New Exploration Frameworks for Temperature-Aware Design of MPSoCs

Prof. David Atienza

Dept Computer Architecture and Systems Engineering (DACYA)
Complutense University of Madrid, Spain

Integrated Systems Lab (LSI)
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

MPSoC '08 **Château St. Gerlach, Valkenburg, The Netherlands**

Evolution of Electronics to Multi-Processor System-on-Chip (MPSoC)

- Roadmap continues: 90→65→45 nm
- Multi-Processor System-on-Chip (MPSoC) architectures: complex HW and SW

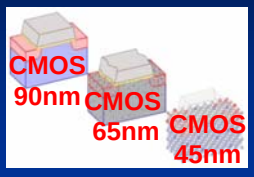

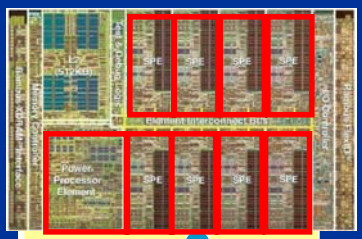

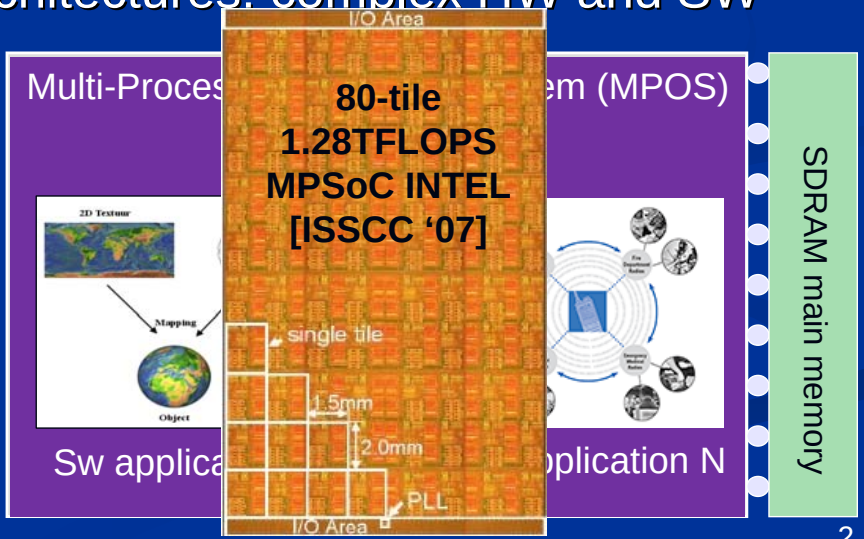

Design Issues in MPSoCs

- MPSoCs have very sophisticated architectures
  - Complex components and CAD tools very expensive
  - Time-closure issues, system speed decreased
- Increasing thermal issues
  - Hot-spots, non-uniform thermal gradients

[Sun, 1.8 GHz Sparc v9 Microproc]

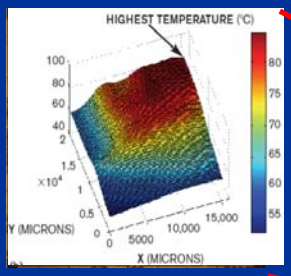

[Santirini, EDN, March '05]

[Sun, Niagara Broadband Processor]

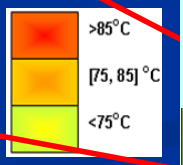

[Coskun et al '07, Sun]

High chances of thermal wear-outs and very short lifetimes!

Advocating Thermal-Aware Design

- Integration of HW/SW modeling and management

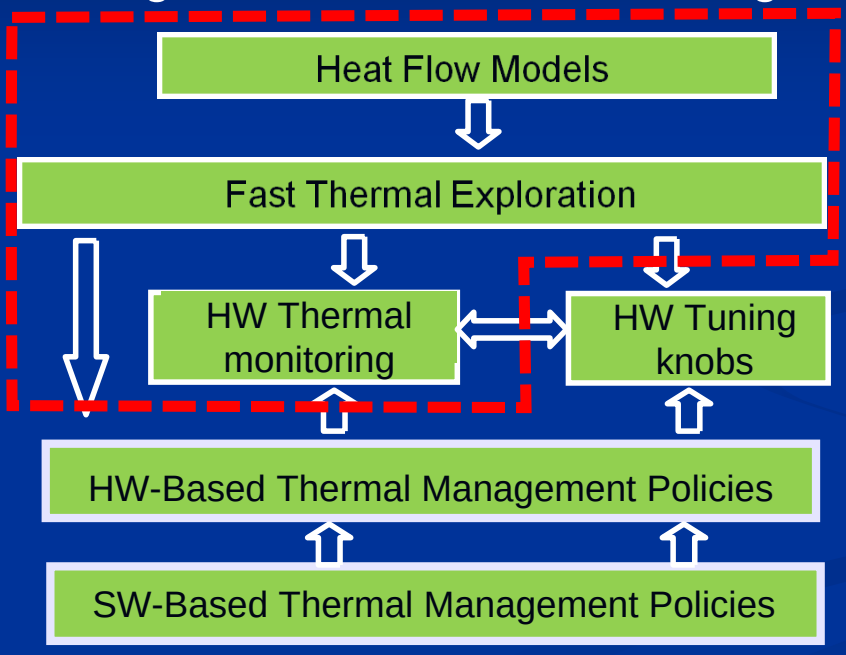

Thermal and architectural information

HW-SW emulation frameworks

HW reconfigurable architecture

HW and compiler-level algorithms

OS and application-level algorithms

Outline

- Introduction
- MPSoC thermal modeling
  - Computational infrastructure
  - Closed-loop statistics extraction system
  - Thermal model
- Case studies
- Summary and conclusions

