ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Thermal-Aware System-Level Design of 2D/3D MPSoC Architectures

David Atienza
Embedded Systems Laboratory (ESL),
Ecole Polytechnique Fédérale de Lausanne (EPFL)

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

- Roadmap continues: 90→65→45→32 nm
- Multi-Processor System-on-Chip (MPSoC) architectures are a reality for a while...

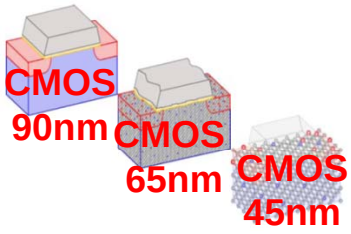

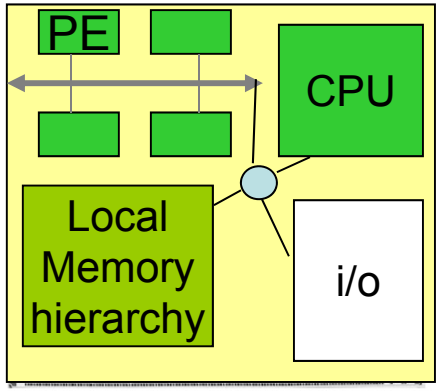

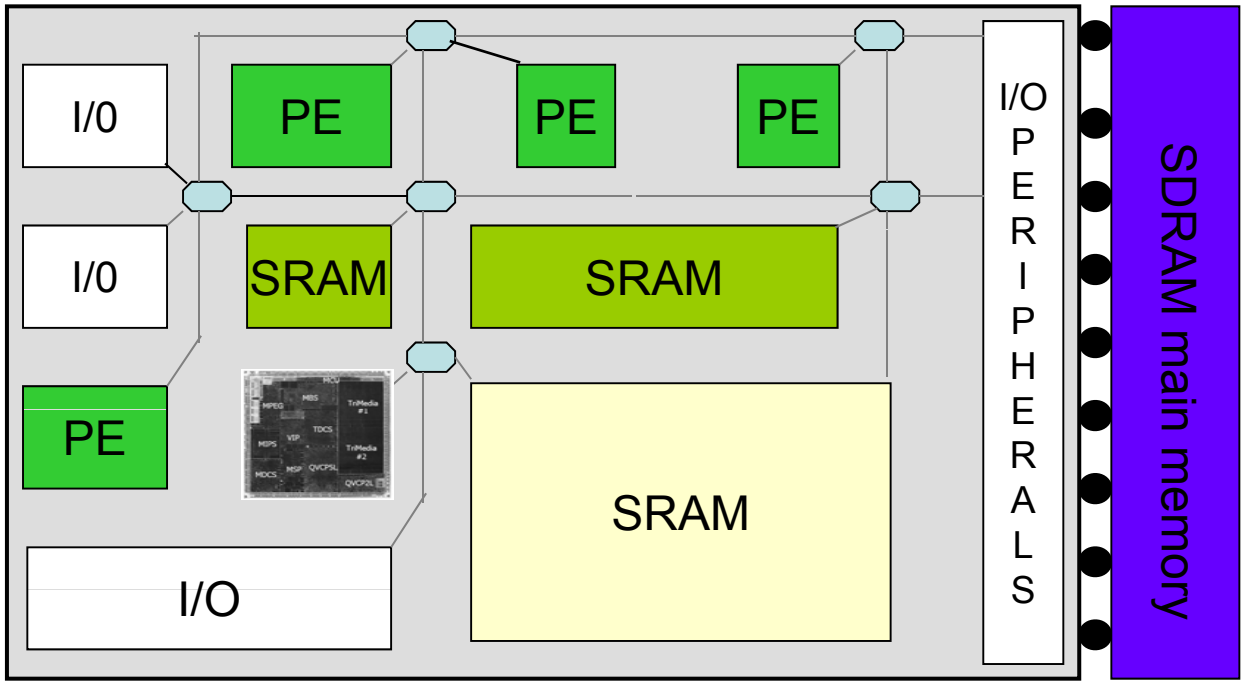

Evolution to Multi-Processor System-on-Chip (MPSoC)

- Roadmap continues: 90→65→45→32 nm
- Multi-Processor System-on-Chip (MPSoC) architectures are a reality for a while...

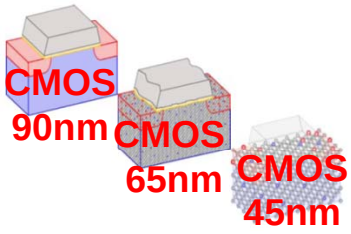

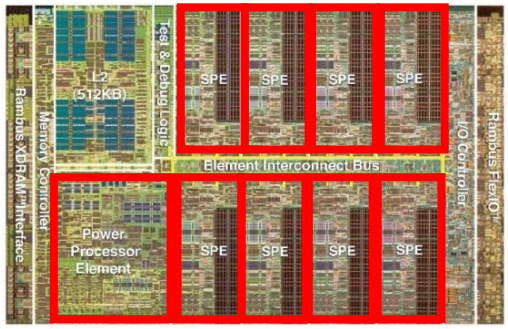

[Cell Multi-Processor – PS3]

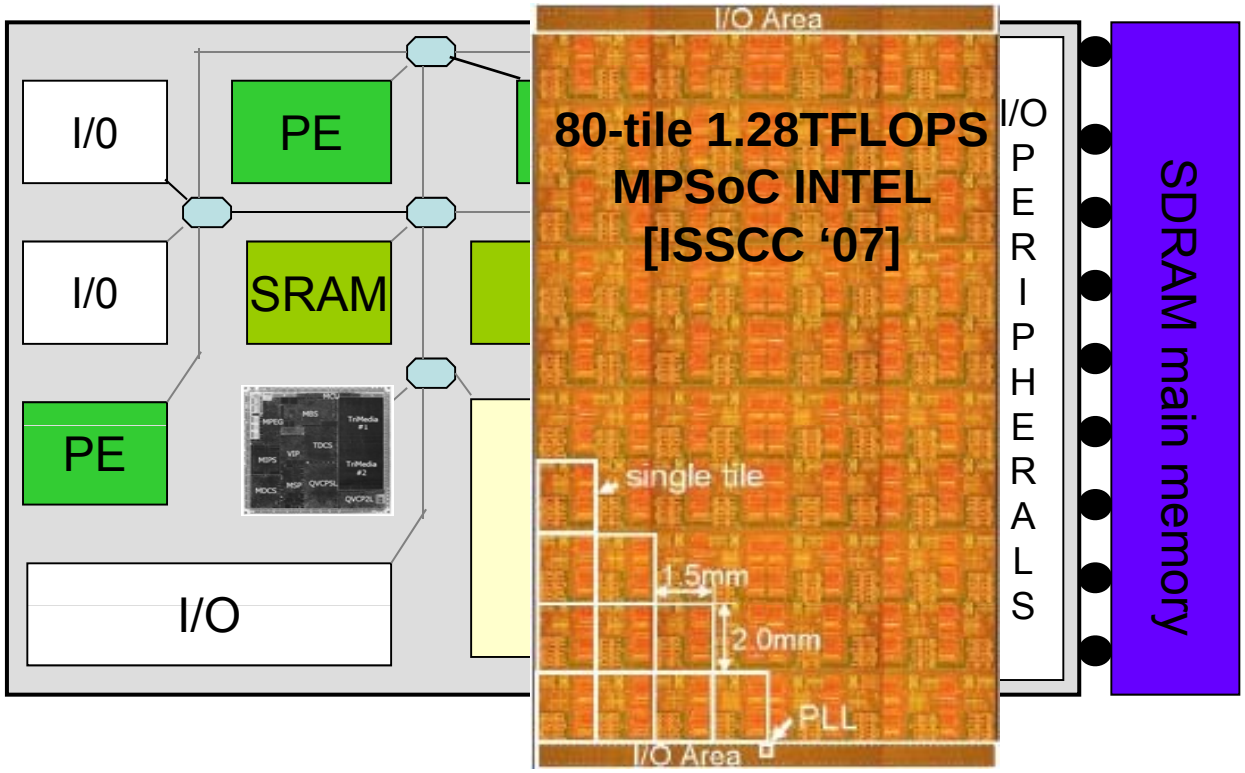

MPSoCs are Spreading Fast

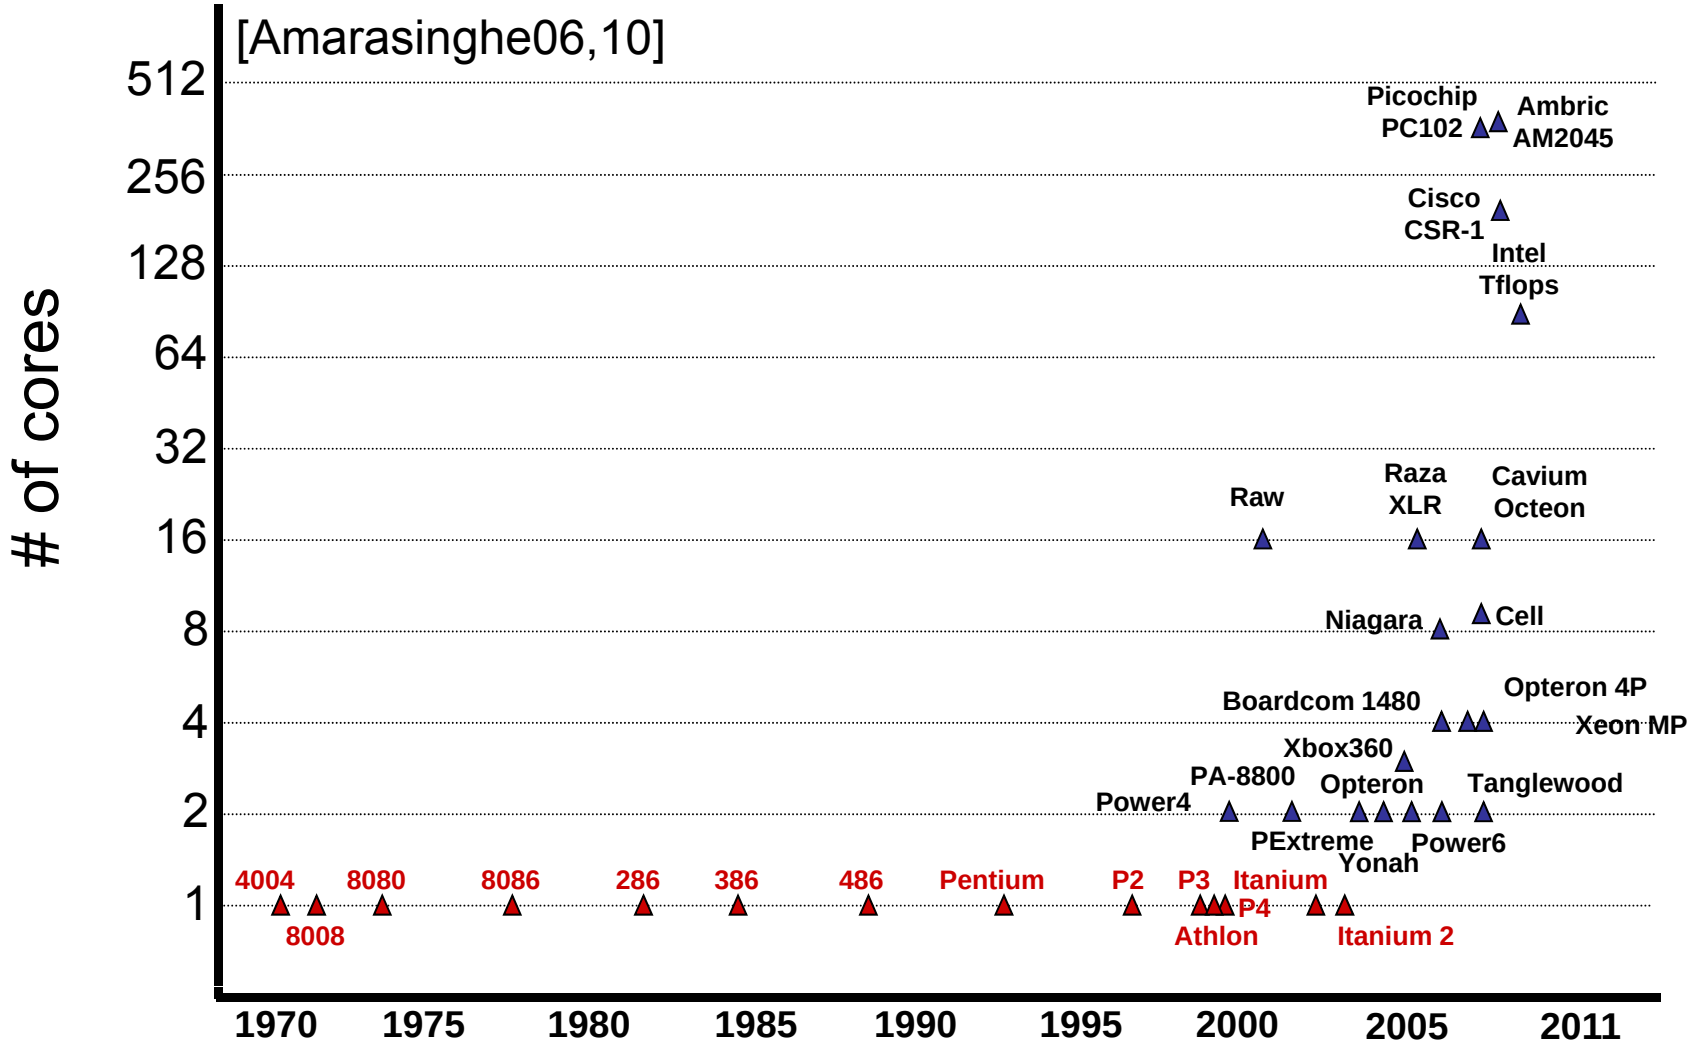

Design Issues in MPSoCs

- MPSoCs have very complex architectures
  - Advanced components and CAD tools very expensive
  - Time-closure issues, system speed decreased
- Aggravated thermal issues: Thermal-Aware MPSoC Layouts
  - Hot-spots, non-uniform thermal gradients

[Sun, 1.8 GHz Sparc v9 Microproc]

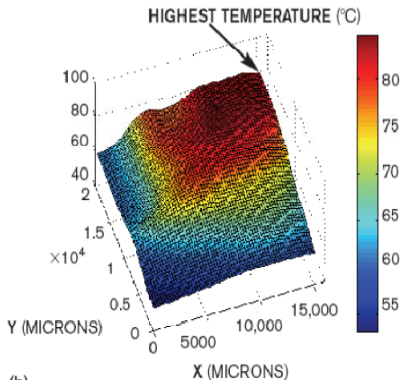

[Santarini, EDN, March '05]

[Sun, Niagara Broadband Processor]

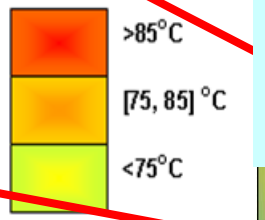

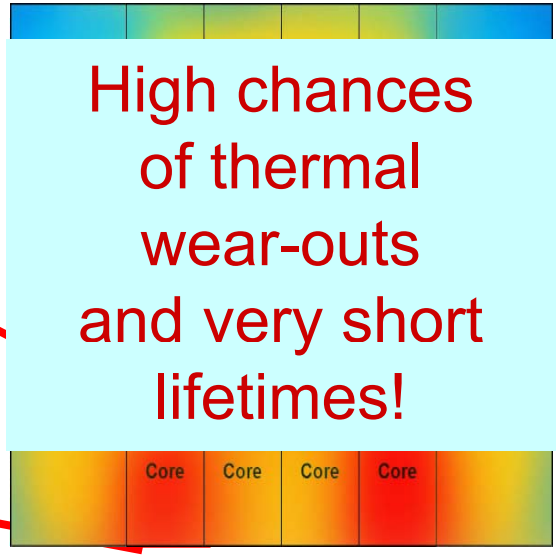

[Coskun et al '07]

Advocating Thermal-Aware 2D/3D MPSoC Design

- Integration of HW/SW modeling and management

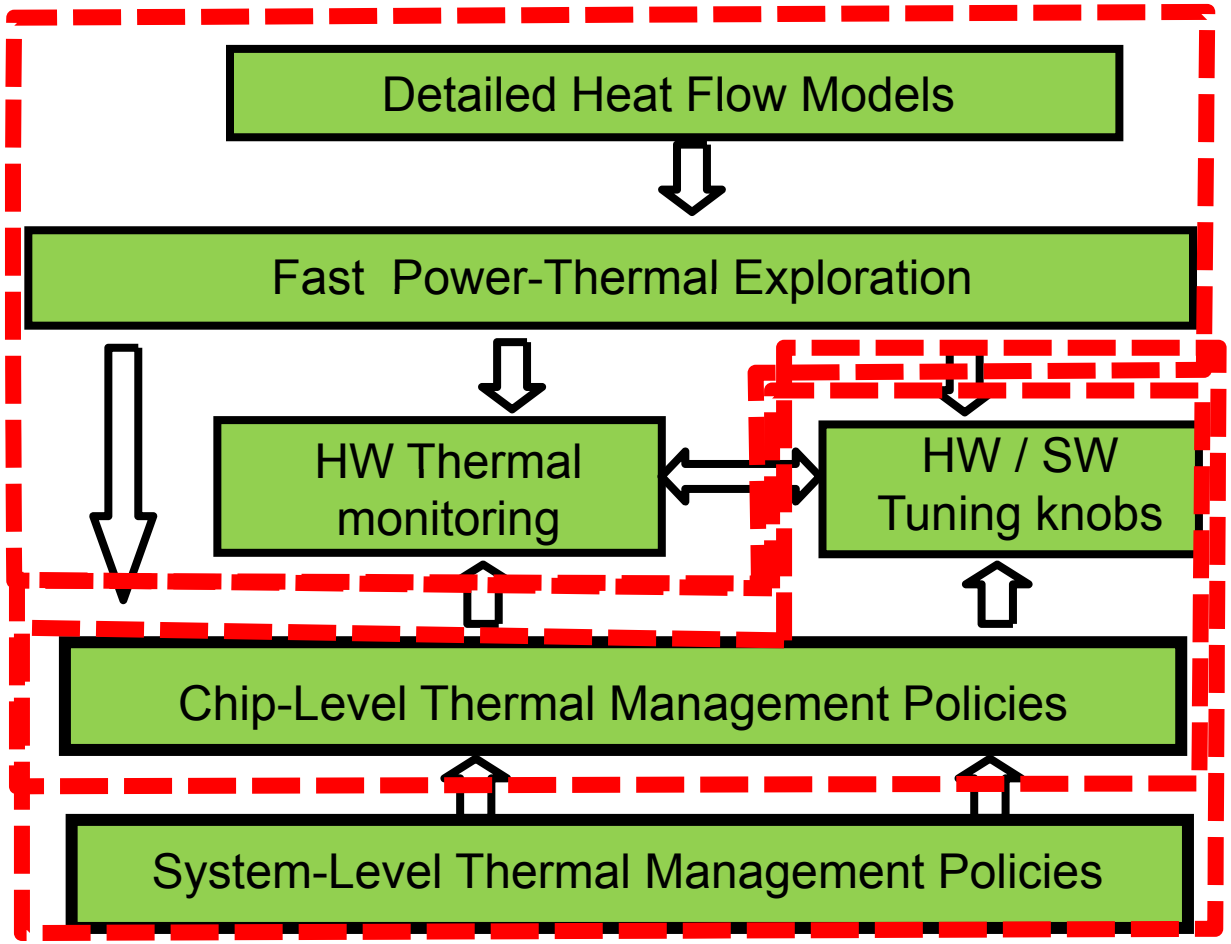

MPSoC Thermal Modeling Problem: Initial Thoughts

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

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

- Heat Flow Modeling
  - Finite-Element simulation
    [COMSOL Multiphysics [FEMLAB]

Too computationally intensive and very complex to tune in MPSoC with limited set of sensing components!

