Application Modeling

Jered Dominguez-Trujillo

Introduction

- Performance modeling and application prediction
  - Run-to-run variation of identical jobs difficult to quantify
- Data required to model application behavior prohibitively large
- Seek to develop modeling method:
  - Minimize data requirements
  - Provide insight into complex application behavior

Contributions

- New modeling approaches for application performance and variation
  - A single non-parametric approach for simple workloads
  - Two parametric approaches for complex workloads
- MPI implementation of the required measurement technique
- Evaluation of these techniques
- Full results presented at the PMBS Workshop at SC20[1]

[1] J. Dominguez-Trujillo *et al.*, "Lightweight Measurement and Analysis of HPC Performance Variability," *2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, GA, USA, 2020, pp. 50-60, doi: 10.1109/PMBS51919.2020.00011.

Modeling Approach: Non-Parametric

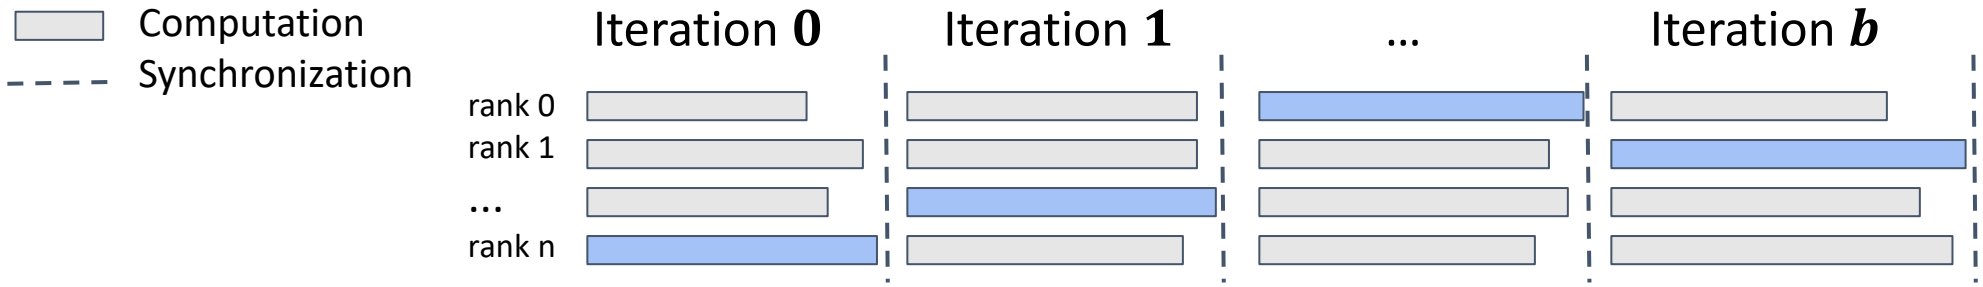

- Goal: Predict performance when scaled by a factor of k
- Gather maxima timing data from each of b iterations
  - Each data point is the maximum of n ranks
- Randomly resample k data points with replacement and take the maximum
  - Perform this b times to generate predicted distribution of maximums at new scale

