Modeling State-of-the-Art (A Case for Ubiquitous Modeling)

Adolfy Hoisie

ASCR Workshop on Modeling and Simulation for Exascale Systems and Applications

August 9-10, 2012

Seattle, WA

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Outline

- A taxonomy of methodology
- Applications of modeling
- Landscape of modeling
- Modeling on the way to Exascale

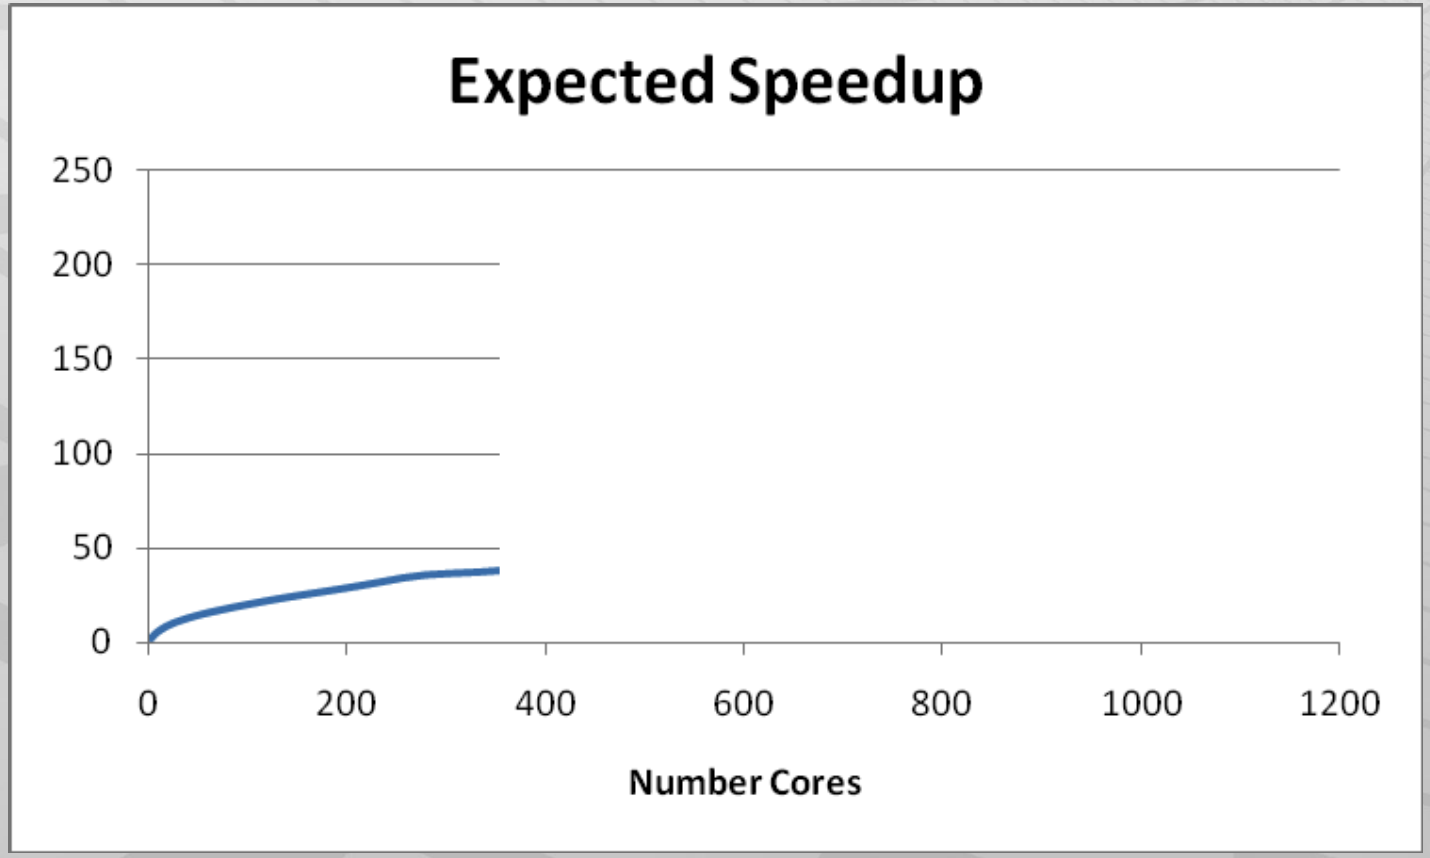
Expected Speedup
250
200
150
100
50
0
0
200
400
600
800
1000
1200
Number Cores

What Makes Performance Prediction Challenging?

