Advanced power and thermal packaging enabling dense integration of compute nodes

Dr. Arvind Sridhar
Smart System Integration
IBM Research – Zurich, Switzerland

https://www.zurich.ibm.com/st/
https://www.zurich.ibm.com/st/smartsystem/

Bell's law in our times

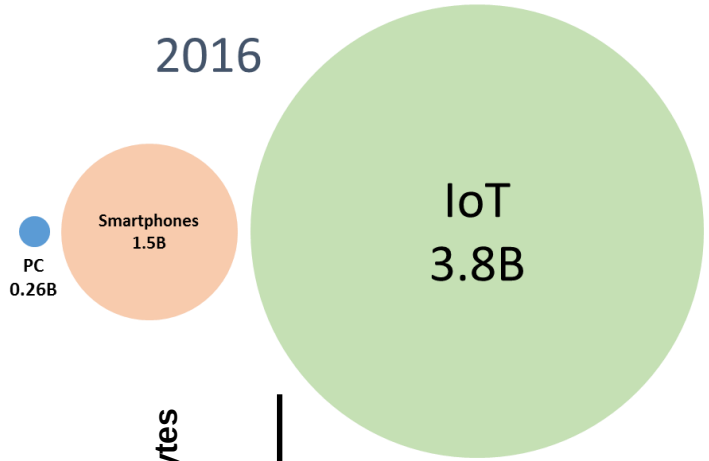

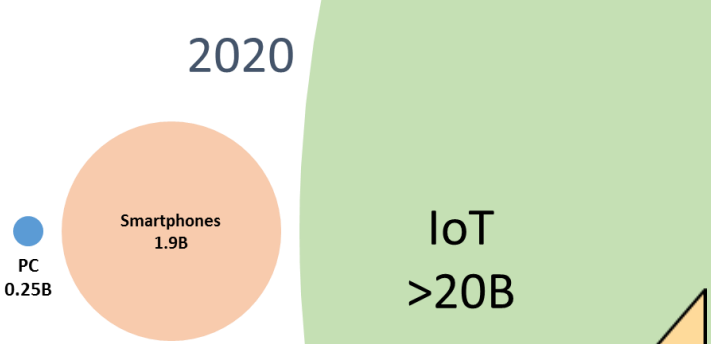

"Roughly every decade a new, lower priced computer class forms based on a new programming platform, network, and interface resulting in new usage and the establishment of a new industry."

- Gordon Bell 1972

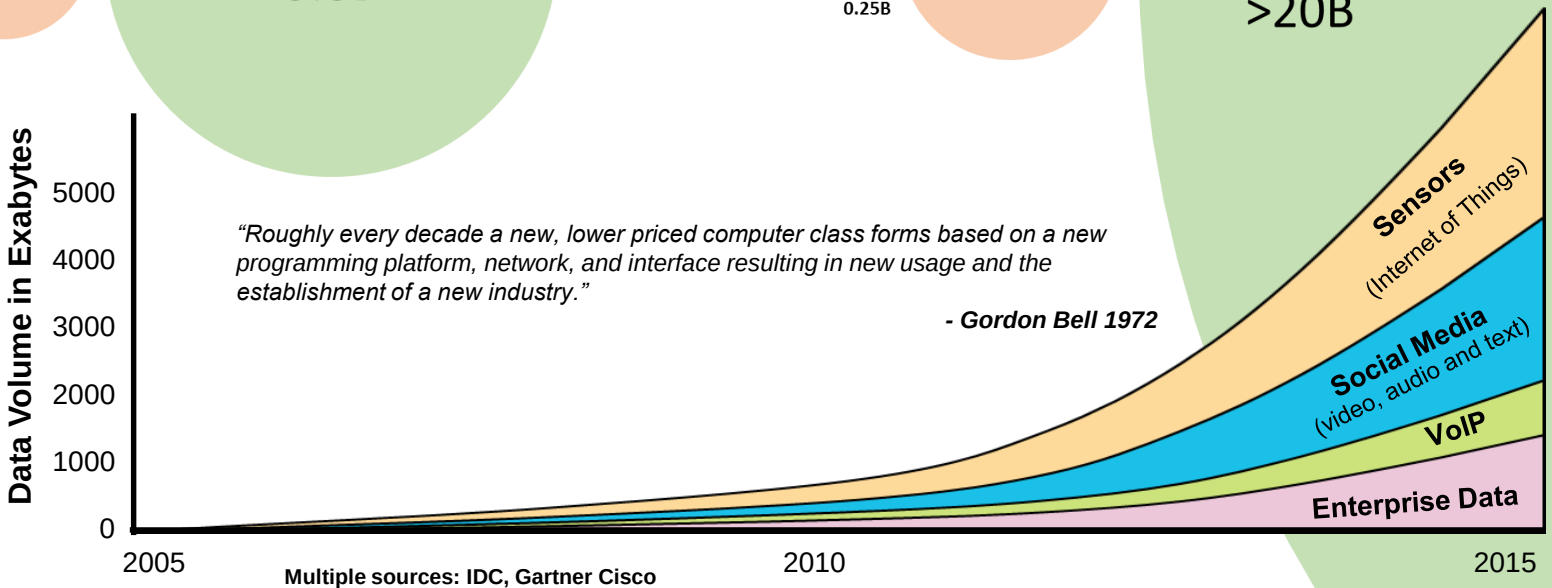

Multiple sources: IDC, Gartner Cisco

Is our computing infrastructure up to the task?

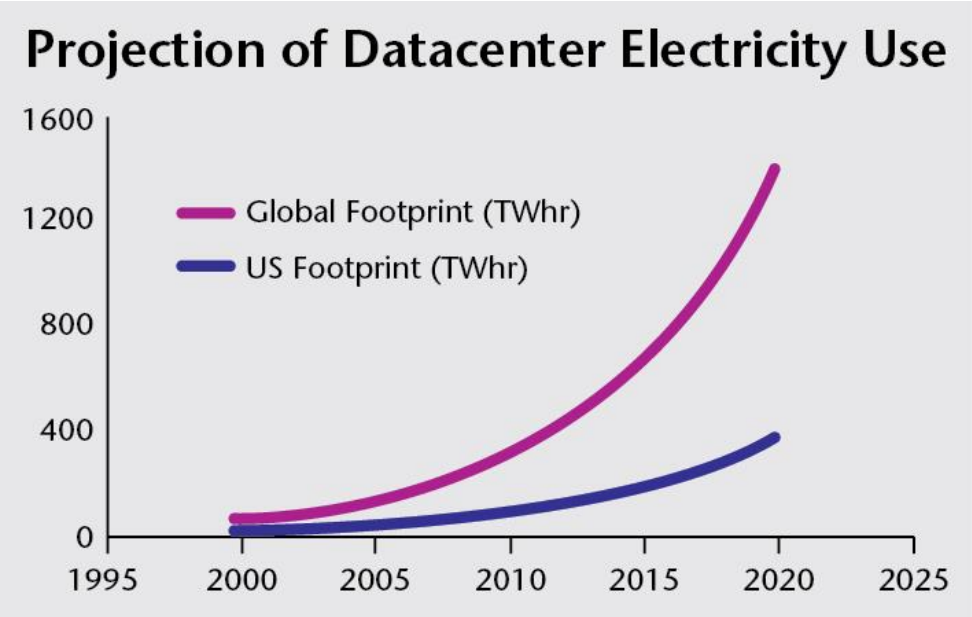

3% of global electricity, 2% greenhouse emissions

Renewables?

→ no where close to running 100%

Efficiency?

→ First energy source

→ 1% improvement in global energy efficiency = 23 nuclear power plants/entire wind power output

Source: Independent, McKinsey Report, New Climate Economy Report 2015

Where do we improve efficiency?

Via CMOS scaling?