Modeling Approach: Non-Parametric

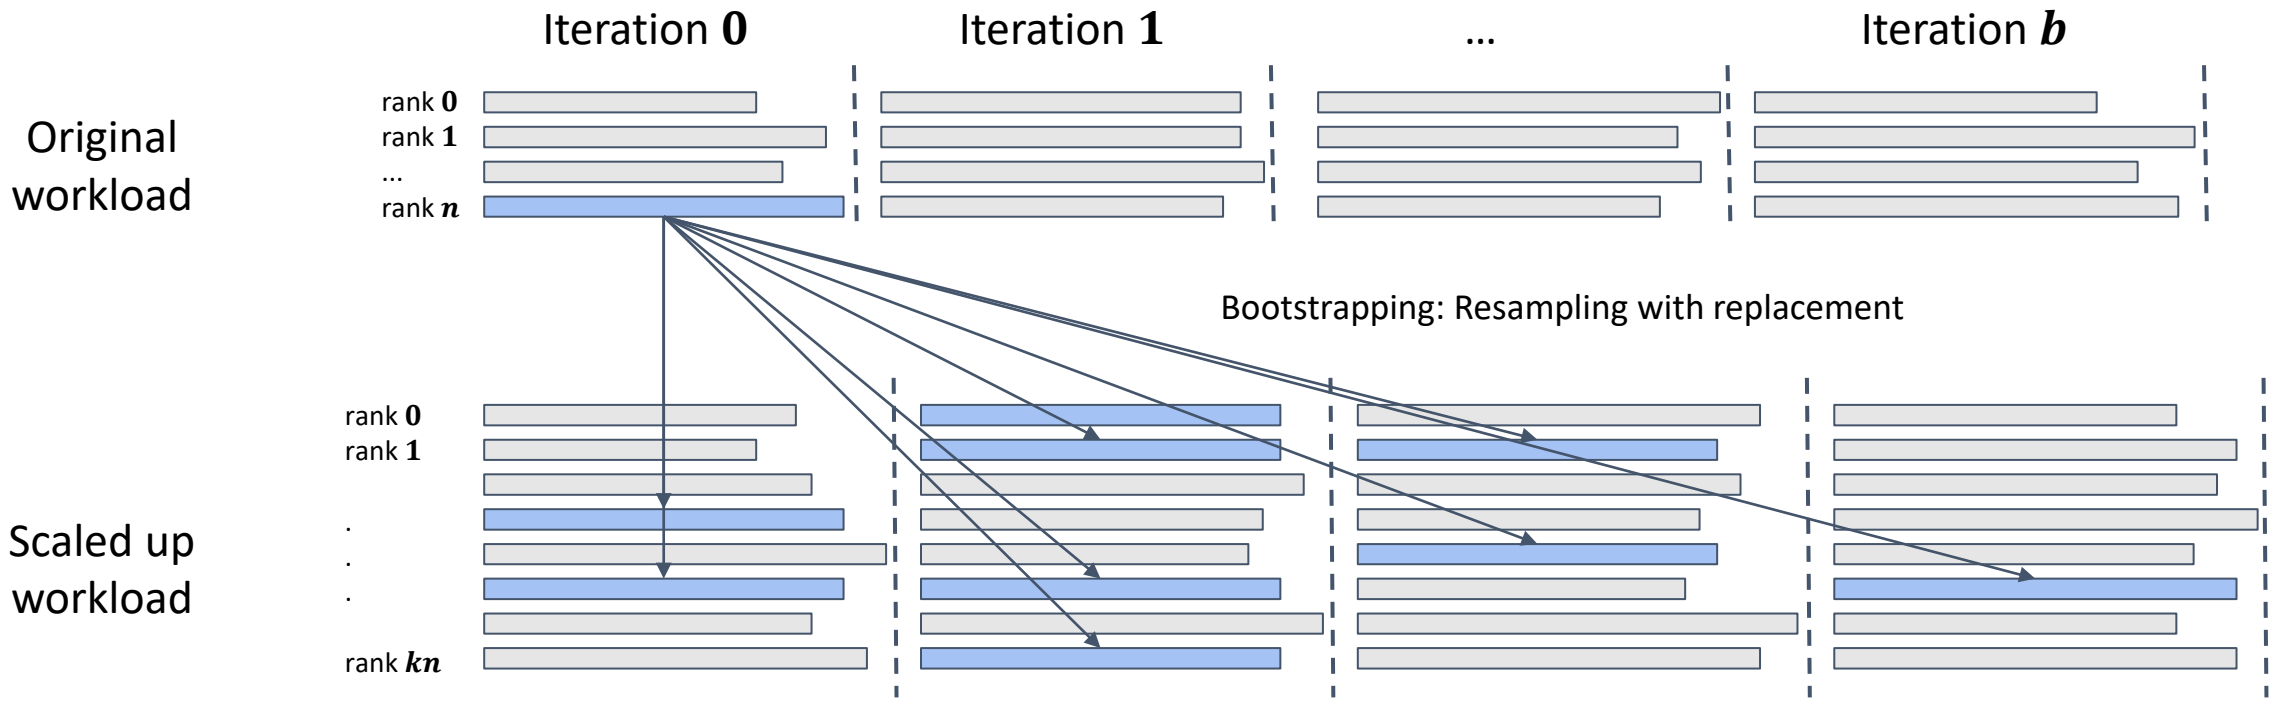

Modeling Approach: Parametric

- Gather maxima timing data from subsets of ranks
- Fit GEV distribution to samples
  - Outputs shape, location, and scale parameters
  - Probability weighted moments (PWM) to predict lower-bound parameters
  - Method of moments (MOM) to predict upper-bound parameters
- Expected Mean Maximum Approximation (EMMA) to predict scaling behavior

