Issue 53

GAṆANAM (गणनम्)

April 2025

HIGH PERFORMANCE COMPUTING NEWSLETTER
INSTITUTE FOR PLASMA RESEARCH, GANDHINAGAR

Are GPUs the Future of HPC Workloads?

High-Performance Computing (HPC) plays a pivotal role in solving some of the most complex and computationally demanding problems across a broad range of domains - including molecular dynamics, fluid dynamics, artificial intelligence (AI), etc.. In recent years, Graphics Processing Unit (GPU), originally developed for accelerating image rendering, is now widely used in HPC to speed up workloads that involve large amounts of parallel processing. This article explores the evolving role of GPUs in the HPC landscape, examining the advantages they offer, the challenges they present, etc.

The Growing Dominance and Advantages of GPUs

1) Highly Parallel Architecture

GPUs are built with thousands of lightweight cores that can perform many calculations at the same time. This makes them ideal for high-performance workloads like molecular dynamics, computational fluid dynamics, Particle-in-Cell Simulations etc. where massive parallelism can significantly reduce execution time.

2) Superior Energy Efficiency

GPUs offer more computational output per watt compared to CPUs if code has higher GPU affinity than CPU. In power-hungry HPC environments, such as large-scale clusters, this efficiency helps reduce operational costs and carbon footprint.

3) Improved Software Ecosystem

Programming tools like CUDA, OpenACC, and modern OpenMP versions have made it easier to develop or adapt HPC applications for GPUs. These tools help optimize performance without needing to completely rewrite legacy codes.

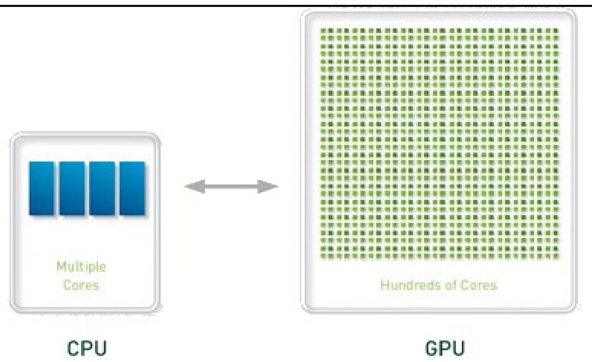

Image Source: *WHAT IS GPU COMPUTING?* Accessed via https://www.boston.co.uk/info/nvidia-kepler/what-is-gpu

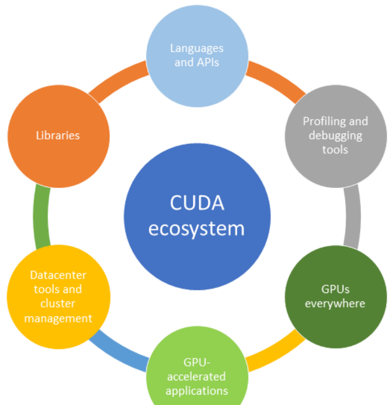

Image Source: *CUDA Refresher: The GPU Computing Ecosystem* Accessed via https://developer.nvidia.com/blog/cuda-refresher-the-gpu-computing-ecosystem/

Current Limitations and Considerations

1) Legacy Code Requires Refactoring

Adapting existing CPU-optimized applications for GPUs often demands significant code modification. For long-established scientific codes, this can be a substantial investment in time and resources. Users may use profiling to identify GPU-suitable hotspots and selectively offload those hotspots. Instead of rewriting the entire program for GPUs all at once, user can start by speeding up just a few parts - step by step using tools like OpenACC or OpenMP by using their directives.

2) Not All Workloads Benefit Equally

Some applications are not inherently parallelizable or are memory-bound. In such cases, GPUs may not offer a performance advantage over well-optimized CPU implementations. Identify the workload suitability by performing scaling studies on GPUs to check compute and memory usage intensity on GPU as compared to CPU.

3) Incomplete Software Compatibility

Although many modern libraries support GPU acceleration, some legacy HPC software packages have yet to be fully ported or optimized for GPUs. Users can identify and engage with open source communities related to the software to stay updated on any recent advancements related to GPU.

While GPUs have become a powerful part of modern HPC, they are not a one-size-fits-all solution. Their strength lies in handling highly parallel tasks like simulations and large-scale modelling. However, not all workloads benefit equally. Memory-bound, sequential, or highly specialized tasks may perform better on CPUs . Porting older code to GPUs can also be resource-intensive. The future of HPC lies in heterogeneous computing - a mix of GPUs, CPUs, for best and optimal system performance.

