Modeling and Simulation Challenges

Sudhakar Yalamanchili

School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA. 30332

Sponsors: NSF, DoE (SNL, LANL), SRC

Motivation

"Remember that all models are wrong; the practical question is how wrong do they have to be to not be useful."

Box, G. E. P., and Draper, N. R., (1987), *Empirical Model Building and Response Surfaces*, John Wiley & Sons, New York, NY.

George E. P. Box, 2011

George E. P. Box, 2011

Simulation Infrastructure Challenges

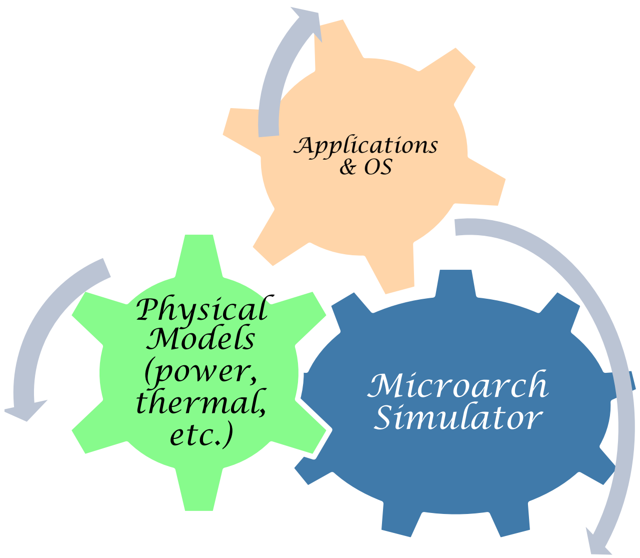

- Scalability
  - Processors are parallel and tools are not → not sustainable
- Multi-disciplinary
  - Functional + Timing + Physical models
- Need to model complete systems
  - Cores, networks, memories, software at scale
- Islands of expertise
  - Ability to integrate point tools → best of breed models
- Composability
  - Easily construct the simulator you need

Needs and Capabilities

Need to distinguish between modeling and engineering

Modeling

- Performance models of complex phenomena
- Abstract behaviors of interest
- Draw upon a palette of mathematical and simulation techniques

Engineering

- Construction of software or hardware implementations
- Modularity, composition, interoperability
- Practical determinant of ease of use

Challenge

Composition

Application   Microarchitecture   Physical Behaviors

Common APIs!

Example: Thermal Coupling

I. Paul and A. Vanderheyden

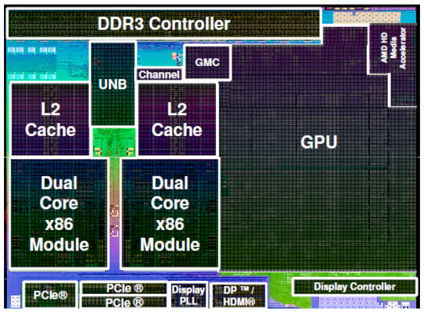

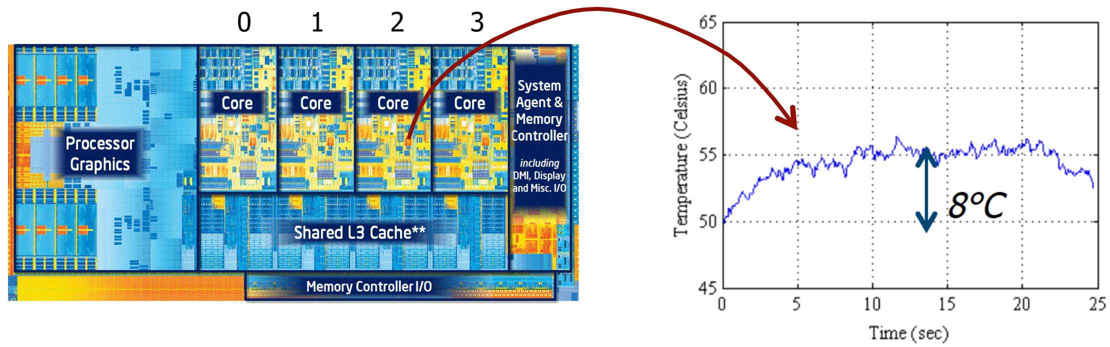

Temperature on Core 2 when Core 3 is busy and remaining cores are idle

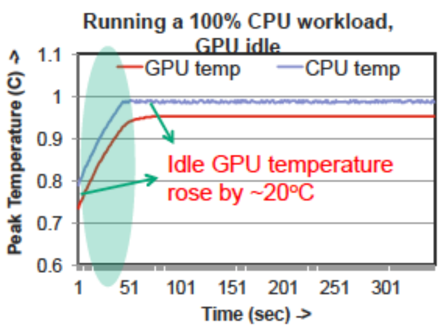

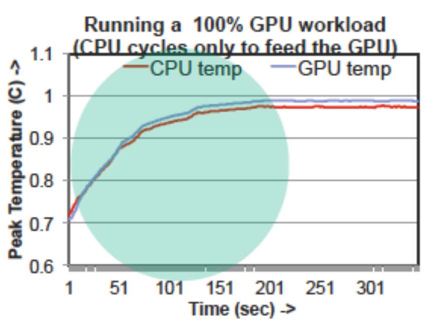

- Significant rise in temperature of the idle component due to thermal coupling and pollution
- CPU cores consume thermal headroom more rapidly (4X faster)
- Better management for significant gains in measured energy efficiency are possible
- Power management ≠ thermal management