Since 2006/07: Circuits still become smaller and cheaper, BUT not faster and not more efficient.
Since 2015: Circuits still become smaller but not cheaper ➔ **Moore's economic "law" is dead**

Transistors for computing occupy much less than 1 ppm of the system volume.
Rest for power supply, cooling, and communication

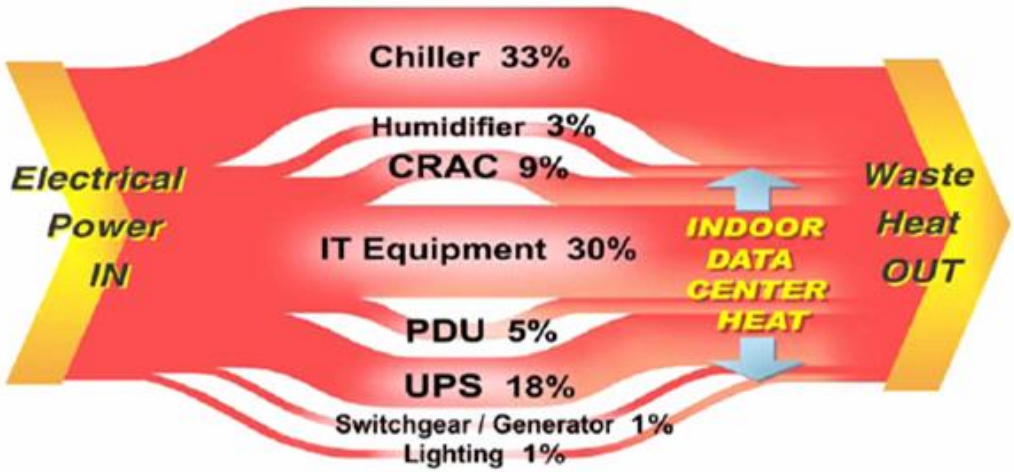

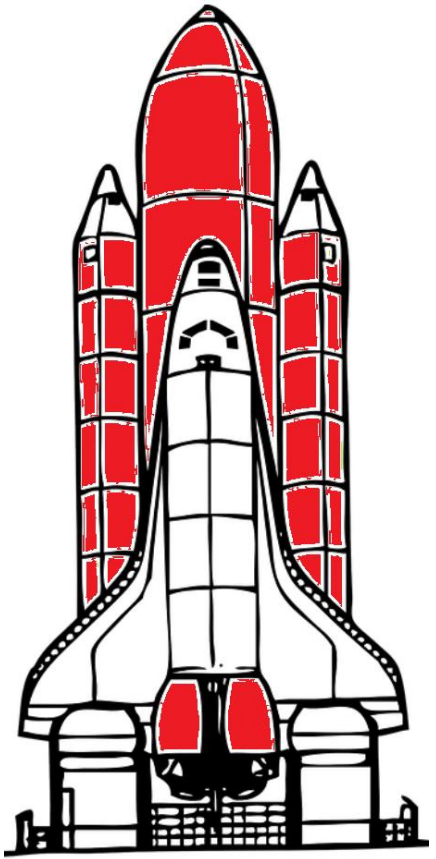

 Brouillard, APC, 2006

(R)evolution of Information Technology

Device centric view

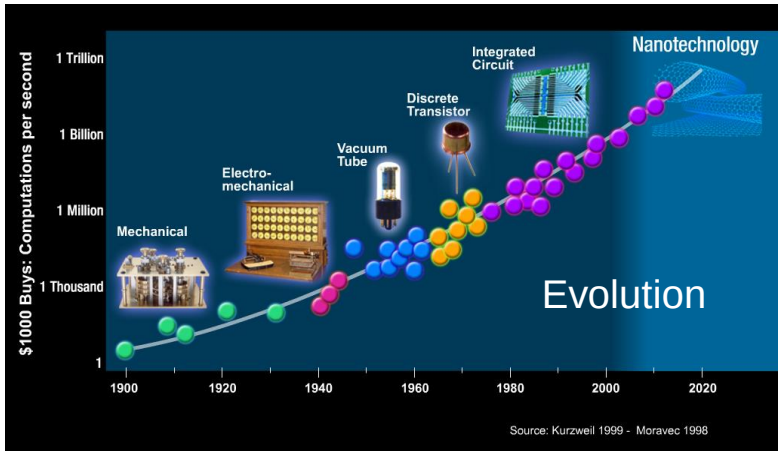

Density centric view

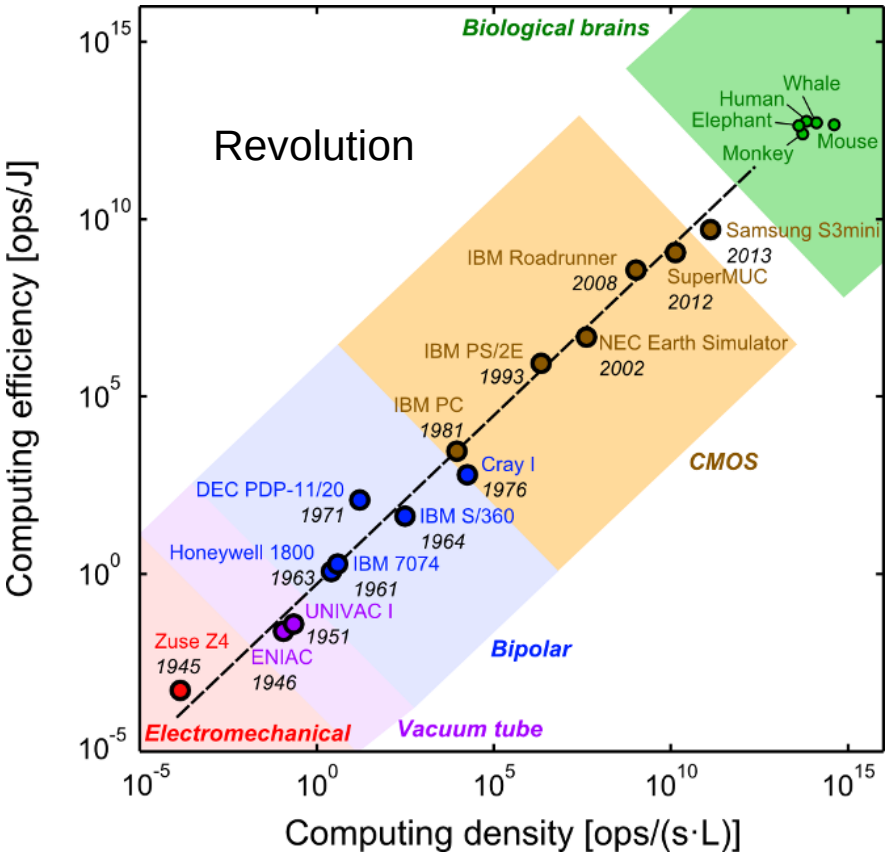

Allometric Scaling in Biology

- Basal Metabolic rate of animals scales sub-linearly with mass

$\rightarrow R \propto M^{3/4}$

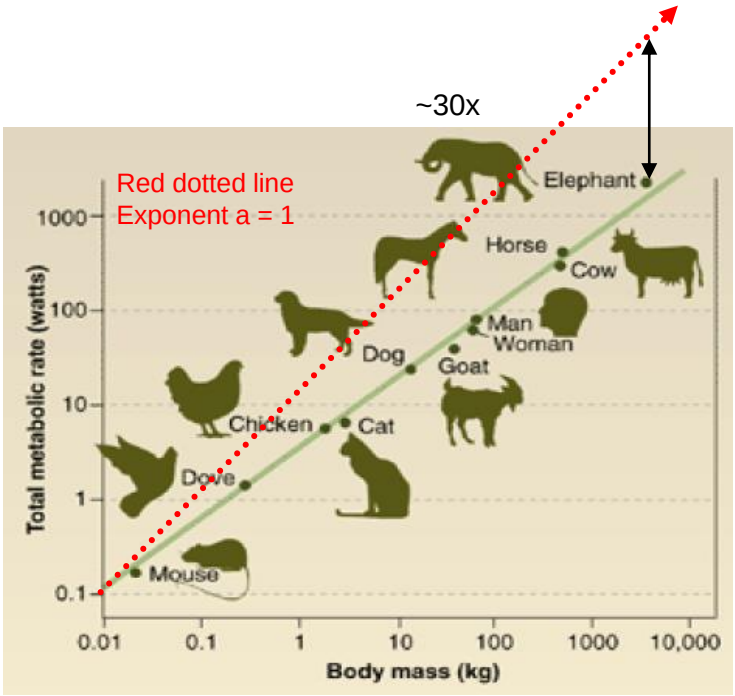

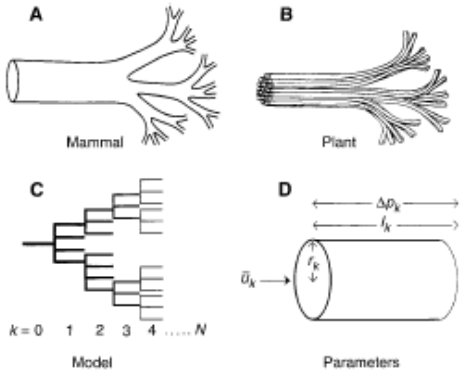