ANTYA Utilization: MARCH 2025

ANTYA DAILY OBSERVED WORKLOAD

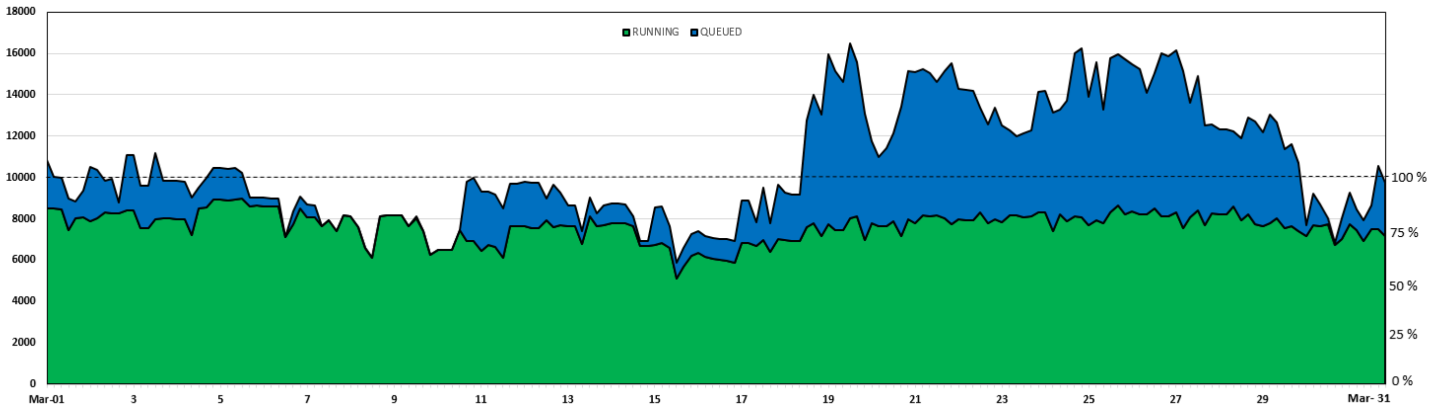

ANTYA HPC Users' Statistics March 2025

Total Successful Jobs~ 1581

♦ **Top Users** (Cumulative Resources)

- **CPU Cores** — Amit Singh
- **GPU Cards** — Suruj Kalita
- **Walltime** — Kaushal Parikh
- **Jobs** — Jugal Chowdhury

ANTYA Usage, Updates and News

- **Scheduled Downtime**: There was no downtime of ANTYA for March 2025.
- **Job Submissions**: The highest job loads were observed in the *serialq*, regular, *mediumq*, and *longq* queues, reflecting sustained user activity across multiple workloads in various queues.
- **Cluster Utilization**: The system maintained an average utilization of approximately 79%.
- **Packages/Applications Installed**: Ion Beam Simulator (ibsimu) module has been installed on ANTYA. To use ibsimu in ANTYA:
  $ module load ibsimu/1.0.6

Other Recent Work on HPC

Thermal design and optimization of an ECRH launcher mirror for long pulse operation	Hardik Girishbhai Mistry
Thermo-hydraulic and artificial intelligence-based off-design performance characterization of a	Dr. Manoj Kumar
On the Transition Mechanism to Multi mode Operational Regime in Rotating Detonation Engine Operation	Sunil Bassi
Excitation of Non-planar Pinned Solitons in a flowing Dusty Plasma	Prasanta Amat
Study of Reversed Magnetic Shear Configuration in ADITYA-U TOKAMAK	Gopal Krishna M
3-Dimensional Vacuum Field Modeling and Edge Plasma Response to Applied Radial Magnetic Perturbations of ADITYA-U TOKAMAK	Ananya Kundu
Plasma Transport Study with 3D Shaped First Wall For Limiter Ramp-Up Phase of ITER	Arzoo Malwal
Influence of Ion Temperature on the Dynamics of Unidirectional Current Carrying Filamentary ELM BLOBS in the Edge Region of a Tokamak	Souvik Mondal
Study of radial field dependent flows in inboard limited Aditya-U plasmas using EMC3-Eirene simulations	Arzoo Malwal
Development of AI Framework for Plasma Equilibrium Parameters Generation for virtual Tokamak environment	Agraj Abhishek

Acknowledgement

Join the HPC Users Community
hpcusers@ipr.res.in
If you wish to contribute an article in GAṆANAM, please write to us.

GAṆANAM (गणनम्)

Contact us
HPC Team
Computer Division, IPR
Email: *hpcteam@ipr.res.in*