  - High-order RC-level heat flow models
    [Hotspot, Link et al.]

Not close-loop interaction at run-time with inputs from MPSoC components!

MPSoC Thermal Modeling Problem: Initial Thoughts

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

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

- Heat Flow Modeling
  - Finite-Element simulation

Request: Fast (and relatively accurate) thermal model for MPSoCs that enables close-loop run-time interaction

with limited set of sensing components!

  - High-order RC-level heat flow models
    [Hotspot, Link et al.]

Not close-loop interaction at run-time with inputs from MPSoC components!

RC-Based Thermal Modeling for MPSoC

- Model interface
  - Input: power model of tier components, geometrical properties
  - Output: temperature of tier components at run-time
- Thermal circuit: 1st order RC circuit
  - Heat flow ~ Electrical current ; Temperature ~ Voltage
  - Metal and Si layers composed of elementary blocks

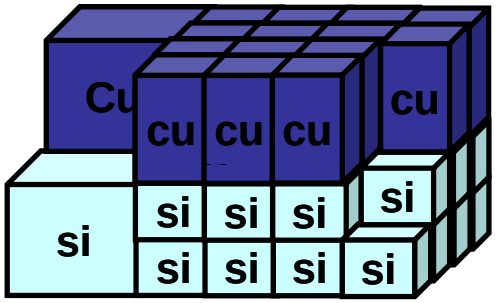

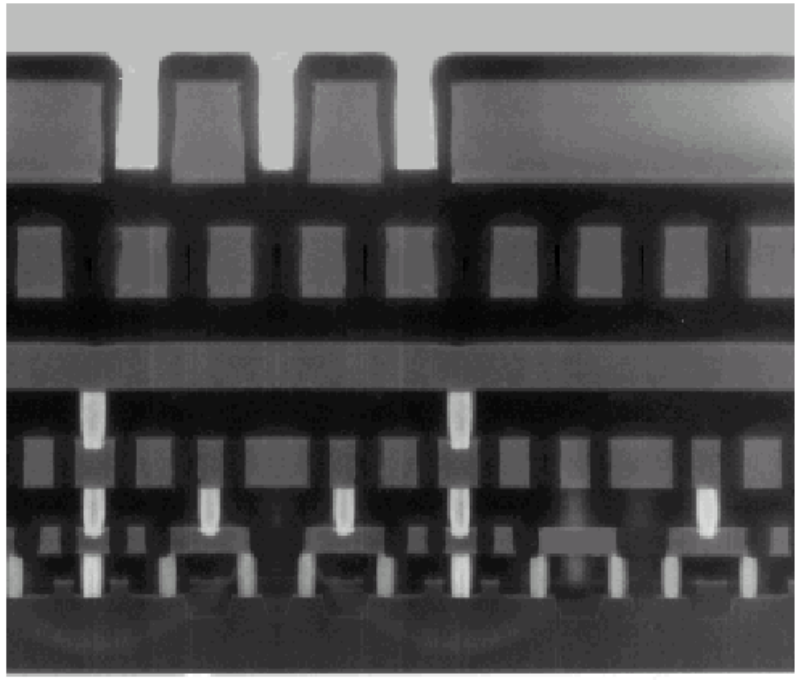

RC-Based Thermal Modeling for MPSoC

- Model interface
  - Input: power model of tier components, geometrical properties
  - Output: temperature of tier components at run-time
- Thermal circuit: 1st order RC circuit
  - Heat flow ~ Electrical current ; Temperature ~ Voltage
  - Metal and Si layers composed of elementary blocks

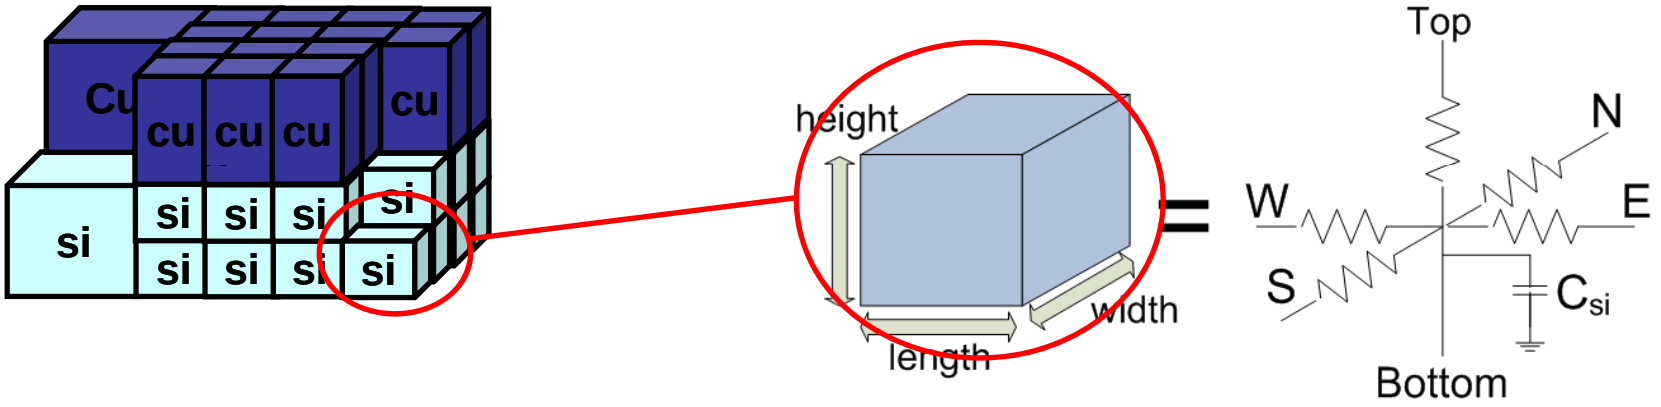

RC-Based Thermal Modeling for MPSoC

- Model interface
  - Input: power model of tier components, geometrical properties
  - Output: temperature of tier components at run-time
- Thermal circuit: 1st order RC circuit
  - Heat flow ~ Electrical current ; Temp
  - Metal and Si layers composed of ele

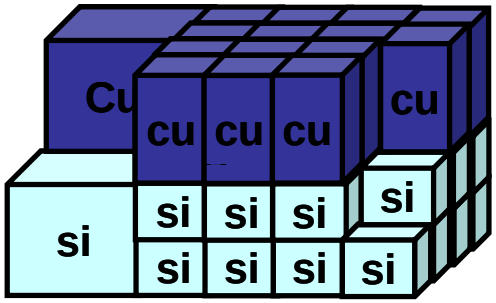

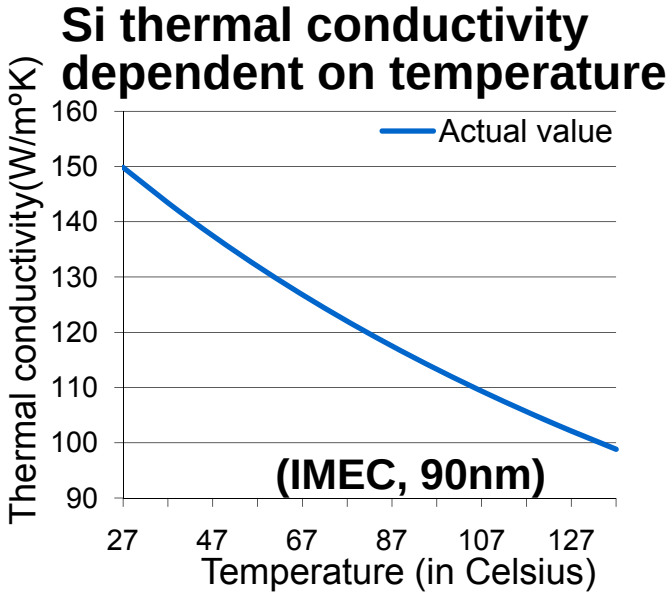

Thermal conductance matrix

Thermal capacitance matrix

[Atienza, et al., *TODAES 2007*]

$$C(t_k) = -G(t_k)t_k + p_k(t_{k-1}); k = 1..m$$

Temperature change

Temperature vector at instant k

power consumption vector

Discrete RC-Thermal Estimation Tool for tiers of 3D Chips

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

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

$$t_{k+1} = A(t_k) t_k + B p_k \; ; \; k = 1..m$$

$$A(t_k) = (I - d_t C^{-1} G(t_k)) \; ; \quad B = d_t C^{-1}$$

Time step chosen small enough for convergence

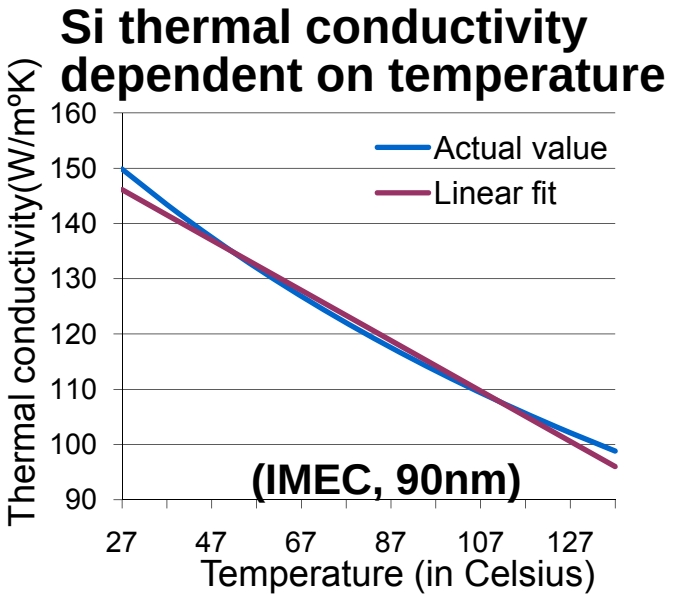

Discrete RC-Thermal Estimation Tool for tiers of 3D Chips

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

- 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 $\dot{t}_k$:

$$t_{k+1} = A(t_k)t_k + Bp_k \; ; \; k =$$

$$A(t_k) = (I - d_t C^{-1} G(t_k)) \; ; \quad B$$

Complexity scales linearly with the number of modeled cells (simulated on Xeon Server)

Thermal library validated against finite element model (IMEC and EPFL)

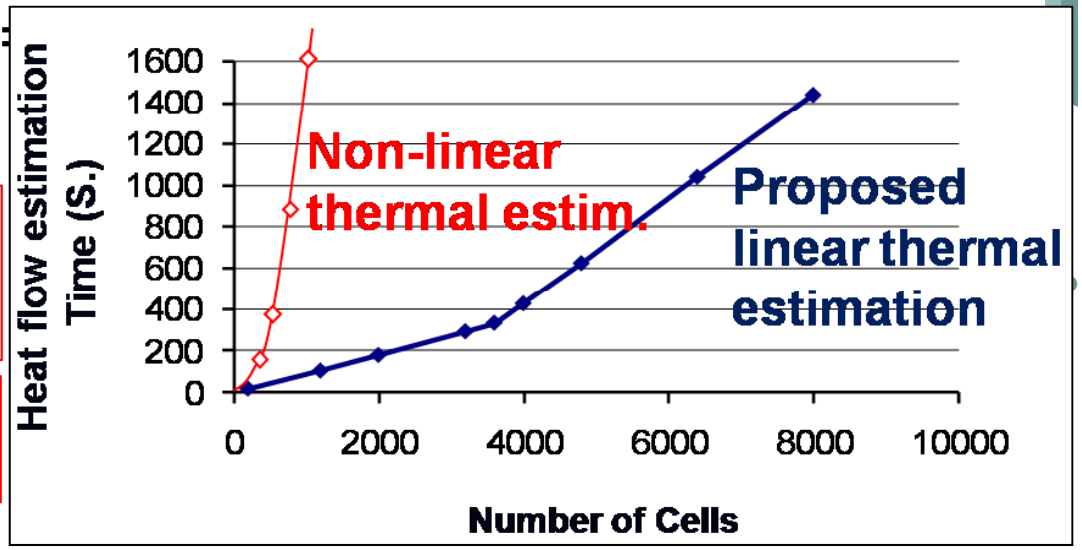

MPSoC Thermal Library Validation

- Extensible set of layers in MPSoC designs
  - Pre-defined material layers and components:
    Silicon, copper (10 layers), packaging, interposer, bumps, etc.
- Configurable nr. of cells and iterations per tier
  - Initially 10ms thermal interval (1000 iterat./tier)
- Test chips manufactured at EPFL:
  - Three types of layouts

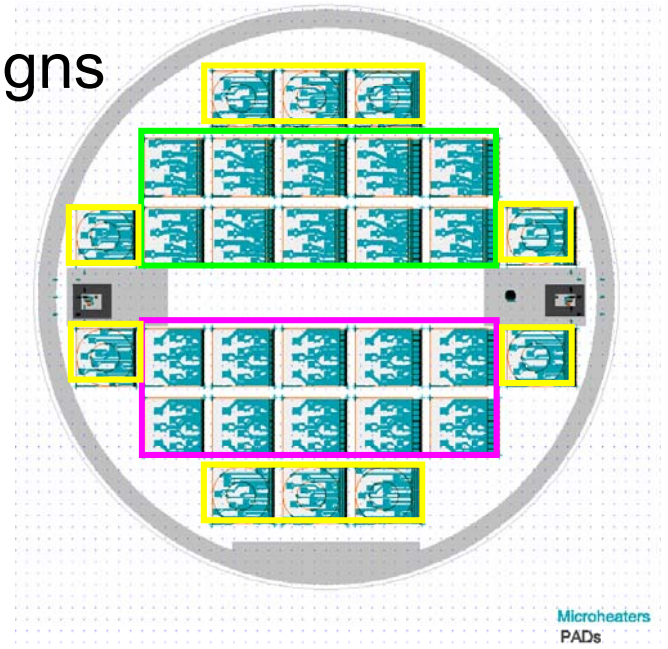

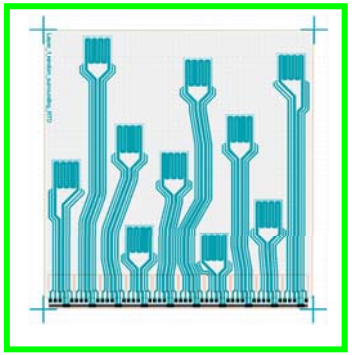

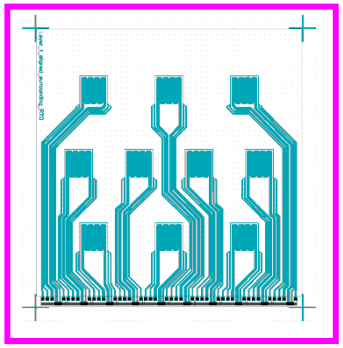

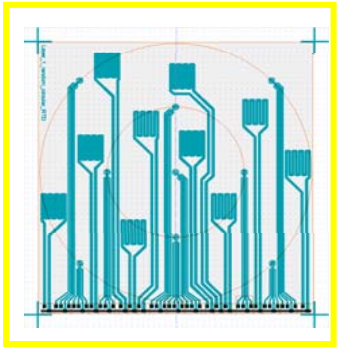

Correlation Results: Intra-Tier Heat Transfer

- Lateral heat flow
  - Tested range: 0.5W to 10W per heater
  - Similar accuracy results at different tiers
- Measurements/simulations in case ~9W:

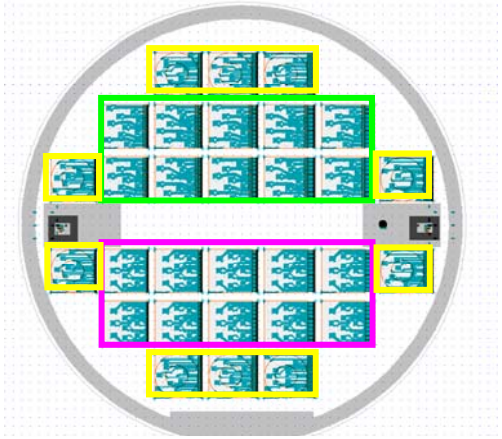

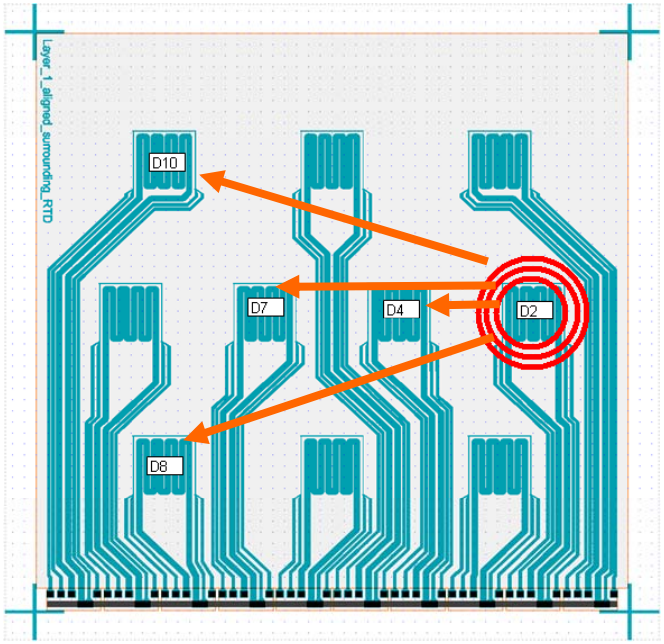

Heater and sensor at the same point (Max temp)

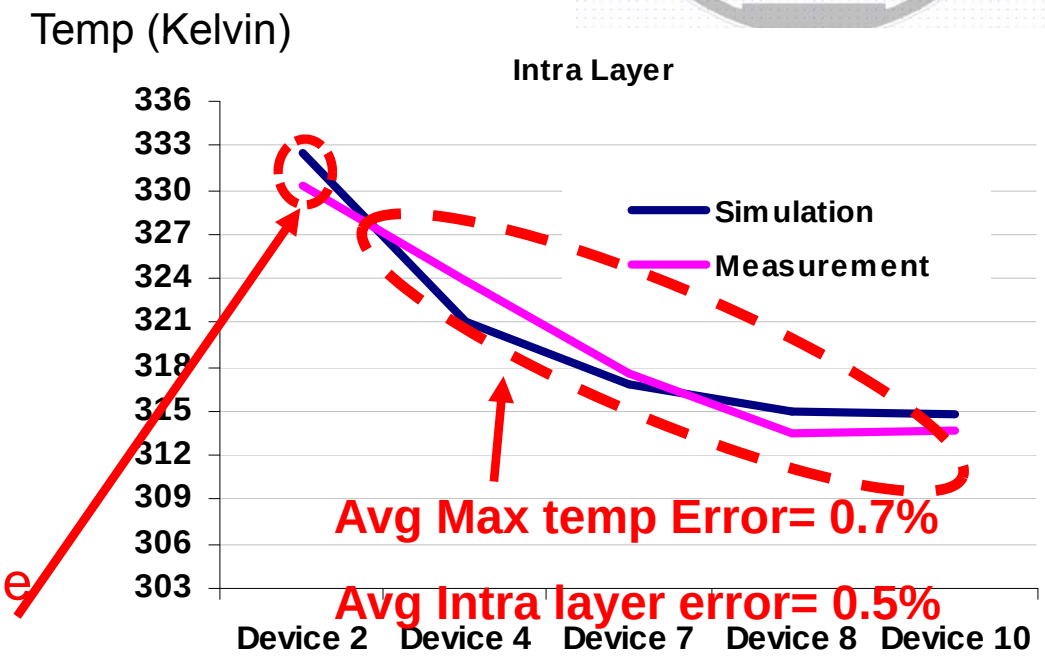

Correlation Results: Intra-Tier Heat Transfer

- Lateral heat flow
  - Tested range: 0.5W to 10W per heater
  - Similar accuracy results at different tiers
- Measurements/simulations in case ~9W:

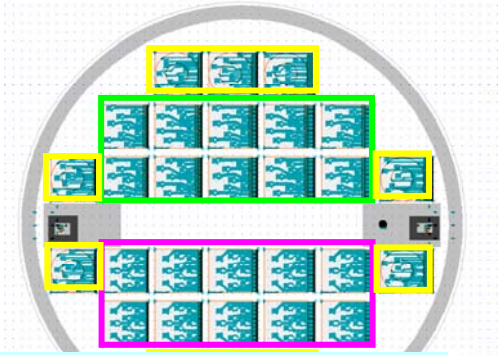

Variations of less than 1.5% between 3D chip measurements and RC-based 3D thermal model

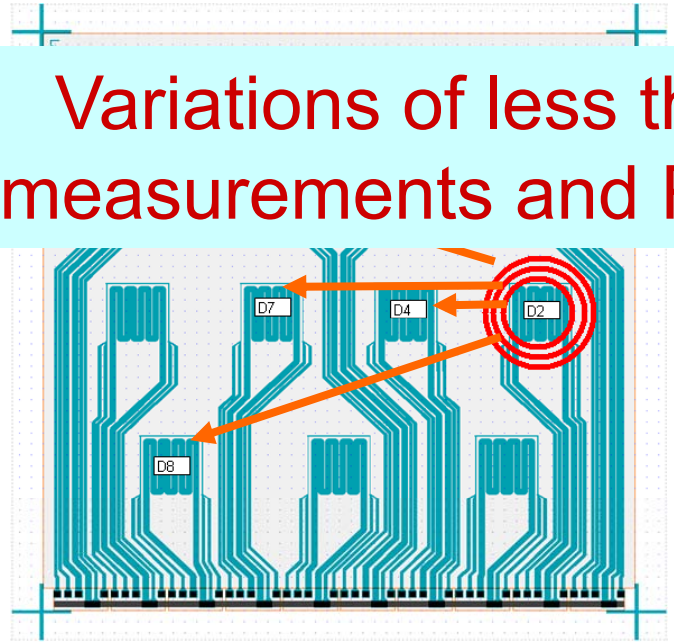

Heater and sensor at the same point (Max temp)

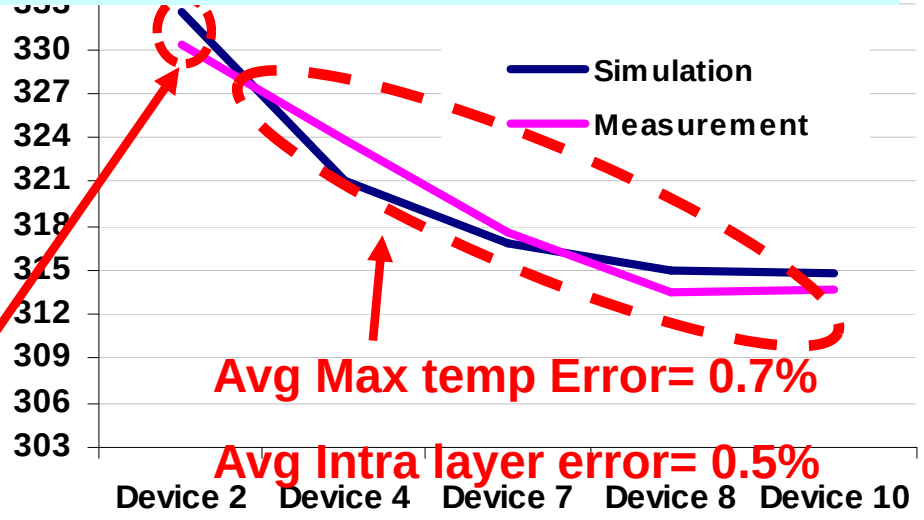

Temperature Management is Power Control under Thermal Constraints

- Power consumption of cores determines thermal behavior
  - Power consumption depends on frequency and voltage
  - Setting frequencies/voltages can control power and temperature
- Optimization problem: frequency/voltage assignment in MPSoCs **under thermal constraints**
  - Meet processing requirements
  - Respect thermal constraint at all times
  - Minimize power consumption

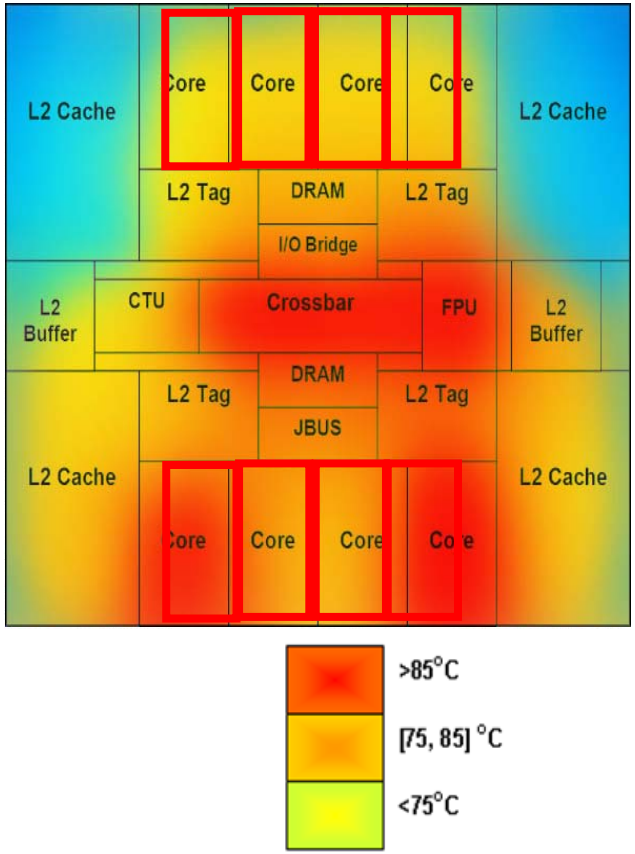

Thermal Management: Initial Thoughts

- Static approach: thermal-aware placement to try to even out worst-case thermal profile [Sapatnekar, Wong et al.]
  - Computationally difficult problem (NP-complete)

Not able to predict all working conditions, and leakage changing dynamically, not useful in real systems

No formalization of thermal optimization problem!

- Dynamic approach: HW-based dynamic thermal management
  - Clock gating based on time-out [Xie et al., Brooks et al.]
  - DVFS based on thresholds [Chaparro et al, Mukherjee et al,]
  - Heuristics for component shut down, limited history [Donald et al]

Techniques to minimize power, they only achieve thermal management as a by-product…

Formalization of Thermal Management Problem in MPSoCs

- Control theory problem
  - Obse
    - H
    - P
  - Cont
    - Tuning knobs: frequencies/voltages of the system

Optimal frequency assignment module, 2-phase approach:

1) **Design-time phase:** Find optimal sets of frequencies for the cores for different working conditions

2) **Run-time phase:** Apply one of the predefined sets found in phase 1 for the required system performance

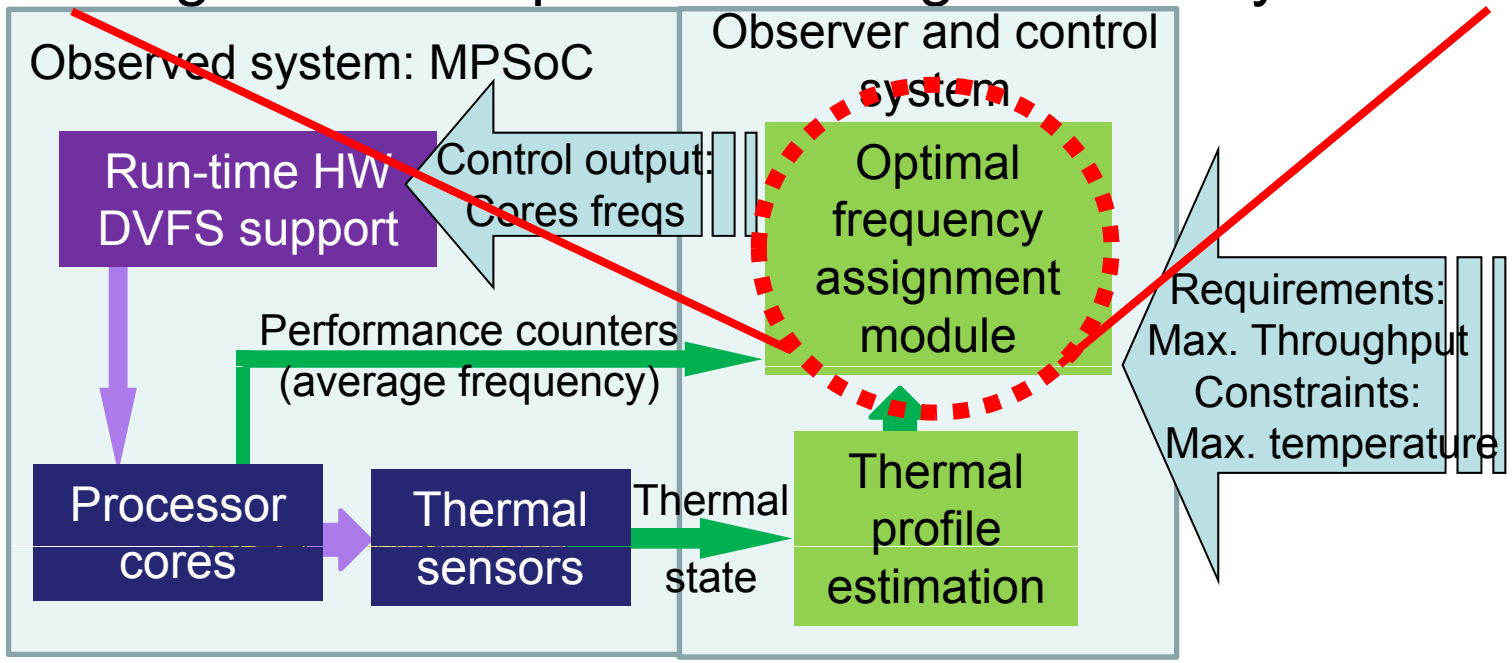

Pro-Active Based Thermal Control: Phase 1 – Design-Time

- Predictive model of thermal behavior given a set of frequency assignments

Allowed core power values and frequencies

Chip floorplan

Packaging, heat spreader information

Phase inputs

Non-linear offline problem

ze sum of power consumption of cores

Constraints:

$$\sum_{k=1}^{m} 1^T f_k \geq m \times n \times f_{avg}$$

Performance constraint: on average, freq. is f_{avg}

$$t_{k+1} = A(t_k)t_k + Bp_k, \quad k = 1, \ldots, m$$

Thermal equation

$$t_k \leq t_{\max}, \quad k = 1, \ldots, m$$

Meet temp. constraints at all time points

$$p_{\max} f_{i,k}^2 / f_{\max}^2 = p_{i,k}, \quad i = 1, \ldots, n, \quad \forall k$$

Power equation based on frequency

$$f_{\min} \leq f_k \leq f_{\max}, \quad k = 1, \ldots, m.$$

Frequency in predefined range

Method

outputs

Table of cores frequencies assignments

Pro-Active Based Thermal Control: Phase 1 – Design-Time

- Predictive model of thermal behavior given a set of frequency assignments

Allowed core power values and frequencies

Chip floorplan

Packaging, heat spreader information

Phase inputs

Non-linear offline problem

ze sum of power consumption of cores

Constraints: $\sum_{k=1}^{m} 1^T f_k \geq m \times n \times f_{avg}$

$$t_{k+1} = A(t_k)t_k + Bp_k, \quad k = 1, \ldots, m$$

Thermal equation: Si conductivity depends on temp

$$t_k \leq t_{\max}, \quad k = 1, \ldots, m$$

$$p_{\max} f_{i,k}^2 / f_{\max}^2 = p_{i,k}, \quad i = 1, \ldots, n, \quad \forall k$$

Power equation: quadratic dependence on freq.

$$f_{\min} \leq f_k \leq f_{\max}, \quad k = 1, \ldots, m.$$

Method

outputs

Table of cores frequencies assignments

Making Power and Thermal Constraints Convex

- Power constraint adaptation
  - Change non-affine (quadratic equality):

$$p_{max} (f_{i,k})^2 / (f_{max})^2 = p_{i,k} ; i = 1,..,n, \forall k$$

  - To convex inequality:

$$p_{max} (f_{i,k})^2 / (f_{max})^2 \leq p_{i,k} ; i = 1,..,n, \forall k$$

- Thermal constraint adaptation
  - Use worst case thermal conductivity in the range of allowed temperatures, and iterate (if needed) to optimum

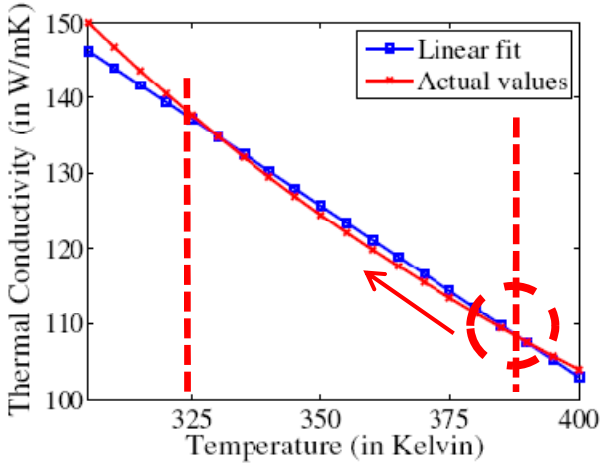

Making Power and Thermal Constraints Convex

- Power constraint adaptation

Solve convex problem and get table of optimal frequencies for different working conditions in polynomial time (number of processors)

Required average frequencies	Starting Temperatures <= 30 °C	35 °C	▪ ▪ ▪	100 °C
<= 100 MHz	<120,80,80,120>			
150 MHz				
▪ ▪ ▪				
1000 MHz				

Pro-Active Based Thermal Control: Phase 2 - Run-Time, Putting It All Together

- Use table of frequencies assignments and index by actual conditions at regular run-time intervals

Targeted operating frequency of cores

Current temperature of cores

Method inputs

Run-time optimal DVFS assignment module

1) Index table output of pl
with current working conc

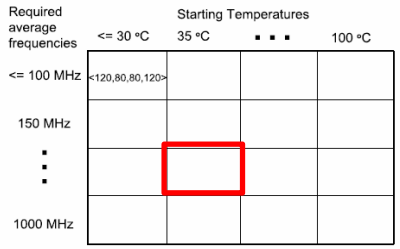

2) Compare to current assignment to cores and generate required signaling to modify DVFS values

Phase output

Run-time DVFS changes for processors

Case Study: 8-Core Sun MPSoC

- MPSoC Sun Niagara architecture
  - 8 processing cores SPARC T1
- Max. frequency each core: 1 GHz
  - 10 DVFS values, applied every 100ms
- Max. power per core: 4 W
- Execution characteristics of workloads [Sun Microsystems]:
  - Mixes of 10 different benchmarks, from web-accessing to multimedia
  - 60,000 iterations of basic benchmarks, tens of seconds of actual system execution

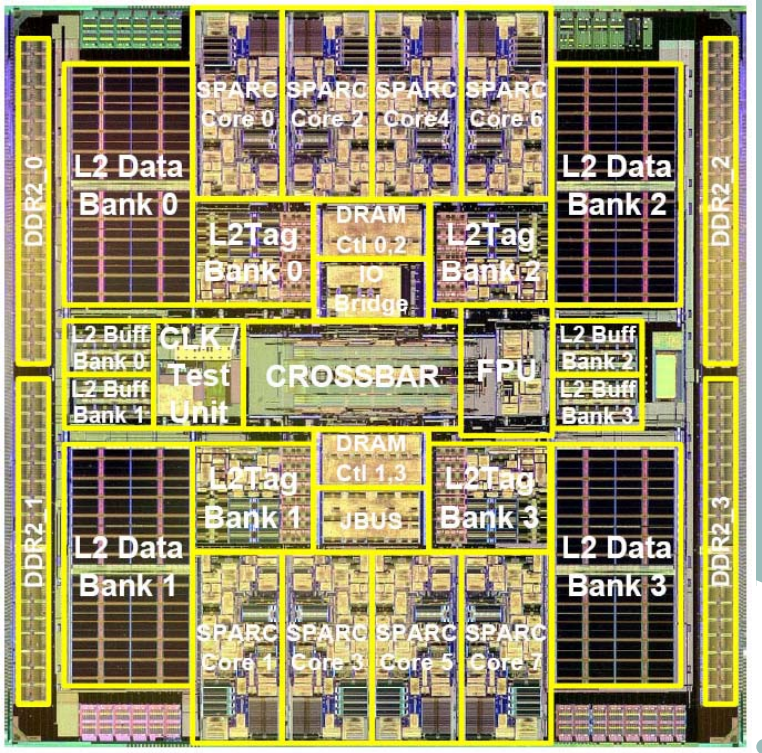

Sun's Niagara MPSoC

Thermal Constraints Respected… And Faster overall!

DVFS:

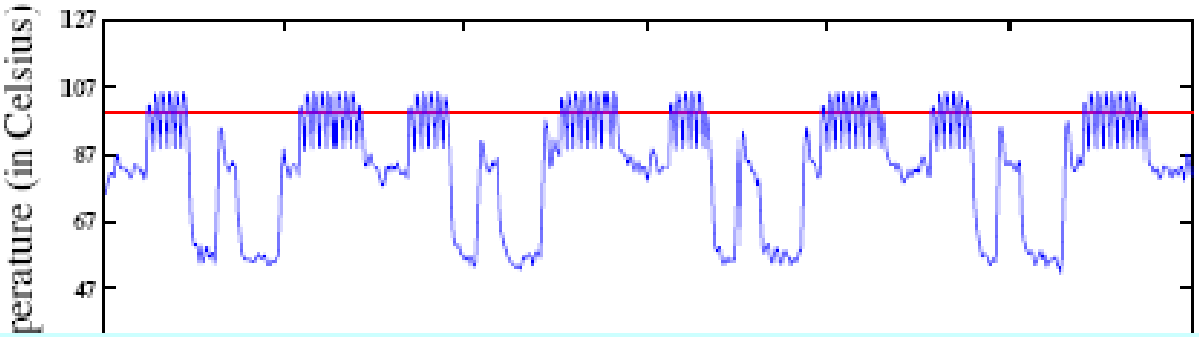

Total run-time of benchmarks

180 sec

Proposed method achieves better throughput than standard DVFS while satisfying thermal constraints

2-phase Convex method:

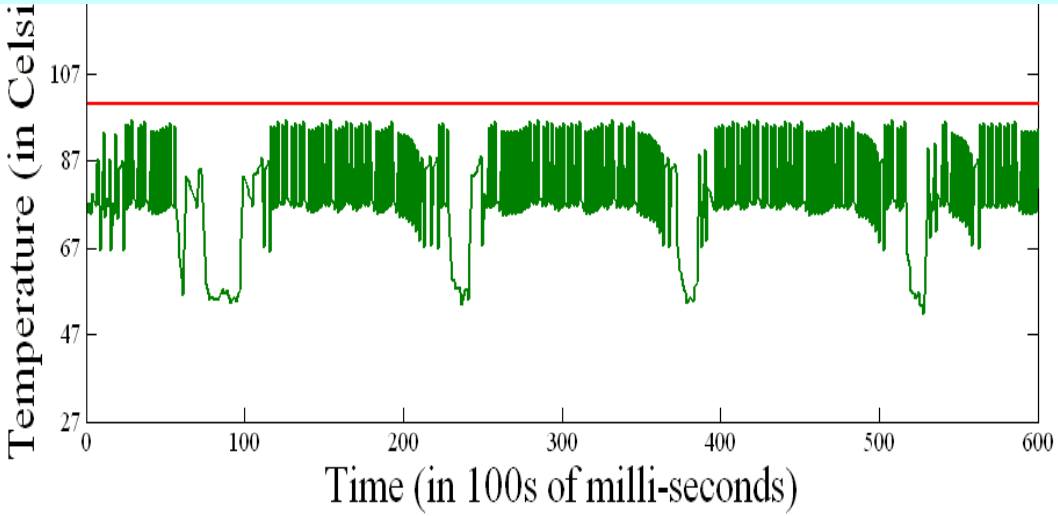

106 sec (45% less exec. time)

The Big Picture for Thermal-Aware Design: Large-Scale Computing in Datacenters

- Area is expensive, we try to get denser infrastructures
  - New containers: many more servers each, >10x density

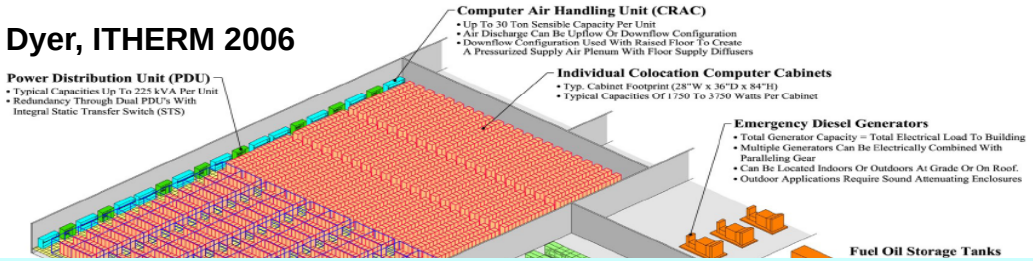

45% of energy overhead in cooling, how to get higher computational densities (with lower cooling costs)?

