Thermal-Effective Clustered Microarchitectures

P. Chaparro, J. González and A. González

Intel Labs - UPC

Motivation

- Removing heat is expensive
- Design point is set for worst case temperatures
  - Expensive thermal solution guarantees peak performance
    - Usually temperatures are lower
  - A localized hotspot may…
    - trigger global emergency mechanisms: But it could be avoided by focusing only on that hotspot
    - not be detected: Sensors covering wider areas
- Clustered architectures give new opportunities for temperature reduction
  - Peak temperature 33%
  - Average temperature 12%

Overview

- Introduction
- Processor Architecture
- Simulation Infrastructure
- Thermal Analysis of Clustered Architectures
- Cluster Hopping
- Conclusions

Introduction

- Clustering opens new opportunities for temperature reduction
  - Distribution of resources
    - Activity distribution
  - Hopping schemes
  - Layout flexibility
    - Trade off unit location vs. wire delay
  - Resource grouping into clusters
    - Voltage and clock domains
    - Leakage control
    - V_{dd} gating

Processor Architecture

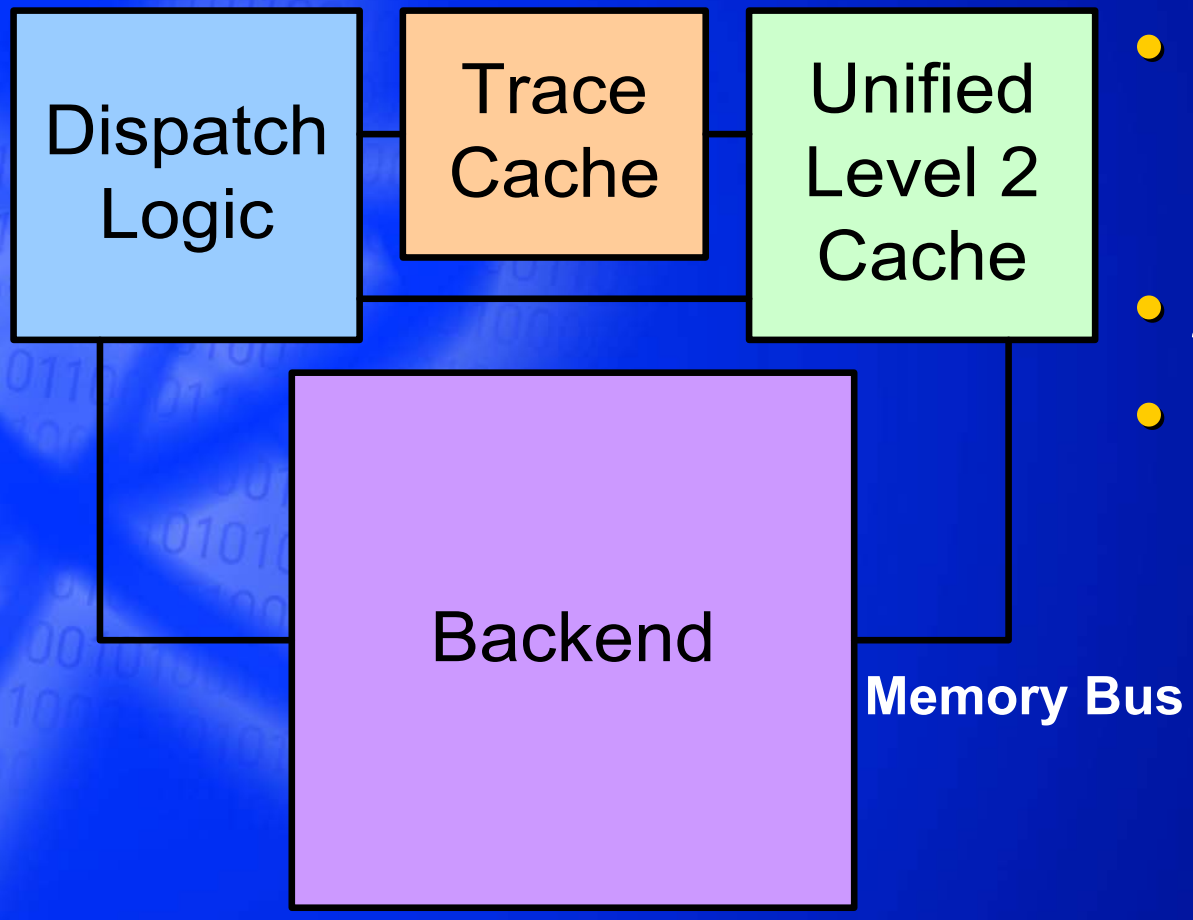

- Large frontend
  - 32Kuop trace cache
  - dispatch 8 uops/cycle
- 2MB L2 cache
- Highly OOO
  - 80-entry issue queue
  - 384-entry MOB
  - 4 int + 3 fp + 4 ld/st
  - 544+544 physical regs
  - 64KB, 2-way L1

