POLITECNICO DI MILANO

October 11rh, 2012

SoC 2012 – *Tampere, FINLAND*





Thermal/Performance Trade-off in Network-on-Chip Architectures

Davide Zoni, Simone Corbetta and **William Fornaciari**

fornacia@elet.polimi.it

http://home.dei.polimi.it/fornacia

Politecnico di Milano
Dipartimento di Elettronica e Informazione
Via Ponzio 34/5, 20133, Milano, ITALY

Outline

- **Introduction**
  - From single-core to multi-core
  - ITRS projections
- **Key observations**
  - The ring-based organization
  - The proposed methodology
- **Results**
  - Simulation framework
- **Future perspective**

What's next

- **Introduction**
  - From single-core to multi-core
  - ITRS projections
- **Key observations**
  - The ring-based organization
  - The proposed methodology
- **Results**
  - Simulation framework
- **Future perspective**

Power trend: The 2004 inflection point

- **From single-core to multi-core processors**

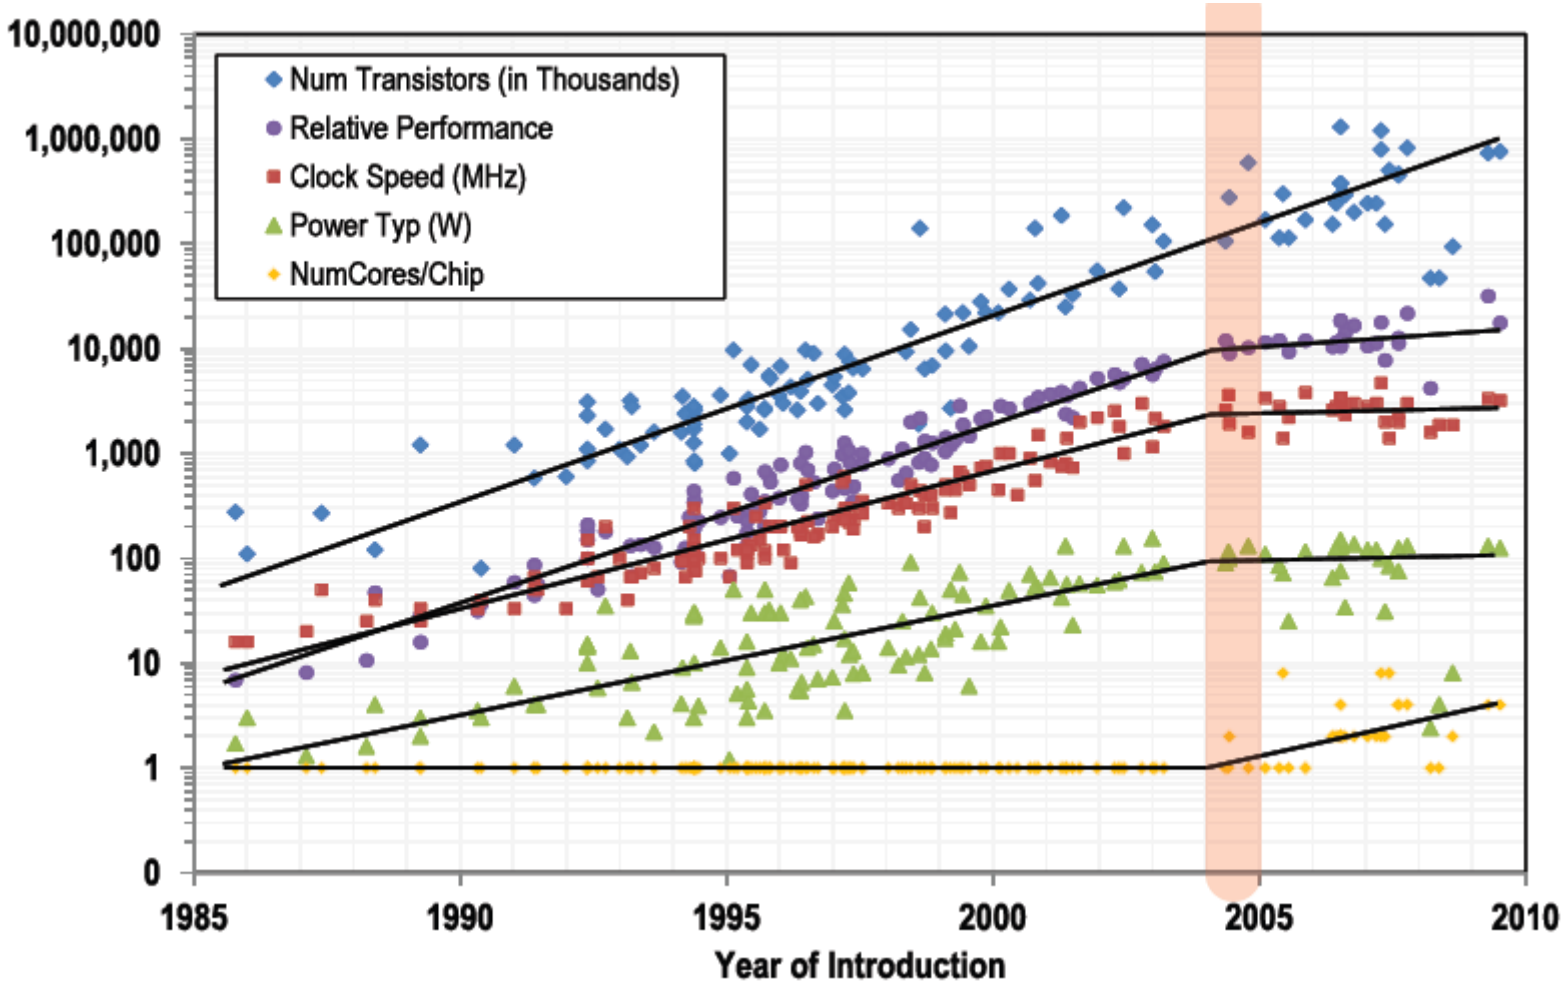

Source: B. Falsafi. "Reliability in the Dark Silicon Era". IOLTS2011 Keynote, July 2011.

Multi-core architectures era

- **Number of cores is expected to increase**
  - Die area is not (constant to 260mm^2)
  - Higher MPU power density

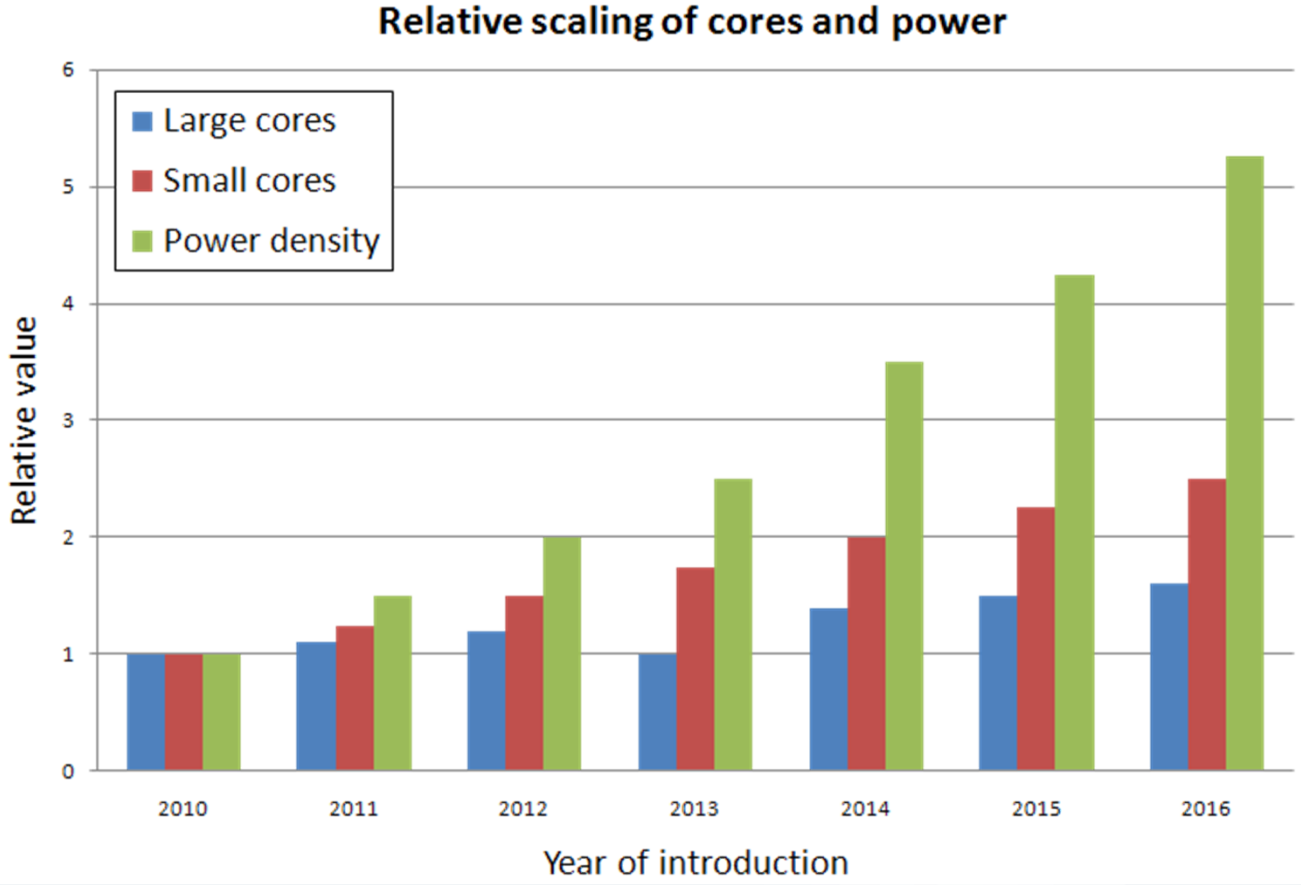

What's next

- Introduction
  - From single-core to multi-core
  - ITRS projections
- **Key observations**
  - The ring-based organization
  - The proposed methodology
- Results
  - Simulation framework
- Future perspective

