

Inside this issue:

(4 pages)

Plasma simulation on ROBIN Ion Source geometry 1

Part 2: Running LAMMPS in Containers – Multi-Node, Multi-GPU with PBS 2

ANTYA Usage Statistics and Observed workloads - Sep 2025 4

In large fusion devices such as ITER and DEMO, the Negative Ion-based Neutral Beam Injector (NNBI) is a vital system for plasma heating and current drive. Each NNBI ion source in ITER is designed with eight inductively coupled plasma (ICP) drivers, powered by a 1 MHz RF generator, to generate negative ions for the beam.

ITER-India at IPR is engaged in building and characterizing the Diagnostic Neutral Beam (DNB) at the India Test Facility (INTF), which serves as a full-scale test bed before deployment in ITER. To develop operational expertise with high-power RF ion sources, IPR operates two dedicated systems, ROBIN (with a single RF driver) and TWIN (with two RF drivers). For ROBIN, detailed simulation studies are being carried out using COMSOL Multiphysics (Plasma Module) to understand the mechanism of RF power coupling to the plasma and its dynamics, which strongly affects the efficiency and stability of negative ion production for the ion beam. The simulation was performed on the ANTYA high-performance cluster at IPR using a 16-core CPU in visualization mode, which greatly reduced computation time. Each run finished in about 45 minutes, compared to 5–8 hours on a single-core system. This facility enables faster and more efficient plasma simulations at IPR. The simulation domain on the ROBIN ion source and its driver is shown in Fig. 1. The size of the ROBIN ion source is one-eighth of the size of ITER's negative neutral beam injector (NNBI) source.

In the present simulation work, Maxwellian EEDF is utilized. This function is significant because it affects the interactions of electrons with other plasma particles, their energy absorption from the electric fields, and their contributions to ionization, excitation, and radiation. A cross-section file containing all the reaction cross-sections corresponding to ionization, different excitations, elastic collision, etc., for a range of electron energies (0.1eV to 100eV) is included in the simulation. Integrating the cross-section data allows the simulation to precisely predict the rate coefficients that govern the formation and loss of different charged particle species in the plasma. The simulation for the ROBIN ion source is investigated in two orthogonal 2D planes. The reasons for choosing two orthogonal planes for the simulation are computational time for 3D and the design structure of the ROBIN ion source. The ion source geometry is



Plasma simulation on ROBIN Ion Source geometry

Ram Swaroop (PDF, NNB Systems Development Division, IPR)
Email: ram.swaroop@ipr.res.in




Figure 1: Cross-sectional view of the ROBIN ion source and its essential components.




Figure 2: 2D-plane study of plasma parameters (Electron density, Electron Temperature, and Plasma Potential) using COMSOL Multiphysics.






Figure 3: Variation in plasma parameters (maximum values) (Electron density, electron temperature, and Plasma potential) with Pressure (0.3-1.0) Pascal, 50kW RF power keeping Ar gas temperature 1000K.

such that the upper part, i.e., the driver section, is cylindrical, while the lower half is of race-track geometry.

The initial work is started with argon plasma with the Faraday shield in the RF driver, a set of external permanent magnets on the wall for plasma confinement, and also a set of filter field magnets. The profiles of plasma parameters with 50kW RF power, 0.3Pa Argon gas pressure are shown in Fig. 2. The simulation is performed for different Ar pressure conditions (0.3-1.0 Pa). The outcomes of electron density, electron temperature, and plasma potential for different pressure conditions are recorded and plotted as shown in Fig. 3. The influence of the Faraday shield and the external DC magnetic field profile generated by all the permanent magnets is observed in the simulated electron density, electron temperature, and plasma potential across a pressure range of 0.3 to 1.0 Pa.

This argon-based study serves as a learning platform, enhance expertise in plasma modelling and preparing for the complexities of hydrogen plasma such as additional volume reactions (e.g., ionization, dissociation, and recombination) and wall interactions (e.g., secondary electron emission and surface recombination). These sophisticated phenomena will be modelled using the plasma module in COMSOL Multiphysics, and simulated on ANTYA HPC Cluster in future again for actual 3D-geometry of ROBIN source.