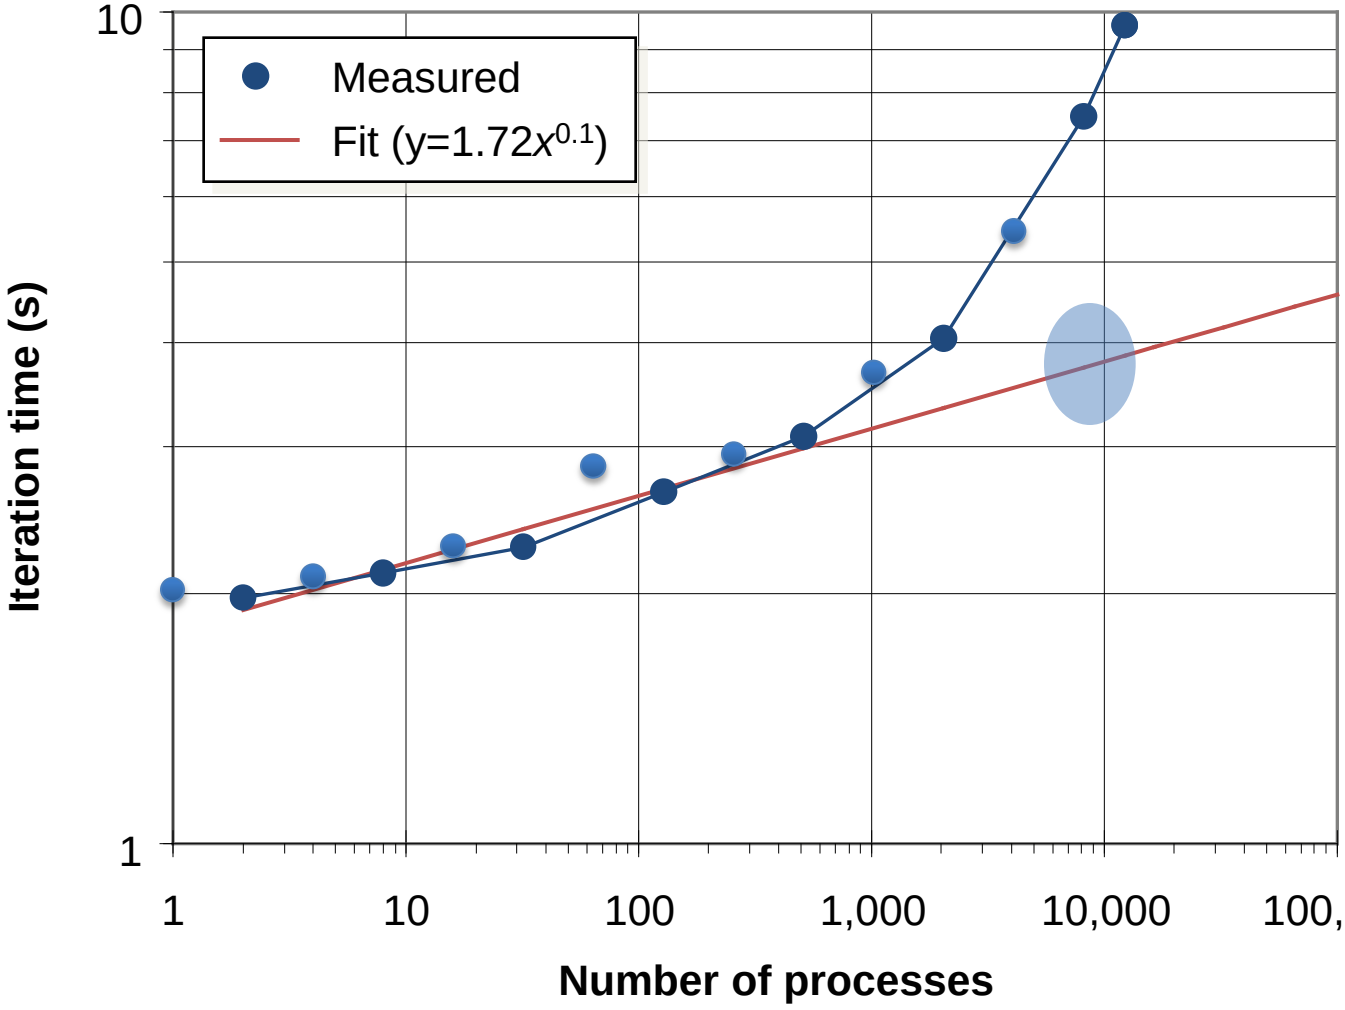

What's the expected iteration time of an 8192-process run?

Challenge #1
Performance characteristics may change at scale.

Challenge #2
Nonlinear behavior may be caused by either the system *or* the application.

Curve fitting does not provide performance *insight*!

Methodology

- (Quasi-)Analytical (PNNL, SDSC, others)
- Statistical: machine learning, simple statistical models, MC, curve fitting (LBL)
- Analytical/dynamic models (PNNL,SDSC)
- Bag-of-tools (PNNL)
- Hybrid simulation-emulation (UIUC, GT)
- Multiscale (GT)
- Cycle accurate simulation (Industry, Sandia, many others)
- Model generation (LLNL, PNNL, others)
- Low-fidelity model generation (UIUC)

- Let's complete it and populate it by the end of the workshop!

