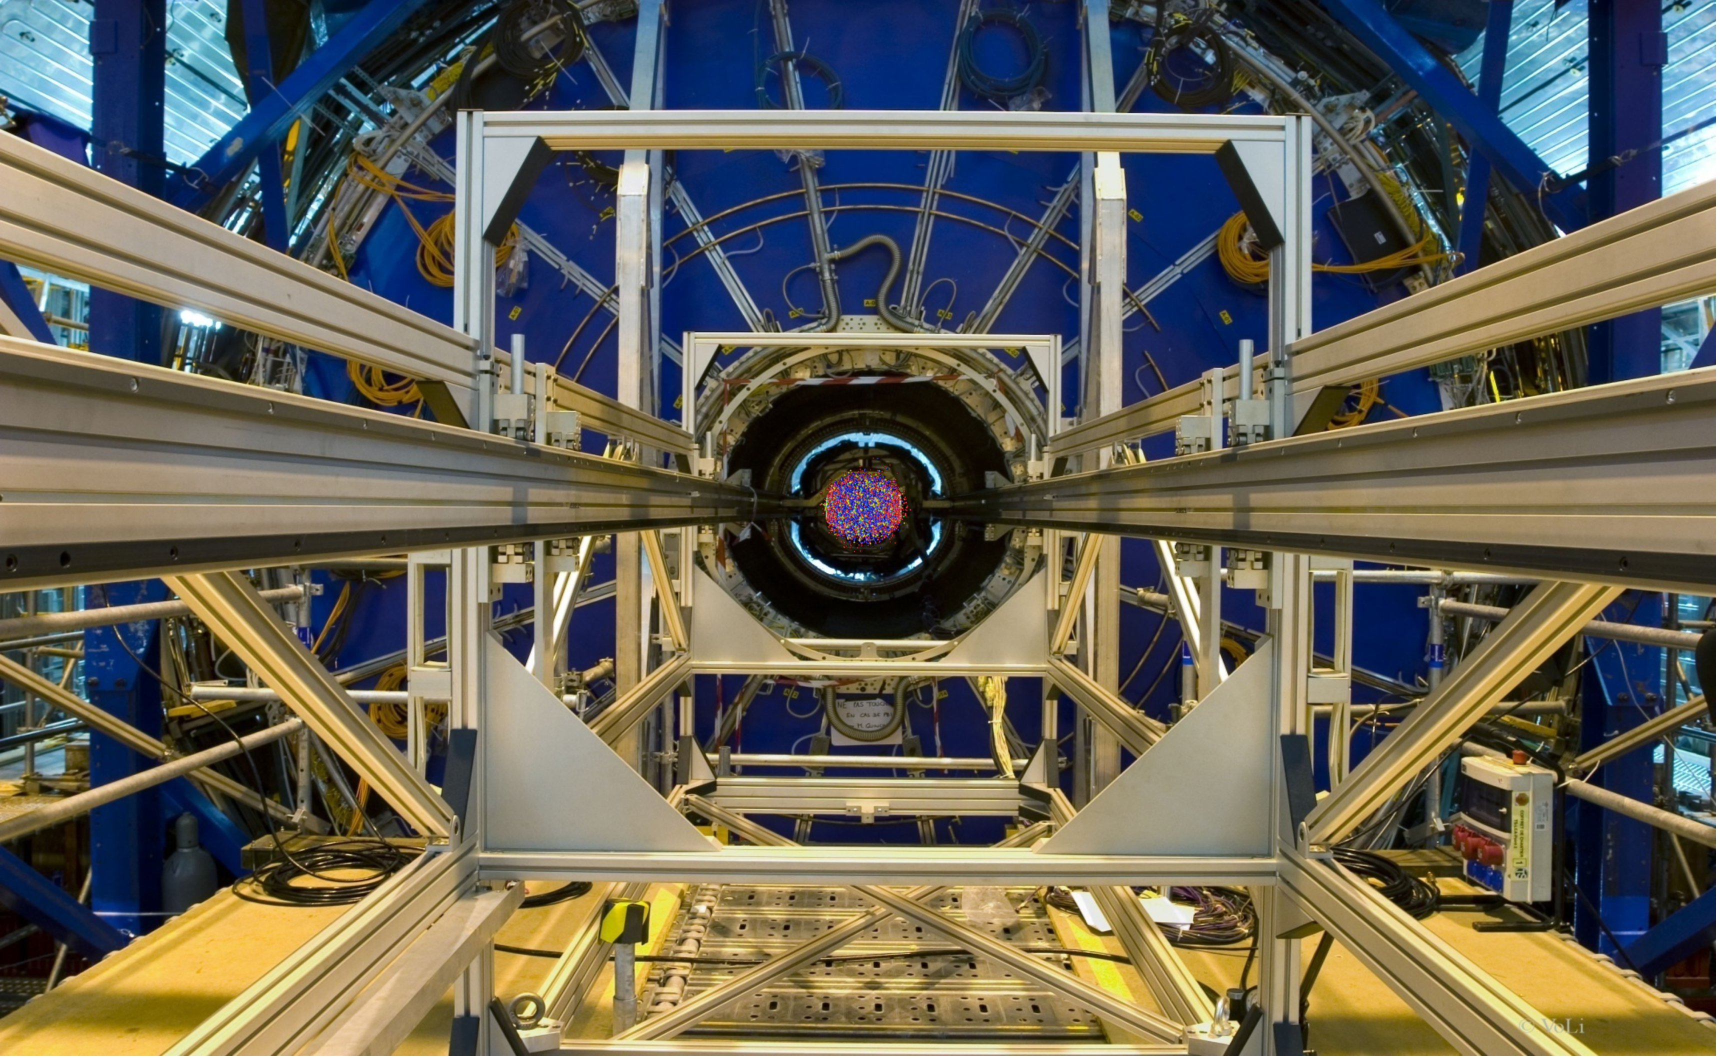

EPS Technology and Innovation Workshop

High Performance Computing

Volker Lindenstruth
GSI, FIAS, LOEWE Professor
University Frankfurt, Germany
GSI Helmholtzcenter
Phone: +49 69 798 44101
Fax: +49 69 798 44109
Email: voli@compeng.de
WWW: www.compeng.de


INTERNATIONAL JOURNAL OF HIGH-ENERGY PHYSICS
CERNCOURIER
VOLUME 52 NUMBER X MONTH 2012
Headline headlin headline hea
EYEBROW
Strapline strapline strapline strapline strapline strapline pXX
EYEBROW
Strapline strapline strapline strapline strapline strapline pXX
HEADLINE IN CAPS
Strapline strapline strapline strapline pXX