MPSoC Thermal Modeling Problem

- Continuous heat flow analysis
  - Capture geometrical characteristics of MPSoCs
  - Explore different packaging features and heat sink characteristics
- Time-variant heat sources
  - Transistor switching depends on MPSoC run-time activity
  - Dynamic interaction with heat flow

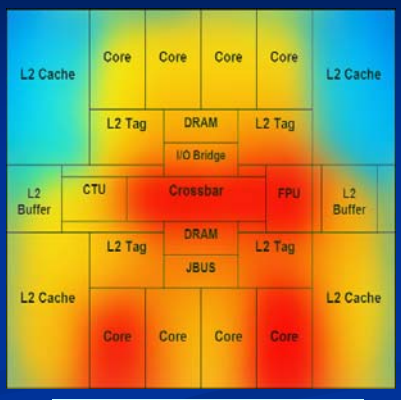

Very complex computational problem!

MPSoC Thermal Modeling State-of-The-Art

- MPSoC Modeling and Exploration:

1. SW simulation: Transactions, cycle-accurate (~100 KHz)
   [Synopsys Realview, Mentor Primecell, Madsen et al., Angiolini et al.]

At the desired cycle-accurate level, they are too slow for thermal analysis of real-life applications!

2. HW prototyping on FPGAs: Core dependent (~MHz):
   [Cadence Palladium II, ARM Integrator IP, Heron Eng, Marescaux et al.]

Combination of cycle-accurate MPSoC behavior and run-time IC heat flow modeling is unheard of

- Heat Flow Modeling:

1. Software thermal/power models [Skadron et al., Kang et al.]

Too computationally intensive, not able to use detailed run-time inputs from MPSoC components!

Orthogonalizing MPSoC Thermal Modeling and Analysis

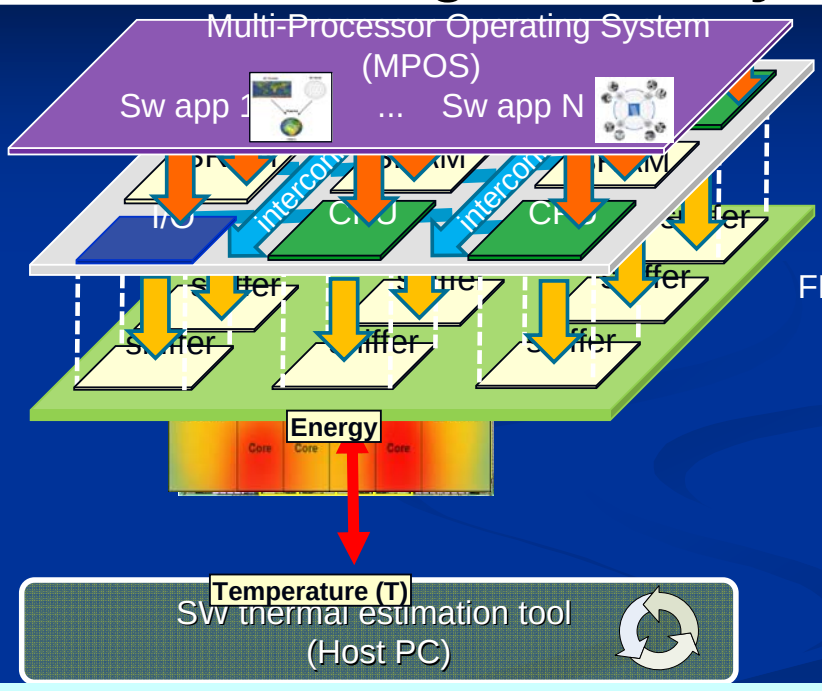

Framework: MPSoC behavioral model on reconfigurable HW interacting with efficient thermal estimation

Emulation vs. Simulation/Prototyping

- *"To emulate an electronic system is to build a platform capable of imitating its behaviour in an accurate and analyzable way"*

