



# Automated Calibration of a Simulator of MPI Application Executions

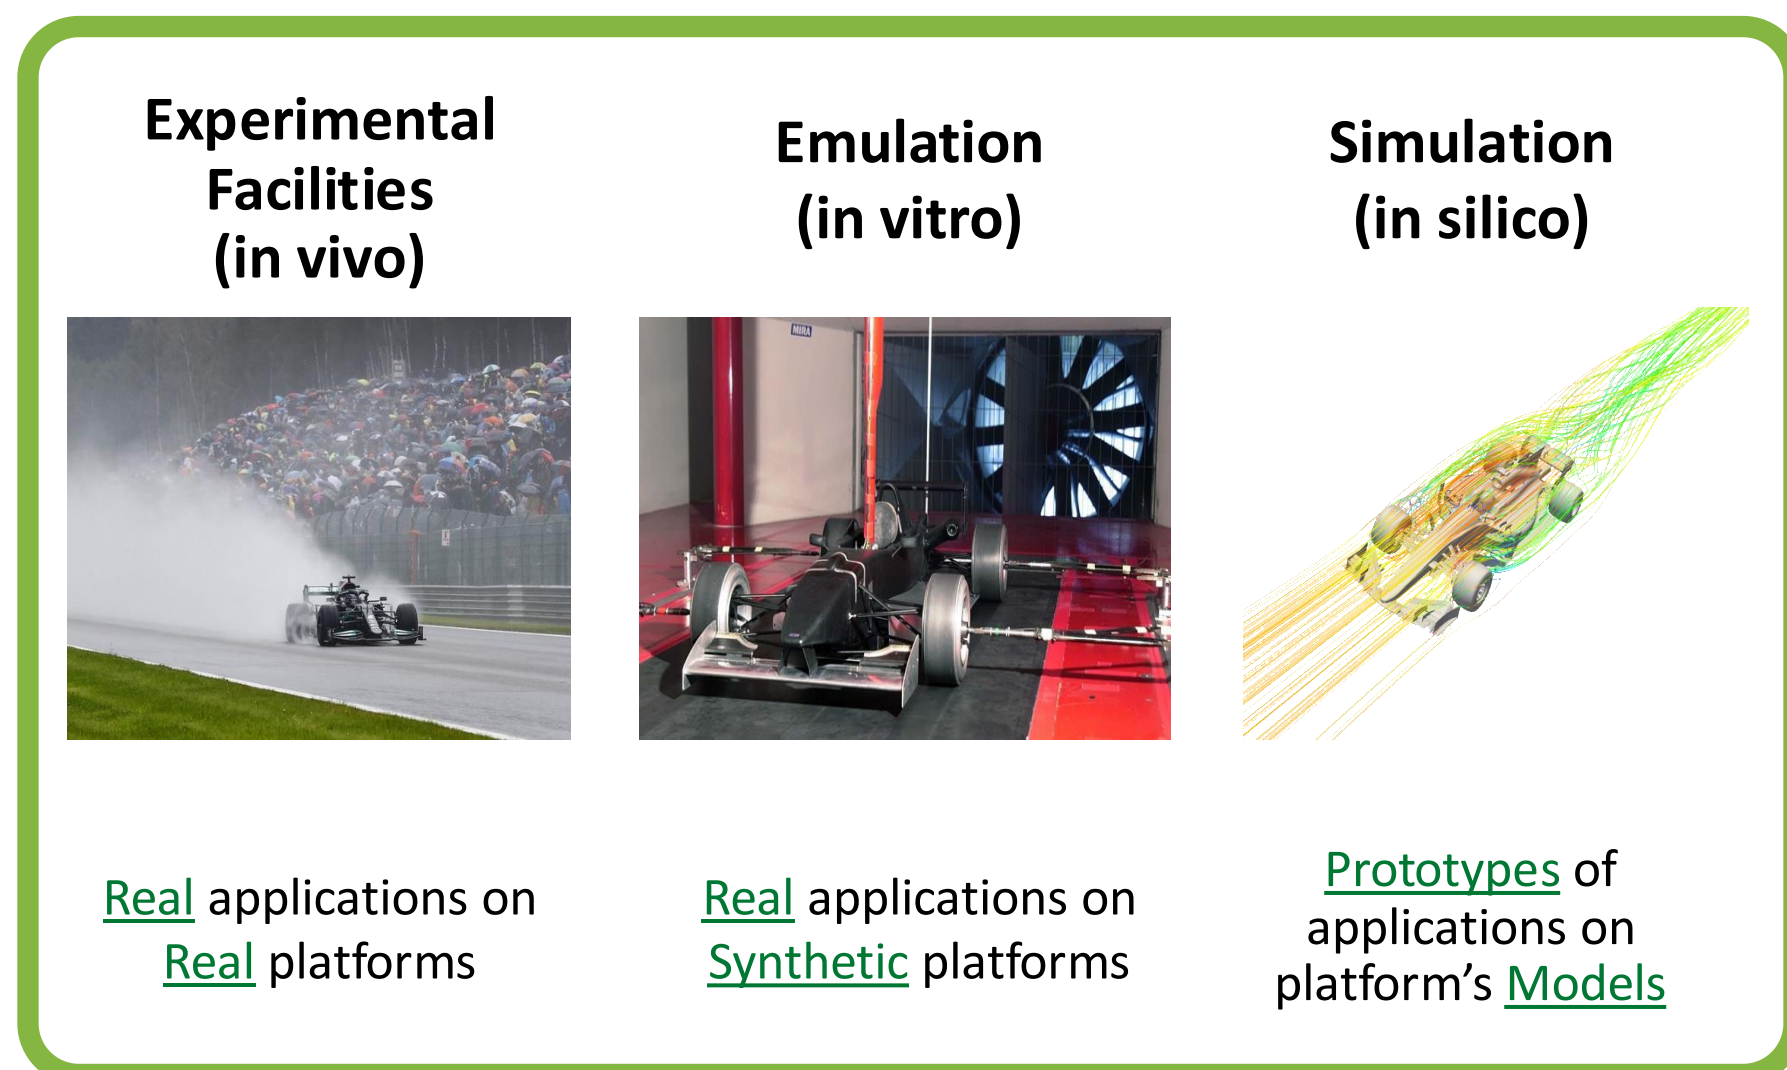
Yick Ching Wong<sup>1</sup>, Jesse McDonald<sup>1</sup>, Fred Suter<sup>2</sup>, Kshitij Mehta<sup>2</sup>, Henri Casanova<sup>1</sup>

