Ultimately Dense and Efficient Future Computers


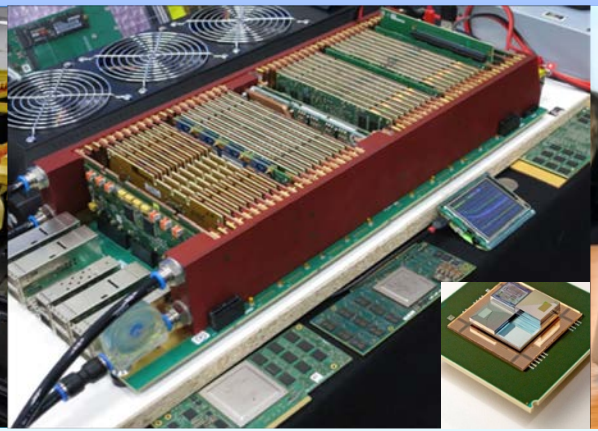

Bruno Michel, Mgr. Smart System Integration, Fellow IEEE
Member US National Academy of Engineering and Member IBM Academy of Technology
https://www.zurich.ibm.com/st/ https://www.zurich.ibm.com/st/smartsystem/

Mega Trends with Implications

Datacenter Carbon Footprint

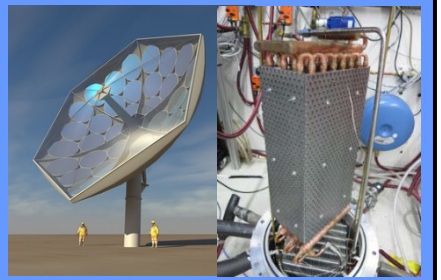

Renewables Energy Cost and Climate

Big Data

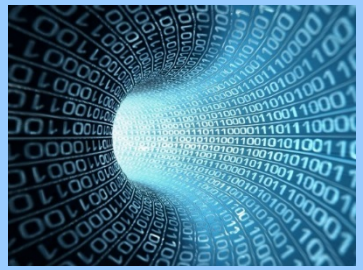

Cognitive Computing

The End of Transistor Scaling

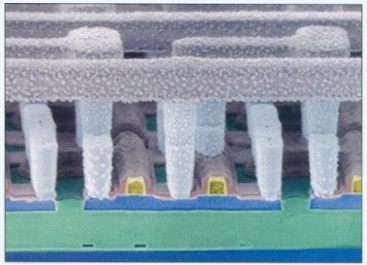

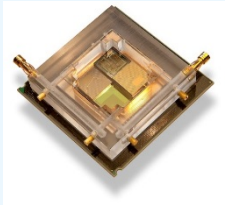

Dense and Efficient Systems Roadmap

The Internet of Things IoT

Wearables and Healthcare

Bruno Michel, bmi@zurich.ibm.com

Low Datacenter Efficiency

- **Efficiency is 0.000'004%**
- **Volume used for compute is <1 ppm (part per million)**

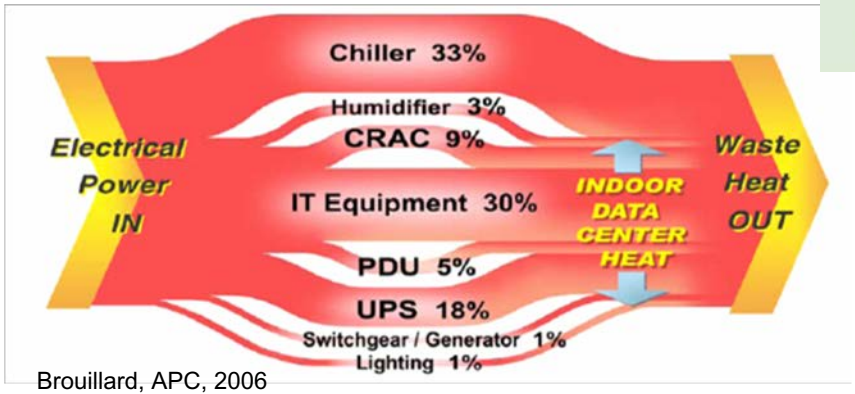

Brouillard, APC, 2006

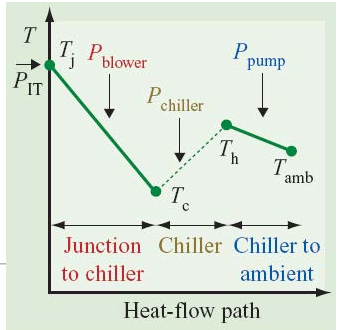

Chillers (Refrigeration)

Evaporative Tower Fans

Racks & Fans

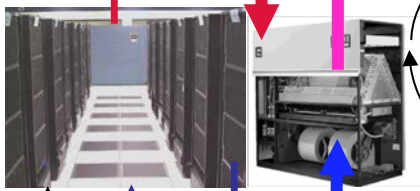

Electrical Power

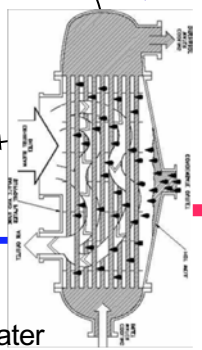

Chilled Water CRAC/ CRAHs

Condenser

- **A computer is an inefficient "joule heater" "producing" 10-20ºC "heat".**
  The output is 0.000'004% ICT and 99.999'996% "negative grade" heat
- **Consequence: We have a lot to do to change energetics and density!**

Bruno Michel, bmi@zurich.ibm.com

Biological Concepts Transform IT Industry

Hot-Water-Cooled Zero Emission Datacenters „Aquasar“

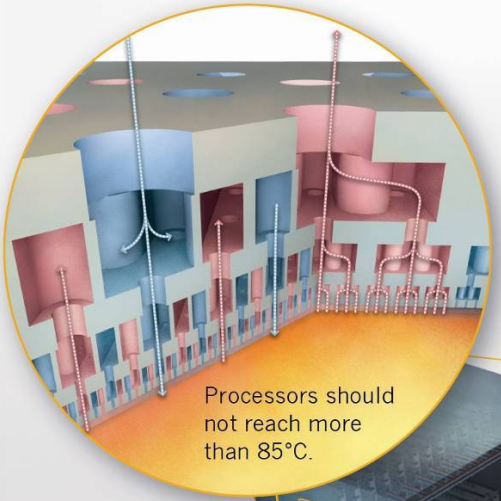

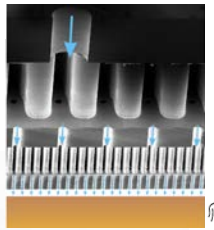

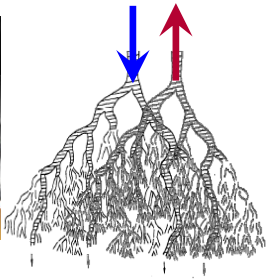

Micro-channel liquid coolers

Heat exchanger

Water Pump

CMOS 80ºC

Heat exchanger

Direct „Waste“-Heat use e.g. heating

Water Out 65ºC

Water In 60ºC

Biological inspired:
Vascular systems optimized for low pressure transport

SuperMUC in Munich and Summit in Oak Ridge

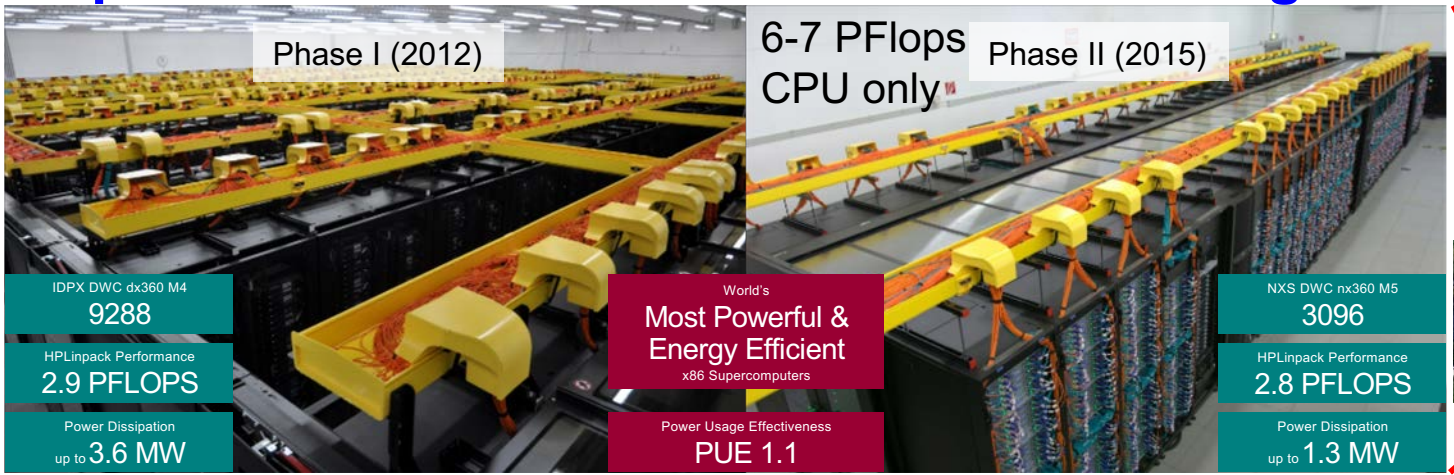

CoolMUC-2 (2016)