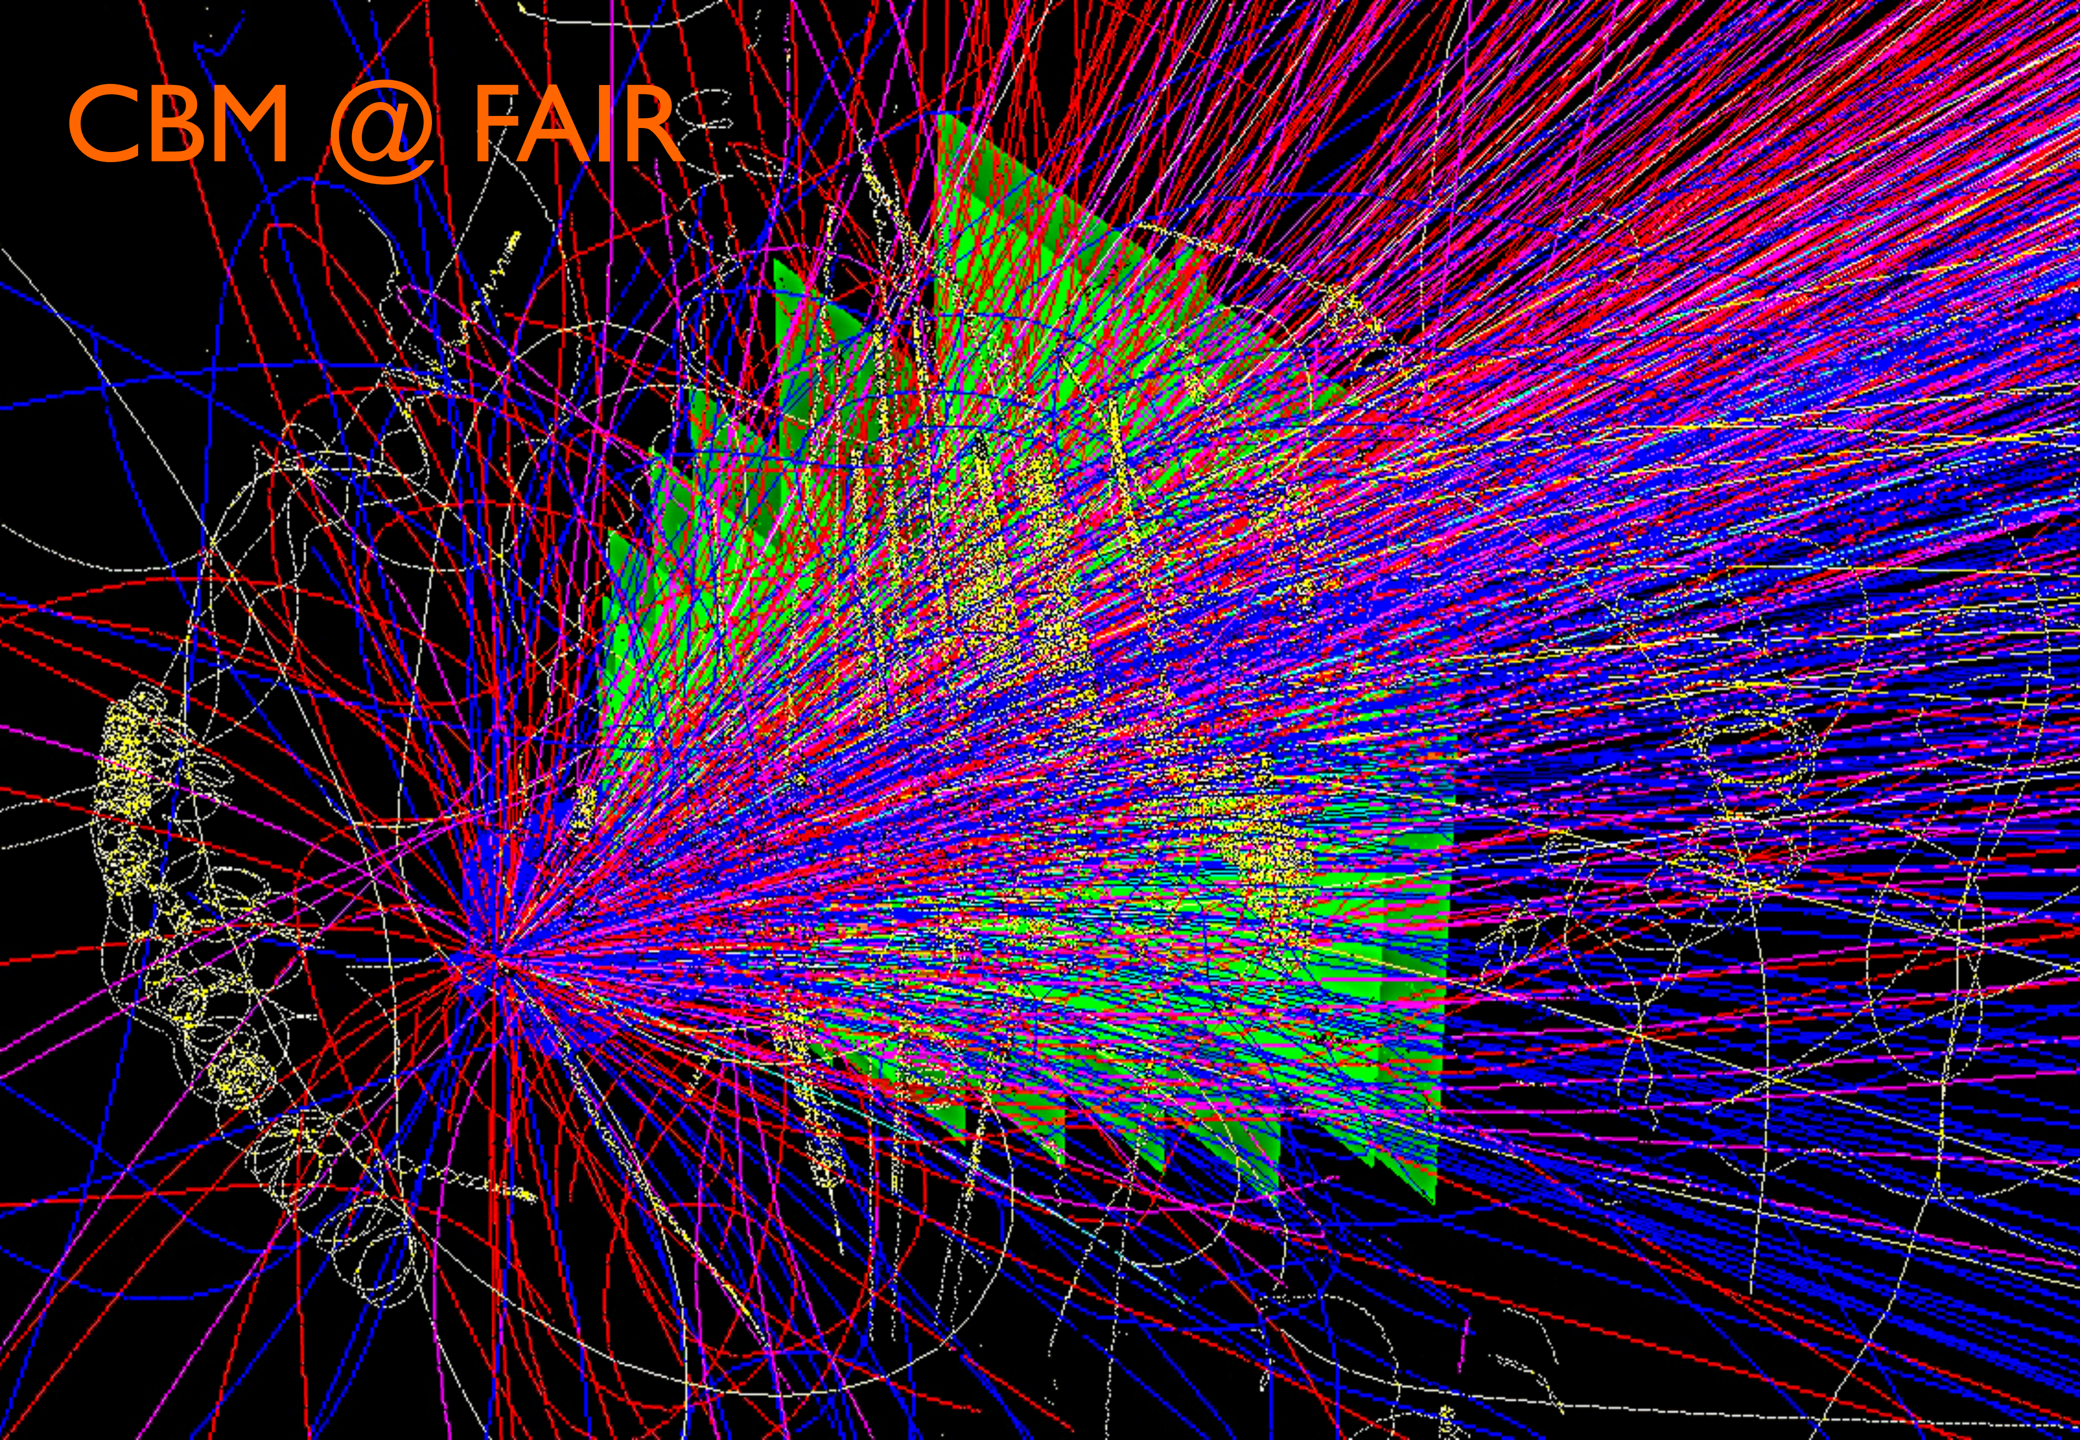
CBM @ FAIR

Data Center State of the Art

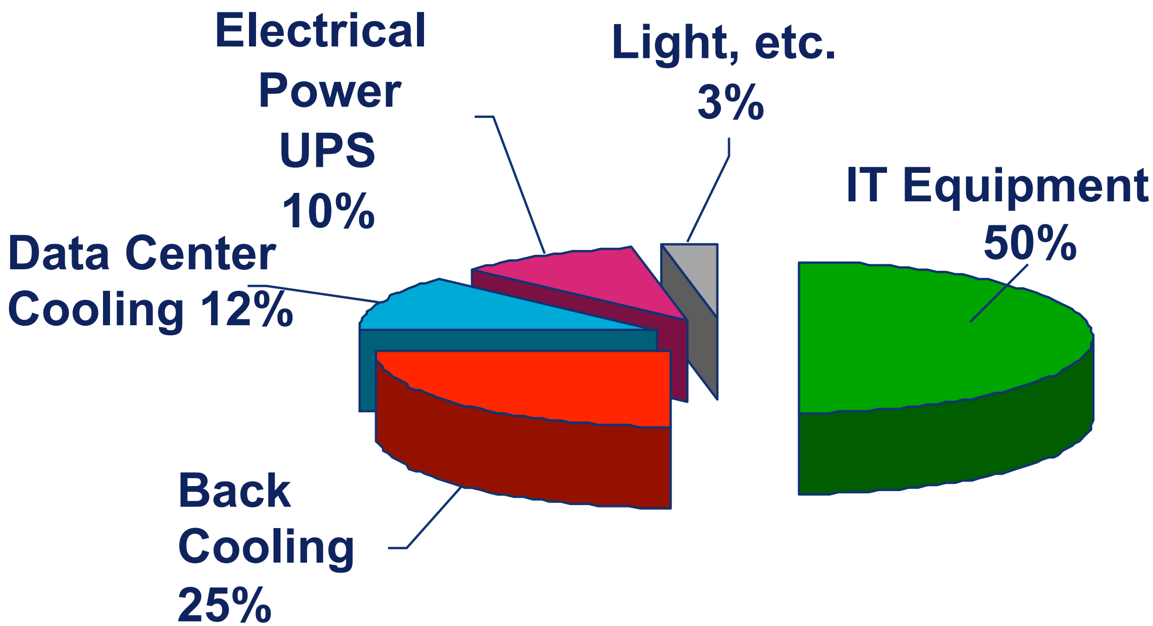

the green grid

$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

Power **U**sage **E**ffectiveness

PUE = 2

(typical PUE = 1,6 to 3,5)

Note: 12C/kWh → 1W = 1,05€/a

Source: EYP Mission Critical Facilities Inc., New York

Heat Transmission via Air and Water

Required Volumetric Current: $Q = \dot{V} = \dfrac{P}{c_\mathrm{p} \cdot \rho \cdot \Delta T}$

P : **Thermal Power Loss**
ΔT : **Temperature Difference**

Air Specific Heat Capacity: $c_\mathrm{p} = 1{,}005\,\frac{\mathrm{kJ}}{\mathrm{kg \cdot K}}$ Density: $\rho = 1{,}184\,\frac{\mathrm{kg}}{\mathrm{m^3}}$ (Standard Conditions)

Example: Notebook-Computer (30 W)

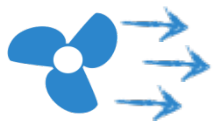

Room: 25 °C

max. 50 °C

$\Delta T = 25\,\mathrm{K}$
$\Rightarrow Q = 1{,}0\,\frac{\mathrm{l}}{\mathrm{s}}$

Example: Data Center (1 MW)

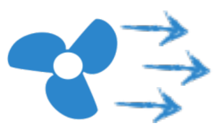

Supply Air: 18 °C

Exit Air: 43 °C

$\Delta T = 25\,\mathrm{K}$
$\Rightarrow Q = 33.600\,\frac{\mathrm{l}}{\mathrm{s}}$

Water Specific Heat Capacity: $c_\mathrm{p} = 4{,}183\,\frac{\mathrm{kJ}}{\mathrm{kg \cdot K}}$ Density: $\rho = 997{,}0\,\frac{\mathrm{kg}}{\mathrm{m^3}}$ (Standard Conditions)

Example: Data Center (1 MW)

$\Delta T = 5\,\mathrm{K}$
$\Rightarrow Q = 48{,}0\,\frac{\mathrm{l}}{\mathrm{s}}$

