

## Dynamic Resource Management for HPC

Format: Half day

### Abstract:

Traditional batch scheduling models, which rely on static resource allocation, are ill-suited for the growing scale, heterogeneity, and workload variability of current and emerging HPC systems. Currently, applications must typically acquire their maximum resource footprint upfront and retain it for their entire runtime, leading to inefficient utilization, elongated queue times, and reduced system throughput. To address these limitations, dynamic resource management (DRM) and dynamic applications (malleable and elastic) have emerged as key strategies to enable malleability and improve responsiveness in modern supercomputing environments.

This tutorial introduces the foundations, frameworks, and practical methodologies for DRM in HPC, with a particular focus on dynamic MPI applications and their interaction with resource managers and schedulers. Participants will explore the motivations and system-level implications of DRM, learning how dynamic resizing of parallel jobs can improve both user experience and overall system efficiency. We will present and compare two representative tools: Dynamic Process with PSets (DPP), and the Dynamic Management of Resources (DMR) framework (pronounced ‘dimmer’), which enable runtime resource adaptation through different mechanisms, programming abstractions, and integration strategies.

The tutorial will be divided into conceptual and practical blocks. We begin by examining the core concepts of MPI malleability and scheduler interaction models, covering both fundamental principles and recent research advances. Next, we present each representative approach for dynamic resource management in detail, outlining its architecture, programming interfaces, supported resource managers, and specific use cases. Participants will gain insights into how these frameworks address challenges such as process migration, data redistribution, reconfiguration latency, and coordination with job schedulers.

A significant portion of the tutorial will be dedicated to hands-on exercises, in which participants will execute applications with dynamic resource utilization on provided virtual or cluster environments. Through guided examples, they will learn how to launch, monitor, and dynamically resize applications observing in real time how resources are expanded or released in response to external events. These exercises will help bridge the gap between theoretical understanding and practical deployment.

By the end of the tutorial, attendees will be able to (1) understand the benefits and limitations of current DRM approaches, (2) evaluate which framework best fits their workloads and environments, and (3) start adapting their own applications to exploit runtime MPI malleability. This tutorial is intended for HPC practitioners, scientific application developers, system administrators, and researchers interested in improving system productivity, increasing system utilization, and performance and energy efficiency through DRM. No previous experience with DRM frameworks is required, although a working knowledge of MPI and job schedulers is recommended.