I. Paul, et.al., "Cooperative boosting: needy versus greedy power management", *ISCA 2013.*

Performance Improvements: Cooperative Boosting

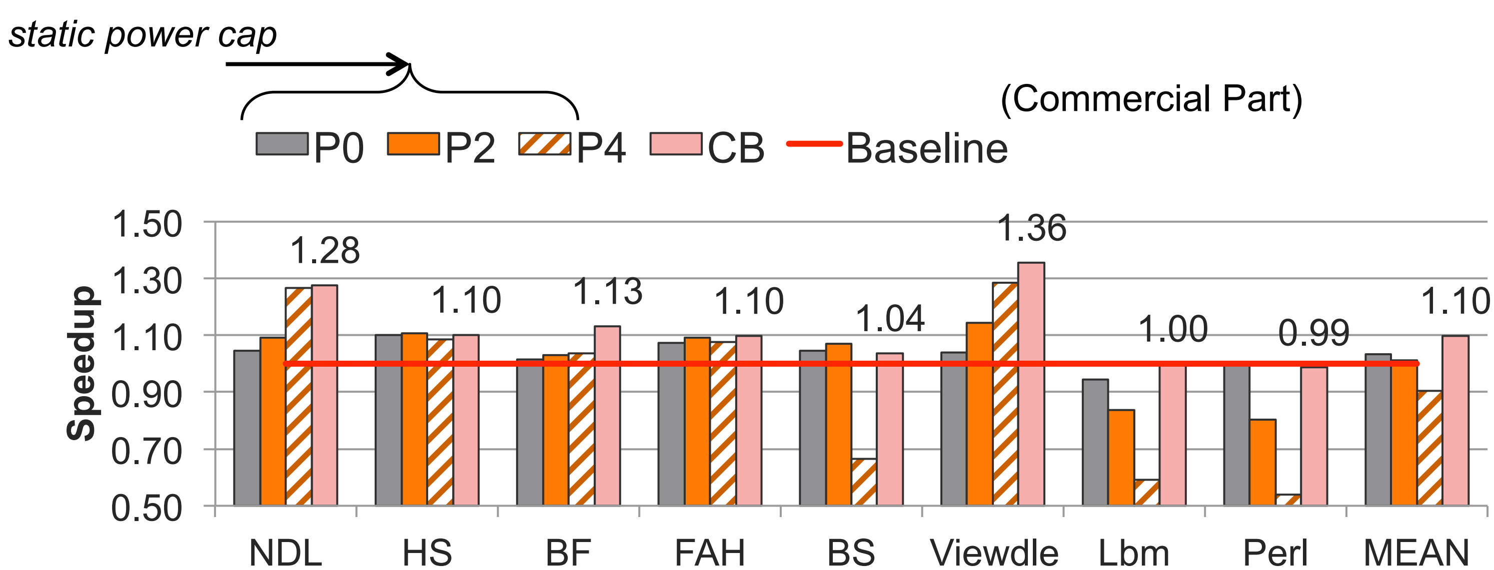

CPU	DVFS-state
HW Only (Boost)	Pb0
	Pb1
SW-Visible	P0
	P1
	P2
	- - -
	Pmin

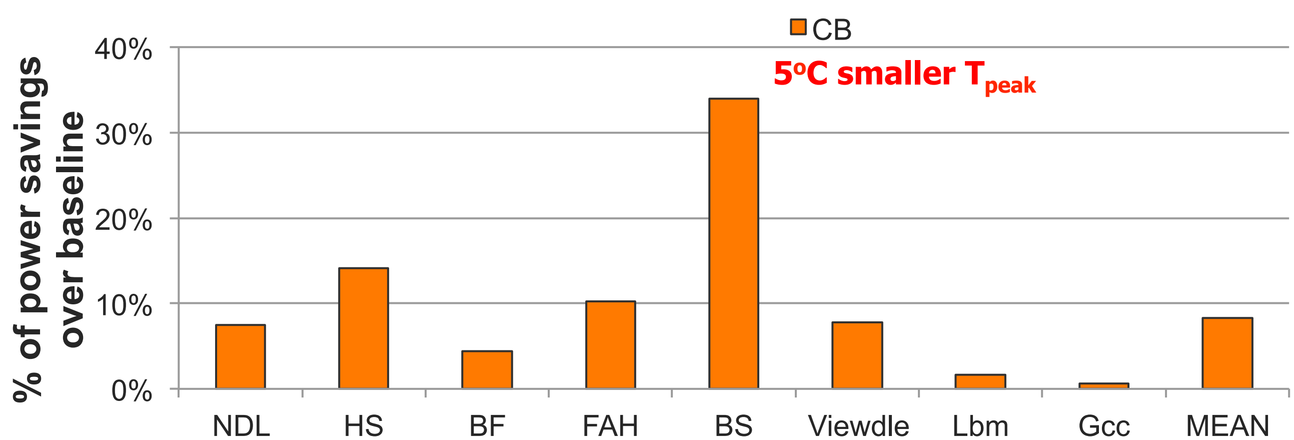

Manifold: Execution Model (Socket/Blade)

