ISSUE 40 MARCH 2024

GAṆANAM (गणनम्)

HIGH PERFORMANCE *COMPUTING* NEWSLETTER

INSTITUTE FOR PLASMA RESEARCH, INDIA

INSIDE THIS ISSUE
(3 Pages)

A special case of superiority of a GPU solver over a CPU solver: A Cost vs Accuracy test

Shishir Biswas (Research Scholar, PDTSD, IPR)
Email: shishir.biswas@ipr.res.in

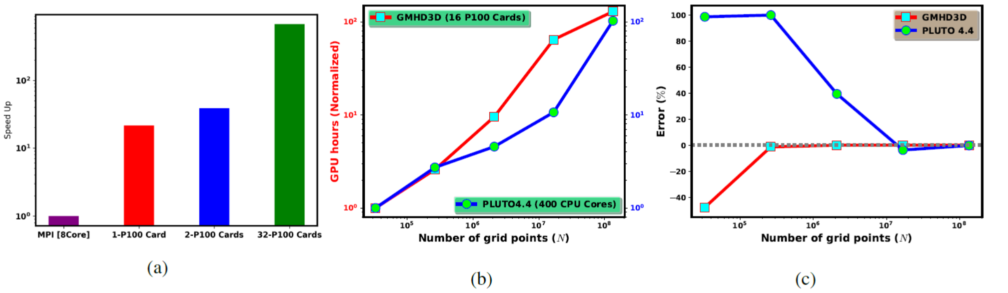

Fig. 1 : (a) Speed Up of GMHD3D code. (b) Cost metric comparison (c) Comparison of cost vs accuracy for GMHD3D and PLUTO4.4

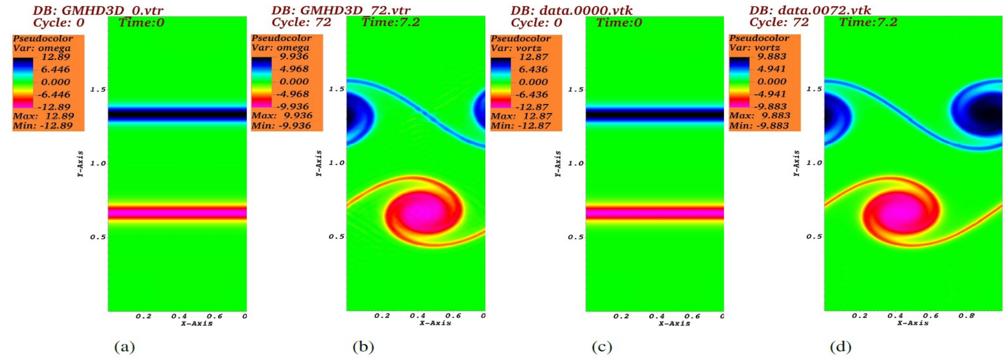

Fig. 2: Time evolution of vorticity for two oppositely directed KH unstable jets (broken jet) from GMHD3D code [(a & b)] and PLUTO4.4 code [(c & d)]

Plasma is a collection of charged particles that is frequently represented as a fluid on which electromagnetic body forces operate. It has been observed that a spatially averaged model termed the "fluid model" is particularly efficient in predicting the behavior of the plasma when a large number of charged particles is present. Therefore, Maxwell's equations coupled with the equations of hydrodynamics become the primary governing equations for the motion of the charged-fluid element in the presence of an self-generated electric and magnetic field. The subject that studies the self-consistent evolution of such a magnetized plasma fluid is known as MagnetoHydroDynamics (MHD). Because of nonlinear interactions across different scales of length, energy can cascade through different modes in a fully formed turbulent plasma medium. It has long been a problem in fluid dynamics to characterize the nature of the cascade of kinetic energy for a given initial spectrum. Understanding plasma turbulence is essential for controlling the disruption of plasma in such experimental devices, thereby enhancing plasma confinement for fusion plasmas and enabling the prediction of extreme events in astrophysical objects and stellar matter. An important challenge in astrophysical plasmas is the creation of multi-scale magnetic fields, which occurs in the Sun, newborn stars, accretion disks, and other astronomical entities. The "Dynamo Theory" of Parker [1] is one of the first explanations for the generation of such magnetic multi-scale fields. Such large or intermediate-scale magnetic field is generated at the expense of the plasma kinetic energy, which primarily governs the dynamics of the charged fluid (plasma) via a time-dependent Lorentz force (back-reaction) term added to the Navier-Stokes equation, thereby self-consistently influencing the dynamics of the fluid flow. In general, one needs to solve the set of coupled partial differential equations in order to deal with the complex astrophysical MHD phenomena outlined above. For this reason, high-performance numerical solvers are required to accurately represent the physics problems occurring in plasmas. In order to simulate the MHD systems on a wide scale, including astrophysical entities and laboratory scenarios, it is necessary to develop highly scalable codes.