Modeling: many uses

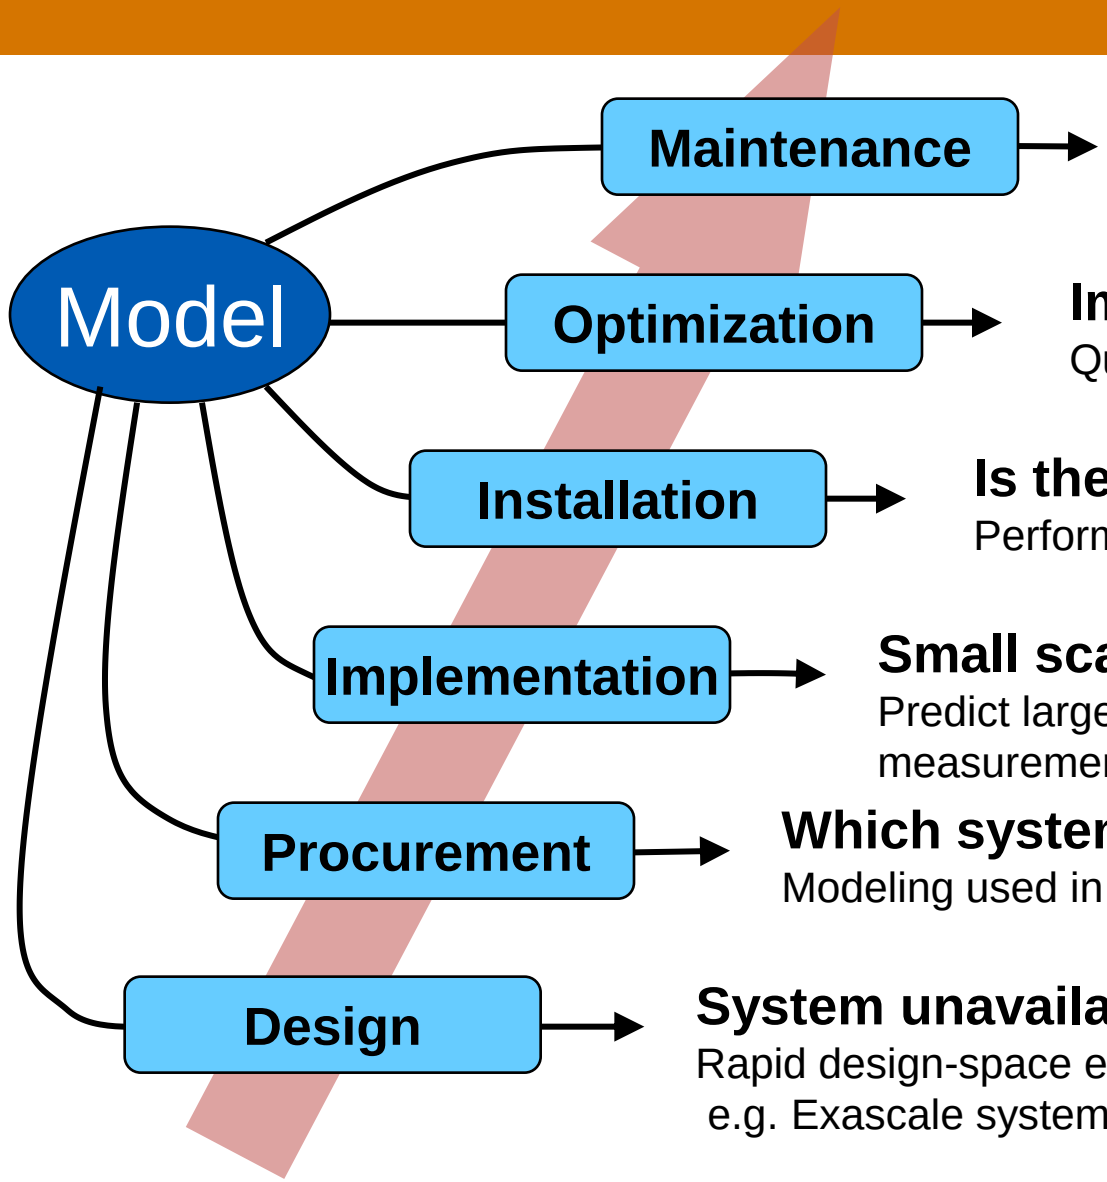

Runtime operation
Is the system healthy today?

Improvements
Quantify impacts prior to implementation

Is the machine working?
Performance should be as expected

Small scale (nodes) available
Predict large-scale system performance using measurements @ small-scale

Which system should ABC buy ?
Modeling used in procurements for almost a decade

System unavailable for measurement
Rapid design-space exploration.
e.g. Exascale system performance

More details on uses of modeling

- Modeling as co-design tool
- What-if scenarios
- Exploration of architecture space
- System and application optimization
- System design
- Algorithm design
- Application design

Landscape of modeling

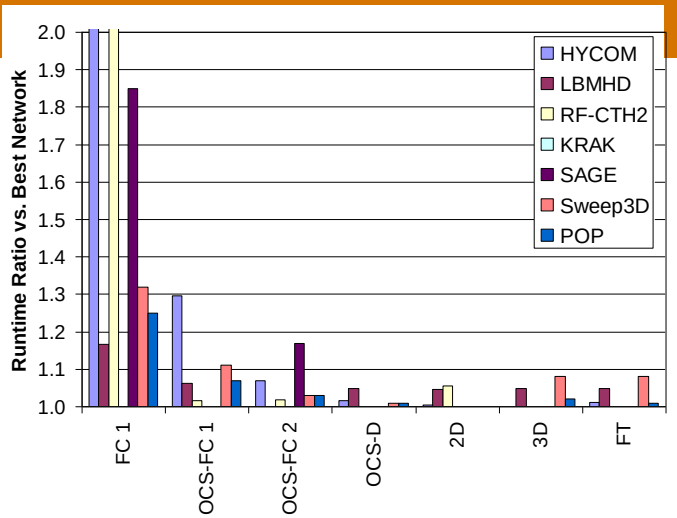

Design space exploration
(e.g. P7 and topology)

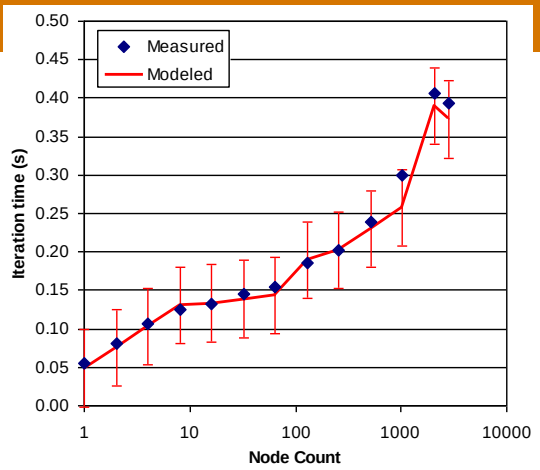

Acceptance testing
(Validate system)

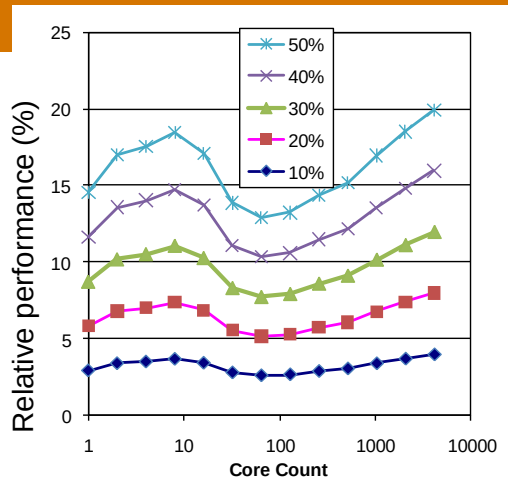

Changes in application

Explore in advance & Optimize at run-time

- **Change of metrics to power**
- **Key attributes for models:**
  - Rapid evaluation (e.g. alternate execution paths)
  - Accuracy (within reason)

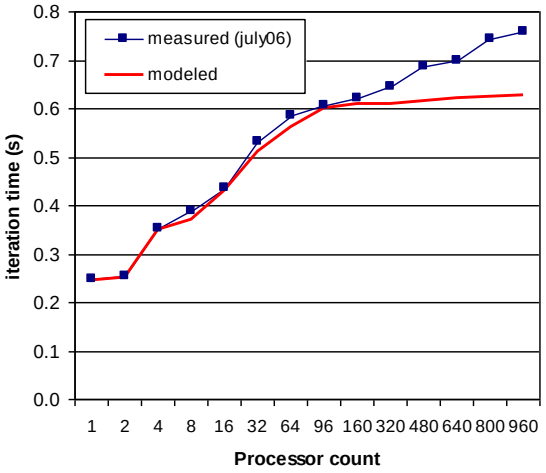

Model driven
Optimizations

Modeling in action: architecture exploration

- How would different compute and communication parameters impact performance, prior to the implementation?