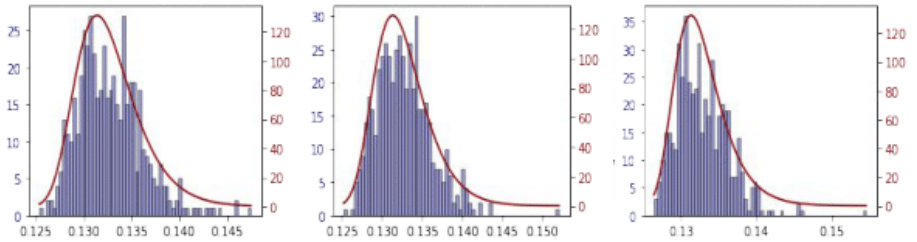

a) FTQ workload variation and MOM fitting

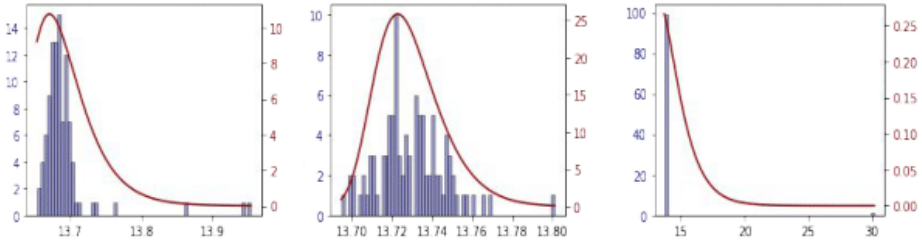

b) HPCG workload variation and MOM fitting