Thermal management

- **Design-time analysis and optimization**
  - Provides thermal map analysis of a given scenario
  - Compile-directed optimization requires full knowledge of application
  - Do not consider dynamic operating conditions
- **Run-time approaches**
  - Adaptive to application requirements
  - Expensive
  - Often, requires monitor/sense/actuate
- **Hybrid approach**
  - Design-time suggested and run-time exploited

Temperature distribution in MPSoCs

- **Questions**
  - How multi-programmed workload impacts temperature profile in a 2D-mesh architecture?
  - Which is the impact of different workloads?
  - How to reduce hot-spot?

- **Preliminary analysis can be carried out**
  - Using simulation flows, providing temperature estimation beforehand
  - Formal models can be used to optimize temperature profile and asses validity

Temperature distribution in MPSoCs (cont'd)

- **Power consumption of different applications**
  - Low variability, due to microarchitecture design (in-order pipeline)
- **Implications**
  - Hot-spot is located at the center of the chip!

Core #	Placement		Instructions			Power
	Row	*Col*	*Int*	*FP*	*Mem*	[W]
1	0	0	12.8%	6.6%	80.6%	7.132
2	0	1	34.6%	36.5%	28.9%	6.810
3	0	2	64.4%	30.4%	5.3%	7.185
4	0	3	40.8%	29.6%	29.6%	6.717
5	1	0	40.8%	29.6%	29.6%	6.956
6	1	1	64.4%	30.4%	5.3%	6.630
7	1	2	40.8%	29.6%	29.6%	6.849
8	1	3	67.7%	32.3%	0.0%	6.632
9	2	0	66.1%	33.9%	0.0%	6.909
10	2	1	80.7%	19.2%	0.1%	6.790
11	2	2	99.1%	0.9%	0.0%	6.717
12	2	3	68.0%	31.9%	0.0%	6.848
13	3	0	69.4%	30.6%	0.0%	6.626
14	3	1	66.1%	33.9%	0.0%	7.131
15	3	2	34.6%	36.5%	28.9%	7.042
16	3	3	12.8%	6.6%	80.6%	6.790

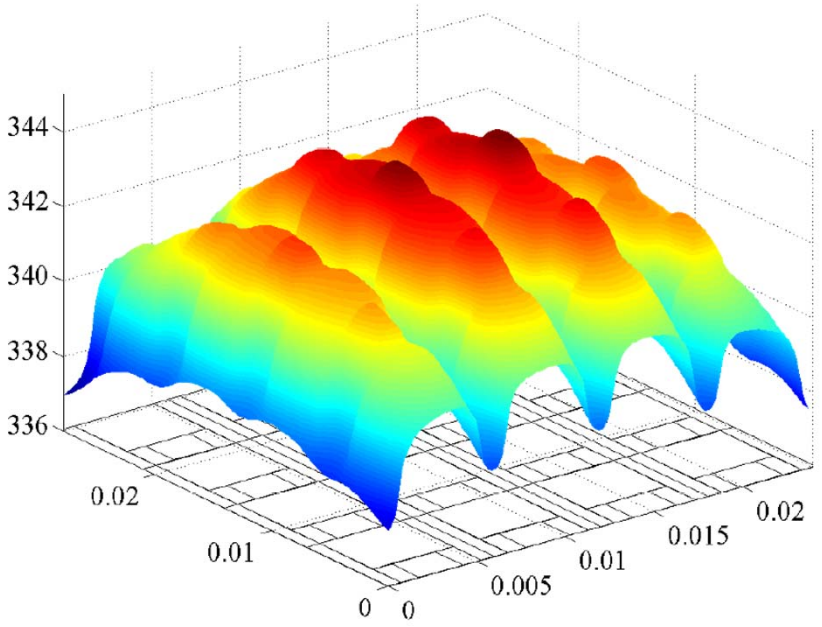

Ring organization of NUCA architectures

- **Tile-based architecture grouping tiles with similar thermal profiles and relative location**

16-cores, 2D mesh and 2 rings

Basic idea: inner rings will receive fewer CPU time to reduce power consumtpion, hence temperature

