

Towards Zero-Emission Datacenters through Direct Re-use of Waste Heat

I. Meijer, T. Brunschwiler, S. Paredes and B. Michel, Advanced Thermal Packaging, IBM Research GmbH, Säumerstrasse 4, 8803 Rüschlikon, Switzerland, bmi@zurich.ibm.com

Overview

- Green Datacenter market drivers and trends
 - *Inefficient air cooled data centers: Waste of energy and exergy*
- Thinking global about carbon footprint and energy usage
 - *Role of IT to tackle climate change*
 - *Economic and political interest*
- Thermal packaging and liquid cooling
 - *Improved thermal conductivity and heat transfer*
 - *Water cooling and refrigerant cooling*
 - *Hotspot cooling*
 - *History of liquid cooling: Cold – Warm - Hot*
 - *Future interlayer cooling of 3D stacked chips*
- Energy re-use in liquid cooled data centers
 - *Reduction of carbon footprint with efficiency increase and community heating*
 - *Value of heating and cooling in different climates*
 - *Joint project with ETH: Aquasar*
- Main messages and next steps
 - *Efficiency investments have a short payback time*
 - *Concentrated photovoltaics with energy recovery*


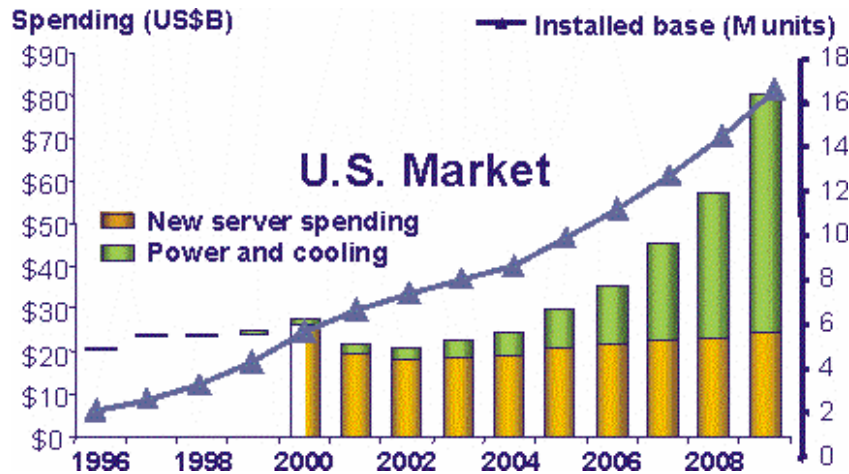


Green Datacenter Market Drivers and Trends

- Increased green consciousness, and rising cost of power
- IT demand outpaces technology improvements
 - Server energy use doubled 2000-2005; expected to increase 15%/year
 - 15 % power growth per year is not sustainable
 - Koomey Study: Server use 1.2% of U.S. energy
- ICT industries consume 2% ww energy
 - Carbon dioxide emission like global aviation



Real Actions Needed



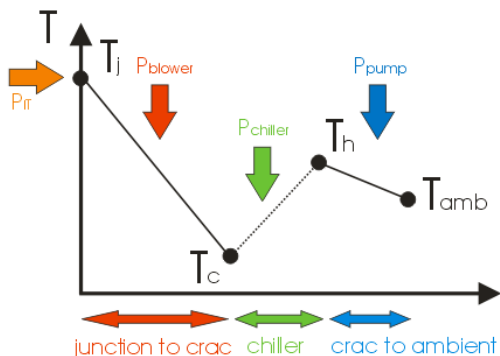
Brouillard, APC, 2006

Future datacenters dominated by energy cost; half energy spent on cooling

Source IDC 2006, Document# 201722, "The impact of Power and Cooling on Datacenter Infrastructure, John Humphreys, Jed Scaramella"


Datacenter: Cooling Infrastructure

DC1:




Chillers
(Refrigeration)

Evaporative Tower Fans



Electronic Power

Chilled Water
CRAC/
CRAHs



Condensator




What's Driving Demand


- Mobility** – cell/pda, broadband services, security, merchandising (13.6% CAGR from IT Mobile Tracker)
- Digital Media** – streaming, iptv, music, Radiology, Visualization (200% CAGR from e-week Jan 2005)
- 24x7 global e-commerce:** financial & manufacturing sectors (26% CAGR
 Source: United States Department of [Commerce](#), 11/ 22/05)
- HPC:** BioMed – Pharma, genetic research, Oil & Gas, Virtualization in structural dynamics, Weather (10.5% CAGR Source: IDC)
- Real time BI:** Walmart, Amazon, Yahoo/Google/MSN (Google Growth 79% Source: Wall Street Journal 4/21/06)
- Compliance** – Hipaa, Sox – ILM (“SOX may be your biggest information-technology expenditure this decade” source: e-week 8/1/03).



From the Individual Transistor to the Globe



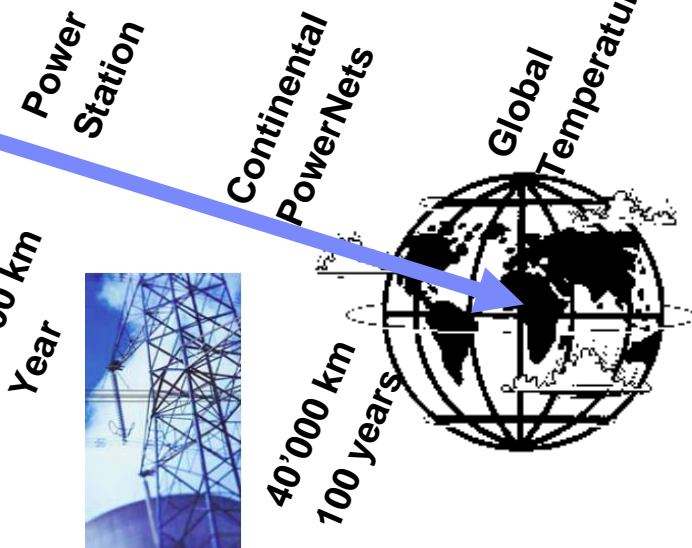
Morteratsch Glacier, Switzerland, 1985




Morteratsch Glacier, Switzerland, 2007



Glacier recession 1985-2007: 400 meters!




- Thermal/power issues propagate up to the world climate
 - Global length- and decade-long time-scales involved
 - A truly holistic view is required to solve these problems
 - IT efficiency improves more than 15% per year
- ➔ IT to become part of the solution!



Energy Consumption of Transistor (Leakage Current)

Table PIDS2a High-performance Logic Technology Requirements—Near-term



Year of Production	2007	2008	2009	2010	2011	2012	2013
MPU/ASIC Metal 1 (M1) 1/2 Pitch (nm)(contacted)	68	59	52	45	40	36	32
MPU Physical Gate Length (nm)	25	22	20	18	16	14	13

Equivalent Oxide Thickness

Extended planar bulk (Å)	11	9	7.5	6.5	5.5	5			
UTB FD (Å)				7	6	5.5	5	5	5
DG (Å)					8	7	6	6	6

Maximum gate Leakage Current

Extended Planar Bulk (A/cm ²)	8.00E+02	9.09E+02	1.00E+03	1.11E+03	1.25E+03	1.43E+03			
UTB FD (A/cm ²)				1.11E+03	1.25E+03	1.43E+03	1.54E+03	1.82E+03	2.00E+03
DG (A/cm ²)					1.25E+03	1.43E+03	1.54E+03	1.82E+03	2.00E+03

Source/Drain Leakage Current

	Leakage Current [8]								
Extended Planar Bulk (μA/μm)	0.34	0.71	0.70	0.64	0.74	0.68			
UTB FD (μA/μm)				0.33	0.52	0.62	0.56	0.55	0.60
DG (μA/μm)					0.2	0.34	0.37	0.38	0.38