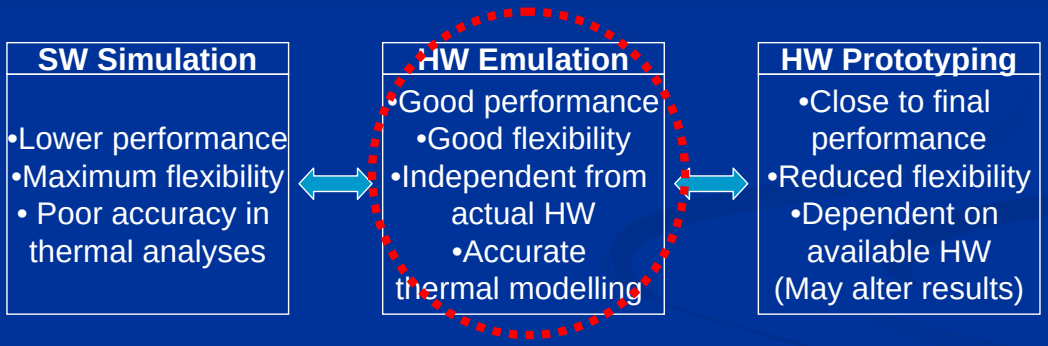

- More efficient than SW simulation
- More flexible than pure HW prototyping
- Thermal behavior analysis earlier than HW prototyping

MPSoC Architecture Emulation

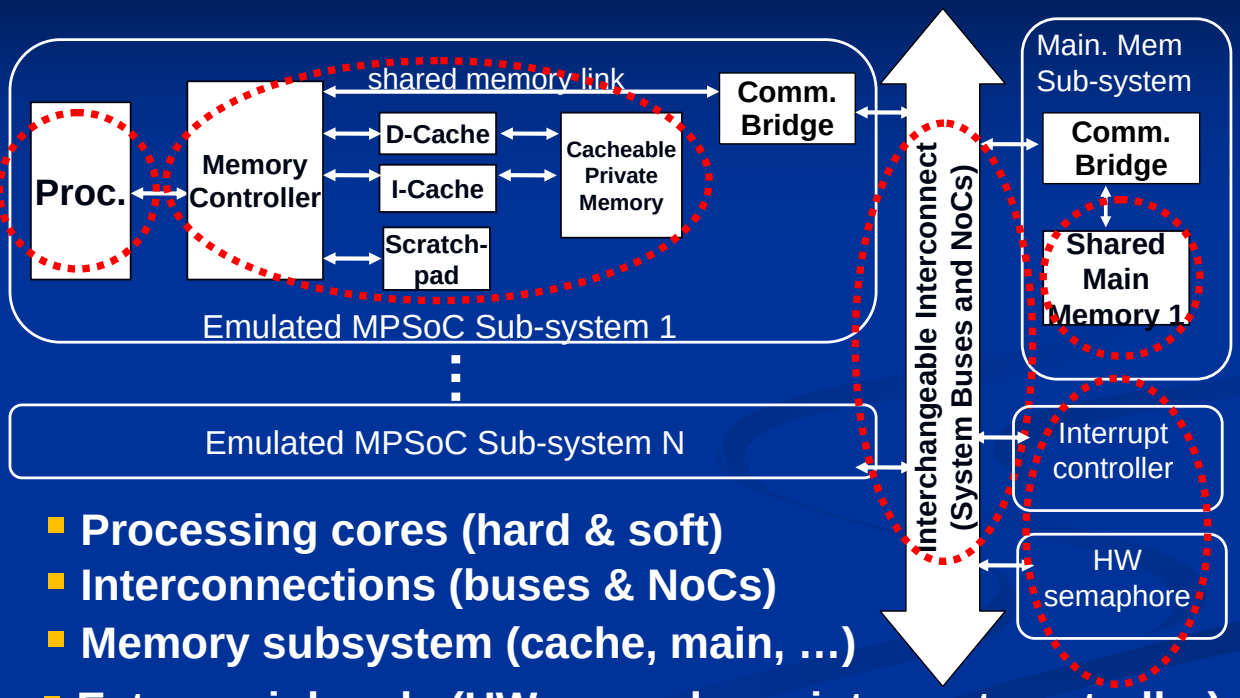

- **Processing cores (hard & soft)**
- **Interconnections (buses & NoCs)**
- **Memory subsystem (cache, main, …)**
- **Extra peripherals (HW semaphore, interrupt controller)**