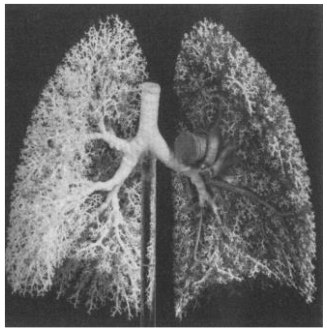

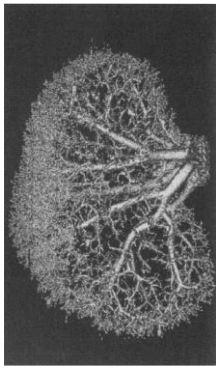

Biology achieves this by:

1. efficiency “densification” of supply networks by branching
2. combining power delivery and heat removal

G. West et al, Science, 1999

Paradigm Change: Vertical Integration

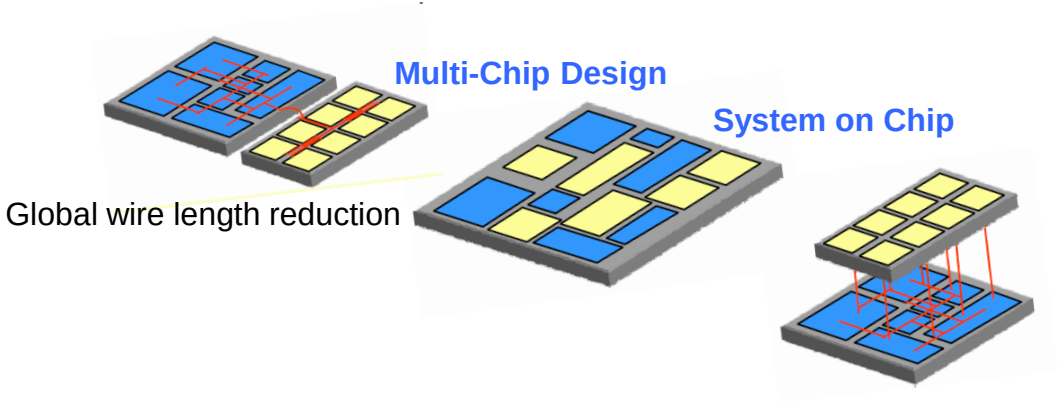

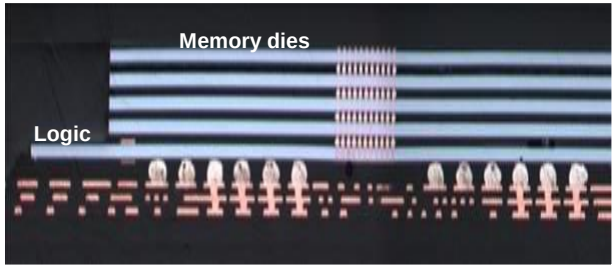

Source: Micron-IBM HMC development

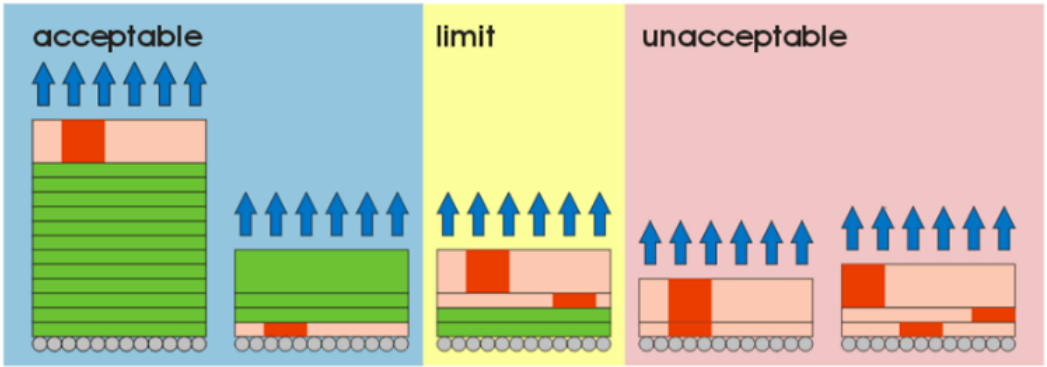

Limitations to further scaling:

1. **Heat removal**
2. **Power delivery**

Heat Removal

Liquid cooling in Datacenters

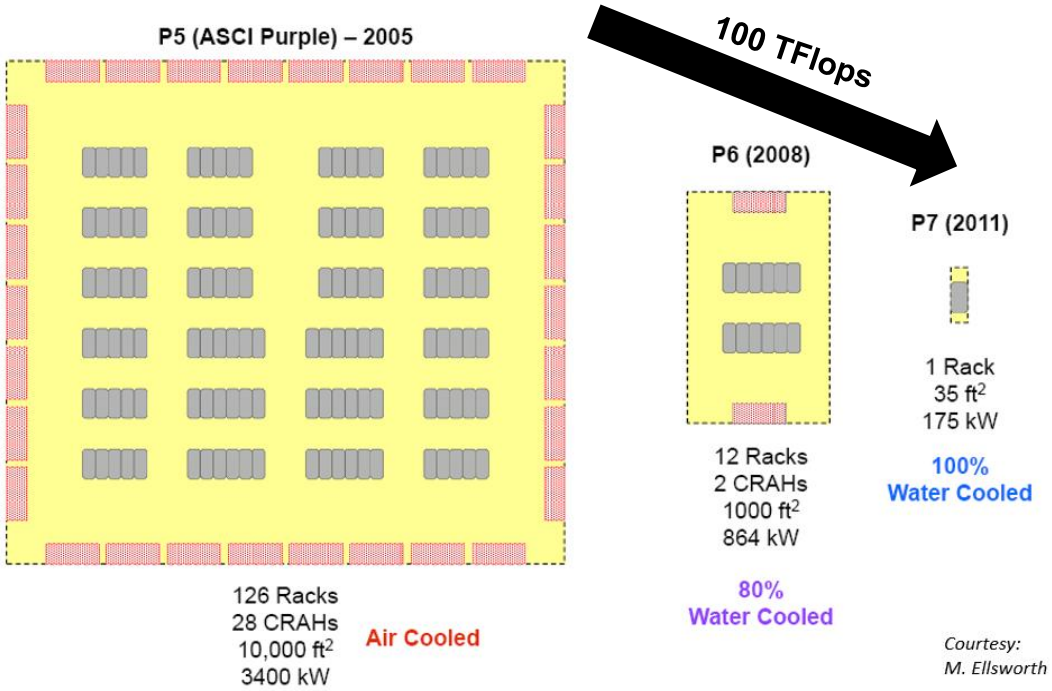

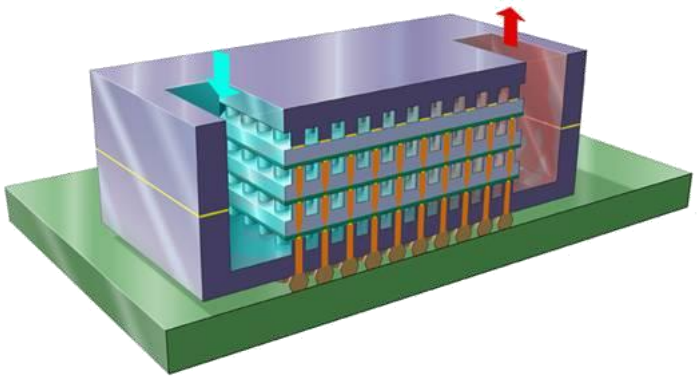

286x footprint reduction

→ 16x by node scaling

→ 18x by packaging technology

Interlayer Liquid Cooling (2005-2015)