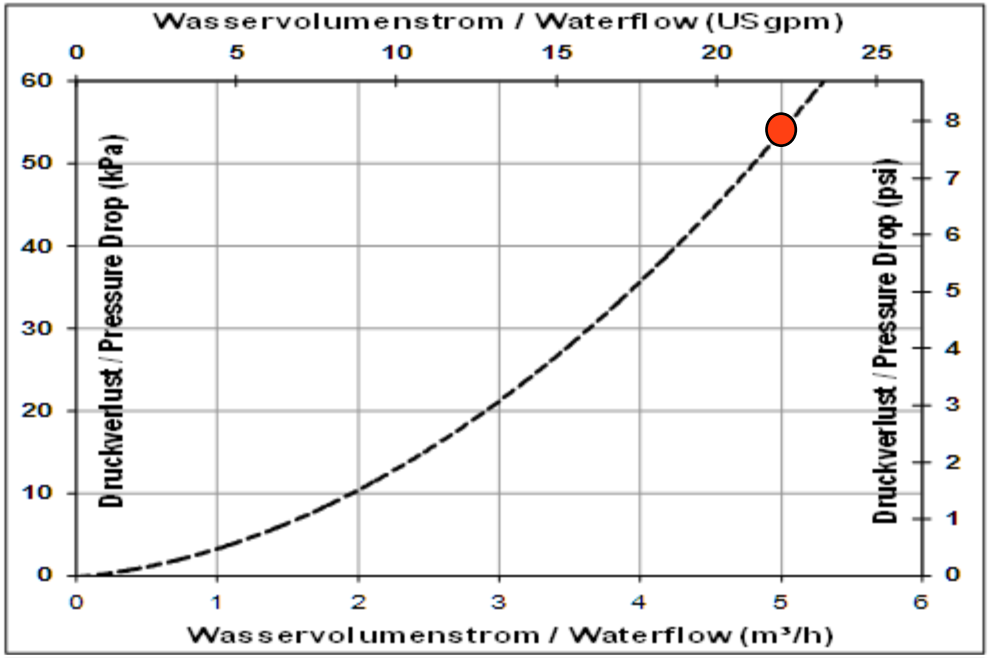

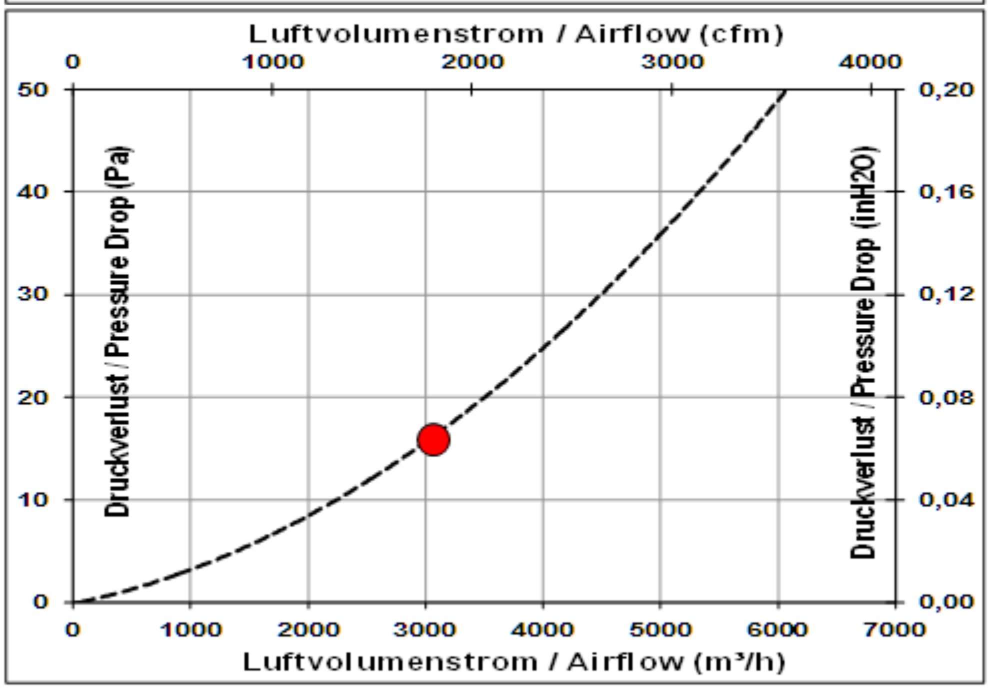

knürr environments for electronics

Heiko Ebermann

LOEWE-CSC 2011

Green500:	8	(2010)
Top500:	33/21	(2010/11)
Cost:	200 €/core (2010)	