200 Racks ↑ ↓ 240 Racks ↑ 40 Racks

~1 Mio t CO_2 saved

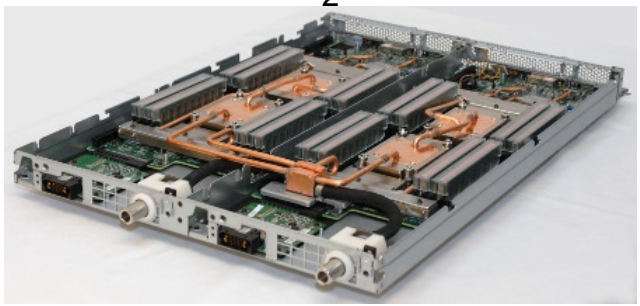

Bruno Michel, bmi@zurich.ibm.com

Data Center Evolution … @ 100 TFLOPS

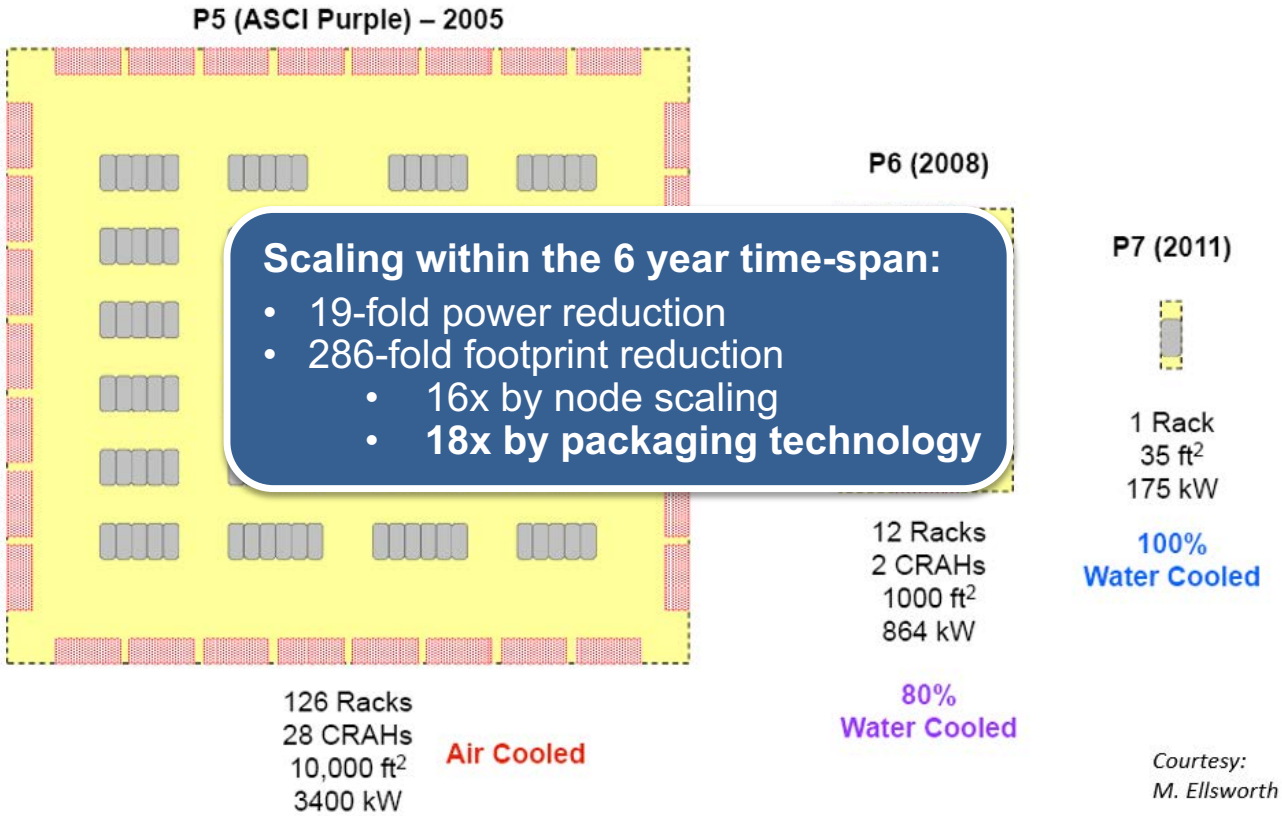

Bruno Michel, bmi@zurich.ibm.com

(R)evolution of Information Technology

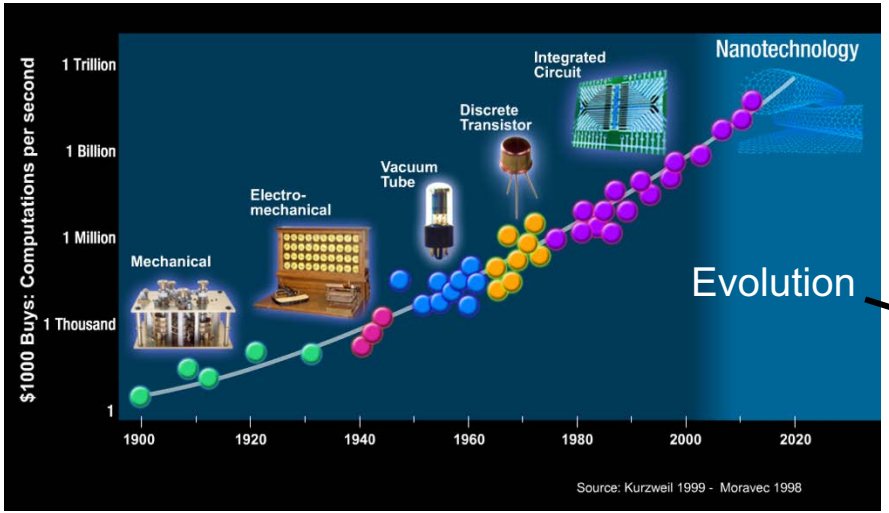

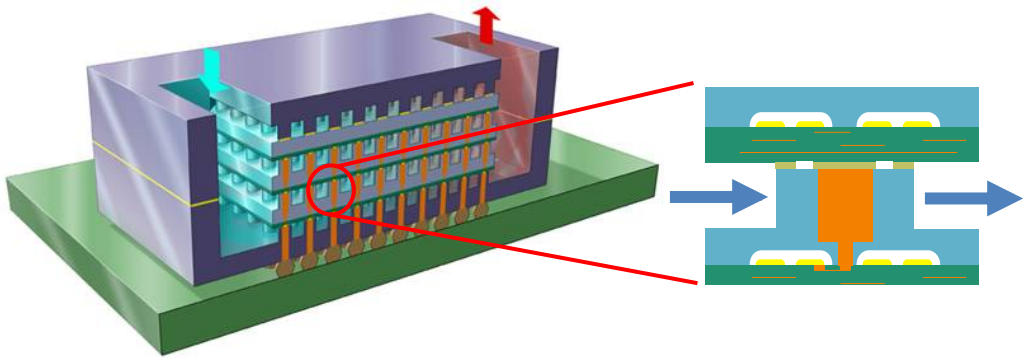

3D Integration

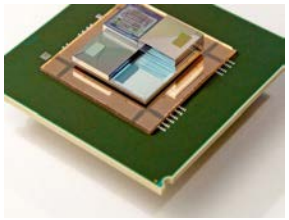

Brain: synapse network

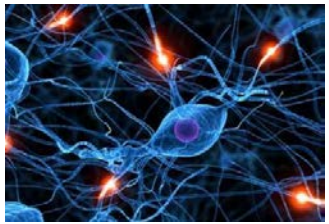

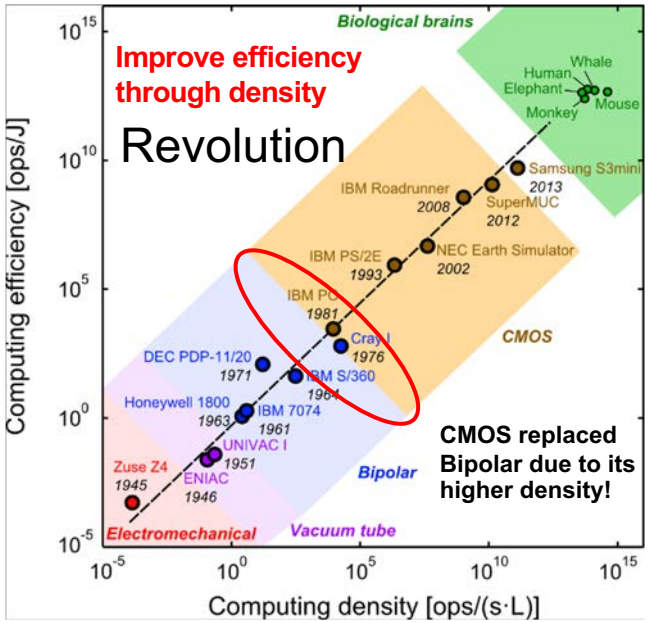

Bruno Michel, bmi@zurich.ibm.com

How are Thermal Resistances Massively Reduced?

