

Exploration of Simulation Calibration on Intel MPI benchmarks using SimCal

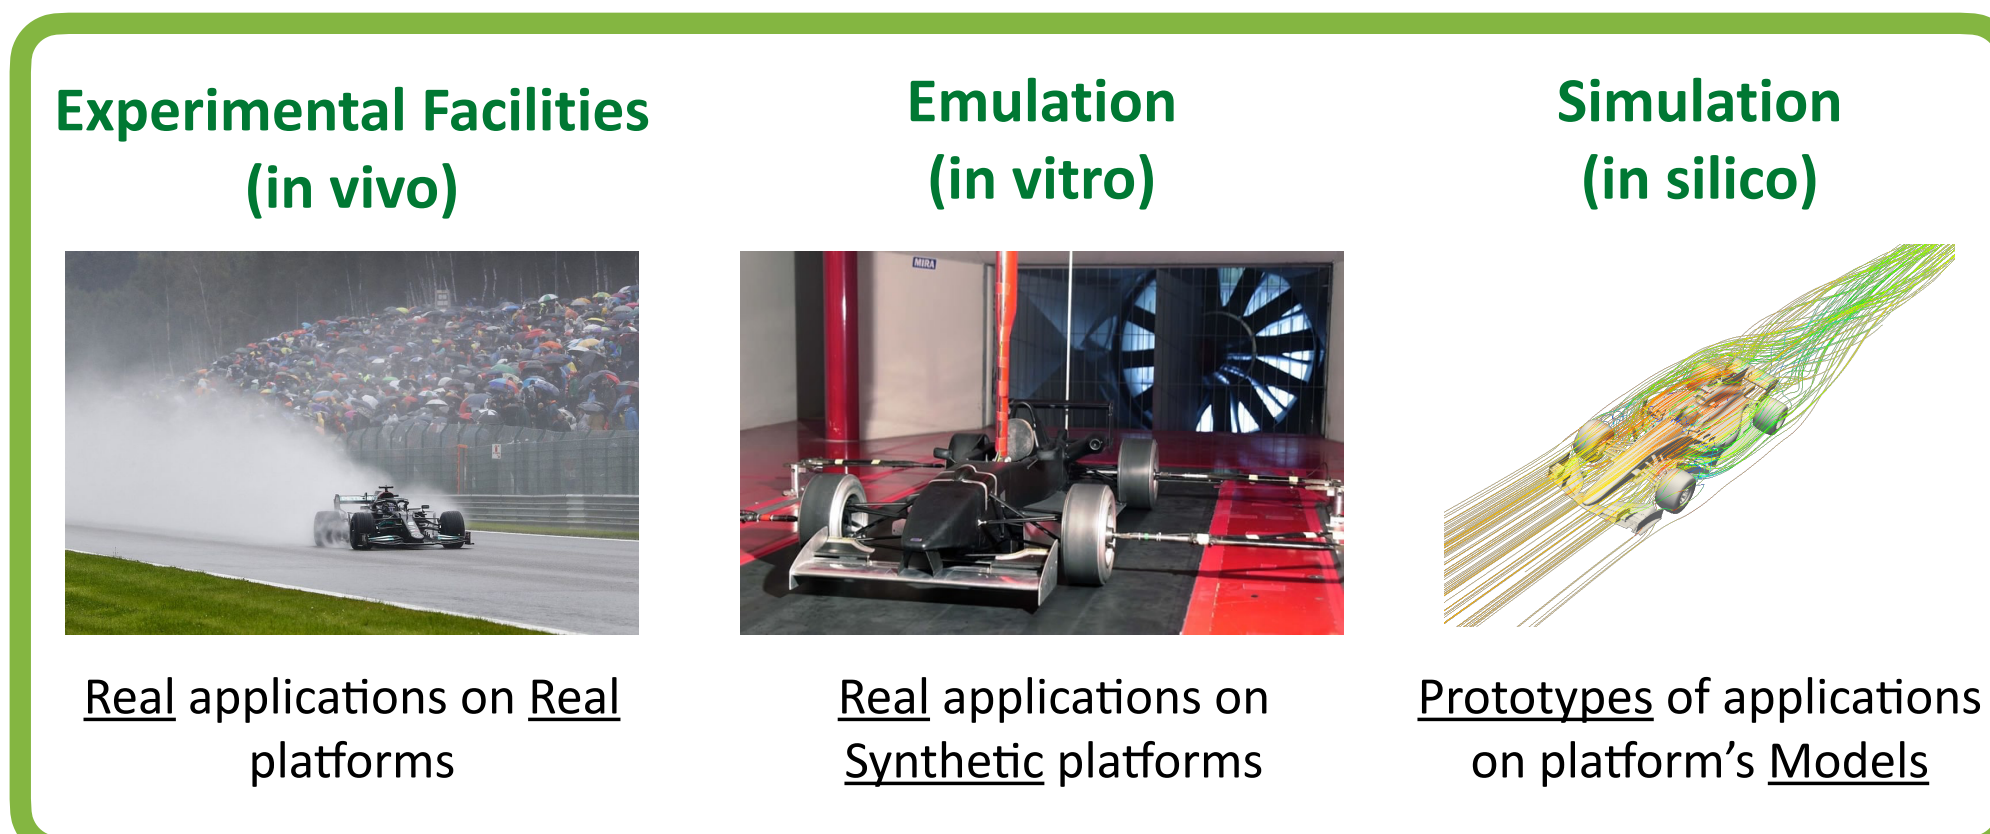
Yick Ching Wong¹, Fred Suter², Kshitij Mehta², Henri Casanova¹, Jesse McDonald¹