> ***Upgradation of a 3D compressible single GPU MHD solver to a multi-node, multi-card GPU architecture (GMHD3D), resulted in a remarkable 675.5x speedup across 32 P100 GPU cards compared to the MPI version and a 32x speedup compared to the single-GPU version.***

At the Institute for Plasma Research [IPR], INDIA, we have recently upgraded an already existing threedimensional compressible single GPU MHD solver to multi-node, multi-card GPU architecture [GMHD3D] [2, 3] using OpenAcc & MPI. After multi-GPU upgrade, we obtain a 675.5x speedup across 32 P100 GPU cards in comparison to the MPI version, and a 32x speedup in comparison to the single-GPU version (See Fig. 1a) [2, 3]. The solver currently employs OpenMPI/4.0.1 for its multi-node communication and the AccFFT library for FFT operations. PyEVTK, a data converter (ASCII to BINARY) written in Python, is designed to dump data in VTK binary format for the sake of visualization. We also provide a comprehensive comparison between the aforementioned in-house pseudo-spectral MHD solver (GMHD3D) [3] and an open source grid based MHD solver PLUTO4.4 [4] for some specific physics problems. The primary goal of current investigation is to validate the precision of the recently developed GPU solver and to compare the superiority of a pseudo-spectral solver to that of a grid-based solver atleast for certain class of physics problems. All simulations using the PLUTO4.4 and GMHD3D codes were executed on the 1 PetaFlop ANTYA supercomputer located at the Institute for Plasma Research in India. PLUTO4.4 utilizes a dual configuration of 20 CPU cores, namely the Intel Xeon 6148 model, operating at a clock speed of 2.4 GHz. The system is equipped with a total of 384 GB DDR4 RAM. For simulations utilizing GMHD3D code we have used GPU nodes of ANTYA cluster with similar specification along with two NVIDIA tesla P100 GPU cards in a single node with 16 GB RAM each.

To conduct a cost metric comparison between the two solvers, a series of simulation runs have been performed, varying the number of resources (CPUs and GPUs) and grid points. We have plotted the normalized computational costs for the GPUs and CPUs in relation to the grid resolutions (refer to Fig. 1b). From Fig. 1b it is readily understood that the computational expenses increases linearly for both CPUs and GPUs as grid points increases. It can be observed from Fig. 1b that the normalized computational time of 16 GPUs is nearly similar to the computational time of 400 CPUs in cases when the computational workload is significant. It is widely acknowledged that CPUs exhibit a higher power consumption, as compared to CPUs, GPUs which have a lower power consumption because of their shared memory architecture. Therefore, the maintenance of 400 CPUs would result in higher computational expenses, including electrical power consumption, cooling, and rack space, as compared to the maintenance of 16 GPU cards. This, in turn, signifies the cost-effectiveness of the GPU solver [3].

The calculation of the accuracy (value of error %) is determined by the following formula.