Rethink: Green Datacenter Metric

New ranking list for supercomputers: Green 500

- Ranking of supercomputers which are listed in Top 500
- Energy needed for one floating point operation: **MFLOPS / Watt**
- Rank 1 to 20: PowerXCell, BlueGene, and MD GRAPE Accelerator


Earth Simulator 2002

- Linpack: 40.9 TFlops
- Peak power: 11.9MW
- 3.4 MFlops/Watt

Roadrunner 2008

- PowerXCell 3.2GHz
 - Linpack: 1100 TFlops
 - Peak power: 2.5MW
 - 445 MFlops/Watt
- No. 1 Top500, No.4 Green500

27x computing performance, 1/5 power consumption, but harder to program



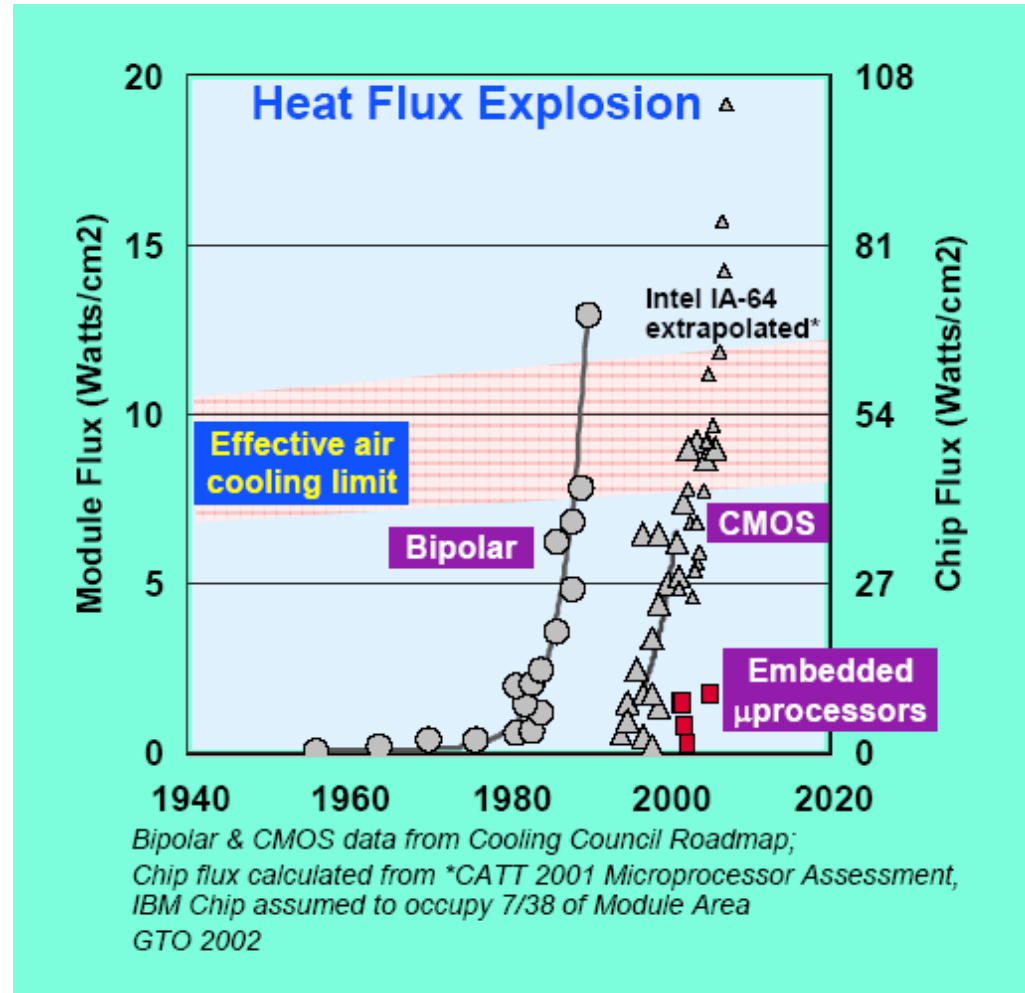
New metric needed that includes the site & facility energy consumption

Historic Heat Flux Trends

- Module and Chip Heat Flux Explosion
- Number of transistors doubles every 18 month
- CMOS scaling was power density invariant 1980 – 1995 → reduction of V_{dd}

$$P_{sw} = \frac{1}{2} C \cdot f \cdot V_{dd}^2$$

- Energy per operation still shrinks
- Challenge:
Transistor leakage



Thermal Packaging History

Time line




8088
Natural convection
bare die



Pentium III
Forced convection
50x50x40 mm³




Pentium IV
Forced convection
80x80x70 mm³





Apple G5
Liquid cooling
170x170x170 mm³

→ Increase 10x in volume, 10x in cost, 10x in complexity

History of Water Cooling at IBM


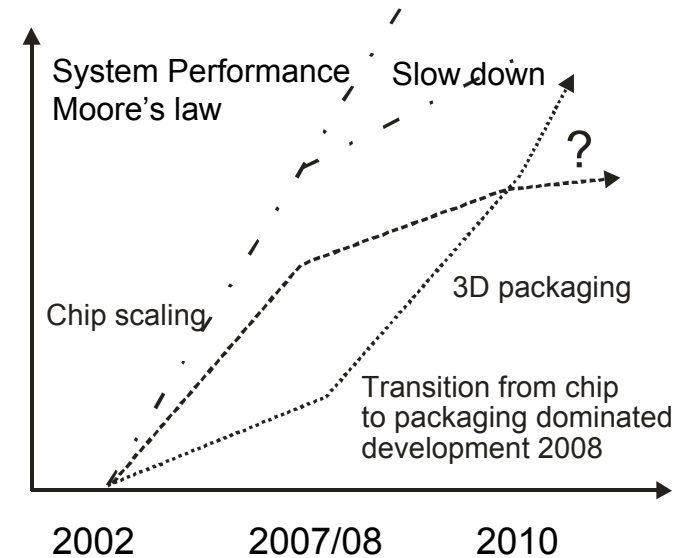


- Thermal conduction module for 3090 water cooled systems
- Used conventional water cooling for power densities up to 100 W/cm^2



Motivation for Packaging Research

- Scaling beyond 22 nm decreases performance ~30% per generation
 - Needs to be compensated by new technologies like high k dielectrics, air gap, multicore
 - Will get more difficult with every generation
- ITRS reports “Acceleration of pace in assembly and packaging” and extensively revised roadmap in 2008 upgrade
- Packaging compensates for slower chip efficiency improvements



Why Thermal Packaging?



Performance

- Increased mobility at lower temperatures
- Leakage current depends exponentially on temperature

Reliability



- Most failure mechanisms are accelerated with temperature
- Catastrophic failure can occur due to thermo-mechanical stress

Typical processor package



Packaging and Thermal Interfaces

- Interfaces are large portion of total resistance
 - Thermal interface materials TIM 1 and TIM 2 cause almost half the overall thermal resistance in a high performance processor package
- Particle filled materials have cost benefit
 - Easier processing, no metallization, flexibility for many applications
- Conductivity increase with higher particle loading
 - Viscosity and shear strength also increase
 - If bondline thickness increases – **No Gain!**
- Assembly loads cannot be too high
 - C4 crushing, chip cracking
 - Substrates bend trapping thick TIM
- Hierarchical Nested Channel (HNC) creates **thinner bondlines** with **higher conductivity** materials using **low assembly forces...**






High Performance Thermal Interface Technology

- IEEE Harvey Rosten Award for Excellence in Thermal Sciences
2008 R. Linderman, T. Brunschwiler, U. Kloter, H. Toy and B. Michel


Directed self-assembly

- Fluid-shear driven self-assembly
- Control of stacking with channel pattern



High performance thermal interface

- Increased particle density
- High performance with matched paste $R_{th} (<5 \text{ mm}^2\text{K/W})$
- Quick integration into products possible




Motivation for Liquid Cooling

- Increase in heat removal performance:
Superior thermal properties of liquids compared to air
- Design flexibility: Sensible heat transport to locations with available space
- Centralized secondary heat exchanger
- Efficient water-water heat exchanger

