IAA Interconnects Workshop Panel

Performance Prediction and Simulation for Exascale Interconnection Networks

N. Binkert (HP Labs), D. Chiou (UT Austin),
A. Rodrigues (SNL), S. Yalamanchili (GTech)

The Challenge

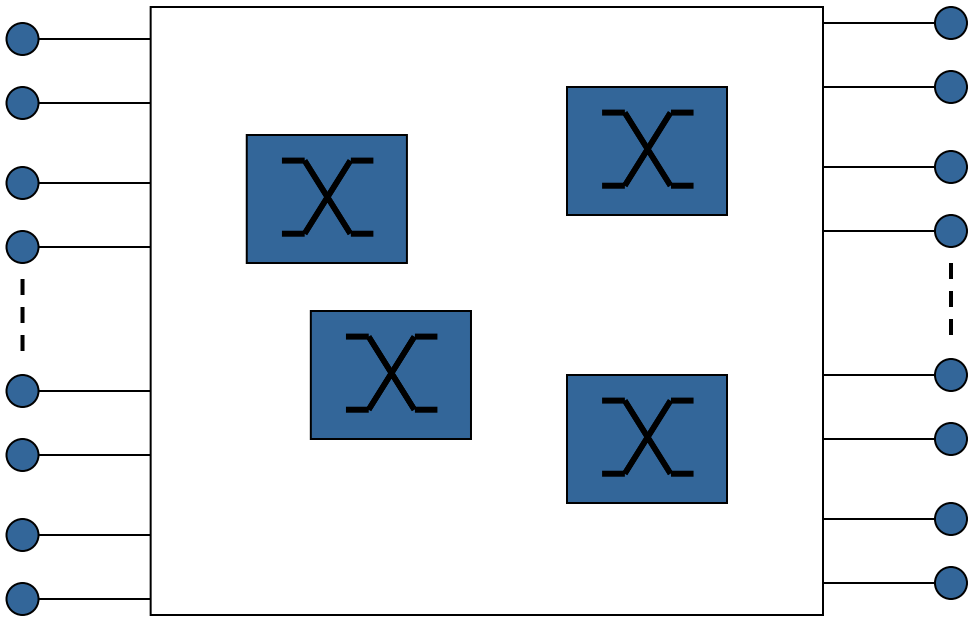

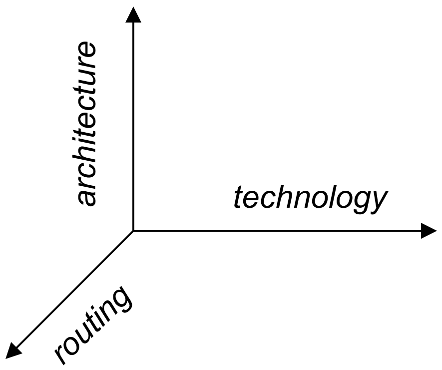

- Model communication between 10M to 20M hardware entities
  - Cores, memories, I/O devices
- An additional factor of 2-16X in terms of producers and consumers
  - For example threads, parcels, bundles, etc.
- Subject to technology (e.g., power), architecture (e.g., switching), and behavioral (e.g., routing) attributes

Some things to consider

- Complexity: number of components being modeled
  - Proportional to (K * #terminals)
- Granularity
  - Events/port/switch/clock
  - Analytic vs. simulation
- Design Space
  - F(topology, physical parameters, microarchitecture)
- Do we build only what we can analyze/predict?

Questions

- Challenges for performance prediction and simulation
- Promising avenues for mitigating these challenges
- Are we constrained by what we can predict?

Challenges – Beyond Point Tools

- Parallelism
- Interactions
- Automation
- Modeling Complexity

1. Parallelism

- Use parallel machines to design Exascale machines
- Parallelization of models
  - Discrete event vs. time-stepped vs. hybrid
  - Partitioning and global virtual time management is key
  - Some recent examples: ROSS (RPI), COTSON (HP Labs), CAPSTONE (GTech)

2. Interactions

- How does the network design impact the core, I/O and memory system design and vice versa?
  - Cache sizes impact network demand
  - Latency impacts number of concurrent transactions that need to be supported
- Correlate applications and traffic
  - Program phase behavior and traffic
- Modeling reliable networks

3. Automation

- Model description, composition, deployment, and experiment specification
- Platforms for leveraging point tools → customization of models
- Data explosion
  - Mining performance data for analysis

4. Modeling Complexity

- Lessons from the architecture community
  - Statistical sampling, simpoints, etc
  - Works better for networks?

