

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Green Considerations for System-Level Computing Architectures Design

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

Acknowledgements:
LSM-EPFL, Advanced Packaging Group-IBM Zürich,
ECE-Boston University

Smart Energy Day (EPFL, Lausanne, December 13th, 2010)

© ESL/EPFL 2010

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Large-Scale Computing Systems: Energy-Efficiency in Datacenters

- Area is expensive, we try to get denser infrastructures
 - New containers: 2500 servers each, >10x density

Dyer, IThERM 2006

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 10 MW datacenter :
~US\$ 4M wasted per year

Electrical Power IN

INDOOR DATA CENTER HEAT

Waste Heat OUT

Datacenter energy overhead, ASHRAE

© ESL/EPFL 2010 2

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Energy-Centric Design for Datacenters: System-Level View

- Multi-scale energy management solution needed
 - Architecture: middleware, power management, network design, etc.

One server
DRAM: 16GB, 100ns, 20GB/s
Disk: 2TB, 10ms, 200MB/s

Local rack (30 servers)
DRAM: 1TB, 300ns, 100MB/s
Disk: 160TB, 11ms, 100MB/s

Cluster (30 racks)
DRAM: 30TB, 500ns, 10MB/s
Disk: 4.80PB, 12ms, 10MB/s

Barroso & Hölzle, 2009

© ESL/EPFL 2010 [Courtesy: Prof. T. Simunic, UCSD] ³

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Advantages of 3D vs. 2D Multi-Processor System-on-Chip (MPSoCs) ICs

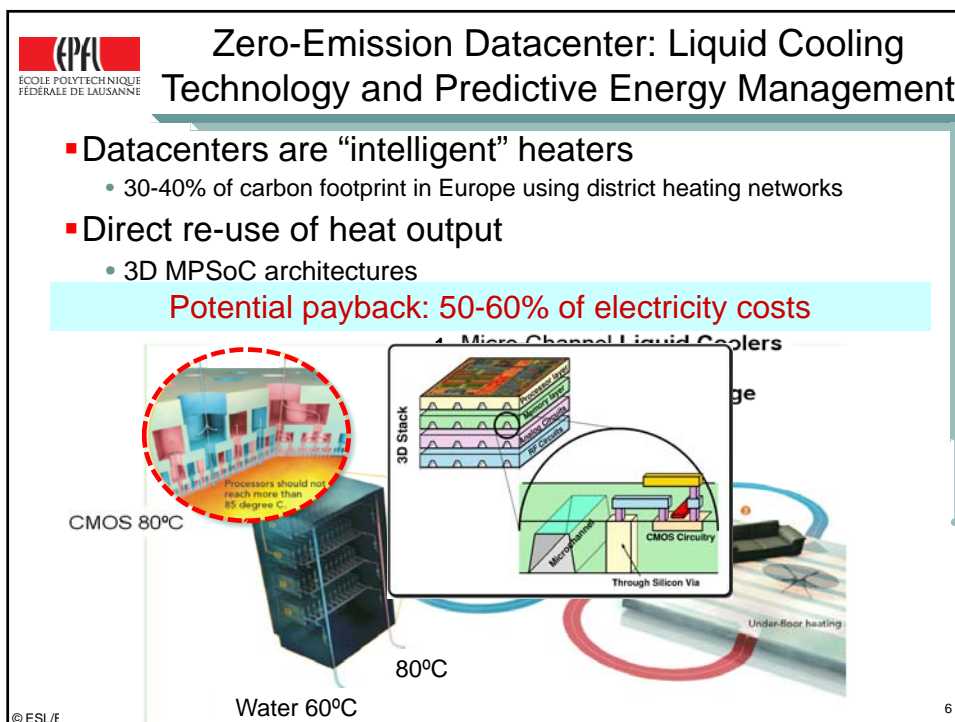
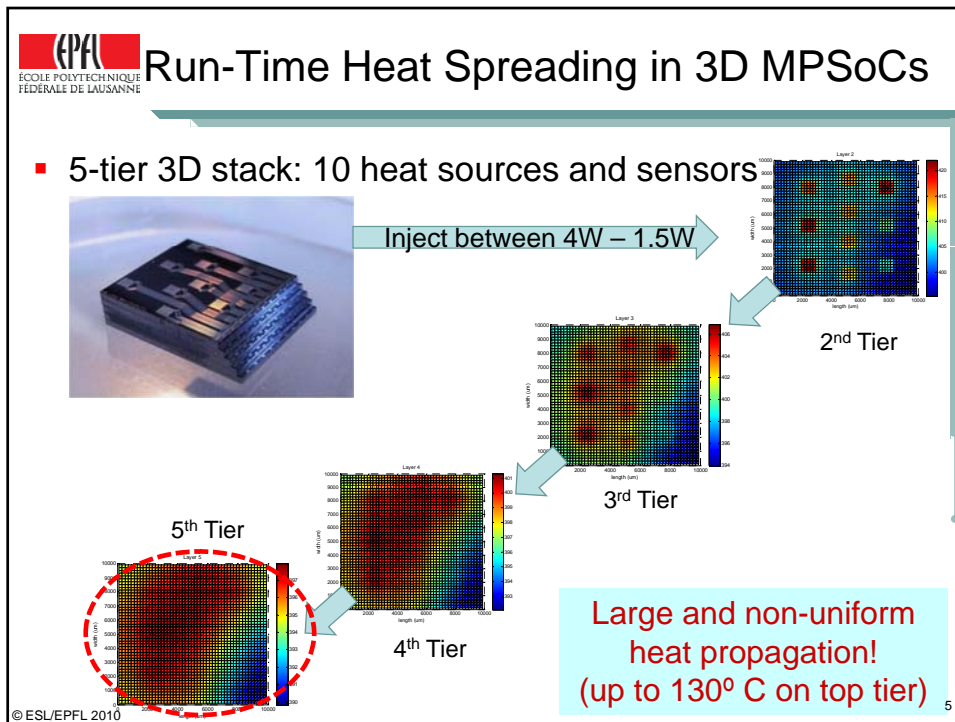
- Promises
 - Reduce average length of on-chip global wires
 - Increase number of devices reachable in given time budget
 - Greatly facilitate massive storage integration (i.e., logic-DRAM stacks)