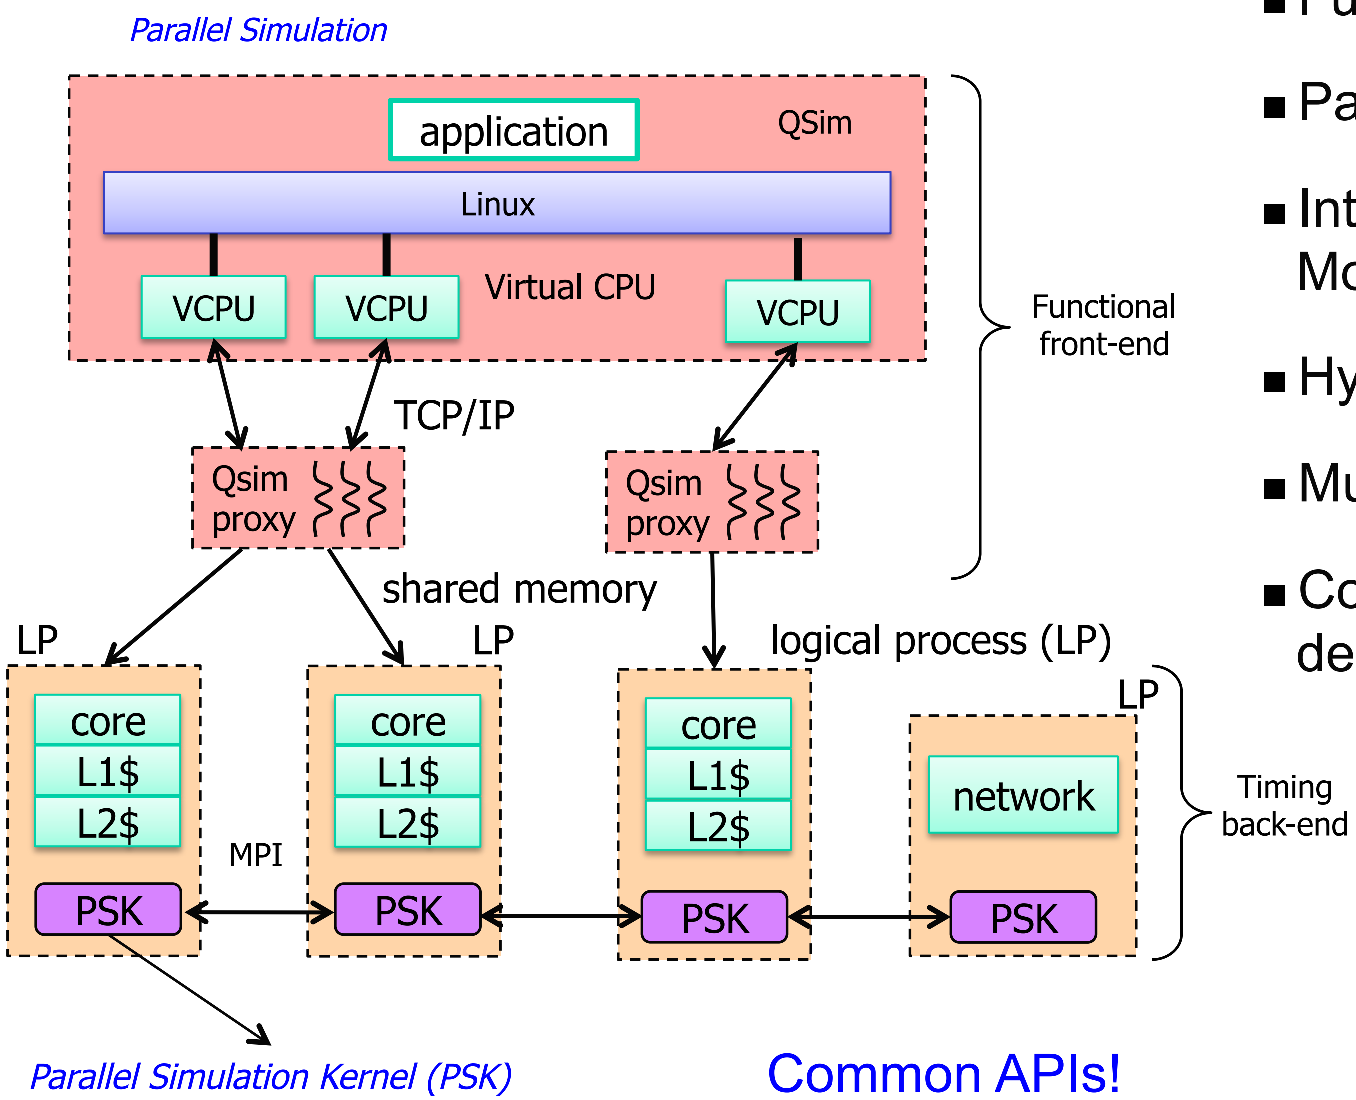

- Full-system simulation
- Parallel Simulation
- Integrated Physical Models
- Hybrid timing model
- Multiscale
- Component-based design

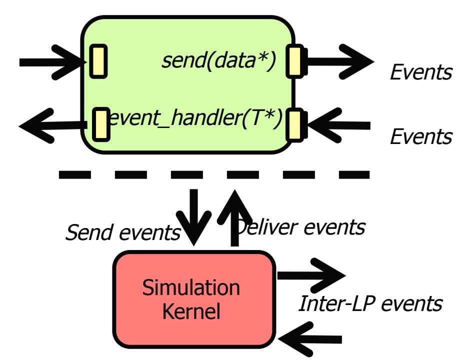

Common APIs!