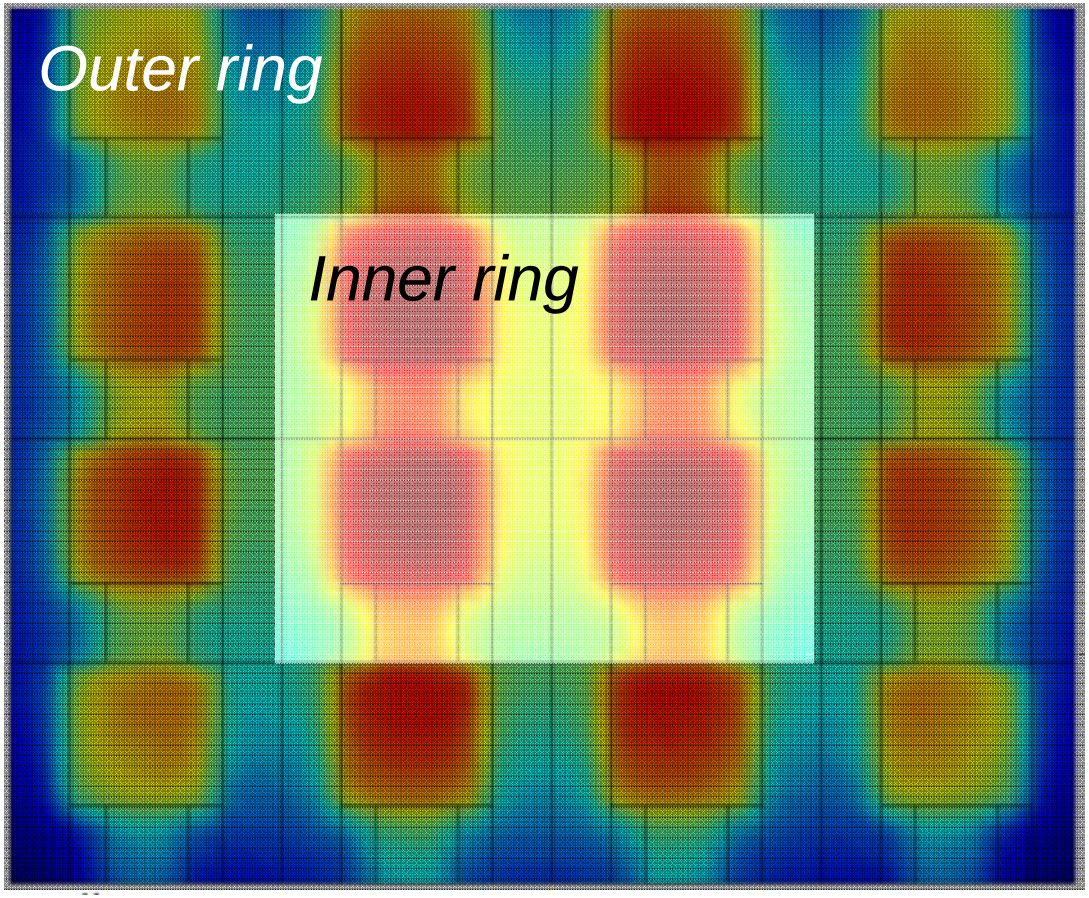

What's next

- Introduction
  - From single-core to multi-core
  - ITRS projections
- **Key observations**
  - The ring-based organization
  - The proposed methodology
- Results
    - Simulation framework
- Future perspective

Linear formal model

- **From clock-toggling to temperature**

$$t_{i,j} := \sum_{d \in D} (\alpha_d \cdot r_d) \qquad \forall \ (i,j) \in R \times C,$$

r_d is the duty-cycle for island d

α_d is the coefficient to be determined

$t_{i,j}$ is the temperature on tile (i,j), determined as a linear function over all thermal islands

- **From clock-toggling to performance**

$$p_{i,j} = r_{f(i,j)} \qquad \forall \ (i,j) \in R \times C,$$
$$f := (i,j) \rightarrow D,$$

$P_{i,j}$ is the performance level of (i,j) tile

$r_{f(i,j)}$ is the duty-cycle for thermal island that owns tile i,j

f(i,j) is a mapping function from i,j to the thermal island

Optimization: LP formulation

- **Objective**
  - Maximize the lower tile performance (fairness)

$$max\ q$$
$$q \leq p_{i,j} \qquad \forall\ (i,j) \in R \times C$$

- **Constraints**
  - Linear performance/duty-cycle relation

$$p_{i,j} = r_{f(i,j)} \qquad \forall\ (i,j) \in R \times C,\ f := (i,j) \rightarrow D,$$

  - Linear thermal/duty-cycle relation

$$t_{i,j} := \sum_{d \in D} (\alpha_d \cdot r_d) \qquad \forall\ (i,j) \in R \times C,$$

  - Thermal constraints on maximum allowed temperature

$$t_{i,j} \leq T_{max} \qquad \forall\ (i,j) \in C \times R$$

- **Variables**
  - The duty-cycle to each thermal island d

$$r_{f(i,j)} \qquad \forall\ (i,j) \in R \times C,\ f := (i,j) \rightarrow D,$$

What's next

- Introduction
  - From single-core to multi-core
  - ITRS projections
- Key observations
  - The ring-based organization
  - The proposed methodology
- **Results**
  - Simulation framework
- Future perspective

Experimental results

- **HANDS estimation framework**
  - Design-time analysis on application performance and thermal impact