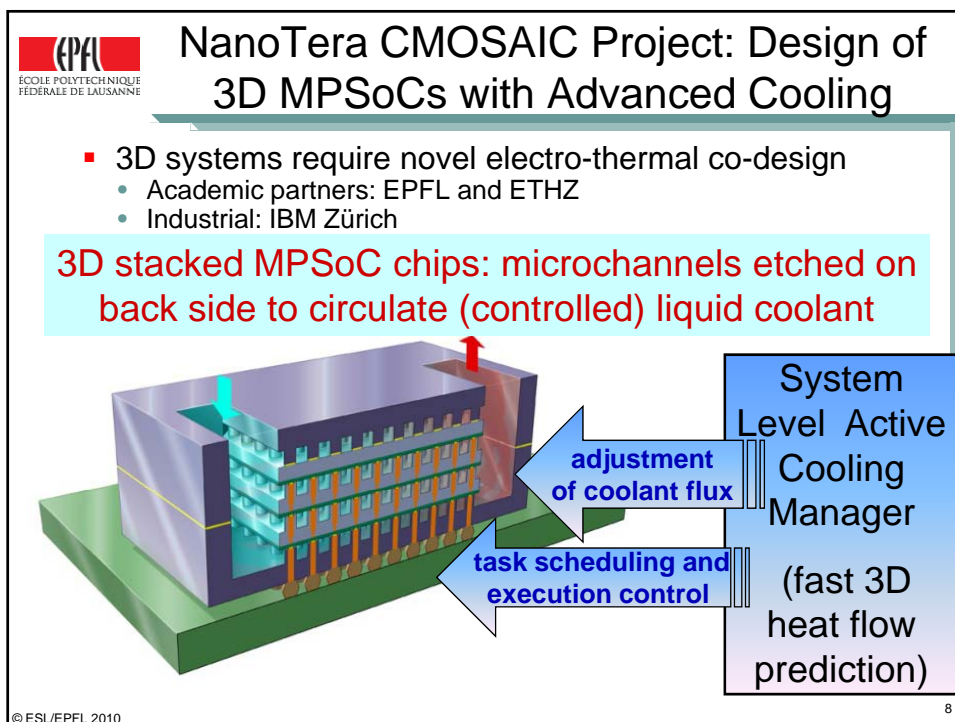
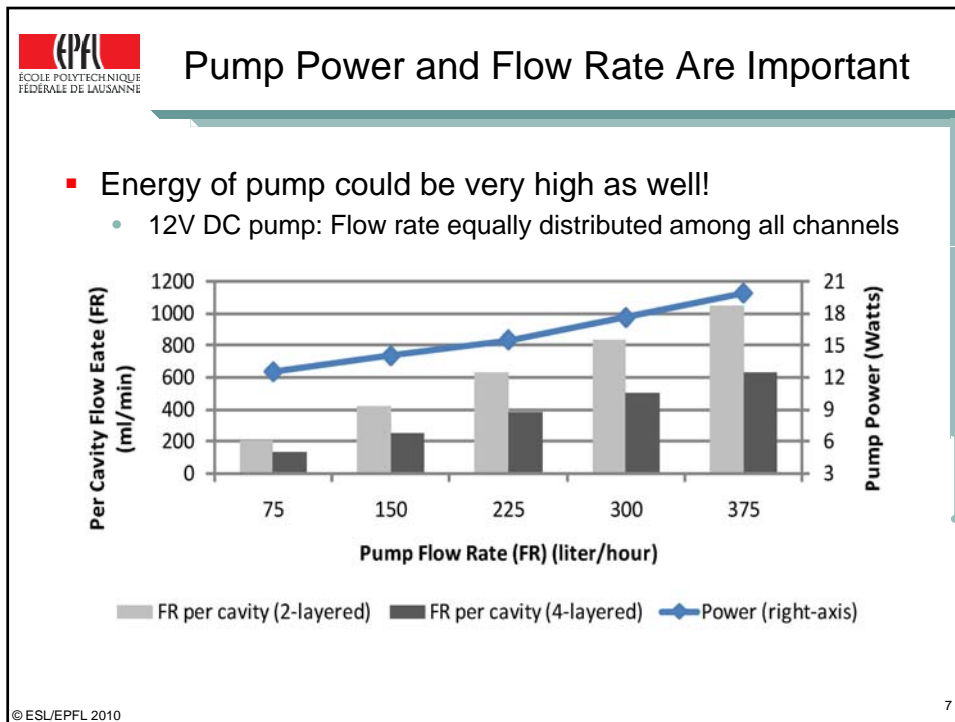
2D Routing (large chip) 3D Routing (small chip)