A Manifold Socket/Blade Simulation

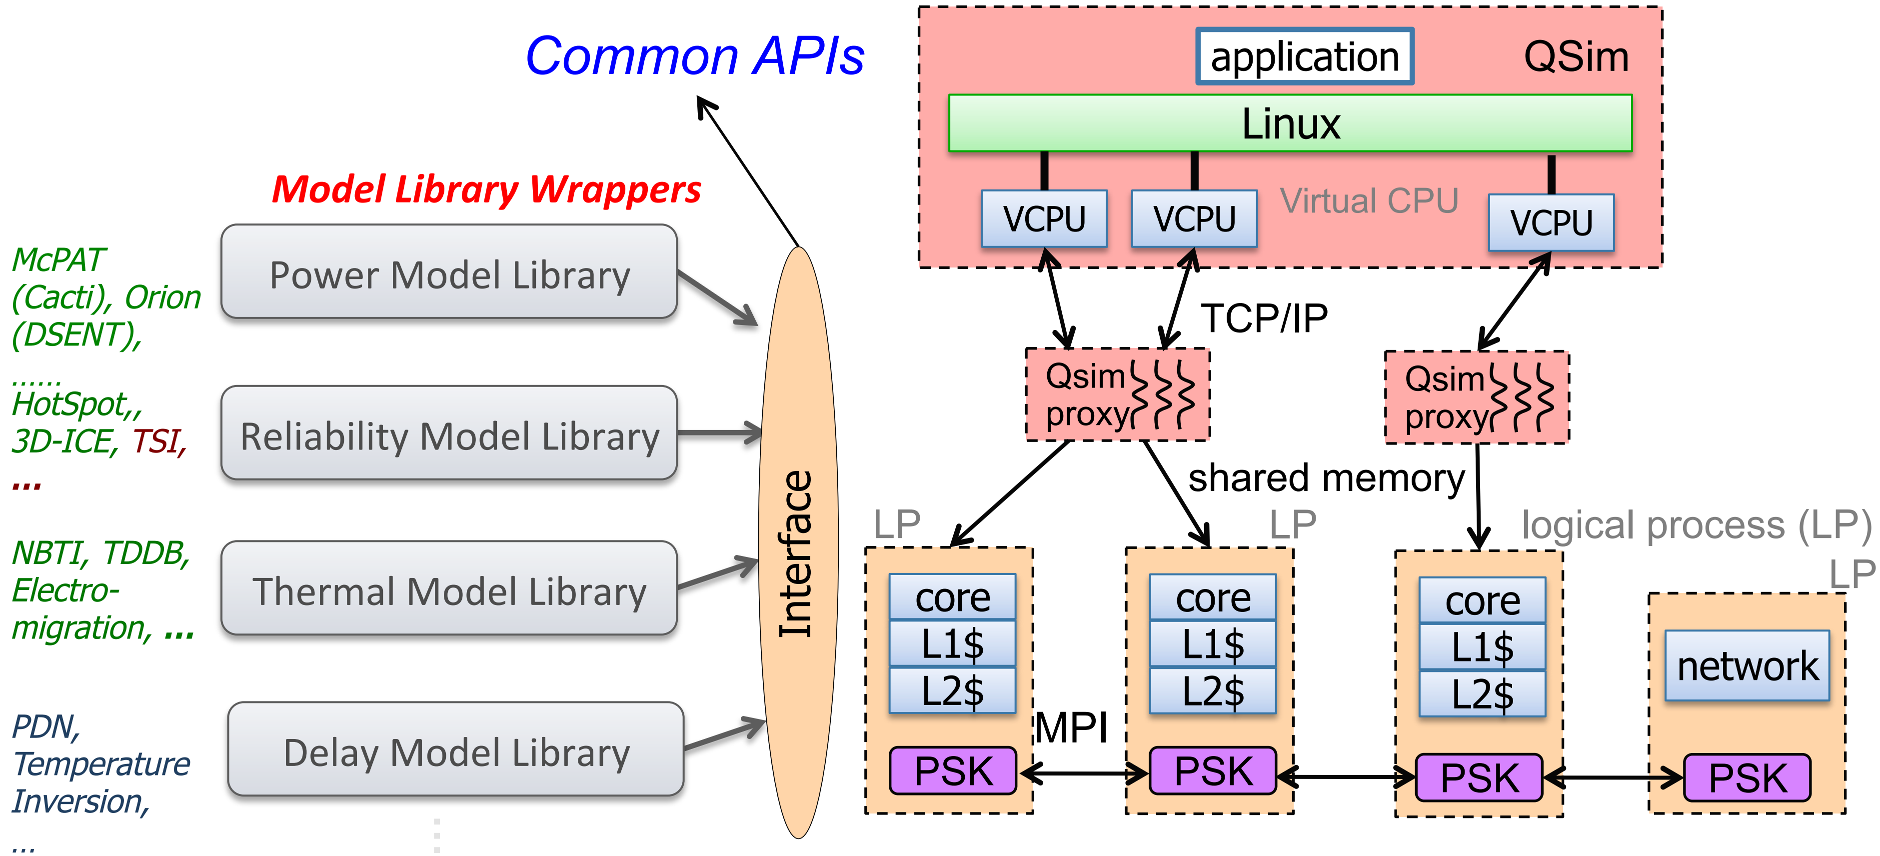

- Composition of application models (full emulation to skeletal), timing models, and physical models