- **Air-cooled datacenters** are **very inefficient**
  - Cooling needs as much energy as IT… and thrown-away
- For a 10MW datacenter **~US$ 4M wasted** per year

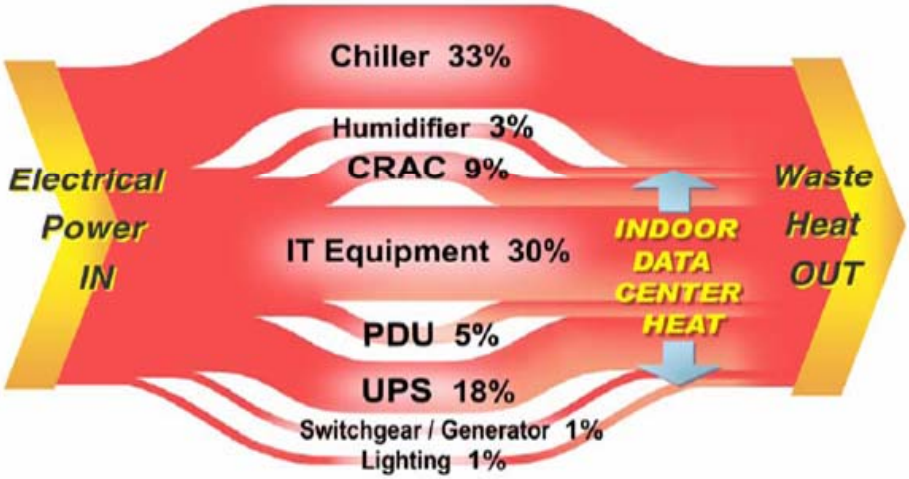

Datacenter energy overhead, ASHRAE

Processing Trends: Single to Multi-Core

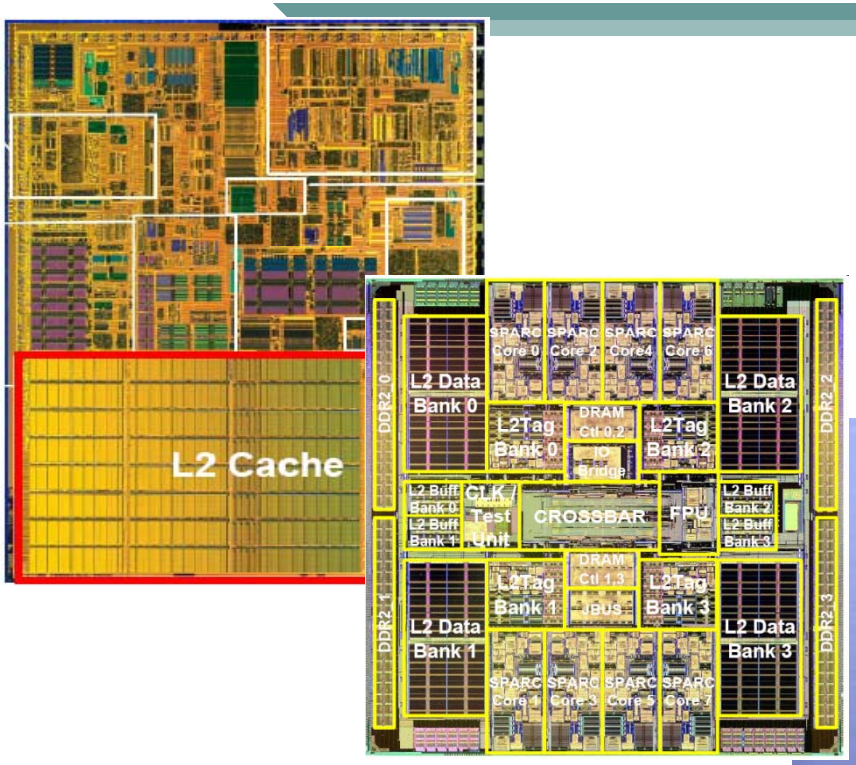

How do you feed so many cores?
Memory bandwidth wall!

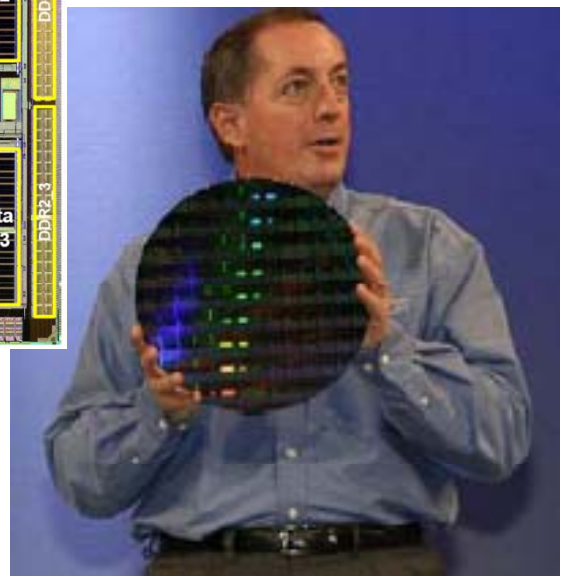

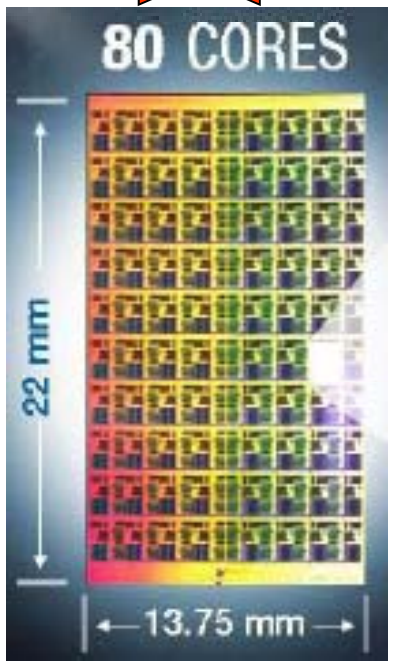

[Courtesy: Yuan Xie, ICCAD 2010]

Why not using 3rd Dimension?

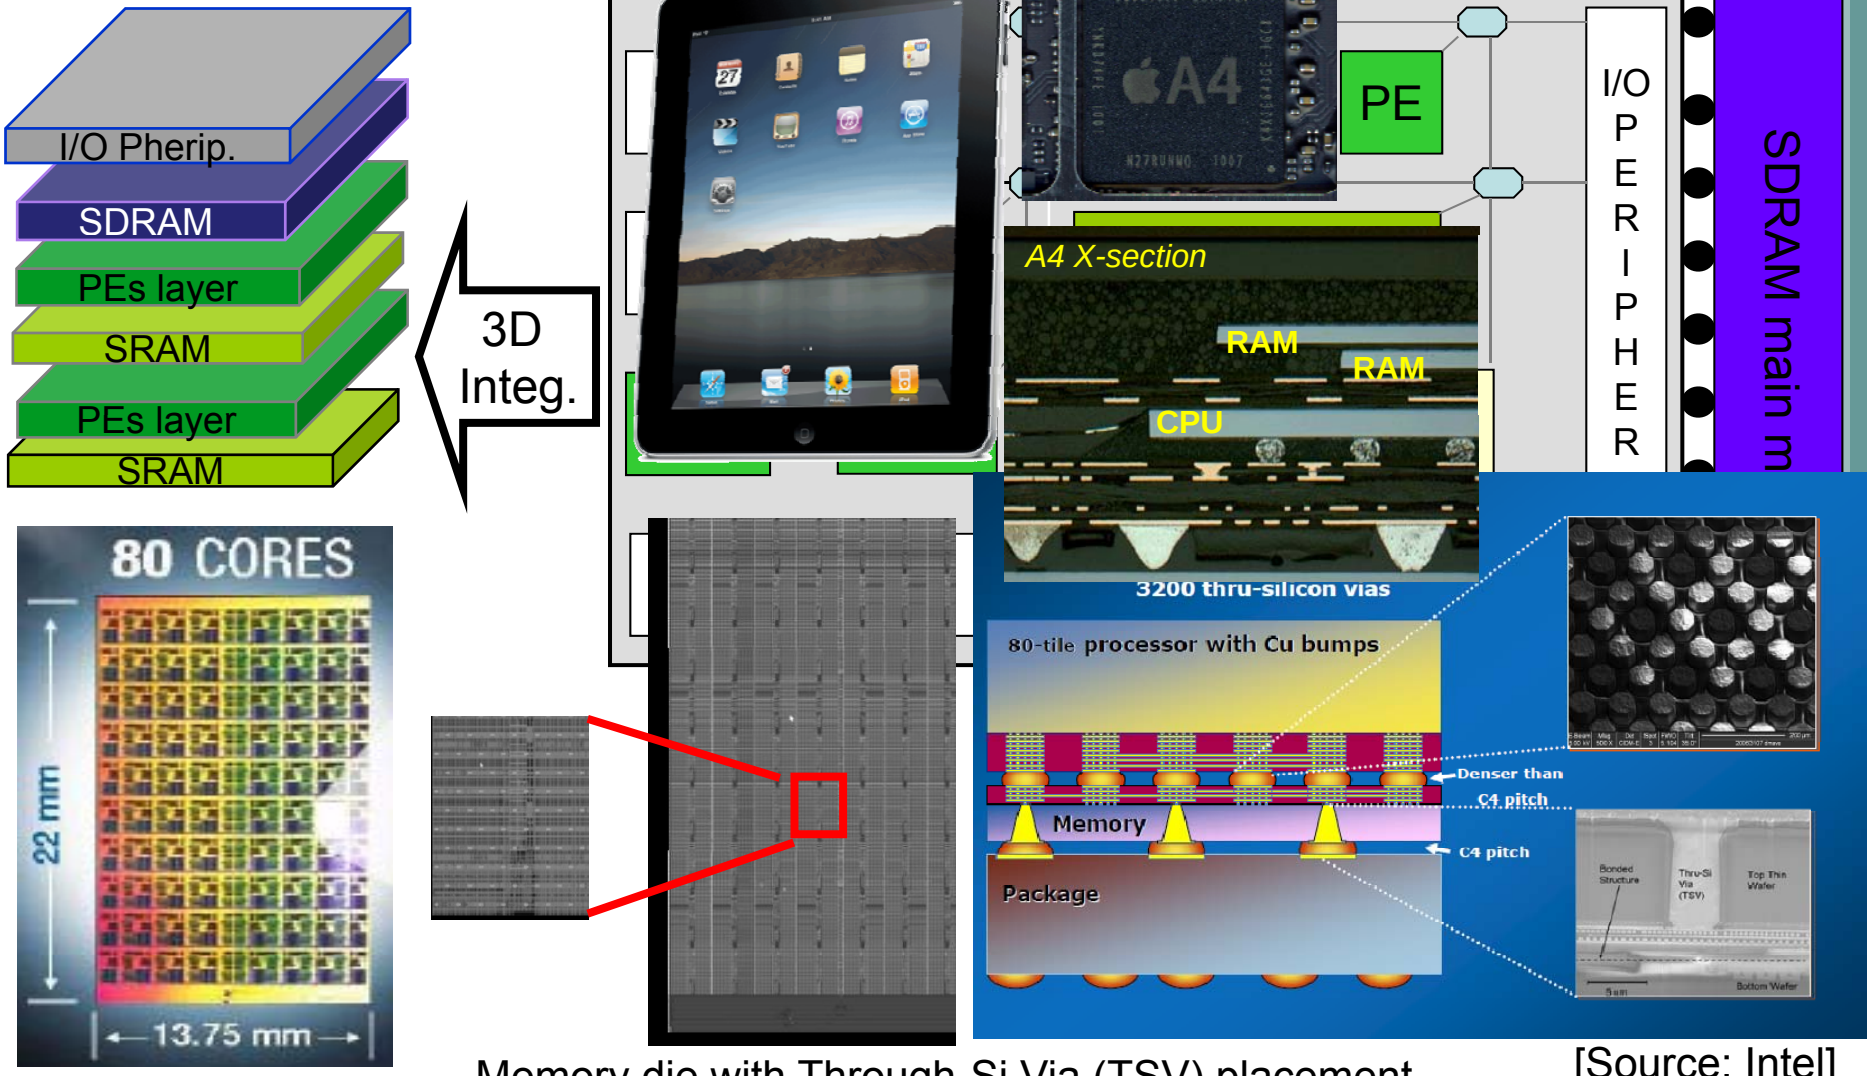

Memory die with Through-Si Via (TSV) placement

[Source: Intel]

Run-Time Heat Spreading in 3D MPSoCs: More Complex Cooling Needs!

- 5-tier 3D stack: 10 heat sources and sensors

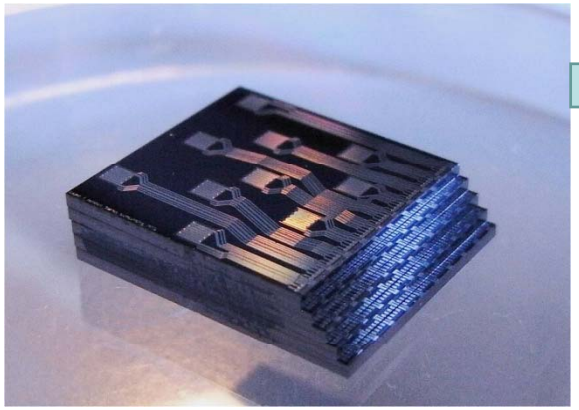

[Garcia and Atienza, *Microelectornics Journal 2010*]

Inject between 4W – 1.5W

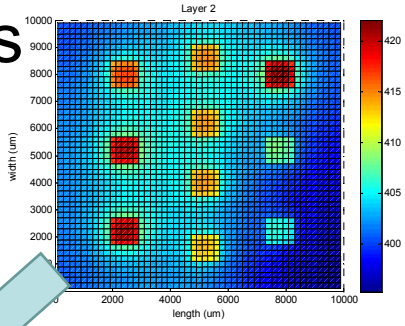

2nd Tier

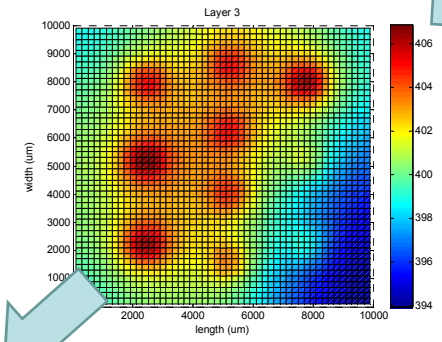

3rd Tier

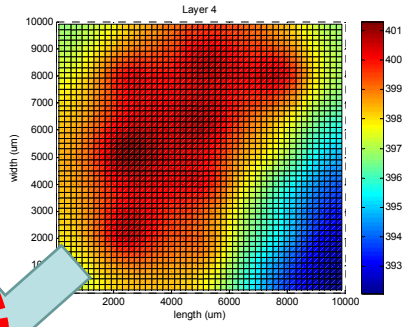

4th Tier

5th Tier

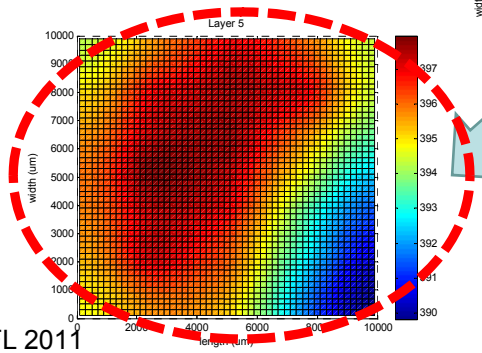

Large and non-uniform heat propagation! (up to 130º C on top tier)

Zero-Emission Datacenter: Liquid Cooling Technology and Predictive Energy Management

- Datacenters are “intelligent” heaters
  - 30-40% of carbon footprint in Europe using district heating networks
- Direct re-use of heat output
  - 3D MPSoC architectures

Aquasar datacenter server: 80% payback of electricity costs

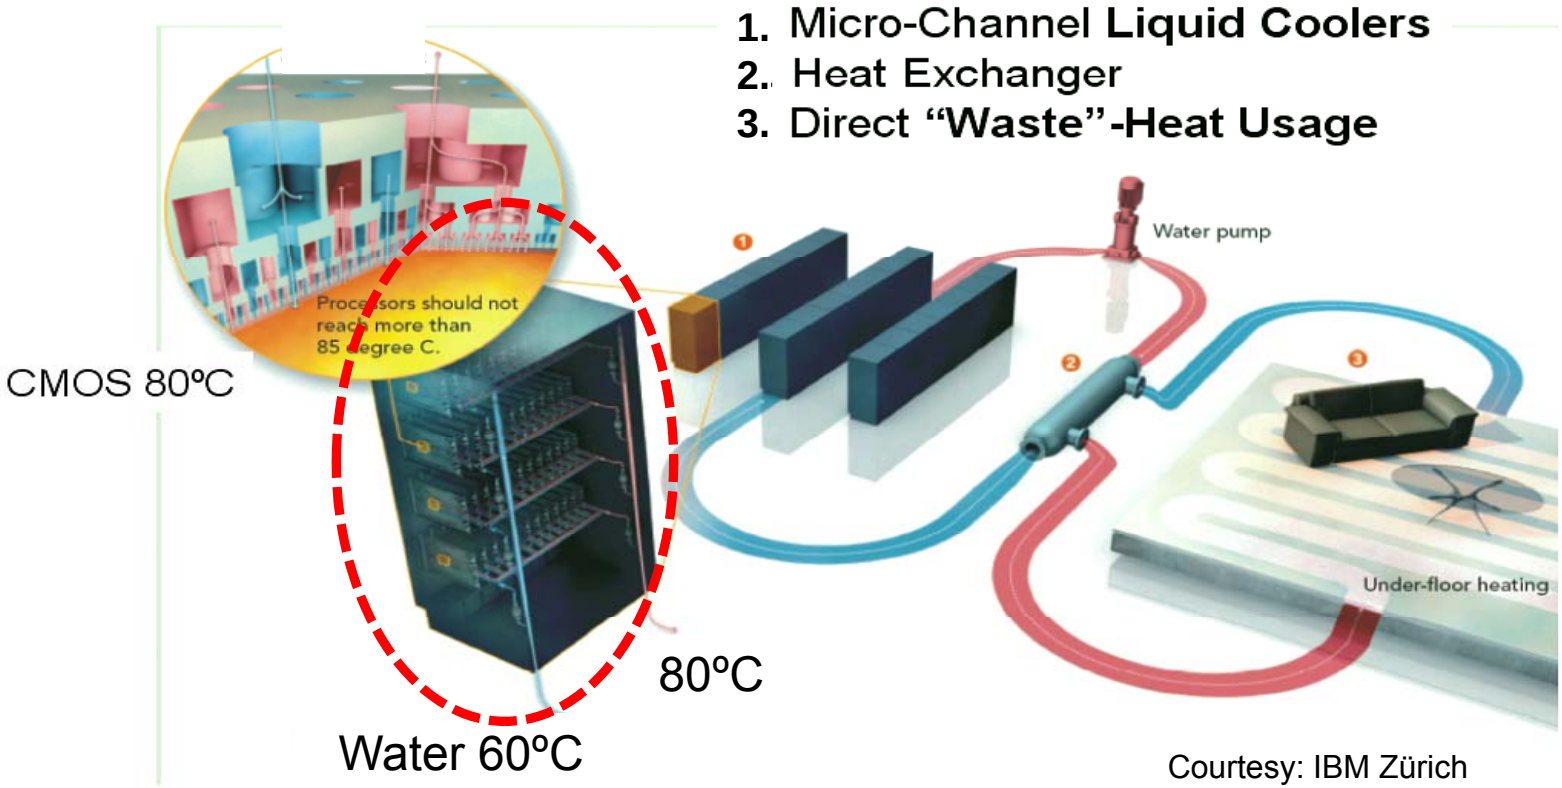

Courtesy: IBM Zürich

Zero-Emission Datacenter: Liquid Cooling Technology and Predictive Energy Management