	Thermal conductivity [W/(m*K)]	Volumetric heat capacity [kJ/(m ³ *K)]
Air	0.0245	1.27
H ₂ O	0.6	4176

Disadvantage: Increased complexity




Long distance transport possible




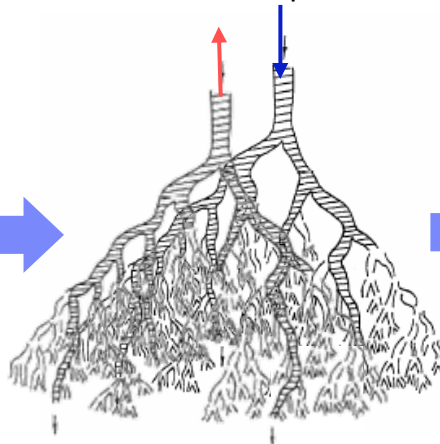
Limited heat transport due to fin efficiency

Chip Scale Liquid Cooling



Arrayed jets, distributed return



Biological vascular systems are optimized for the mass transport at low pressure



Cooling of up to 350 W/cm²




SEM cross-section of two-level jet plate with diameter of 35μm

Direct Liquid Jet-Impingement Cooling with Micron-Sized Nozzle Array and Distributed Return Architecture, T. Brunschwiler et al., ITHERM 2006

Ultra Thin High Efficiency Heat Sinks


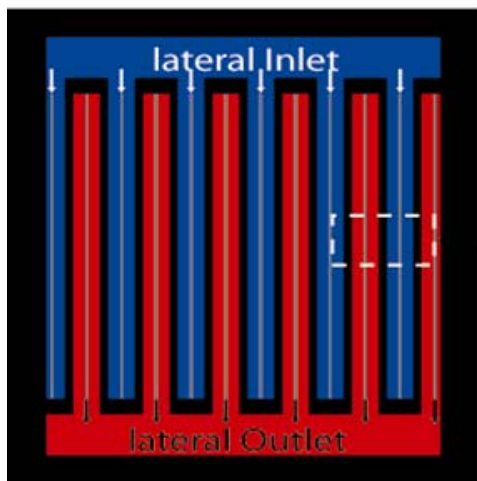
- Motivation: Find the best coolant and the best structure for ultra-compact heatsinks (thickness < 2 mm)
 - Nanofluid thermal properties explained by effective medium theory which means they cannot 'magically' improve heat transfer
- ➔ Water provides the best combination of material properties





- Flat heatsinks reduce the board pitch of future systems from >30 mm (1U) to 3 mm (1/10 U)
- Optimum design provides a total thermal resistance of $0.09 \text{ cm}^2\text{K/W}$ @ $V = 1.3 \text{ l/min}$, $\Delta p = 0.22 \text{ bar}$
⇒ maximum power density > 700 W/cm^2 for $\Delta T = 65 \text{ K}$
- Increasing inlet temperature to 70°C (190 F) enhances the heat sink efficiency >40%



Manifold Micro-Channel Heat sink





W. Escher, B. Michel and D. Poulikakos, 'A novel high performance, ultra thin heat sink for electronics', International Journal of Heat and Mass Transfer, 2009.



W. Escher, T. Brunschweiler, B. Michel and D. Poulikakos, 'Experimental Investigation of an Ultra-thin Manifold Micro-channel Heat Sink for Liquid-Cooled Chips', ASME Journal of Heat Transfer, 2009.



RADially OScillating Flow Hybrid Cooling System - RADIOS



Schematic operation of fluid pumping and temperature oscillations

- Self-contained, thin form factor
- Spreader plates in base of air heat sinks
- Basic principle:
 - Cold plate on the chip (Heat Absorber, HA)
 - Heat transport to larger area
 - Cold plates in periphery (Heat Dissipator, HD)
- Keep chip area free, heat transfer to air at unpopulated area
- Heat shuttling via fluidic branches (N)
- Displacement actuator with membranes
 - Periodic phase difference ($\varphi = 2\pi/N$)
- Fluid periodically dispensed in all directions
- At the center:
 - Constant flow speed with radially oscillating direction


Oscillating Flow Liquid Cooling




Unit cell stack

- Hybrid cooling: internal fluid external air
- Intermediate step – does not require external fluid connection
- Hermetically sealed, low risk
- Displacement pumps, low fluid volume
- Multiple pumping schemes

Blade-level




: Heat flow
 : Heat and fluid flow



Rack-level


Hot Spots are Everywhere

Power map of a dual core microprocessor



Hot spot heat flux 4x higher than mean

Temperature map of a data center



Current thermal management infrastructure is **over-dimensioned** to keep hot spots cool

Improved efficiency:

- Thermal aware chip and datacenter design → reduced hot spot peak heat flux
- Hot spot adapted cooling architectures → minimal pumping power and thermal mixing

Phonon Transport Engineering


Investigation:

- Phonon relaxation times in channels
- Vibrational matching at interfaces
- Phonon tunneling in superlattice structures



Molecular Dynamics Modeling:

- Phonon transport considered, no electrons → valid in dielectrics

Hot spot mitigation in SOI channel:



Solid to fluid phonon transport engineering:



Improved phonon coupling at hydrophilic interfaces

Slip Flow Induced Pressure Drop Reduction

Super-hydrophobicity:

- Micron / nano-sized topography with low surface free energy
- Water droplet contact angle close to 180°C
- Pressure stability: according to Laplace pressure $P \sim 1/r$


Pressure drop reduction:

- Slip length as a result of finite velocity at fluid-air interface
- If slip length 1/10 of the hydraulic diameter
→ 60% reduction in pressure drop


Lotus-Effect





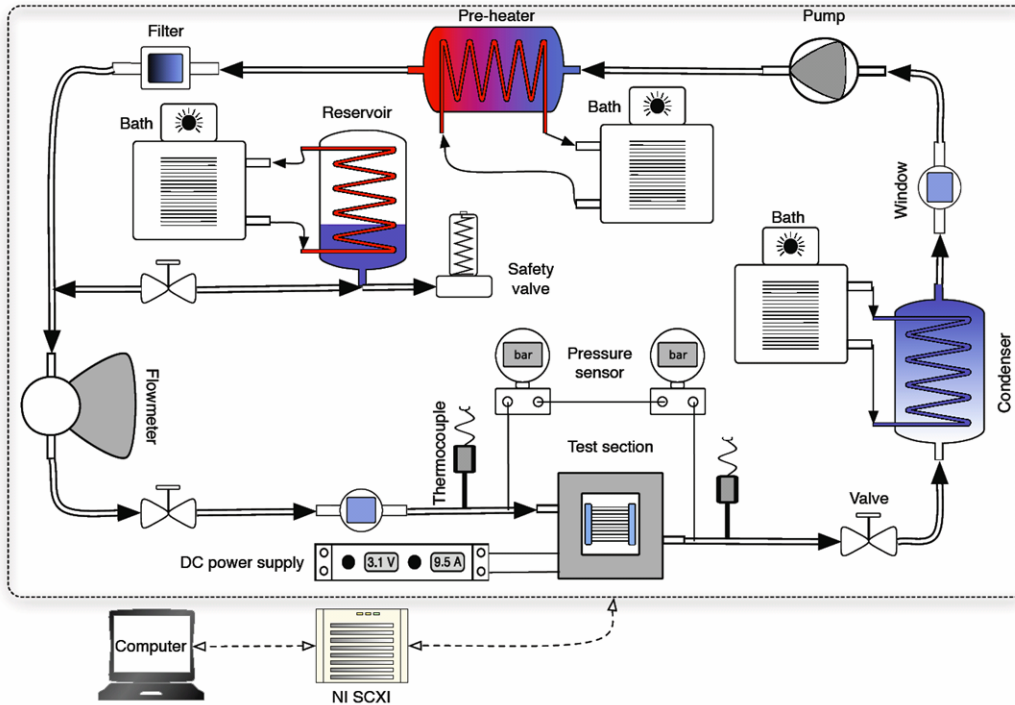
Three-phase Cassi-state



(Ou 04)