Processor Representation

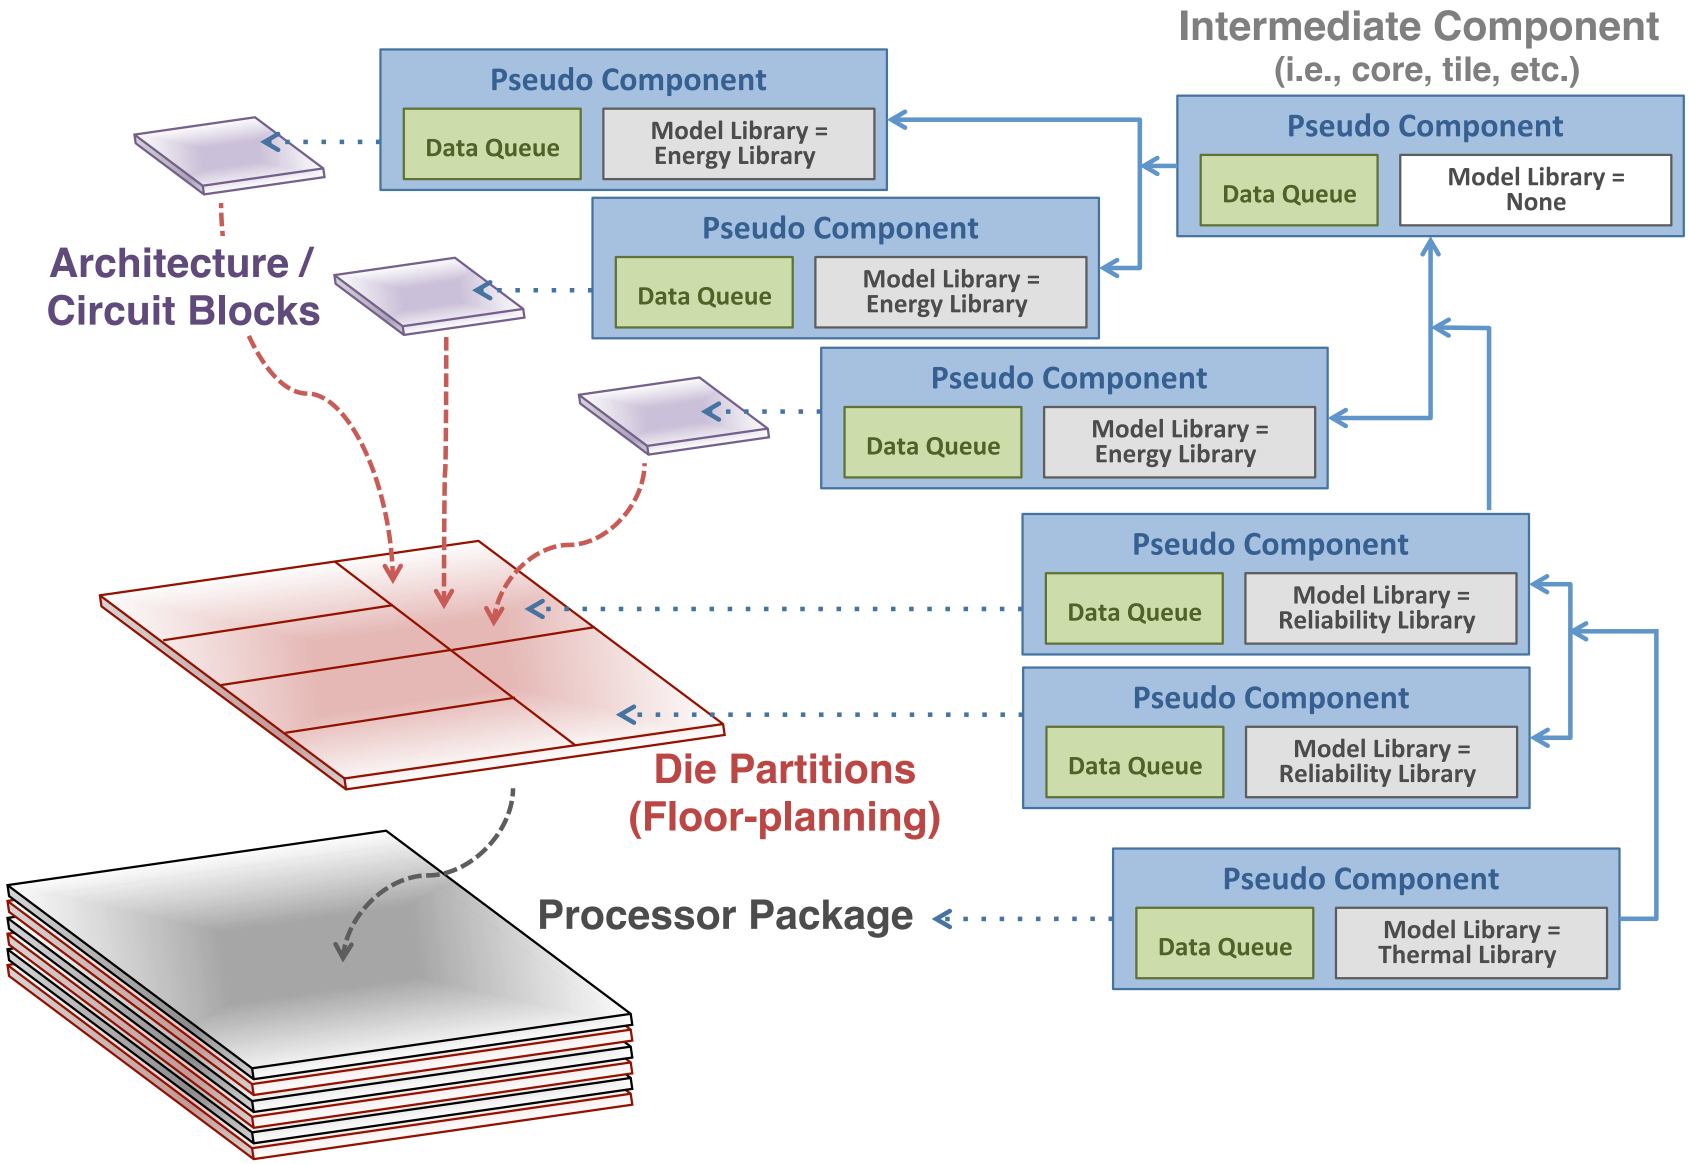

Example: Workload Cooling Co-Design

2.1mm x 2.1mm

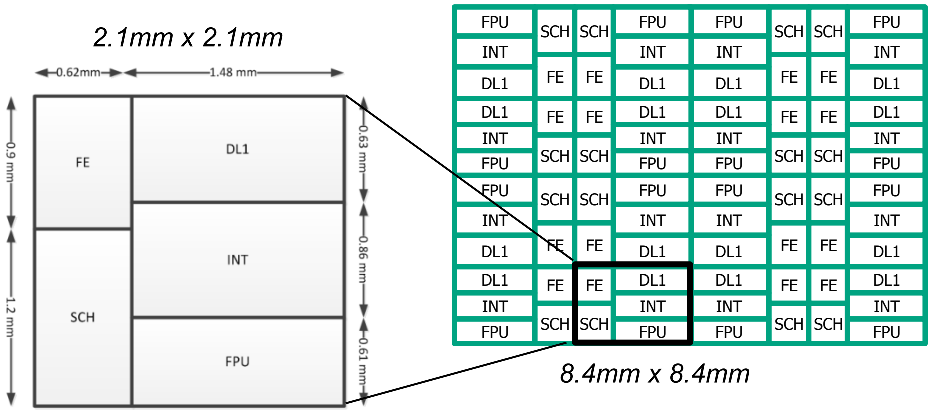

8.4mm x 8.4mm

16 symmetric cores

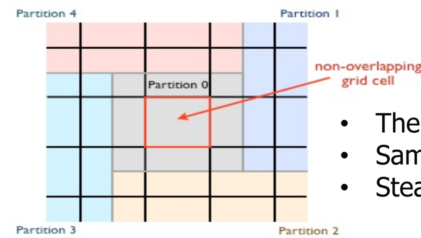

- Thermal Grids: 50x50
- Sampling Period: 1us
- Steady-State Analysis

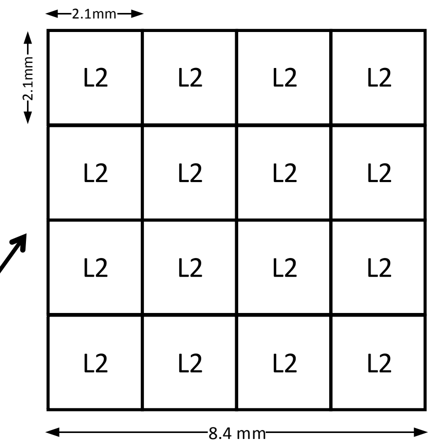

Nehalem-like, OoO cores;
3GHz, 1.0V, max temp 100°C
DL1: 128KB, 4096 sets, 64B
IL1: 32KB, 256 sets, 32B, 4 cycles;

Ambient:
Temperature: 300K

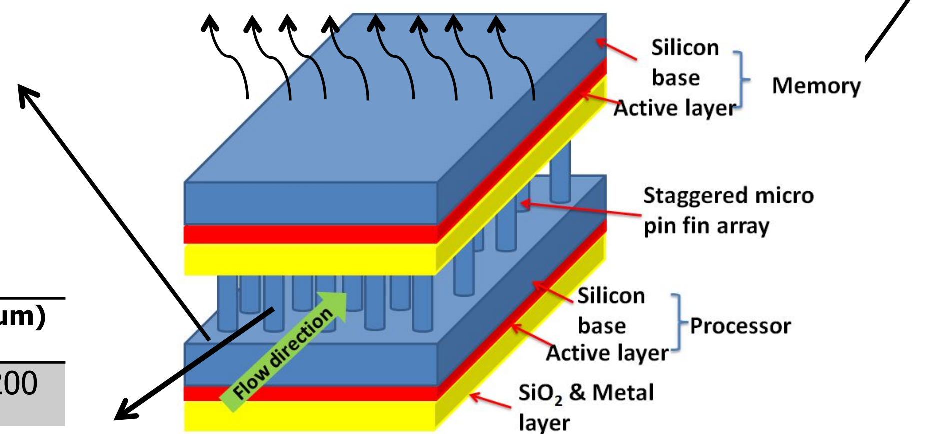

L2 & Network Cache Layer:
L2 (per core): 2MB, 4096 sets, 128B, 35 cycles;
DRAM: 1GB, 50ns access time (for performance model)

	DP (um)	PS (um)	HP (um)
baseline	100	200	200
optimized	180	320	400

DP: diameter, PS: pitch spacing, HP: height

H. Xiao, Z. Min, S. Yalamanchili and Y. Joshi, "Leakage Power Characterization and Minimization over 3D Stacked Multi-core Chip with Microfluidic Cooling," *IEEE Symposium on Thermal Measurement, Modeling, and Management (SEMITHERM)*, March 2014

Architecture Performance: Example

- Results from an example simulation
- Optimized pin fin structure
  - Energy Per Instruction (EPI) 40% over the worst case.