¹University of Hawai'i at Mānoa, Honolulu, HI, USA

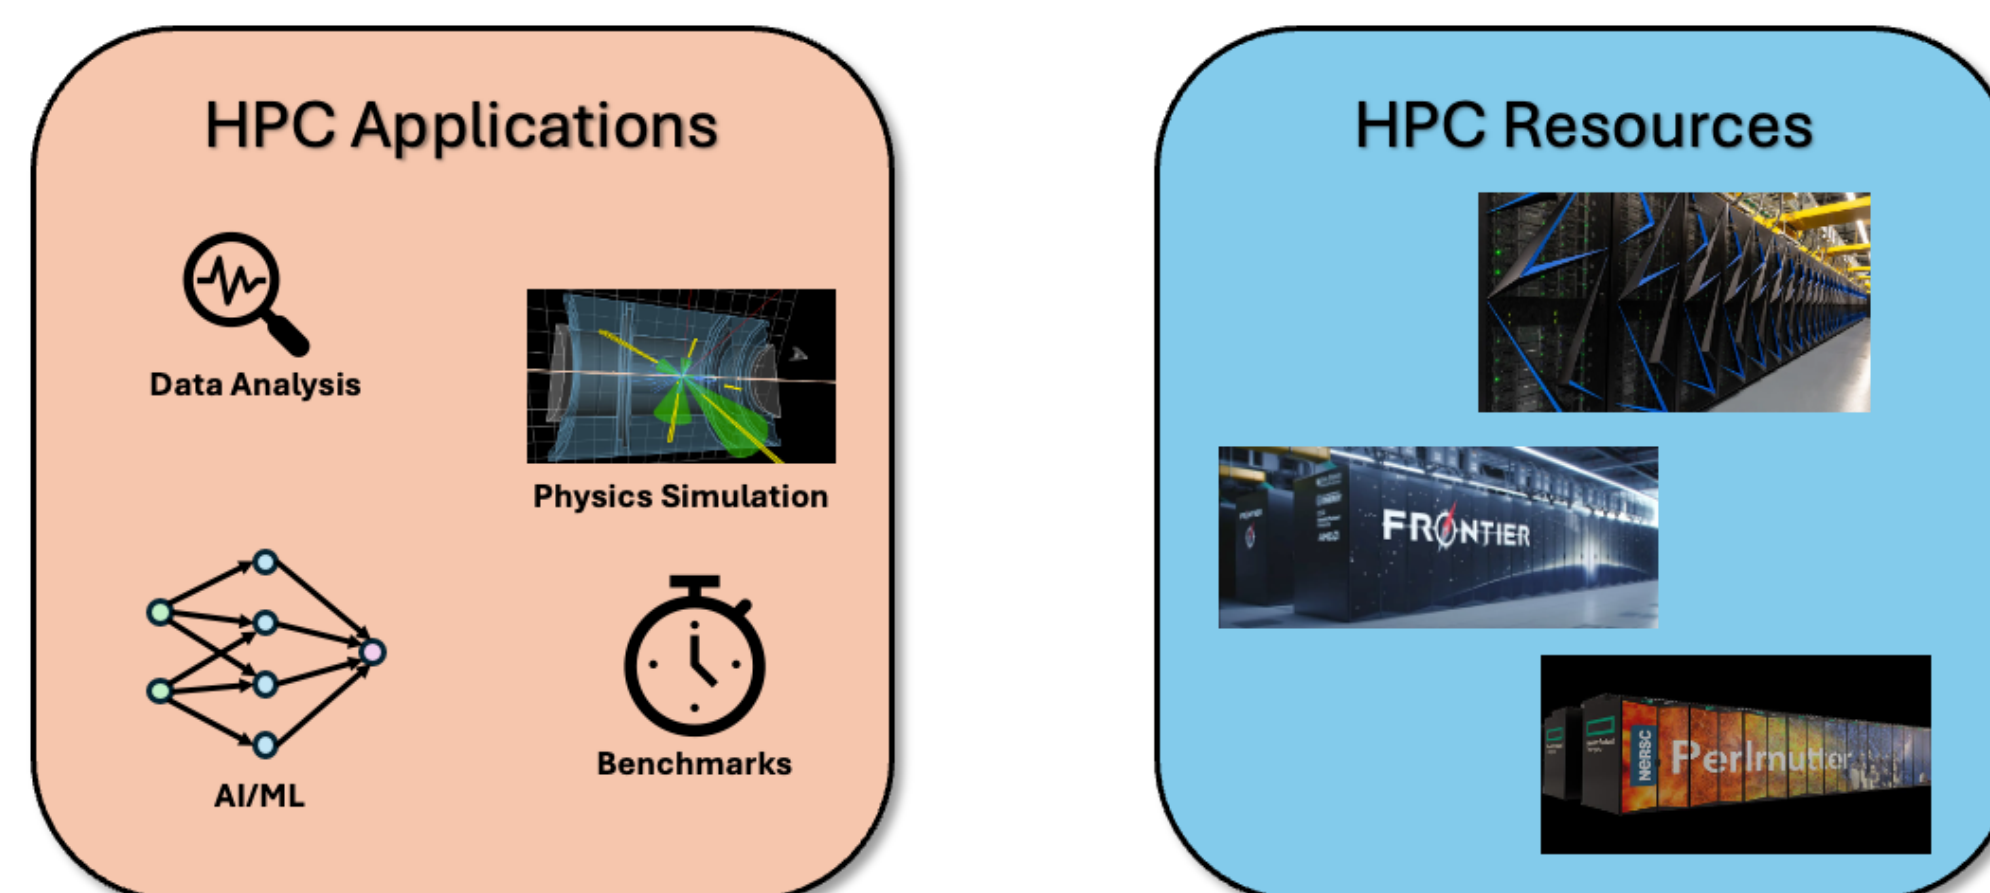
²Oak Ridge National Laboratory - Computer Science and Mathematics Division, Oak Ridge, TN, USA