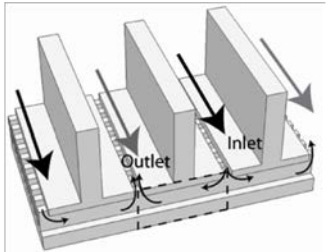

Transition from air cooling to liquid cooling

- Water with 4000x better heat capacity and 30x better thermal conductivity

A) Massively reduce convective thermal resistance

- **Microchannels (Manifold Microchannel)**
  - High aspect ratio ➔ massive surface enlargement
- **Branched hierarchical transport**
  - Best mass transport with minimal pumping power

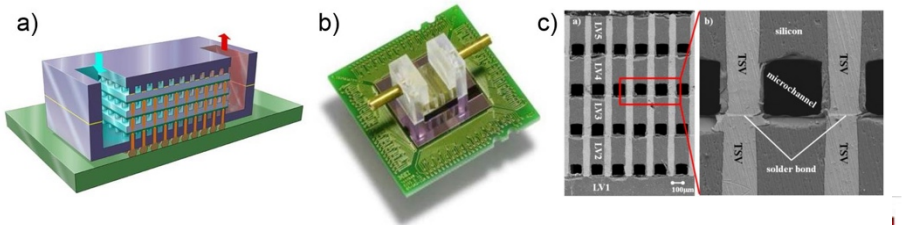

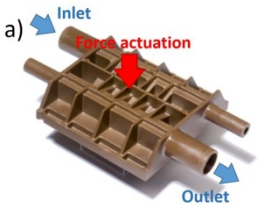

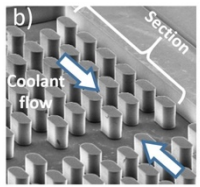

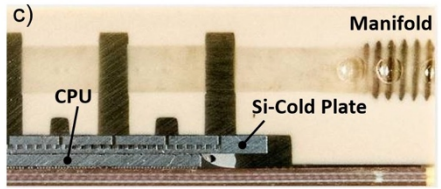

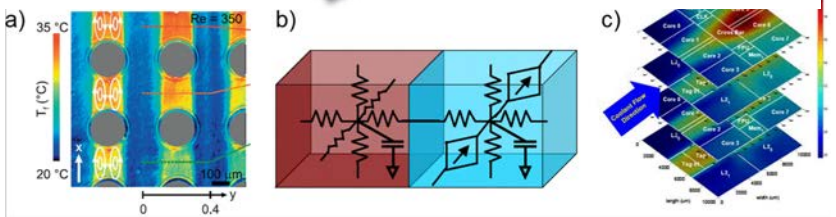

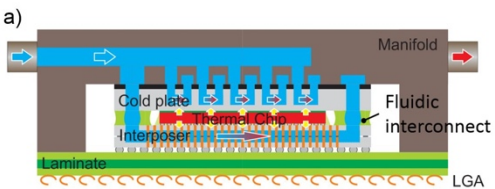

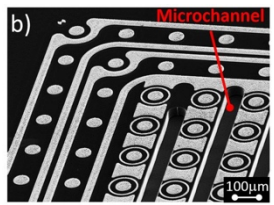

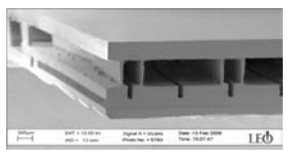

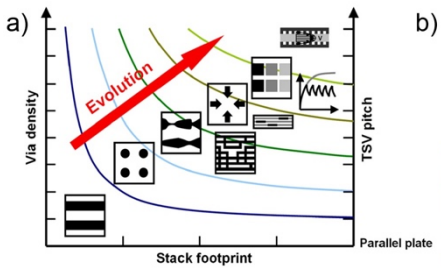

Bruno Michel, bmi@zurich.ibm.com

How are Thermal Resistances Massively Reduced?

B) Massively reduce conductive thermal resistance

- **Radical miniaturization and use of silicon**
  - Good thermal conduction with less thermal interfaces
- **Better interfaces**
  - Filling materials/gaps with percolation and improved overall thermal conduction due to necking

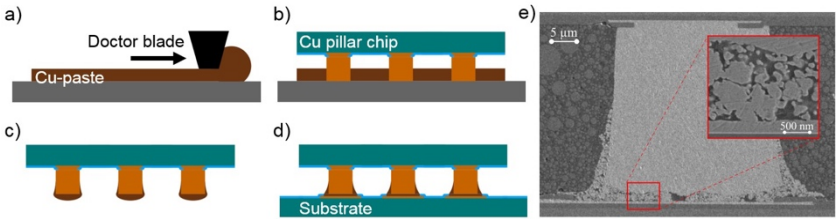

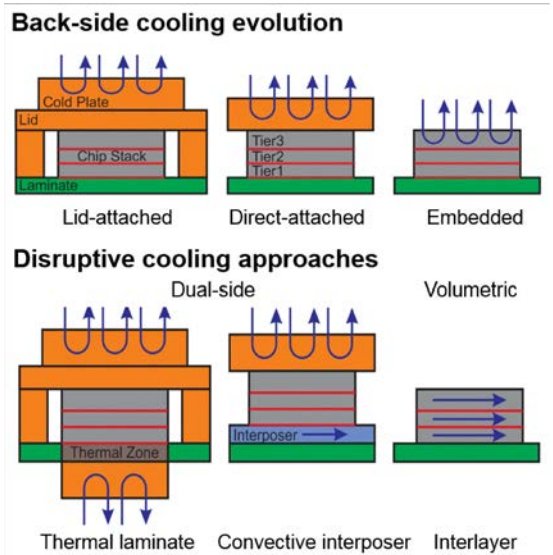

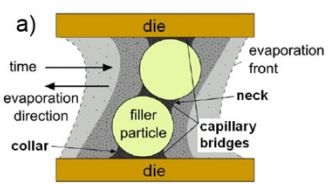

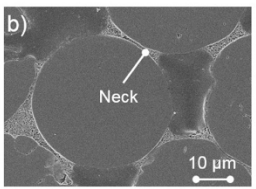

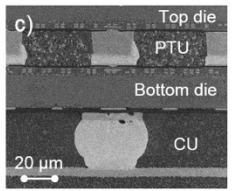

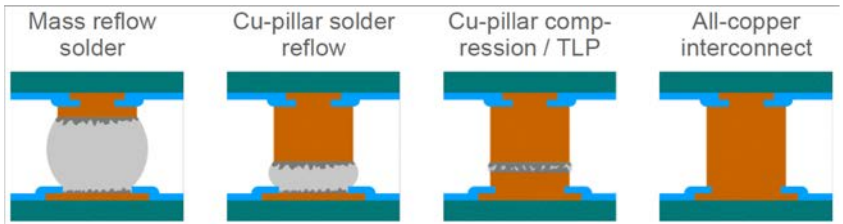

Bruno Michel, bmi@zurich.ibm.com

All Copper Interconnects Reduce Electrical resistance

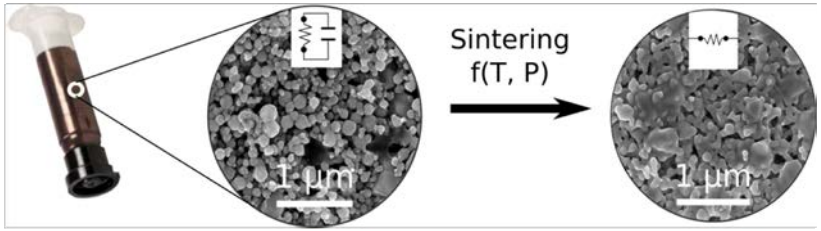

Syringe (5 ml)

- Tdwell > 130 ˚C
- $d_{particle}$ < 2 µm

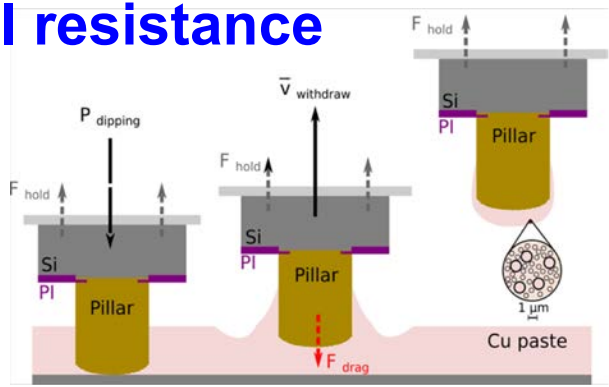

Widely applicable all-copper interconnects by sintering copper nanoparticles.