<sup>1</sup>University of Hawai'i at Mānoa - Information and Computer Sciences, Honolulu HI, USA

<sup>2</sup>Oak Ridge National Laboratory - Computer Science and Mathematics Division, Oak Ridge TN, USA

## Why Simulation-based Performance Assessment?

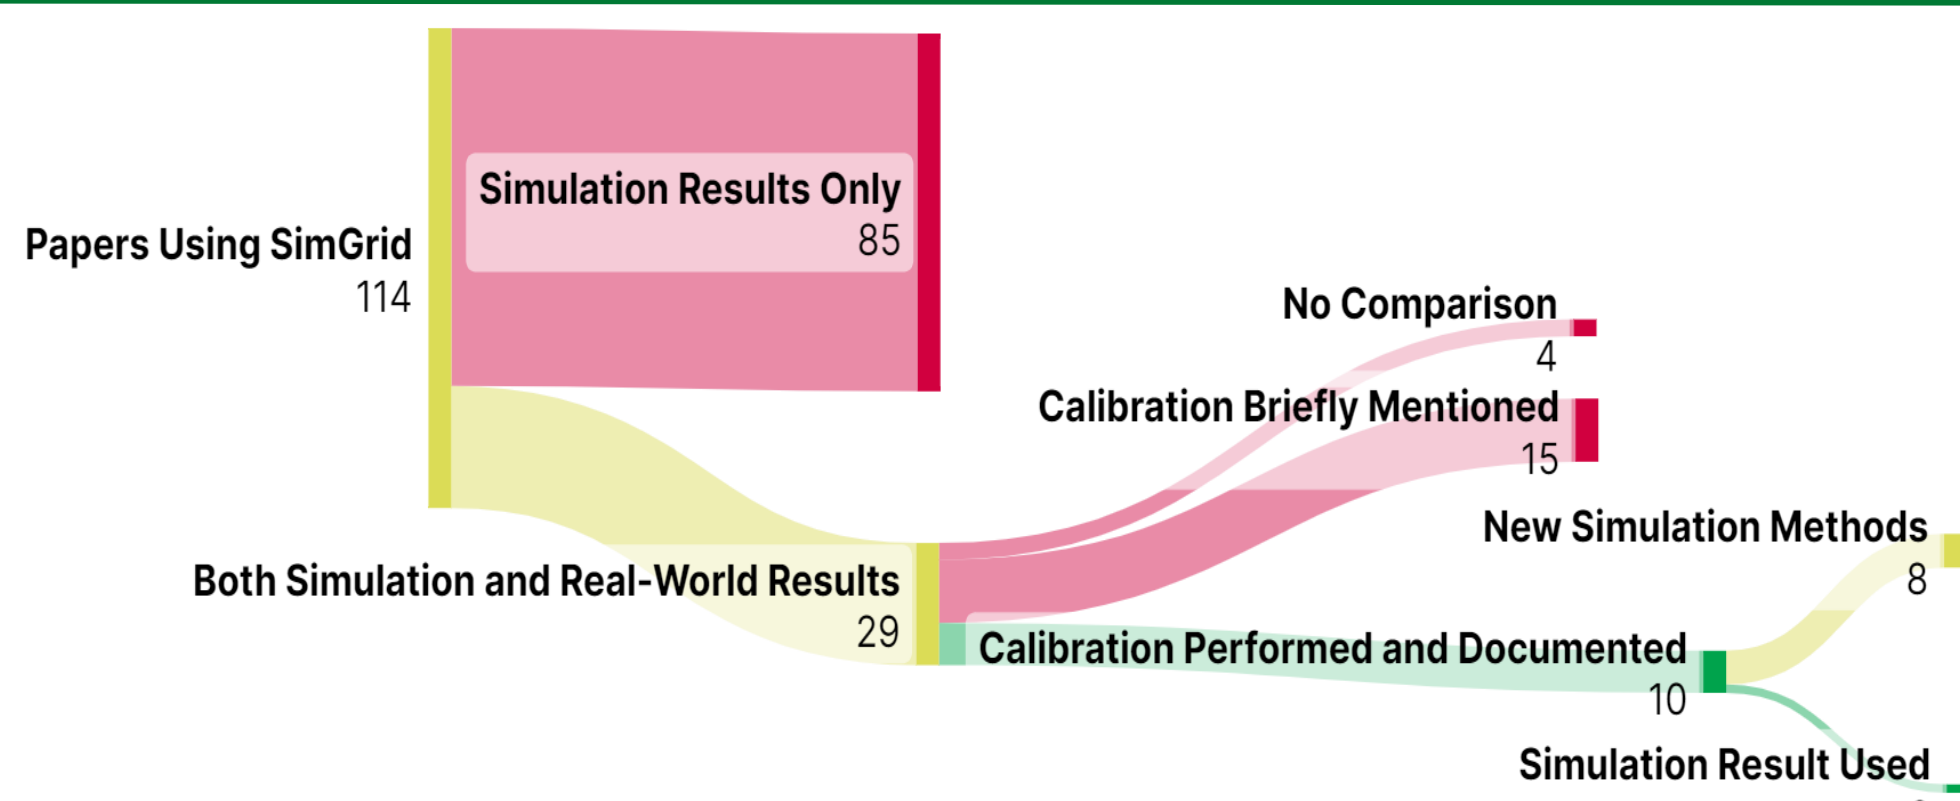
- Theoretical studies are not enough to account for the unprecedented complexity of modern systems
- Simulations are repeatable and fully observable
- Simulations enable experiments with arbitrary platform/application configurations
- Often less time-, labor-, and cost-intensive than real-world experiments