Designing the Emulated Architecture

What we want **VS.** **What we have**

Desired:
A fast memory
Latency: 1 cycle

Mapping

Physical resource:
On board BRAM
Latency: 1 cycle

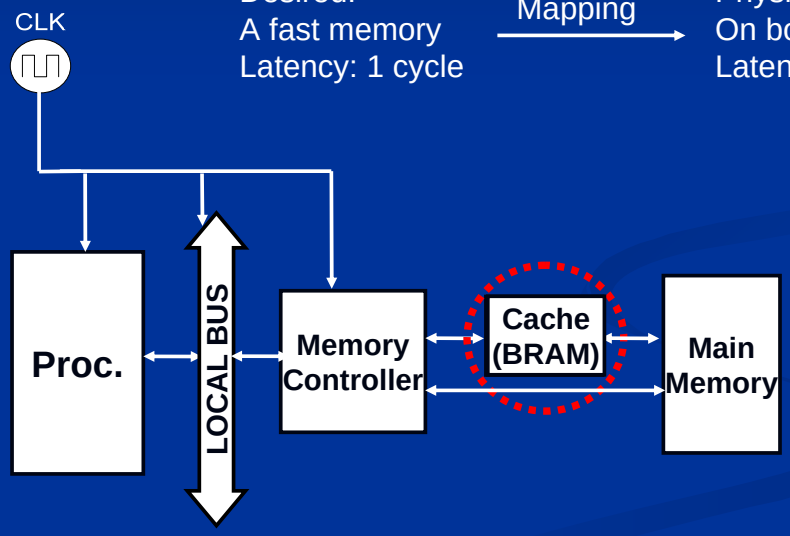

Designing the Emulated Architecture

What we want **VS.** **What we have**

Desired:
Big memory
Latency: 3 cycles

~~Mapping~~

Physical resource:
On board DDR RAM
Latency: 5 cycles

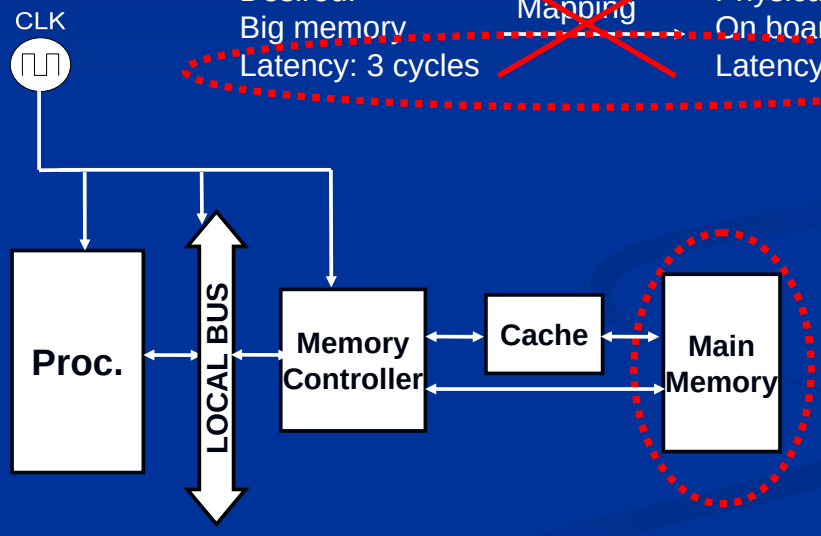

Designing the Emulated Architecture

2) Clock management

Proposed solution:

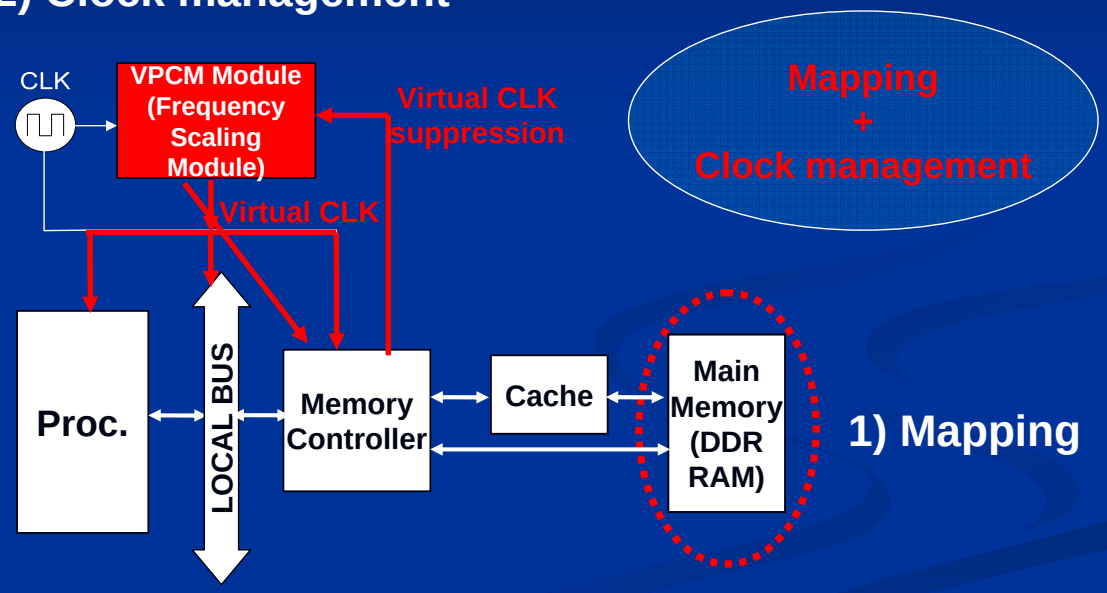

Statistics Extraction Subsystem

Another clock domain (100 MHz)