Robust dip-transfer for 20 µm pillar pitches.

Densification by bonding pressure during sintering reducing electrical resistance close to bulk copper and increasing shear strength by 4x.

1.2±0.4 mΩ and 24 Mpa shear strength obtained at 200°C with 50.6 MPa bonding pressure.

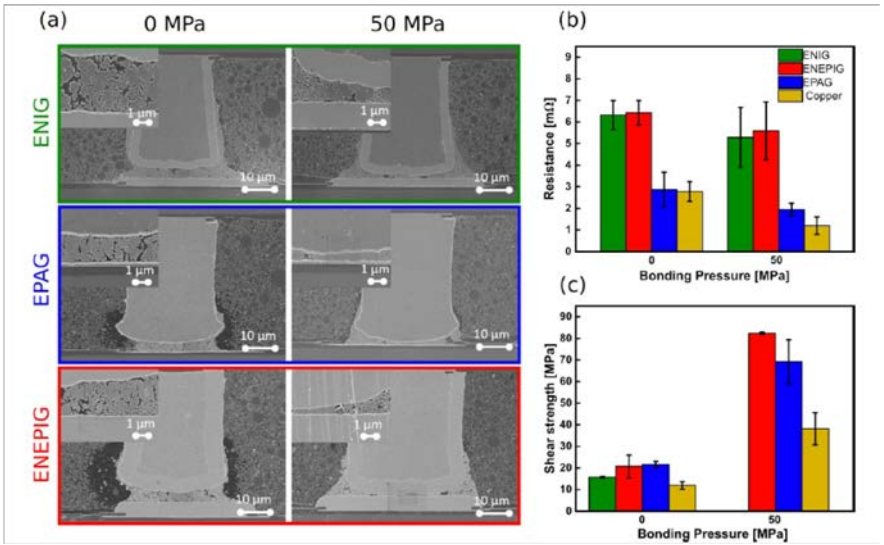

Bruno Michel, bmi@zurich.ibm.com

Scaling to 1 PFlops in 10 Liters

- **Efficiency comparison**
  - 1PFlops system currently consumes ~10MW
  - 0.1 PF ultra-dense system consumes 20 W
- **Ultra-dense Bionic System**
  - Stack ~10 layers of memory on logic
  - Stack several memory-logic stacks to stack of stacks
  - Combine several blocks of stacks to MCM (MBM)
  - Combine MCMs to high density 3D system
- **Key enabling technologies**
  - Interlayer cooling
  - Electrochemical chip power supply
- **Impact**
  - **5'000x smaller power**
  - 50'000'000x denser
  - Scalability to zetascale

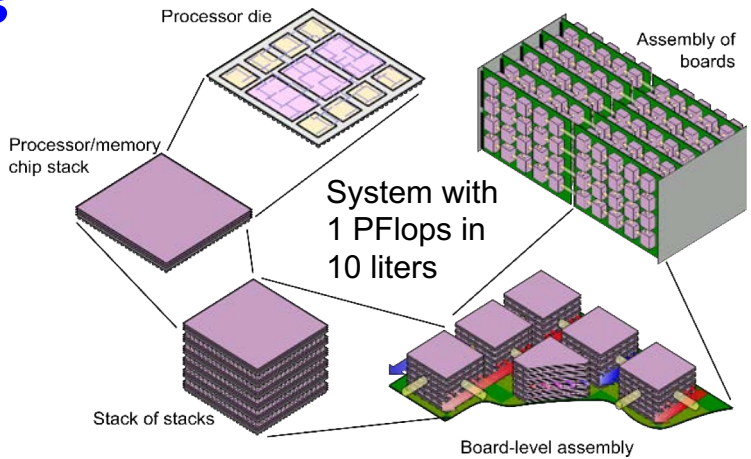

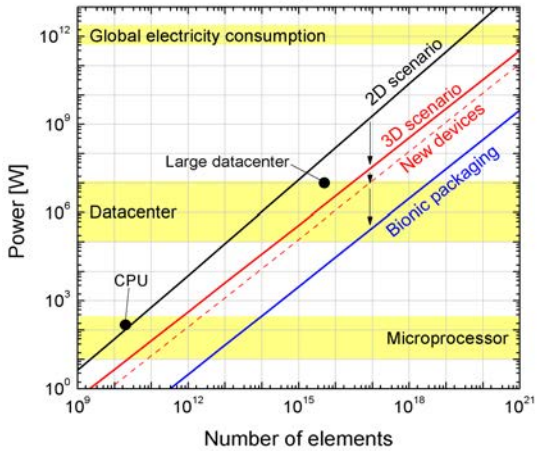

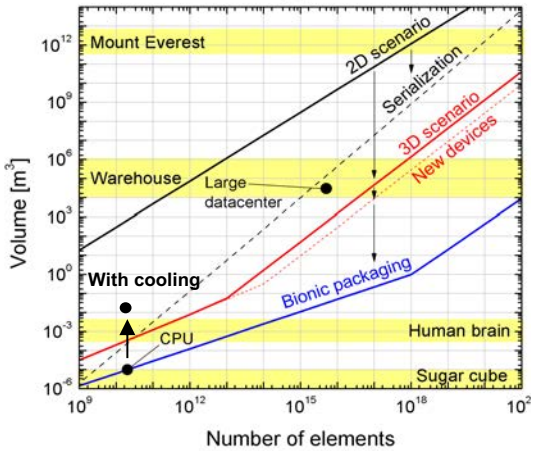

P. Ruch, T. Brunschwiler, W. Escher, S. Paredes, and B. Michel, "Towards 5 dimensional scaling: How density improves efficiency in future computers", IBM J. Res. Develop. 55 (5, Centennial Issue), 15:1-15:13 (2011).

Bruno Michel, bmi@zurich.ibm.com

Brain Inspired Computing

- Stepwise introduction of brain inspired concepts: Form – Function – Material
- Step 1 (Form): Brain inspired packaging with classical CMOS ➔ Now
- Step 2 (Function): Brain-inspired, non-von Neumann architecture ➔ Ongoing
- Step 3 (Material): DNA computing … ➔ Later
- Chess, Jeopardy, and Go games against humans were unfair why?
- Computers are challenged more by “open” (Jeopardy) than “closed” system (Chess, Go)

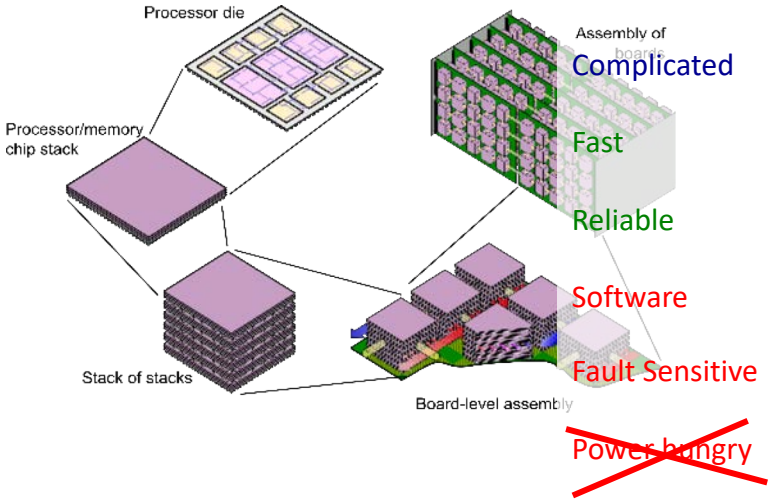

Computer

vs.

Human in Chess and Jeopardy

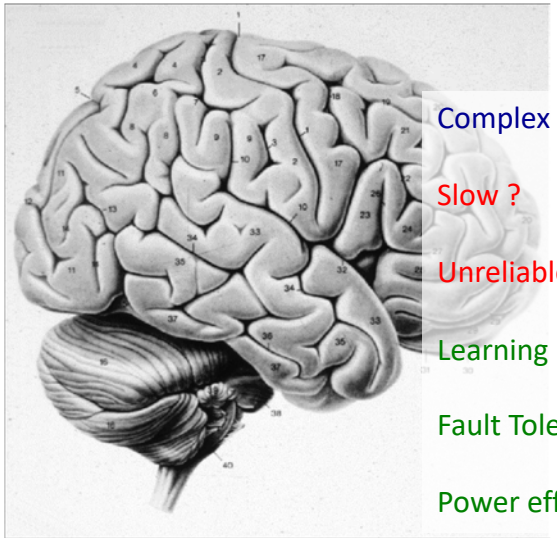

Bruno Michel, bmi@zurich.ibm.com

CarrICool Project (EU-FP7, 2014-2017)

PMIC
32nm IBM SoI process/
14nm GF Bulk

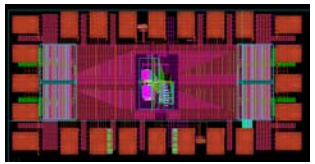

A silicon interposer that:
provides liquid cooling
delivers power using a VRM
provides a high-speed optical link