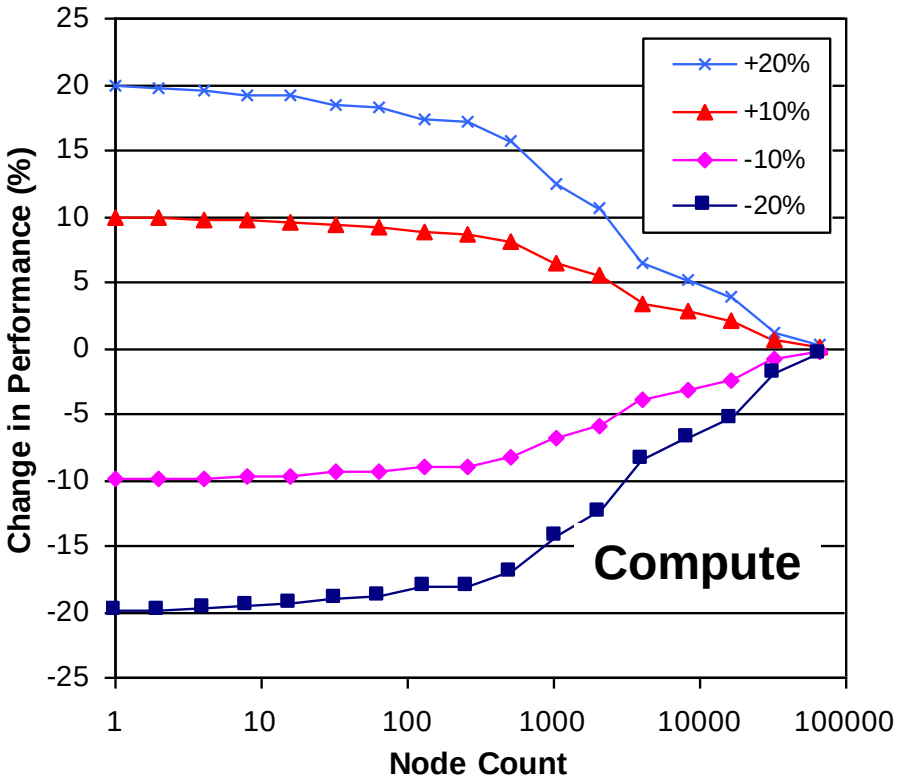

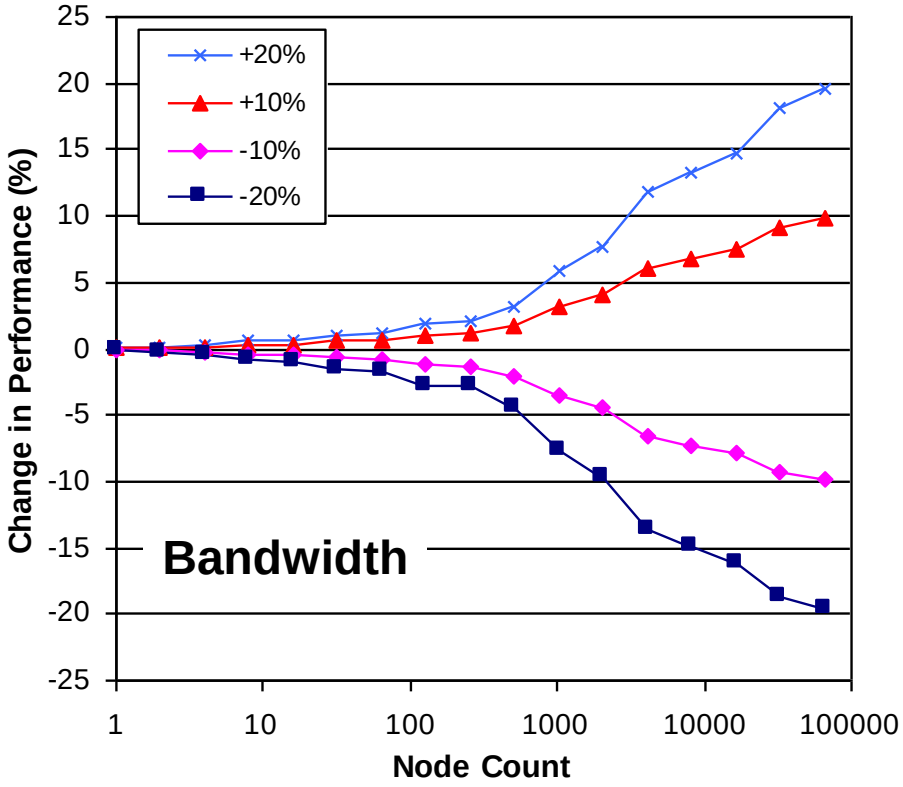

Application optimized and mapped to different architectures before implementation