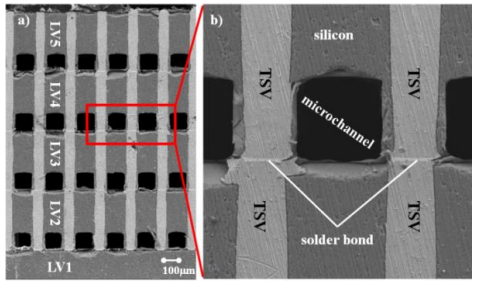

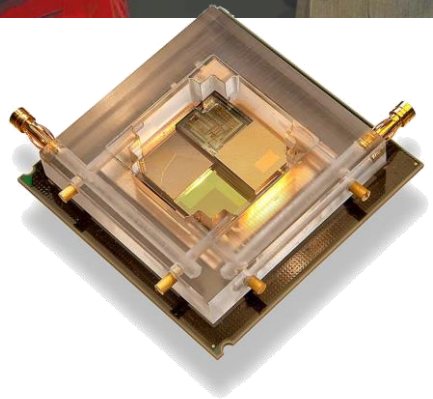

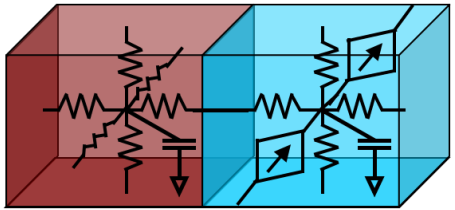

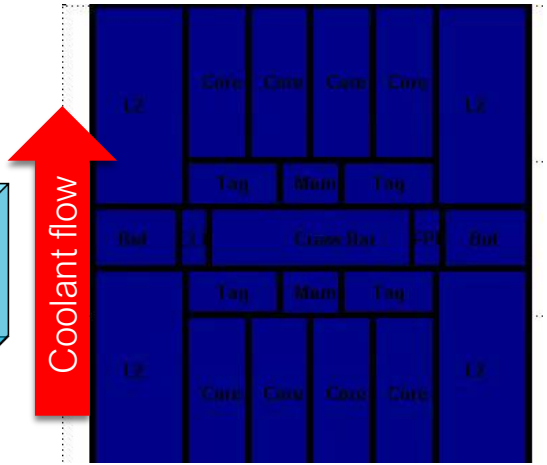

Water cooling today, and tomorrow…

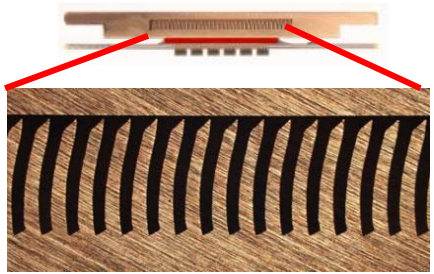

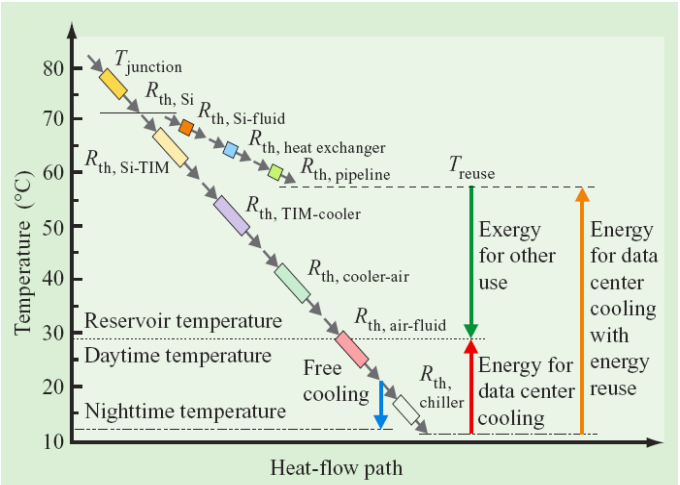

- Save chiller energy: Cool with **T > 60ºC hot water → Free cooling**
- Cold and moderate climates: **energy re-use**
- Hot climates: **desalination**

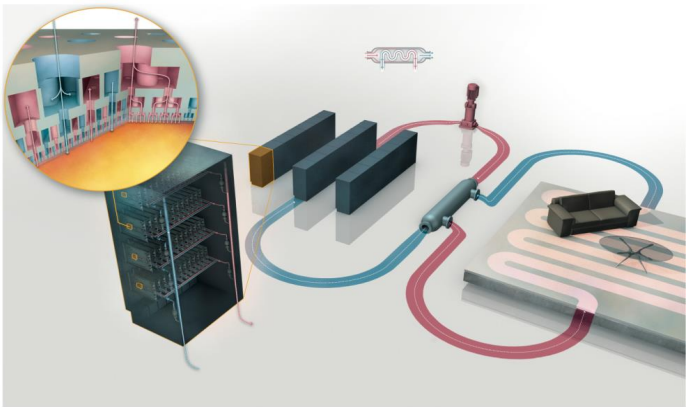

Power Delivery

The Power “bottleneck”: need for VR granularity

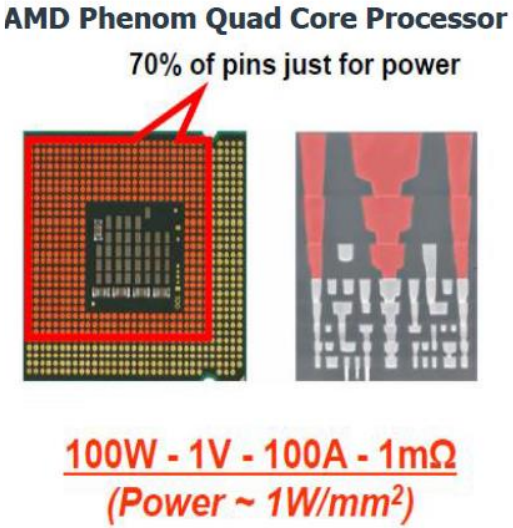

Source: C. Mathuna, APEC 2014

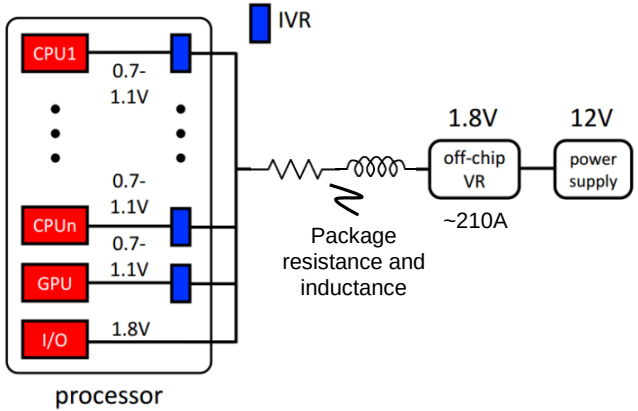

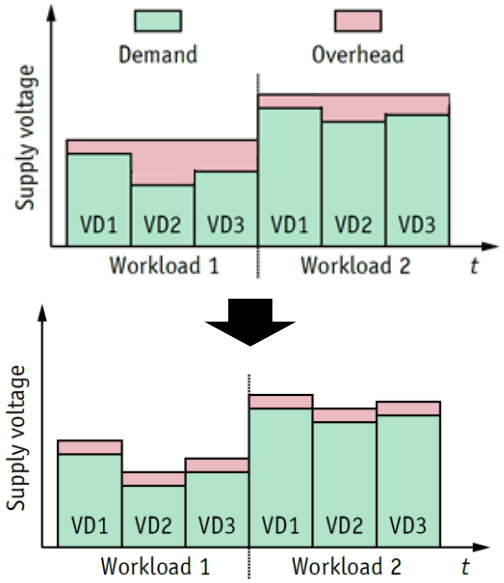

Goals:

- Bring the point of load (PoL) VRM as close to the processor as possible
- Perform the largest DC-DC down conversion possible
- Have the highest spatial granularity possible
- Follow the fast changes in CPU activity

Different approaches to iVRs: separated, fully and hybrid

Separated VR

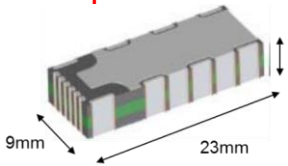

SMChiP VRM Current Multiplier

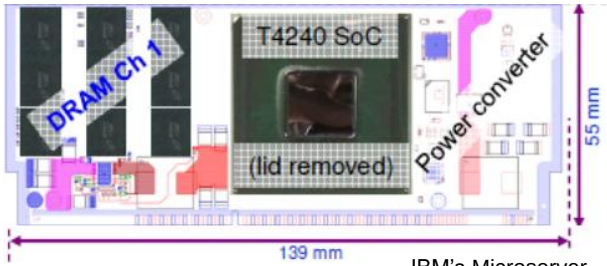

IBM's Microserver
[Luijten et al. ISSCC 2015]

Fully Integrated VR

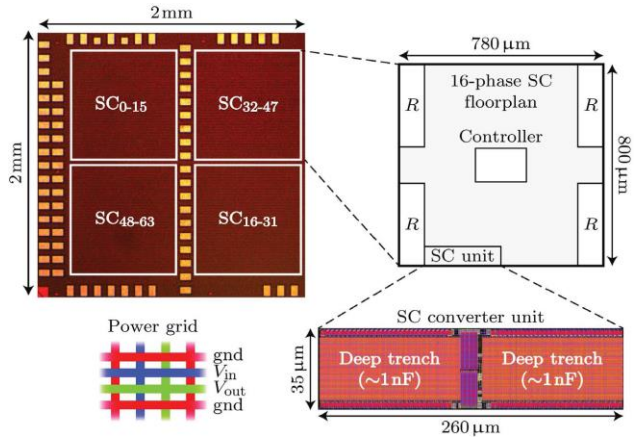

IBM/ETHZ's Switched Capacitor
IBM 32nm SoI
[Andersen et al. ISSCC 2015]

Hybrid Integrated VR

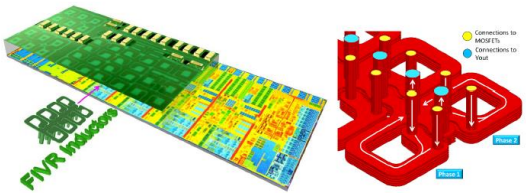

Intel's FIVR buck converter
CMOS 22nm air-core L on laminate
[Burton et al. APEC 2014]

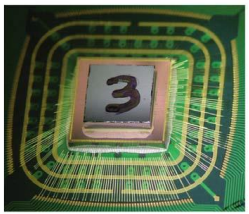

IBM/Columbia U's iVR25D buck converter
CMOS 45nm magnetic-core L on Interposer
[Strucken et al. JSSC 2013]

Intel FIVR
CMOS 14nm discrete L post processed [ISSCC 2017]

CarrICool Project (EU-FP7, 2014-2017)

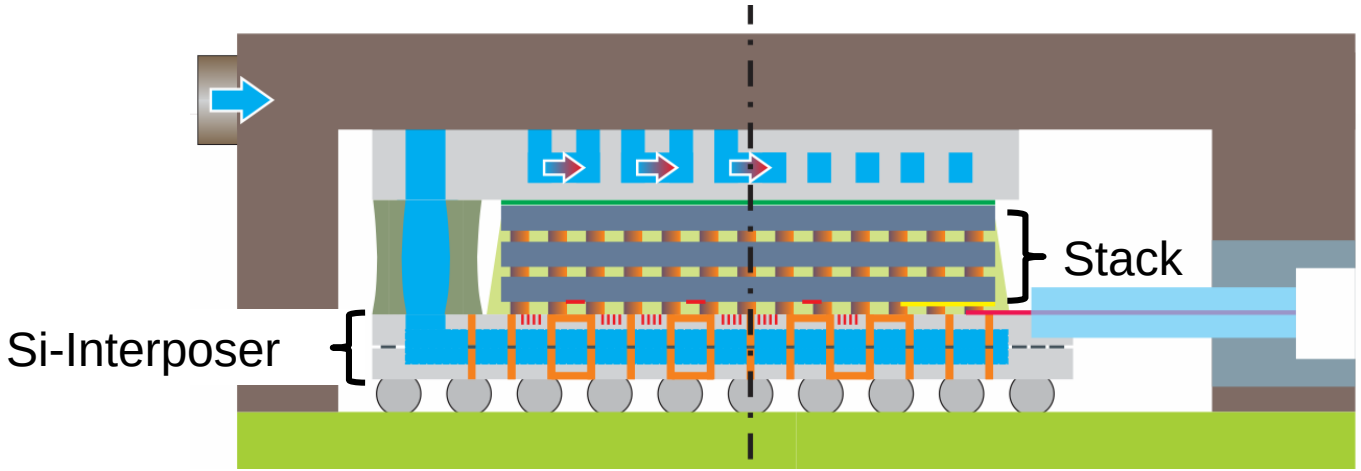

A silicon interposer that:

- provides liquid cooling
- delivers power using a VRM
- provides a high-speed optical link

CarrICool Hybrid iVR Implementation

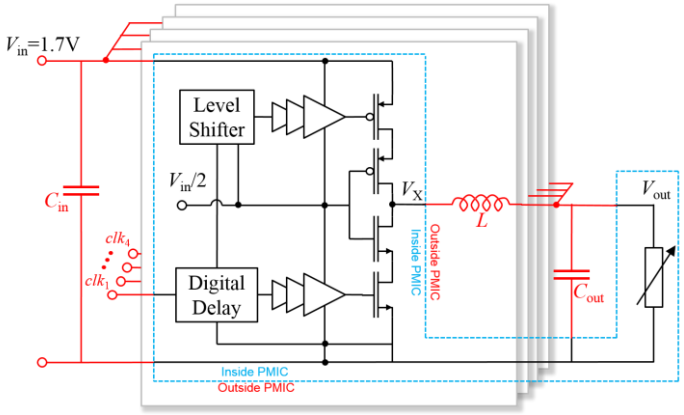

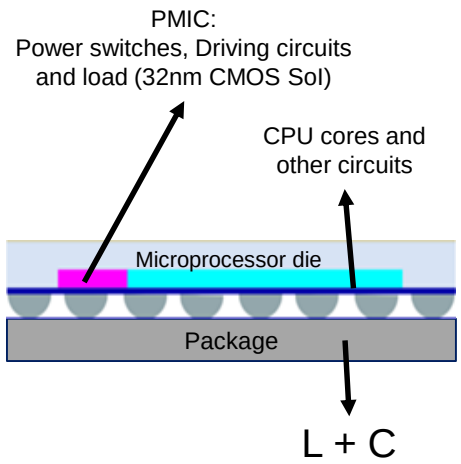

- Provide accurate voltage regulation for conversion ratios < 1
- Very simple/robust topology
- IVR challenges:
- Switches and control circuits: perform better if low loss. Can make use of the latest CMOS technology
- Inductors (and capacitors): require large structures, can't be compact enough for silicon. Don't need cutting-edge process.

