Received: XX.XX.2025
Received in revised form: XX.XX.2025
Accepted: XX.XX.2025
https://doi.org/10.62020/svet.str.asXXXXXXX

Using High-Performance Computing for Engineering Innovation: The HPCFS Infrastructure and 2025 Upgrade

Pavel Tomšič, Leon Kos, Matjaž Šubelj

University of Ljubljana, Faculty of Mechanical Engineering, Aškerčeva 6, 1000 Ljubljana (pavel.tomsic@fs.uni-lj.si)

High Performance Computing (HPC) has become a key enabler of innovation in modern engineering. At the Faculty of Mechanical Engineering, University of Ljubljana, the HPCFS infrastructure provides researchers and students with the computational power necessary for advanced simulations, modelling, and data analysis. This paper presents the motivation for using HPC in engineering, an overview of the HPCFS system, and the planned 2025 upgrade. The selected projects demonstrate how HPC supports complex analyses ranging from plasma simulations to AI-driven modelling.

Introduction

High Performance Computing (HPC) enables engineers and scientists to solve complex problems that exceed the capabilities of conventional computers. In mechanical engineering, HPC allows advanced modelling, simulation, and optimization of systems with high computational demands. At the Faculty of Mechanical Engineering, University of Ljubljana, the HPCFS infrastructure provides access to large-scale parallel computing resources that support both research and education.

General about HPC cluster

HPC clusters are specialized computing systems designed to solve complex scientific and engineering problems that exceed the capabilities of standard desktop computers. As presented in Figure 1, by linking multiple compute nodes through a high-speed interconnect, HPC clusters operate as a single unified system, allowing researchers to perform parallel computations on large datasets or highly detailed numerical models. Each node typically contains several multicore processors and large memory banks, while the network fabric ensures efficient communication between processes running across nodes.

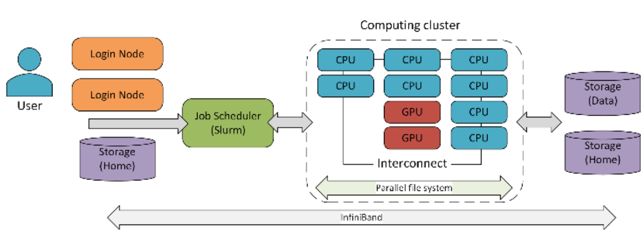

Figure 1: Generalised HPC Cluster Architecture

HPC clusters serve as a cornerstone of innovation. They enable a wide range of applications, from thermomechanical modelling, computational fluid dynamics (CFD), finite element analysis (FEA), to modern data analytics, and artificial intelligence. The ability to simulate physical systems at high resolution reduces the need for experiments and accelerates the development of new technologies. HPC thus provides both scientific insight and industrial competitiveness.

Presentation of HPCFS with 2025 upgrade

The continuous growth of computational requirements makes regular upgrades of HPC infrastructure essential. The 2025 upgrade, presented in Figure 2, will ensure the platform remains capable of supporting advanced simulations, data-driven research and modern AI workloads.

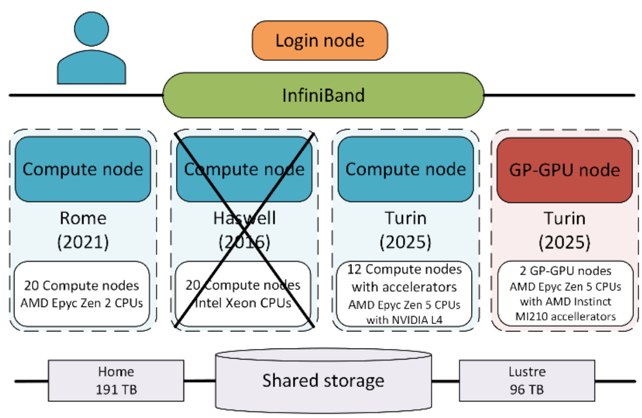

Figure 2: Upgraded HPCFS cluster setup

The HPCFS system consists of two main compute partitions (Haswell and Rome) and will be upgraded in 2025 to include new CPU and GPU nodes, fast storage, and enhanced InfiniBand interconnect. The upgrade will improve computational performance and support for heterogeneous workloads. Software updates will include containerization, Spack-based package management, and improved user tools.

Hardware

The upgraded HPCFS will feature next-generation compute nodes with high-core-count CPUs and dedicated GPU accelerators, connected through a high-speed, low-latency interconnect. A tiered storage system with an expanded parallel file system, will offer both speed and capacity.

Software environment

The system will run a current Linux distribution with the Slurm scheduler for resource management and GPU job execution. Software deployment will be handled via module systems, complemented by container support using Docker. This setup simplifies application management and improves portability and reproducibility of workflows.

Expected performance and user experience

The upgrade will result in better computational and storage performance. Workloads, particularly in CFD, FEA, and materials modelling will benefit from reduced runtimes and improved scalability. For users, this will mean faster job execution, more efficient data handling, and access to GPU acceleration for modern AI and machine learning tasks.

