

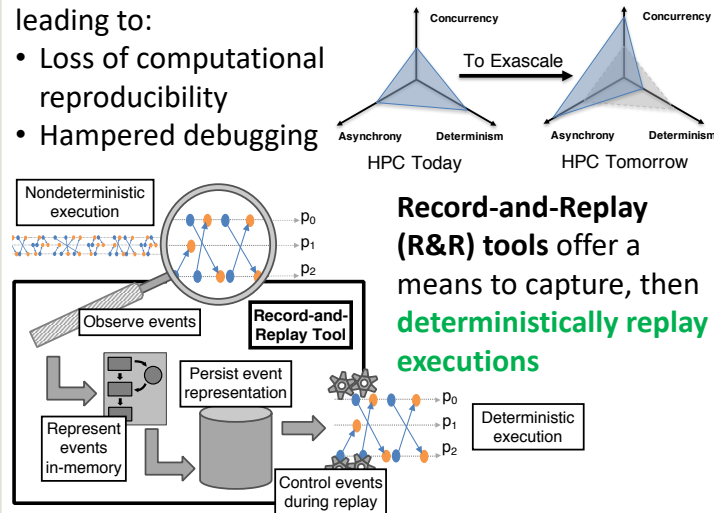
Modeling Nondeterminism for Reproducibility of HPC Applications

Dylan Chapp, Danny Rorabaugh, Michela Taufer

Project Overview

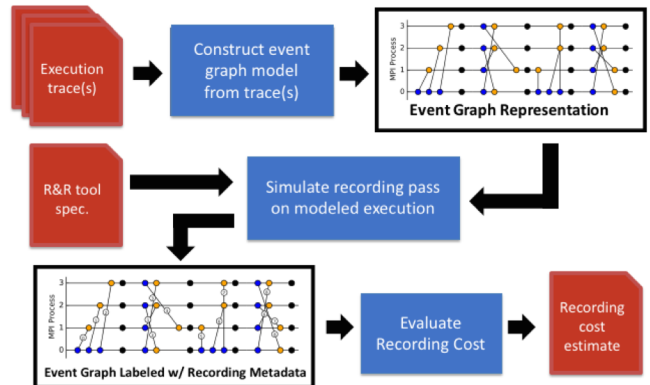
HPC applications at exascale must contend with an **unprecedented degree of nondeterminism**, leading to:

- Loss of computational reproducibility
- Hampered debugging



Modeling Methodology

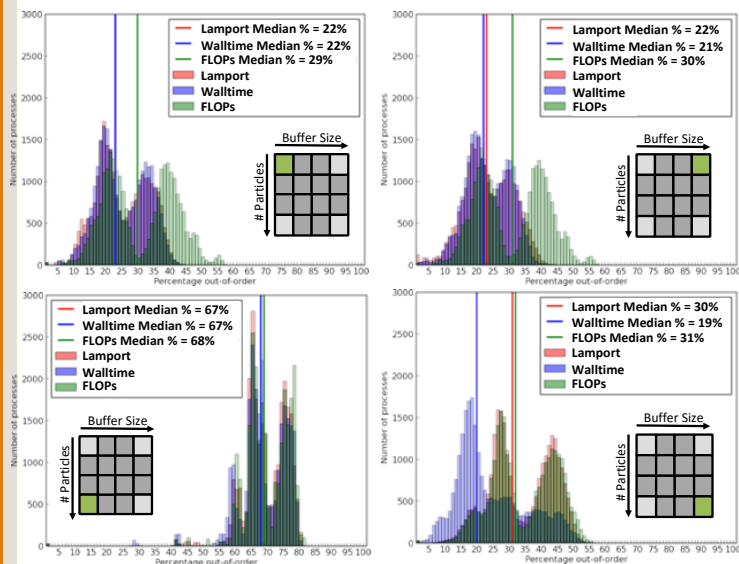
We model parallel executions as directed acyclic event graphs and simulate the recording phase of R&R tools on the modeled executions.




Our modeling methodology enables study of the links between message-passing patterns and recording cost

Motivating Case Study

We study the cost of recording the **Monte Carlo Benchmark (MCB)** using the R&R tool **ReMPI** under varied application and tool configurations



Initial Results and Future Directions



We extract the communication patterns that contribute most to recording cost from event graph models

Publication

D. Chapp, D. Rorabaugh, and M. Taufer: **Modeling Record-and-Replay for Nondeterministic Applications on Exascale Systems**. Poster at Workshop on Modeling and Simulation of Systems and Applications (ModSim), 2018.