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

CarrICool iVR technologies

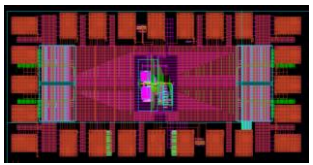

PMIC
32nm IBM SoI process/
14nm GF Bulk

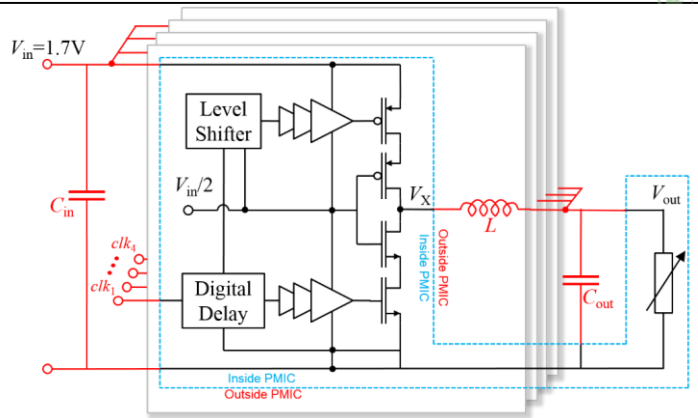

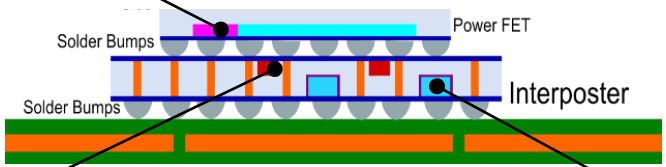

Deep Trench Capacitors
IPDIA MIMIM (>200nF/mm^2)
[Bunel and Lengignon, Device Packaging, HiTEC, HiTEN, & CICMT 2013]

Air-core Inductors
Fabricated using the RDL layers on the interposer

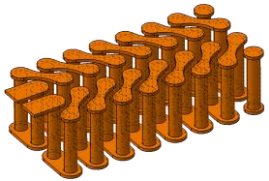

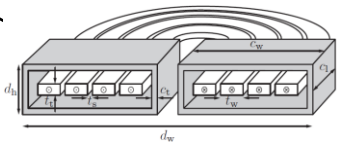

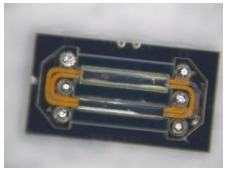

Inductors
Planar or Volumetric Inductors using TSVs with Magnetic core material added using thin film deposition

Final assembly

Multiple generation of PMICs

PMIC 1 – IBM 32nm SoI 2015
PMIC 2 – GF 14nm LPP 2016
PMIC 3 – GF 14nm LPP 2016
PMIC 4 – GF 14nm LPP 2016

Open-Loop regulation

PMIC 5 – GF 14nm LPP 2016

Analog control

PMIC 6 – GF 14nm LPP 2017
PMIC 7 – GF 14nm LPP 2017
PMIC 8 – GF 14nm LPP 2017

Digital Control

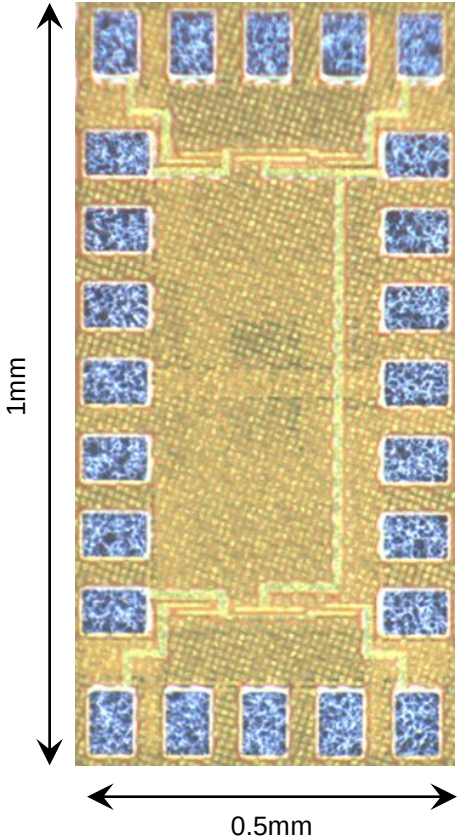

[COBEP SPEC 2015, ECPE 2015, PwrSoC 2014/2016, COMPEL 2017]

CarrICool iVR demonstrators

Multiple demonstrators built on a single interposer tile

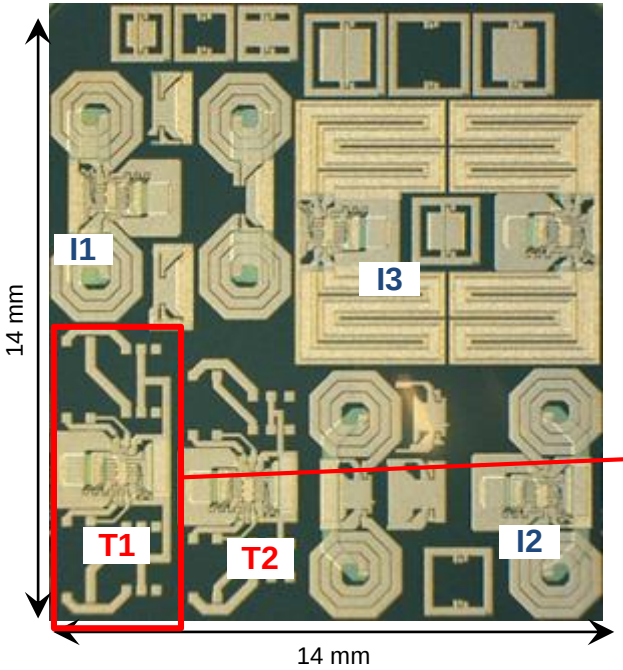

Au stubbing + Thermocompression

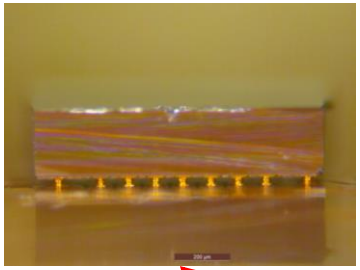

Low T solder balling + reflow

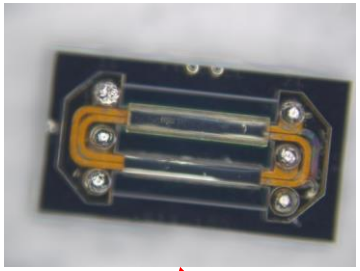

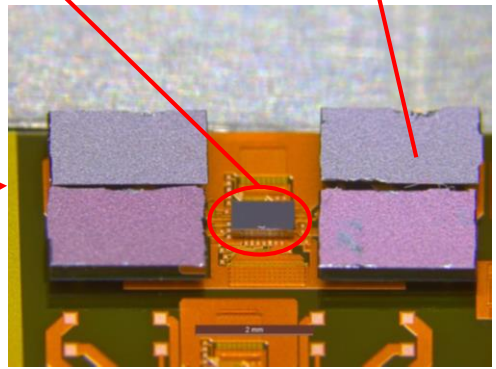

Fraunhofer IZM

Detailed Characterization pending

Outlook: Exploration of Passive Platforms

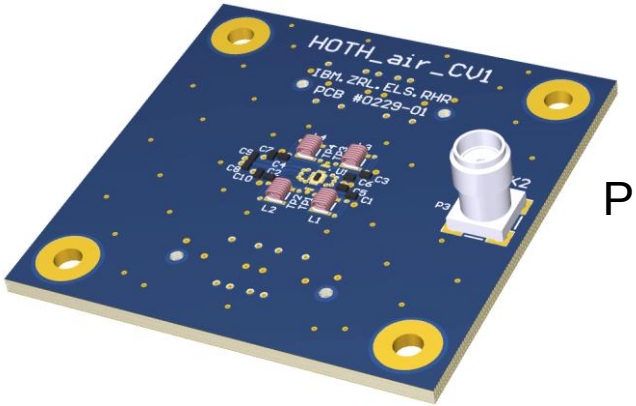

PCB/ organic substrate

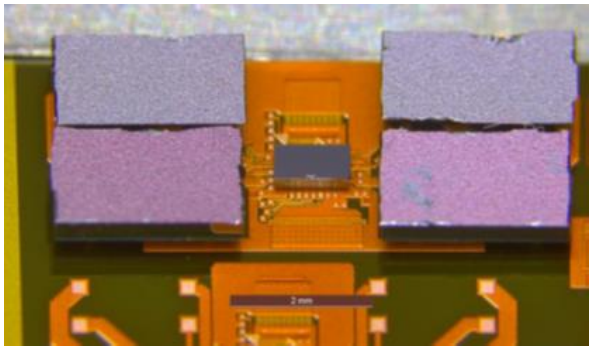

Silicon Interposer

CarrICool

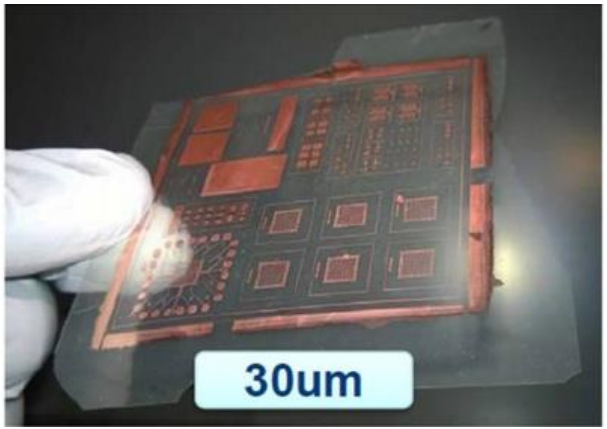

Glass Interposer

GaNonCMOS Project (EU-H2020, 2017-2020)