- Datacenters are “intelligent” heaters
  - 30-40% of carbon footprint in Europe using district heating networks
- Direct re-use of heat output
  - 3D MPSoC architectures

Aquasar datacenter server: 80% payback of electricity costs

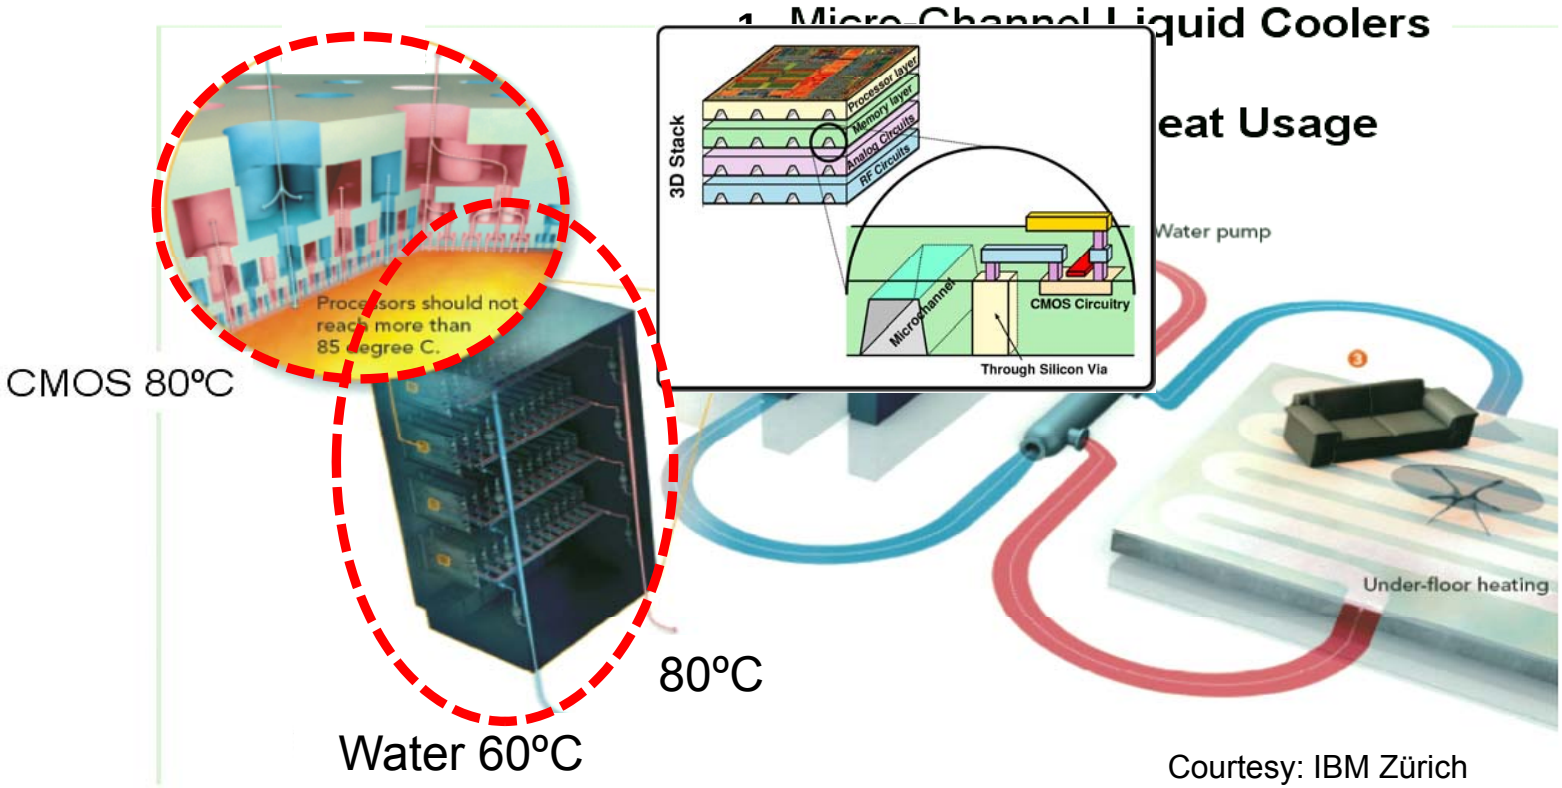

Courtesy: IBM Zürich

Zero-Emission Datacenter: Liquid Cooling Technology and Predictive Energy Management

- Datacenters are “intelligent” heaters
  - 30-40% of carbon footprint in Europe using district heating networks
- Direct re-use of heat output
  - 3D MPSoC architectures

Aquasar datacenter server: 80% payback of electricity costs

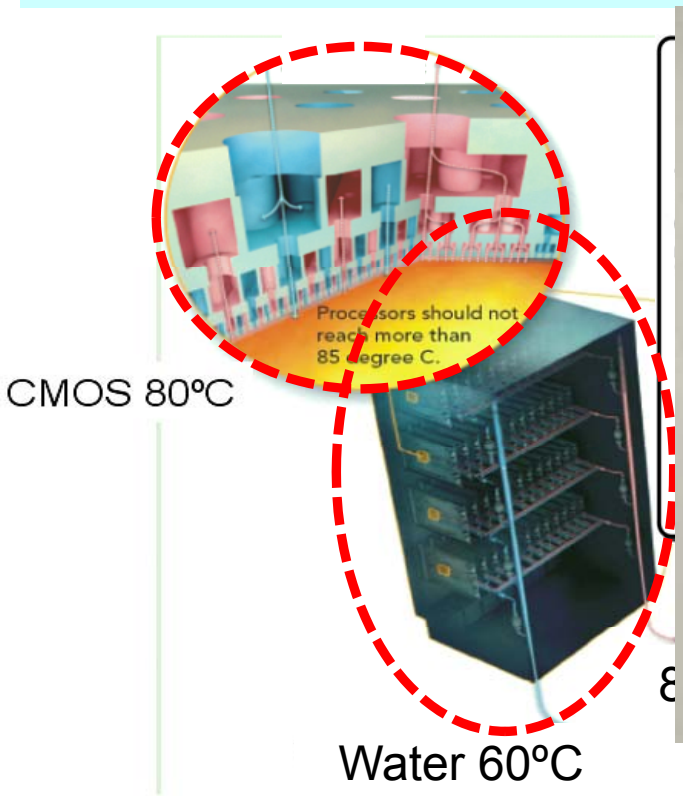

Courtesy: IBM Zürich

NanoTera CMOSAIC Project: Design of 3D MPSoCs with Advanced Cooling

- 3D systems require novel electro-thermal co-design
  - Academic partners: EPFL and ETHZ
  - Industrial: IBM Zürich and T.J. Watson

3D MPSoC datacenter chip: microchannels etched on back side to circulate (controlled) liquid coolant

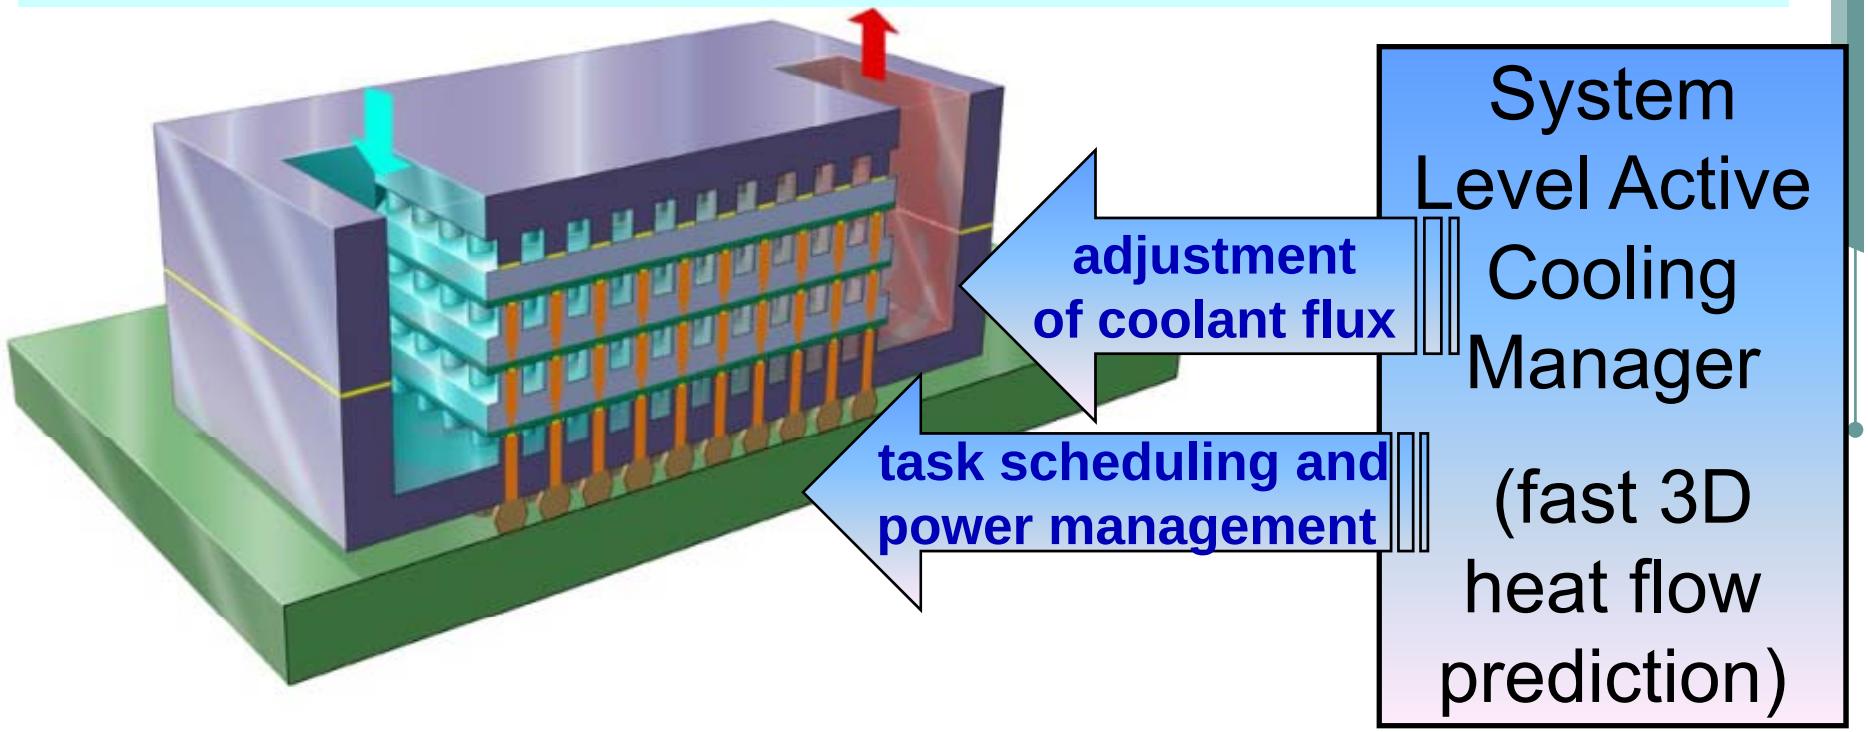

Creating a Fast Thermal Model: Compact RC-Based Stack Model with TSVs

- Chip-Level thermal model

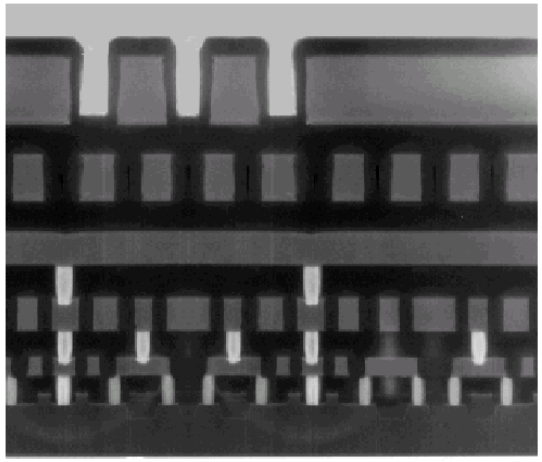

RC Network of

Si/metal layer cells

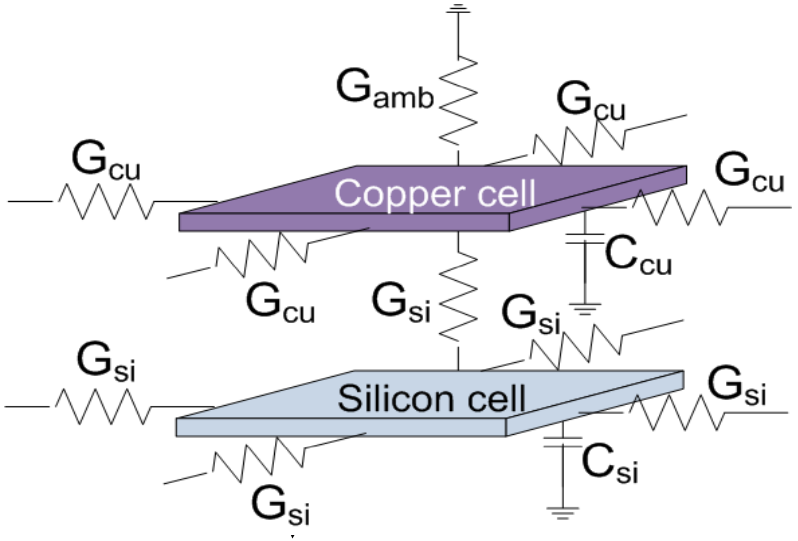

- Convective boundary conditions between layers **for each tier separately**

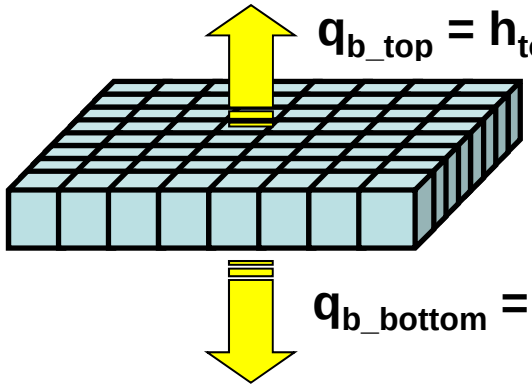

- TSVs change resistivity of interlayer material

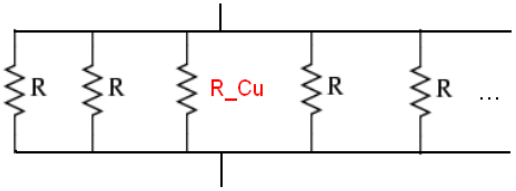

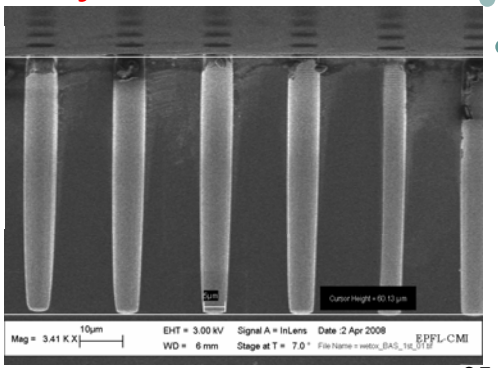

3D MPSoC Thermal Library Deployment

- Extensible set of layers in 3D stack
  - Up to 9 tiers and heat spreader
  - Pre-defined layers:
    - Silicon, copper (10 layers), glue, overmold, interposer, bump
- Configurable nr. of cells and iterations per tier
  - Also 10ms thermal interval (1000 iterat./tier)
- Multi-tier test chip manufactured at EPFL:

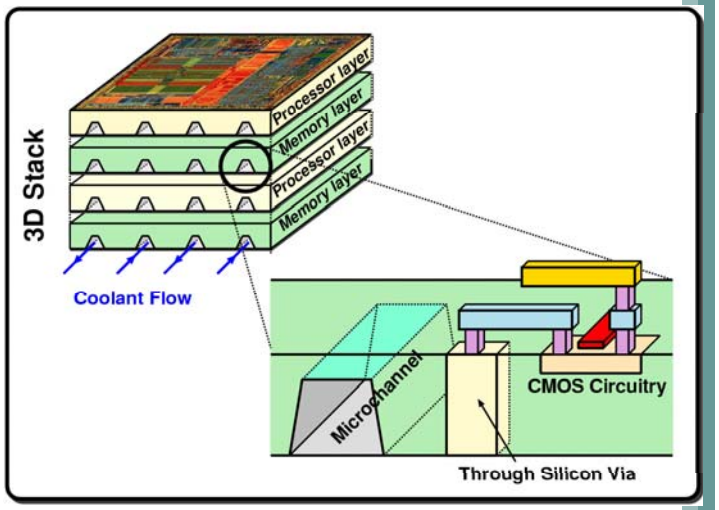

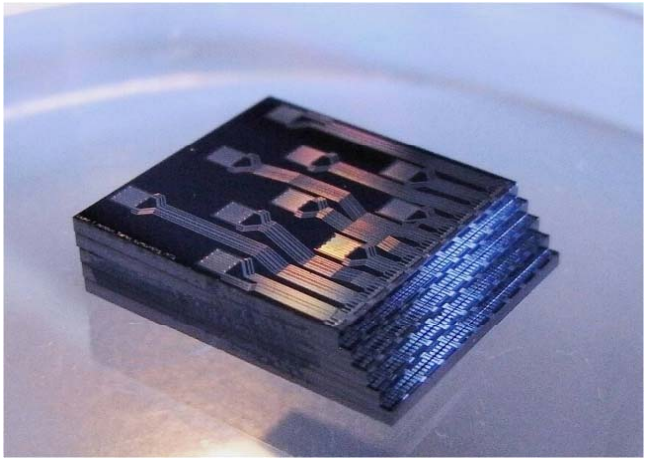

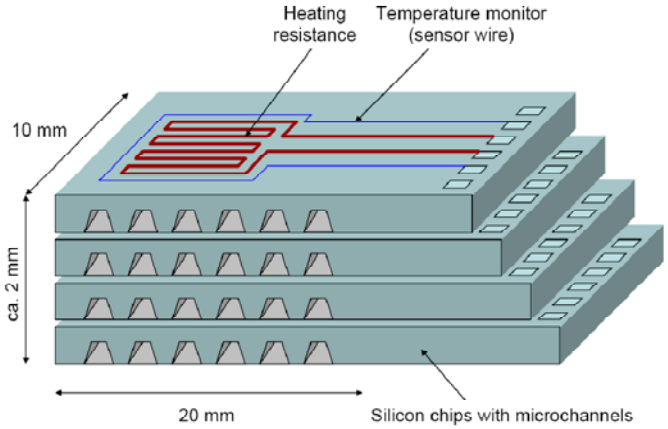

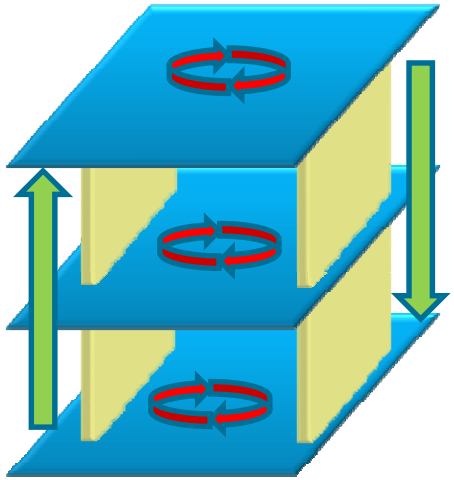

Correlation Results: Inter-Tier Heat Transfer

- Vertical heat flow (multi-level measurements)
  - Tested range: 0.5W to 10W per heater
  - Variations only of global temperatures trend
  - Tier-2 measurements/simulations (~9W):