Modeling Approach: Parametric

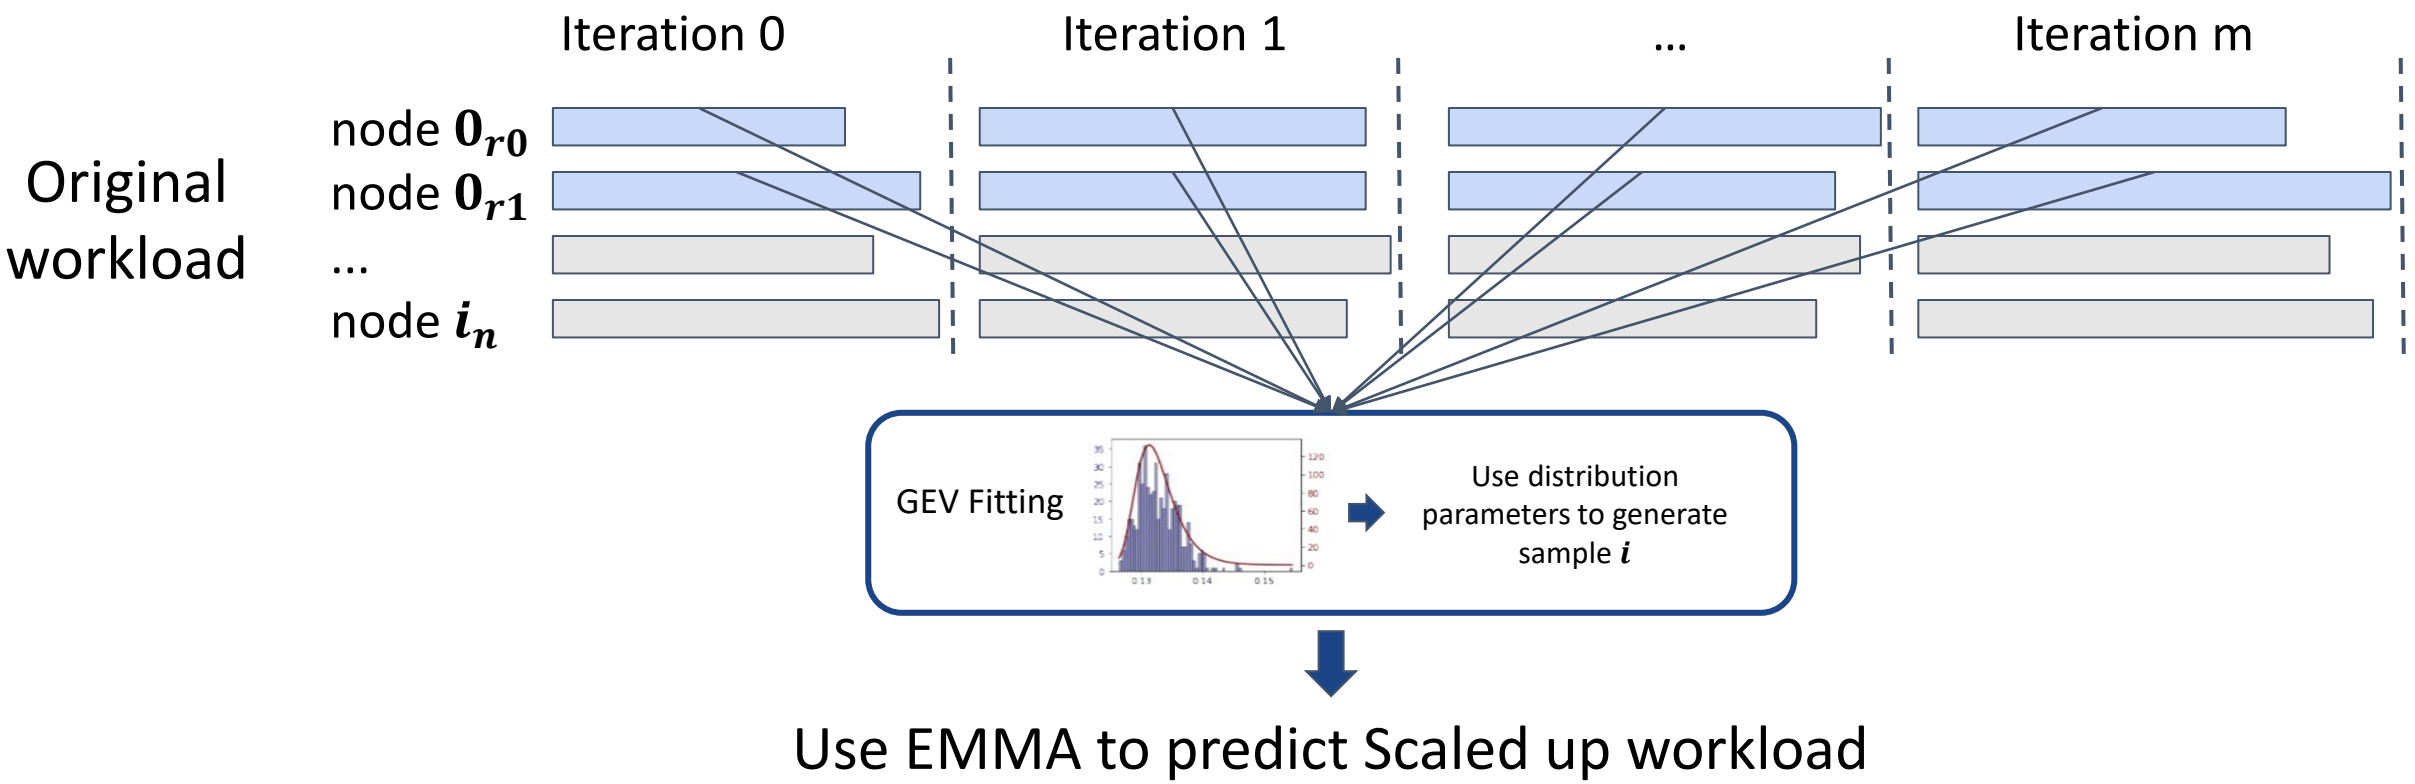

Complex Workloads

Attaway: Parametric Per Rank

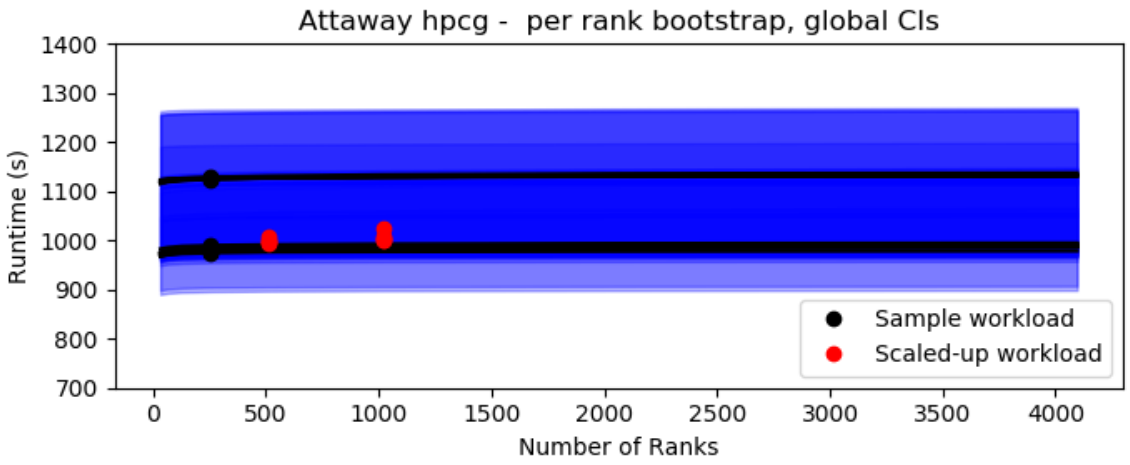

Cori: Parametric Per Rank