Processor Architecture

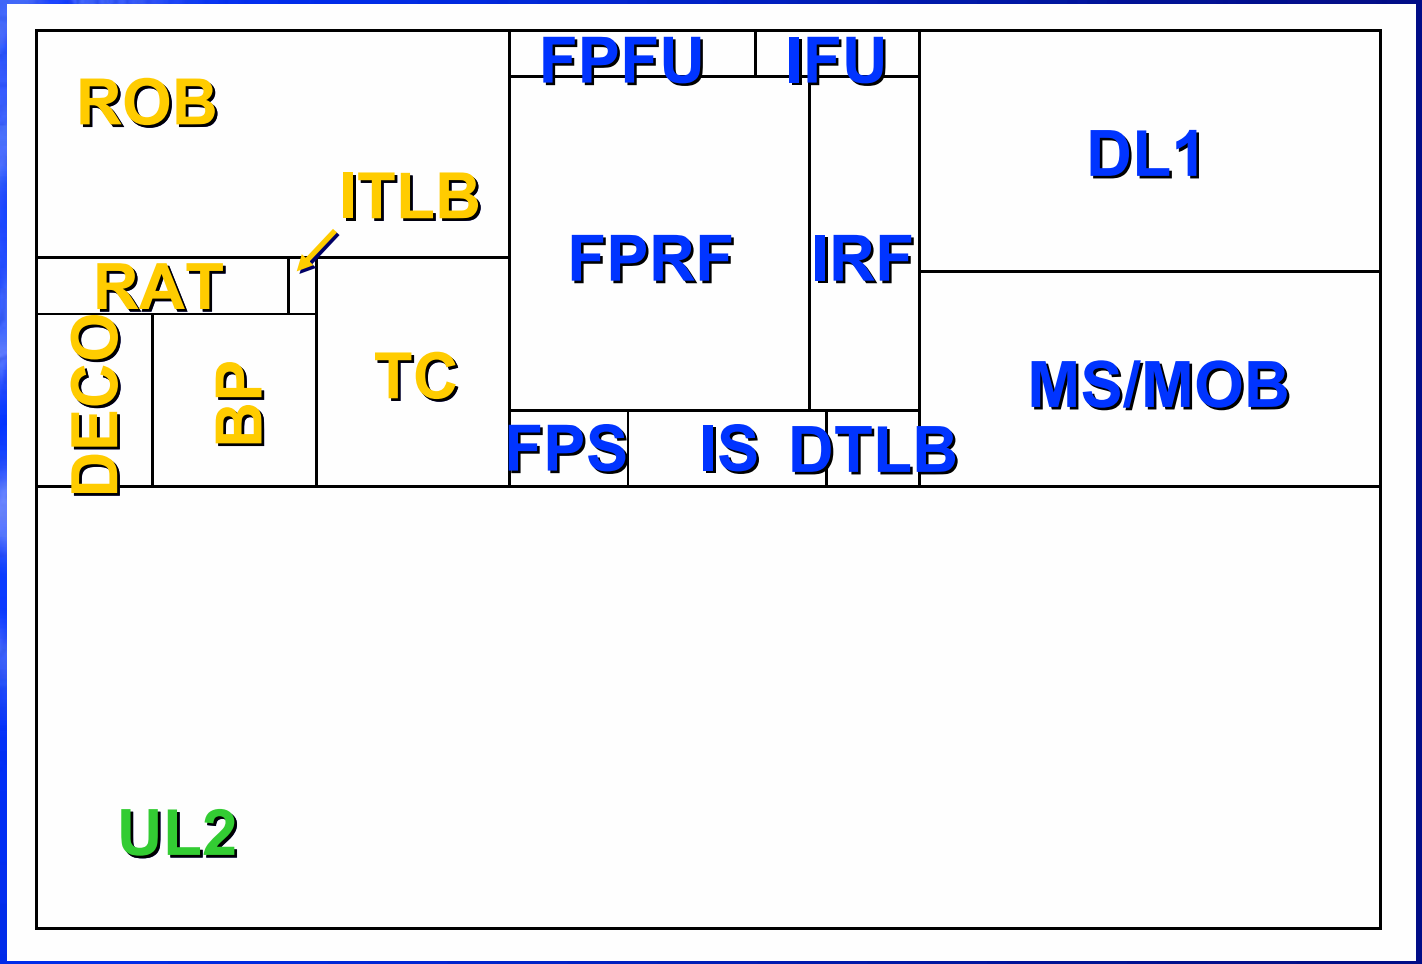

Processor Architecture

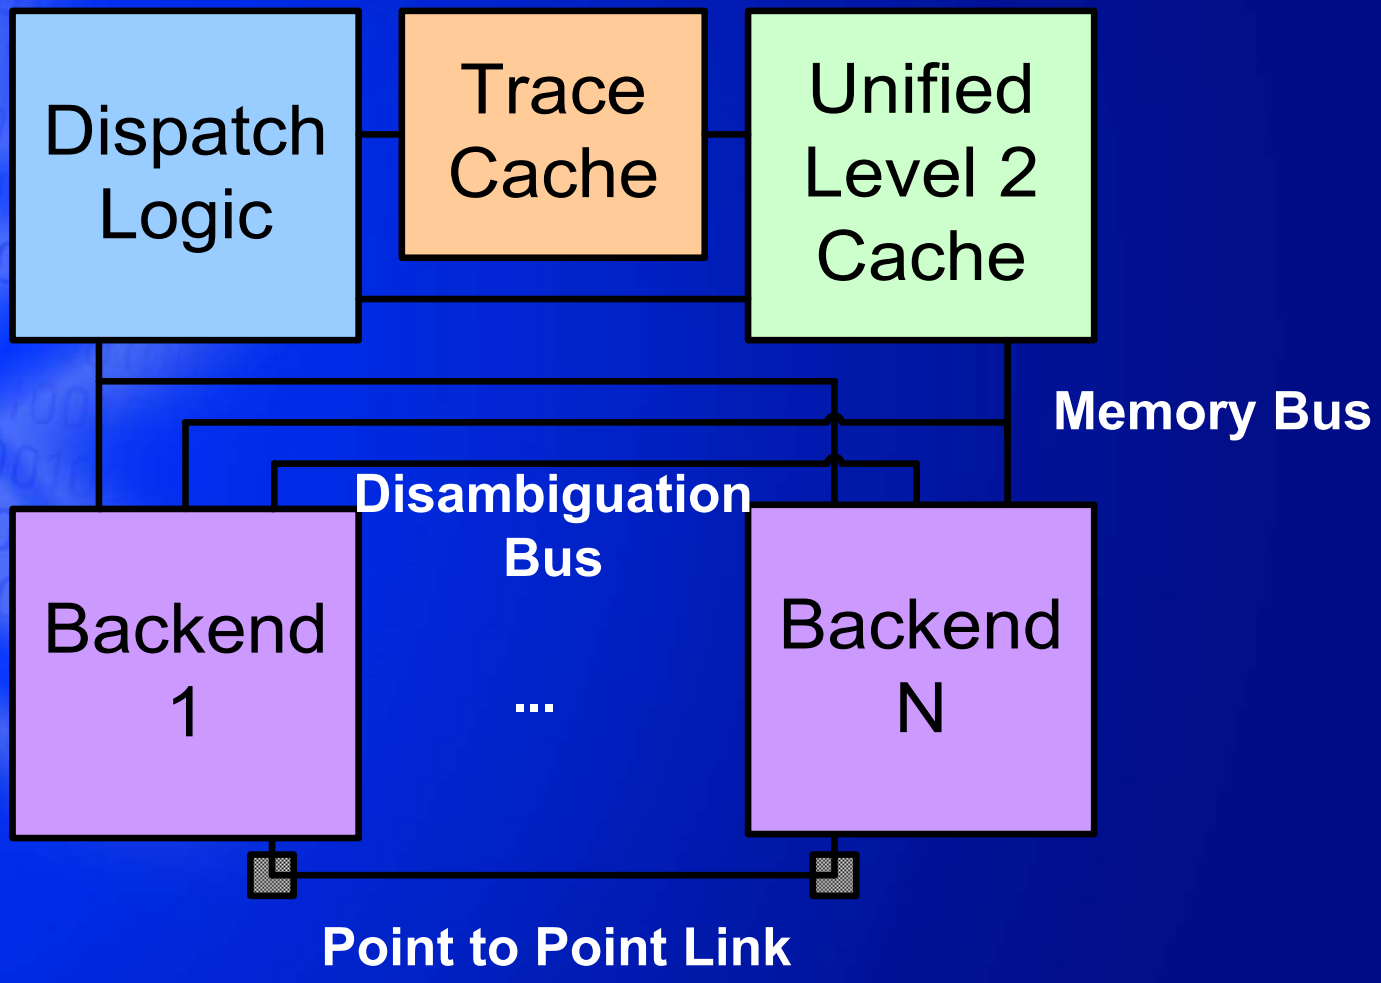