CO_2 neutral


fan power 6%
of server power

Innovative Cooling System Architecture

Cooling Overhead
5% of computer power

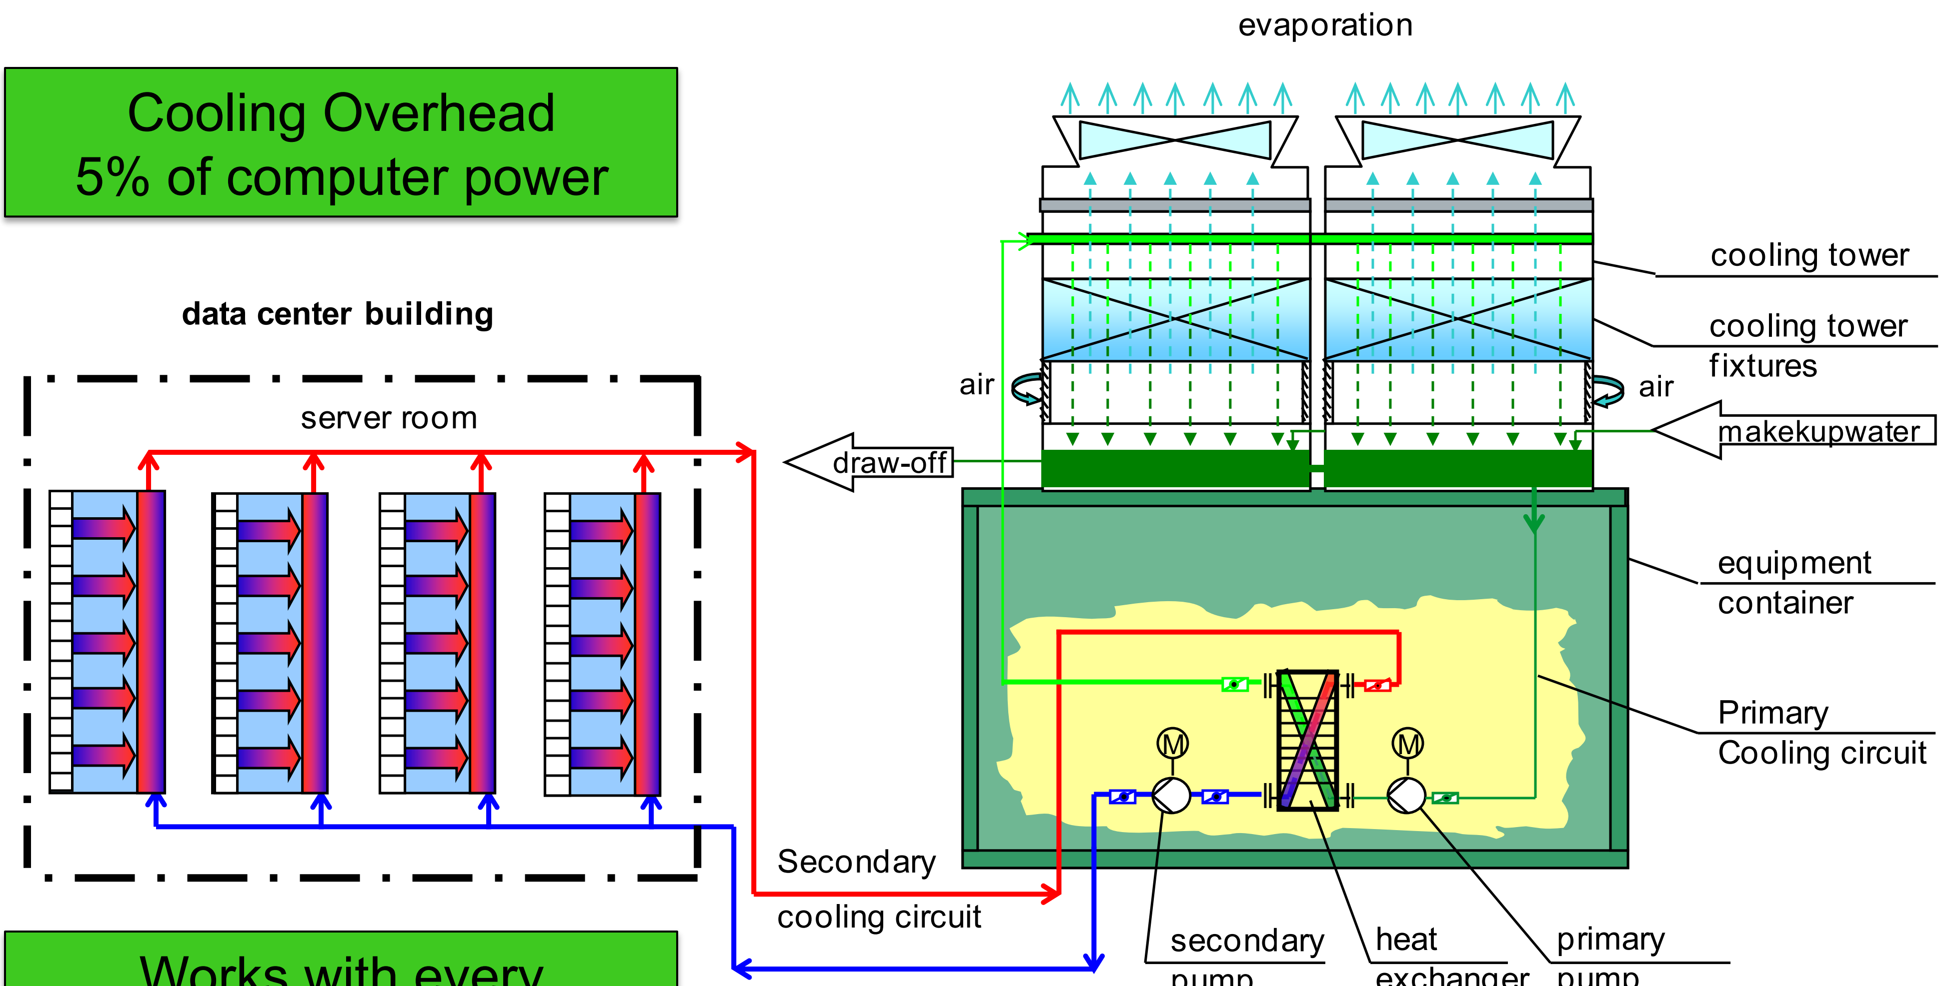

Works with every
commercial computer

Wet bulb temperature 2010

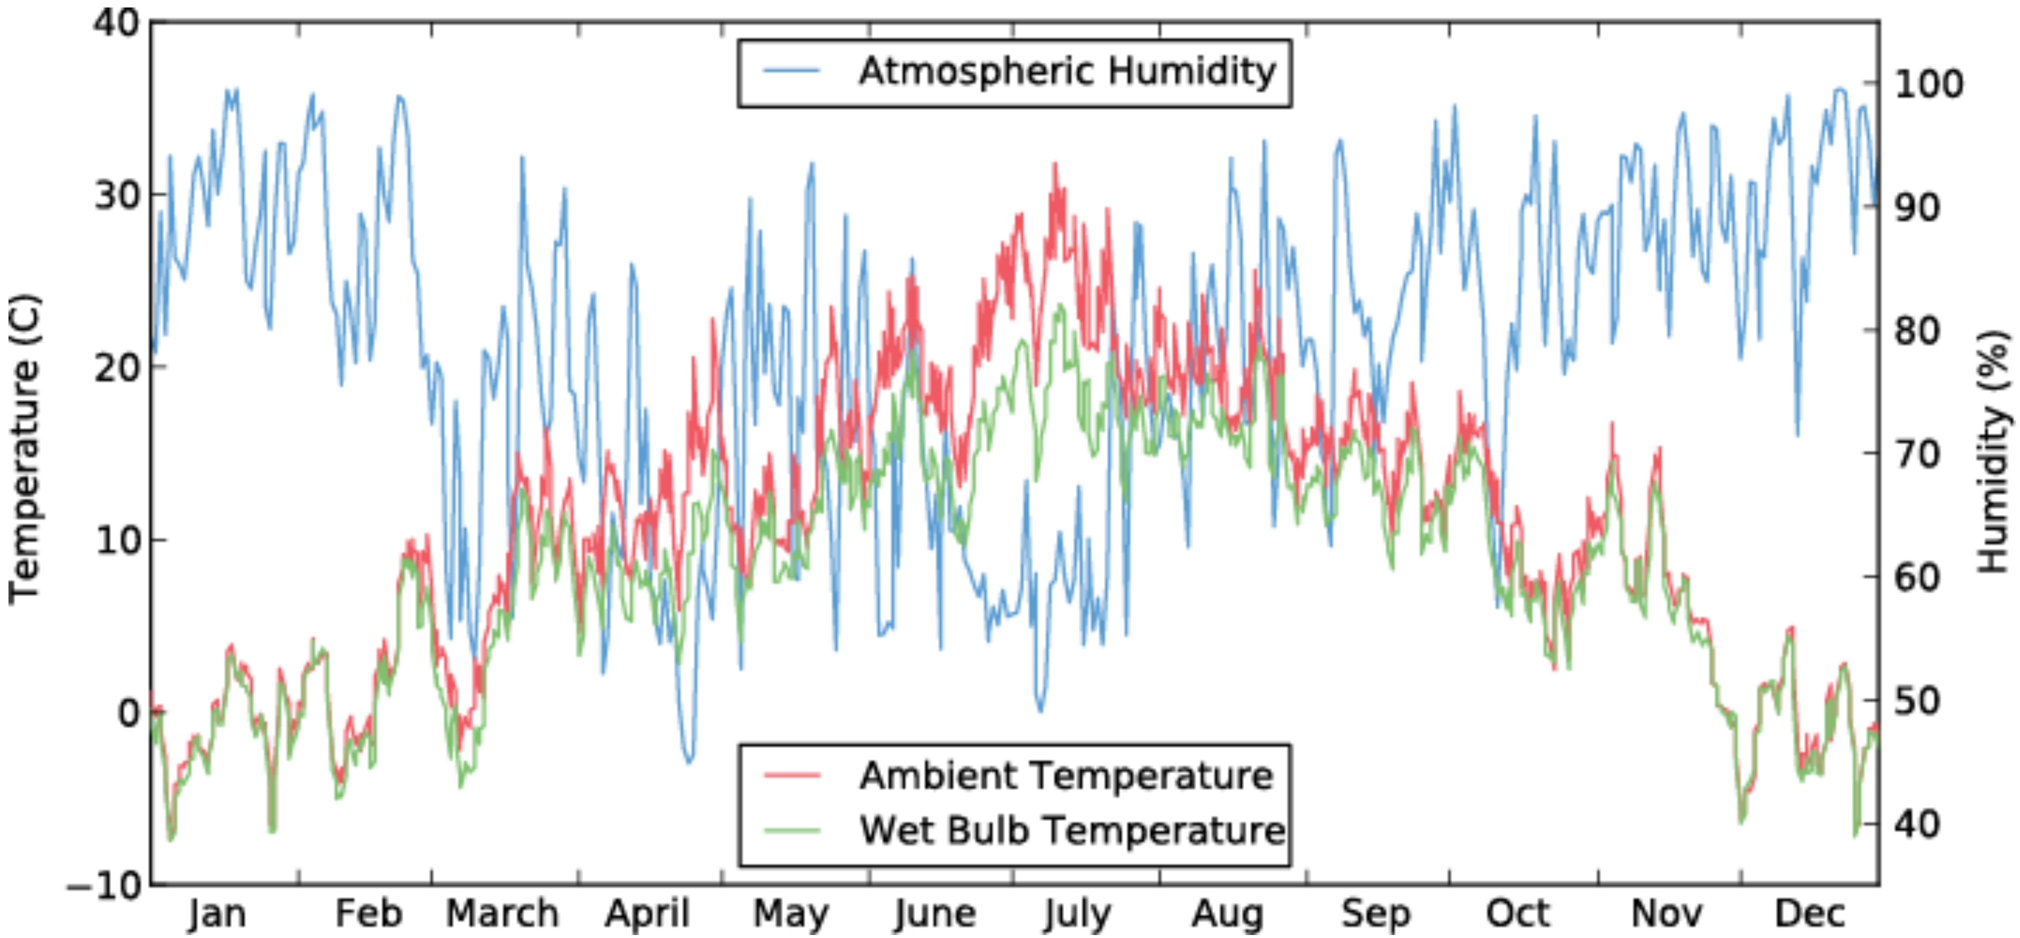

LOEWE-CSC Cooling (operational >2a)

• Max cooling power	900 kW
• Secondary pump	28 kW
• Primary pump	6 kW
• Cooling tower Fans not required if outside temperature below 15 °C	2x4.5 kW

→PUE (best)	1.05
→PUE @ 450 kW	1.07

• Server fan overhead	6%

FAIR Russia Research Center
Data Center Container

- Space for 13 19“ racks
- 300 kW cooling capability
- All inclusive but back cooler
- Cooling power 5% of computer power
- Turn key
- 190 k€ including racks and electric
- AB redundant cabling
- Redundant pumps
- Operable up to -40°C
- Just arrived in Moscow (10/2012)