C4 Bump
Package
Adhesive Layer
TSV
Metal ILD
Active Silicon
Substrate
Thermal Interface Material
Heat Spreader
Heat Sink

Three-Dimensional Multi-Processor System-on-Chip (MPSoC) IC Concept

© ESL/EPFL 2010 4





EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Creating a Fast Thermal Model: Compact RC-Based Tier Model

- Chip-Level thermal model
 - RC Network of Si/metal layer cells
 - 2D tier modeled as heat flux moving between adjacent cells
 - Convective boundary conditions between layers in tier

© ESL/EPFL 2010 9


EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Complete 3D Chip Thermal Modeling

- Multi-level execution for thermal convergence in 3D
 - Local (2D-tier with liquid channels in parallel), and global (3D) propagation at the end

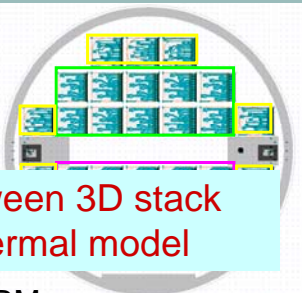
Speed-ups of 800x vs. finite-element thermal simulations

© ESL/EPFL 2010 10



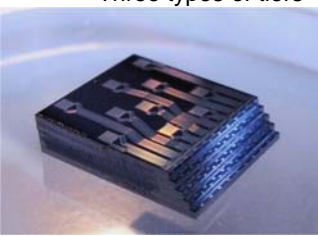
3D Chip Heat Flow Model Validation

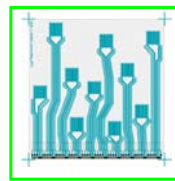
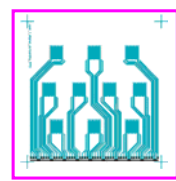

- Extensible set of layers in 3D stack
 - up to 9 tiers and heat spreader
 - Pre-defined layers:
 - Silicon, copper (10 layers), glue, overmold, interposer, bump




Variations of less than 1.5% between 3D stack measurements and new 3D thermal model

- Multi-tier test chip manufactured EPFL-IBM:
 - Three types of tiers



© ESL/EPFL 2010 11



Modeling Through Silicon Vias (TSVs) in 3D Stacks

- TSVs:
 - Size: 5-10um x 10-100um
 - TSVs change resistivity of interlayer material (IM)

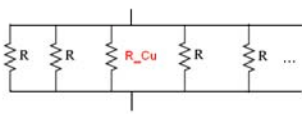
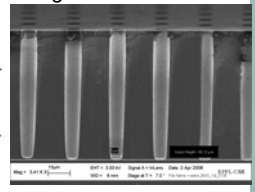


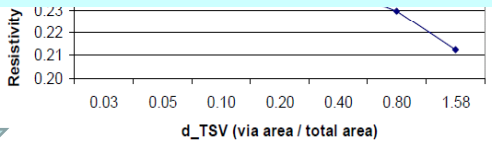
Figure: LSM-EPFL



Chosen to model TSV groups in localized positions of 3D MPSoCs (less than 2% error)

1. value for the IM
2. Different R value per unit (core, cache, etc.)
3. Exact locations of TSVs

• Higher accuracy
• Higher complexity



Resistivity

d_TSV (via area / total area)

Source: IBM Zürich and Y.Heights

© ESL/EPFL 2010 12

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Modeling Liquid Cooling as RC-Network

- Local junction temperature modeled as 4-resistor based compact transient thermal model (4RM-based CTTM)
 - $R_{tot} = R_{cond} + R_{conv} + R_{heat}$

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

4 convective resistances 2 voltage-controlled current sources

© ESL/EPFL 2010 12

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

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

Variations of less than 4% between measurements and RC-based 3D thermal model with liquid cooling

Multi-tier active cooling technology feasible for 3D-stacked MPSoCs

© ESL/EPFL 2010 14

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

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 elements to achieve:
 - Thermal balance
 - Energy efficiency

15

© ESL/EPFL 2010


EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Active cooling management for 3D MPSoCs

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

16

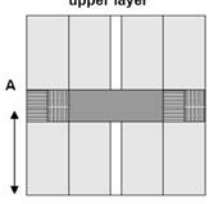
© ESL/EPFL 2010



Experiments 3D Thermal Management: 3D MPSoCs with Microchannels

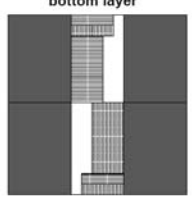
- Target 3D systems based on 3D version Sun UltraSPARC T1+
 - Power values and workloads from real traces measured in Sun platforms (multimedia players, web servers, databases, etc.)
- Cores and caches in separate layers

upper layer



A

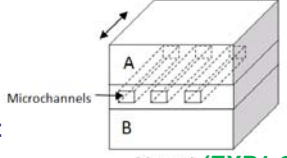
bottom layer



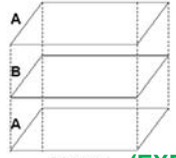
B

sparc core
 crossbar
 scdata
 other

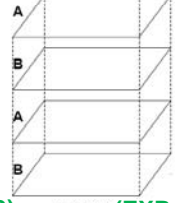
- Channels:
 - Width 100um,
 - Depth 50um.
 - Four flow rate settings, default at 15ml/min.