Processor Architecture

Bicluster

Each cluster has half the resources of the original monolithic backend

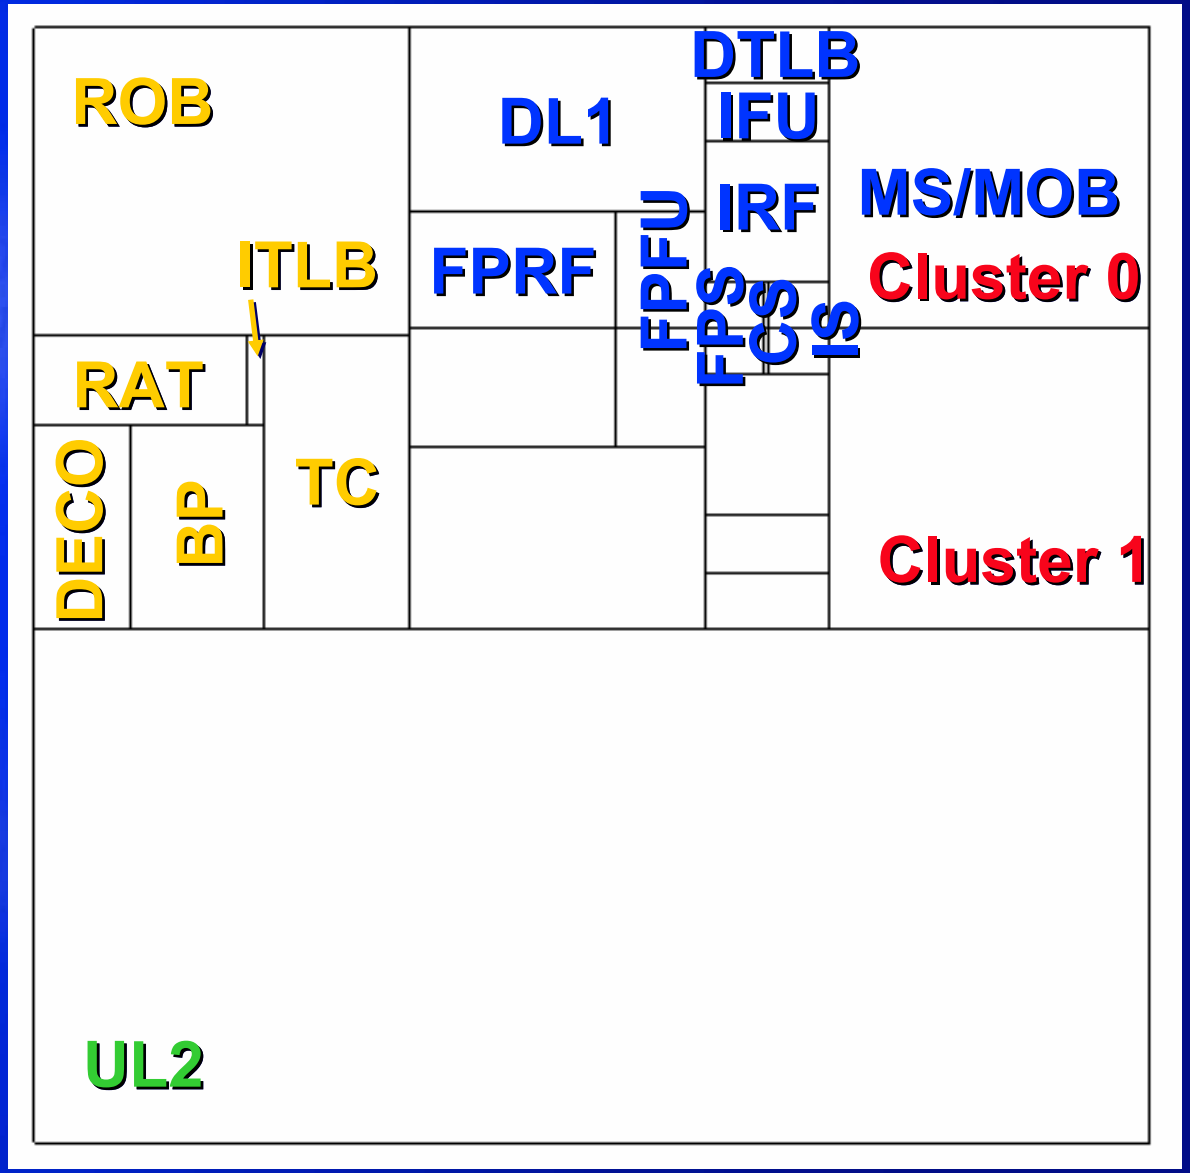

Processor Architecture

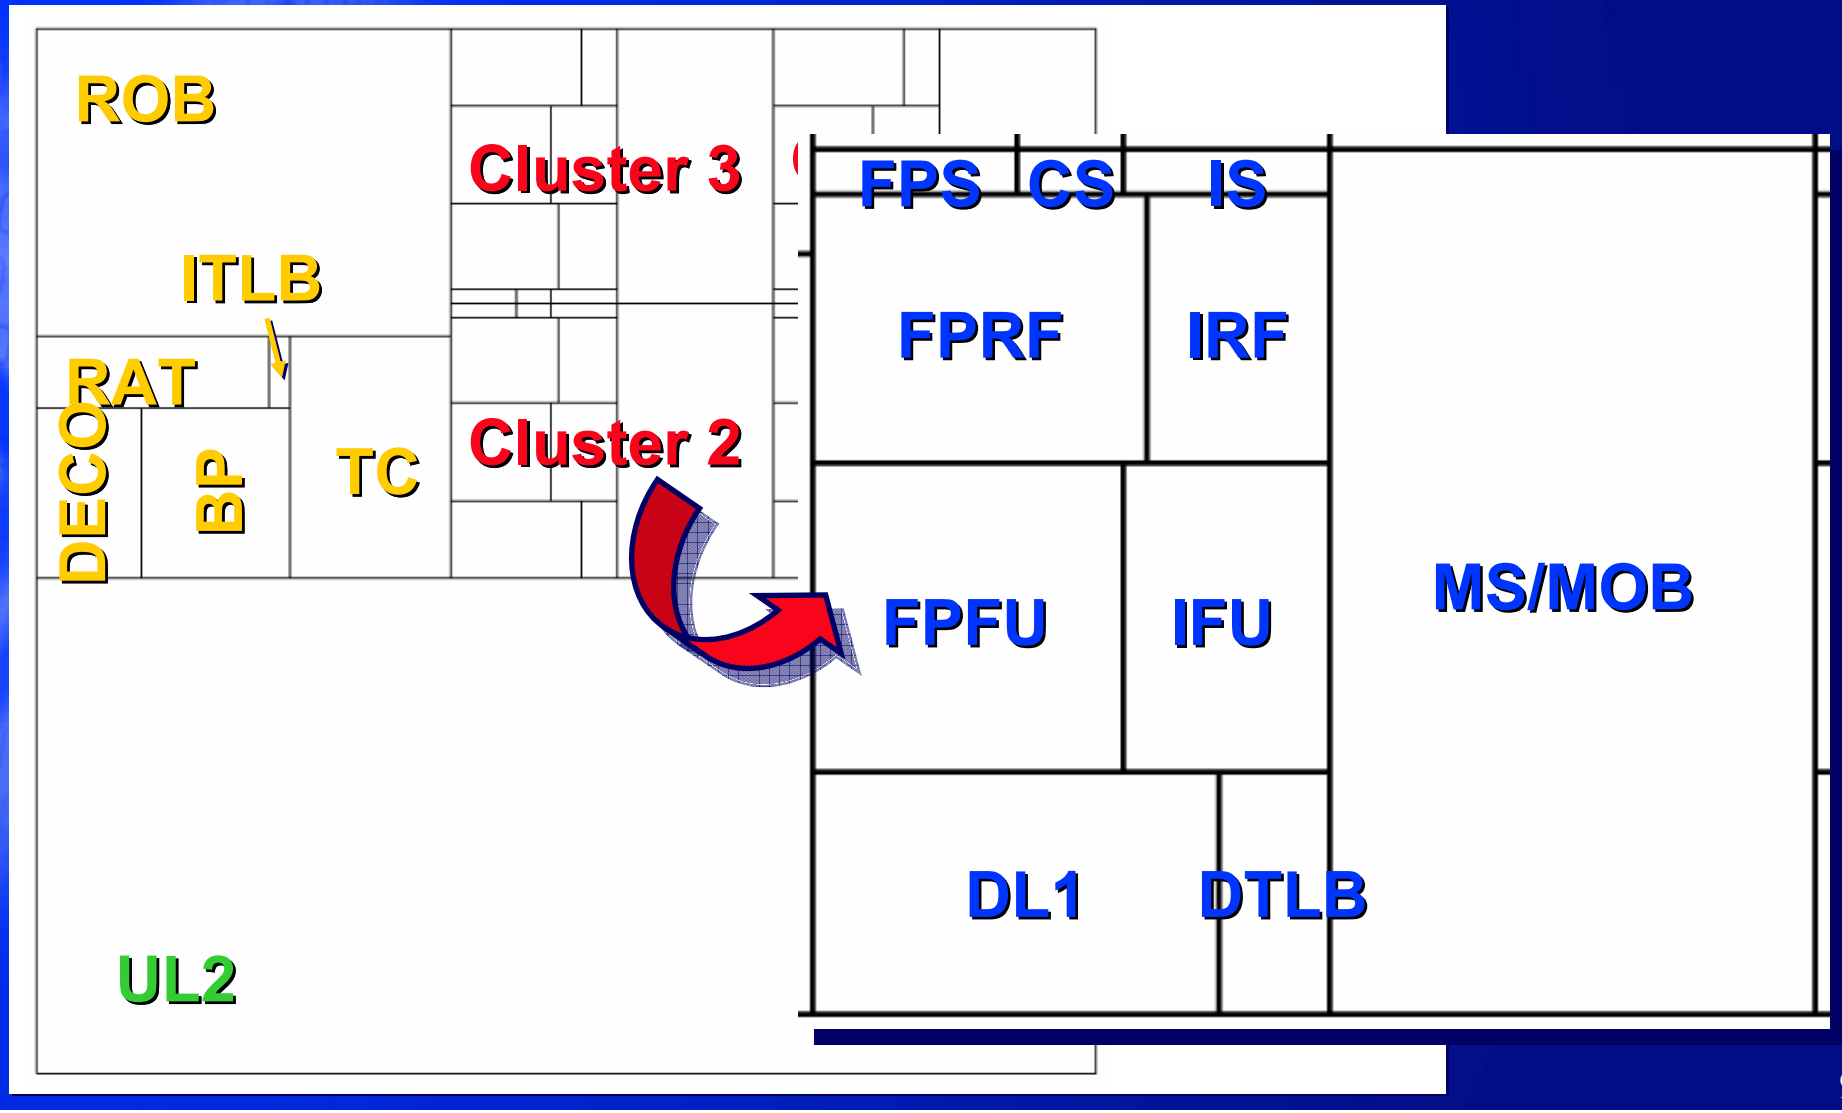

Simulation Infrastructure

- Computes dynamically the temperature of selected functional blocks (emulates thermal sensors)
- Integrated in a microarchitectural simulator