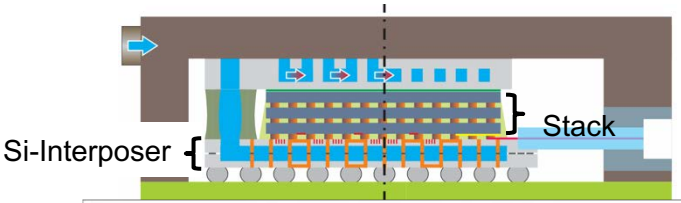

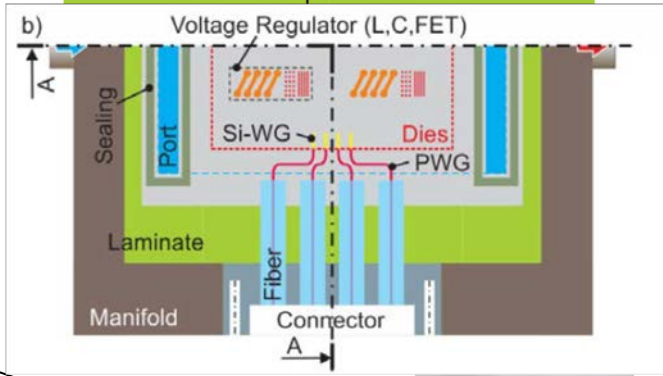

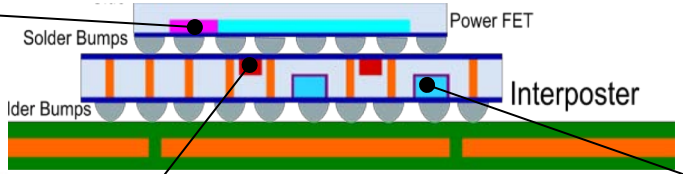

Deep Trench Capacitors
IPDIA MIMIM (>200nF/mm²)

Final assembly

Air-core Inductors
RDL layers on the interposer

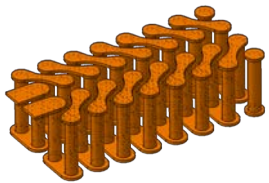

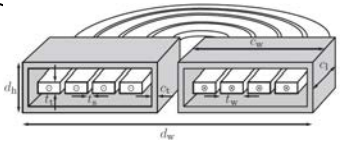

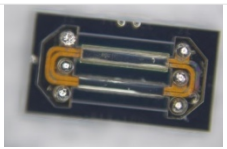

Inductors
Planar or Volumetric Inductors using
TSVs with Magnetic core material added

Densely integrating WBG switches like GaN with CMOS driving/control can enable higher conversion ratios while maintaning small form factors

CarrICool Hybrid iVR Implementation

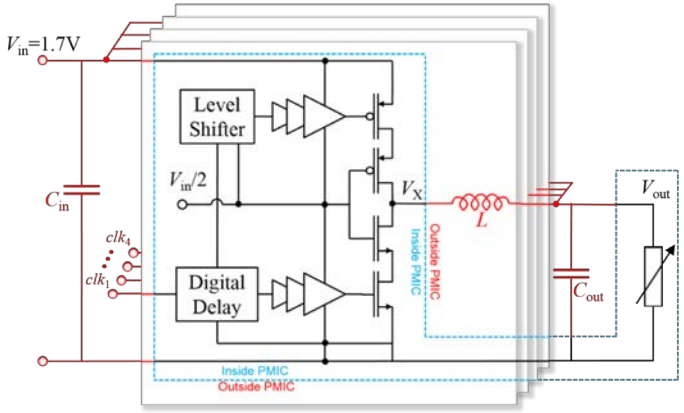

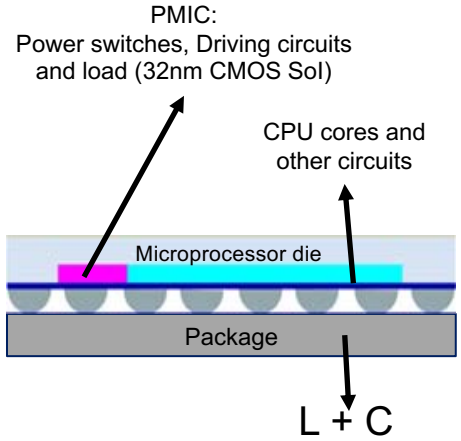

Parameter	Description	Value
V_{in}	Input voltage	1.6 V
V_{out}	Nominal output voltage	0.8 V
$P_{out, max}$	Max output Power	800mW
η	Efficiency	≥ 90%
PD_{Chip}	Power density on PMIC	≥ 30W/mm^2
$PD_{Interposer}$	Power density on Interposer	≥ 1W/mm^2
$\Delta V_{out, pp}$	Max ripple	8 mV
$t_{sett, load}$	Settling time	40 ns

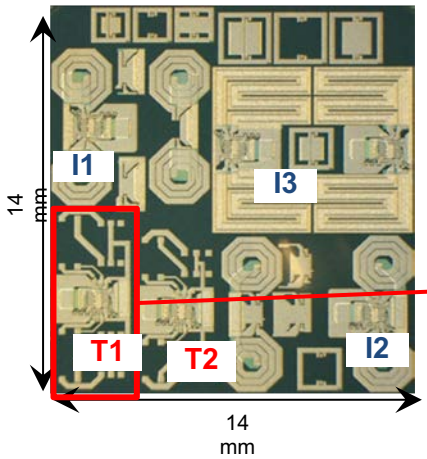

Multiple demonstrators built on a single interposer tile

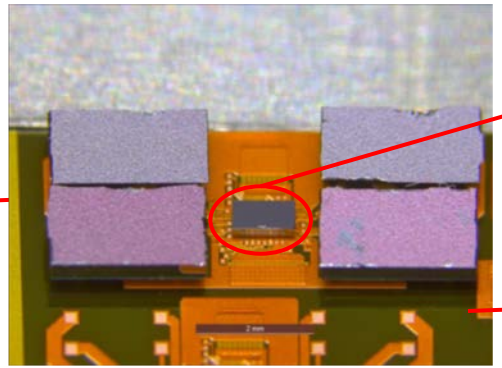

Au stubbing + Thermocompression

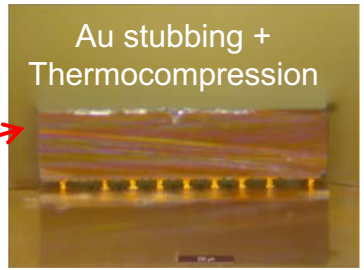

Low T solder balling + reflow

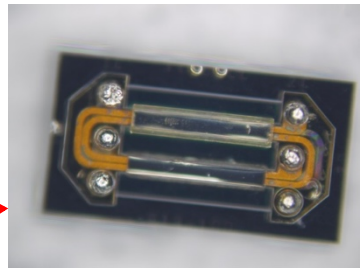

Datacenter in a Box: Top down Densification

Leading cooling and power delivery concept for Dome project

Increase density using hot water cooling structure for power delivery
➔ Density: Key differentiator

1000x denser and 10x more efficient!!

Density enabled by reduced thermal and electrical resistance

Efficiency and low cost due to Bell's law

Target system
- T4240 1.8GHz node
  - 12 cores 2 threads
  - 48GB of DRAM
  - **220 GFLOP**
- **2U rack unit**
  - 128 nodes
  - 3072 Threads
  - 6 TB DRAM
  - **28 TFLOP**