GaNonCMOS

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

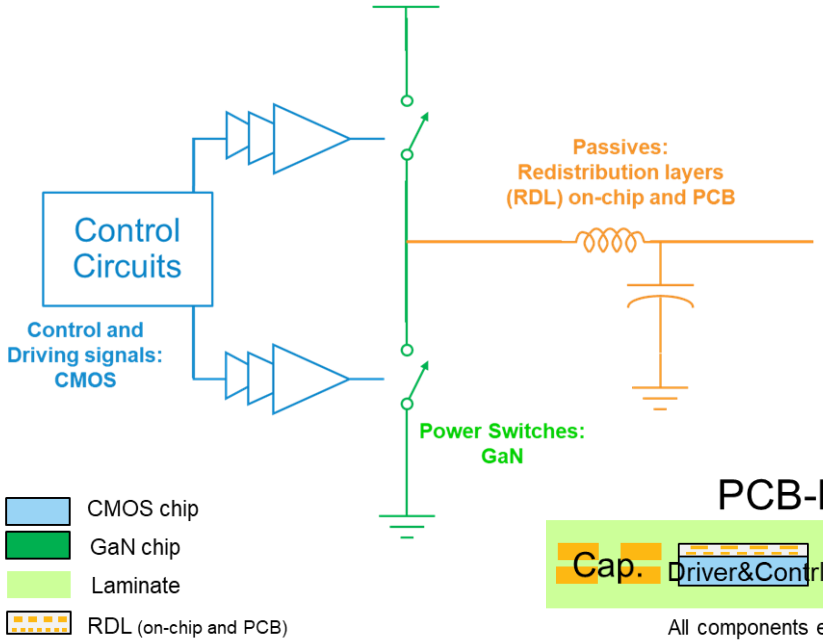

More details in the Poster!

Chip Level

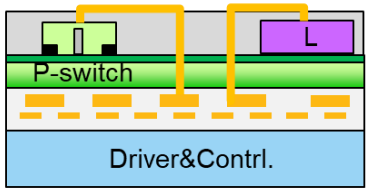

GaN Power Switches and Thin-film core Inductors attached to the CMOS chip using DWB
High-risk, High-Frequency Integration

Stack-level

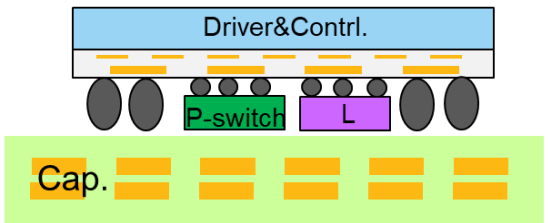

Components bonded together using C4 solder
Mid-level risk, Mid-Frequency Integration

PCB-Level

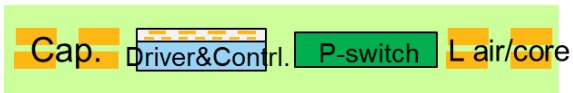

All components embedded in PCB
Low-risk, Low-Frequency Integration

Evolution of GaN on CMOS Integration

Combining Power and Heat

IBM Microserver

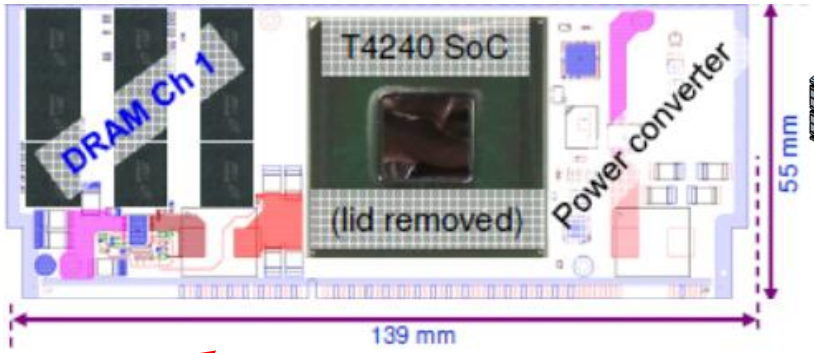

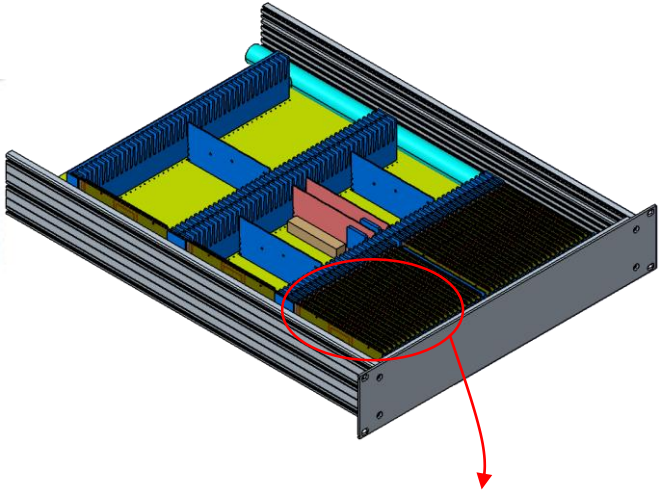

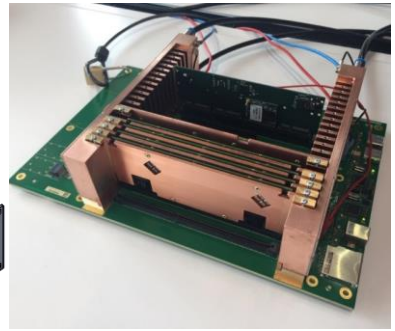

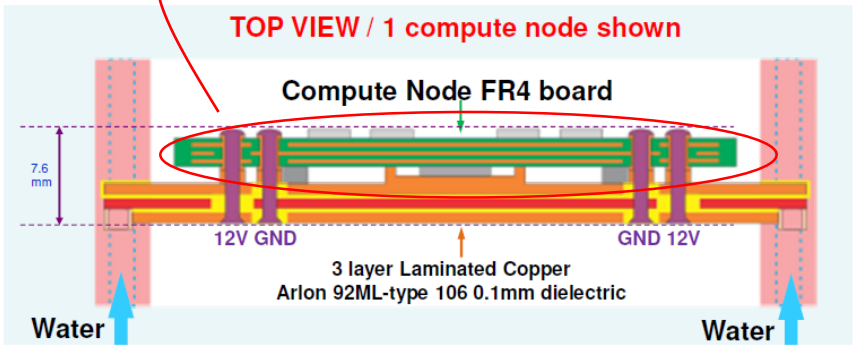

[ISSCC 2014]

