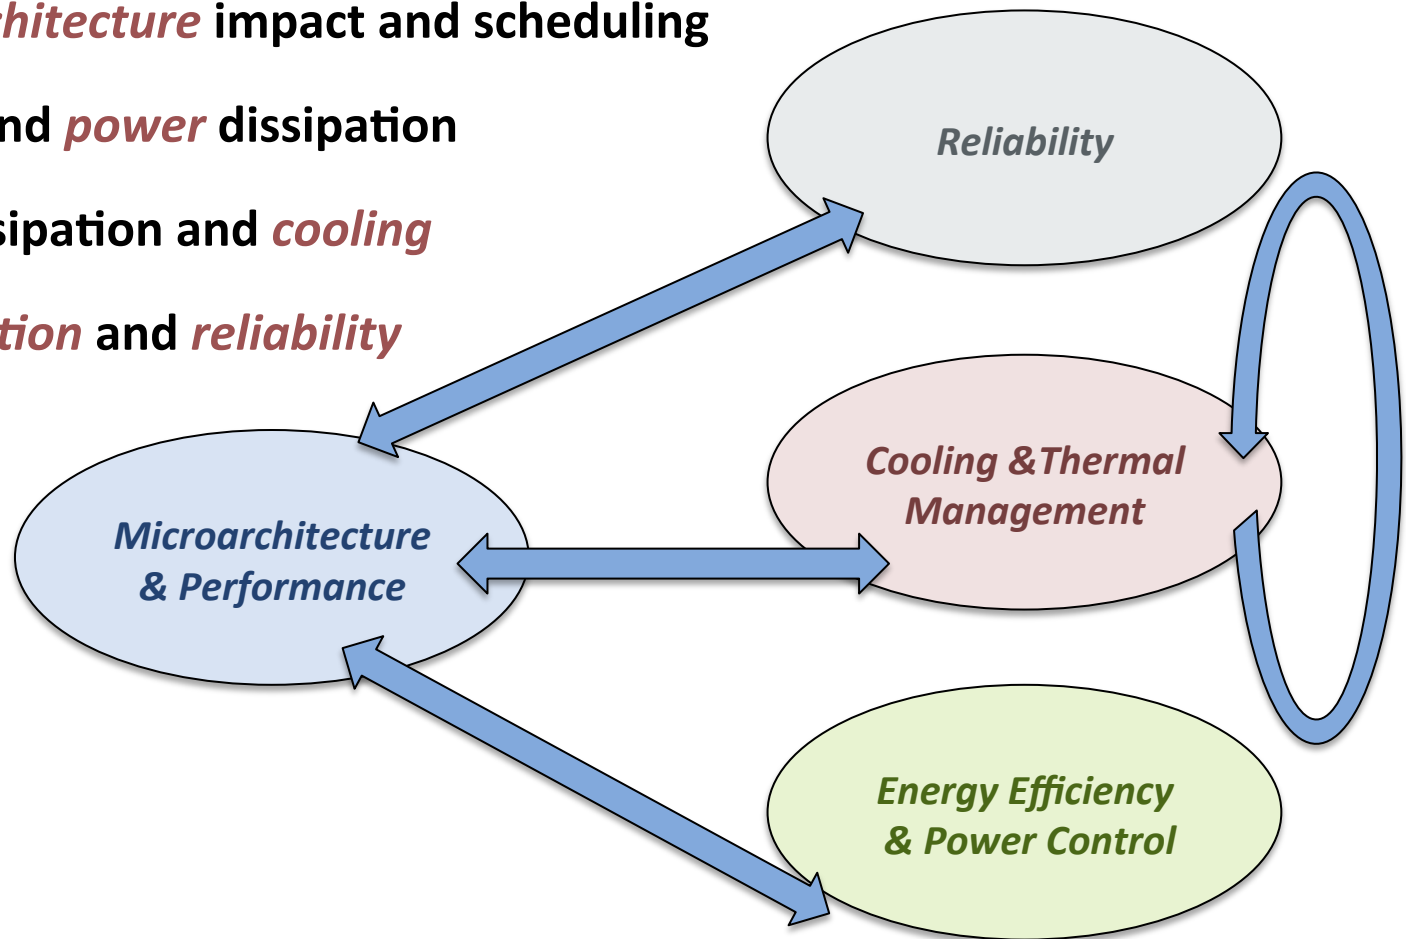


# ***Coordinated Architecture – Multi-Physics Modeling and Reliability Analysis***

**William Song, Saibal Mukhopadhyay, and Sudhakar Yalamanchili**  
*School of Electrical and Computer Engineering*  
*Georgia Institute of Technology*

# Microarchitecture and Physics Interactions

- **Multi-physics modeling:**
  - **Workload** dynamics
  - **Microarchitecture** impact and scheduling
  - **Energy** and **power** dissipation
  - **Heat** dissipation and **cooling**
  - **Degradation** and **reliability**