3D data center

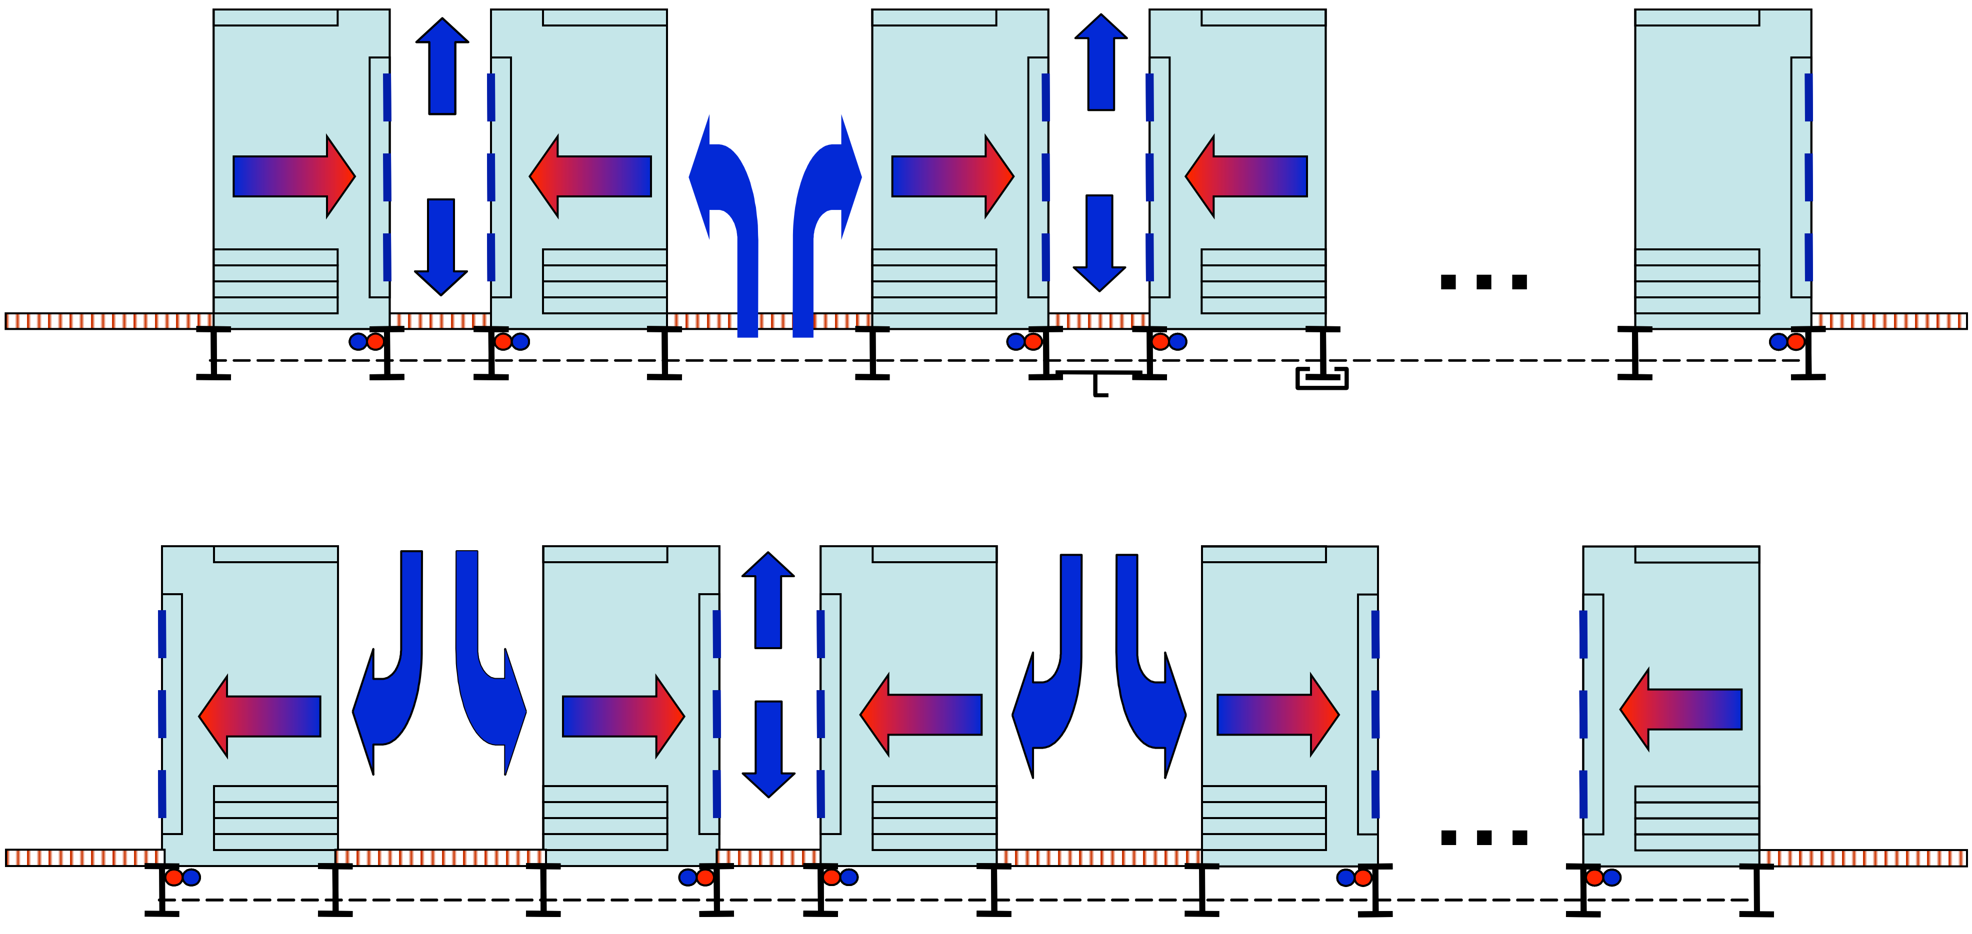

FIAS Patent submitted @ DPMA/PCT 5.7.2008 (Lindenstruth, Stöcker)


PDR 3 (Reihe 3)

Cooler for GSI MiniCube

The FAIR Tier-0 Green Cube

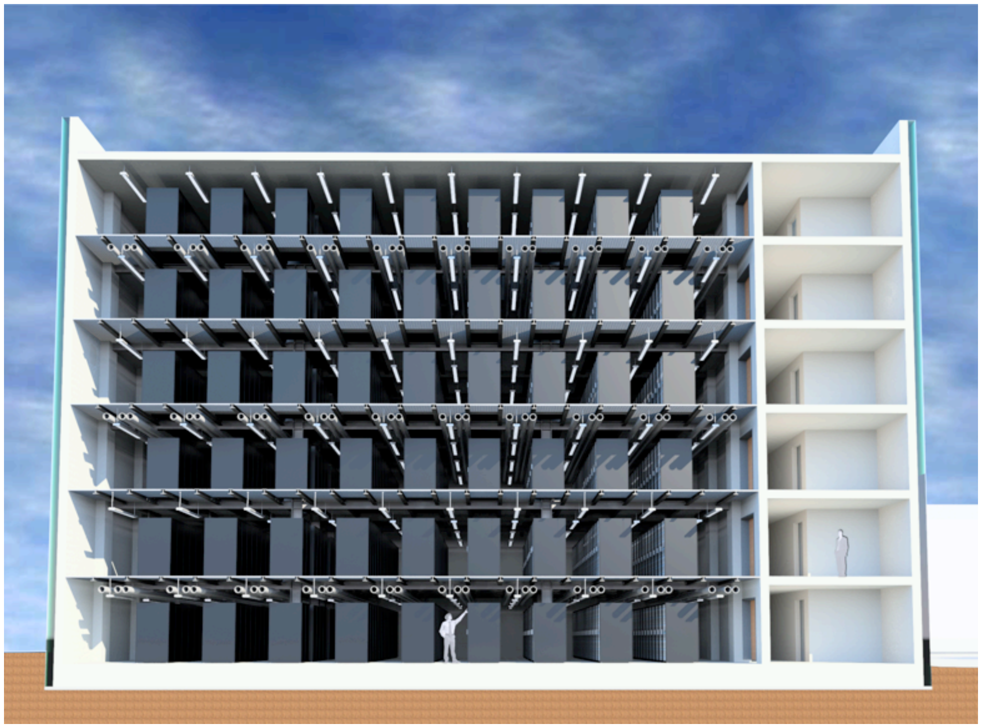

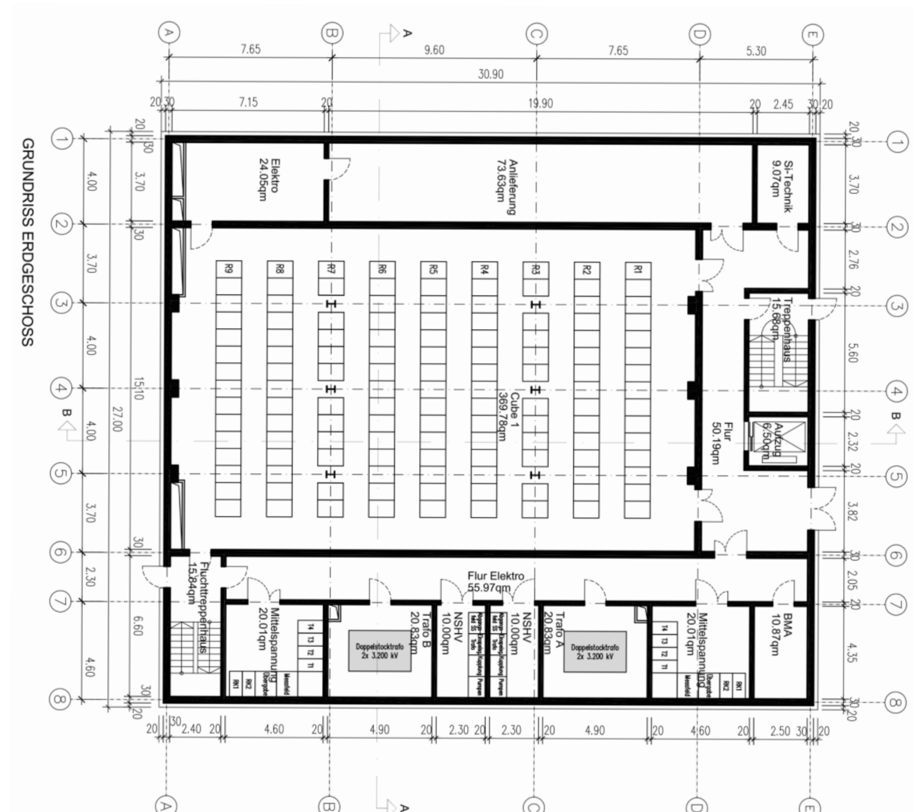

- Space for 780 19” racks (2,2m)
- 6 MW cooling (baseline)
- Max cooling power 16 MW
- Fully redundant (N+1)

- Can be used for any commercial IT
- No battery backup required
- Target PUE 1.05
- Baseline building cost 11 M€

Prices

- **Green-IT Award of German federal government „Visionary concepts“ – November 2011 for LOEWE-CSC**
- **3. rank DataCenterDynamics Awards in December 2011 category „Green IT“ for LOEWE-CSC**
- **2. rank DataCenterDynamics Awards in December 2011 category „DataCenterBlueprint“ for GSI MiniCube**
- **„Land der Ideen“ January 2012 for LOEWE-CSC**
- **German Datacenter price 2012 category energy efficient cooling – March 2012 for GSI MiniCube**
- **„Green Cube“ BMBF project of the month April 2012**