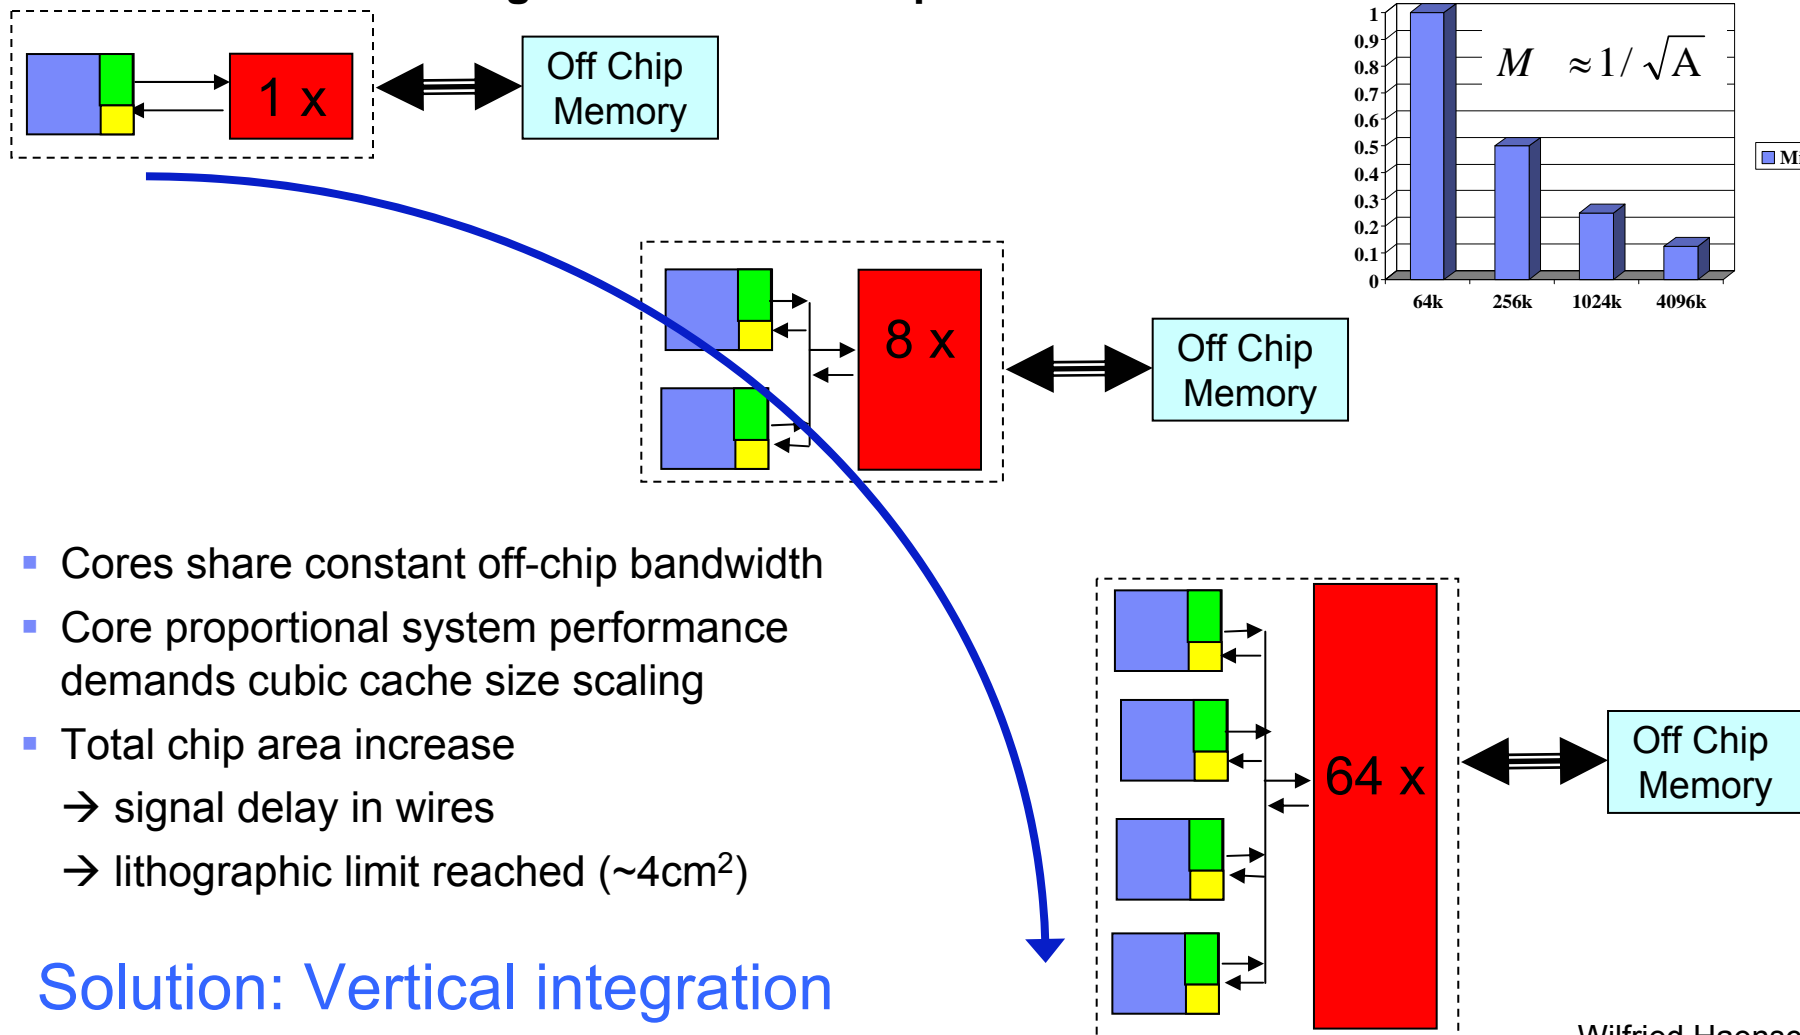
Flow-Boiling Loop



- Cooling with dielectric refrigerants R236fa and R245fa
 - Remove same chip heat flux
 - Eliminates risk for electronics
- More complex and more expensive to build and run
 - High system pressure (> 2 bar / 28 psi)
 - Many more control points needed

Multi-Core Architecture: Communication Bandwidth Limit

Cache – core balancing at constant off-chip bandwidth




- Cores share constant off-chip bandwidth
- Core proportional system performance demands cubic cache size scaling
- Total chip area increase
 - signal delay in wires
 - lithographic limit reached ($\sim 4\text{cm}^2$)

Solution: Vertical integration

Wilfried Haensch 08

Limits of Traditional Back-Side Heat Removal



Guidelines from a thermal perspective:

- MPU as close as possible to the cold plate
 - Lower peak temperature → high heat flux is conducted through minimum number of layers
 - Memory can handle 15K higher junction temperatures
- Non-identical hot spot locations

Unacceptable:


- Two identical MPU's with overlapping hot spot
- More than two MPU layers



Heat removal limit constrains electrical design

Interlayer Thermal Management

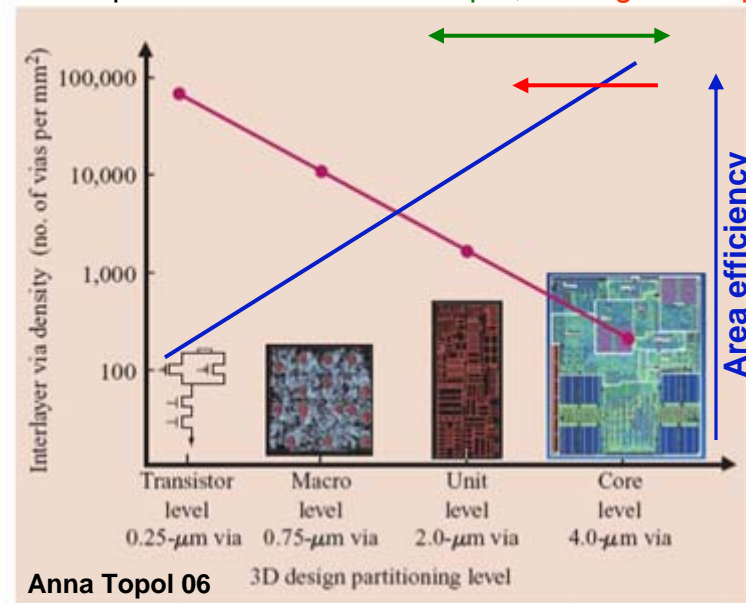
Interlayer cooled chip stack




Hydraulic diameter: 25-100 μ m
 Flow rate 1/10 of traditional cold plates

Interconnect density

Pitch optimum: electrical 1 - 100 μ m, cooling 50-200 μ m





Interconnect compatible heat transfer structures



Convective interlayer heat removal scales with the number of stacked tiers

Scalable Heat Removal by 3D Interlayer Cooling

- 3D integration will require interlayer cooling for stacked logic chips
- Bonding scheme to isolate electrical interconnects from coolant
- Heat removal scales with the number of dies



Solder functionality

sealing } Thermo-
electrical } mechanical

Through silicon via with bonding scheme

Interlayer cooling with lateral feed manifold

- Cool between logical layers with optimal vias
 - Best performance with 200 μm pin fins
 - Through-silicon via height limit, typically 150 μm
 - Microchannel, pin fins staggered/in line, drop shape

- Interlayer cooling of 3D stacked chips
 - Remove 180 W/cm² per layer or
 - Remove 7.2 KW from 10 layers with 4 cm²

Interconnect compatible heat transfer structures

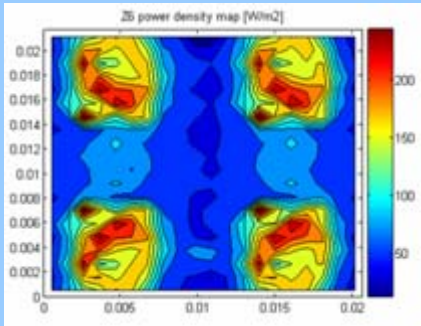


- Microchannel
- Pin fin inline / staggered

Electro-Thermal Co-Design

Chip design

Power map




Electrical interconnects




Heat Transfer Building Blocks

Efficient heat removal

- Heat transfer structure




- Modulation of heat transfer structure

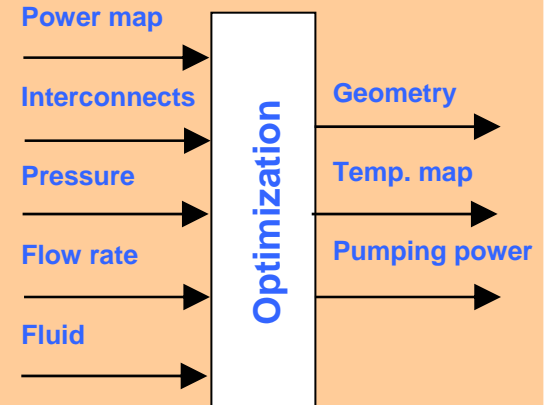


Increase in local hot spot flow rate

- Fluid focusing



Heat Transfer Structure Design



Feedback


Experimental Validation: Pyramid Chip Stack

Thermal Demonstrator:

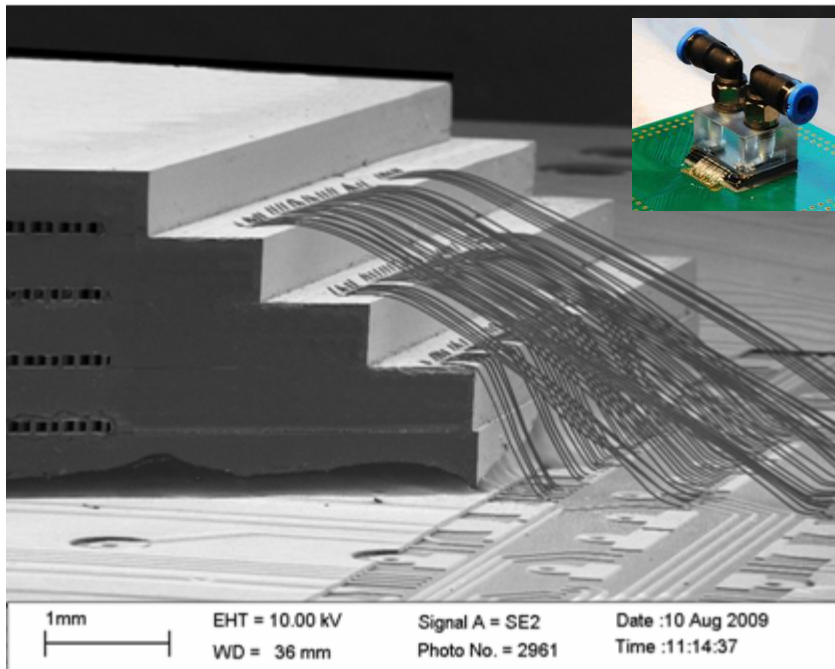
- Three active tiers, cooled with four cavities
- Polyimide bonding → represents wiring levels
- Multi-scale modeling accuracy validated (+/-10%)

Realistic Product Style Stack:

- Aligned hot-spot heat flux of 250W/cm² possible



Pyramid chip stack




Interlayer cooled chip stack

Cross-section: central



Top view: middle





Random power map

Zero-Emission Data Centers

- **High-performance chip-level cooling** improves **energy efficiency** AND **reduces carbon emission**:
 - Cool chip with $\Delta T = 20^\circ\text{C}$ (previously 75°C)
 - and cool datacenter with $T > 60^\circ\text{C}$ (170 F) **hot water**; **no chillers** are required anymore
 - **Re-use waste energy** in moderate climate, e.g., heat 700 homes with waste heat from 10 MW datacenter
- Necessity for **carbon footprint reduction**
 - EU, IPCC, Stern report targets
- Note:
 - Chillers use $\sim 50\%$ of datacenter energy
 - Space heating $\sim 30\%$ of carbon footprint


T. Brunschwiler, B. Smith, E. Ruetsche, and B. Michel, "Toward zero-emission datacenters through direct reuse of thermal energy", IBM JRD 53(3), paper 11.