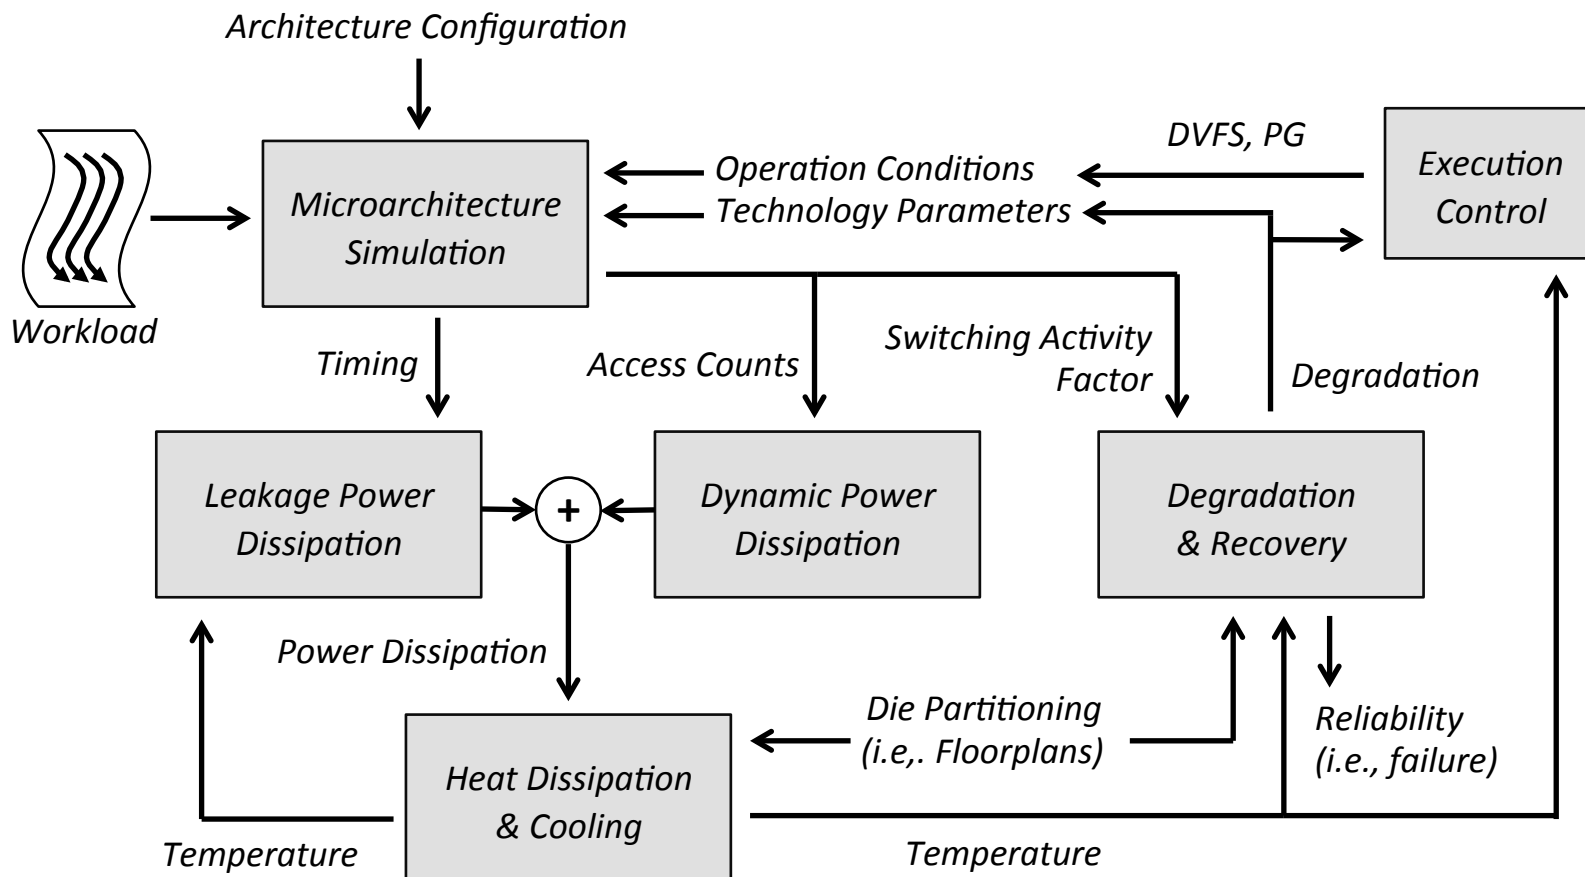


# Architecture-Level Modeling

- ***Single phenomenon Modeling*** (conventional):
  - **Power modeling:**
    - Circuit-level breakdown (i.e., functional units)
    - Measurement-based regression models
    - **Thermal impacts? Process variation?**
  - **Thermal modeling:**
    - Package-level analysis (i.e., differential equations)
    - Source-layer floorplanning
    - **Temperature-power interactions? Performance impacts?**
  - **Reliability modeling:**
    - Device-level characterization (i.e., NBTI)
    - **Turbo boosting/core? Race/idle computing?**
    - ***+ Dynamic control techniques:***
      - *DVFS*
      - *Power gating*
      - *Thread migration*

# Coordinated Architecture Modeling

- Abstract representation of *Microarchitecture-Physics Interactions*:

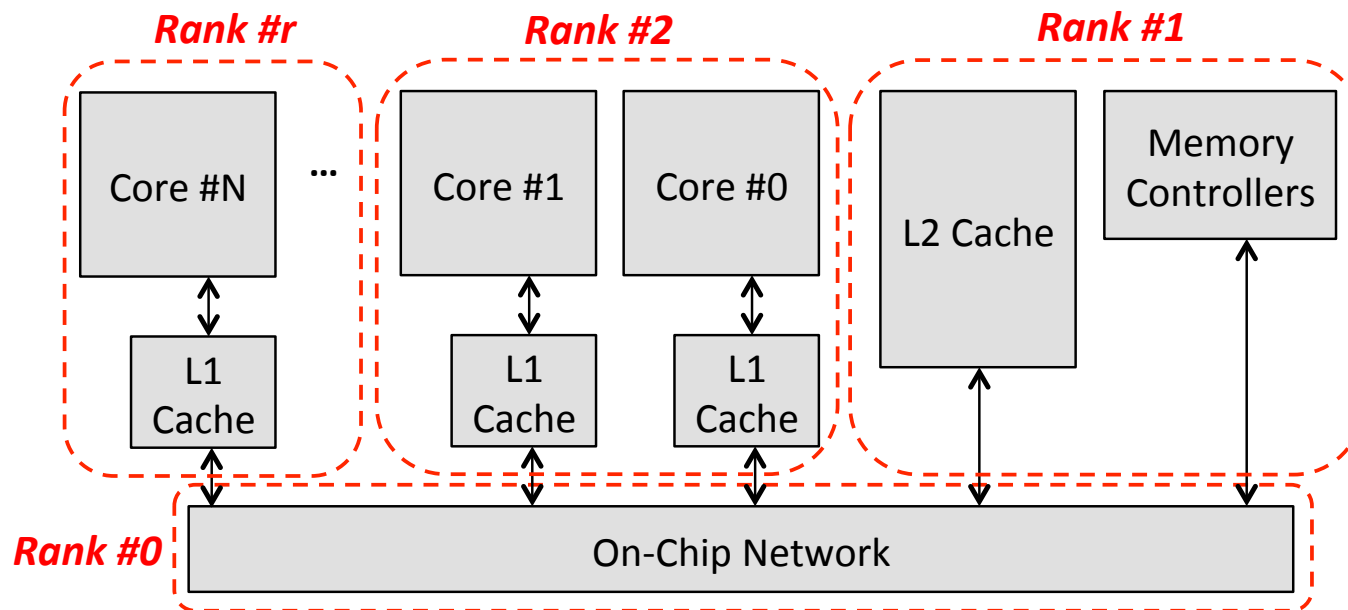


# Proposed Architecture Simulation Framework

- **Energy Introspector (EI):**
  - **Compatibility:**
    - Integration of various C/C++ models already (or being) developed, validated by different research groups.
  - **Usability:**
    - Model-independent interface and handy user functions.
  - **Flexibility:**
    - Adaptation to different microarchitecture, technologies, and designs.
  - **Coordination:**
    - Interactions between integrated models.
  - **Scalability:**
    - Large core-count processor modeling.

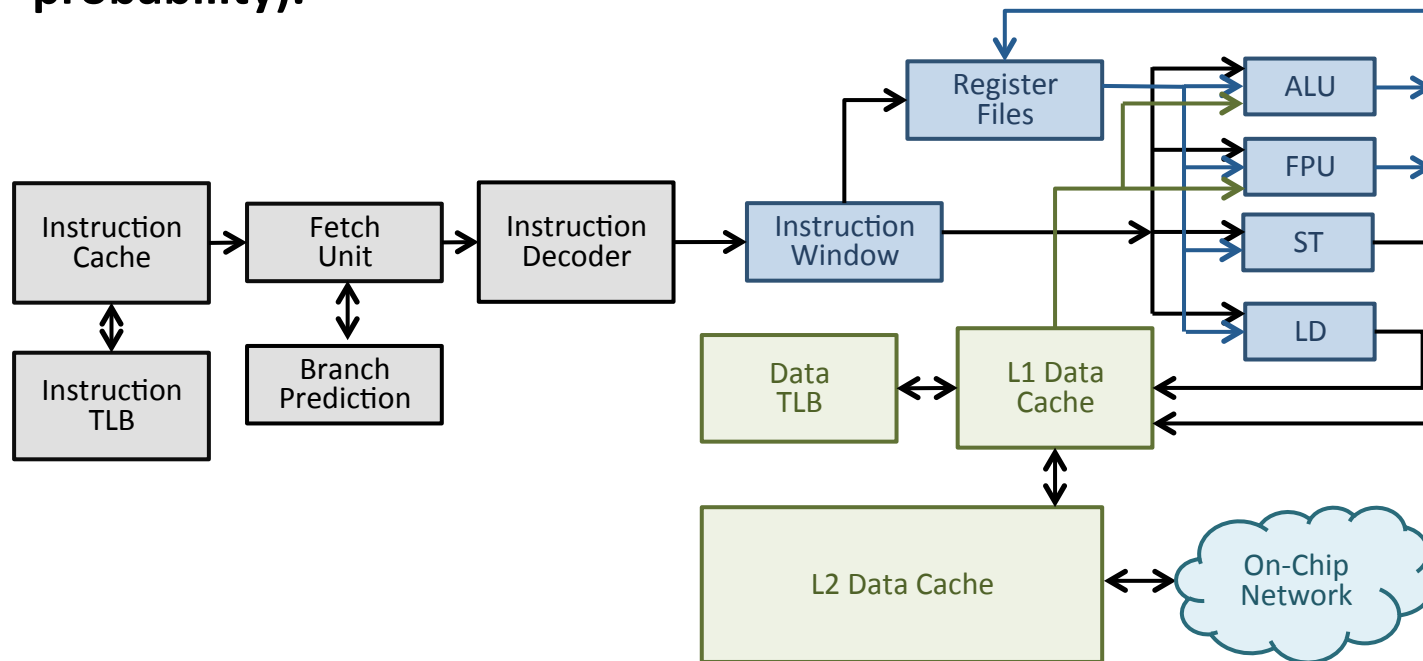
# Microarchitecture Modeling

- Scalable simulation framework:
  - *Parallel, scalable architecture simulation* via MPI implementations.
  - Structural Simulation Toolkit (SST) from Sandia National Labs
  - Manifold from Georgia Tech



# Microarchitecture Breakdown

- **Microarchitecture characterization:**
  - **Statistics (i.e., performance counts) are collected at functional architecture blocks (*sources in EI term*).**
  - **Collected statistics are used in the EI to characterize switching activities and compute energy (and power) and reliability (i.e., failure probability).**