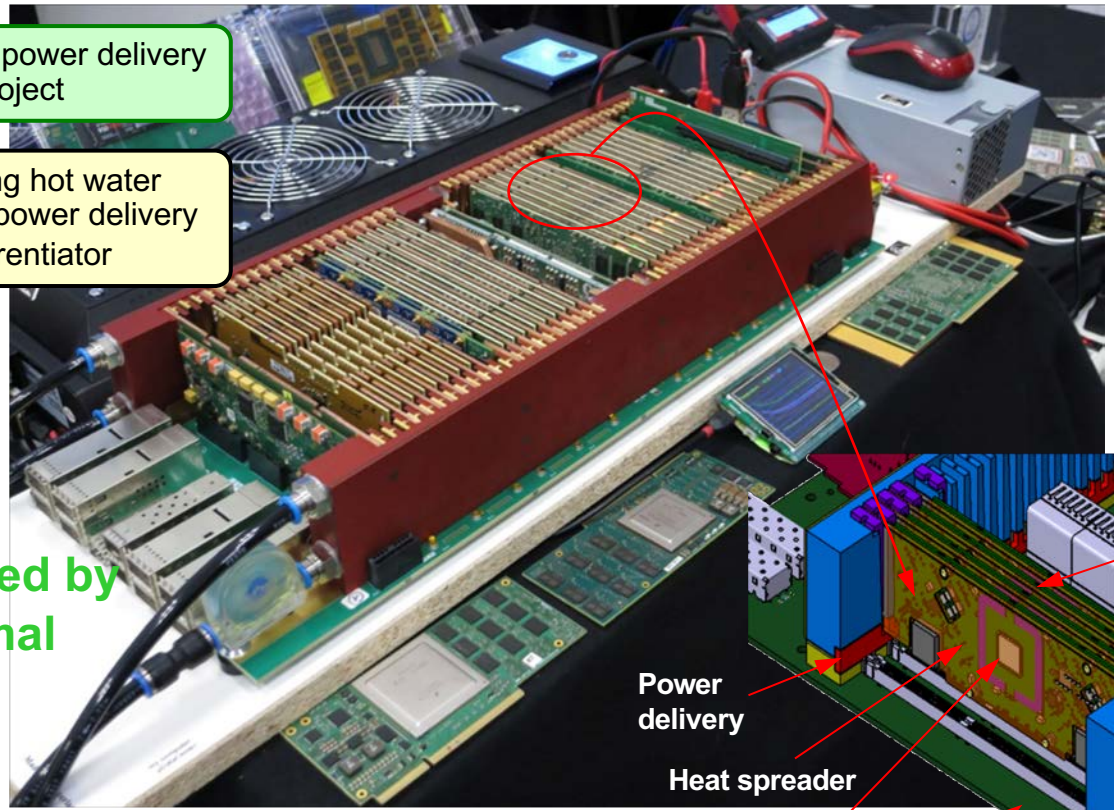

Bruno Michel, bmi@zurich.ibm.com

Bell's Law Demands more Integration

- Every 12-15 years restart of a new generation
- Hardware cost fraction decreases from 100% (mainframe) to <10% while adding functionality


Figure. Jeff Holter with his original 38-kg radio-electrocardiograph recording device in 1947.

- Sensing and communication were also miniaturized
- Mainly driven by packaging
- Sensing and computing meet in wearables
- **Remember: proximity improves efficiency!**

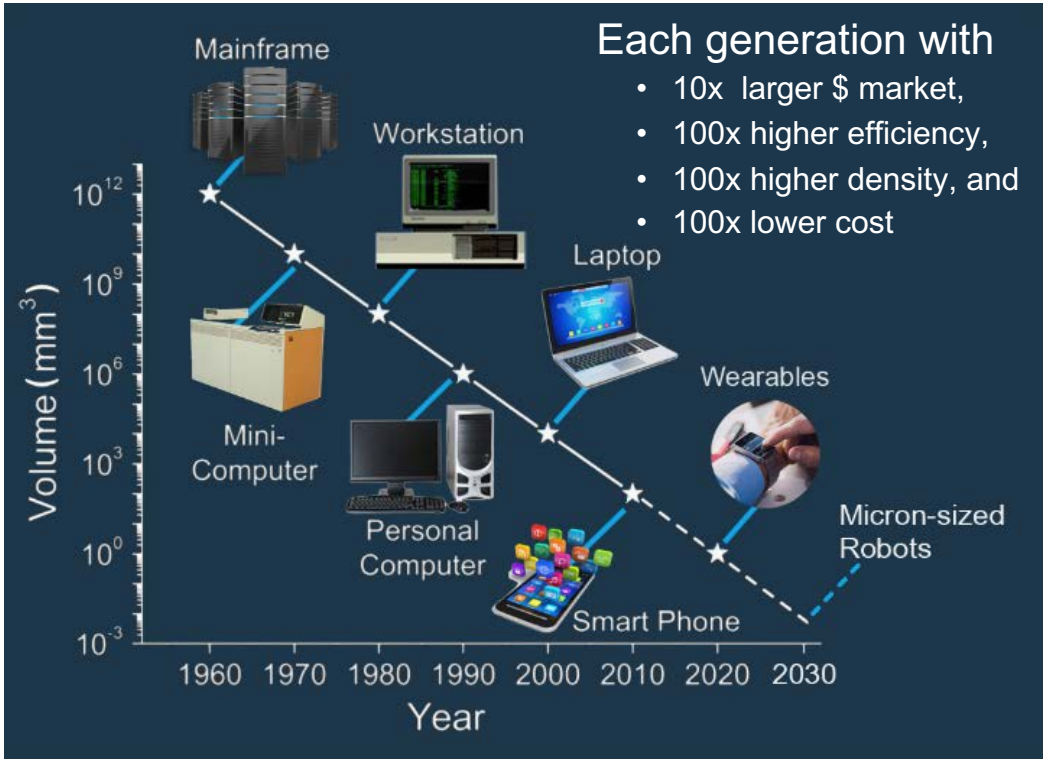

- **Data not transferred to cloud but AI transferred to the edge to improve efficiency**

Bruno Michel, bmi@zurich.ibm.com

Human Centric Sensing and Computing

- Cost and accessibility of healthcare, blockbuster drugs not personalized
- Stress strong link to human wellbeing
- **Human Centric Sensing and Computing:** Context key for relevant personalized cognitive services
- Personalized cognitive services in preventive medicine / coaching; work safety; wellbeing; elderly care
- Miniaturization for low-cost non-intrusive monitoring to reduce cost in acute and preventive medicine
- **Move intelligence to the edge instead of data to the cloud for solutions to be relevant to people**

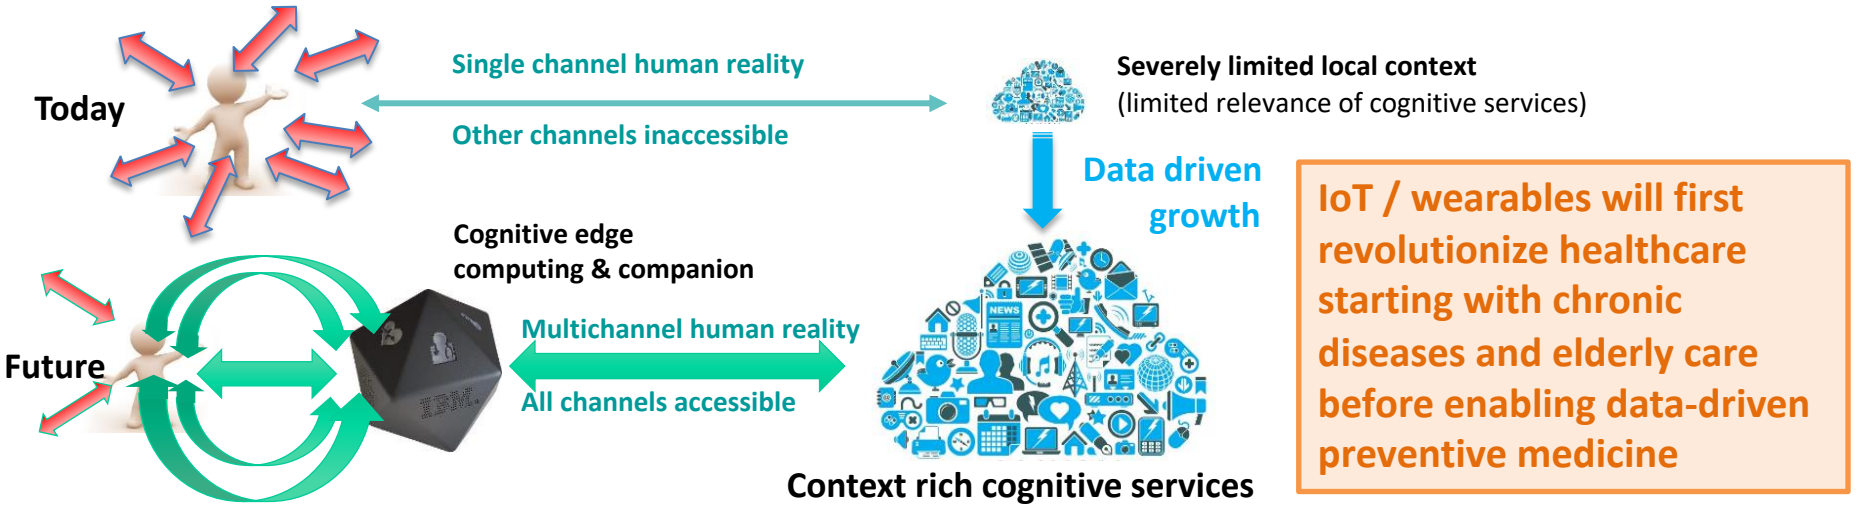

Bruno Michel, bmi@zurich.ibm.com

IoB Platform Building blocks

Humans as largest data source and largest consumer of cognitive services

***modular, flexible and scala-ble** platform that adapts to use cases and strengthens base of AI pyramid*

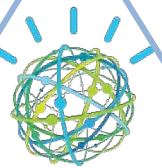

Data Science
- SML Algorithms
- Deep Learning
- microservices

Data Collection/Preparation
- Gathering protocol
- Reliable Data
- Data Storage
- Data cleaning
- Segmentation, labeling
- Feature selection

Human in Center and Pyramid Base
- Physiology / Medicine
- Psychology / Behavioral change
- Acceptance / Usability / User experience