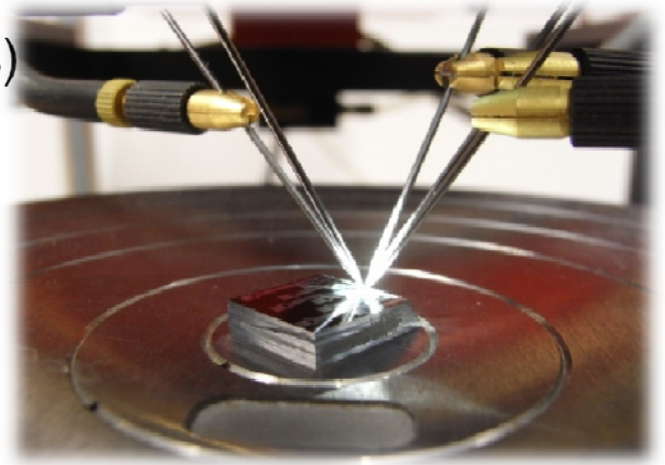

Inter Layer

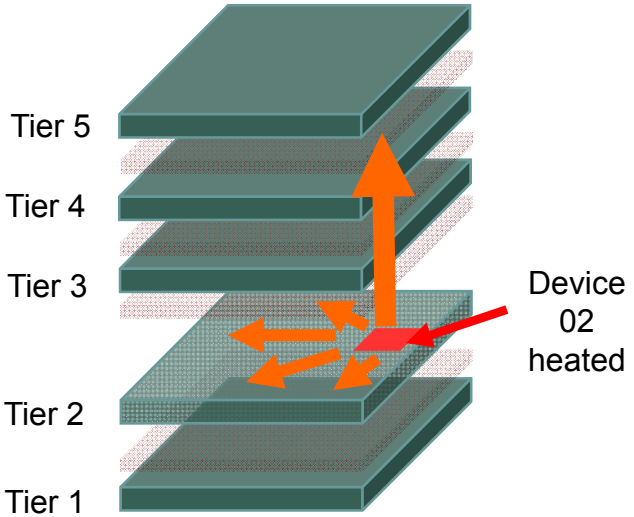

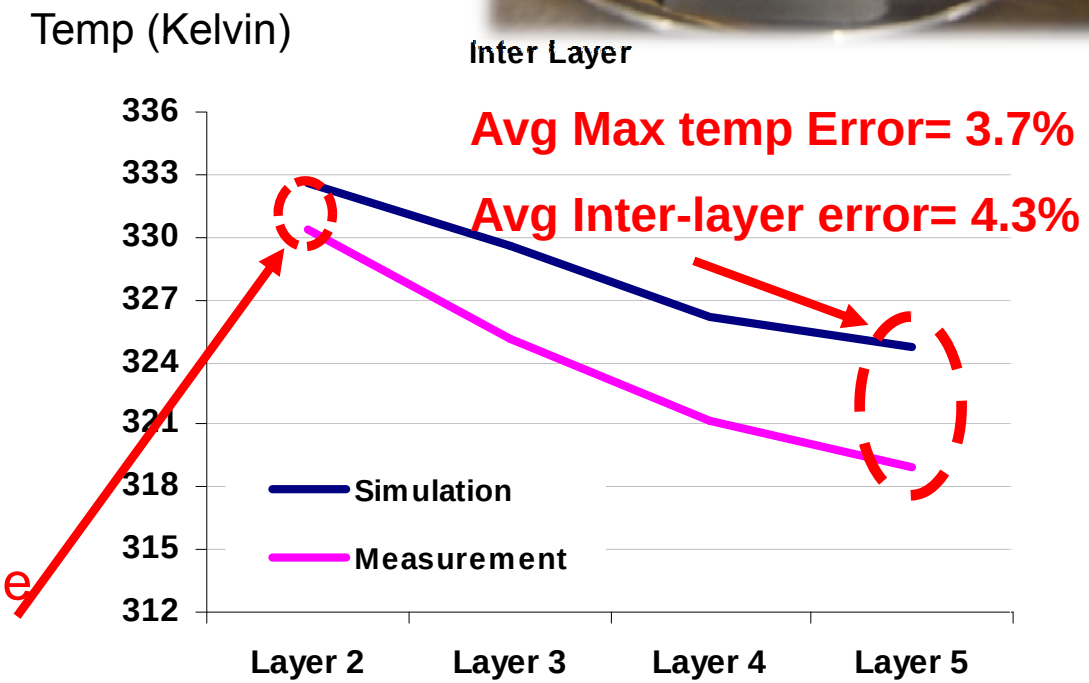

Heater and sensor at the same point (Max temp)

Correlation Results: Inter-Tier Heat Transfer

- Vertical heat flow (multi-level measurements)
  - Tested range: 0.5W to 10W per heater
  - Variations only of global temperatures trend
  - Tier 2 measurements/simulations (~9W):

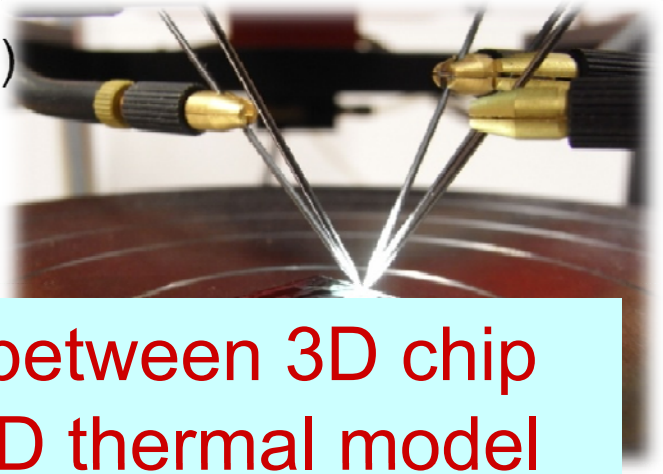

Variations of approximately 5% between 3D chip measurements and RC-based 3D thermal model

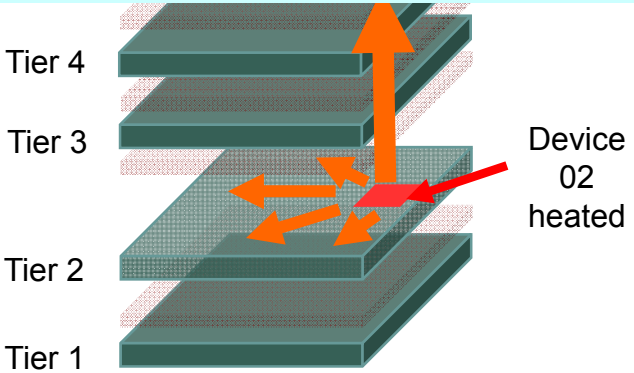

Heater and sensor at the same point (Max temp)

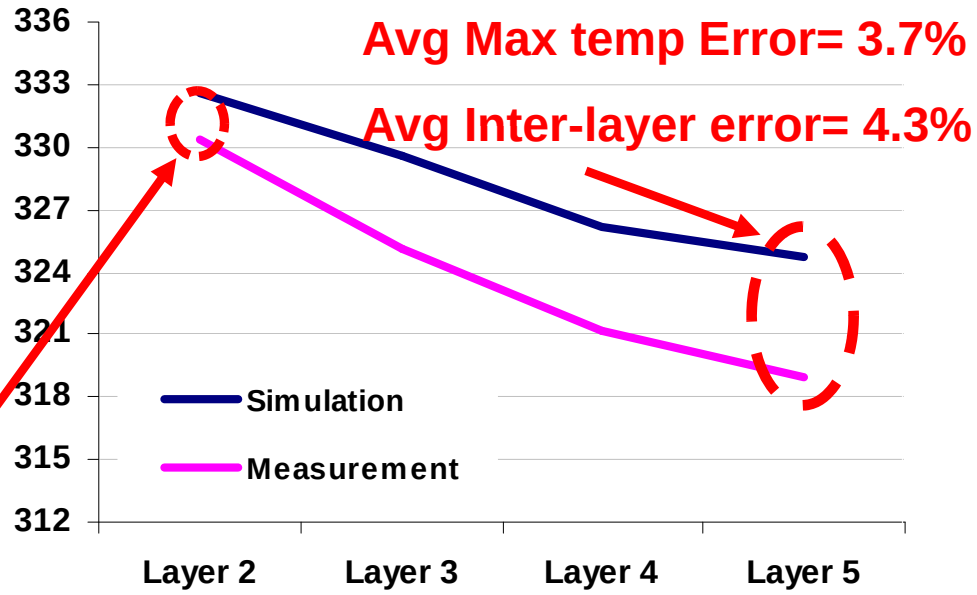

Modeling Liquid Cooling as RC-Network in 3D MPSoC stacks

- Local junction temperature modeled as 4-resistor based compact transient thermal model (4RM-based CTTM)
  - Rtot = Rcond + Rconv + Rheat

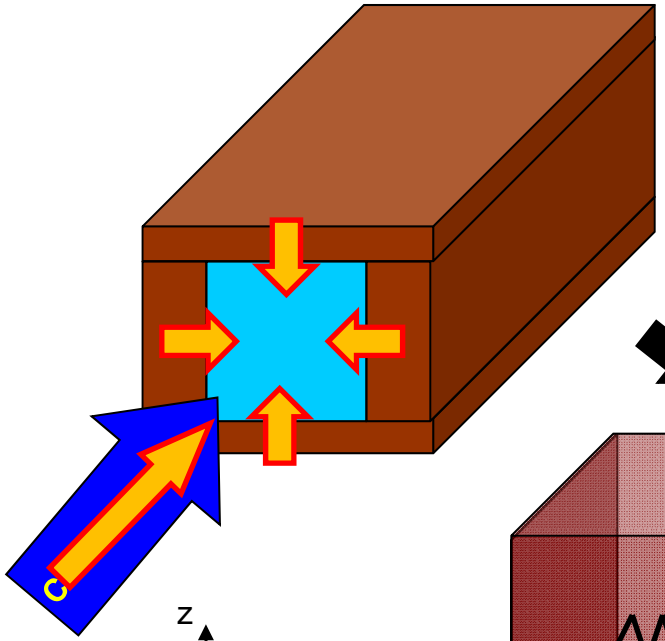

[Sridhar, et al., *ICCAD 2010*]

3D-ICE: 3D IC Emulator with liquid cooling
http://esl.epfl.ch/3d-ice.html

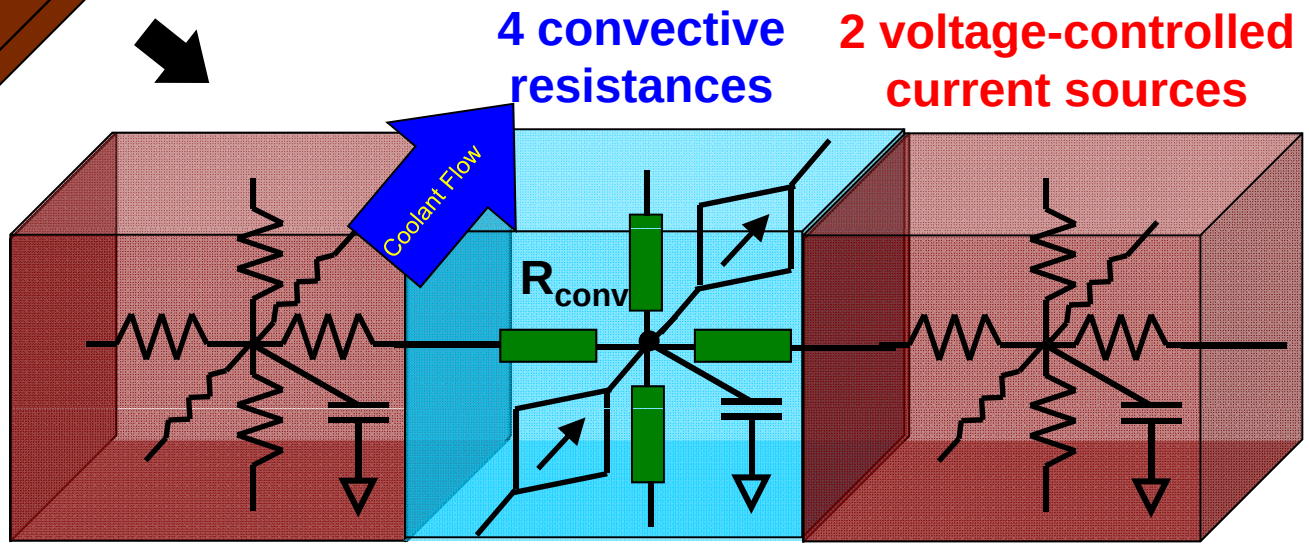

Manufacturing of 5-Tier 3D Chips with Liquid Channels in Multiple Tiers

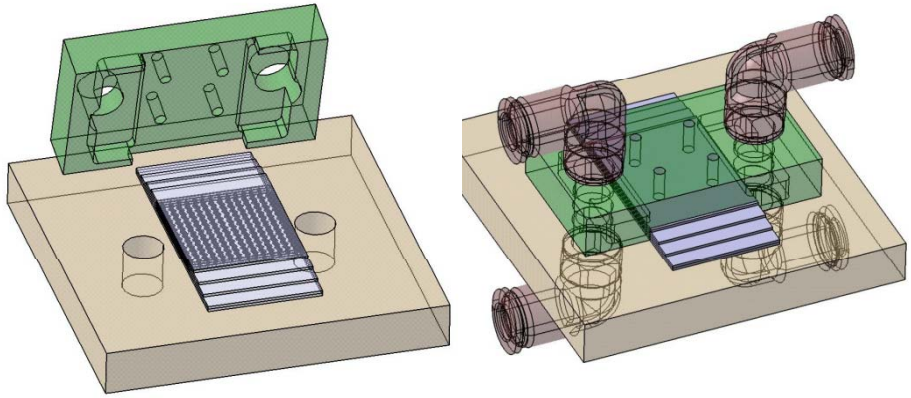

Adding multi-tier liquid cooling in-/out-lets

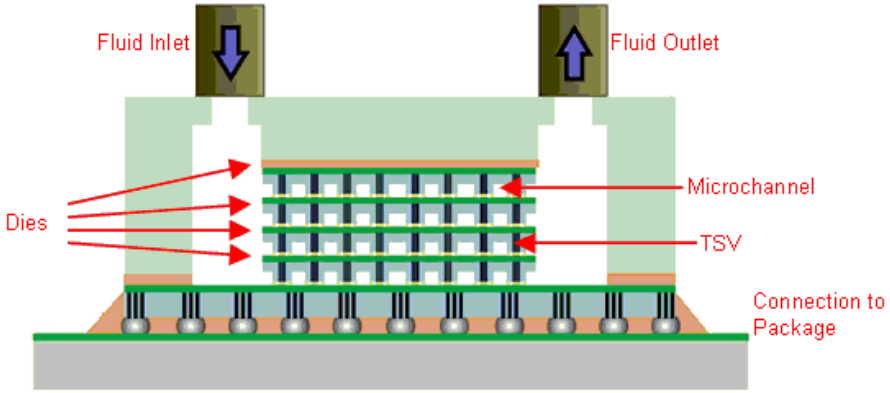

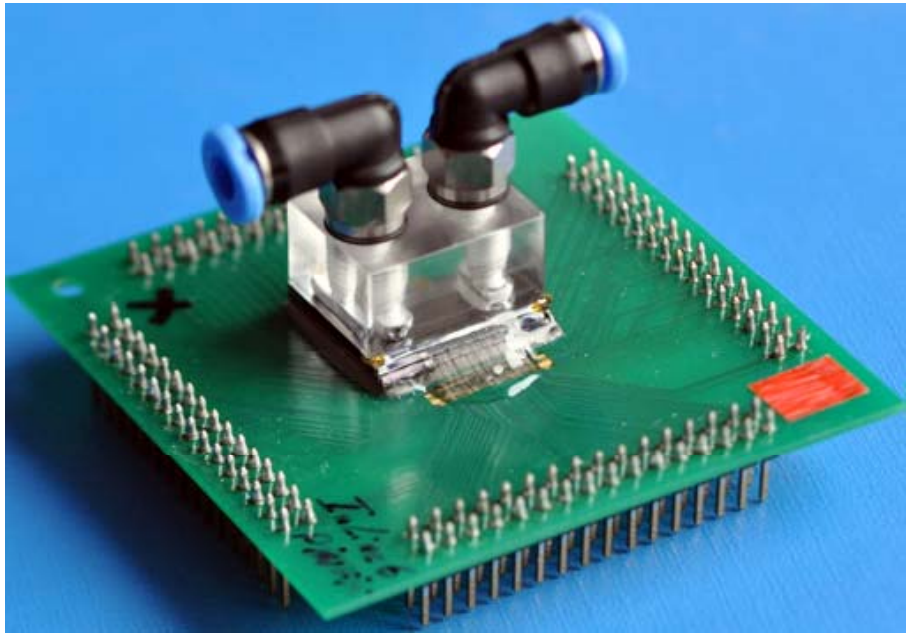

Manufacturing of 5-Tier 3D Chips with Liquid Channels in Multiple Tiers

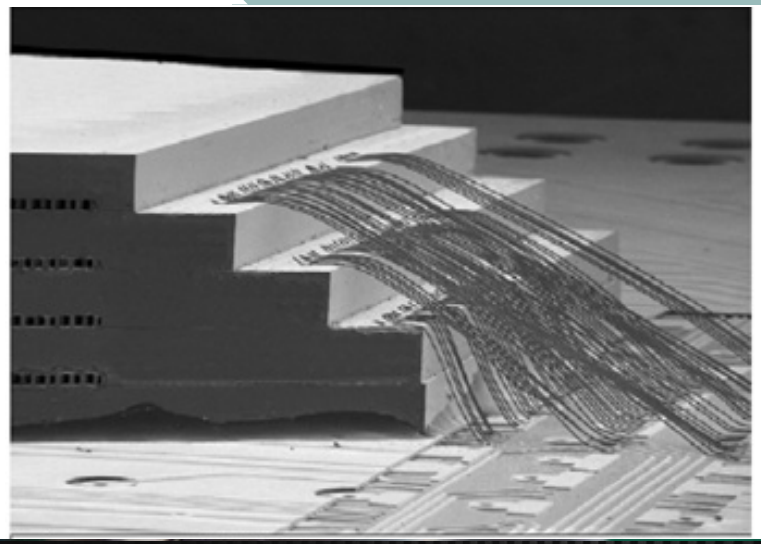

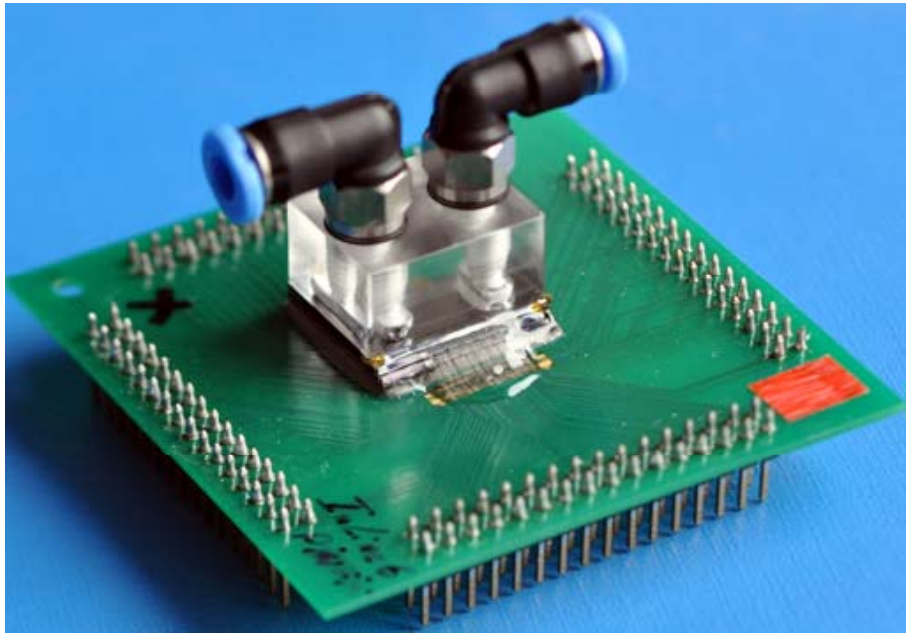

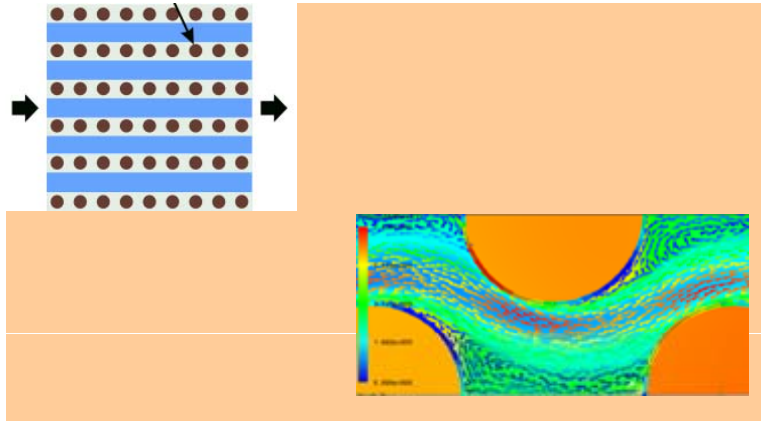

Manufacturing of 5-Tier 3D Chips with Liquid Channels in Multiple Tiers

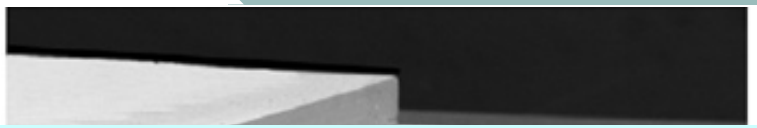

Maximum error of 3.5% between measurements and RC-based 3D thermal model with liquid cooling