Deutschland
Land der Ideen
Ausgewählter Ort 2012

FAIR HPC Backbone 120 GB/s →1TB/s

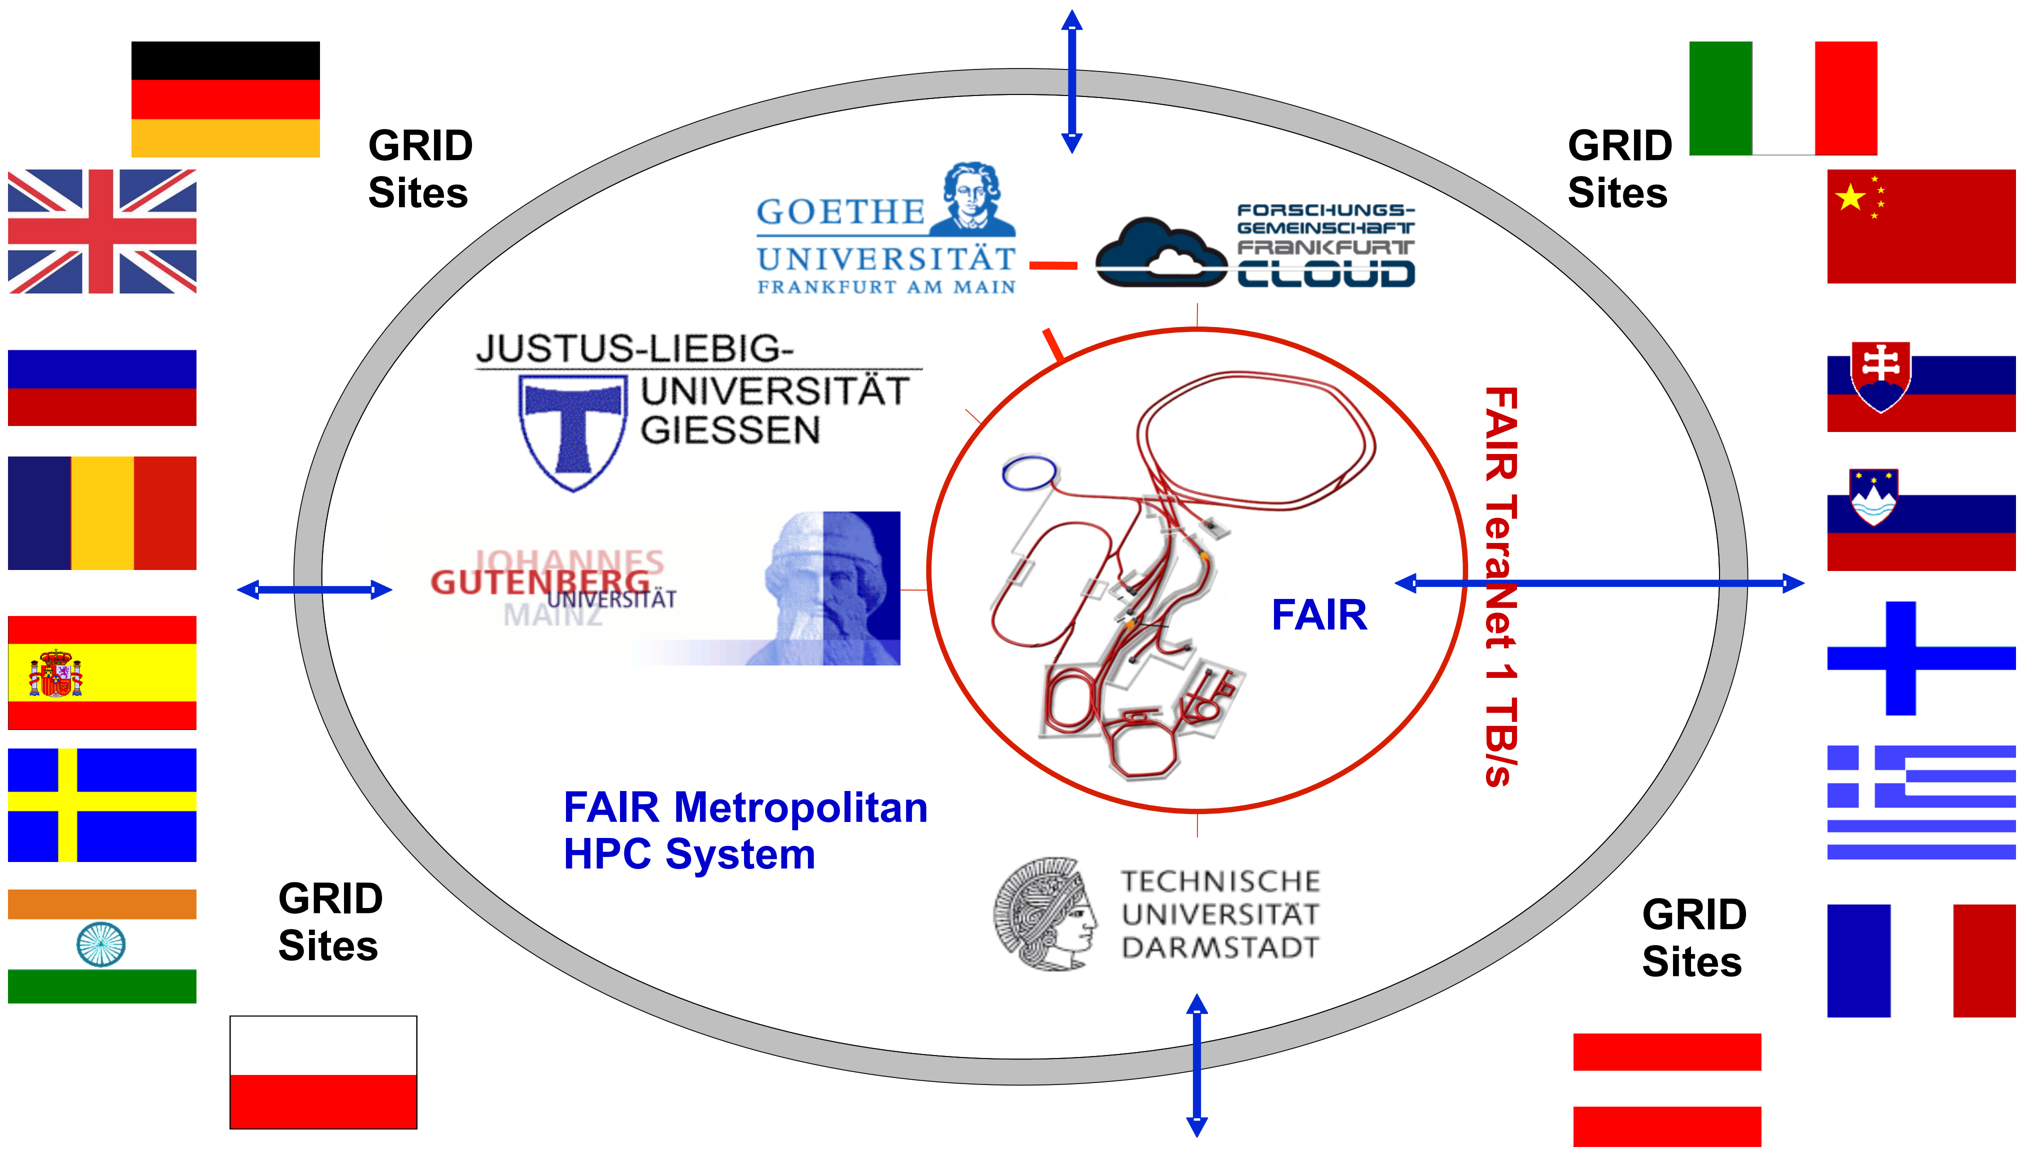

Vector Classes (Vc)

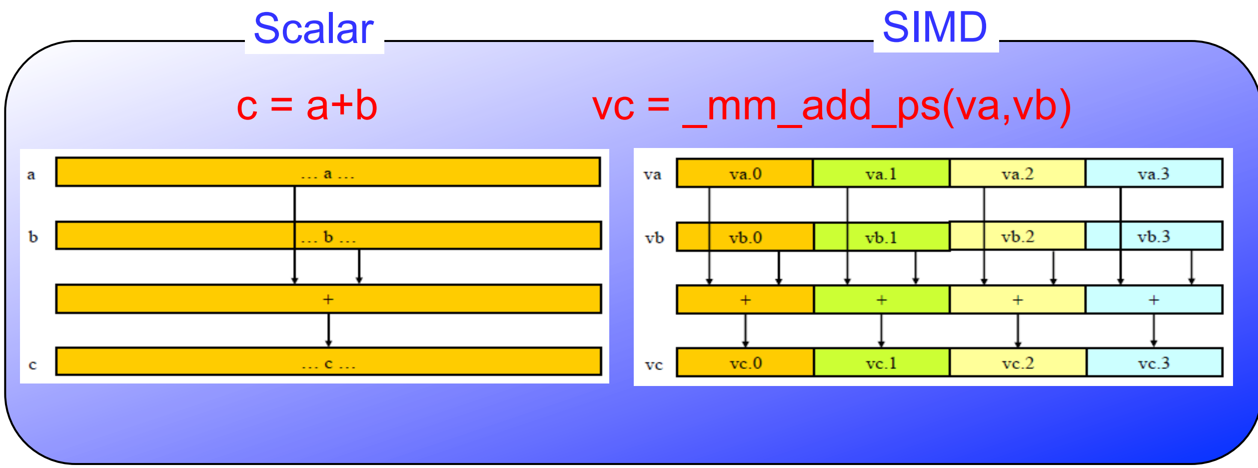

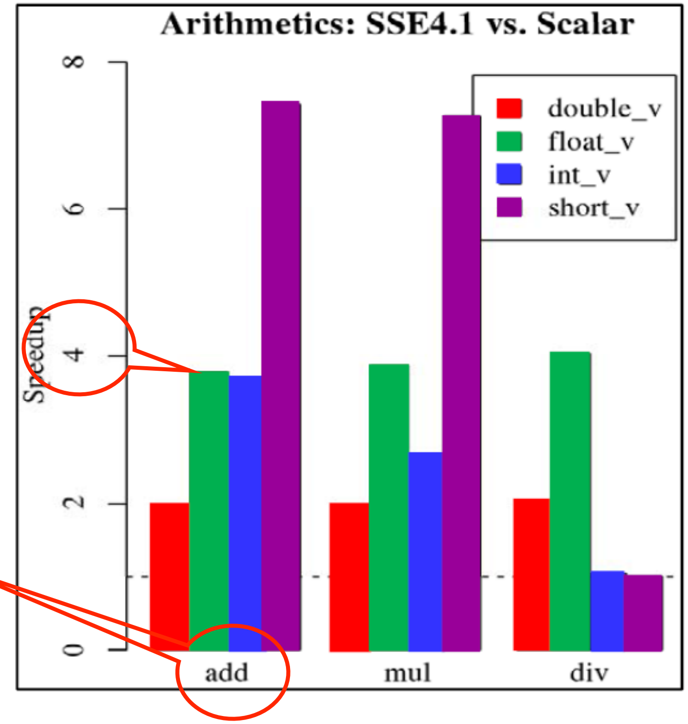

Vector classes overload scalar C operators with SIMD/SIMT extensions

Vector classes:

- provide full functionality for all platforms
- support the conditional operators

phi(phi<0)+=360;

Vc increase the speed by the factor:

- ✓ SSE2 – SSE4 4x
- ✓ future CPUs 8x
- ✓ MICA/Larrabee 16x
- ➢ NVIDIA Fermi research

Vector classes enable easy vectorization of complex algorithms

Modern Programming

CBM

Stage	Description	Time/track	Speedup
	Initial scalar version	12 ms	–
1	Approximation of the magnetic field	240 μs	50
2	Optimization of the algorithm	7.2 μs	35
3	Vectorization	1.6 μs	4.5
4	Porting to SPE	1.1 μs	1.5
5	Parallelization on 16 SPEs (2 Cells)	0.1 μs	10
	Final SIMDized version	0.1 μs	120,000

Summarized stages of the porting procedure

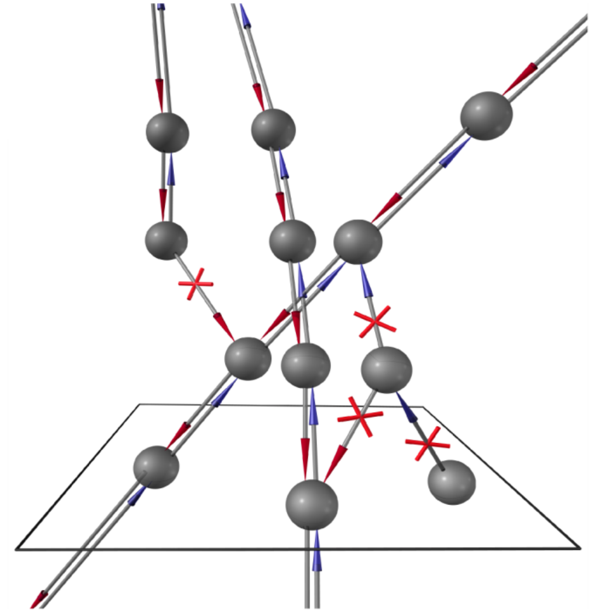

Fast simdized Kalman filter based track fit
U. Kebschull, I. Kisel, V. Lindenstruth, W.F.J. Müller
Computer Physics Communication **2007**

