon-Phi coprocessors is precisely same as the programming on shared address space platforms as well as message passing clusters. The traditional programming models such as MPI, combination of MPI with Pthreads or Intel TBB or OpenMP 3.0/4.0, Cilk Plus can run on Xeon-host multi-core or Xeon Coprocessor using Compiler Vectorization techniques or offload pragmas.

Programming environments on GPUs are intended to make it easier for developers to take advantage of the parallel processing capabilities of man-core processors. Programming environments (CUDA, OpenACC, OpenCL 4.0 and OpenCL) on GPUs are intended to make it easier for developers to take advantage of the parallel processing capabilities of GPU's. However, the programming models are sufficiently different from traditional parallel programming models (MPI, Pthreads, OpenMP), that taking advantage of Intel Co-processors or GPU's requires the developer to rethink their application to extract maximum performance from system.

We recognize that HPC application users require good exposure to most popular hybrid computing platforms based on Intel Xeon Phi Coprocessors and GPGPUs. We believe that this technology workshop will be useful for future endeavors of application experts. In this regard, we've taken new initiative to bring application users to re-design the existing applications and explore new algorithms and use different programming paradigms that can extract maximum performance on multi-core platform based message passing cluster with different devices.

We believe that it is our commitment to achieve performance improvement based on Multi-Core processor Clusters with Intel Xeon Phi Coprocessors or GPU accelerators that continue to contribute to our success in performance of Scientific and Engineering applications.

To quickly adapt to programming on Heterogeneous computing platforms (HPC GPU Cluster) from the application point of view, **Prof. Rajat Moona**, Director General, C-DAC, India, **Dr. Hemant Darbari**, Executive Director, C-DAC, Pune and **Dr. P.K. Sinha**, Senior Director, HPC, C-DAC, Pune & Corporate Office, C-DAC motivated us, to conduct extensive technology proliferation activities i.e. workshops in collaboration with academic institutions in India as well as private sector IT companies, to understand problems of extracting performance of applications on new emerging Hybrid parallel processing platforms based on HPC devices. **Prof. Ramakrishna Ramaswamy**, Vice-Chancellor, UoH and **Dr. S.K.Udgata** Director, CMSD, UoH encouraged us for such joint collaborative efforts on Parallel Processing activities to address grand challenges in Scientific and Engineering applications and make sincere efforts for scientific discovery.

Witnessing the growing demand for Programming on hybrid computing platforms with Intel Xeon-Phi Coprocessors and GPGPUs as well as ARM Multi-Core Systems, 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 the period **October 15-18, 2013.**

The workshop has originated from C-DAC's experience on High Performance Computing Frontier Technologies Exploration (HPC-FTE) projects as well as HPC Core project and important references are included in this.

To meet this goal, this technology workshop **hyPACK-2013** is designed to provide comprehensive coverage on the state-of-the-art of Hybrid Computing Platforms with Intel Xeon Phi Coprocessors and GPUs as devices in a message passing environment with special emphasis on performance of scientific and engineering applications in a clear and concise manner. The **hyPACK-2013** technology workshop is designed for **four** days on Programming on HPC Clusters with different devices. Participants will get an opportunity to walk-through and execute some of the programs and application kernels designed for various programming paradigms on different computing systems as discussed in sections in the form of **Mode-1, Mode-2, Mode-3, Mode-4, Mode-5** and **Mode-6** of **hyPACK-2013** workshop.

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).

The **hyPACK-2013** 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, Cilk Plus, Intel TBB, and 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), other programming paradigms such as CUDA, OpenCL, Intel Xeon-Phi Offload /native pragmas, OpenMP, Cilk Plus, Intel TBB, and 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.

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 coprocessors and devices. The **hyPACK-2013** laboratory session is provided foundation for application user to implement parallel algorithms on Hybrid Computing platforms such as HPC Cluster computing systems with Intel Xeon-Phi Coprocessors and GPU Accelerators. 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** 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.

We have covered only few out of the large number of programs for numerical linear algebra (NLA) algorithms available on the Internet sites and in books in the

**hyPACK-2013** laboratory. In the course of preparing the **hyPACK-2013** technology workshop soft copy and hardcopy proceedings, we consulted numerous original articles, books, important web-sites, attended many tutorials as listed in references related to certain topics and programs on Multi-Core Processors and HPC GPU Clusters. The reader or user is assumed to have a basic understanding of programming languages, and exposure to Parallel Processing Programming is preferable but not essential. The material is prepared from various references that are included in the web-pages.

We have provided a number of practical exercises in the Hands-on Session with detailed step-by-step explanation that aim to serve as examples to quickly adapt to emerging **HPC Cluster with Coprocessors and GPU Accelerators** based parallel processing platforms in order to solve application problems. We apologize for not being able to acknowledge all the resources in our references due to space and time constraints and for any personal bias in the selection of algorithms and software programs used in the hands-on session. 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. 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.

We would like to thank the participation of **Intel, NVIDIA,** and **AMD.** It was a great pleasure working with the pioneer IT companies. Also, thanks to the technical expert software engineers from **Intel, NVIIDA,** and **AMD** who delivered talks and demonstrated emerging technology products on parallel processing. This workshop would not have been possible without support and encouragement from the Department of Electronics and Information Technology (**DeitY**), Govt. of India, the Indian Space Research Organization (**ISRO**), Govt. of India, Ministry of Earth Sciences (**MoES**) Govt. of India, and Council of Scientific and Industrial research (**CSIR**). Govt. of India

It is been a privilege to work with C-DAC and CMSD staff members and contribute to this workshop. We appreciate the efforts of C-DAC members and CMSD, University of Hyderabad, who put in many long hours and assisted us on a number of occasions in producing the final hardcopy of classroom lectures notes and soft-copy of **hyPACK-2013** proceedings.