Dedicated bus & HW: Not intrusive, concurrently running

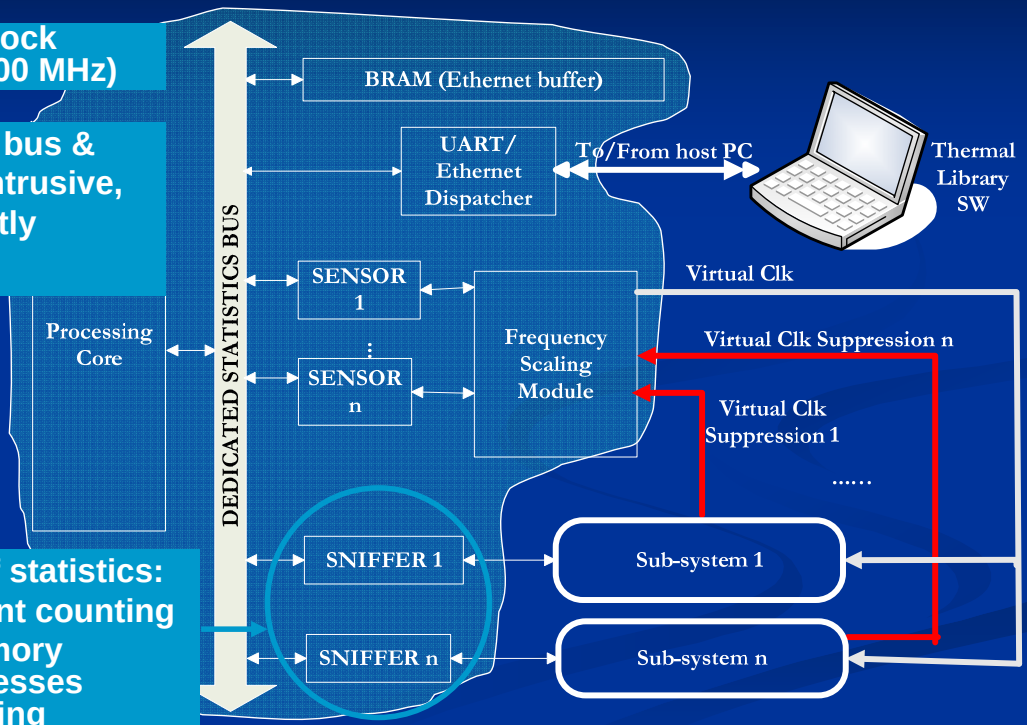

2 Types of statistics:

- Event counting
- Memory accesses tracing

Sending Energy Values

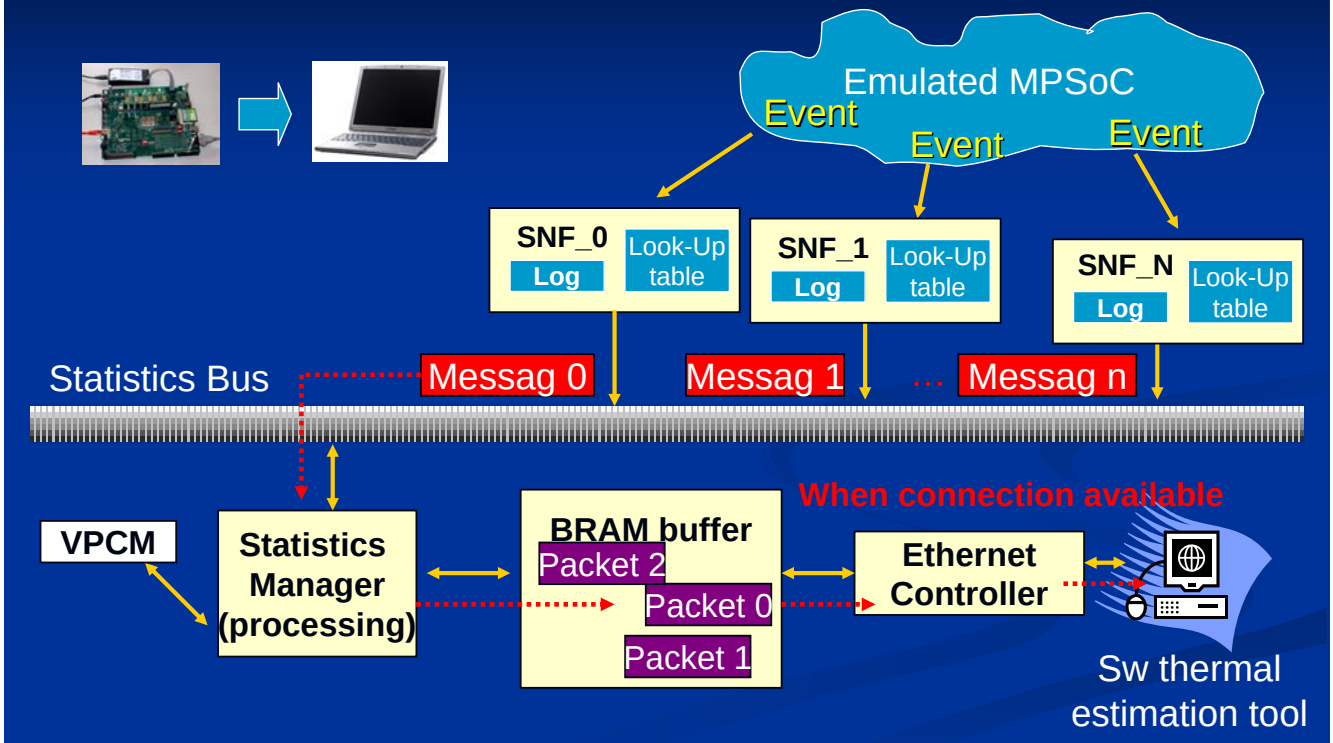

Receiving Temperatures

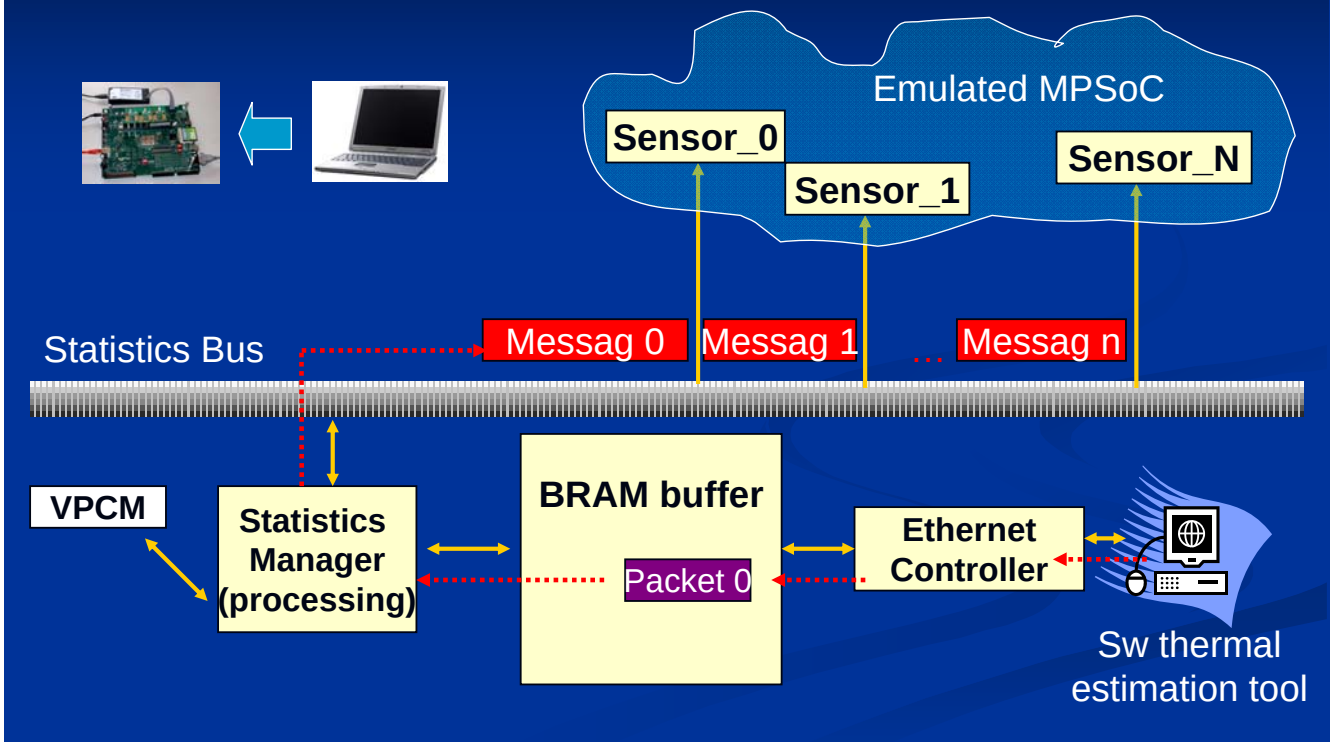

Chip and Package Heat Flow Modeling

- Model interface
  - Input: power model of MPSoC components, geometrical properties
  - Output: temperature of MPSoC components at run-time