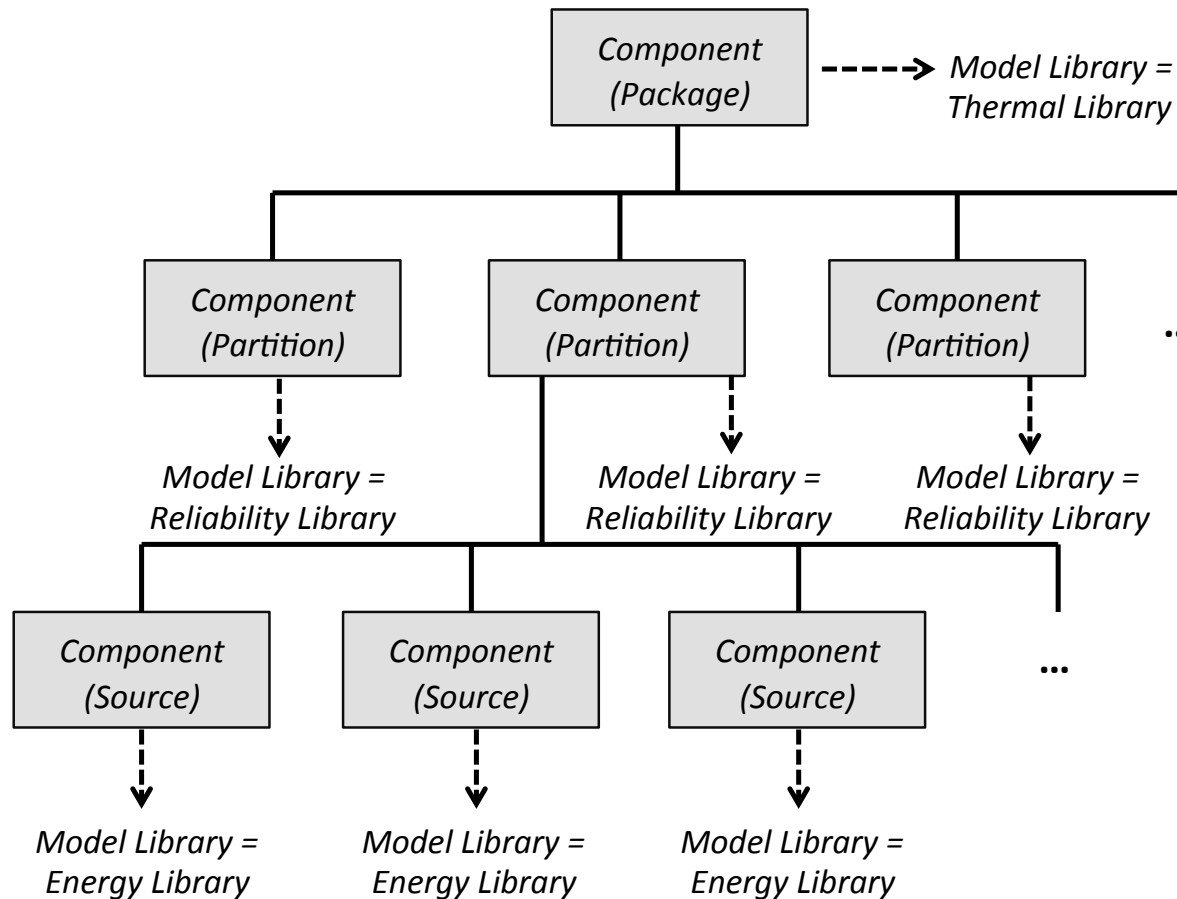
# Calculation of Physical Phenomena

- Physics characterization is via conventional models.
- BUT, the calculations are based on *transient data* dynamically updated via *runtime simulation* (vs conventional trace-driven or offline modeling).
- Coordination problem:
  - *Switching activities* are characterized via microarchitecture simulation.
  - *Energy (or power)* is calculated at basic functional blocks (i.e., circuit-level or block-level granularity).
  - *Temperature* is computed at the package-level.
  - *Reliability* may be characterized at block or floorplan levels.



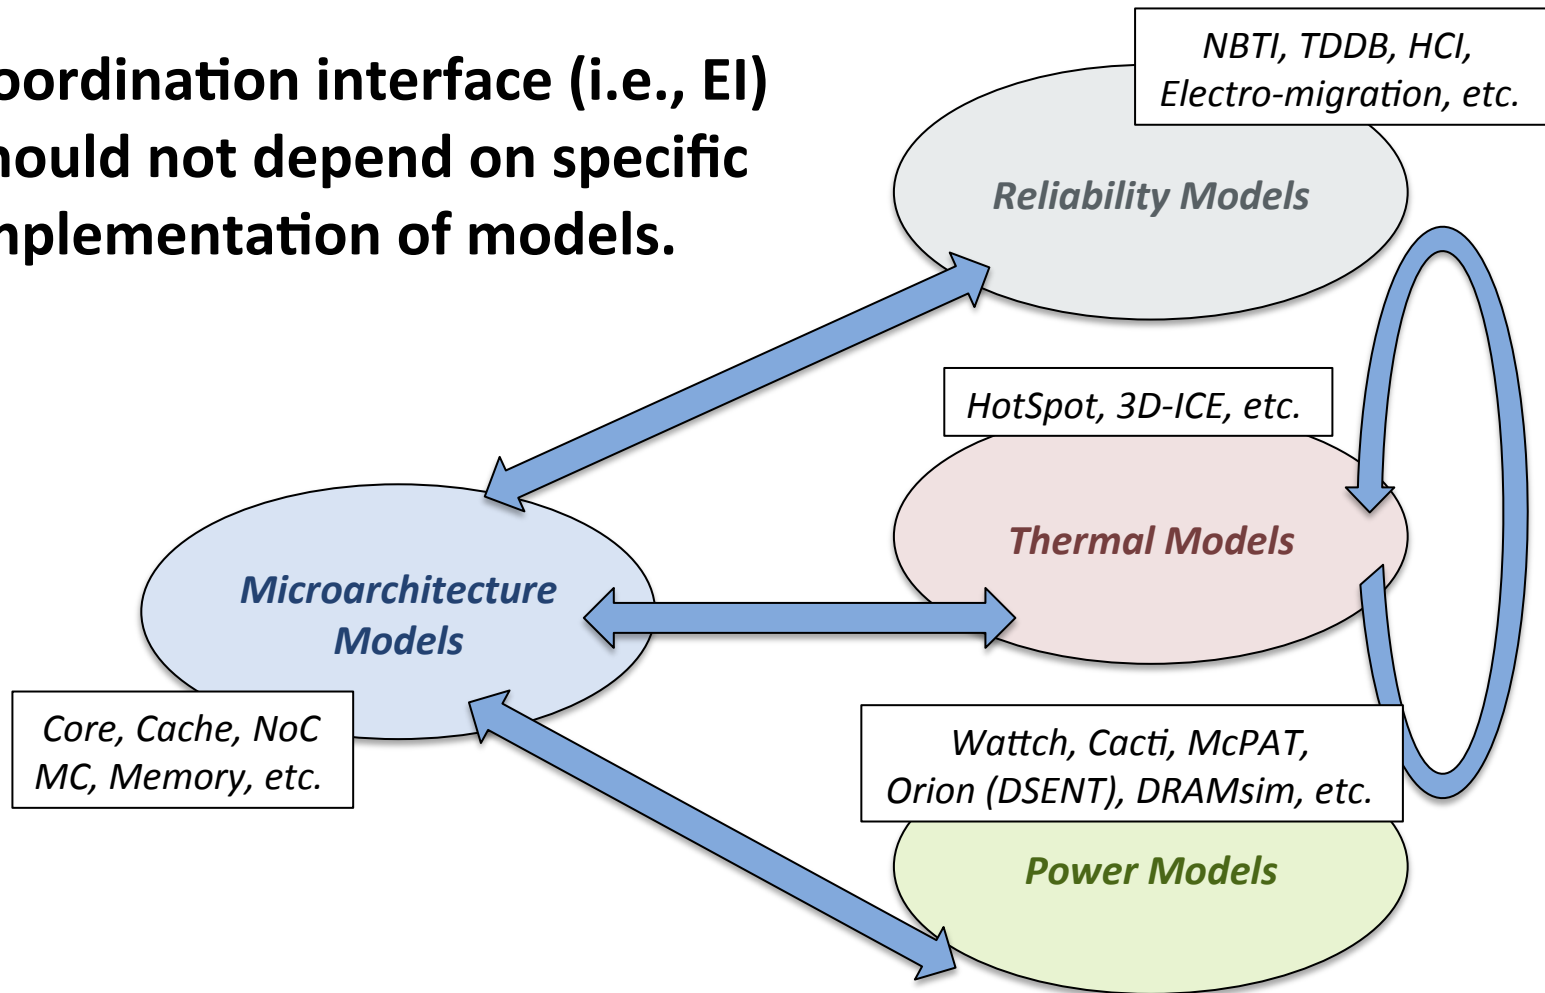
# Abstract Representation of Processor Hierarchy