ALICE

CA, step1 (Pb-Pb collision)**:**
CPU time = 690 ms
GPU time = 39 ms :
- 25.3 ms algorithm,
- 13.7 ms CPU ↔ GPU data transfer

Hardware options for the parallel computation:
- SIMD CPU instructions (Vectorization)
- multi-threading
- multi-core CPU
- many-core hardware (Graphics cards, Larrabee, …)

DGEMM Scalability

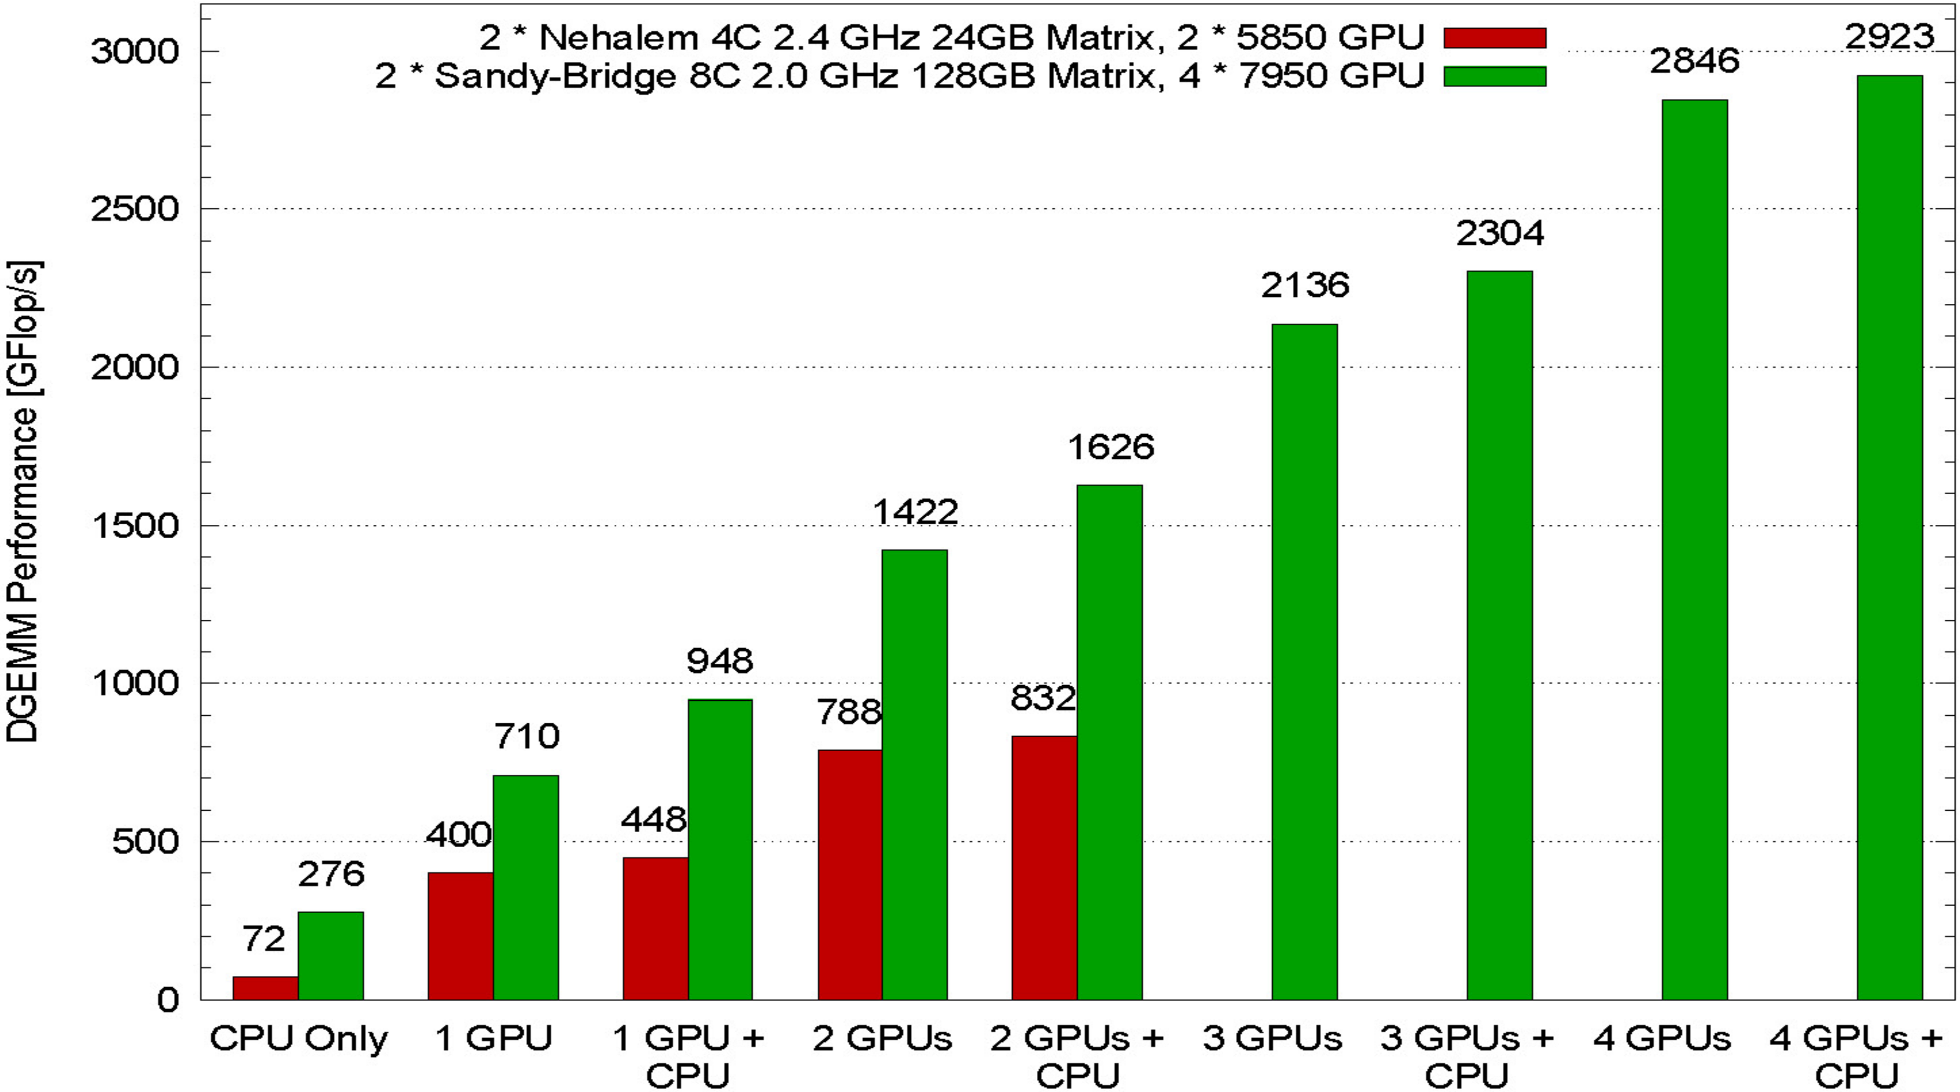

Optimized HPL for AMD GPU and Multi-Core CPU Usage, ISC June 2011

Ultrarelativistic Quantum Molecular Dynamics

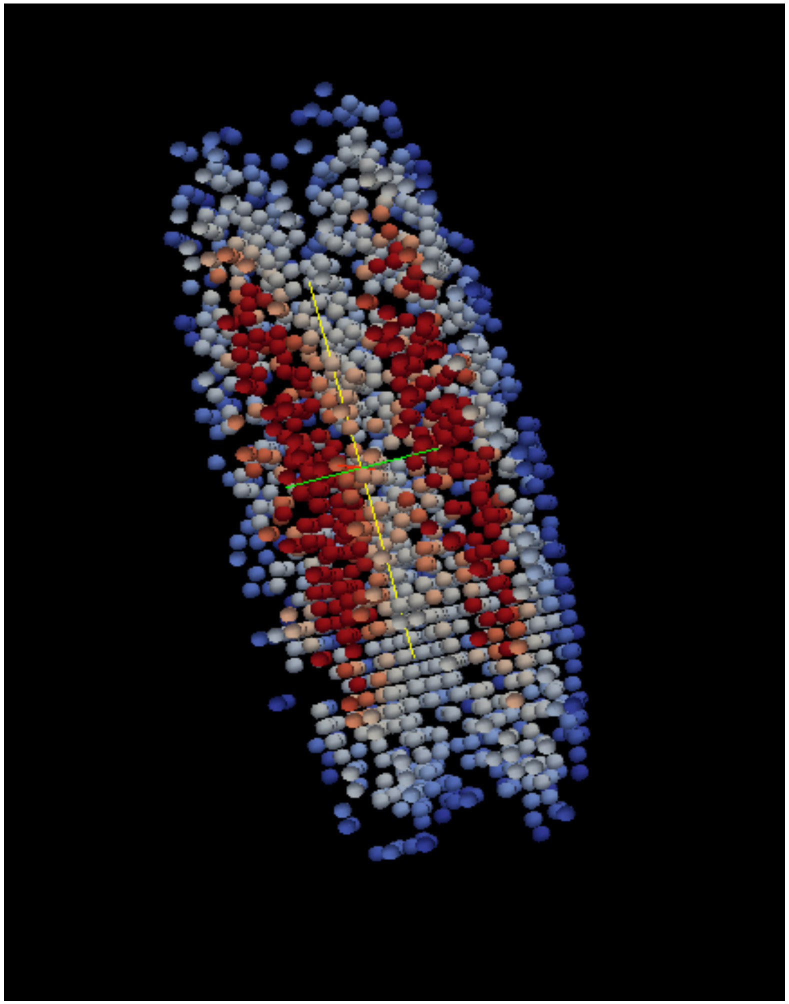

- Simulate Pb+Pb collisions with a gas phase.
- Use grid resolution of 0.2 fm.
- Millions of cells are used to compute quantities, like energy density or baryo-chemical potential.
- A very demanding computation: takes hours ...
- Redesign algorithms to work on modern hardware: LOEWE-CSC has more than 800 GPUs ...
- New code 500 x faster than old code, but uses only 1/5 of memory
- Now unprecedented event-by-event calculations can be carried out.
- Better statistics