- Thermal circuit: 1st order RC circuit
  - Heat flow ~ Electrical current ;
  - Heat spreader and IC compos

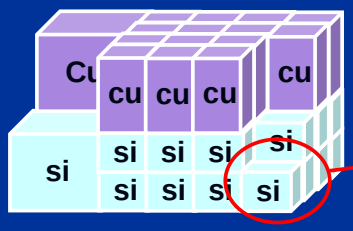

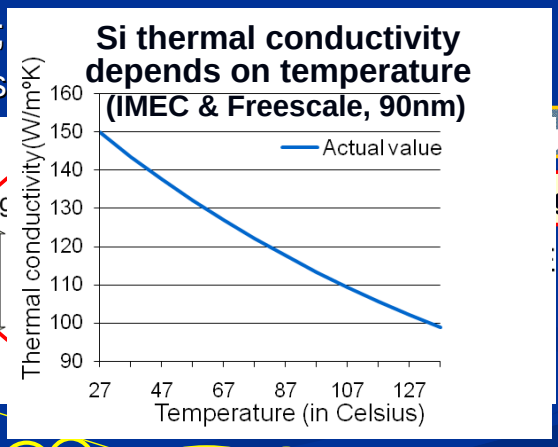

Thermal conductance matrix

Thermal capacitance matrix

$$\begin{bmatrix} C_{si,1} & & & \\ & C_{si,2} & & \\ & & \ddots & \\ & & & C_{cu,n} \end{bmatrix} \begin{bmatrix} G_{1,2}-G_{1,2} & & \\ -G_{2,1} & G_{2,1} & \\ & & \ddots \end{bmatrix}$$