- Processor is modeled as *a hierarchical tree of pseudo components* that represents processor components at different levels.



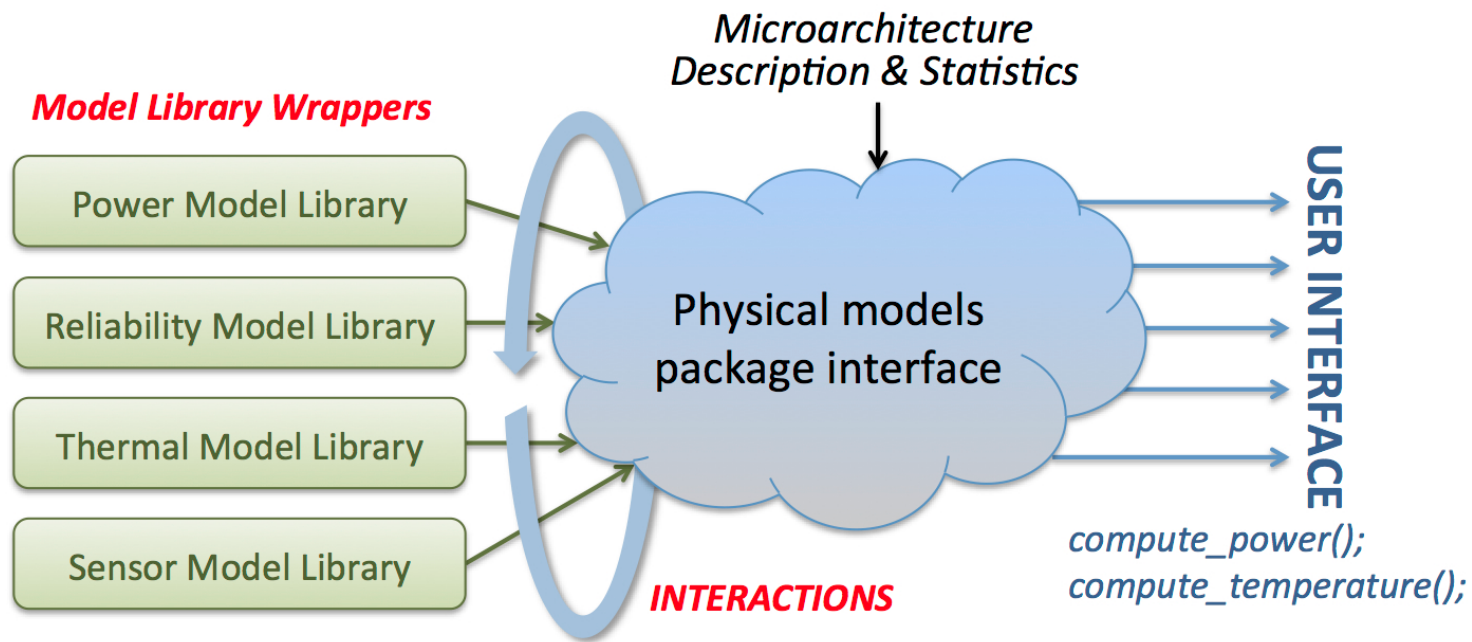
# Revisiting Coordinated Architecture Modeling

- A number of combinations/options to select from each modeling pool.
- Coordination interface (i.e., EI) should not depend on specific implementation of models.