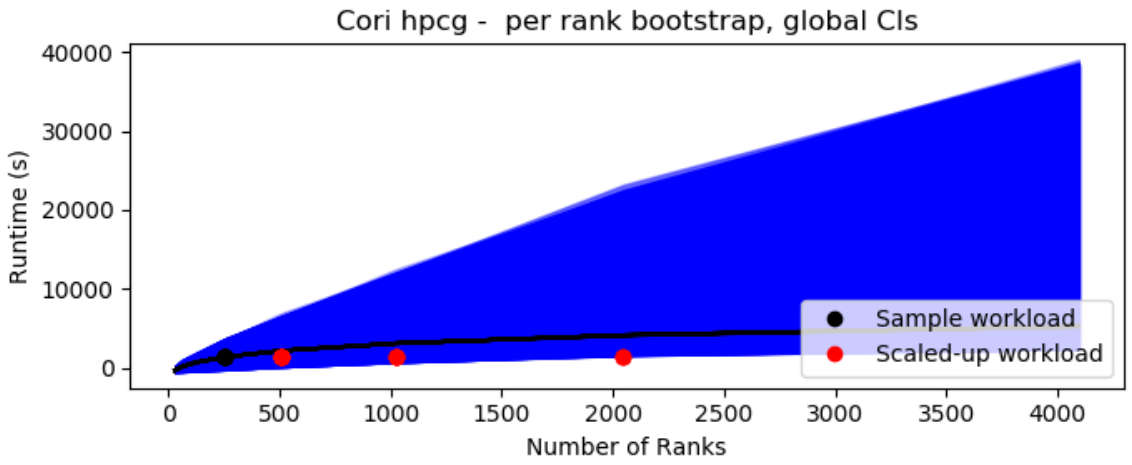

Discussion: Complex Workloads

- Prediction on noisy systems is challenging
- Parametric per rank prediction required