## Do we need to calibrate our simulators?

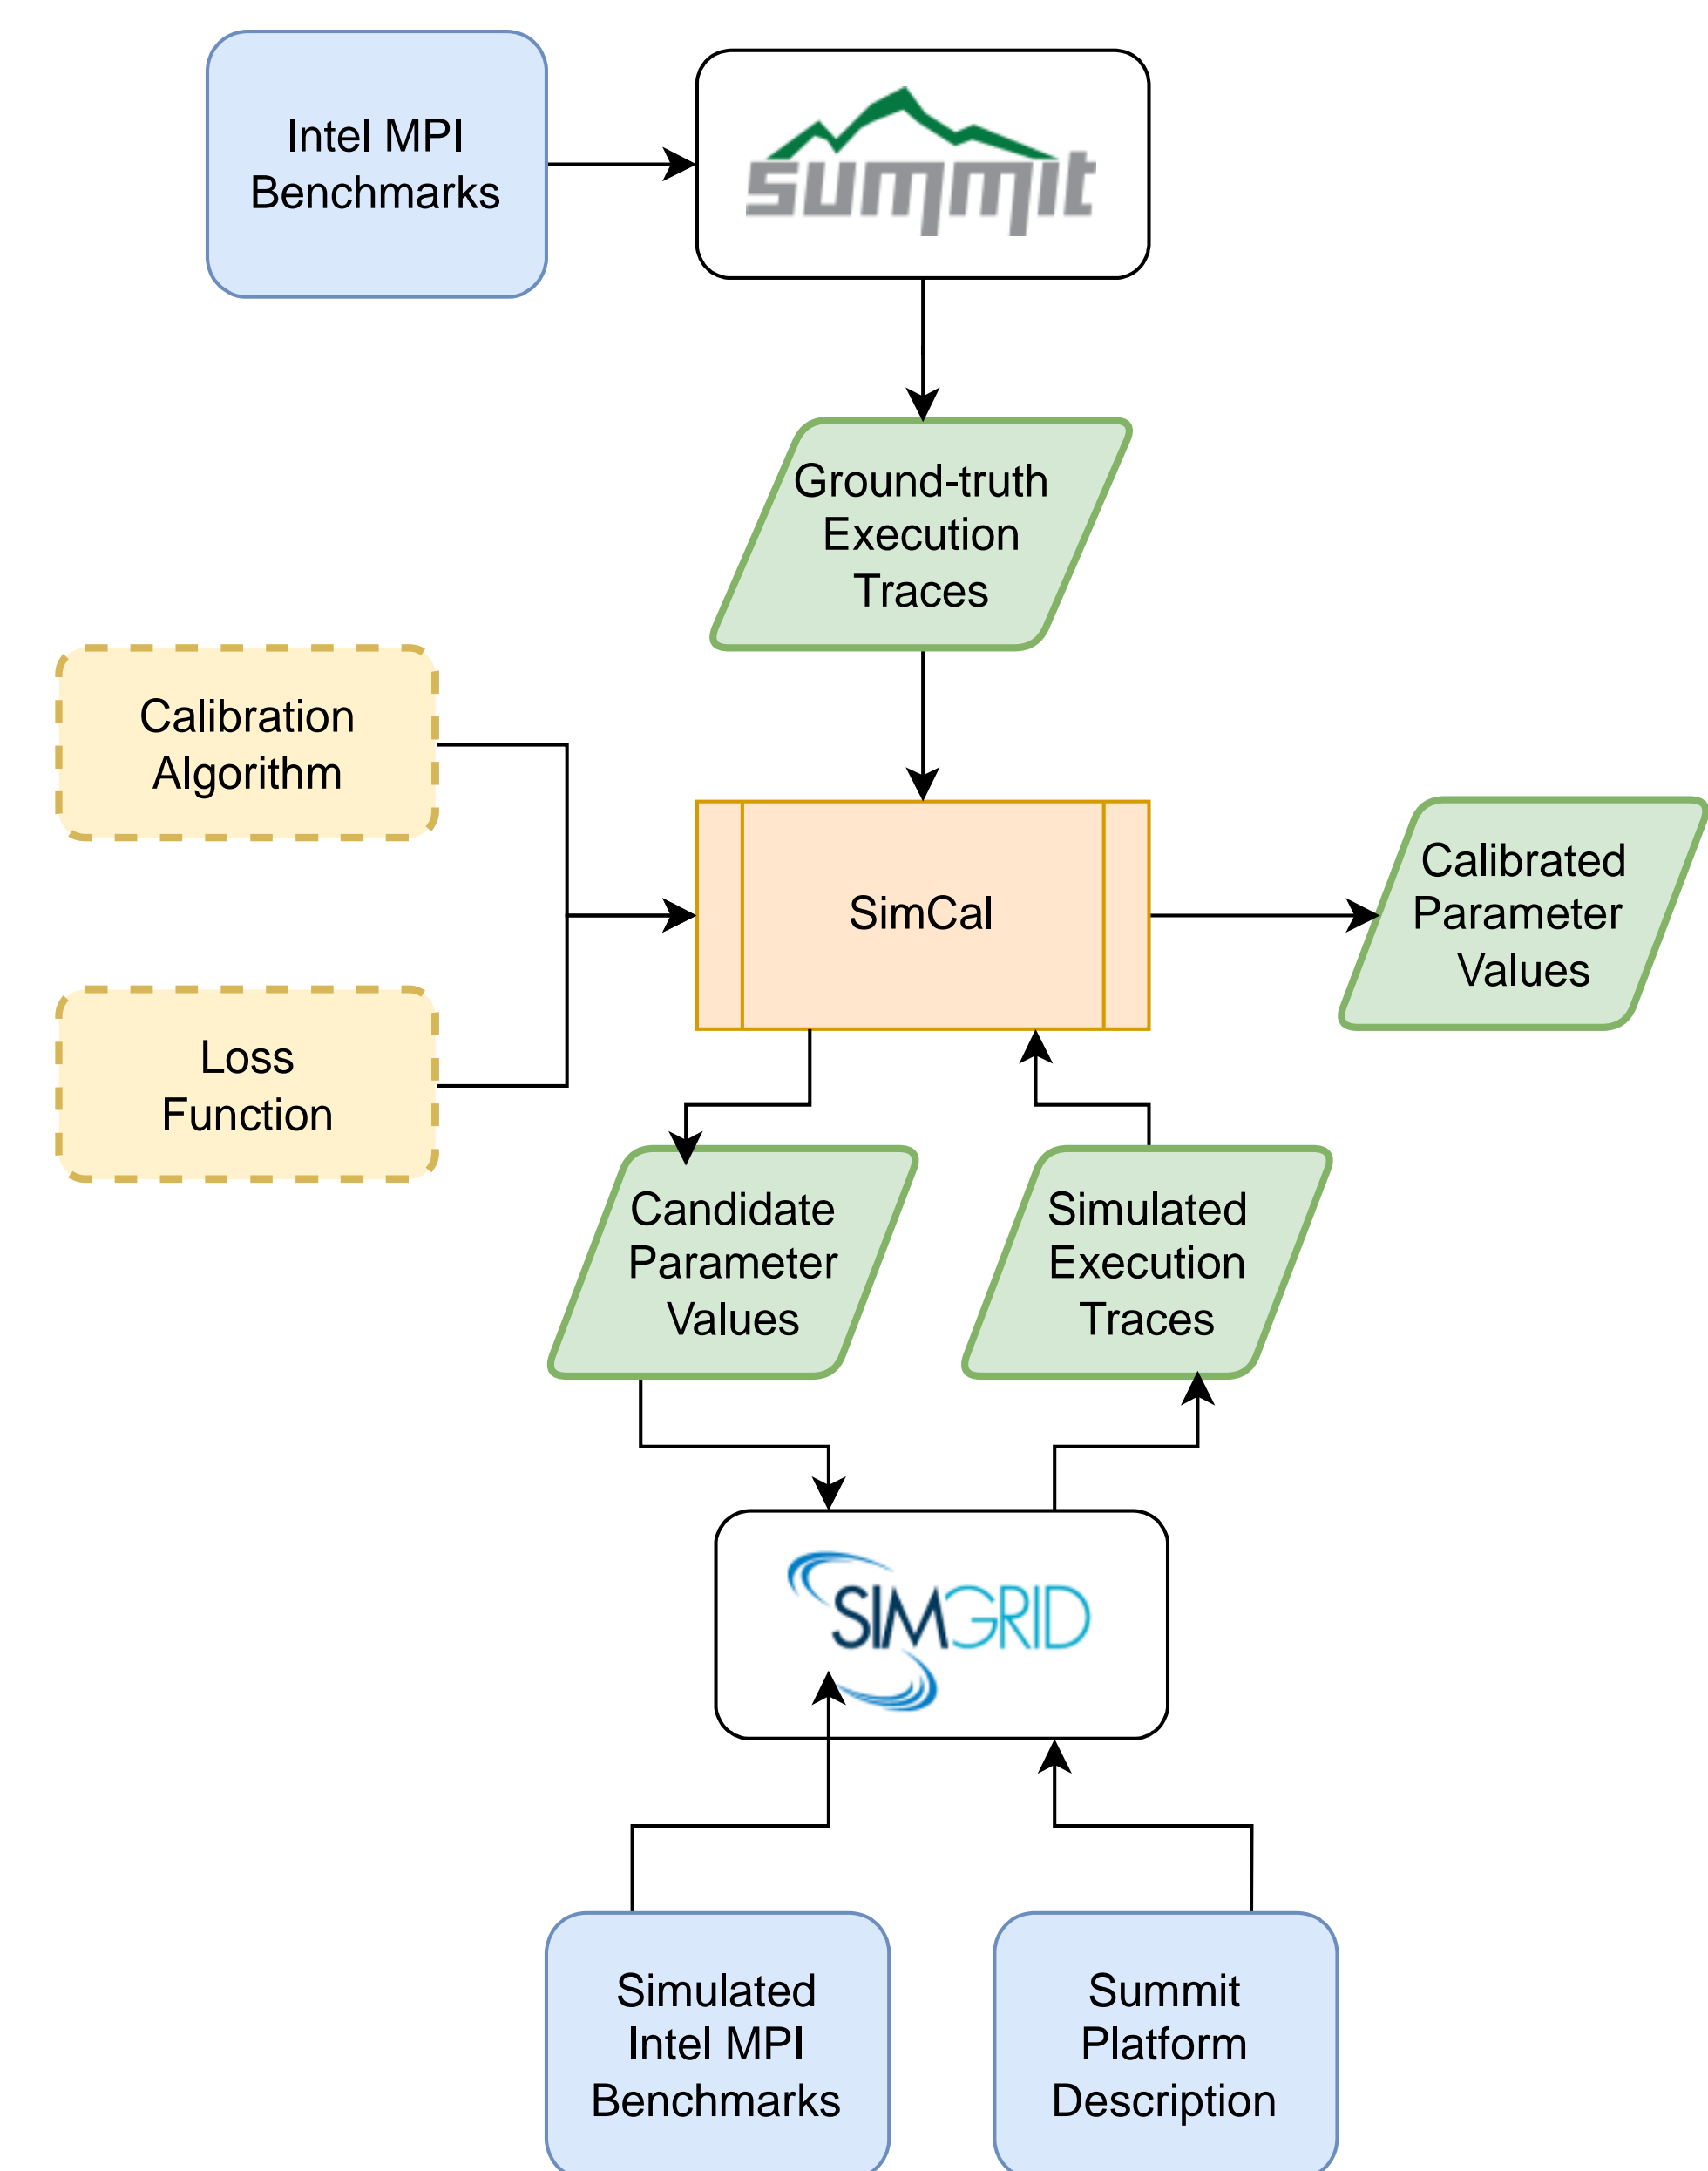
- We want simulators to reflect real-world results but to abstract away real-world behaviors
- To allow versatility, simulation models are configured via parameters
- **Challenge:** Parameters values cannot always be picked just based on hardware specs
- **Approach:** Calibrate simulation parameters with respect to ground-truth real-world execution traces