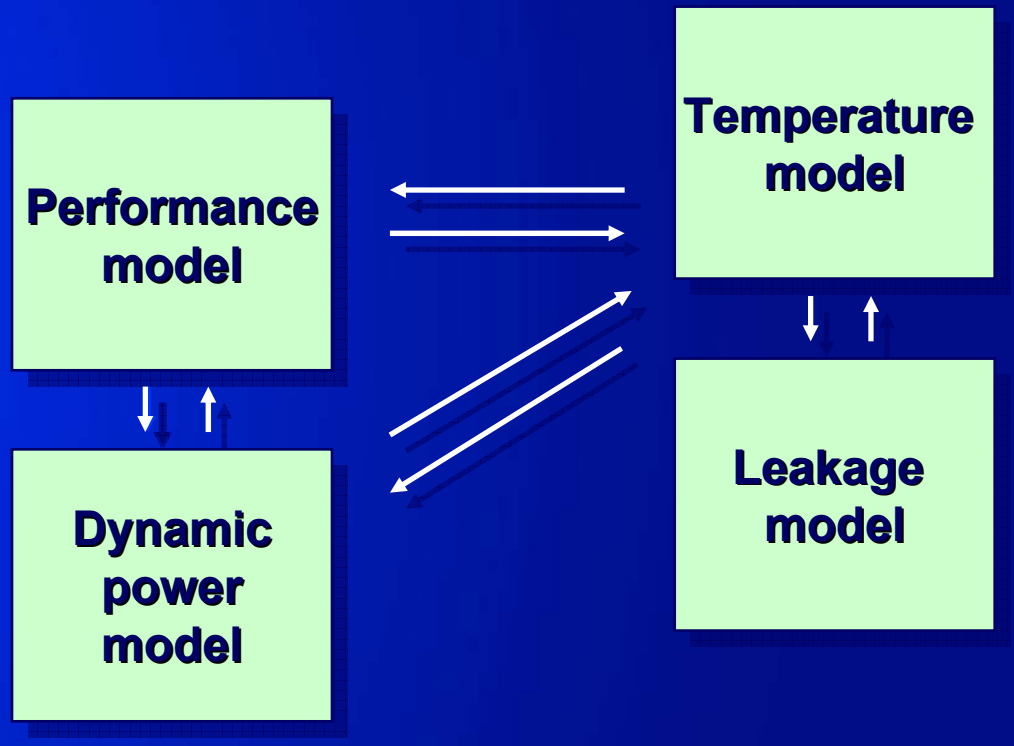

Simulation Infrastructure

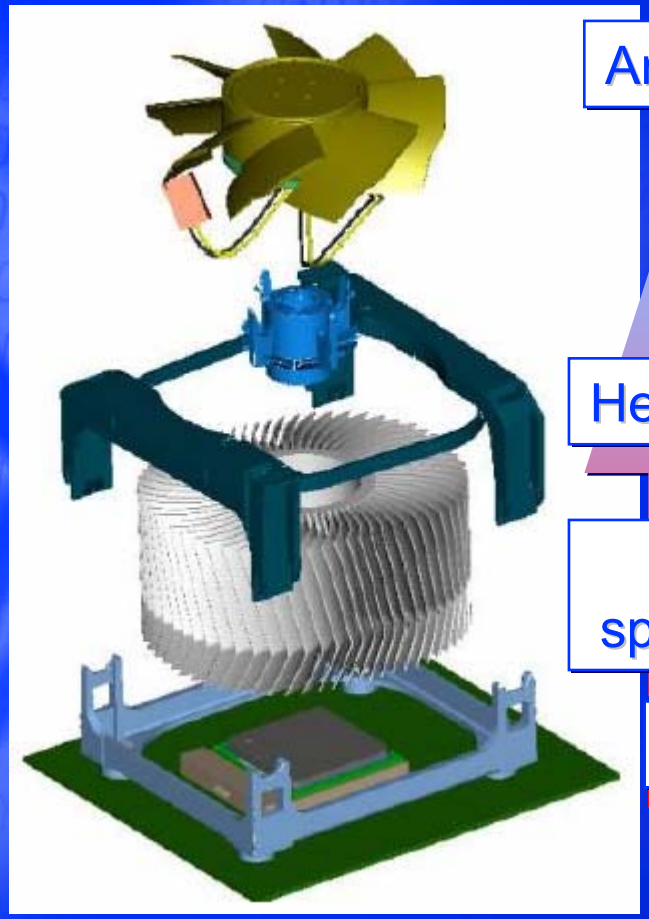

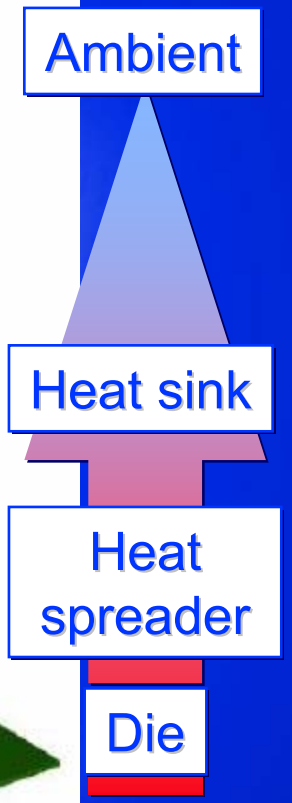

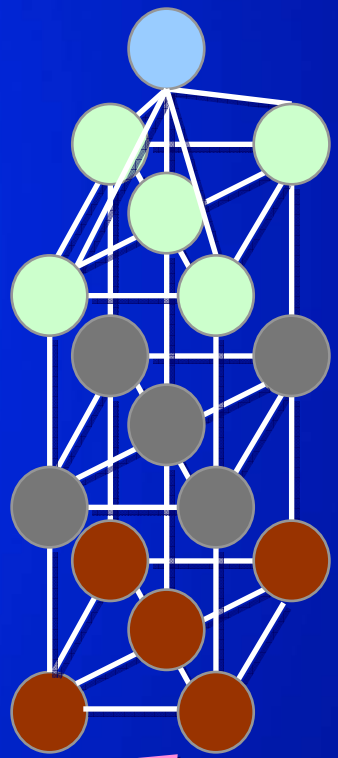

Electrical	Thermal
Voltage (*V*)	Temperature (*K*)
Current (*A*)	Power (*W*)
Resistance (*V / A = Ω*)	Resistance (*K / W*)
Capacity (*J / V = F*)	Capacity (*J / K*)
Time constant τ = R · C (*s*)	

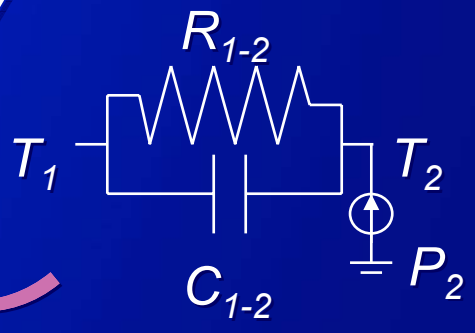

R-C pairs

Thermal Analysis of Clustered Architectures

- Temperature metrics
  - AbsMax
    - Maximum sensed temperature
  - Average
    - Average temperature of the chip area over time
  - AverageMax
    - Average temperature over time of the maximum sensed temperature