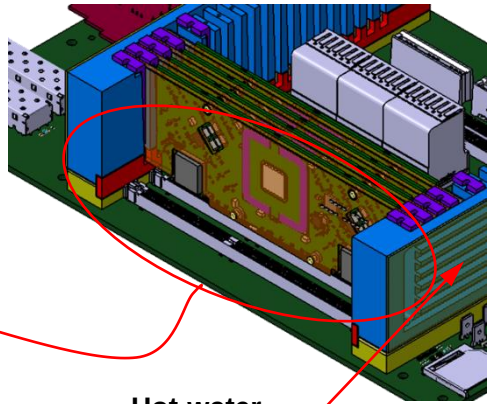

Hot-water cooled heat sink

- Standard 2U Rack
- Low Power/Cost Compute nodes (T4240 1.8GHz, 12 cores)
- 128 nodes – 6TB DRAM
- 28 TFLOP
- 1000x Denser
- 10x more efficient than standard HPC Rack

“Electronic Blood” Electrochemical Power Supply on-chip

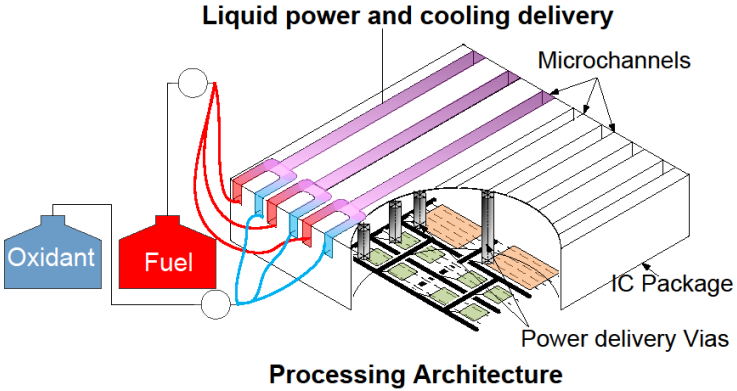

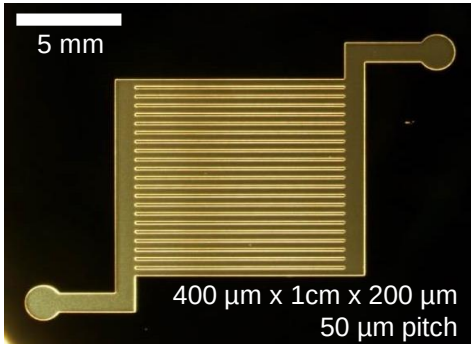

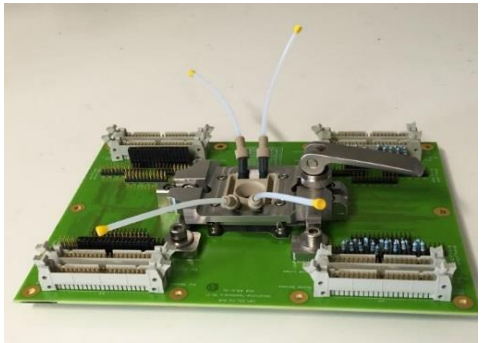

Characteristics

- Chip-level “battery” flow cell array
- Water Soluble redox species
- Co-laminar flow, no membraneless
- Single macroscopic charging unit
- Multiple chip-level discharge units
- Satisfies congruent demand for power delivery and heat removal

Test vehicle demonstration

- Process technology feasibility proven
- Challenge: power density order of magnitude lesser than needed
- PCB level scaled flow cell array can still power/cool the CPU

[IEEE Trans Computers 2018]
[International Flow Battery Forum, 2015 & 2017]

Smarter Energy: Impact Outside of ICT Industry

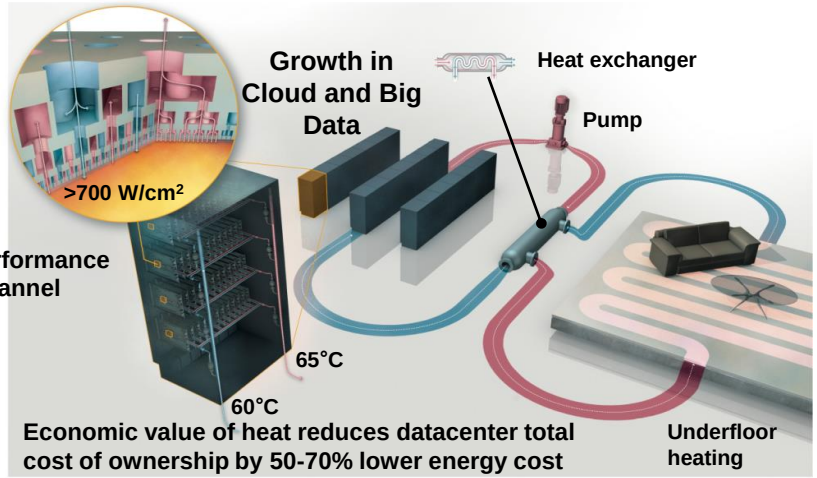

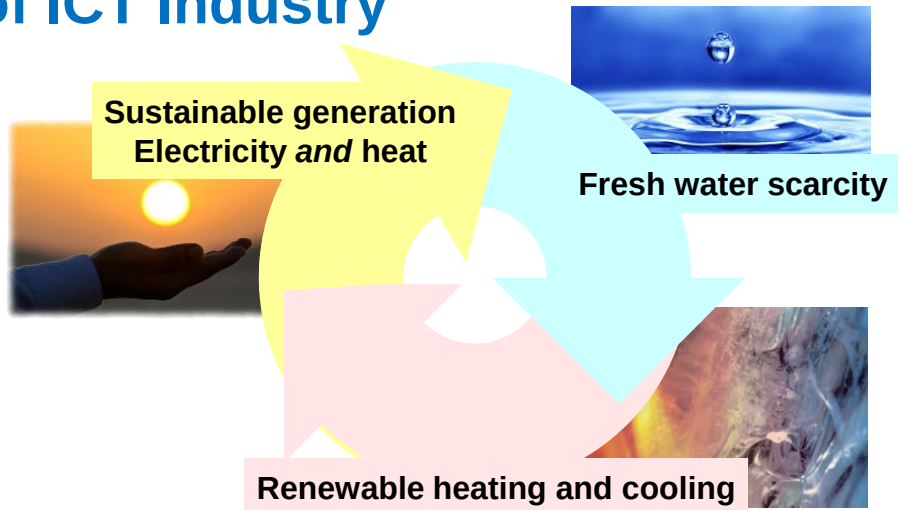

Zero-emission datacenter

High-concentration PV/thermal

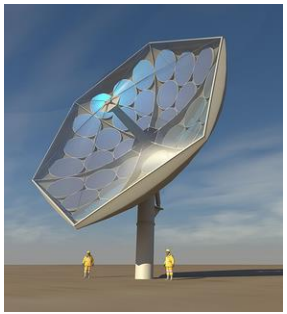

Adsorption heat pump

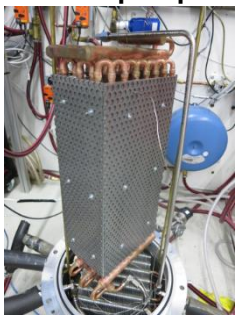

Membrane distillation desalination

Electrochemical redox energy conversion

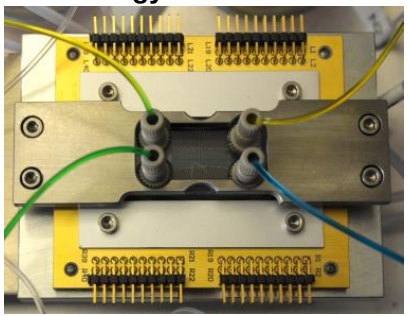

Thank you!

Organizations
- IBM Research Zurich – Smart Systems Integration Group
- IBM Research Zurich – IO Links group
- IBM Research Zurich – Heterogenous Cognitive Computing Systems group
- ETH – PES group, Zurich Switzerland
- ETH – LTNT group, Zurich Switzerland
- EPFL – ESL group, Lausanne Switzerland
- Tyndall National Institute, UCC, Cork Ireland
- IPDIA (Murata), Caen France
- Fraunhofer IZM, Dresden Germany
- Fraunhofer IAM, Freiburg Germany
- IHP, Frankfurt (Oder) Germany

Funding projects
- EU FP7 CarrICool
- EU H2020 GaNonCMOS
- SNSF RepCool
- SNSF CMOSAIC
- ASTRON Dome project