References

- [1] Chen, F. F. (1984). Introduction to plasma physics and controlled fusion (Vol. 1, pp. 19-51). New York: Plenum press.
- [2] Bandyopadhyay, M. (2004). Studies of an inductively coupled negative hydrogen ion radio frequency source through simulations and experiments; PhD thesis.
- [3] Bansal, G., Gahlaut, A., Soni, J., Pandya, K., Parmar, K. G., Pandey, R., ... & Parmar, D. (2013). Negative ion beam extraction in ROBIN. Fusion Engineering and Design, 88(6-8), 778-782.
- [4] Bandyopadhyay, M., Singh, M. J., Pandya, K., Bhuyan, M., Tyagi, H., Bharathi, P., ... & Chakraborty, A. K. (2022). Overview of diagnostics on a small-scale RF source for fusion (ROBIN) and the one planned for the diagnostic beam for ITER. Review of Scientific Instruments, 93(2).

Part 2: Running LAMMPS in Containers – Multi-Node, Multi-GPU with PBS

In [Part-1 of this article series](#), the idea of containers in HPC environments was introduced and demonstrated how a handful of **essential Singularity/Apptainer commands** can help researchers become productive quickly. The article covered how to:

- Pull container images from trusted sources such as DockerHub or NVIDIA NGC.
- Run and execute commands interactively inside containers.
- Bind host directories for data access.
- Enable GPU and MPI support for scientific workloads.

These foundational steps are critical because they allow HPC users to become acquainted with containers before moving on to more complex, real-world applications.

The Part 2 of this article series will move from **theory to practice** by working with molecular dynamics package: **LAMMPS** (Large-scale Atomic/Molecular Massively Parallel Simulator). NVIDIA provides an optimized LAMMPS container in the **NGC Catalog**, preconfigured with **OpenMPI and CUDA support**. This article will cover how to:

- Pull and run the LAMMPS container from NGC.
- Test single-node execution with GPU acceleration.
- Launch a **multi-node, multi-GPU container job** using the PBS scheduler.

This example highlights the real benefits of containerization in HPC: portability, reproducibility, and scalability — all without manually compiling complex software stacks.

1) Pull the LAMMPS Image from NGC Catalog (<https://catalog.ngc.nvidia.com/orgs/hpc/containers/lammps/tags>)

```
## This will create lammps_15Jun2020.sif

[user@login1 ~] $ module load singularity/3.5.3/3.5.3

[user@login1 ~] $ singularity pull docker://nvcr.io/hpc/lammps:15Jun2020

INFO:   Converting OCI blobs to SIF format
WARNING: 'nodev' mount option set on /tmp, it could be a source of failure during build process
INFO:   Starting build...
Getting image source signatures
Copying blob 59492a727a6b done
Copying blob d5c21de1eb64 done
Copying blob c4254b053e79 done
...
2025/09/02 11:51:03 info unpack layer: sha256:70ab4bbd07c794e0c99ac789728dde92685005e7dd59810d42466074bc7ed095
2025/09/02 11:51:03 info unpack layer: sha256:0a5fa0c17551240730ffe538ecb1c67b8433384a03433ac3ffe2fb4094077db4
INFO:   Creating SIF file...
INFO:   Build complete: lammps_15Jun2020.sif
```

2) Check MPI and CUDA Support Inside the Container

```
## Submit an interactive job to verify the details of the MPI and CUDA

[user@gn01 ~] $ singularity shell --nv lammps:15Jun2020

Singularity> which mpirun
/usr/local/openmpi/bin/mpirun

Singularity> nvidia-smi
```