Normalized EPI comparison among all 4 pin fin structures

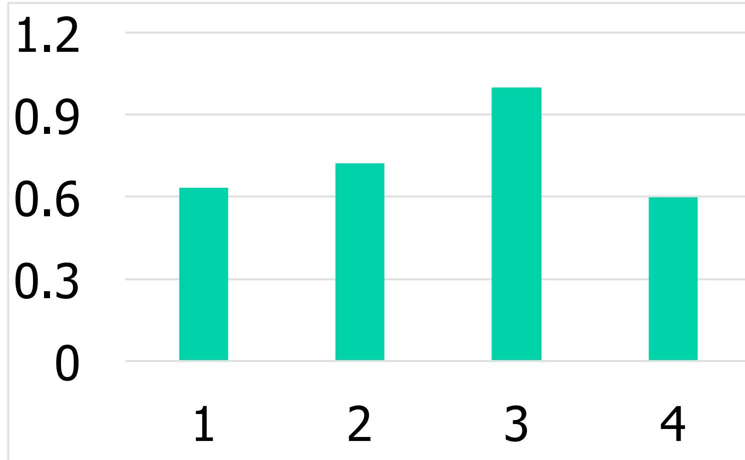

Courtesy L. Zheng (ECE) and Professor Muhannad Bakir (ECE)

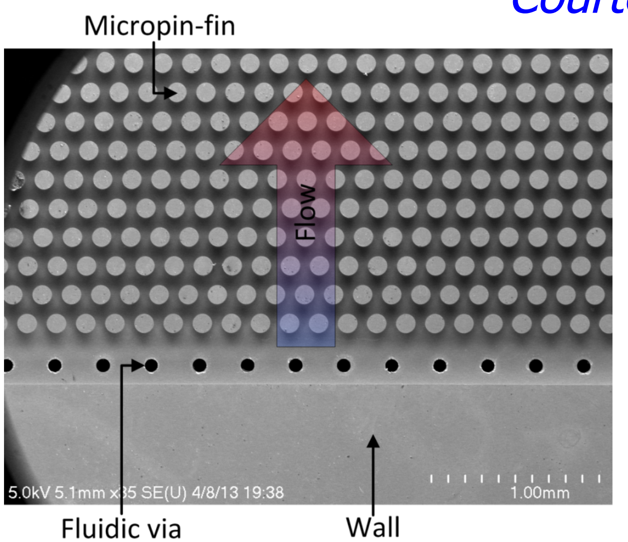

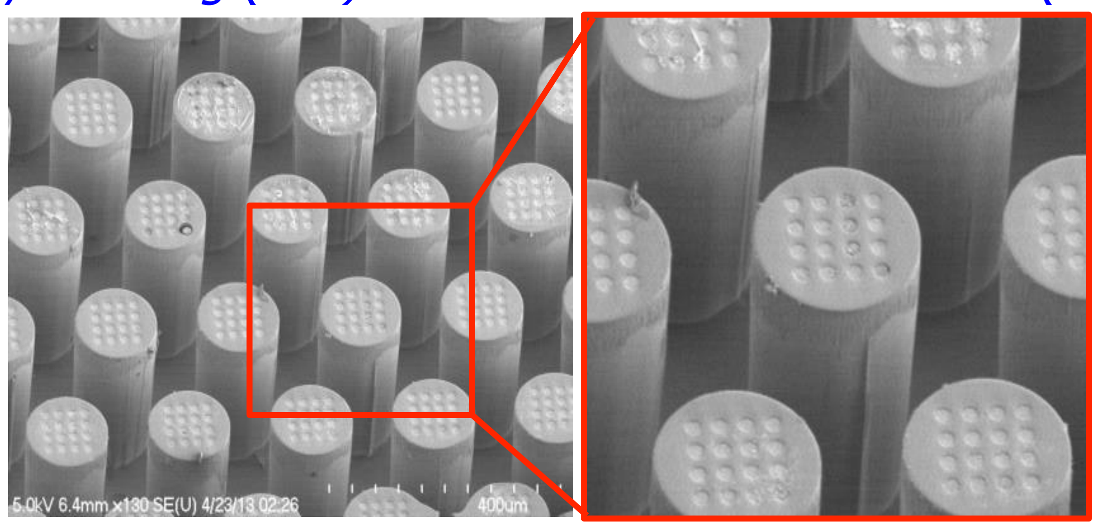

Example: Adaptive Regulation

- Adaptive control algorithms utilize the DVFS capability of microprocessors to regulate power, thermal, or throughput to constant level.
- Energy Introspector provides an interface to apply dynamic execution controls, e.g., DVFS.