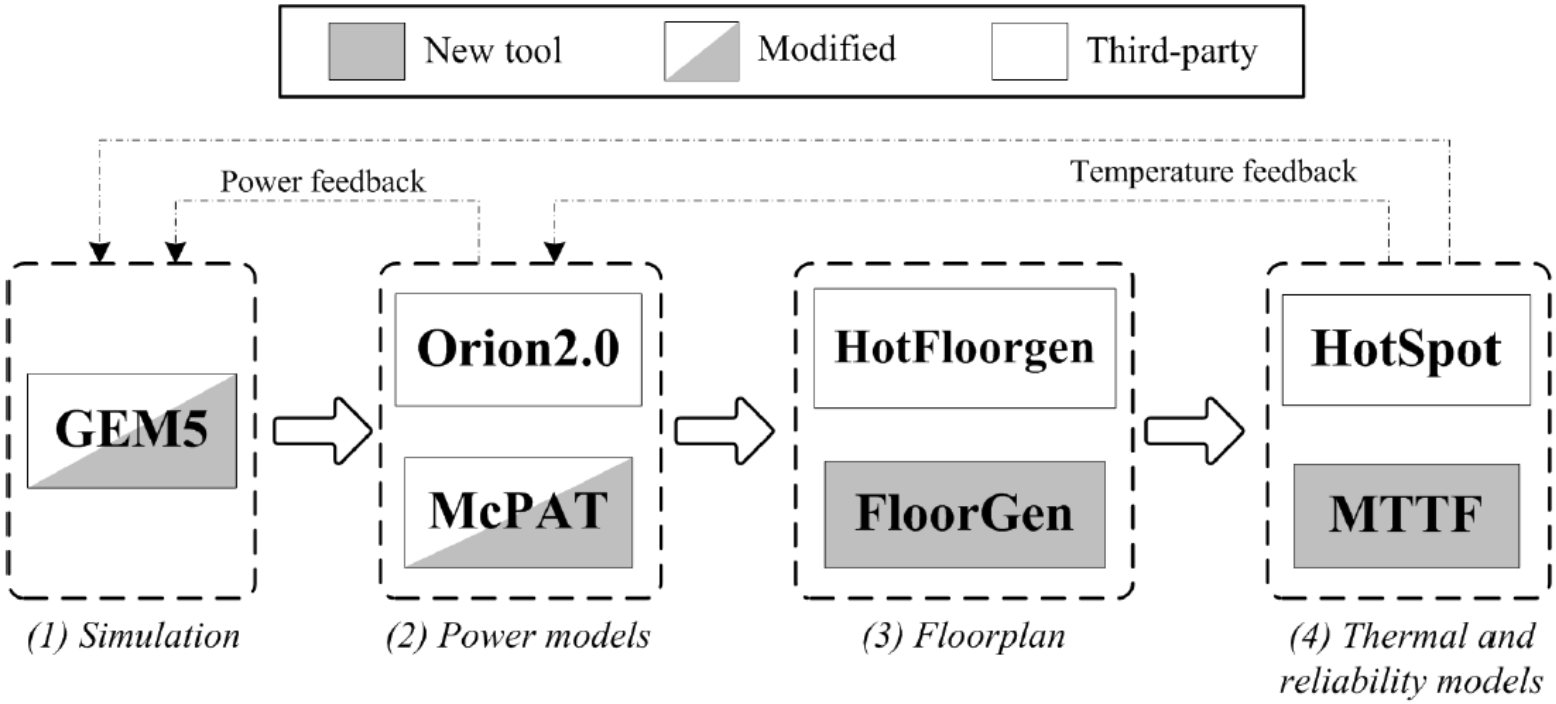

D. Zoni, S. Corbetta and W. Fornaciari. **"HANDS: Heterogeneous Architectures and Networks-on-Chip Design and Simulation",** ISLPED'12

Experimental settings

- **Multiple benchmarks configuration**
  - Different suites WCET, SPLASH2 and MIBENCH
  - Different instruction breakdwon

