



Profiling and Modeling for Application and System Analysis

ModSim22

Heidi Poxon

August 10, 2022



Topics

- Profiling and modeling to assess application behavior
- Lessons learned
- How these approaches transfer to cloud when analyzing performance
- What's needed in the future?

Filesystem

Processor

Node Architecture

Memory

Accelerator

What Impacts Performance?

System and Developer Software

System Size

Application

Network

Analysis Considerations

Learn how an application performs the way it does and why

- Analyze whole program, single node, code segment
- Run in dedicated or loaded environment
- Perform static code analysis, roofline, impairment, or sensitivity studies

Modeling Application Behavior

Learn how an application reacts to different architectural characteristics or in different environments

- Run or profile and see what you get, then extrapolate
- Create equations (simple or complex) to evaluate behavior
- Imitate application or hardware behavior
- Feed traces or applications through performance simulators
- Combine profiling and modeling to increase modeling effectiveness

Performance Assessment Approaches



Perform multiple runs and compare figure of merit

Use system tools (dtrace, turbostat, memstat, etc.)

Insert timers into code (printf)

Use profiling tools

What Makes a Profiling Tool Effective?

- Reduces time investment to port and tune applications
- Analyzes whole-program at scale for bottlenecks
- Provides simple interfaces with capability to target complex HPC applications
- Offers actionable guidance

Challenges

Periodic use

No one tool fits all needs

Those elusive performance counters

Profiling-ready applications

Lessons Learned from a Tool User Perspective

- Choose tool to fit the goal (2x vs 3% speedup, evaluating new architecture)
- Consider the investment
- Understand the why, not just the what
- Analyze the job size and problem size you intend to run
- Aggregated data is good, but need the ability to break it down too
- Understand how data is aggregated by a tool

Lessons Learned from a Tool Developer Perspective

- Offer simple steps to obtaining performance results
- Both high level summaries and details are important
- Choosing the most meaningful information for a single run is difficult
- Sometimes the simplest solutions can be the most effective
- Aggregate data, derive metrics, but also provide the ability to review non-aggregated data
- Describe how data is aggregated
- Balance overhead injection with insight
- Providing more data is not necessarily better

Cloud – a New Performance Capture Environment

Continuity Between Cloud and Local Data Centers

- Another set of resources for computing and data analytics
- Application and system performance still matters!
- Yes, you can use performance counters (but not quite as many)
- You have access to dedicated, high performance resources
- System tools, timers, profilers, modeling are available

Example Differences

Virtual environments and hypervisors

Technology choice with more frequent uplifts

You become the administrator

Topology details and some controls are not available

Bringing it Together to Address Common Problems

- Tool platform portability
- Profiling interfaces for programming layers (libfabric, RAJA, Kokkos, DPC++)
- Data explosion due to large scale systems and converged workloads
- Profiling and modeling require expertise
 - Simplicity and familiarity for wider adoption
 - Performance teams or “Centers of Excellence” for profiling goals, methodology, consistency, knowledge transfer



Thank you!

hpoxon@amazon.com