- New techniques that work for networks
  - For example coupling routing, in-order delivery, and MPI → drastically simpler models for specific applications?
  - Validation pilots on current machines

An Integrated Simulation Environment[1]

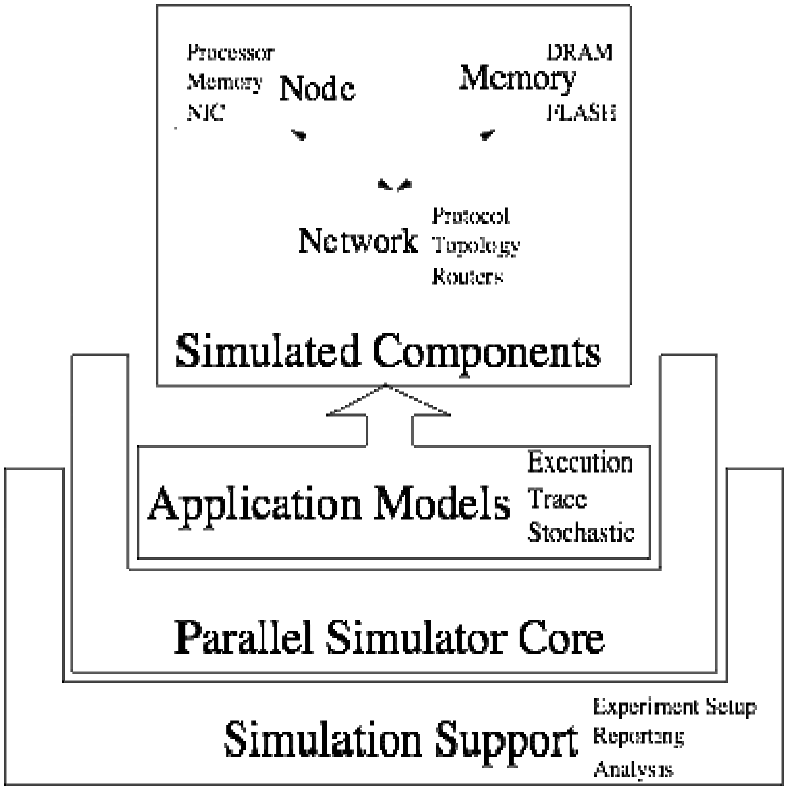

- Core platform for parallel simulation
- Requirements
  - Scalable parallel simulation
  - Multi-scale simulation
  - Hardware acceleration
  - Standardized interfaces to existing tools for nodes, memory and I/O models
    - Workload stimulus
    - Event management
    - Time management
    - Synchronization services

[1]Gordon B. Bell (IBM), Derek Chiou (U.Texas-Austin), Bruce Jacob (U.Maryland), Curtis Janssen (Sandia National Labs), George Riley (GTech), Arun Rodrigues (SNL), Philip Roth (Oak Ridge National Laboratory), Jeffrey Vetter(ORNL), Sudhakar Yalamanchili (Georgia Tech), Weikuan Yu (ORNL)

Automation

- System description language for model construction
- Experiment construction and reporting methodology
  - Facilitate comparisons and interpretation
  - Managed code base
- Export to standardized performance data format
- Determinism
- Partitioning and checkpointing services

Usage Model

Support non-herculean transitions from

- Software serial models to
- Serial hardware-accelerated models to
- Software parallel models to
- Software parallel, hardware-accelerated models

IAA Interconnects Workshop Panel

Performance Prediction and Simulation for Exascale Interconnection Networks

N. Binkert (HP Labs), D. Chiou (UT Austin),
A. Rodrigues (SNL), S. Yalamanchili (GTech)