Thermal Analysis of Clustered Architectures

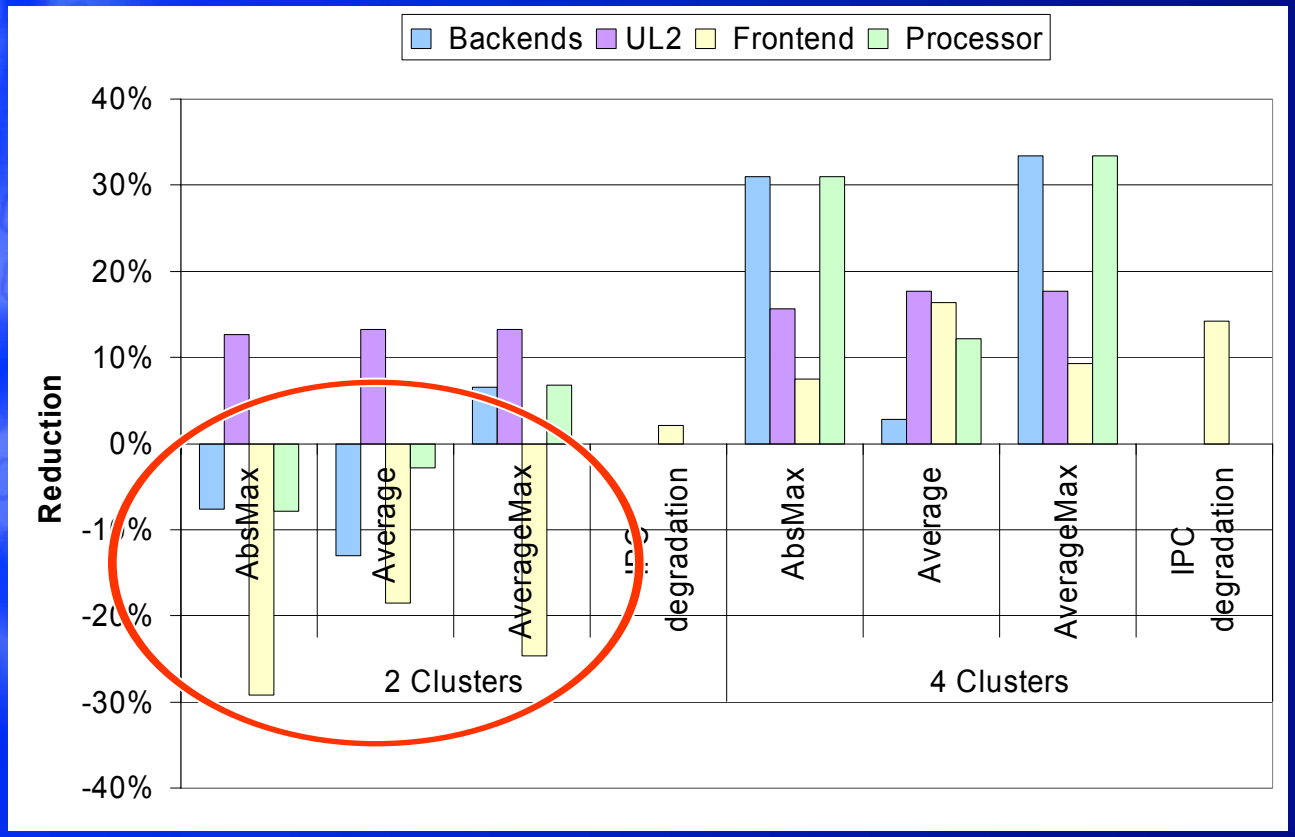

Average temperature reduction for 16 SPEC

Cluster Hopping

- Based on activity migration [Heo, ISLPED 03]
  - V_{dd} gate a subset of clusters
  - Rotate clusters to spread activity along time
  - Gated clusters cannot provide any register value
    - Before gating cluster must be emptied
  - Cache/DTLB contents are lost
  - Proactive and/or reactive behavior
    - Proactive: Per interval basis
    - Reactive: On thermal events