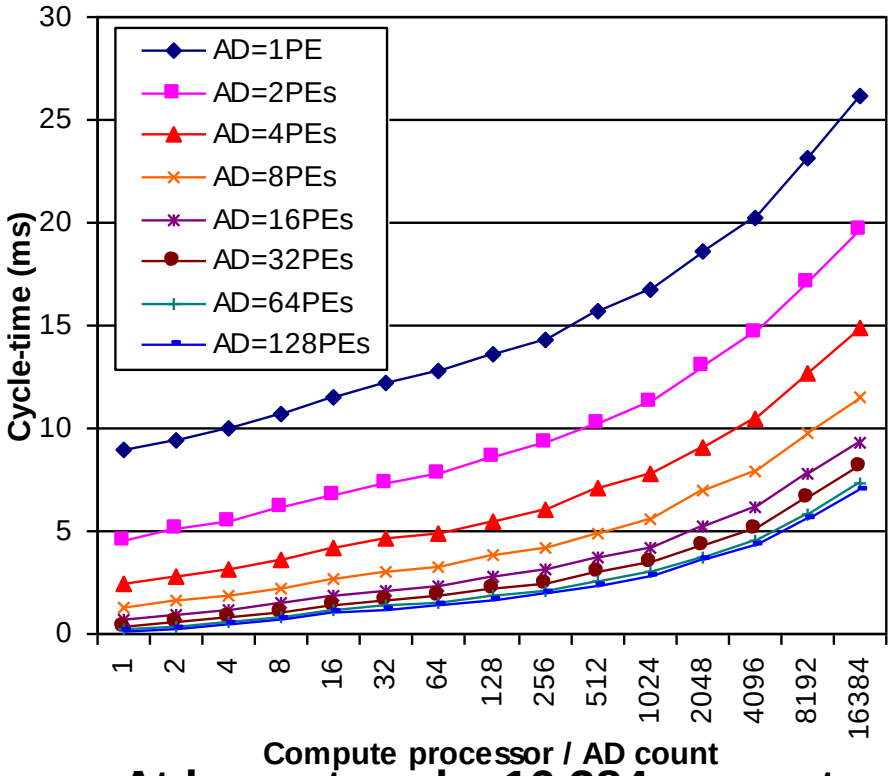

- **System assumptions**
  - Hypothetical accelerated system
  - Inter-PE (on Accelerator)
    - » Bandwidth =1GB/s,
    - » Latency = 50ns
  - Inter-node (MPI)
    - » Bandwidth = 1.6GB/s,
    - » Latency = 4µs
- **Application assumptions**
  - Weak-scaling
  - Problem definition consistent across architectures
  - "Knobs" used to optimize the application trough modeling

- **At largest scale, 16,384 compute processors & 16,384 accelerators**
  - Improvement is ~3.5x when using Accelerators with 128x more PEs

"A Performance Analysis of Two-Level Heterogeneous Processing Systems on Wavefront Algorithms", Kerbyson, Hoisie, Unique Chips and Systems, CRC Press, 2007, pp. 259-279.

Co-design at Petascale: Roadrunner

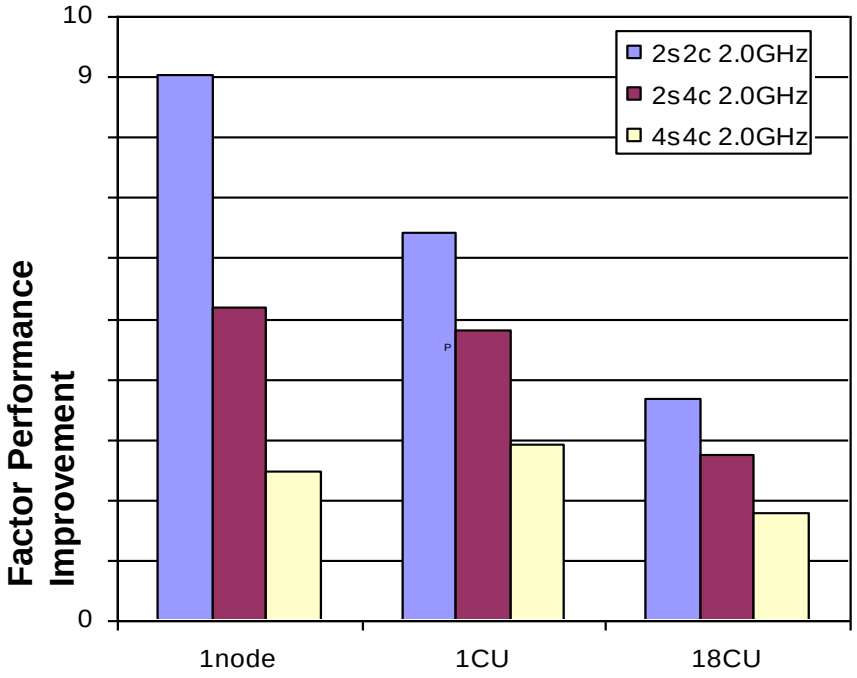

Performance of Roadrunner against hypothetical architectures

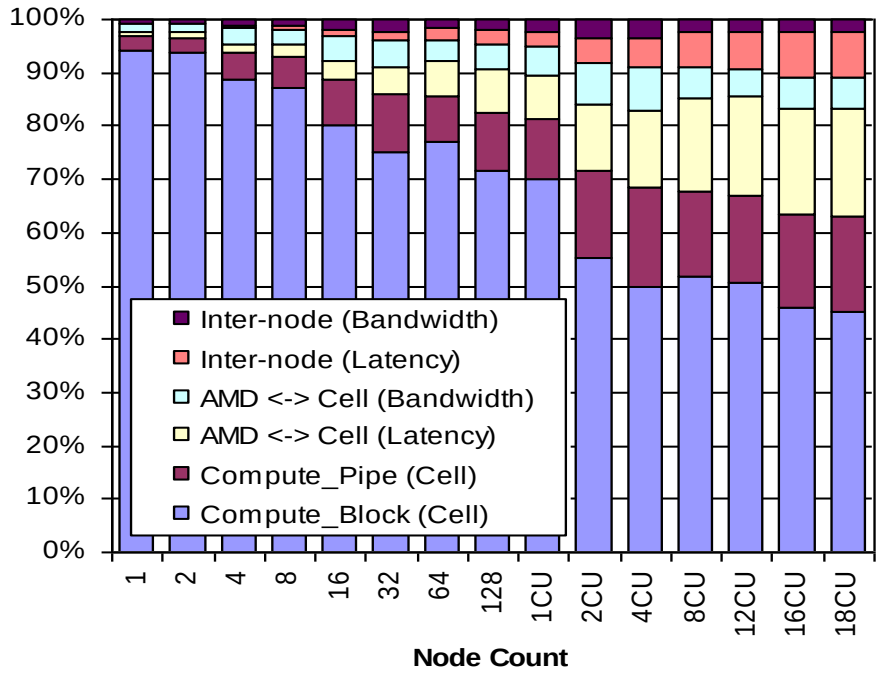

Architectural insight using modeling

- Performance modeling in action!
- Performance/power/reliability modeling in concert

"Entering the Petaflop Era: The Architecture and Performance of Roadrunner"", Barker, Davis, Hoisie, Kerbyson, Lang, Pakin, Sancho, SC'08, 2008.