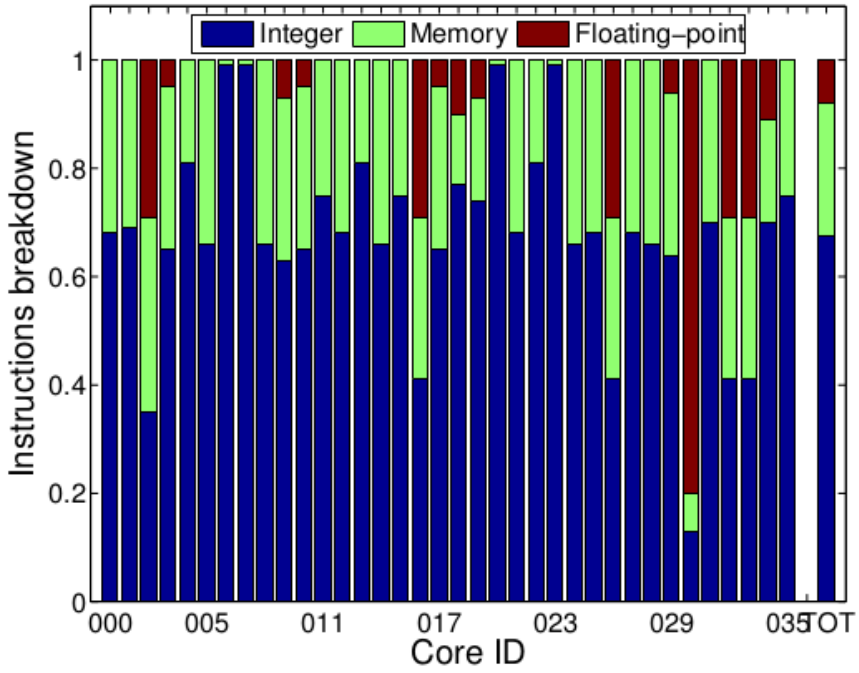

- **Alpha-2136 like architecture**

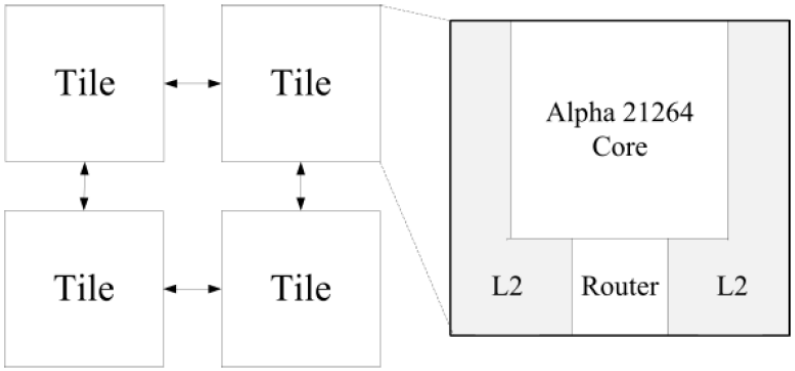

Processor core	3GHz, in-order based on Alpha21264 core
Int-ALU	4 integer ALU functional units
Int-Mult/Div	4 integer multiply/divide functional units
FP-Mult/Div	4 floating-point multiply/divide functional units
L1 cache	64kB 2-way set assoc. split I/D, 2 cycles latency
L2 cache	1.75MB per bank, 8-way associative
Router	2-stage wormhole switched (Garnet network [1])
Topology	2D-mesh based on Alpha21364 network processor
Technology	45nm at 1.1V

Model validation

- **Accuracy of predicted temperature**
  - Very good accuracy irrespective of threshold temperature
  - Low variance, limited to 0.37K