Temperature change

$$C\dot{t}_k = -G(t_k)t_k + p_k \; ; \; k = 1..m$$

Temperature vector at instant k

power consumption vector

SW Thermal Estimation Tool for MPSoCs

$$C\dot{t}_k = -G(t_k)t_k + p_k \; ; \; k = 1..m$$

- Creating linear approximation while retaining variable Si thermal conductivity:
  - Si thermal conductivity linearly approx. : $G_{i,j}(t_k) = I + q\,t_k$
  - Numerically integrating in discrete time domain the

$$t_{k+1} = A(t_k)t_k + B$$

Complexity scales linearly with the number of modeled cells (simulated on P4@ 3GHz)

thermal library validated against 3D finite element model (IMEC & Freescale)

Si thermal conductivity

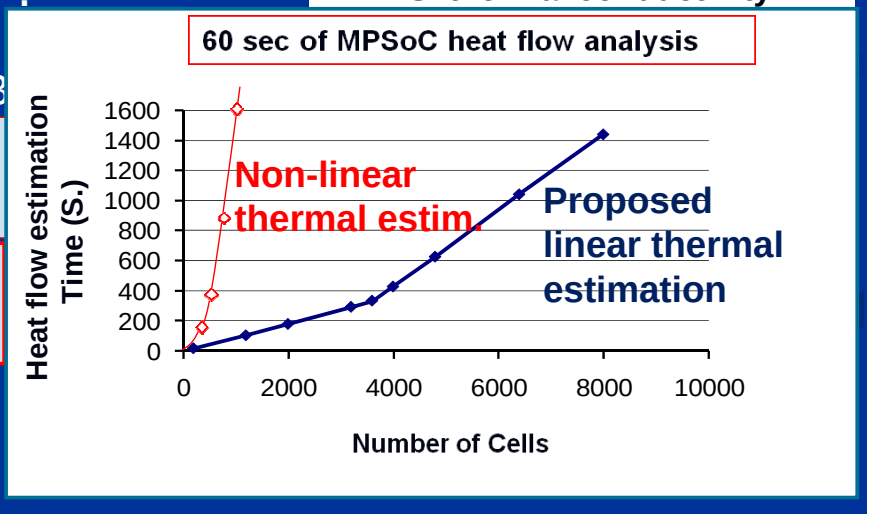

Outline

- Introduction
- MPSoC thermal modeling
  - Computational infrastructure
  - Closed-loop statistics extraction system
  - Thermal model
- Case studies
- Summary and conclusions

Case Study 1: 4-Core MPSoC

- MPSoC Philips board design:
  - 4 processors, DVFS: 100/500 MHz
  - Plastic packaging
- Software:
  - Image watermarking, video rendering
- Power values for 90nm:

Element	Max Power (mW) 100 MHz	Max Power (mW) 500 MHz
Processor	2,92 x 10^2	1,02 x 10^3
D-Cache	1,42 x 10^2	7,10 x 10^2
I-Cache	1,42 x 10^2	7,10 x 10^2
Priv Mem	0,61 x 10^2	2,75 x 10^2
AMBA	0,31 x 10^2	0,68 x 10^2

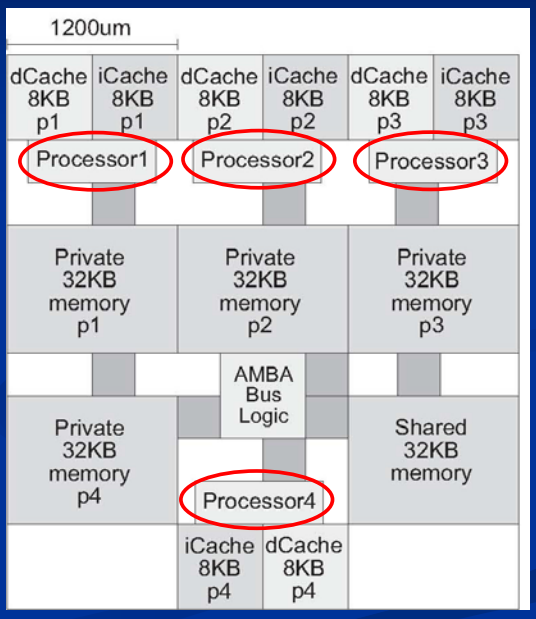

Thermal Validation 4-Core MPSoC

- MPARM: MPSoC SW simulator: power/thermal models tuned for Philips
  - Simulations too slow: 2 days for 0.18 real sec (12 cells)
- HW
  -

Very fast validation of MPSoC
run-time thermal behavior and management

s

on?!

HW/SW Thermal Emulation time 45 sec (128 cells)!

DVFS ON:

500/100 MHz.