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

In the future: ubiquitous modeling

- **Performance & Power & Reliability**
  - together
- **Bag-of-tools approach** –
  - not one for all but all for one.
  - modeling, simulation, and emulation.
- **Lifecycle coverage** –
  - software and hardware,
  - from design space exploration, to analysis of early implementation, to deployment, and to run-time optimizations.
- **Co-design** –
  - modeling need be applied to negotiate tradeoffs at all the boundaries of the Hardware/Software stack
- **Dynamic Modeling** –
  - intelligent and informed decision within runtime software
- **Introspective runtime** –
  - dynamic hardware and software, rapid optimizations.
  - the runtime system is model driven, and the model is actionable

Future uses of modeling

- Model at the center of the hardware/software stack
- New areas: modeling EMs, PMs, boundaries between the app/system SW/HW
- Dynamic modeling
- Model-driven runtime systems
- Modeling energy
- Modeling reliability
- Modeling the triad per/power/reliability *in concert*
- Model as an actionable tool (at runtime)
- Monitoring/mitigation tools
- Lifecycle use for apps and system

The Model as a first class citizen

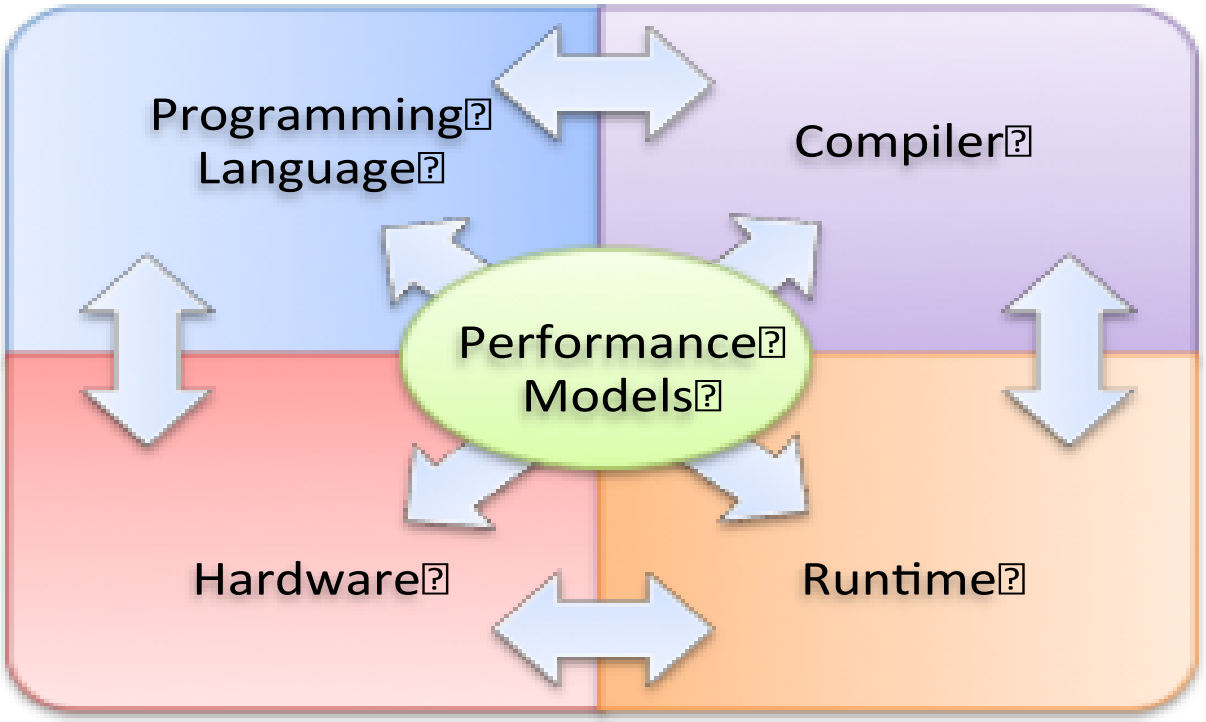

Co-design: Many views are discussed but few are used in practice

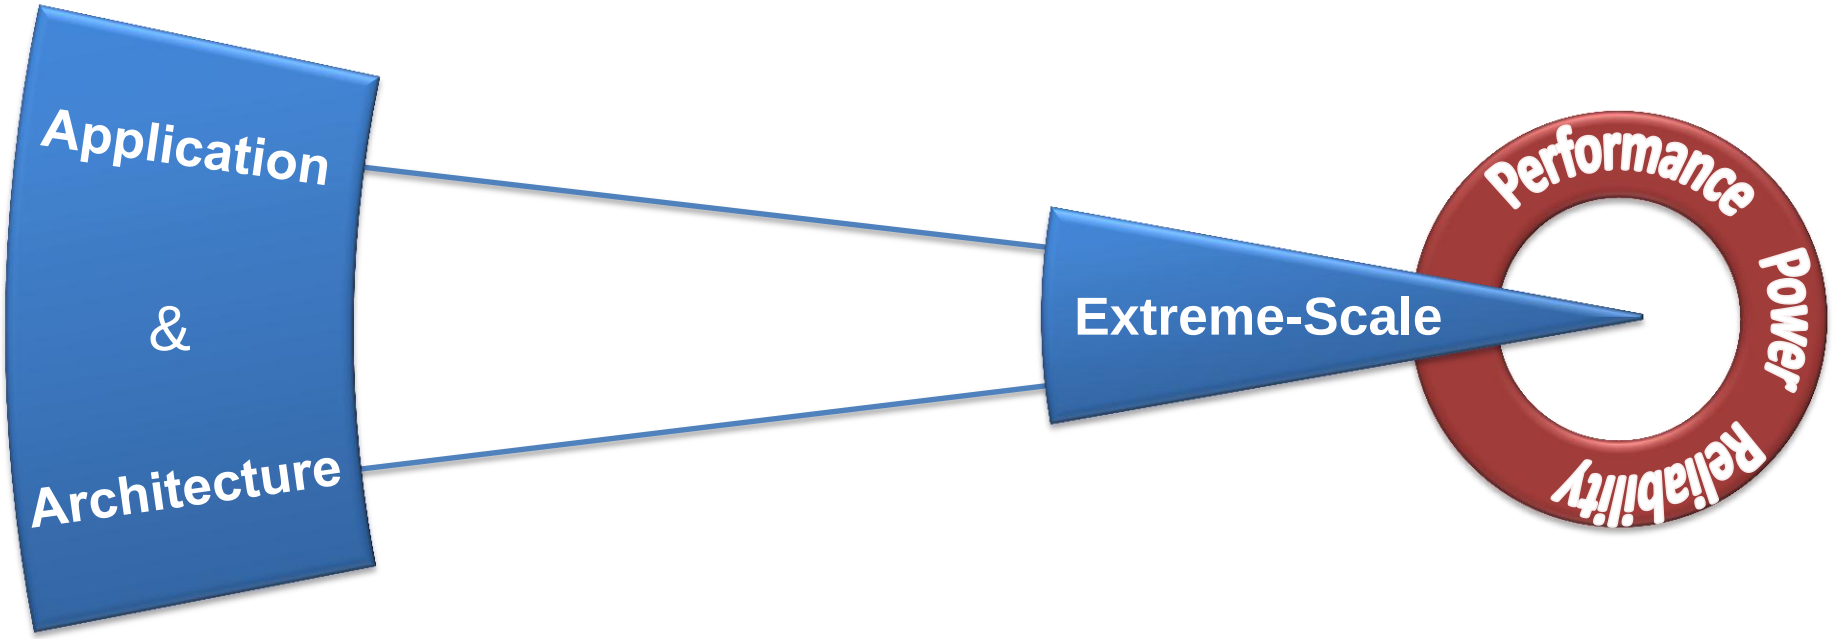

Co-design: Many views are discussed but few are used in practice

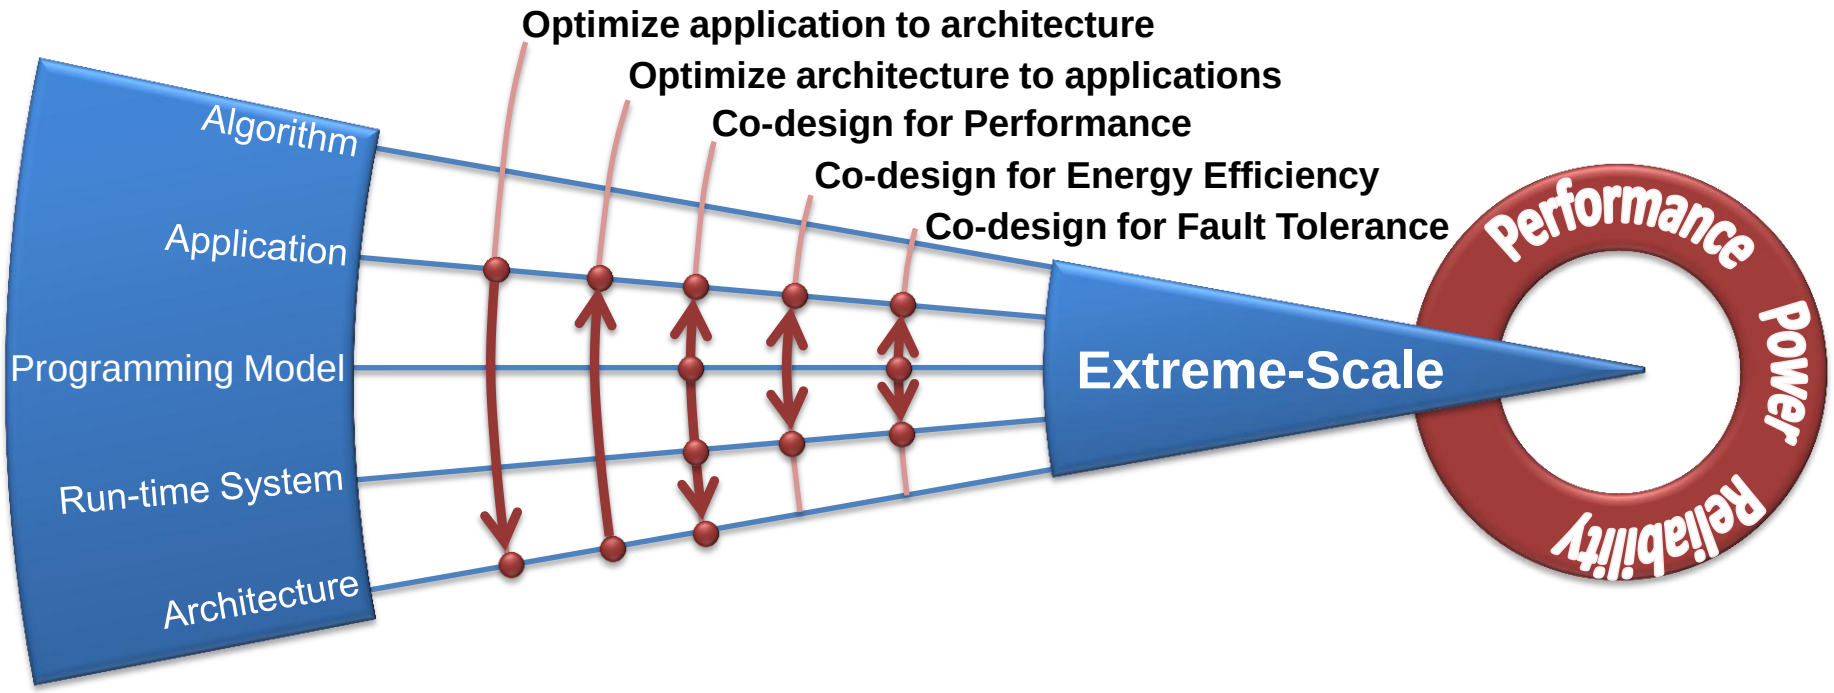

"Co-design for Exascale systems: Performance, Power, Reliability", Kerbyson, Vishnu, Barker, Hoisie, IEEE Computer, Nov. 2011