$$\text{Error}(\%) = 100 * (W1-W2) / W1$$

where, W1 represents the original expected value, while W2 represents the value observed in the numerical simulation. In order to determine the accuracy, expressed as a percentage of error, we have considered an individual test problem that was utilized for the purpose of conducting a cost comparison analysis. The difference in peak values are calculated for various grid points using both codes. We have plotted the value of percentage of error (Error (%)) for the GPU solver (GMHD3D) and CPU solver (PLUTO 4.4) in relation to the grid resolutions (see Fig. 1c). It is evident from Fig. 1c that the GPU solver (GMHD3D) converged rapidly to errors of the order of less than 1%, whereas PLUTO 4.4 requires a more precise grid (higher grid resolution) in order to converge. Fig. 1c demonstrates that the GMHD3D code fulfills the accuracy criteria at the lowest cost, whereas the PLUTO 4.4 code is the most expensive [3]. This indicates the cost-effectiveness and accuracy of the GPU solver [3] compared to the CPU solver being discussed. We have considered some well-known test problems in two- and three-dimensional hydrodynamics and magnetohydrodynamics to accomplish this for the benchmarking purpose. For example, 2-dimensional Kelvin-Helmholtz instability (See Fig. 2), Dynamics of 3-dimensional Taylor-Green (TG) vortex (See Fig. 3), Coherent nonlinear oscillations using 2D Orszag-Tang (OT) Flow (See Fig. 4), Recurrence dynamics in 3D MHD plasma (See Fig. 5) and so on. Our numerical analysis shows that while both algorithms produce comparable answers in most circumstances, the spectral solver surpasses the grid-based solver in periodic domain for a subset of physics-related challenges [3]. We also believe this work highlights the advantages of a spectral solver over a grid-based solver. To the best of our knowledge, this is the first work [3] ever attempted which makes a thorough comparision of a pseudo-spectral code with the freely available grid based MHD solver PLUTO4.4.

References:

1. *Parker, Eugene N. Astrophysical Journal, vol. 122, p. 293 122 (1955): 293*
2. *S Biswas, R Ganesh et al. "GPU Technology Conference (GTC-2022)", https://www.nvidia.com/en-us/ondemand/session/gtcspring22-s41199/*
3. *Biswas, Shishir, and Rajaraman Ganesh. Computers & Fluids (2024): 106207.*
4. *Mignone, Andrea, et al. The Astrophysical Journal Supplement Series 170.1 (2007): 228 .*

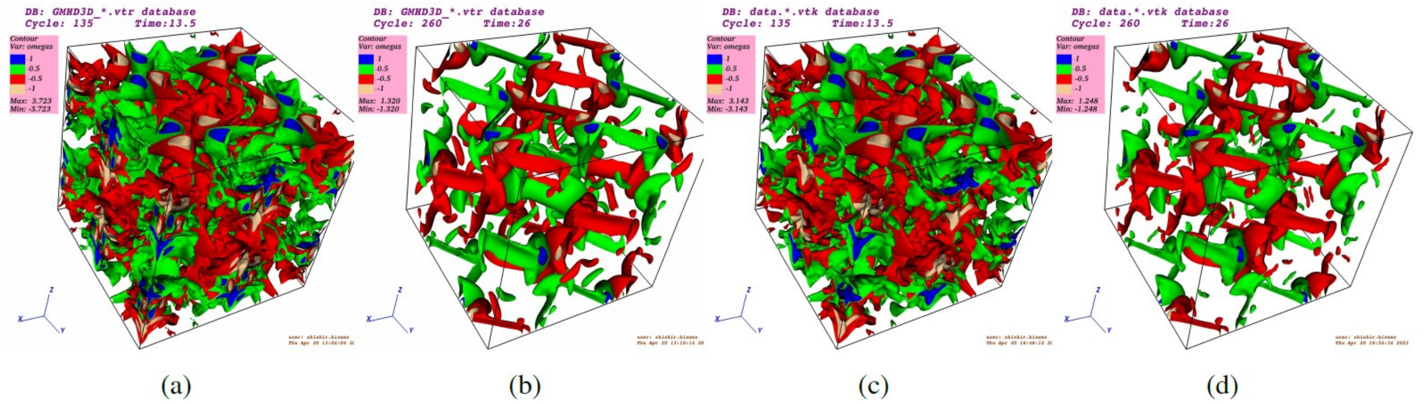

Fig. 3 : The turbulent dynamics of 3D Taylor-Green flow from GMHD3D code [(a & b)] and PLUTO4.4 code [(c & d)]

