

PhD Course Syllabus

High-Performance Computing Environments for Scientific and Engineering Applications

Course Description and Objectives

This course provides a comprehensive introduction to high-performance computing (HPC) environments, focusing on the Bash shell environment and the Slurm workload manager. The primary objective is to equip doctoral students with the practical skills necessary to effectively utilize HPC clusters for scientific and engineering computations. Upon completion, students will be able to navigate and script within the Bash environment, manage computational resources via Slurm, access and install software on a cluster, and execute parallel computations.

Prerequisites

- Basic familiarity with any command-line interface is beneficial.
- No prior programming experience is required, but a willingness to learn scripting is essential.
- **Important:** All course activities will be conducted entirely on the NCBJ computing cluster. Therefore, participants must obtain access to this cluster prior to the start of the course. Information regarding the registration for computational grants on the cluster will be provided several days before the course begins.

Course Contents

Note: All practical work will be conducted on a Linux-based HPC cluster.

1. The System Shell (Bash Environment)
 - Introduction to the Bash Shell:
 - Core concepts and terminology.
 - Command structure and syntax.
 - Getting help: inline help and the Linux manual system (man pages).
 - The File System:
 - Fundamentals: concepts, paths, file attributes, and POSIX permissions.
 - Basic file operations.
 - Wildcards and Globbing.
 - Redirections and Data Streams:
 - Capturing command input/output.
 - Connecting commands with pipes.
 - Environment Variables:

- Variable syntax.
- Controlling the shell and executed commands.
- Shell Scripts:
 - Script structure and the permission system.
 - Using execution scripts with the Slurm system.
- 2. The Computing Cluster and Job Scheduling System (Slurm)
 - Introduction:
 - Cluster access and login procedures.
 - Technical basics/terminology: computational resources, allocations, job types, limits, Quality of Service (QOS).
 - Running and Managing Computational Jobs:
 - Essential Slurm commands and how to get help.
 - GPGPU Jobs.
- 3. Software Management
 - Software available on the NCBJ cluster.
 - Accessing software (using Environment Modules / Spack).
 - Installing custom software.
 - Compiling from source code.
 - Parallel computations with Slurm (MPI).

Hours and Schedule

- Total Hours: 15
- ECTS Points: 1.5
- Schedule: The course consists of 5 sessions, each lasting 3 hours.
- Sessions will blend theoretical concepts with hands-on practical exercises on an HPC cluster.
- Due to the limited number of places, students from outside the Doctoral School will be admitted on a first-come, first-served basis. Students from outside the Doctoral School should contact grad@ncbj.gov.pl regarding enrollment.

Grading Policy

The course is conducted entirely as practical workshops and exercises. Students are expected to actively participate in all sessions and complete practical exercises during the classes to demonstrate acquired competencies in using the Bash shell, writing scripts, submitting and managing jobs via Slurm, and managing software environments. Successful completion of the course requires attendance and active participation in all scheduled sessions.