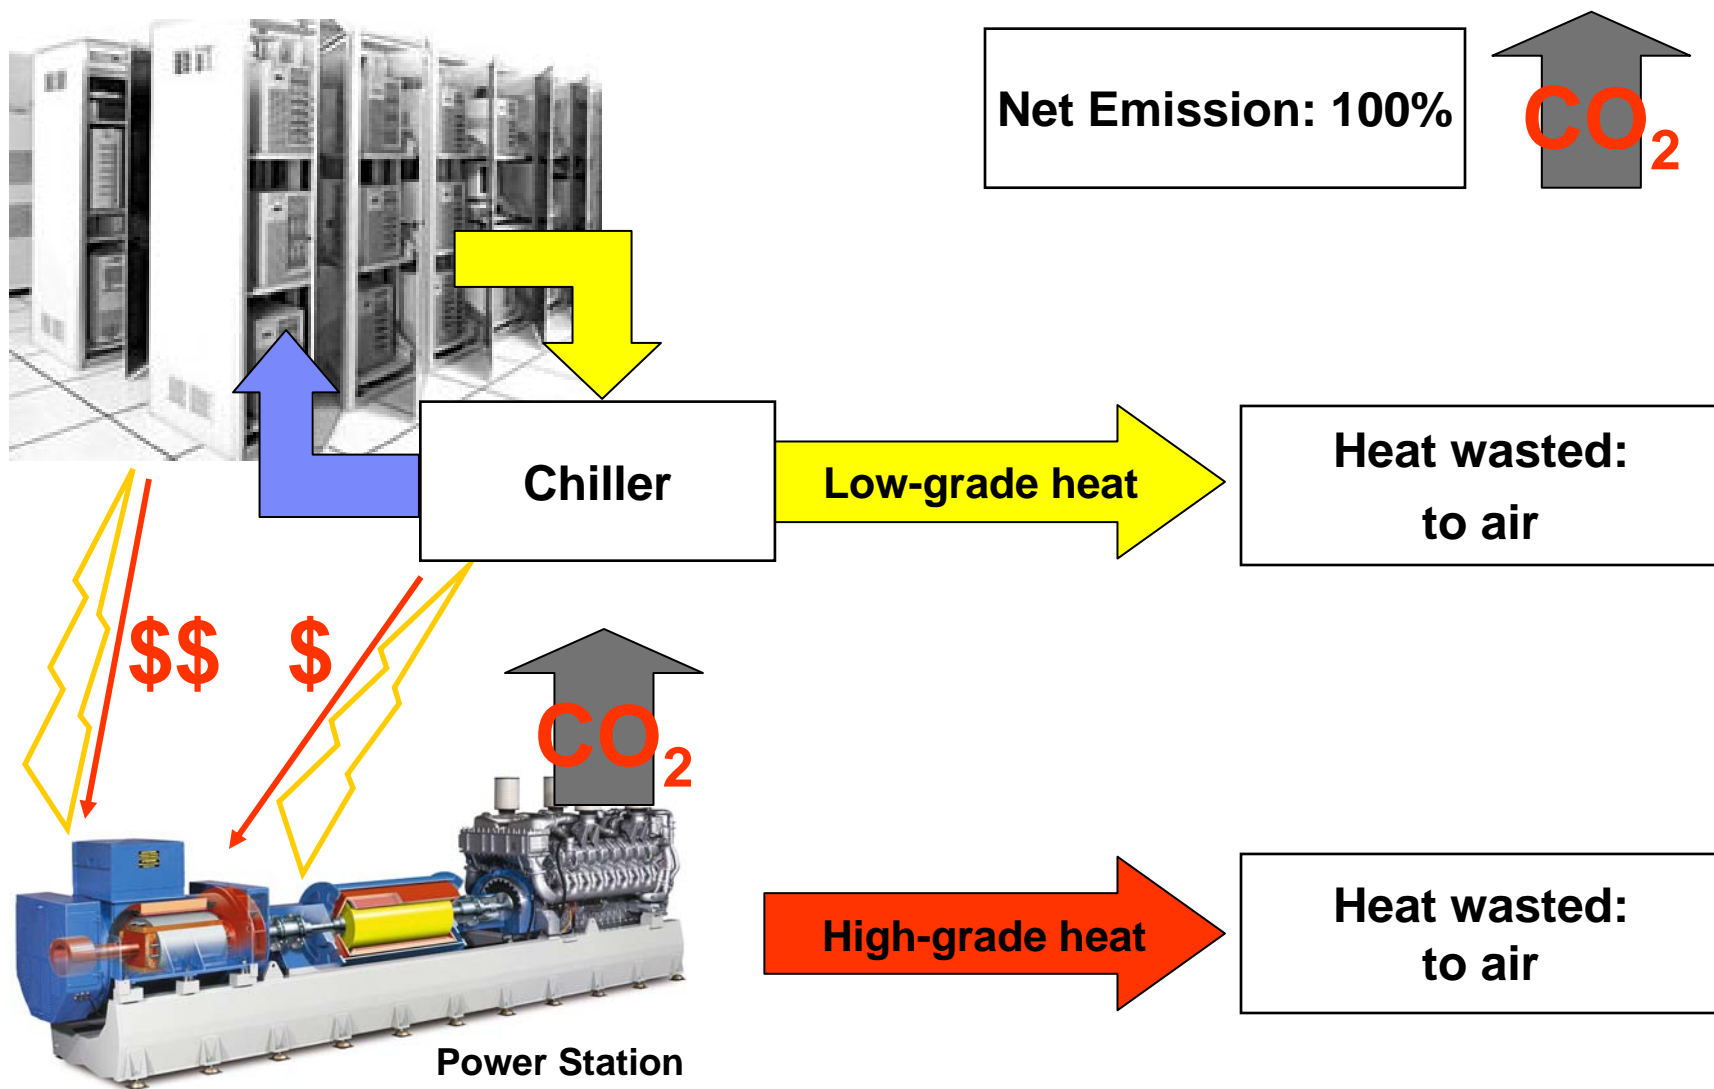
Thermal Energy Re-Use

- Zero-emission concept
valuable in all climates
 - Cold and moderate climates: **energy savings** (no chiller required) and **energy re-use** (for >60°C outlet, district heating)
 - Hot climates: **energy savings “only”** (no chiller required)