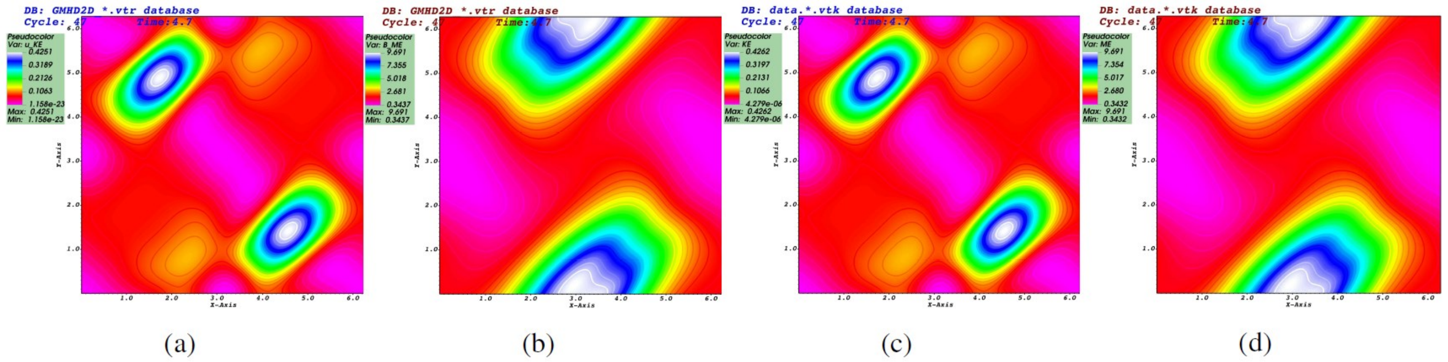

Fig 4: 2D Orszag-Tang Flow kinetic and magnetic energy contour from GMHD3D code [(a & b)] and PLUTO4.4 code [(c & d)].

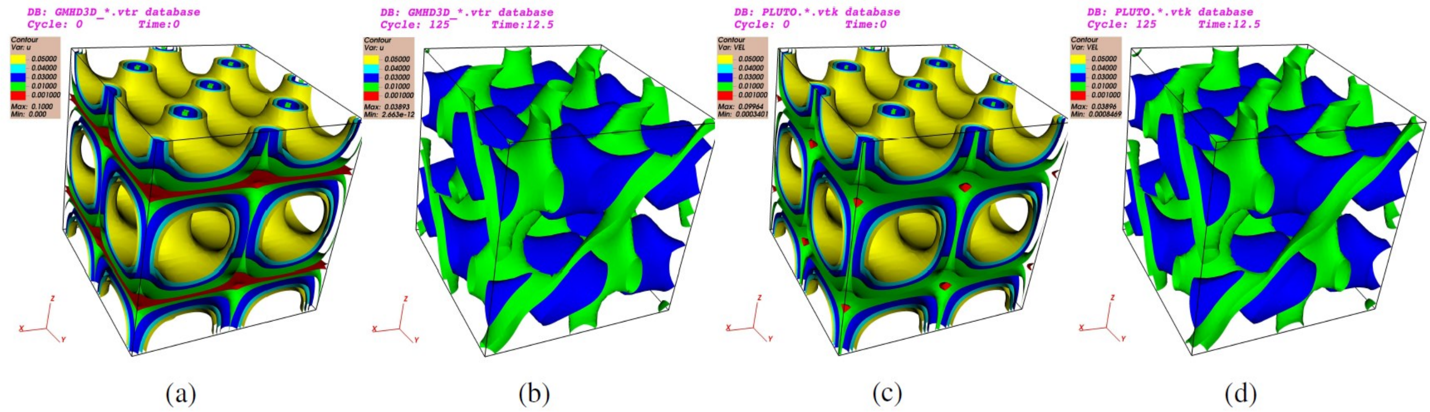

Fig 5: Recurrent dynamics from GMHD3D code [(a & b)] and PLUTO4.4 code [(c & d)].

Utilizing Ready-Made Containers with Singularity: Part 1 General Overview

High-Performance Computing (HPC) environments often require complex software configurations and dependencies. To simplify the deployment of applications in these environments, containerization has become increasingly popular. Docker is a widely used containerization platform with a vast repository of pre-built containers on Docker Hub. However, in HPC clusters, Singularity is often preferred due to its ability to seamlessly integrate with the existing environment. In this article series, we will explore how to leverage ready-made containers from Docker Hub using Singularity in HPC.