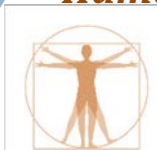

User Interface
- *Data*
- *Thresholds and trends*

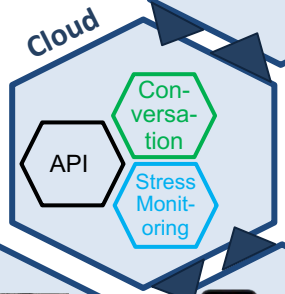

Cloud Services
- *Storage*
- *Big data science*

Data Hub and Sensors
- *Edge*
- *Communication*
- *Alerts*
- *Commercial and new sensors*

Use Cases
- *eCompanion – Hospital of the Future*
- *CAir – Chronic Disease Mgmt.*
- *FireFighter – Stress Detection*

Bruno Michel, bmi@zurich.ibm.com

eCompanion: IoB Platform Implementations

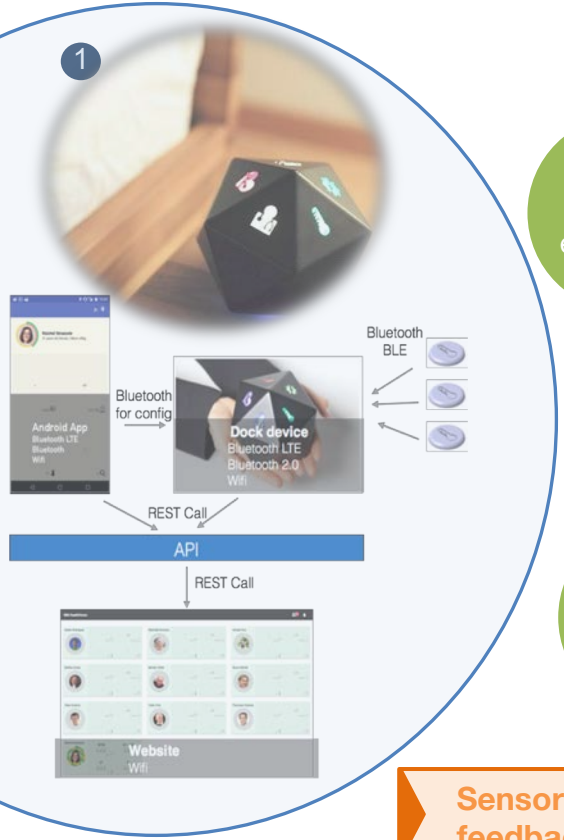

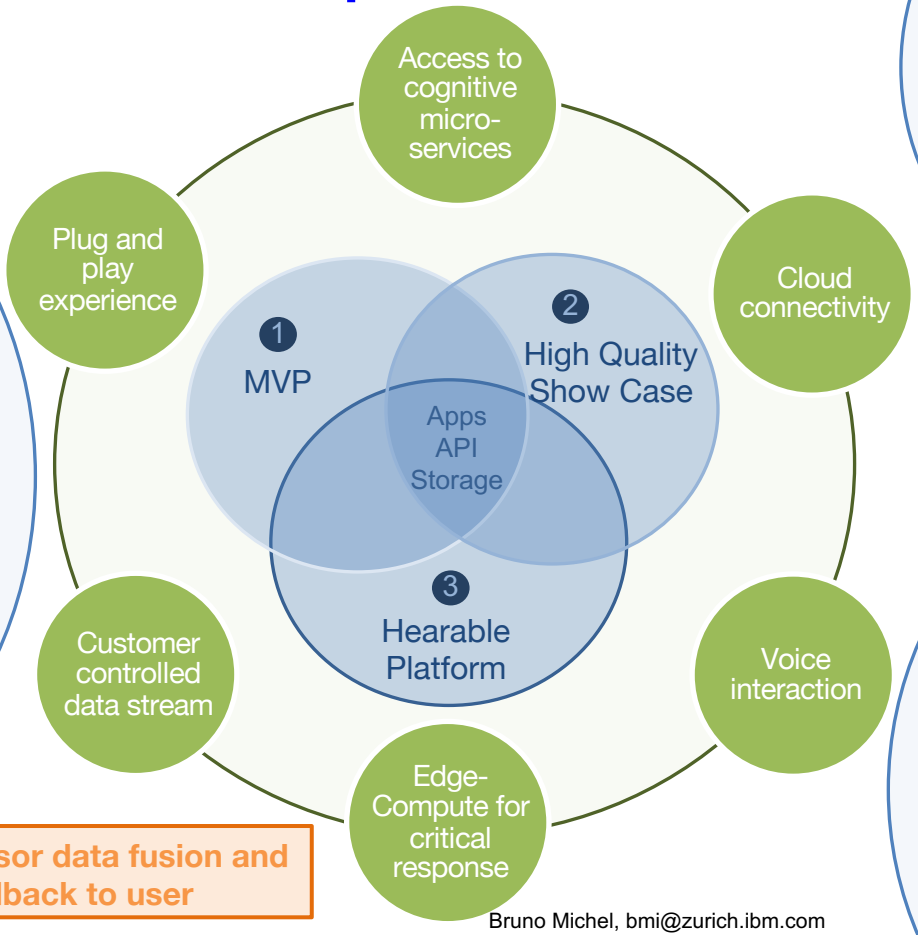

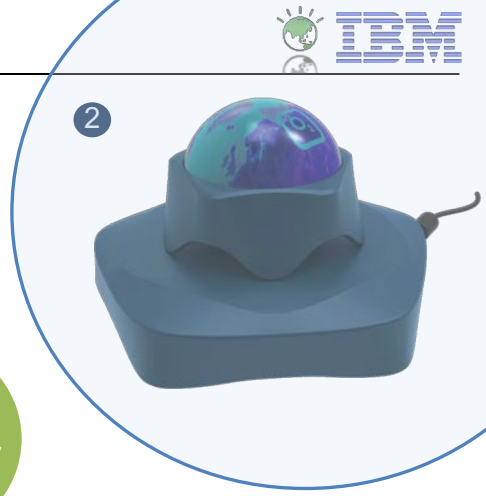

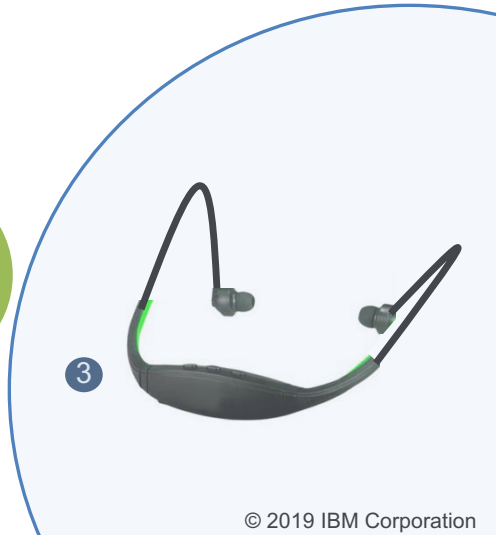

Sensor data fusion and feedback to user

Bruno Michel, bmi@zurich.ibm.com

Summary

- **Thermal resistance key to improve efficiency and reduce carbon footprint**
  - Zero Emission Datacenters – from Idea to largest European Computer in 5.5 Years
  - Application beyond datacenters in all thermally mediated energy conversion processes
  - Thermal and electrical packaging can take lead in many more areas

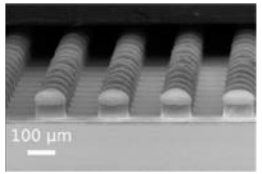

- **Packaging Research**
  - Thermal interfaces, hierarchical microchannel cooling, percolating underfill, all copper interconnect
  - 100 PY investment in reducing convective and conductive thermal resistances
  - Interlayer cooled chip stack with power deliver and optical communication

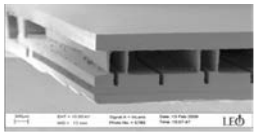

- **Volumetric density scaling to replace Moore's law and transform IT industry**
  - Big Data and Cognitive Computing Drive demand for more efficient IT hardware
  - Brain inspired packaging with combined power supply and cooling ➔ Bionic packaging
  - Top down miniaturization: Datacenter in a box as first step

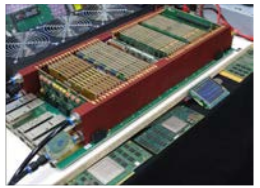

- **Wearables for health and Human Centric Sensing and Computing**
  - **Bell's law:** Mainframe computing – PCs – mobile computing – wearable computing
  - Cognitive companion and the augmented human

Bruno Michel, bmi@zurich.ibm.com

Thank you very much for your kind Attention!