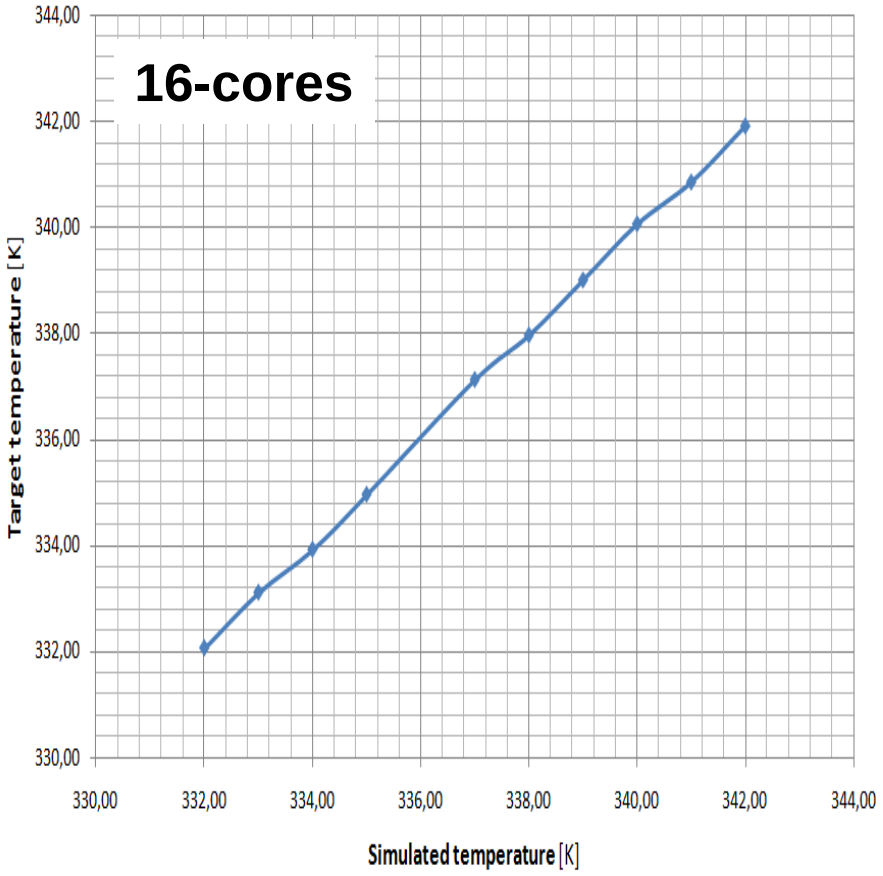

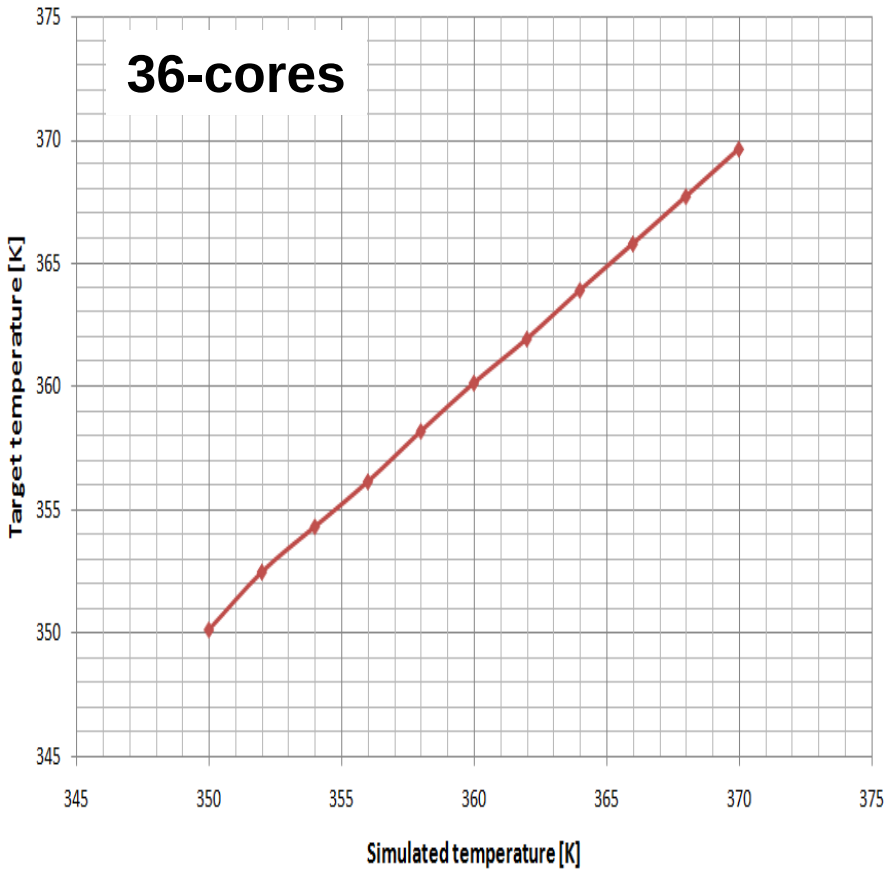

Reliability improvement analysis

- **Theoretical and modeled reliability analysis**
  - Two major (FEOL and BEOL) fault mechanisms ★
    - Stress migration $$MTTF_{SM} \propto |T_0 - T|^{-n} \cdot exp\left\{\frac{E_{SM}}{k\,T}\right\}$$
    - Electromigration $$MTTF_{EM} \propto exp\left\{\frac{E_{EM}}{k\,T}\right\}$$

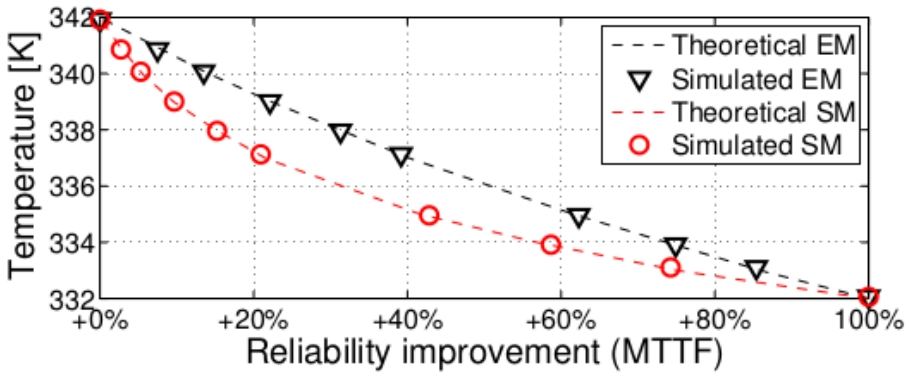

(a) 16-cores

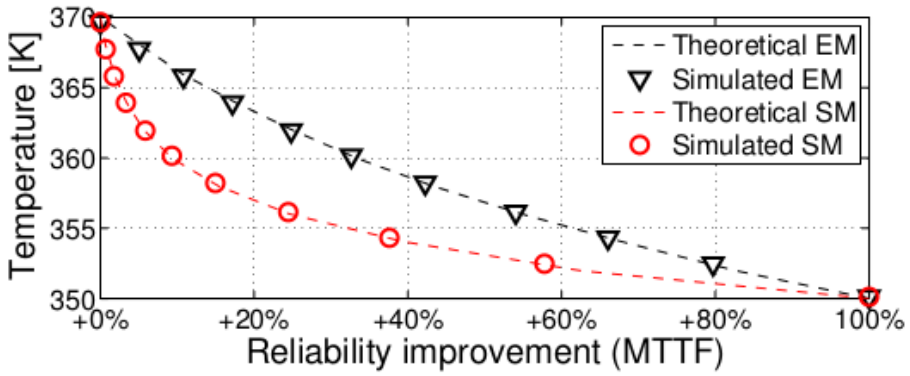

(b) 36-cores

★ J. Srinivasan, S. Adve, P. Bose, and J. Rivers. **"The case for lifetime reliability-aware microprocessors"**. In Computer Architecture, 2004.