This series has been divided into 3 parts, covering 1 part in every issue:

- **Part-1: General Overview**
- **Part-2: Unlocking HPC Potential: A Practical Tutorial on Deploying Docker Hub Containers with Singularity**
- **Part-3: Simplifying TensorFlow Work: Run Jupyter Notebooks Inside Pre-built Containers**

ANTYA Utilization: February 2024

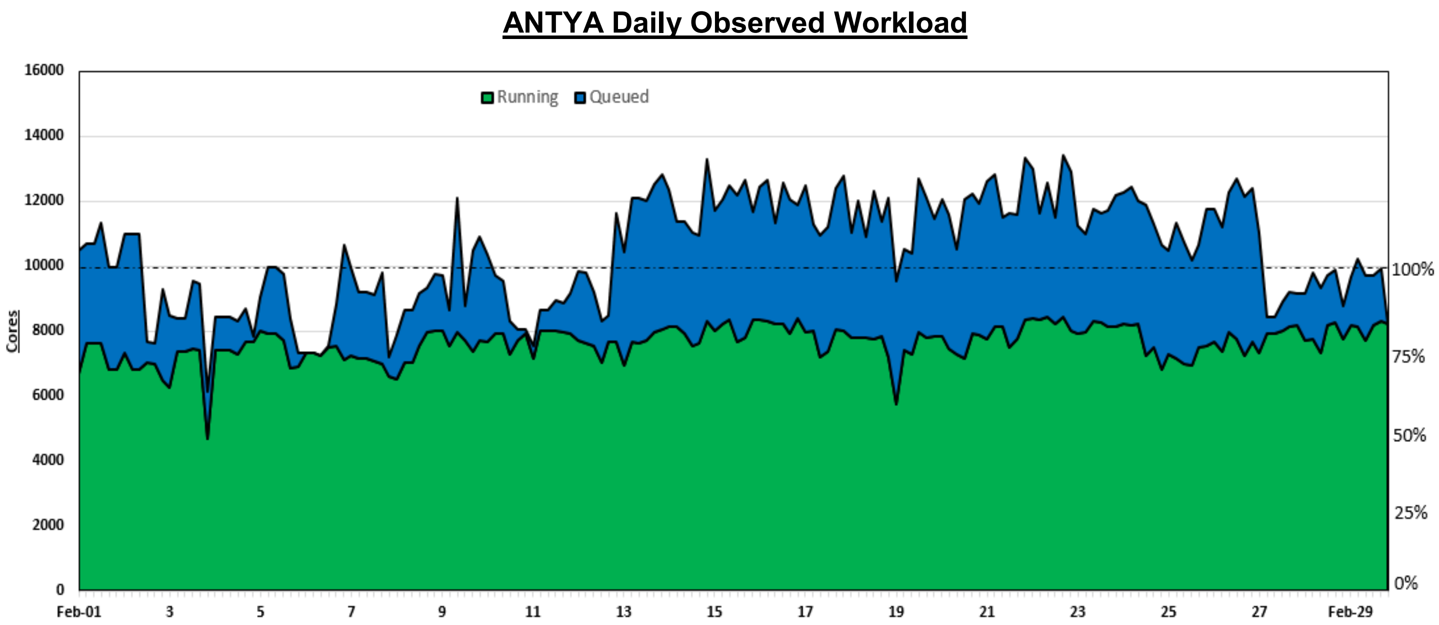

Other Recent Work on HPC (Available in IPR Library)

Aggregate morphing of self-aligining soft active disks in semi-confined geometry	Anshika Chugh
Design and Development of 300 kV, 2 A DC Power Supply for High Power and High Energy Accelerator Based Applications	Ashok D. Mankani
Development of a highly sensitive electromagnetic flowmeter for high temperature liquid metals	Srikanta Sahu
Numerical validation of Yukawa fluid excitations within the Quasi-localized charged approximation theory	Prince Kumar
Design and Analysis of the Components of the Cryogenic Extruder for Producing the Liquid Hydrogen	Vishal Gupta
Design of high-frequency pulse tube cryocooler for cryo-pump applications	Sarvesh Kashyap

ANTYA HPC USERS' STATISTICS— FEBRUARY 2024

Total Successful Jobs~ 2073

♦ **Top Users** (Cumulative Resources)

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

Acknowledgement

The HPC Team, Computer Division IPR, would like to thank all Contributors for the current issue of *GAṆANAM*.

On Demand Online Tutorial Session on HPC Environment for New Users Available
Please send your request to ***hpcteam@ipr.res.in.***

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

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