NVIDIA-SMI 460.32.03		Driver Version: 460.32.03		CUDA Version: 11.2	
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile Uncorr. ECC
Fan	Temp	Perf	Pwr:Usage/Cap	Memory-Usage	GPU-Util Compute M.
					MIG M.
0	Tesla P100-PCIE...	On	00000000:18:00.0	Off	0
N/A	34C	P0	26W / 250W	MiB / 16280MiB	0% Default
					N/A

3) Running LAMMPS on a Single Node

Submit an interactive job and verify the details of the MPI and CUDA

```
[user@gn02 ~] $ module load singularity/3.5.3/3.5.3
```

```
[user@gn02 ~] $ singularity shell --nv lammps:15Jun2020
```

```
Singularity> mpirun -np 4 /usr/local/lammps/sm60/bin/lmp -k on g 2 -sf kk -pk kokkos cuda/aware on neigh full comm device -in in.nemd
```

INFO: Converting OCI blobs to SIF format

LAMMPS (15 Jun 2020)

KOKKOS mode is enabled (src/KOKKOS/kokkos.cpp:85)

will use up to 2 GPU(s) per node

using 1 OpenMP thread(s) per MPI task

Lattice spacing in x,y,z = 1.6796 1.6796 1.6796

Created orthogonal box = (0.0 0.0 0.0) to (167.95961913825073 167.95961913825073 167.95961913825073)

1 by 2 by 2 MPI processor grid

Created 4000000 atoms

create_atoms CPU = 0.488 seconds

...

Per MPI rank memory allocation (min/avg/max) = 160.6 | 160.6 | 160.6 Mbytes

Streaming Multiprocessor (SM) refers to architecture version of NVIDIA GPUs. ANTYA have P100 GPU Cards which have sm60 architecture and compute capability 6.0.

Note: Some Environment variables required to run lammps might not be set by default. So user can export them before using the executable.

```
[shivam.patel@gn02 ~]$ nvidia-smi
Tue Sep  2 16:44:11 2025
```

NVIDIA-SMI 460.32.03 Driver Version: 460.32.03 CUDA Version: 11.2									
GPU Fan	Name Temp	Persistence-M Perf	Pwr:Usage/Cap	Bus-Id	Disp.A Memory-Usage	Volatile GPU-Util	Uncorr. Compute M.	ECC MIG M.	
0	Tesla	P100-PCIE...	On	00000000:18:00.0	Off				0
N/A	37C	P0	41W / 250W	656MiB / 16280MiB		82%	Default		N/A
1	Tesla	P100-PCIE...	On	00000000:AF:00.0	Off				0
N/A	38C	P0	44W / 250W	656MiB / 16280MiB		81%	Default		N/A
Processes:									
GPU	GI ID	CI ID	PID	Type	Process name	GPU Memory Usage			
0	N/A	N/A	63774	C	...local/lammps/sm60/bin/lmp	327MiB			
0	N/A	N/A	63776	C	...local/lammps/sm60/bin/lmp	327MiB			
1	N/A	N/A	63775	C	...local/lammps/sm60/bin/lmp	327MiB			
1	N/A	N/A	63777	C	...local/lammps/sm60/bin/lmp	327MiB			

Fig 1: Lammps Running on GPU verified using nvidia-smi




Fig 2: Block Diagram of Lammps execution from container

3) Multi-Node, Multi-GPU PBS Job Script

Submit an interactive job on Multi Nodes having GPU and submit the test job

Additional Environment needs to be set to allow Multi Node Multi GPU execution

```
[user@login1 ~] $ qsub -I -l select=2:ncpus=40:ngpus=2 -q debugq
```

```
[user@gn02 ~] $ module load singularity/3.5.3/3.5.3
```

```
[user@gn02 ~] $ singularity shell --nv lammps_15Jun2020.sif
```


```
Singularity> export OMPI_MCA_orte_launch_agent="export PATH=/home/application/singularity353/bin:$PATH && singularity run --nv $PWD/${SINGULARITY_NAME} orted"
```

```
Singularity> mpirun -np 64 /usr/local/lammps/sm60/bin/lmp -k on g 2 -sf kk -pk kokkos cuda/aware on neigh full comm device -in in.nemd
```