What's next

- Introduction
  - From single-core to multi-core
  - ITRS projections
- Key observations
  - The ring-based organization
  - The proposed methodology
- Results
    - Simulation framework
- **Future perspective**

Conclusions and future works

- **We proposed a design-time thermal/performance optimization model**
  - Targeting 2D-mesh architectures and NoCs
  - Based on linear formal model suitable for LP formulation and design-time optimization
  - The proposed model can be jointly used with existing dynamic thermal management techniques

- **Future works**
  - Adoption of the proposed methodology to NoC routers
  - Greater degree of freedom: rings are partitioned in multiple area (e.g., cores on the edges or at the corners)

Kiitos!

Any questions?

Take a tour to my web site for more info and tools for:

- *Thermal/power-related reliability (HANDS)*
- *Run-time resource management (BBQ)*
- *Analysis and optimization of Sw energy (SWAT)*

Prof. William Fornaciari
Politecnico di Milano – DEI
fornacia@elet.polimi.it
http://home.dei.polimi.it/fornacia

Validation of the performance clock toggling relation

- We assess the linear relation using a 16-core architecture with 2 rings
- The graph shows the quasi linear behavior of the proposed model, where theoretical (dash-line) and experimental (points) data are very close

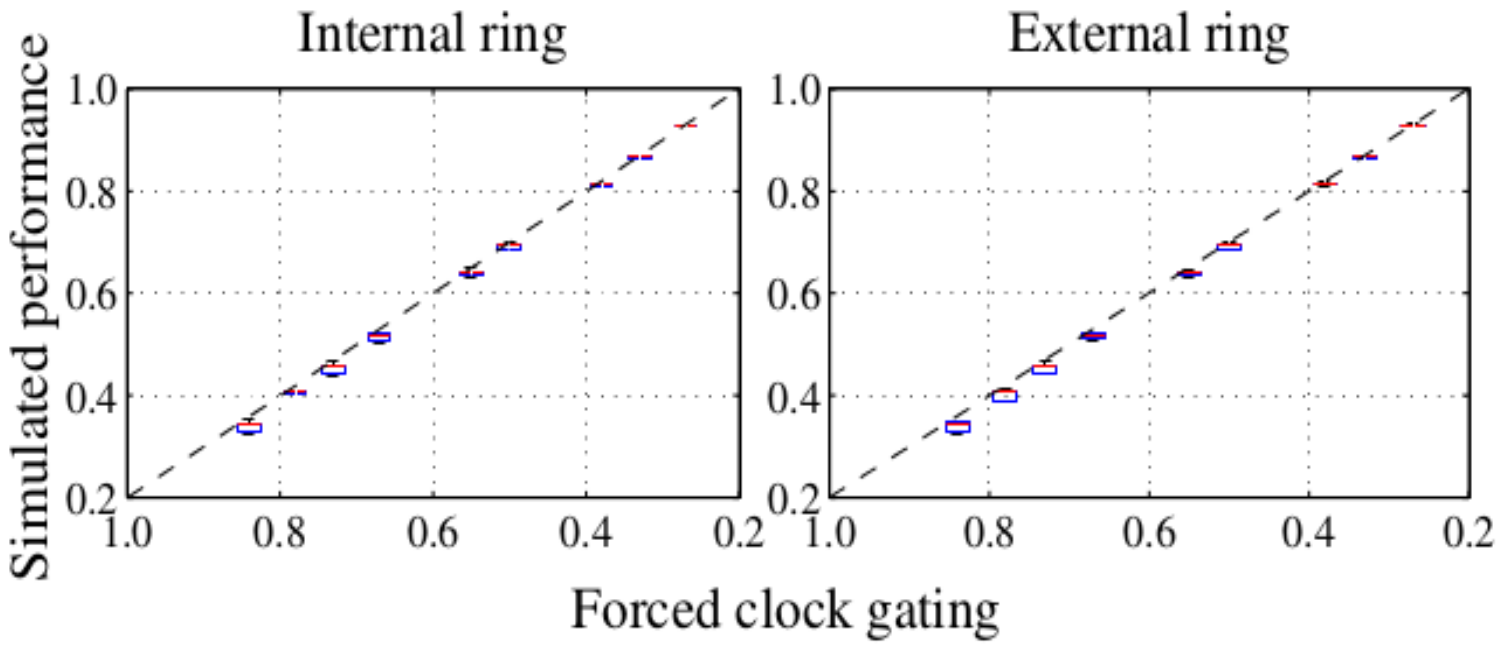

Validation of the thermal clock toggling relation

- We assess linear relation using 16-core architecture with 2 rings
- We sampled a huge space domain (left figure)
- The data are distributed over a virtual plane on 3d space, following a linear function (right figure)