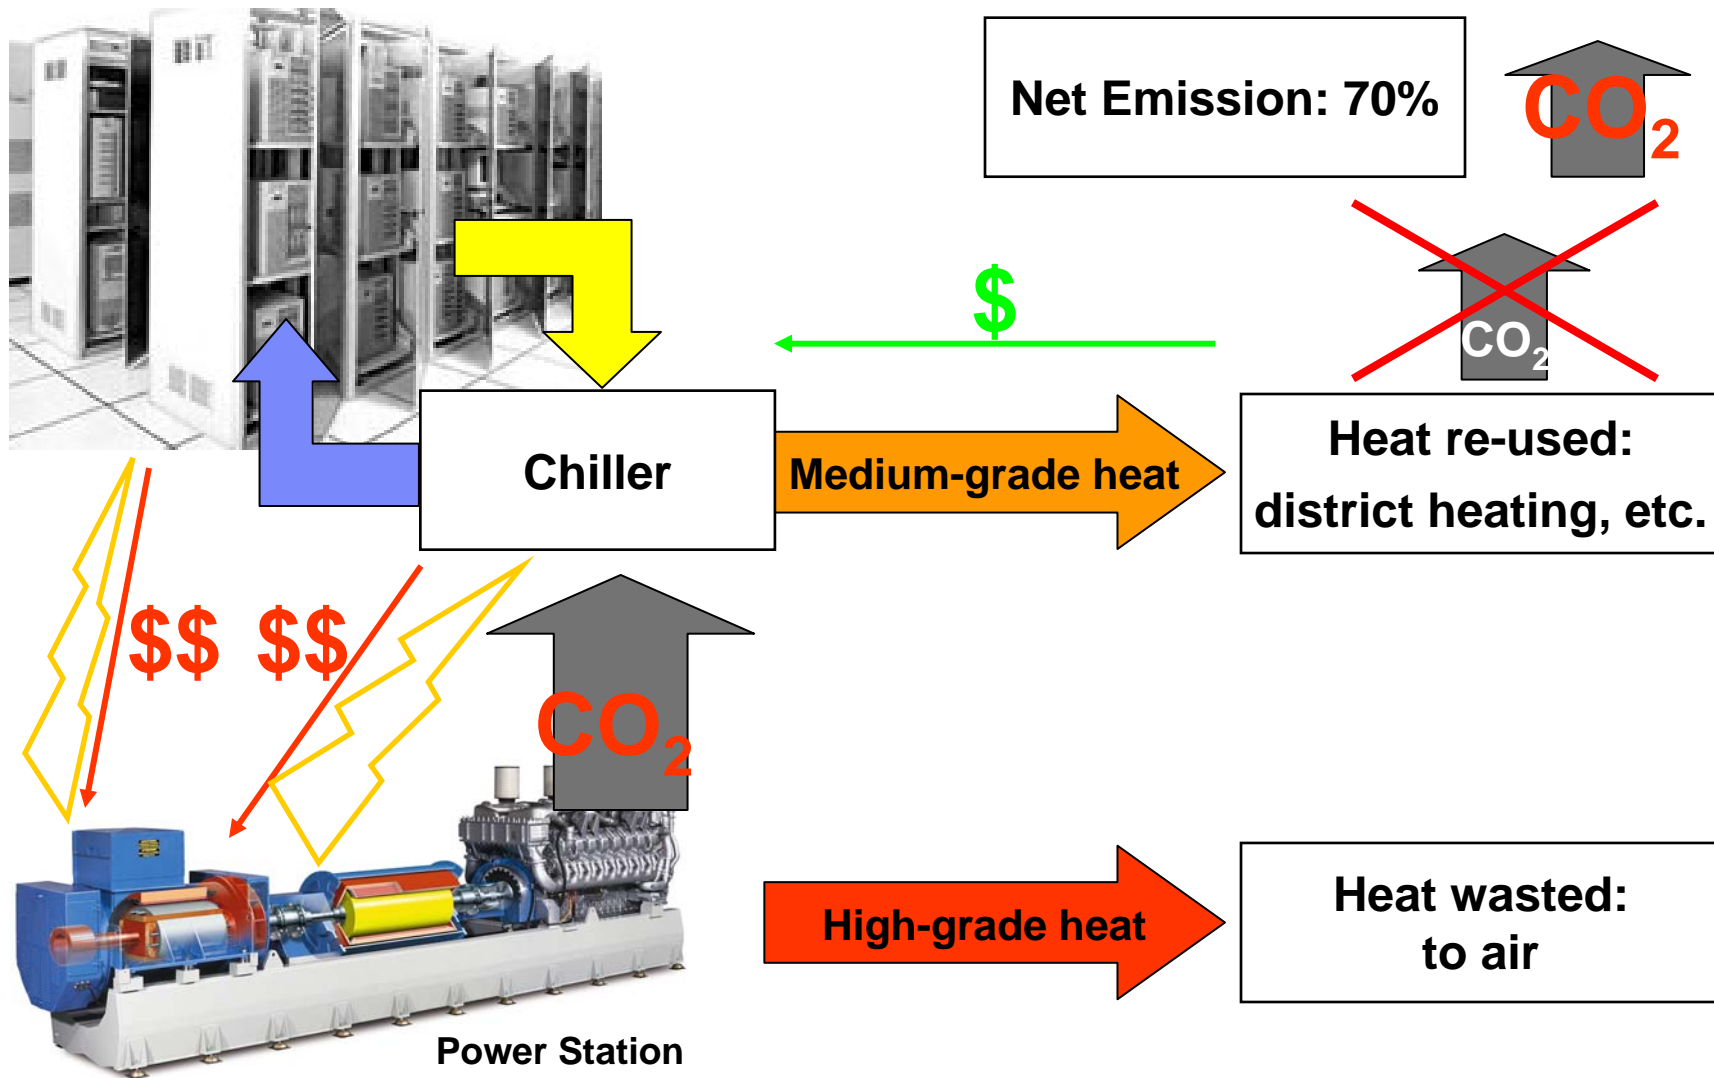
- Europe: 5000 district heating systems
 - Distribute 6% of total thermal demand
 - Thermal energy from datacenters easily absorbed
 - Largest sustainable energy source
 - **Thermal re-use: 3X (wind + solar)**




Energy and Emission Conventional Data Centers



Conventional Data Centers with Energy Re-Use



Zero – Emission Data Centers



Power & Heat Station

Net Emission: <15%

CO₂



\$



Medium-grade heat

Heat re-used:
district heating, etc.

\$



High-grade heat


Heat re-used:
Industry, etc.

Evolution of Liquid Cooled Systems



Left: Rear door cooler removing heat from air with cold water

Right: P575 with cold water processor cooling




- Water slowly approaches chips.....
- Water cooled CRACs / CRAHs in Datacenters
- Rear door coolers, intercoolers etc. in Racks
- The closer water comes the hotter it can be while having the same cooling performance


The QPACE Project

QPACE = Quantum Chromodynamics PArallel computing on CELL



- Research collaboration of IBM Development and European universities and research institutes
- Goal to build a prototype of a cell processor-based supercomputer.
- Funded by the German Research Foundation (DFG – Deutsche Forschungsgemeinschaft) as part of a Collaborative Research Center (SFB – Sonderforschungsbereich [TR55]).



The QPACE System




Warmwater Cooled System





- 56 coldplates/backplanes (8 per rack)
- $32 \times 56 = 1792$ nodecards
- 100% energy recovery with warm water operation 35°C (100 F)
- Medium thermal resistance
- Gottfried Goldrian, Michael Malms, Juergen Marschall, and Harald Pross
IBM Research & Development, Boeblingen, GER

First Prototype at IBM Rueschlikon

- Reduce cooling energy by tailored water cooling system
 - Cooling the chip with "hot" water (up to 60°C / 170 F)
 - Free cooling: no energy-intensive chillers needed
- Reuse waste heat for remote heating
 - The prototype reuses 75% of the energy for remote heating
 - Obtain recyclable heat (60°C) for remote heating.
 - Best in a cold climate with dense population
- Prototype
 - Similar Power of CPU and main board for air / liquid 60°C cooled version
 - Large fan power reduction
 - Liquid pump much more efficient and can vary flow at the rack level




Direct attached / integrated micro-channel cold plate with one interface




Experimental validation: Inlet temperatures up to 60°C / 170 F


Technical Objectives for Aquasar System



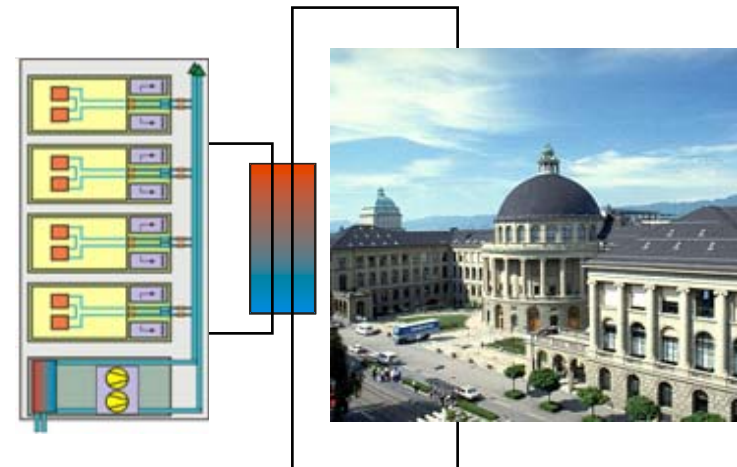
1. QS/HS22 Blades with fluid-loop



2. Fully populated BladeCenter® servers with manifold



3. Populated Rack with Blade Center® servers and pumps



4. Connection to heat distribution system of ETH for 50-60°C hot water

System uses a mixed population of 11QS22 IBM PowerXCell 8i and 3 H22 Intel Nehalem Blades per Blade Center® server.

Two Blade Center® servers are liquid cooled and one is air cooled for reference. The rack also holds communication equipment and a storage server. The closed cooling loop holds 10 liters of water, the coolant flow is 30 liters per minute.

Aquasar: Objectives and Deliverables

Objectives

- Use high-performance “hot”-water cooling to allow 2x energy cost reduction and large reduction of carbon emission
- Show that “zero”-emission datacenter operation is possible & profitable
- Validate concept and accelerate path to commercialization

Deliverables and Next Steps


- Zero-emission datacenter prototype (2 hot water cooled blade centers in a rack)
- Deploy system with 33 QS22 Cell blades + 9 HS22 Intel Nehalem blades at ETH
- Joint IBM-ETH-EPFL CCEM project to started August 2009
- 3-year CCEM project to optimize system for 100% energy recovery at $> 60^{\circ}\text{C}$ / 170 F
- Optimize efficiency and carbon footprint with different loads and clock speeds

- Use experience to create future standards and best practices for datacenter operation with Green Grid

Aquasar: Milestones and Status

Current Status


- Cooling loops for blades and blade centers designed and conversion to hot water cooling pending
- ETH location is being prepared to connect system to water cooling system



Blade with fluid-loop HS22 (Intel) and QS22 (Cell)

Milestones

- December 2009: Hot water cooled components assembled (blade center and rack)
- April 2010: Hybrid Cell/Nehalem System operative and initial parameters study completed and system connected to ETH energy re-use.