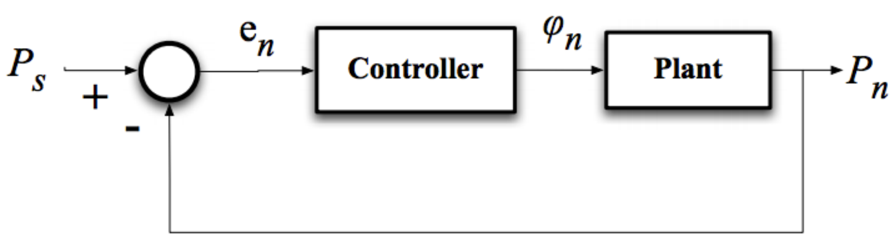

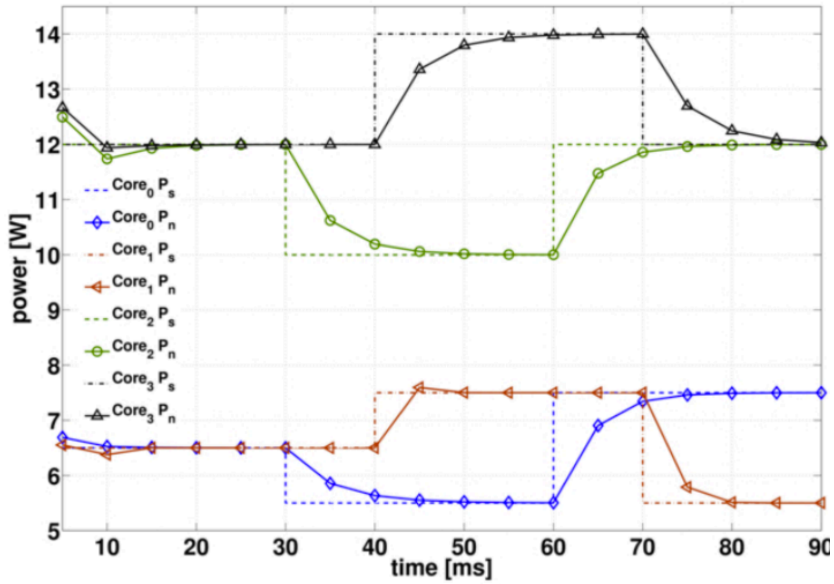

1. N. Almoosa, W. Song, S. Yalamanchili, and Y. Wardi, "Throughput Regulation in Multicore Processors via IPA," CDC, 2012.
2. N. Almoosa, W. Song, S. Yalamanchili, and Y. Wardi, "A Power Capping Controller for Multicore Processors," ACC, 2012.

Scaling Simulations: Fidelity vs. Scale

- Composition
  - Common APIs! – sharing IP models
  - Separation of time, event, and synchronization management
- Hierarchy of fidelity
  - Example: Application skeletons, state machines
- Parallelism
  - Simulation capacity scales with compute capacity
  - Track Moore's Law?
- Integrated Physical Models
- Need to support Co-Design
  - Power delivery and package design
  - Every Joule counts!

Summary

www.manifold.gatech.edu

- Composable simulation infrastructure for constructing multicore simulators
  - Common APIs
  - Parallel execution
  - Integrated physical models
- Provide base library of components to build useful simulators
- Distribute some stock simulators
- Need: Validation Techniques

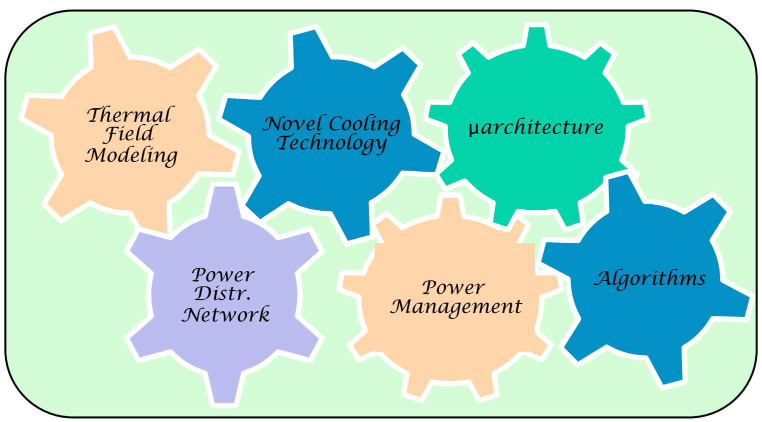

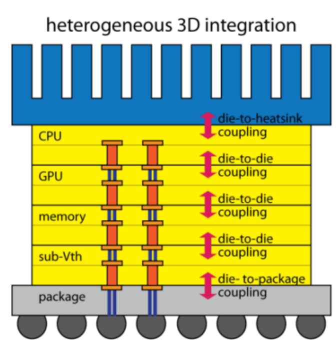

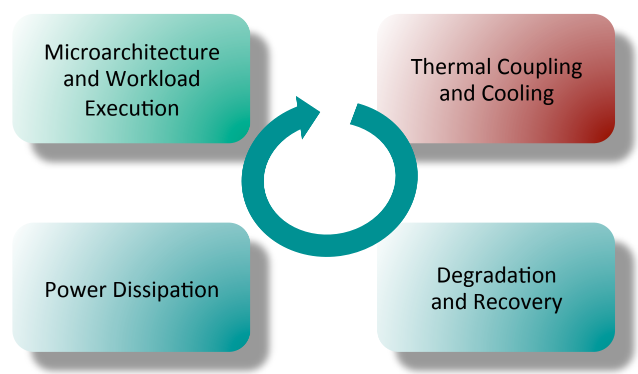

Thank You

Questions?