Why Simulation-based Performance Assessments?

- 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



Components of an HPC Simulation



Why MPI Benchmarks?

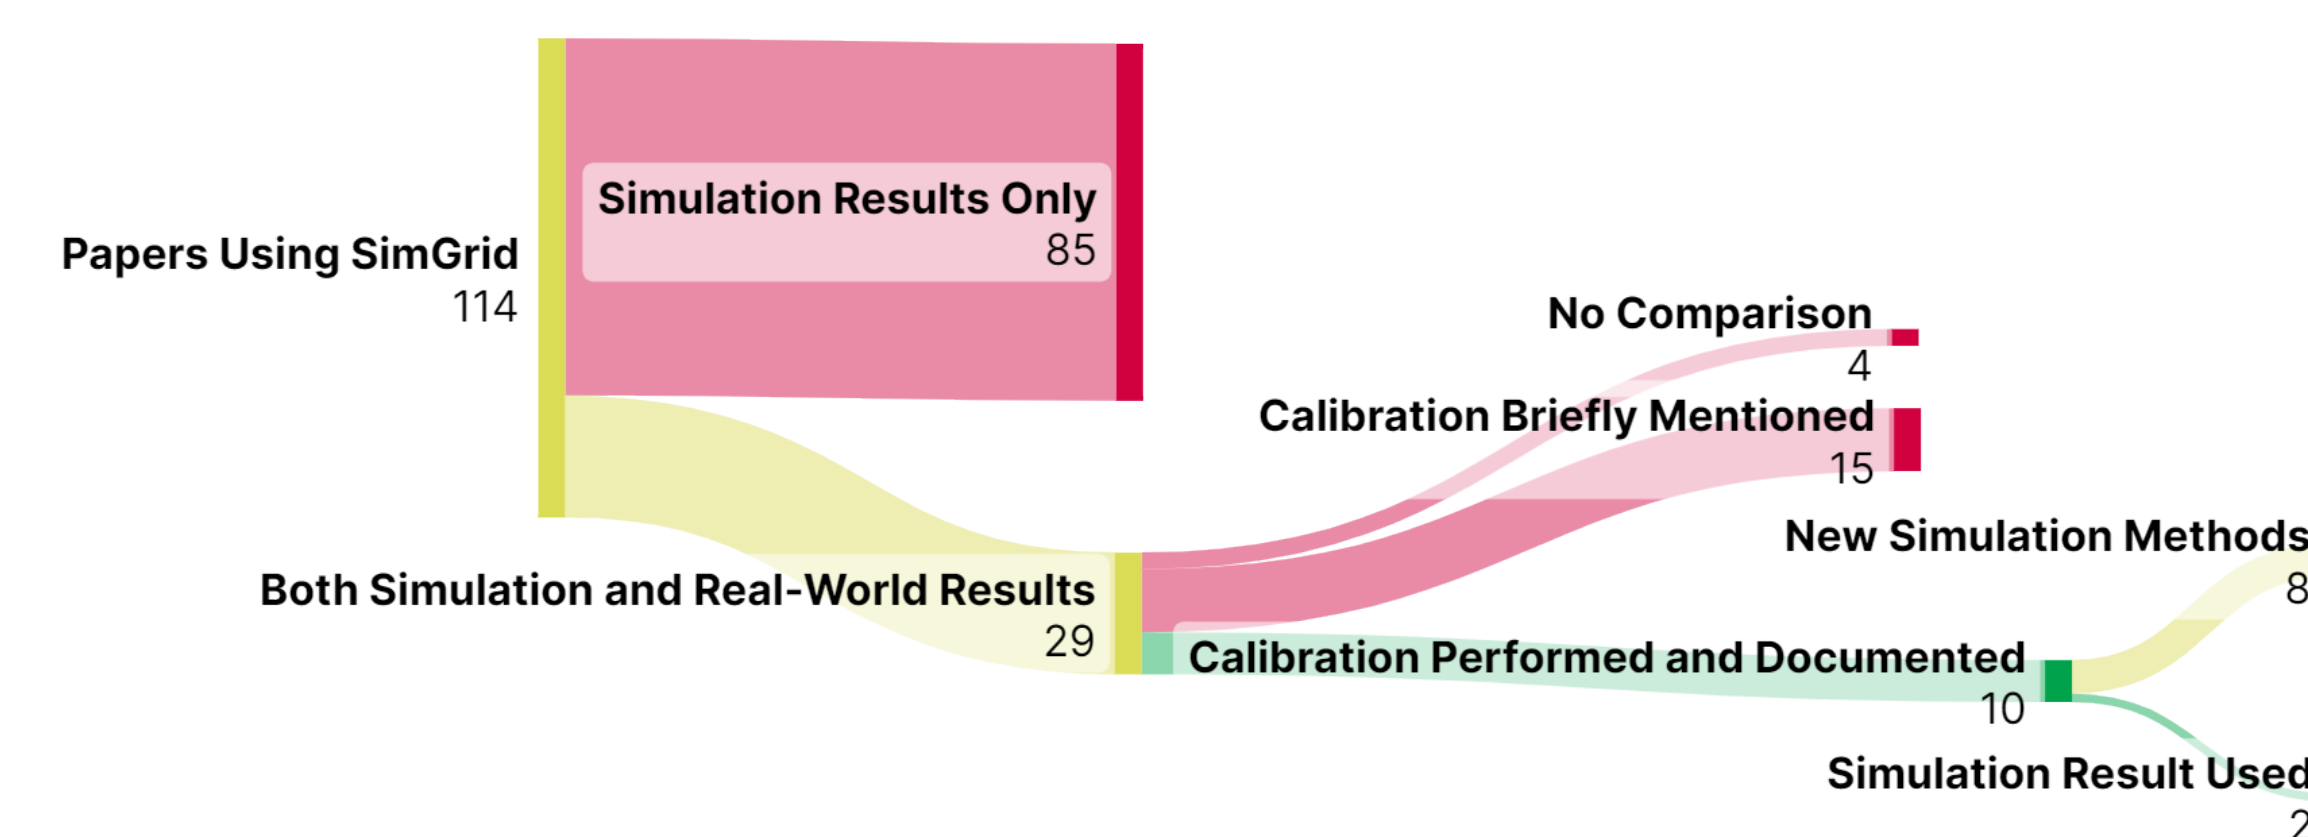
- Many scientific applications relies on MPI for inter-process communication
- Benchmarks can capture the impact of the platform on network communications without the complexity of scientific applications