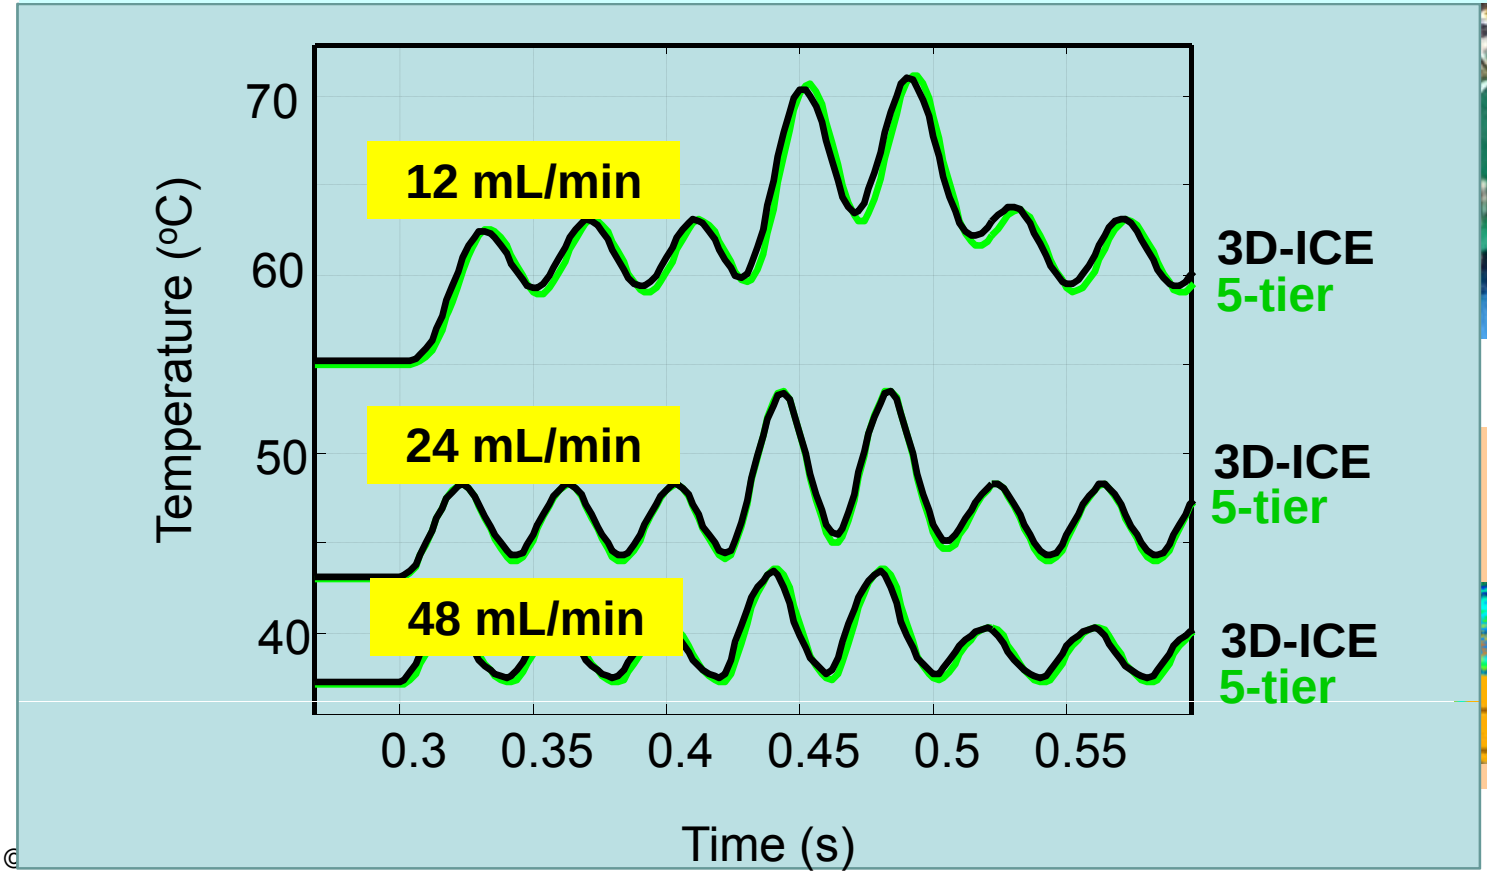

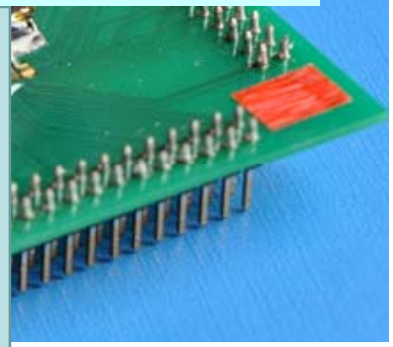

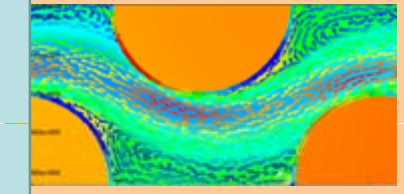

Active-Adapt3D: Active cooling management for 3D MPSoCs

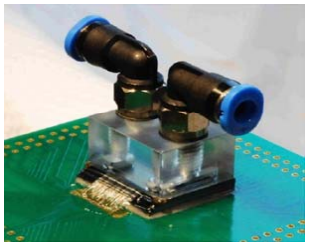

- 3D MPSoC temperature control at system-level:
  - **Electrical based**: task scheduling, and DVFS (µsec or few ms)
  - **Mechanical based**: run-time varying flow rate (hundreds of ms)
- *Fuzzy logic-based controller and thermal-aware scheduler*
  1. **Design-time analysis**: extraction of set of thermal management rules
  2. **Run-time thermal management**: utilization of rules in scheduler and subsequently fuzzy logic controller using both mechanical and electrical methods to achieve:
     - Thermal balance
     - Energy efficiency

[Coskun , Atienza, et al., *MICRO 2011*]

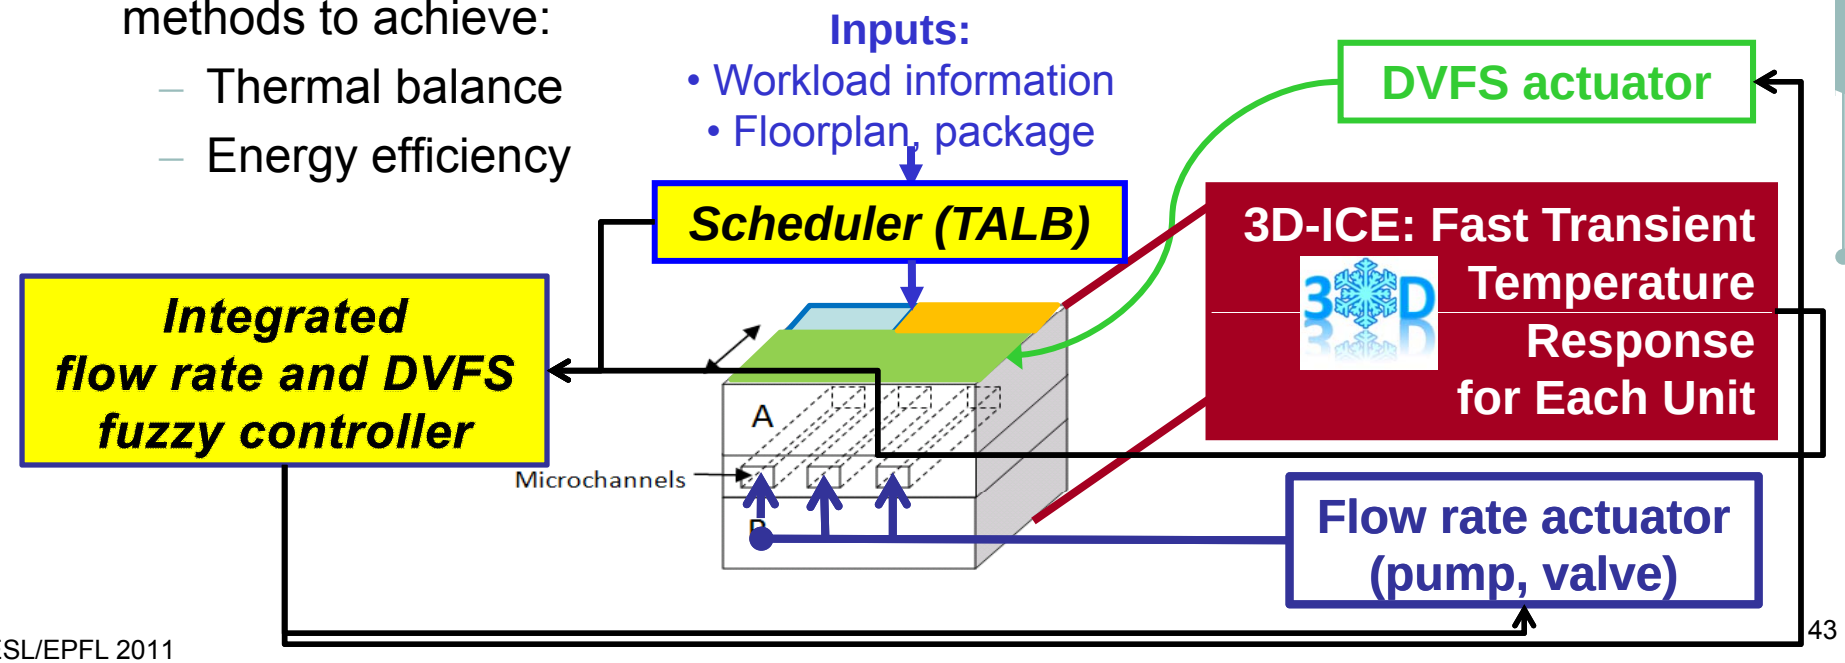

Active-Adapt3D: Active cooling management for 3D MPSoCs

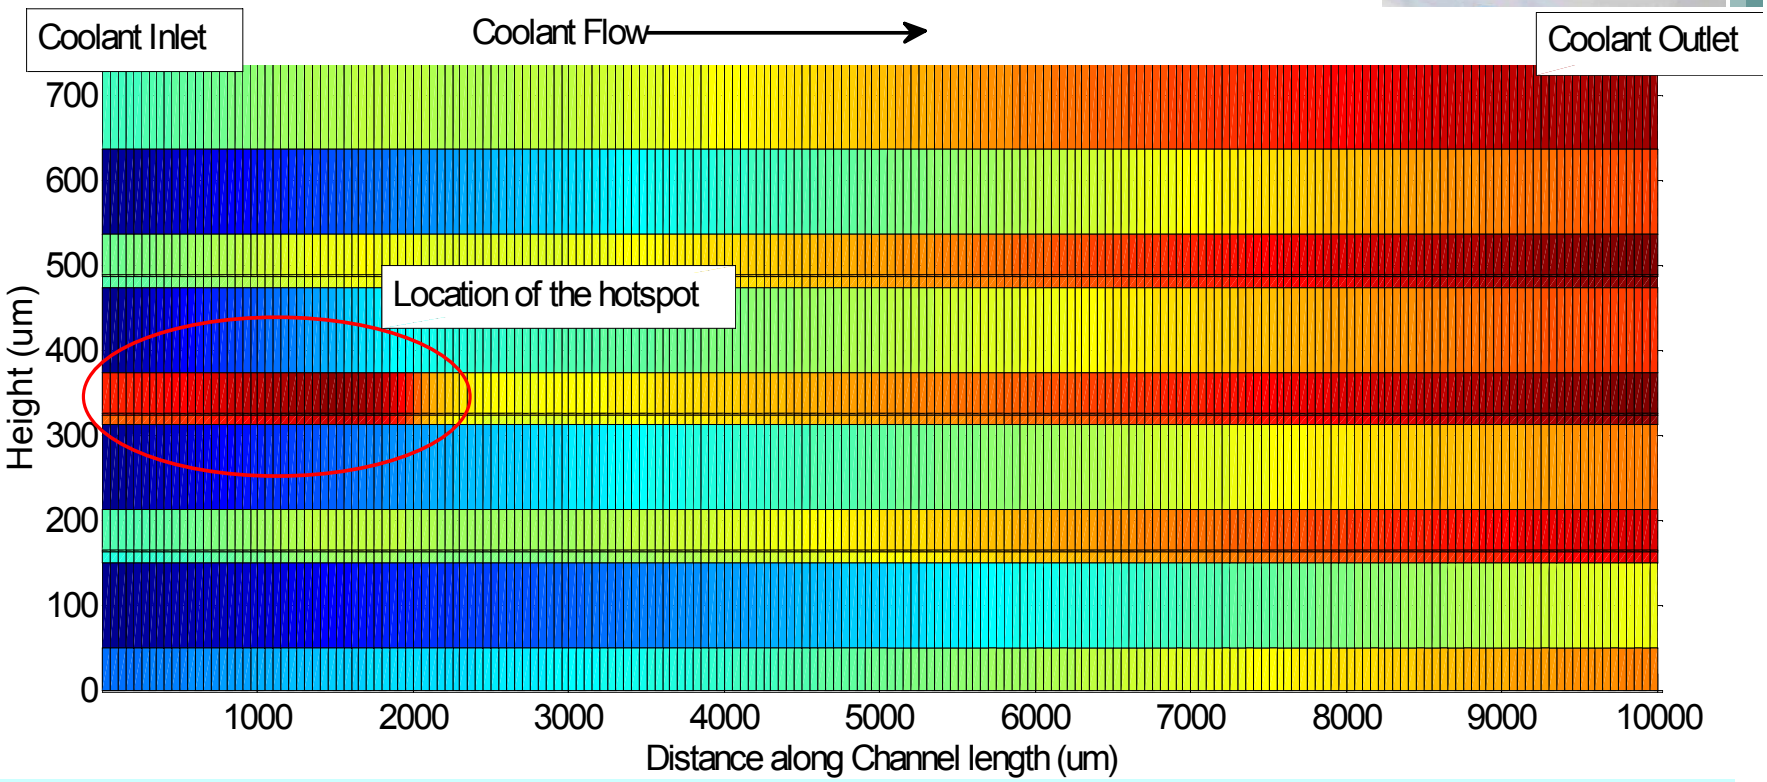

New insights about suitable thermal-aware scheduling and task assignment for 3D MPSoCs!