Cluster Hopping

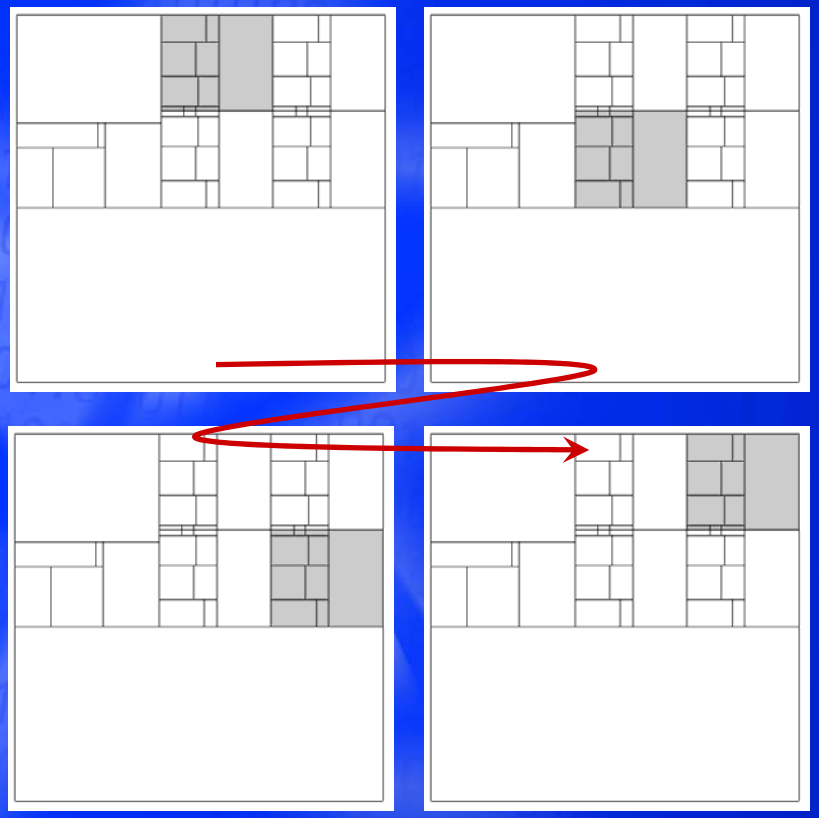

HOP-3

HOP-2

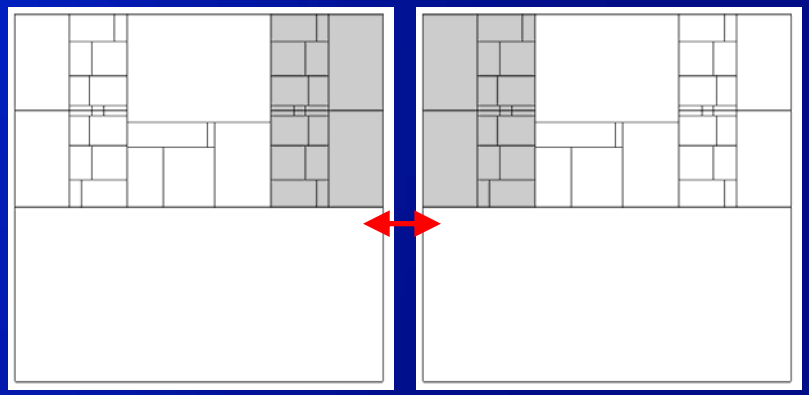

Cluster Hopping

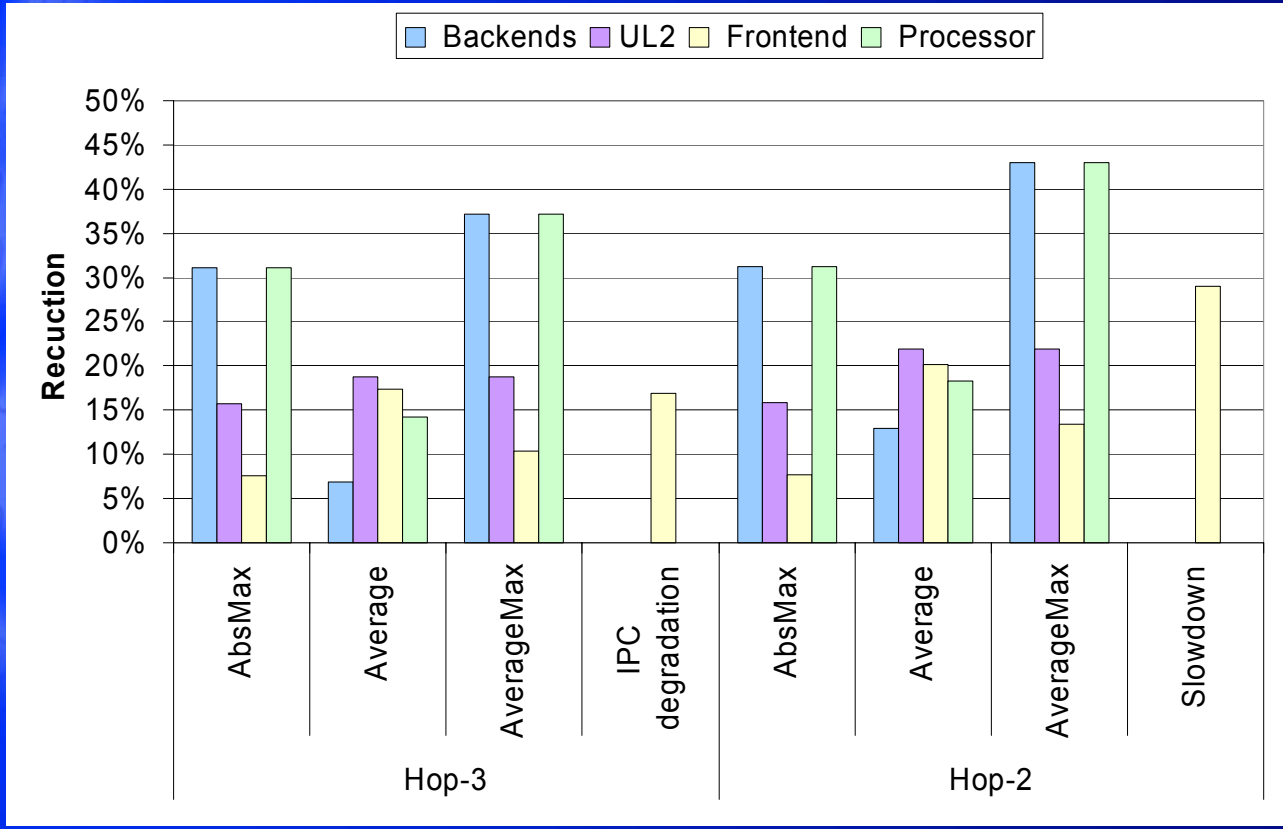

Conclusions

- The analyzed bi-cluster architecture is increasing temperature: Clustering must be applied smartly
- The quad-cluster architecture analyzed is effective reducing temperature:
  - Reduces processor peak temperature 33%
  - Reduces 12% average temperature
  - IPC penalty of 14%
  - Ignored other benefits of clustering for this study
- Improving the quad-cluster architecture with a hopping scheme (*HOP-3*):
  - Peak temperature is reduced 37%
  - Average temperature of the processor 14%
  - Extra penalty of 3%