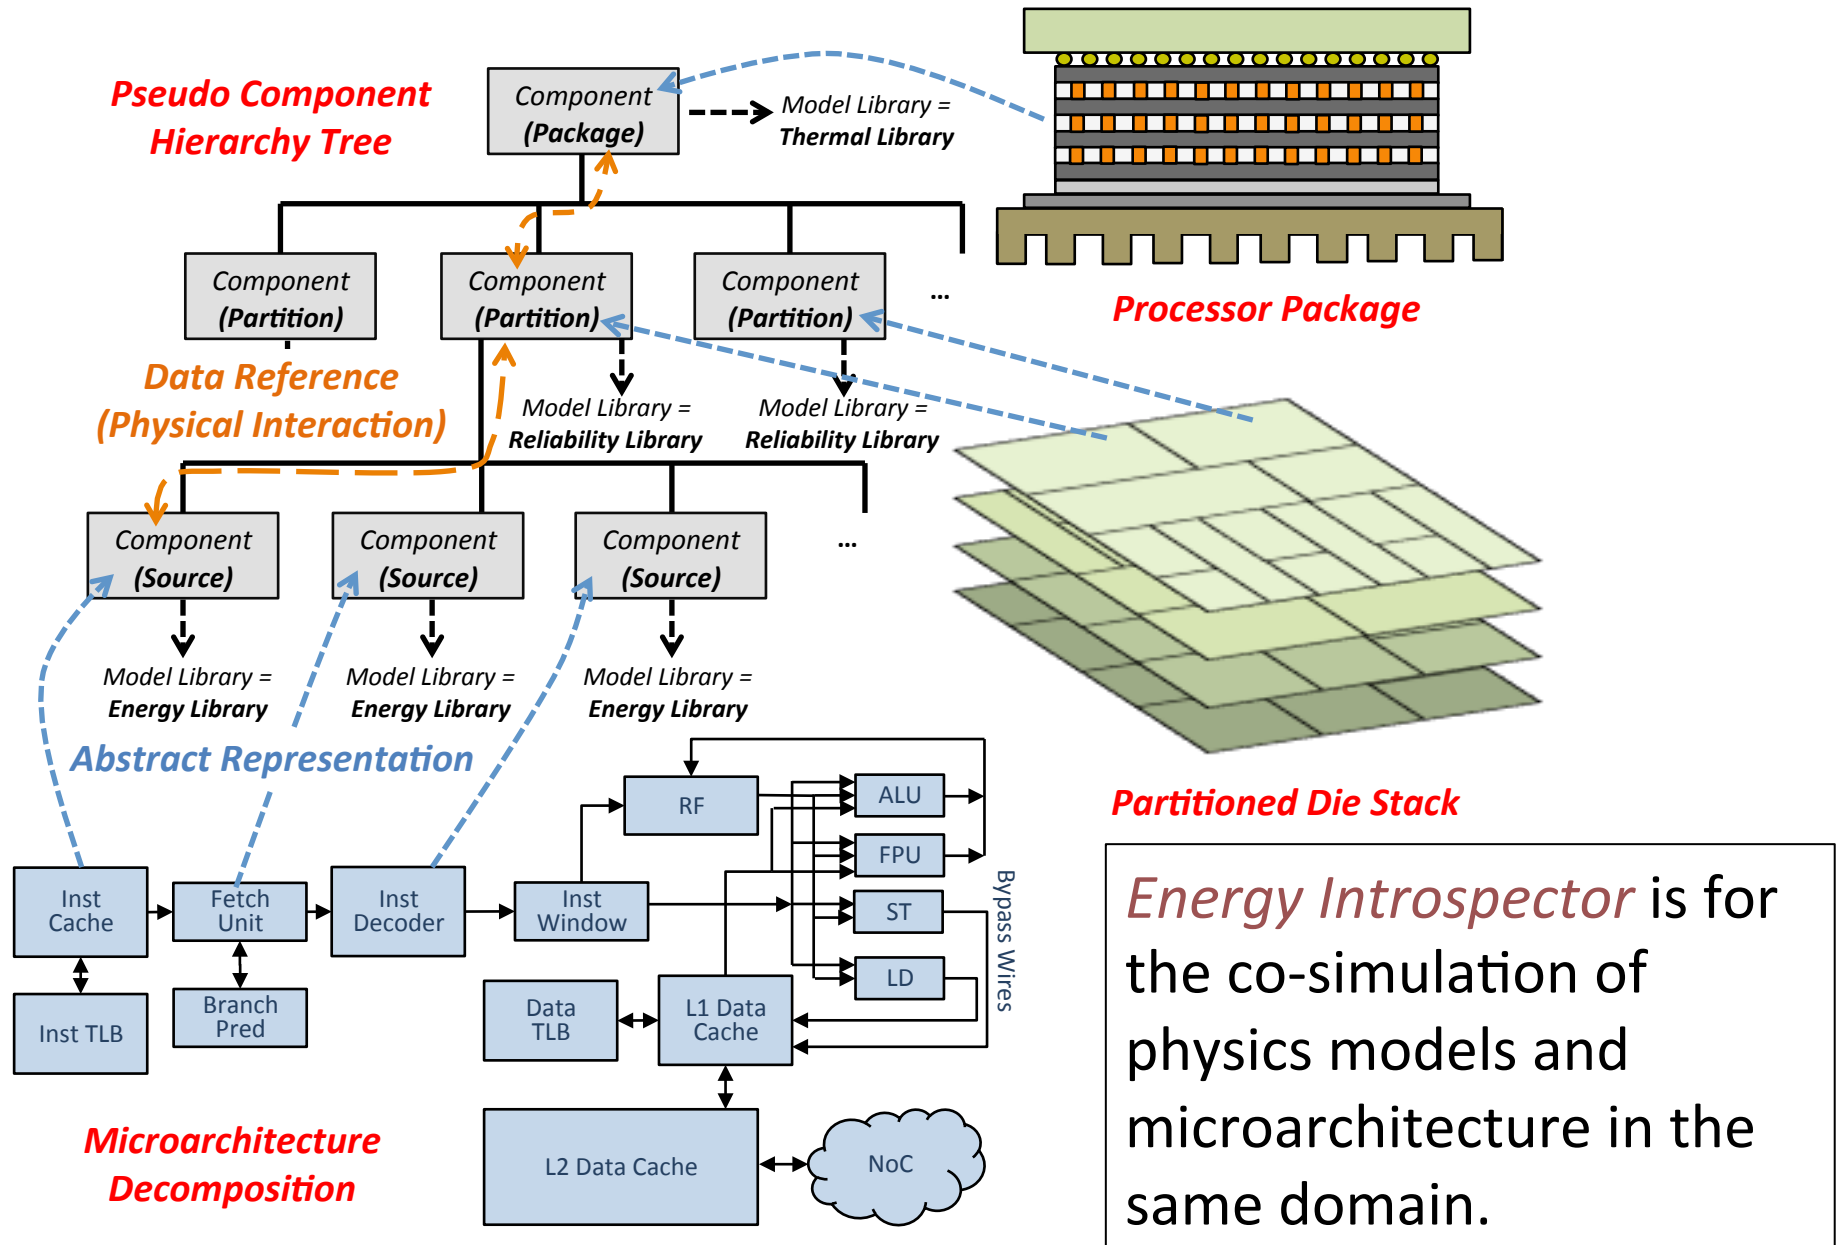


# Library Models

- Similar models are grouped into the same *library*.
  - C++ subclassing, virtual functions, etc.
- The Interface does not handle input parameters.
  - The wrapper class handles input parsing via gcc libconfig.