- Why?
  - Communication occurs during each iteration
  - Currently do not gather data prior to internal synchronizations
    - Consequence: Tail disproportionately weighted
    - Solution: Refine measurement method and model

Future Work

- Examine hybrid methods of parametric and non-parametric bootstrapping
  - Improve complex application behavior model and predictions
- Develop systems for performance anomaly detection in large scale systems
- Expand measurement framework to capture hybrid/thread parallelism
  - Goal: Model realistic applications using MPI+OpenMP or MPI+Kokkos, etc.

Next Steps: Modeling Threads

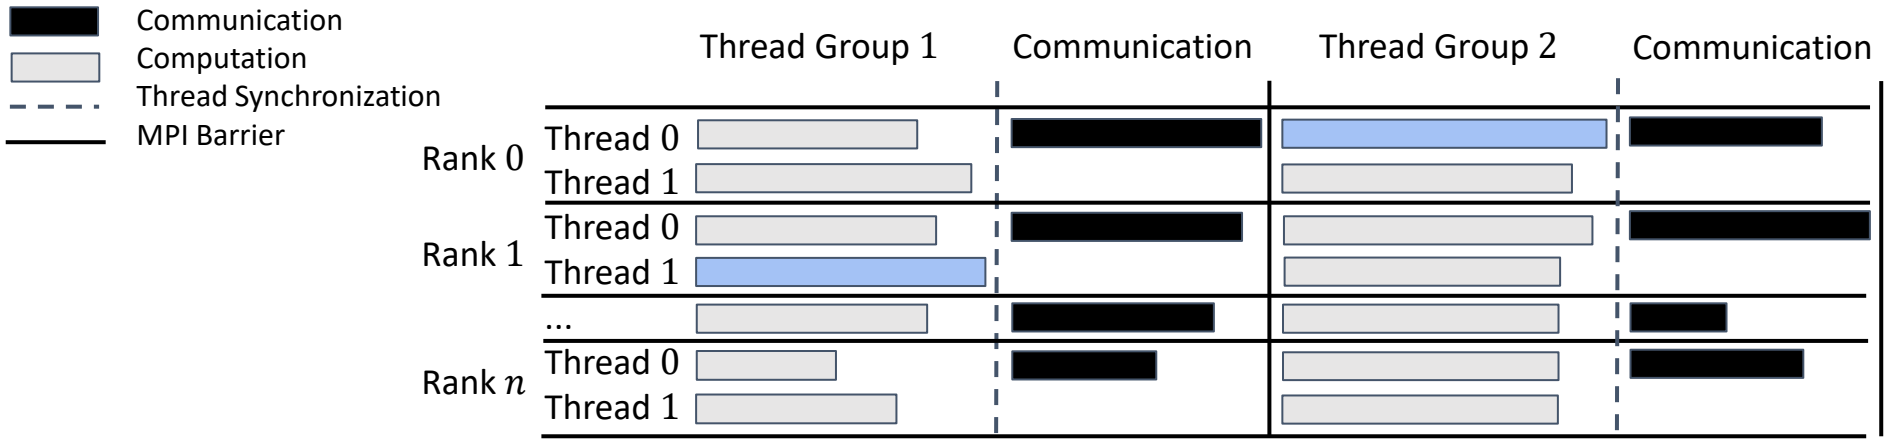

- Why is threading important?
- How can we use the same modeling approach?
- Goal: Utilize similar model to characterize performance and trade-offs of threaded/hybrid codes
  - Predict how partitioned and other communication methods scale in applications
  - Predict scaling performance and trade-offs of MPI only vs. MPI+OpenMP vs. other combinations

Conclusions

- Novel approaches for modeling application performance and variation
  - Significantly reduce the amount of measurement and data required
  - Generalizable framework can be applied to range of HPC workloads and systems

- Parametric and non-parametric methods trade-off data quantity for model fidelity

Questions?