Modeling Execution Models: AntiCiPate framework

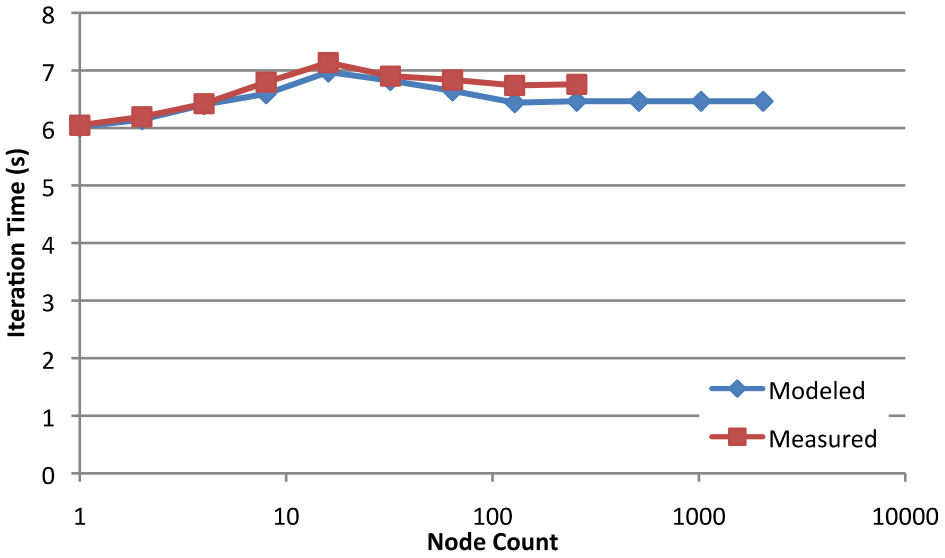

(a) Modeled vs. Measured Performance

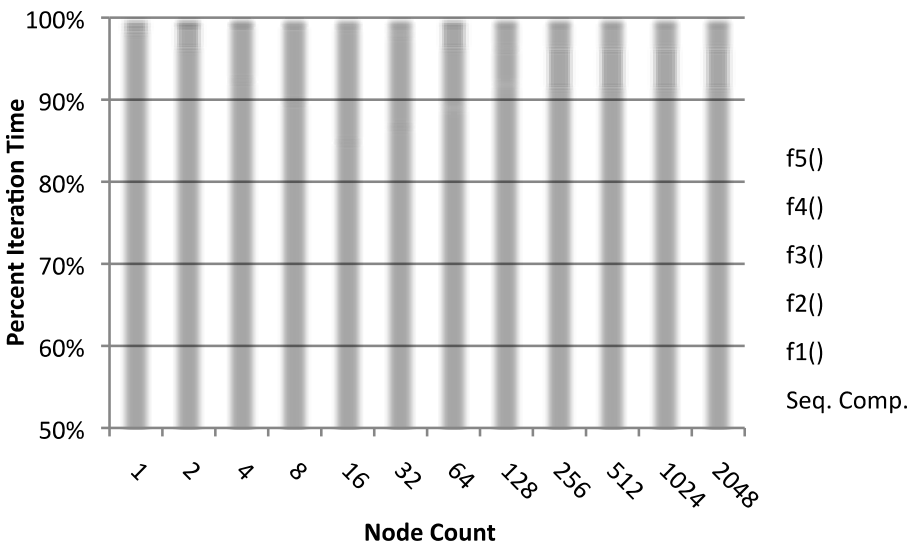

(b) Runtime Breakdown by Function

- First model of an application based on EM knobs
- Modeled vs. Measured performance shows high model accuracy
  - Maximum Error < 5%
- Runtime breakdown observations:
  - Performance is sensitive to load/store performance – Actually TLB miss rate, which varies with the *npartdom* parameter of the input deck
  - Synchronization accounts for < 5% of overall runtime
- Could now modify various EM parameters and analyze the impact predictively

All about the Trade-offs

- **Cannot start too early in design considerations (co-design)**
  - Design space exploration: System & Application
- **Research in each of Performance, Power and Reliability will lead to multi-dimensional optimizations**
  - New metrics
  - Trade-offs
  - Performance at what power
  - Reliability at what power
  - Data-movement costs
    - e.g. Moving instructions to data
  - Power steering
    - e.g. coping with dark-silicon
- **Increased focus on run-time methods**
  - Introspective
  - Need for rapid techniques
  - Performance, power, and reliability

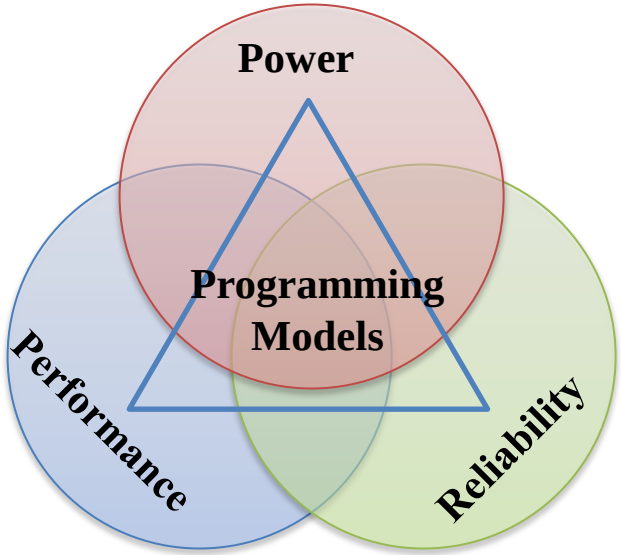

Summary

- The challenges of modeling for Exascale need to be tackled as a community
- Co-design is the key, in the broad definition of the term – modeling is a key technology
- A bag-of-tools approach to co-design is in order – there is no one size fits all
- Progress is being made, but…
- Significant investments are needed in the area, emphasizing coordinated efforts