Acknowledgement

- Ingmar Meijer, Patrick Ruch, Thomas Brunschwiler, Stephan Paredes, Werner Escher, Yassir Madhour, Jeff Ong, Gerd Schlottig and many more for Aquasar and SuperMUC design
- Sebastian Gerke, Rahel Straessle, Arvind Sridhar, Ismael Faro, Yuksel Temiz, Neil Ebejer, Theodore G van Kessel, Keiji Matsumoto, Emanuel Loertscher, Frank Libsch, Hyung-Min Lee, John Knickerbocker, Jonas Weiss, Jonathan E Proesel, Kang-Wook Lee, Mehmet Soyuer, Minhua Lu, Mounir Meghelli, Norma E Sosa cortes, Paul S Andry, Roy Yu, Shriya Kumar, Sufi Zafar, Teodor K Todorov, Yves Martin …. and many more
- PSI Tobias Rupp and Thomas Schmidt
- ETH Severin Zimmermann, Adrian Renfer, Manish Tiwari, Ozgur Ozsun, Dimos Poulikakos
- EPFL Yassir Madhour, John Thome, and Yussuf Leblebici
- Funding: IBM FOAK Program, IBM Research, CCEM Aquasar project, Nanotera CMOSAIC project, SNF Sinergia project REPCOOL, NRP 70 project THRIVE, EUFP7 project CarrICool, EUFP7 project HyperConnect, DARPA IceCool, KTI / DSolar. IBM Research Frontiers Institute: www.research.ibm.com/frontiers

Bruno Michel, bmi@zurich.ibm.com

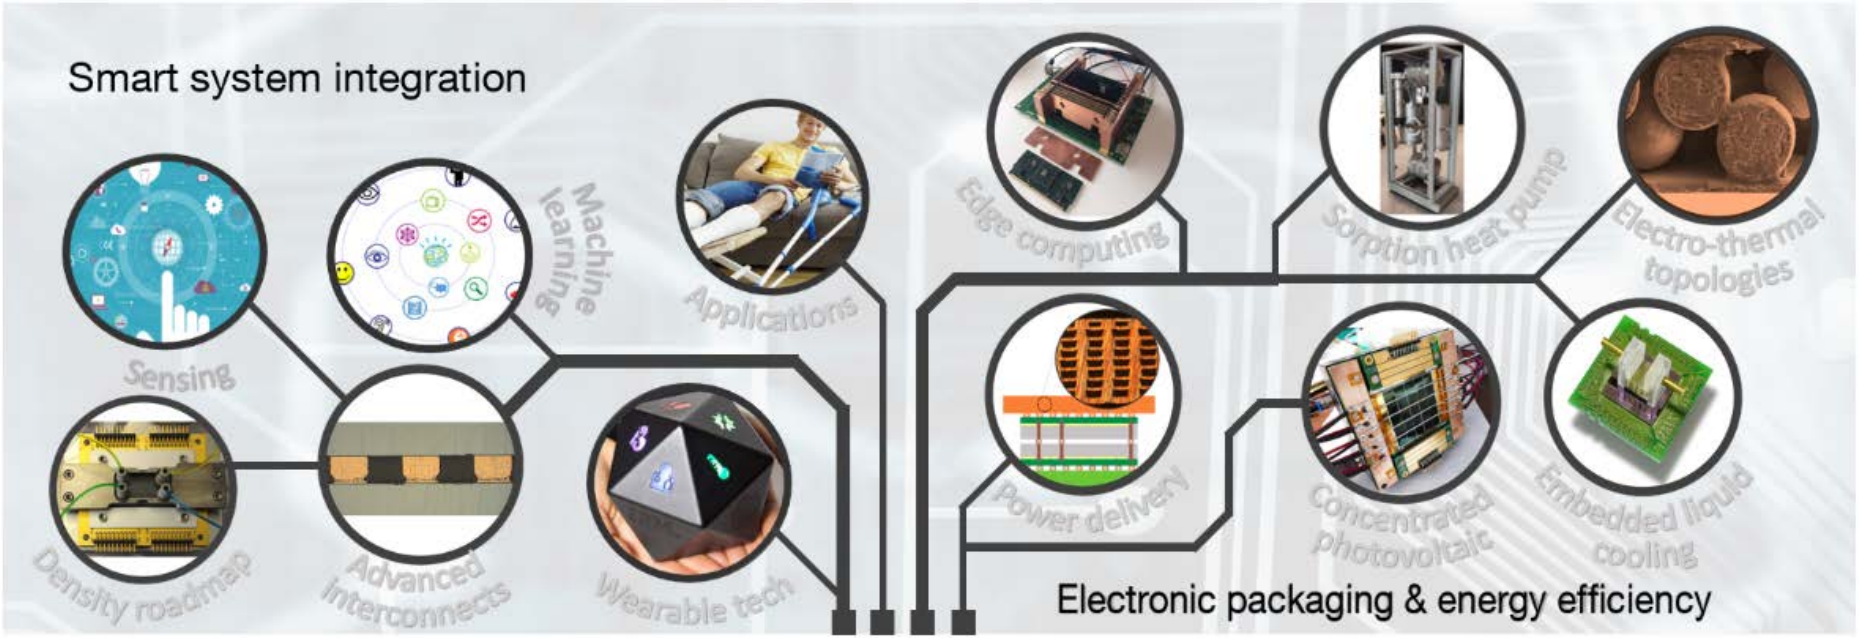

- **Links:**
- **Smart System Integration** https://www.zurich.ibm.com/st/smartsystem/
- **Functional Electronic Packaging** https://www.zurich.ibm.com/st/electronicpackaging/
- **Efficiency and Green Technologies** https://www.zurich.ibm.com/st/energy_efficiency/

Bruno Michel, bmi@zurich.ibm.com

Smarter Energy: Impact Outside of ICT Industry

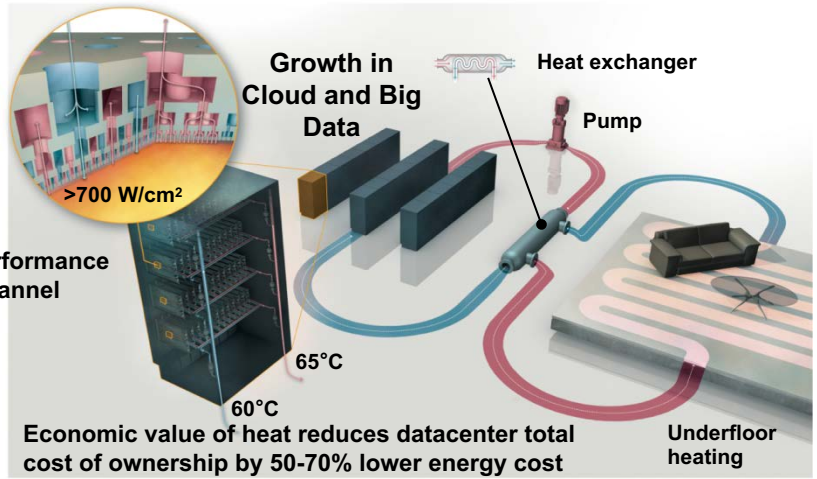

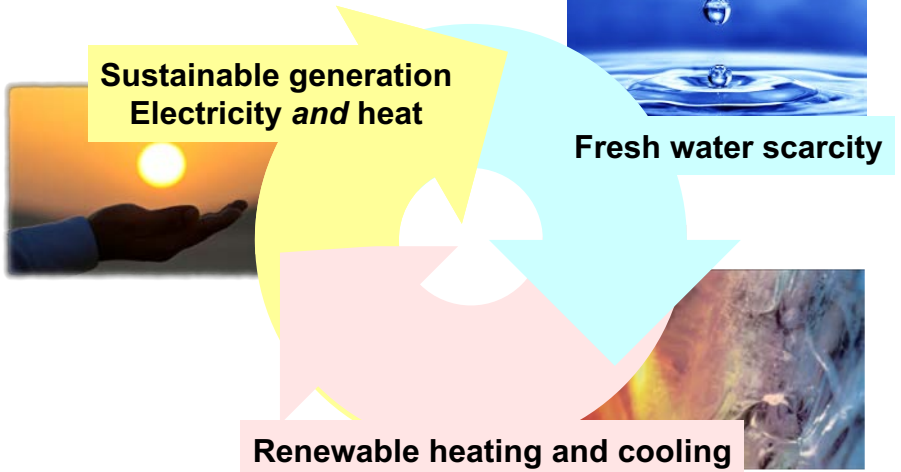

Zero-emission datacenter

High-concentration PV/thermal

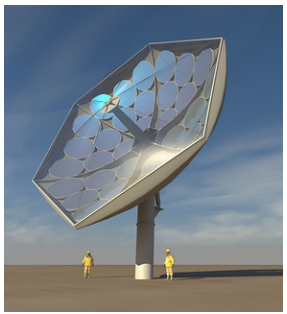

Adsorption heat pump

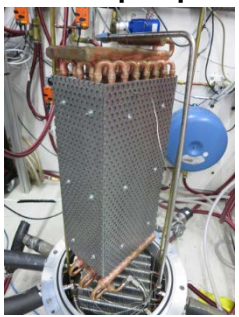

Membrane distillation desalination

Electrochemical redox energy conversion