Methodology of using HPC

The effective use of HPC resources, highlighted in Figure 3, requires a structured approach that combines technical understanding, efficient workflow and resource management. HPCFS system provides an environment that supports these principles, enabling users to run complex simulations and analyses in an optimized and reproducible way.

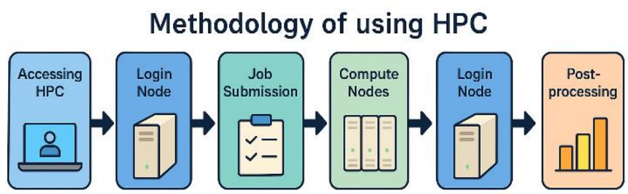

Figure 3: Generalised workflow of using an HPC

Access and user environment

Access to HPCFS is granted to registered users, including researchers, students, and project partners. Each user receives a secure account for login to a Linux-based environment with preinstalled compilers, libraries, and applications. Software modules enable easy loading of required versions and dependencies.

Job submission and scheduling

Jobs are submitted through the Slurm workload manager, which allocates resources based on job requirements and system load. Users define submission scripts specifying the number of nodes, cores, memory, and runtime. Parallel applications typically use MPI for distributed and OpenMP for shared-memory computing, while GPU workloads run on dedicated GPU partitions.

Data management and workflow optimization

Input and output data are stored on a high-performance Lustre file system, while long-term user data are kept on network-attached home directories. This separation ensures both high throughput and data security.

To support reproducibility and portability, HPCFS enables containerized environments. Containers allow to bundle software dependencies into self-contained images, ensuring consistent execution across different systems and projects.

Training and support

Recognizing that efficient HPC usage depends on user knowledge, HPC training sessions, workshops, and tutorials are regularly organized covering topics such as job submission, parallel programming, and performance optimization.

Examples of projects on HPCFS

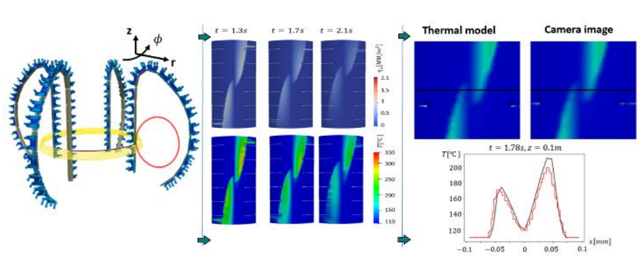

Figure 4: Excellerat P2 – Thermal simulations for Tokamak reactor using HPC [1]

The HPCFS infrastructure supports a diverse range of national and European research projects, providing the computational foundation for advanced simulations, modelling, and training in high performance computing.

Below are selected examples that illustrate the variety of its applications and collaborations:

• **CoE Excellerat P2**: Thermal simulations for the first wall of a tokamak fusion reactor, example in Figure 4. [1]

• **CoE Plasma PEPSC**: Plasma modeling using the BIT particle simulation code. [2]

• **EuroCC 2**: Supporting the Slovenian National Competence Centre for HPC. [3]

• **EuMaster4HPC**: Providing HPC training and education for students across Europe. [4]

• **DESY, Hamburg**: Innovative methods for imaging with the use of X-ray free-electron laser (XFEL) and synchrotron sources: simulation of sample delivery systems [5].

• **CABUM ERC project**, An investigation of the mechanisms at the interaction between cavitation bubbles and contaminants [6]

Conclusion

HPC has become an indispensable tool in modern engineering research. The HPCFS infrastructure provides computational support for demanding simulations, design optimization, and education. The planned 2025 upgrade will significantly enhance performance and user experience, ensuring continued support for scientific and industrial challenges.

Acknowledgements

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101101903. The JU receives support from the Digital Europe Programme and Germany, Bulgaria, Austria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Poland, Portugal, Romania, Slovenia, Spain, Sweden, France, Netherlands, Belgium, Luxembourg, Slovakia, Norway, Türkiye, Republic of North Macedonia, Iceland, Montenegro, Serbia.

The authors acknowledge the HPCFS team for technical support and maintenance of the computing infrastructure and ARIS agency for supporting the 2025 upgrade.

[1] CoE Excellerat P2, https://services.excellerat.eu/en/use-cases/engineering-design-and-digital-twin-of-the-first-wall-of-a-tokamak-fusion-reactor/ (86.10.2025).

[2] CoE Plasma PEPSC, https://plasma-pepsc.eu/bit1/ (26.10.2025).

[3] NCC Slovenia, https://www.sling.si/en/competence-centre/national-competence-centre/ (26.10.2025).

[4] EuMaster4HPC programme, https://eumaster4hpc.eu/ (26.10.2025).

[5] B. Zupan, R. Zahoor, S. Bajt, B. Šarler, Numerical treatment of electrical properties in two-phase electrohydrodynamic systems. Physics of Fluids, vol. 37, 2025353.

[6] J. Zevnik and M. Dular, Cavitation bubble interaction with compliant structures on a microscale: a contribution to the understanding of bacterial cell lysis by cavitation treatment. Ultrasonics sonochemistry, Jun. 2022, vol. 87, pp. 1-20.