In conclusion, containers are rapidly becoming an integral part of the HPC landscape, offering researchers a powerful way to package and share complex software environments. By encapsulating applications and their dependencies, containers make workflows far more portable and reproducible, reducing the challenges often faced when moving between HPC systems or integrating cloud resources. While it is true that containerized applications may not always match the raw performance of natively compiled software loaded through environment modules, the gap is usually acceptable in practice. The additional abstraction layer may introduce minor overheads, particularly in communication-heavy or highly optimized jobs. Portability, environment consistency, and reduced maintenance efforts brings advantages to both users and HPC admin. As technologies like Singularity/Apptainer continue to mature, containers will increasingly serve as a pillar of sustainable, scalable, and future-ready scientific computing.

ANTYA Utilization: SEPTEMBER 2025

ANTYA DAILY OBSERVED WORKLOAD



ANTYA HPC Users' Statistics
September 2025

Total Successful Jobs~ 1001

> Top Users (Cumulative Resources)

- CPU Cores [Gopal Mailapalli](#)
- Walltime [Ritesh Srivastava](#)
- Jobs [Kaushal Parikh](#)

ANTYA Usage, Updates and News

- **Scheduled Downtime:** There was no downtime of ANTYA for September 2025.
 - **Job Submissions:** The highest job loads were observed in the *regularq*, *mediumq*, *serialq*, *longq* and *ansysq* queues, reflecting sustained user activity across multiple workloads in various queues.
 - **Cluster Utilization:** The system maintained an average utilization of ~64.40% and peak utilisation of ~90%.
- Packages/Applications Installed:** No new modules have been installed this month. To view list of available modules.
- > module avail

Other Recent Work on HPC

Quasi-longitudinal Whistler Wave Propagation and Ion Response: A Two-Fluid Simulation Study	Gayatri Barsagade
Dynamic Analysis of Rotor used for Pulse Power Application	Vishal Verma
Resonant Flux and NTV Torque Response to Localized Magnetic Perturbation in the Edge Plasma of ADITYA-U	Ananya Kundu
Estimation of shutdown dose rates in shield mock-up benchmark experiments using OpenMC	Shailja Tiwari
Analysis of Loop and Reflected Voltages in SST-1 Startup Scenarios using Filament Model	Kushagra Nigam
Simulation studies for liquid metal flow through sudden expansion subjected to inclined magnetic field	Arpita Vipat
Interaction and propagation characteristics of precursor solitons in a dusty plasma	Prasanta Amat
Development of an ECRH Launcher Mirror for Long Pulse Operation	Hardik Girishbhai Mistry
Plasma Electron Interaction with Oscillating Sheaths: A Study of Stochastic Effects in CCP Discharges	Rishabh Singh
Numerical Simulation of Disrupting Plasma Discharges using Tokamak Simulation Code	Sunil Bassi
Study of Temperature Gradient Driven Microinstabilities in the Ion and Electron Scales in a Tokamak	Buha Ritesh Arjanbhai
Plasma Based Centrifugal Separation of Cu ³⁺ and Au ³⁺ ions in Crossed Fields Using PIC Simulation – Preliminary Observations	Manu Bajpai
Investigation of Copper to Stainless Steel Friction Welds: Experiments, In-Situ Temperature Measurement, and Simulation	Tapan Patel

Acknowledgement

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

Join the HPC Users Community

hpcusers@ipr.res.in

If you wish to contribute an article in *GANANAM*, please write to us.

GANANAM (गणनम्)

Contact us

HPC Team (hpcteam@ipr.res.in)
Computer Division, IPR

Disclaimer: “*GANANAM*” is IPR’s informal HPC Newsletter to disseminate technical HPC related work performed at IPR from time to time. Responsibility for the correctness of the Scientific Contents including the statements and cited resources lies solely with the Contributors.