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

Toward accurate performance assessment of MPI application executions on HPC platforms using a simulator with automatically calibrated simulation parameters

Are people calibrating their simulators?

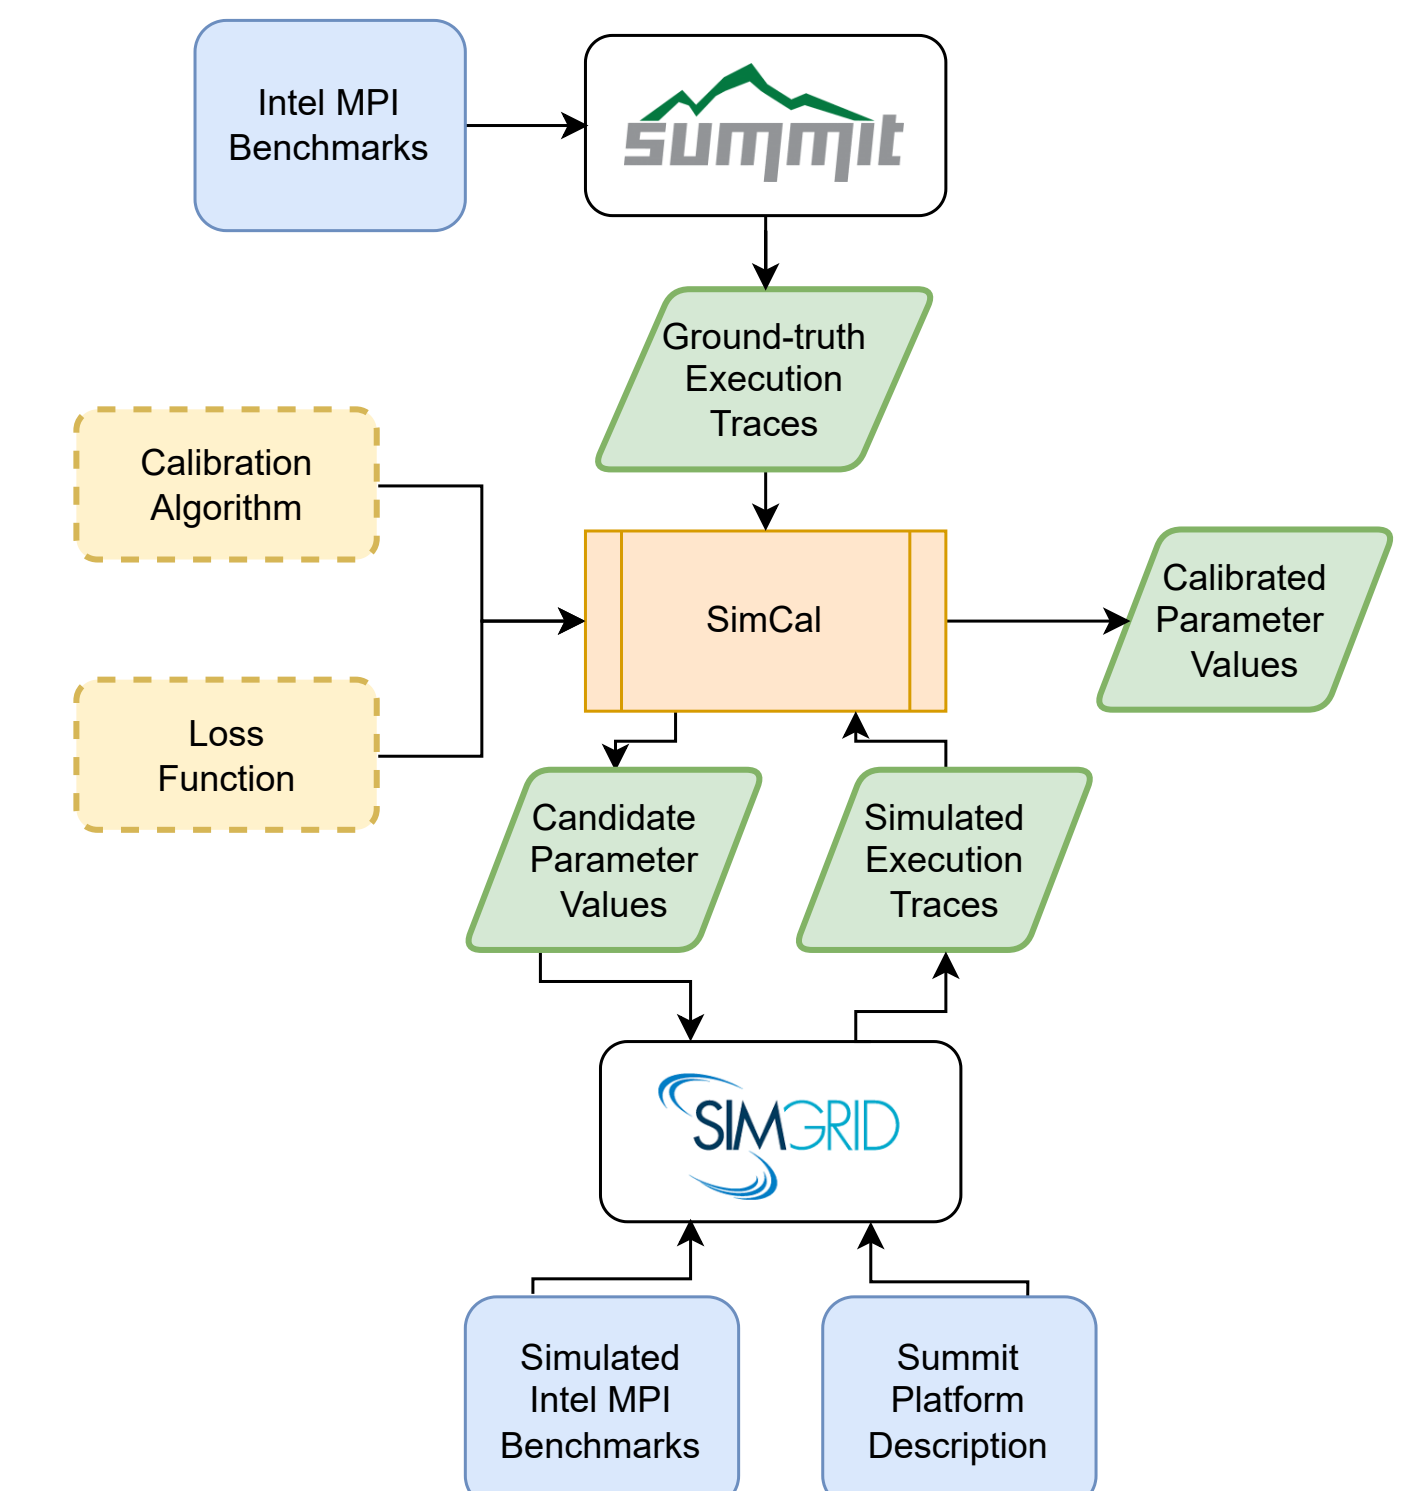


NOT REALLY! In a study of all 114 SimGrid-based papers published over a 5-year period, only **two** papers that do not focus on simulation as the research result perform and document a sound calibration procedure



McDonald, J., Horzela, M., Suter, F., & Casanova, H. (2024). Automated Calibration of Parallel and Distributed Computing Simulators: A Case Study. *arXiv preprint arXiv:2403.13918*.

Methodology



SimCal

A simulation calibration framework written in Python. It includes basic algorithms for calibration such as GridSearch, RandomSearch, and Gradient Descent

SimGrid

A simulation framework for developing simulators of distributed application (including MPI-based applications) executions on distributed platforms

Intel MPI Benchmarks

A series of MPI benchmarks written by Intel. Running via SimGrid's **SMPI** API in the simulator. We will utilize the P2P (Point-to-Point) benchmarks for this calibration study

Summit

Multiple instances of the Intel MPI benchmark were ran on Summit, which provides the ground-truth data we used for this study

By leveraging SimGrid's **programmatic resource description** API, we were able to create a simulated platform of Summit in less than 80 lines of code

We now have all the pieces in place to calibrate for accurate simulations of MPI applications on production HPC platforms