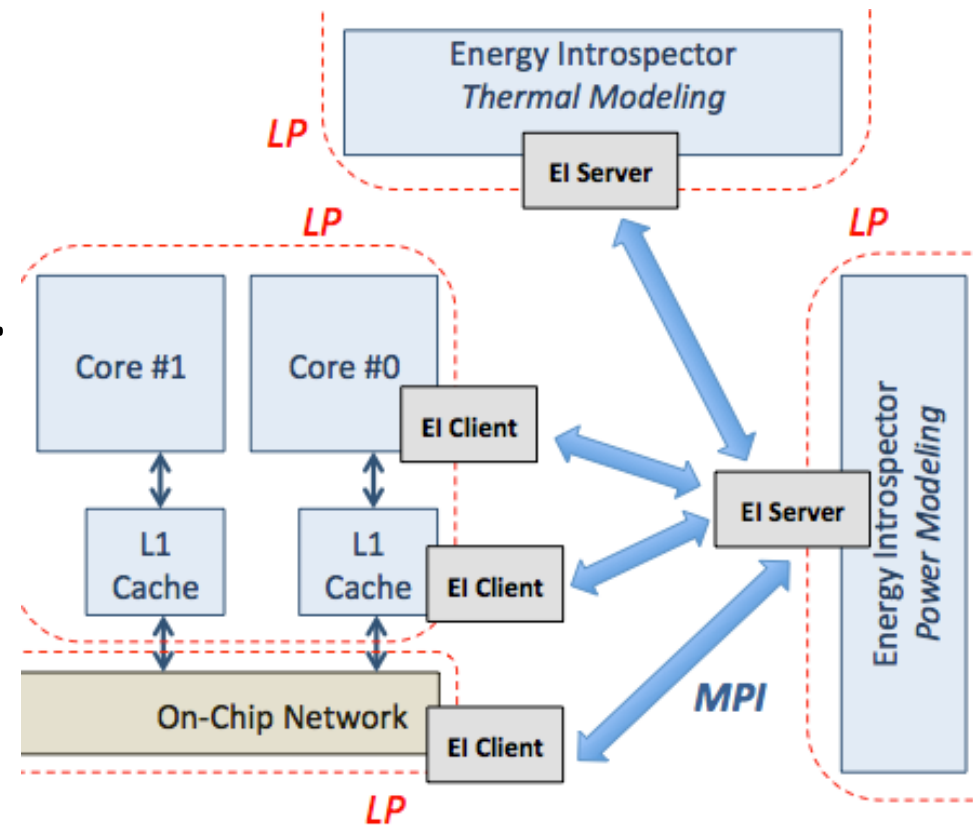
# Overview of Energy Introspector



*Energy Introspector* is for the co-simulation of physics models and microarchitecture in the same domain.

# MPI-based Multi-Process Simulation

- *Single-threaded* simulation of architecture simulation is practically limited to *a few cores*.
- *MPI*-based implementation enables scalable simulation.
- Architecture simulators and Energy Introspector run on *multiple MPI ranks*.
- Energy Introspector spawns *server threads* that wait for client node requests.



***Application of Coordinated Architecture Simulation  
to Reliability Analysis***

# Race-to-Idle Execution and Reliability

- **Race**
  - The execution of a core is boosted for a short period of time to increase performance.
  - Performance improvement is traded with increased power and heat dissipation and accelerated degradation.
- **Idle**
  - Idle period following the race mitigates increased temperature and failure rate.
  - Leakage energy is saved by turning off cores.
- Reliability is believed to be worse for race-to-idle than normal executions?

# Simulation Setup

- 64-core Asymmetric Chip Multiprocessor:

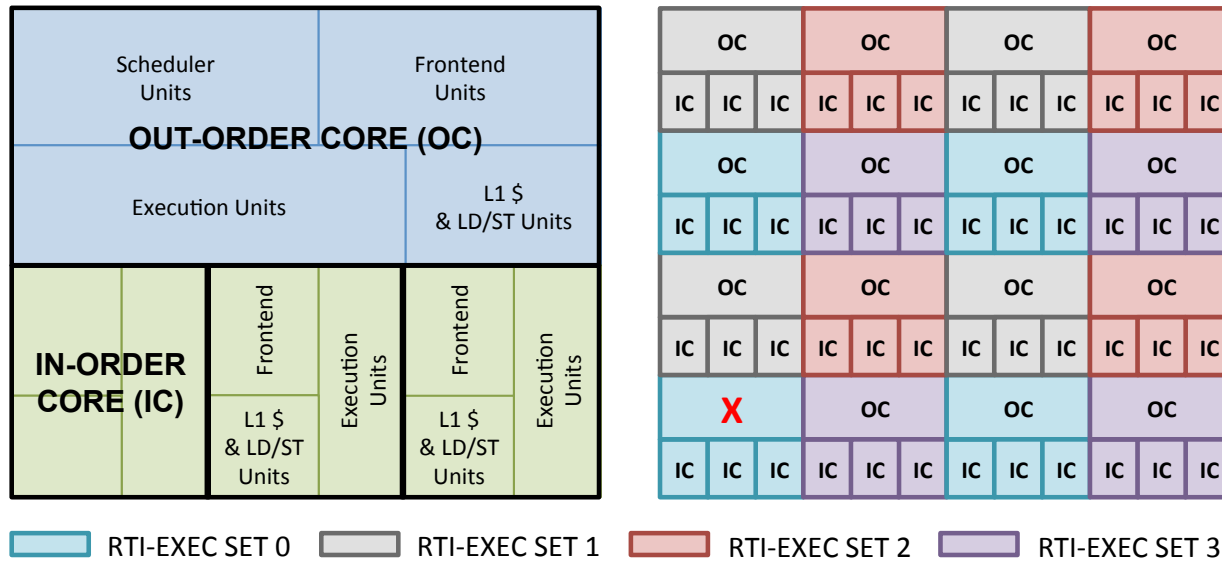


TABLE I. EXPERIMENT CONFIGURATION FOR COORDINATE ARCHITECTURE SIMULATION

Configuration	Description	
Simulator	Manifold 64-core simulation [2]	
Benchmarks	Multi-programmed execution of SPEC2006 suite	
Cores	Out-of-order	In-order
Core counts	16	48
Issue width	4	1
Reorder buffer size	128	N/A
L1 Cache	4-way assoc, 64-byte line, 32KB size	
L2 Cache	8-way assoc, 64-byte line, 256KB size, private L2	
Voltage/frequency levels	0.8V/2.0GHz for NE, 1.2V/4.0GHz for RTI	
Feature size	16nm technology projection to ITRS guideline	



# Failure Probability Modeling

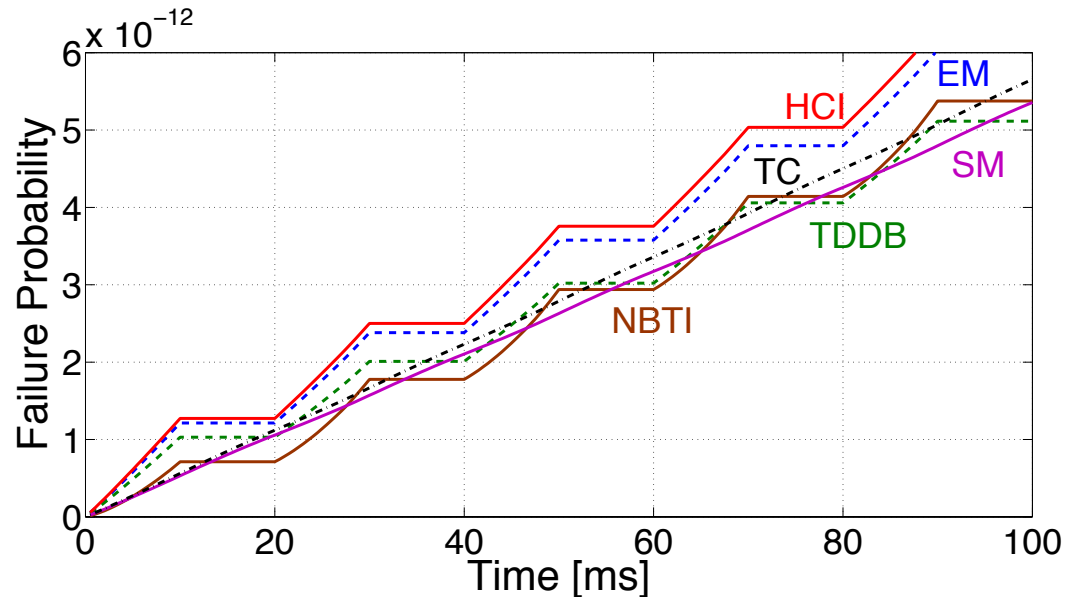

- **Failure rates** computed with 1) NBTI, 2) TDDDB, 3) HCI, 4) electro-migration, 5) thermal cycling, 6) stress migration.
  - **Exponential distributions** with 6 failure mechanisms are used to calculate runtime failure probability.

$$P_{total}(t) = 1 - P_0 \prod_{i=1}^n \prod_{r \in Risks} \left( \begin{array}{c} 1 - P_r(t_i - t_{i-1}) \\ |C_i(T_i, F_i, V_i, A_i, G_i) \end{array} \right)$$

- Each exponential curve is **fitted to be equally likely** at the target condition (i.e., 3.0GHz, 65°C operation, etc.)
- Operation conditions (i.e., temperature, frequency, etc.) are dynamically adjusted via coordinated architecture simulation.


# Transient Failure Probability of Race-to-Idle

- *Periodic race-to-idle* compared with *continuous normal execution*.
- Breakdown of each failure mechanism.
- Dominance of failure mechanisms depends on operation condition.

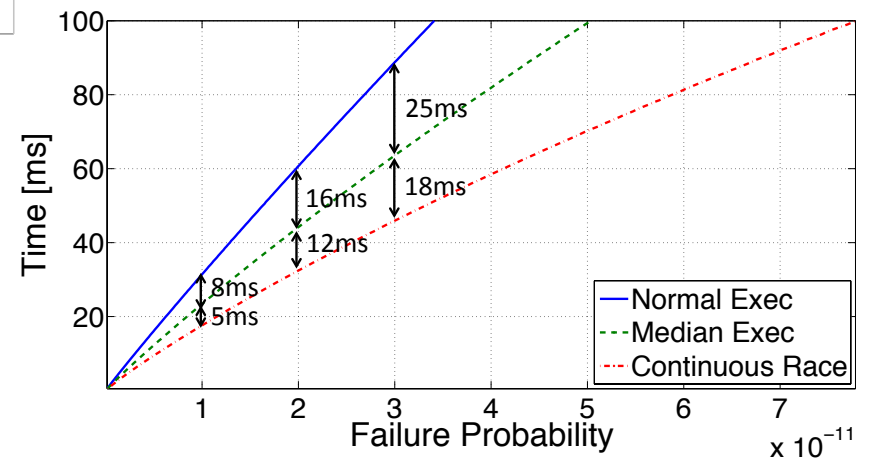


# Race / Idle Time Balancing

- Finding a good ratio of race/idle periods:
  - Race-to-idle execution is controlled such that the failure probability is equalized to pre-generated failure probability of normal execution.



- *Race-to-idle is better with:*
  - *Complex cores than simple cores.*
  - *Power-consuming workloads (computational vs memory)*



# Summary

- Architecture modeling and analysis have become more complicated and need coordinated infrastructure for future designs.