2-layered (EXP1-2)




3-layered (EXP3)



4-layered (EXP4)

© ESL/EPFL 2010 17




Thermal Management for 3D Chips: Active-Adapt3D Comparisons

- Predictive task scheduling, active cooling and floorplan-aware DVFS achieves **having less than 5% hotspots**

Configuration	EXP1 (Hot Spots)	EXP2 (Hot Spots)	EXP3 (Hot Spots)	EXP4 (Hot Spots)	Perf (Performance)
Default	15	15	23	24	1.0
CGate	8	9	13	14	0.95
DVFS_TT	4	6	7	8	0.95
DVFS_Util	3	5	6	7	0.9
DVFS_FLP	3	5	6	7	0.95
Migr	8	9	13	14	1.0
AdaptRand	6	7	8	10	1.0
Adapt3D	6	7	8	10	0.95
Adapt3D & DVFS-TT	4	5	6	7	0.9
Adapt3D & DVFS_Util	3	4	5	6	0.95
Adapt3D & DVFS_FLP	3	4	5	6	0.9

Promising figures for thermal control in 3D-MPSoCs, thermal gradients of less than 5 degrees/tier


© ESL/EPFL 2010 18



Conclusions

- Green datacenters require novel system-level thermal-aware design methods
 - Application of new 3D chips MPSoC architectures and novel liquid cooling methods for energy-efficient solutions
 - Handling complexity in cooling and control implies an interdisciplinary research work
 - Fast RC thermal models applicable even for variable liquid fluxes in 3D stacks (**errors of less than 4%**)
 - Multi-scale thermal management opportunities (**mechanical and electrical methods**) need to be managed at system-level
- New generation of thermal managers: proactive controllers adjusts flow rate considering task assignment and DVFS
 - Proactive control improves the **hot spot reduction to 95%** for systems with variable flow rates, and **reduces thermal variations**
 - Holistic control reduces significantly the energy cost for the whole system (**67% power savings!**)

© ESL/EPFL 2010 19





Key References and Bibliography

- Thermal modeling and FPGA-based emulation
 - **“HW-SW Emulation Framework for Temperature-Aware Design in MPSoCs”**, David Atienza, et al. *ACM Trans. on Design Automation for Embedded Systems (TODAES)*, Vol. 12, Nr. 3, pp. 1–26, August 2007.
 - **“Emulation-Based Transient Thermal Modeling of 2D/3D Systems-on-Chip with Active Cooling”**, David Atienza, *Proc. of the 15th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC '09)*, IEEE Press, Belgium, October, 2009.
- Thermal management for 3D MPSoCs
 - **“Dynamic Thermal Management in 3D Multicore Architectures”**, Ayse K. Coskun, et al., *Proc. of Design, Automation and Test in Europe (DATE '09)*, France, April 2009.
 - **“Modeling and Dynamic Management of 3D Multicore Systems with Liquid Cooling”**, Ayse K. Coskun, et al., *Proc. of 17th Annual IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC '09)*, Brazil, October 2009. (**Best Paper Award**)
 - **“Through Silicon Via-Based Grid for Thermal Control in 3D Chips”**, Jose L. Ayala, et al., *Proc. of the 4th International ICST Conference on Nano-Networks (Nano-Net '09)*, Switzerland, October 2009.


© ESL/EPFL 2010

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Thank You




QUESTIONS ?




nano-tera.ch

Nano-Tera.ch Swiss
Engineering
Programme



European
Commission



Swiss National Science
Foundation

© ESL/EPFL 2010