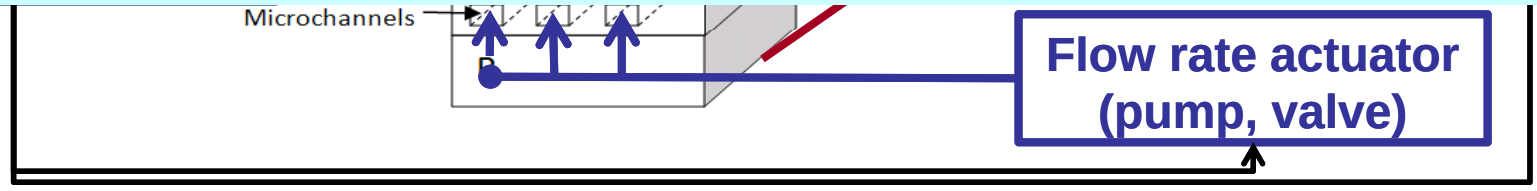

Temperature-Aware Load Balancing (TALB) Scheduler for 3D MPSoCs

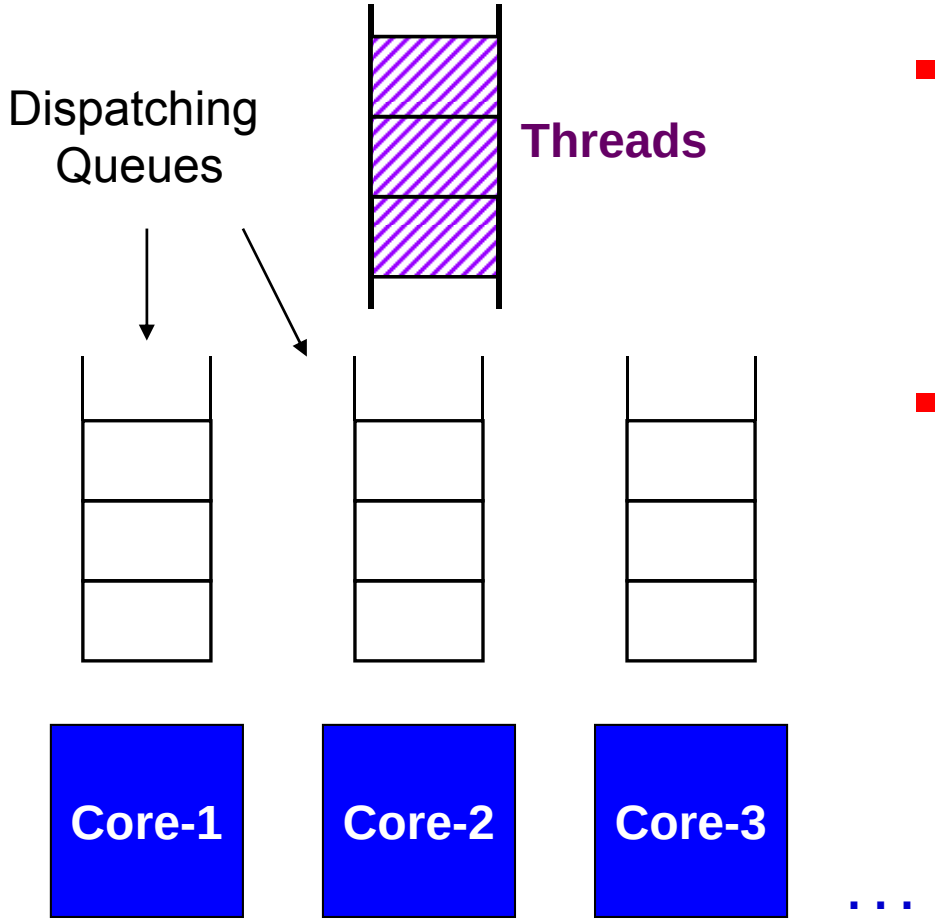

For cores at locations 1, 2 and 3:

$$\alpha^1 > \alpha^2 > \alpha^3$$

- Load balancing (Default):
  - Balances queue-length, but does not consider core locations
- TALB: [Coskun , et al., *DATE 2010*]
  - Thermal index (α) per core

$$l^i_{weighted} = l^i_{queue} \cdot \alpha^i_{thermal}(T(k))$$

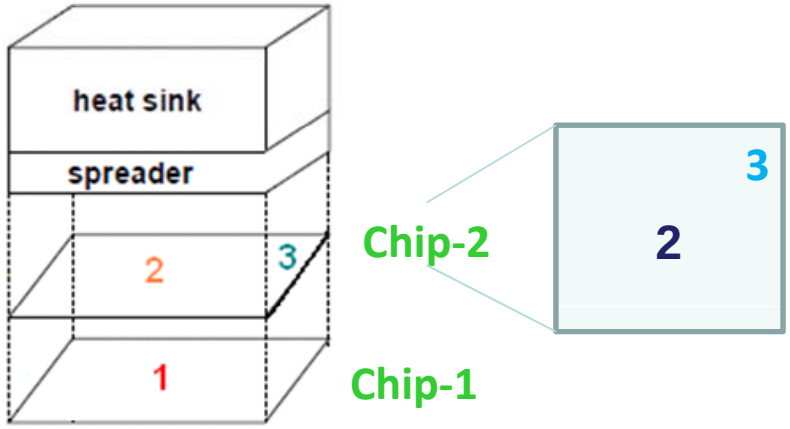

Integrated Flow Rate and DVFS Fuzzy Controller for 3D MPSoCs

- 3D-ICE: predicted max. core temperat.
- TALB: Workload
- Pump power

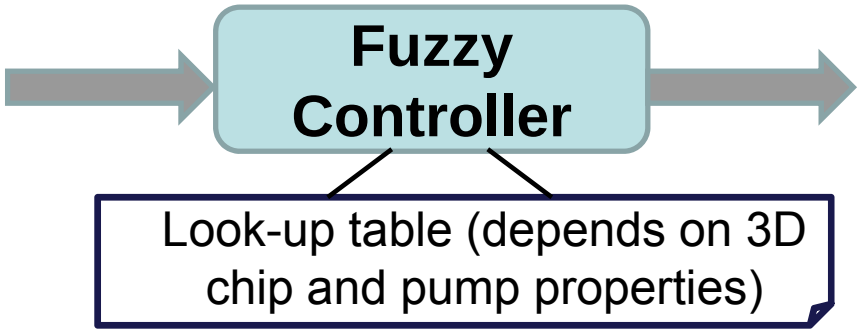

Adjusted Flow Rate and DVFS

[Sabry, et al., *ICCAD 2010*]

- Takagi-Sugeno fuzzy controller

IF $x_1 = a_1$ **AND** $x_2 = a_2$ … **THEN** $y_1 = f_1(x_1, x_2, ..)$

  - Extract cooling rules at design time
  - Stable (circle criteria, BIBO stability,…)

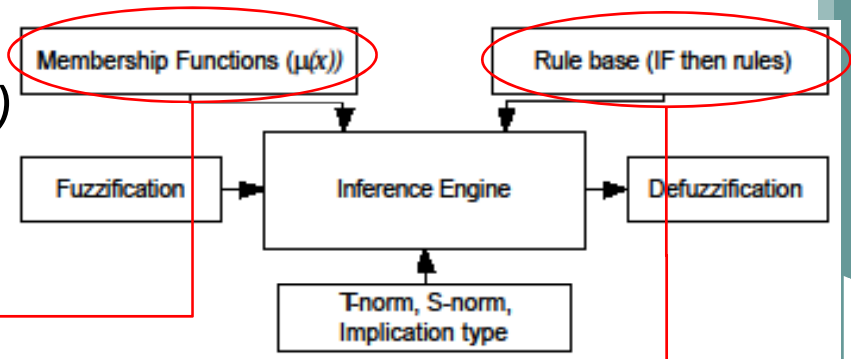

- Rules for multi-outputs at run-time
  - Electrical: DVFS (VF) per core
  - Mechanical: Flow rate (FL)

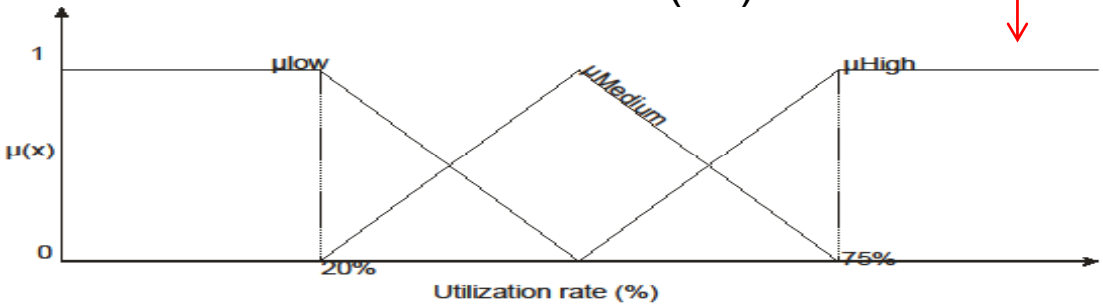

- Cores next to the inlet: no thermal reduction needed

IF D is L **THEN** VF is H **AND** FL is L

4-tier chip: ~150 rules

Experiments Active Thermal Management 3D MPSoCs with Microchannels

- Target 3D systems based on 3D ICs with Sparc-Power cores
  - Power values and workloads from real traces measured in Sun platforms (database queries, web services, etc.)
- Cores and caches in separate layers
  - 3D crossbar as interconnect
- Channels:
  - Width 100µm and height 50µm
  - Three flow rate settings, default at 32ml/min

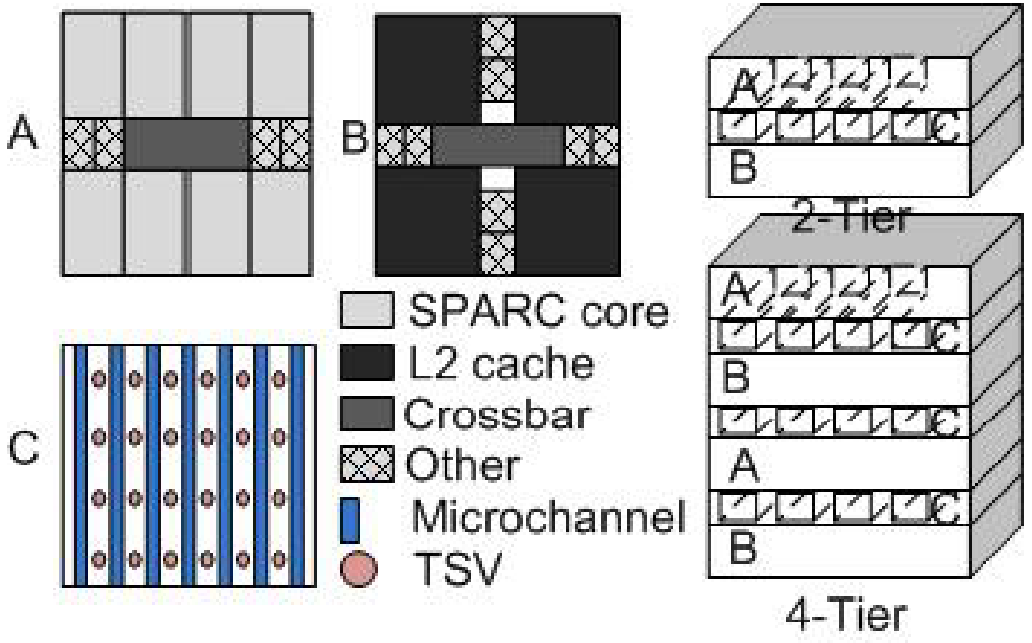

Run-time thermal Management for 3D Chips: thermal evaluation

- For hot spot threshold 85ºC, thermal violations: **0%**
- Energy reduction:
  - **70%** average coolant energy (**max. savings: 77%**)
  - **52%** average total system energy (**max. savings: 85%**)

[Coskun , Atienza, et al., *MICRO 2011*]

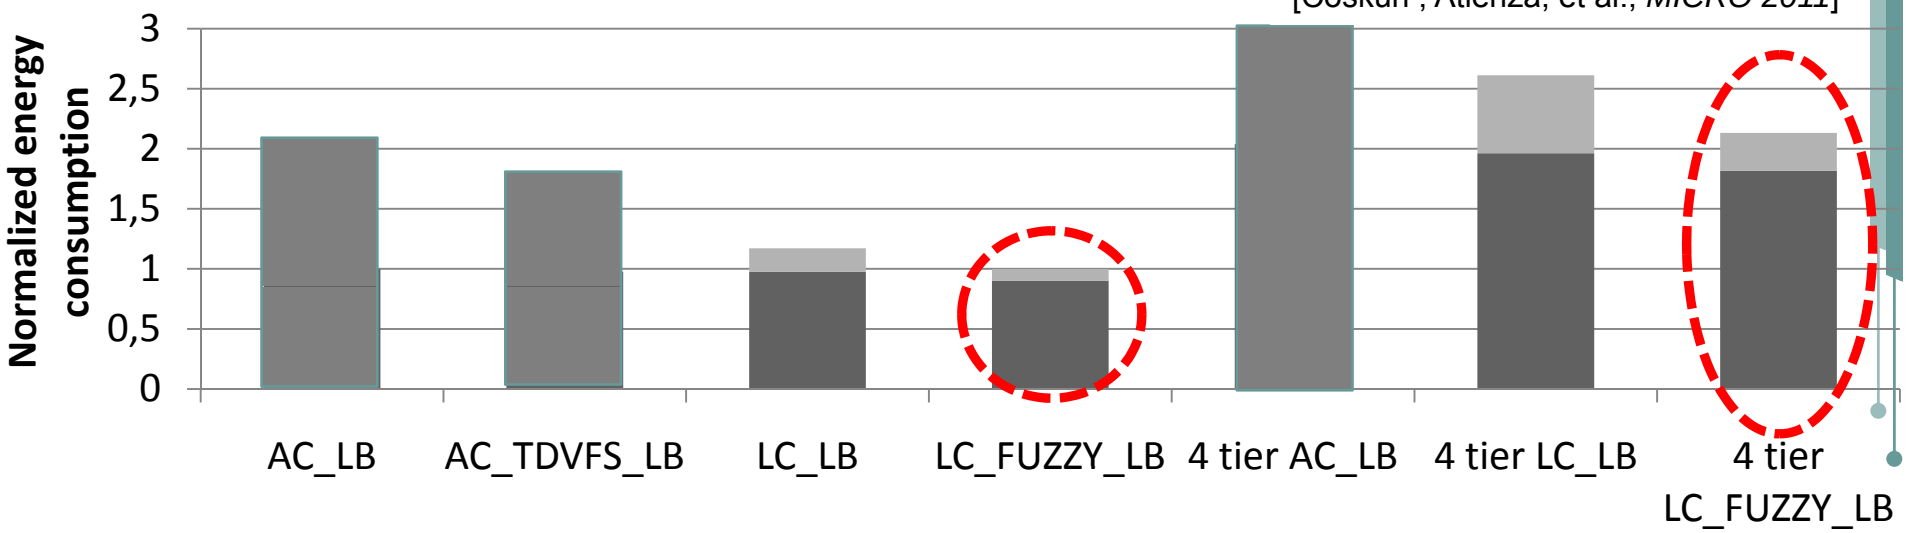

Promising figures for thermal control in 3D MPSoCs, thermal gradients of less than five degrees/tier

Conclusions: Aquasar 2010
First Chip-Level Liquid Cooled Server

- MPSoCs: Interdisciplinary work
  - Fast RC thermal models for 2D/3D MPSoC with inter-tier variable liquid fluxes (**less than 5% error**)
  - Layout combining **electrical and mechanical** constraints-modeling
- Next generation of thermal-aware proactive controllers (task control, flow rate and DVFS)
  - Holistic control reduces significantly the thermal issues and improves energy cost (**80% energy savings**)
  - "Green" datacenters: energy efficient
    - Roadrunner: 445 Mflops/Watt
    - **Aquasar: 2250 MFlops/Watt**

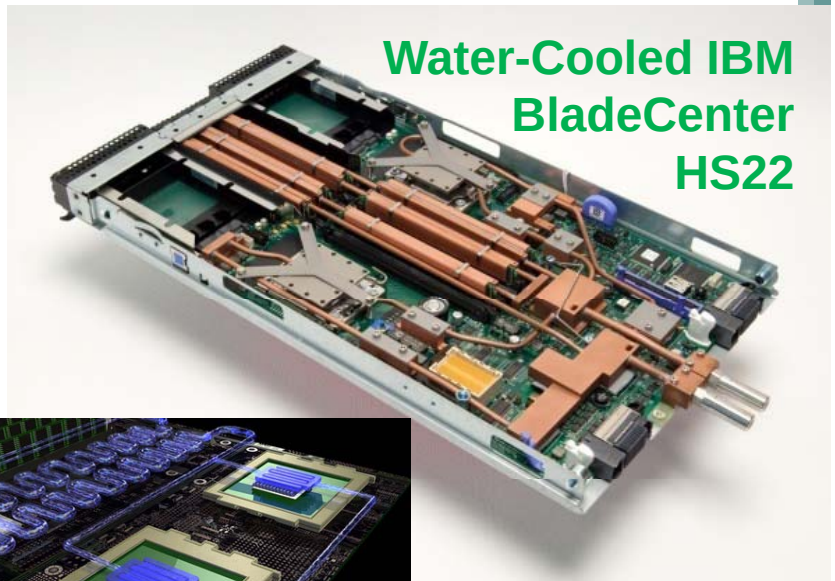

Back side

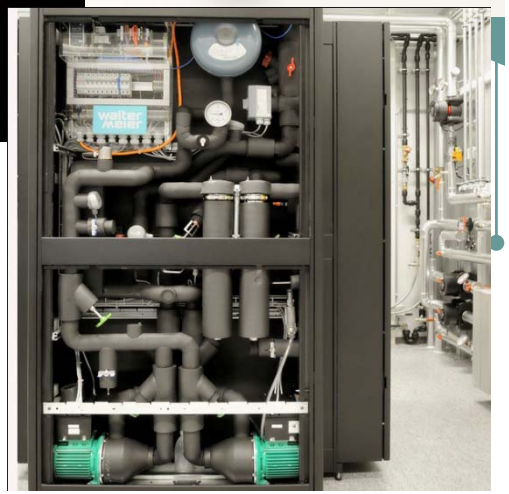

Water conditioning

Courtesy: IBM Zürich

References and Bibliography

- 2D and 3D Thermal modeling
  - "**HW-SW Emulation Framework for Temperature-Aware Design in MPSoCs**", D. Atienza, et al. *ACM TODAES*, August 2007
  - "**3D-ICE: Compact transient thermal model for 3D ICs with liquid cooling via enhanced heat transfer cavity geometries**", A. Sridhar, et al. *Proc. of ICCAD 2010*, USA, November 2010 (http://esl.epfl.ch/3d-ice.html).
  - **"Emulation-based transient thermal modeling of 2D/3D systems-on-chip with active cooling"**, Pablo G. Del Valle, David Atienza, *Elsevier Microelectronics Journal*, December 2010.
  - **"Compact transient thermal model for 3D ICs with liquid cooling via enhanced heat transfer cavity geometries"**, A. Sridhar, et al., *Proc. of THERMINIC 2010*, Spain, October 2010.
  - **"Fast Thermal Simulation of 2D/3D Integrated Circuits Exploiting Neural Networks and GPUs"**, A. Vincenzi, et Al., *Proc. of ISLPED 2011*, Japan, August 2011.

References and Bibliography

- Thermal management for 2D MPSoCs
  - **"Thermal Balancing Policy for Multiprocessor Stream Computing Platforms"**, F. Mulas, et al., *IEEE T-CAD*, December 2009.
  - **"Processor Speed Control with Thermal Constraints"**, A. Mutapcic, et al., *IEEE TCAS-I*, September 2009.
  - **"Online Convex Optimization-Based Algorithm for Thermal Management of MPSoCs"**, F. Zanini, et al., Proc. of *GLSVLSI 2010*, USA, May 2010.
  - **"Temperature Control of High-Performance Multi-core Platforms Using Convex Optimization"**, S. Murali, et al., Proc. of DATE 2008, Germany, March 2008.
  - **"A Control Theory Approach for Thermal Balancing of MPSoC"**, F. Zanini, et al., *Proc. of ASP-DAC 2009*, Japan, January 2009.
  - **"Multicore Thermal Management with Model Predictive Control"**, F. Zanini, et al., Proc. of ECCTD 2009, Turkey, August 2009.

References and Bibliography

- Thermal management for 3D MPSoCs
  - **"3D Stacked Systems with Active Cooling: The Road to Single-Chip High-Performance Computing**", Ayse K. Coskun, et al., *IEEE MIC*RO, August/September 2011.
  - **"Hierarchical Thermal Management Policy for High-Performance 3D Systems with Liquid Cooling"**, F. Zanini, et al., IEEE JETCAS, September 2011.
  - **"Fuzzy Control for Enforcing Energy Efficiency in High-Performance 3D Systems"**, M. Sabry, et al., *Proc. of ICCAD 2010*, USA, November 2010.
  - **"Energy-Efficient Variable-Flow Liquid Cooling in 3D Stacked Architectures"**, Ayse K. Coskun, et al., *Proc. of DATE 2010*, Germany, March 2010.
  - "**Modeling and Dynamic Management of 3D Multicore Systems with Liquid Cooling**", Ayse K. Coskun, et al., *Proc. of VLSI-SoC 2009*, Brazil, October 2009. **(Best Paper Award)**
  - **"Dynamic Thermal Management in 3D Multicore Architectures"**, Ayse K. Coskun, et al., *Proc. of DATE 2009*, France, April 2009.

Thank you!

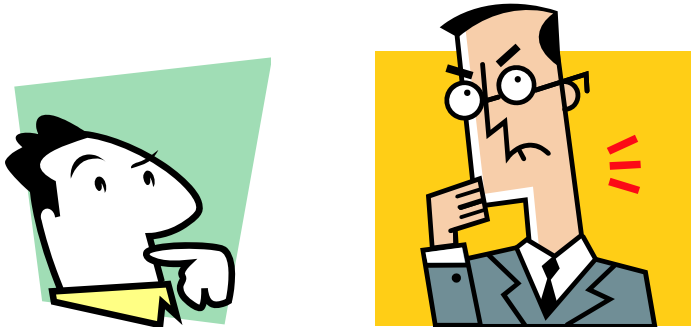

QUESTIONS ?

Acknowledgements: *LSM and LTCM-EPFL, ECE-Boston University*

Nano-Tera.ch Swiss Engineering Programme

Swiss National Science Foundation

European Commission

UCSD / Sun Microsystems

IMEC

IBM Zürich