## Are people calibrating their simulators?



**NOT REALLY!** In a study of all 114 SimGrid-based papers published over a 5-year period, only **two** papers that do not focus on simulation as the research result perform and document a sound calibration procedure

## Methodology

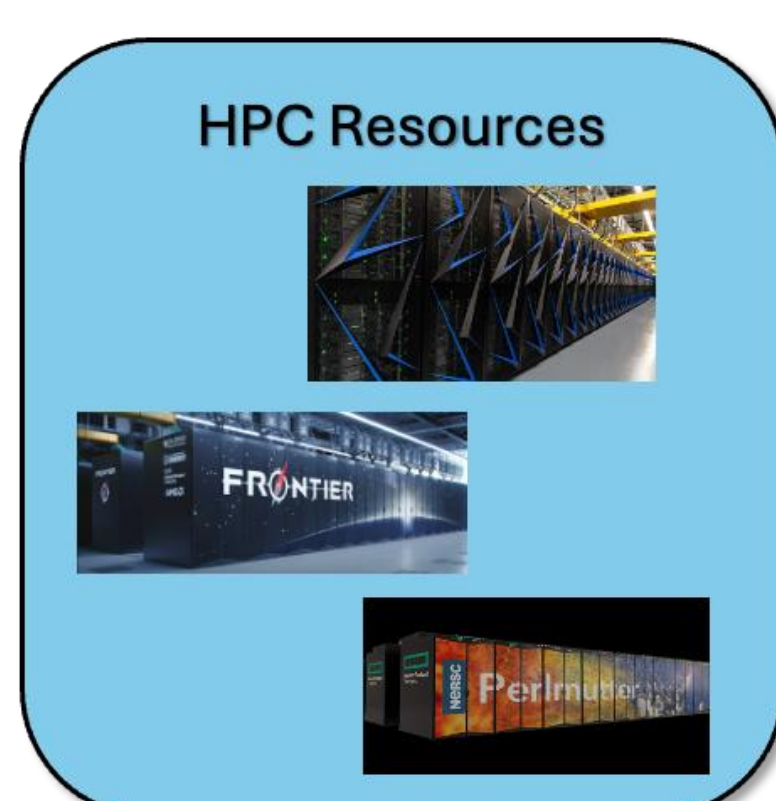
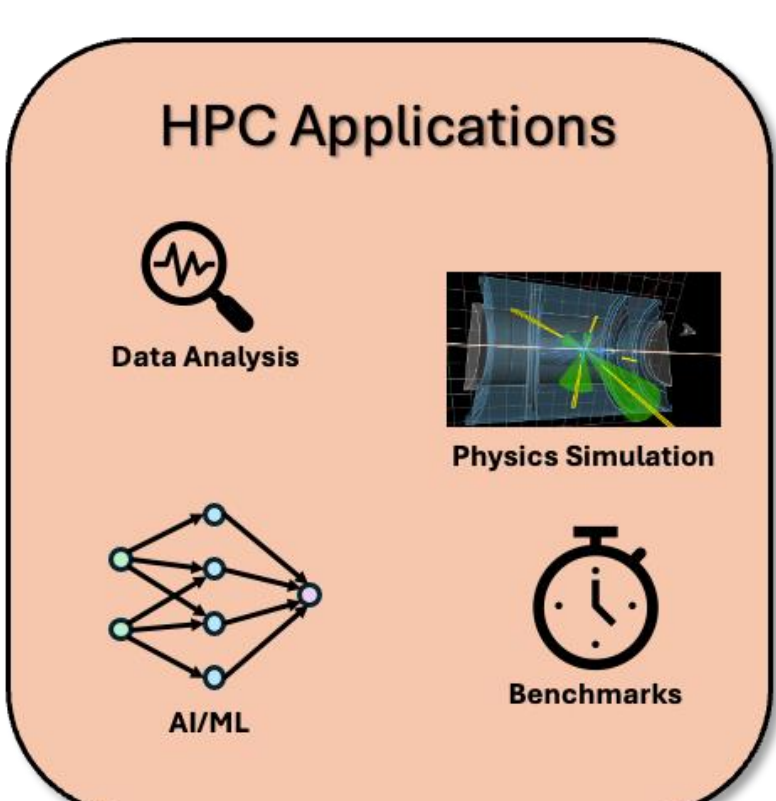


**Investigate whether simple algorithms can be used to enable automatic calibration of simulators to provide domain scientists with more accurate simulators for performance assessment**

## Why MPI Benchmarks?

- Many scientific applications relies on MPI communication
- Benchmarks can capture the impact of the platform on network communications, without the complexity of scientific applications

## Components of a Simulation



## Preliminary Results