BladeCenter with manifold

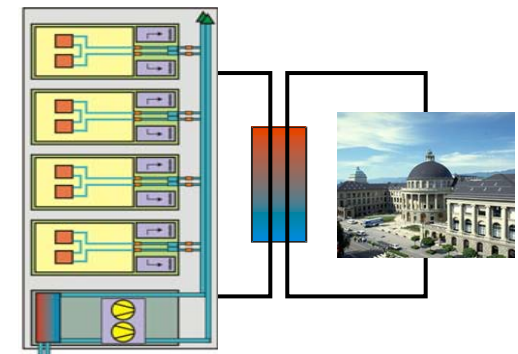
Target

- Reach world record in performance (MFlops/W) and low emission (MFlop/gCO₂)
- Lead standardization for future datacenters
- PUE_{reuse} less than 1

➔ This FOAK will deliver an innovative solution to run future datacenters



Populated rack with pump and metrology




Connection to ETH

Going Green Impact Tool




- Assess a clients data center with two basic scenarios:
- (A) Status quo and no green solutions are implemented
- (B) Green solutions are implemented
- Assess energy efficiency and financial impact of scenario A and B
- Simplified) total cost of ownership (TCO) analysis
- Run a ROI analysis of a project
- Contact: Hannes Engelstaedter



Summary

- Thinking global about energy usage
 - Total cost of ownership perspective
 - Demand and supply of sensible heat
 - Thermal energy re-use
- Energy re-use in liquid cooled data centers
 - Reduction of thermal resistance from junction to coolant by chip water cooling
 - Optimization of exergy efficiency
- Aquasar low-emission demonstrator
 - Cooling chips with “hot” water to obtain recyclable heat (65°C / 178 F) for remote heating
 - Best in a cold climate with dense population
 - Free cooling in all climates: No energy-intensive chillers needed
- Aquasar specifications
 - Saves 40% of energy
 - Reduces emission by 85% through heat re-use on ETH campus



Main Messages and Next Steps

- **Key Component for Future Datacenters**
 - *Chip cooling technology exists but needs to be combined with current computers*
 - *Centralized computing more efficient and emission free*
- Roadmap for large efficiency increase in 10 years
 - *3D interlayer cooling and electrical-thermal co-design*

Next Steps


- Optimization of coolant temperature as function of demand and supply
- **Scale up to full size HPC and business data centers: We are ready are you ready as well?**
- Thermal energy re-use in solar collectors



Thermal Packaging Leverage in Concentrated Photovoltaic

Concentrated photovoltaic

- Expensive triple junction cells with 41% peak cell-efficiency
- Chip cost leverage by sun concentration (today 50 to 200x)
- Concentration limited by junction temperature (efficiency, reliability)



Heat flux (@1000x): $100\text{W}/\text{cm}^2$

Liquid-cooled CPV-unit:
 high-performance thermal management
 enables **1000x sun concentration**
 → Potential for cost reduction

Re-Use: Concentrated Photovoltaic serves Desalination

Worldwide socio-ecological challenge of the 21th century

- Energy demand
- Fresh water supply



External source


CPV-array

electricity

heat

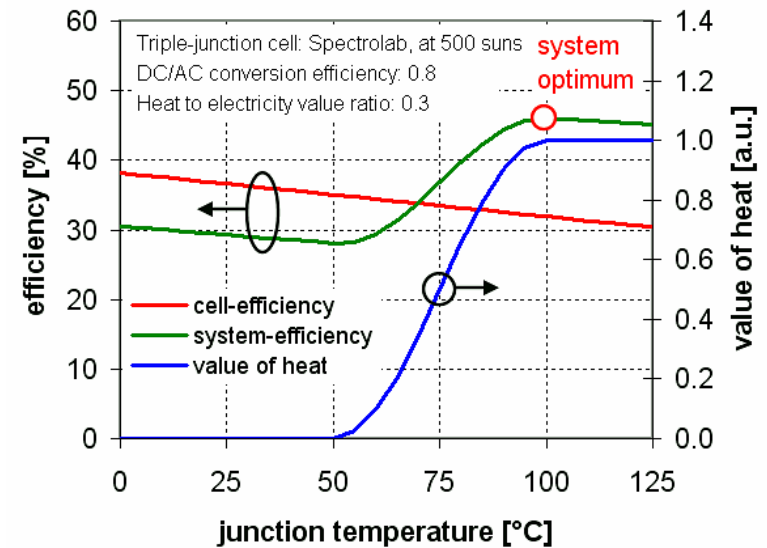


Power grid



Multi-effect boiling (MEB)

System-efficiency: CPV with MEB



→ Locations with high sun irradiation and scarce fresh water resources coincide

Acknowledgment

- Advanced Thermal Packaging Group Members
Thomas Brunschwiler, Werner Escher, Wulf Glatz, Javier Goicochea, Ingmar Meijer, Stephan Paredes, Brian Smith, Reto Waelchli, Ryan Linderman
- Microfabrication team
Rene Beyeler, Daniele Caimi, Ute Drechsler, Urs Kloter, Richard Stutz, Kurt Wasser, and Martin Witzig
- IBM Boeblingen Research and Development GmbH (Germany)
Gottfried Goldrian, Michael Malms, Juergen Marschall, Harald Pross, and Wolfgang Zierhut
- IBM Research Yorktown (USA)
Paul Andry, Evan Colgan, Claudius Feger, Winfried Haensch, Hendrik Hamann, Theodore vanKessel, Ken Marston, Yves Martin, John Maegerlein, and Thomas Theis
- IBM Server and Technology Group in East Fishkill and Poughkeepsie (USA)
Pepe Bezama, Kamal Sikka, Michael Ellsworth, Roger Schmidt, and Madhu Iyengar
- IBM Austin and other locations (USA)
Dave Frank, Vinod Kamath, Hannes Engelstaedter
- Financial Support:
IBM Zurich Research Laboratory, IBM Research FOAK program, Swiss Government KTI Projects, EU FP7 project NanoPack

Thank you for your attention

Literature and Links

- Youtube Video on Zero-Emission Datacenter: <http://youtube.com/watch?v=1J7KpgozpRs&feature=user>
- T. Brunschwiler, B. Smith, E. Ruetsche, and B. Michel, “Toward zero-emission datacenters through direct reuse of thermal energy”, IBM JRD 53(3), paper 11; see: <http://www.research.ibm.com/journal/>
- H. Engelstaedter, “Finding the Right Green IT Strategy Is a Great Challenge—The ‘Going Green Impact Tool’ Supports You,” Smart Energy Strategies: Meeting the Climate Change Challenge, vdf Hochschulverlag an der ETH Zurich, 2008, pp. 48–49; see http://www.vdf.ethz.ch/service/Smart/Smart_Energy_Strategies.pdf
- T. Brunschwiler, H. Rothuizen, M. Fabbri, U. Kloter, B. Michel, R. J. Bezama, and G. Natarajan, “Direct Liquid Jet-Impingement Cooling with Micro-Sized Nozzle Array and Distributed Return Architecture,” Proceedings of the IEEE ITherm Conference, San Diego, CA, 2006, pp. 196–203.
- “Climate Change 2007—The Physical Science Basis: Contribution of Working Group I to the Fourth Assessment Report of the IPCC,” Cambridge University Press, Cambridge, ISBN 978 0521 70596-7.
- S. Hamm, “It’s Too Darn Hot,” Business Week, March 31, 2008, pp. 60–63.
- “Liquid Logic”, The Economist Technology Quarterly, September 2008.