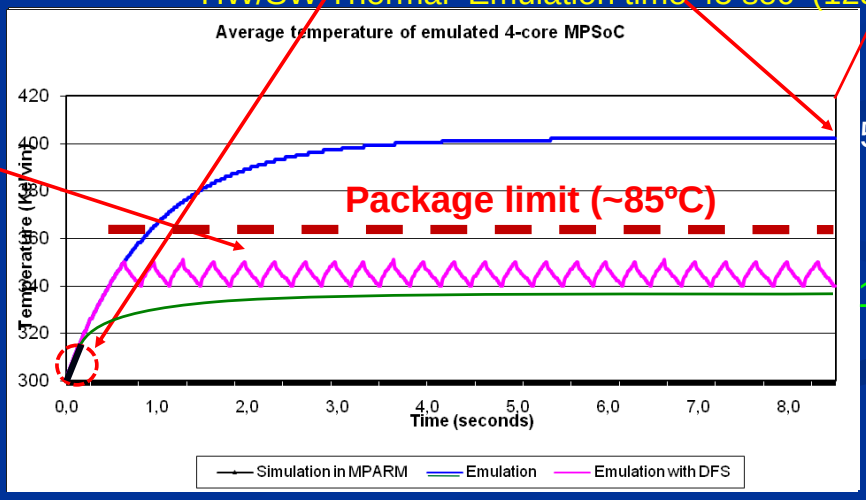

500 MHz.

100 MHz.

Case Study 2: Multi-Core Thermal Control

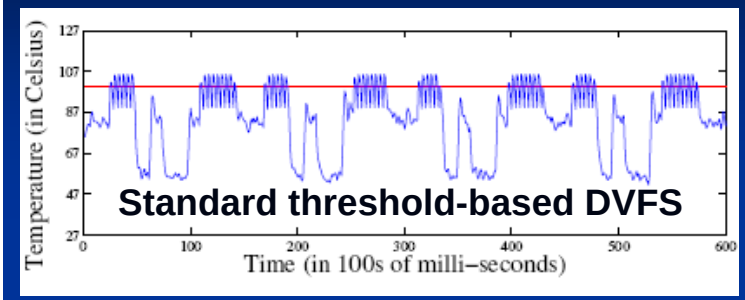

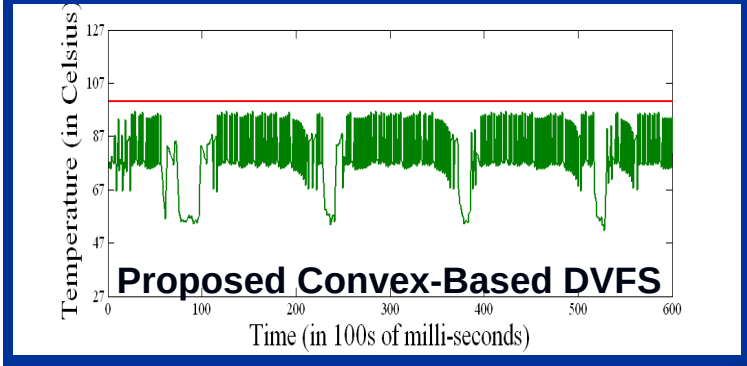

- HW: Sun 8-core Niagara
- MPOS: Solaris Multi-Core
- SW: Web & signal processing

Total run-time of benchmarks:

180 sec

106 sec (45% performance improvement)

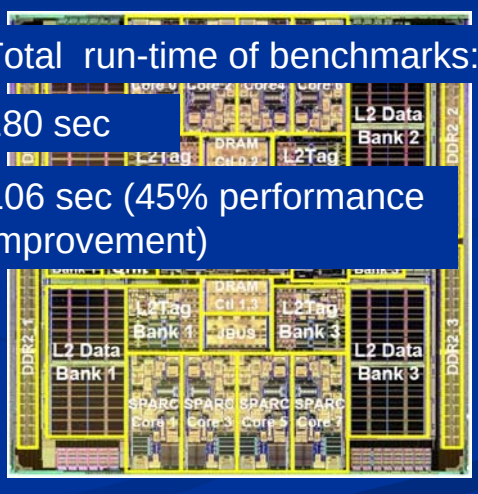

Enhancing thermal control in 90nm multi-cores

Outline

- Introduction
- MPSoC thermal modeling
  - Computational infrastructure
  - Closed-loop statistics extraction system
  - Thermal model
- Case studies
- Summary and conclusions

Conclusions

- Shrinking is providing new processing capabilities through MPSoCs, but creating critical thermal issues
- Need for new thermal-aware modeling methods
  - Orthogonalization of behavior extraction and thermal modeling
  - Novel MPSoC emulation on reconfigurable HW and SW thermal modeling to efficiently extract the behavior of time-variant heat sources
  - Closed-loop thermal evaluation framework enabling run-time testing of thermal management in real MPSoC platforms
- Validation with commercial MPSoC platforms:
  - Fast exploration of thermal behavior of complex MPSoCs
  - Effective tuning of thermal management
    - Thermal run-away avoided with HW-based policies, max. throughput
    - Thermal balancing with negligible performance overhead

Thank you!

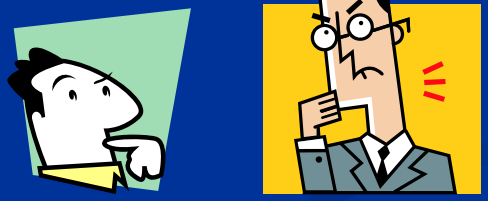

QUESTIONS ?

Acknowledgements:

UCSD / Sun

IMEC / Philips

Freescale semiconductors