Lattice QCD @ LOEWE-CSC

- **Lattice QCD**
  - **QCD = Strong Interaction**
  - **Lattice QCD = Only a priory approach for QCD**
  - **Discretize Space in 4-dim Lattice, extrapolate solutions from different lattice spacings**
  - **Computationally sparse**

- **GPUs Excell at Bandwidth**

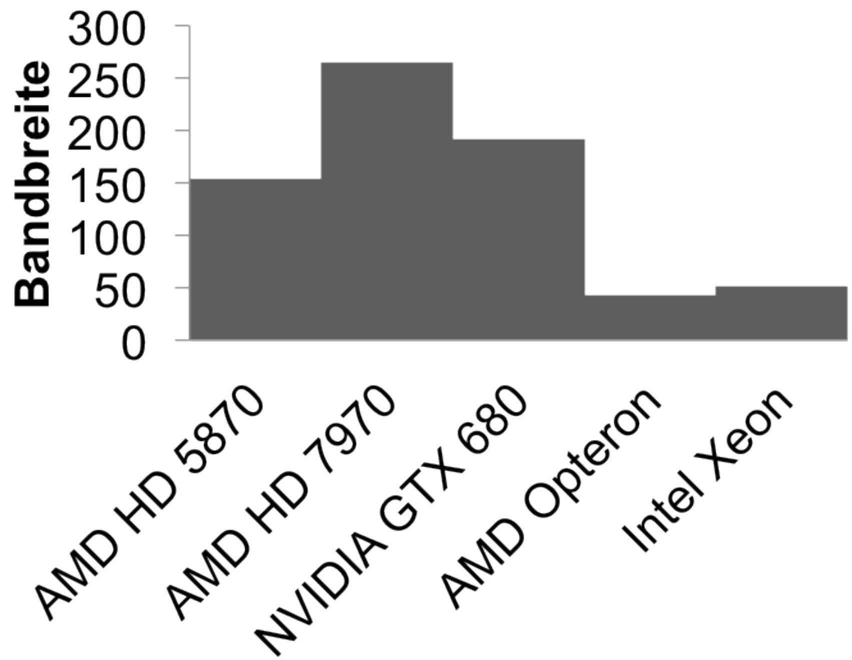

Bach, M., Philipsen, O., Pinke, C., Schäfer, C., & Zeidlewicz, L. (2011). LatticeQCD using OpenCL. http://arxiv.org/abs/1112.5280

FAIR @ KACST HPC System

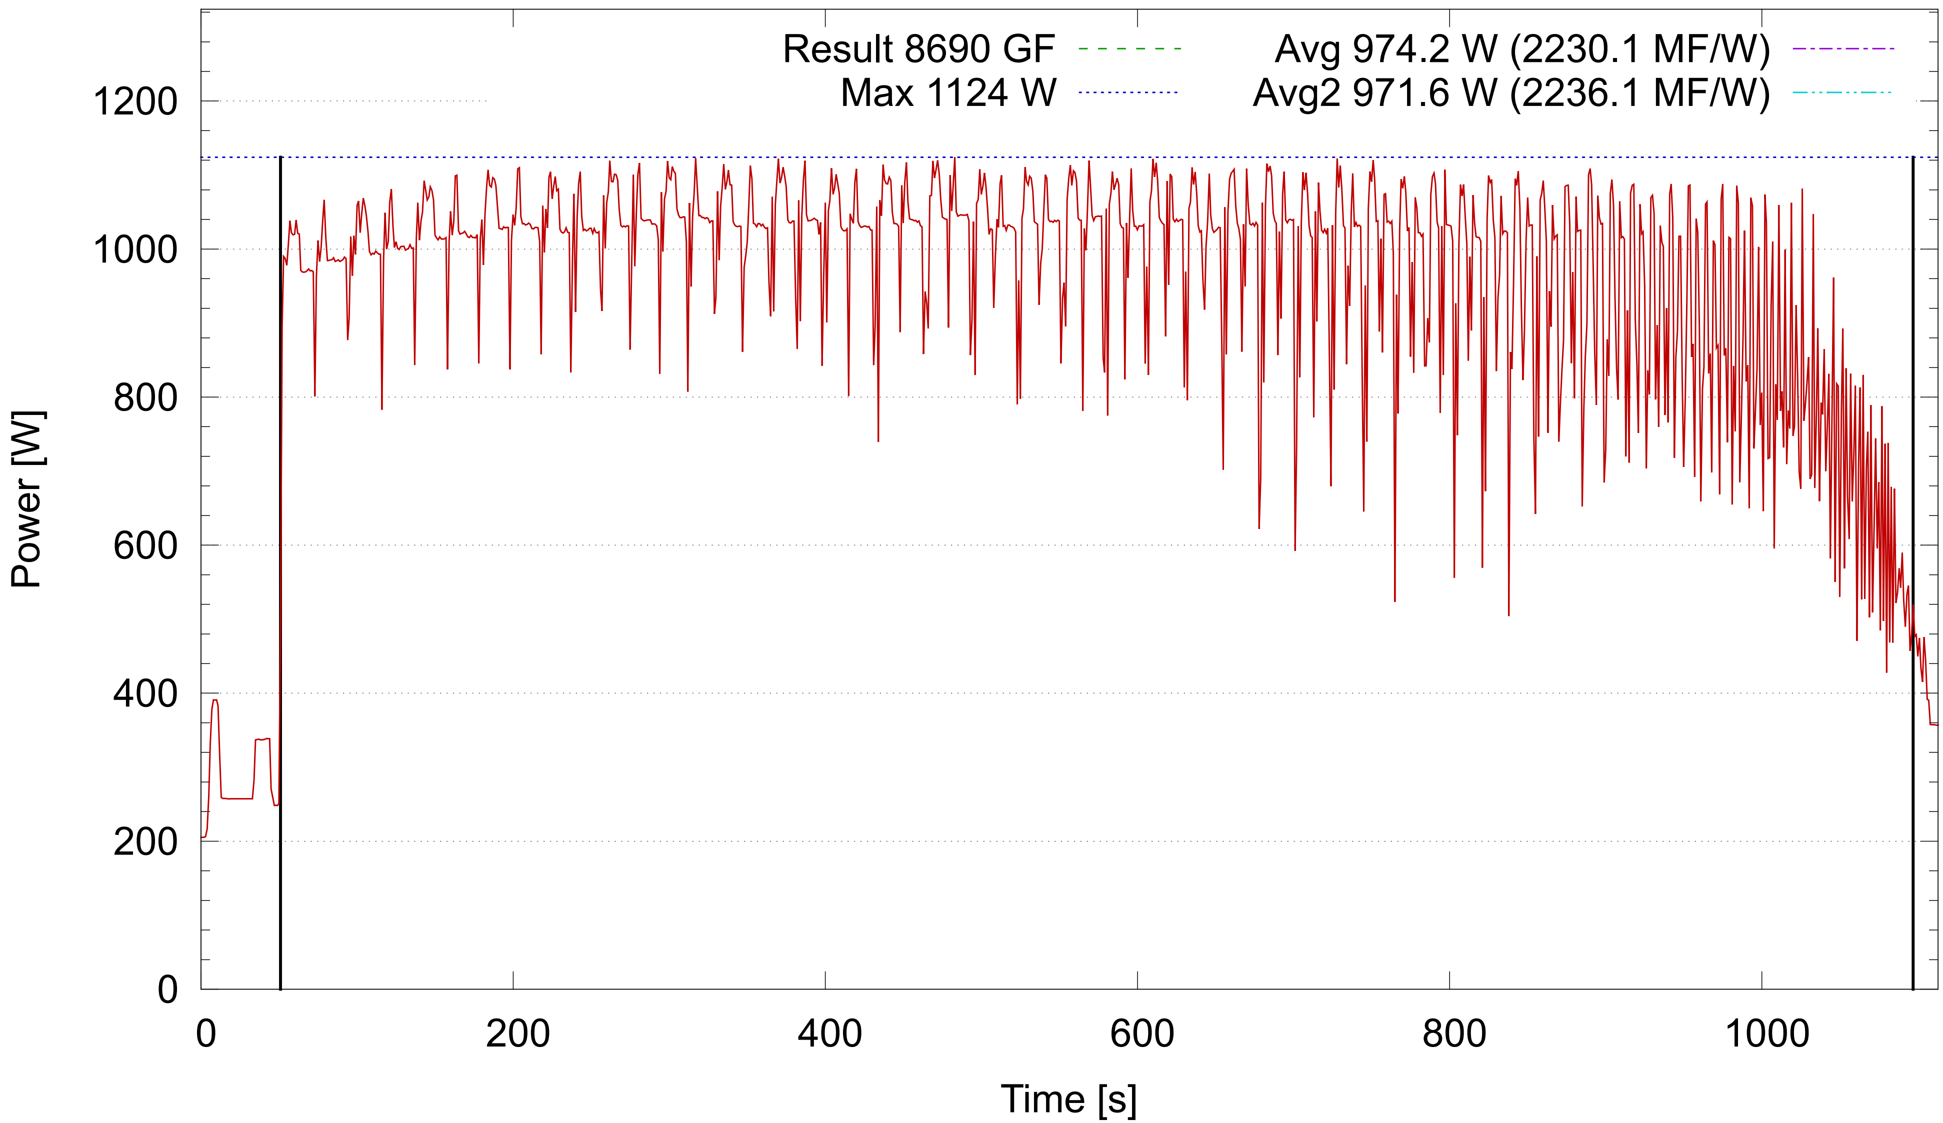

Conclusions, Outlook

- Indirect passive cooling generates very efficient cooling of PUE<10%
- Ambient temperatures in Europe support entire year free cooling operation
- Data center heat can be used to heat buildings with floor heating capability
- Any commercial servers are supported
- No active components or air ducts in the data center
- Highest power density of 20kW/m^2 per floor
- Efficient programming one of the largest savings potentials
- HPL, UrQMD, Lattice-QCD, high speed tracking already working well on GPGPUs
- Efficient algorithms have to consider computer architecture


thank you