TECHNISCHE UNIVERSITÄT DRESDEN

Dresden Database Systems Group

HAEC

ScaDS DRESDEN LEIPZIG

Energy-efficient Computing -
The Impact of Software

Prof. Dr.-Ing. Wolfgang Lehner

Are we done?

Dresden Database Systems Group

ENERGY TECHNOLOGIES A

HOME | RESEARCH AND DEVELOPMENT | RES

All Publications | Authors | Organization | Research Areas | Keywords

Title: United States Data Center Energy Usage Report

Publication Type: Report

Year of Publication: 2016

Authors: Arman Sh, Herrlin, Jo, Azevedo, L

Date Published: 06/2016

Abstract: This repor relying on consumpt In 2014, d representi that data large shift

Related Files

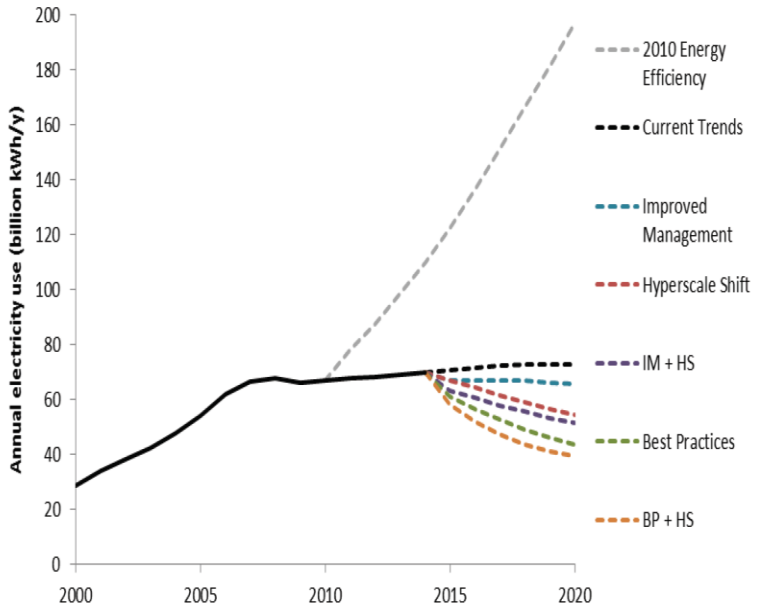

United States Data Center Energy Usage Report (2016)

- estimate of total U.S. data center electricity use (servers, storage, network equipment, and infrastructure) from 2000-2020
- In 2014: consumption of an estimated 70 billion kWh
  - about 1.8% of total U.S. electricity consumption
- Increase of electricity consumption
  - about 4% from 2010-2014
  - about 24% from 2005-2010
  - nearly 90% from 2000-2005
- Energy use is expected to continue slightly increasing in the near future, increasing 4% from 2014-2020
  - data centers are projected to consume approximately 73 billion kWh in 2020

Data is produced continously

Dresden Database Systems Group

Smart Everything

- Smart „things“
- Smart places
- Smart networks
- Smart services
- Smart solutions

→ „Smart-*“ infrastructure

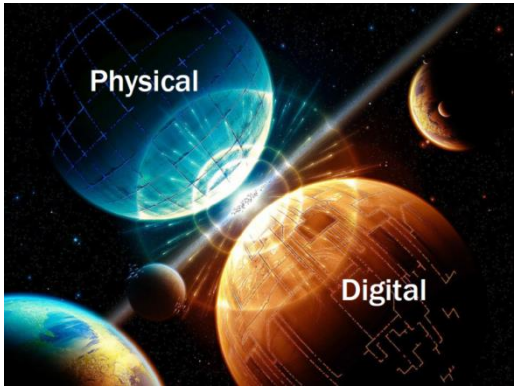

Physical and digital worlds collide!

- need to make things Smart...!
- Requirements for “Smart Everything”
  - Interactive (“tangible”) → low latency
  - High volume → high throughput

...from smart phone to smart lenses

your personal coupon arrived!!!
Buy x get y free

Dresden Database Systems Group

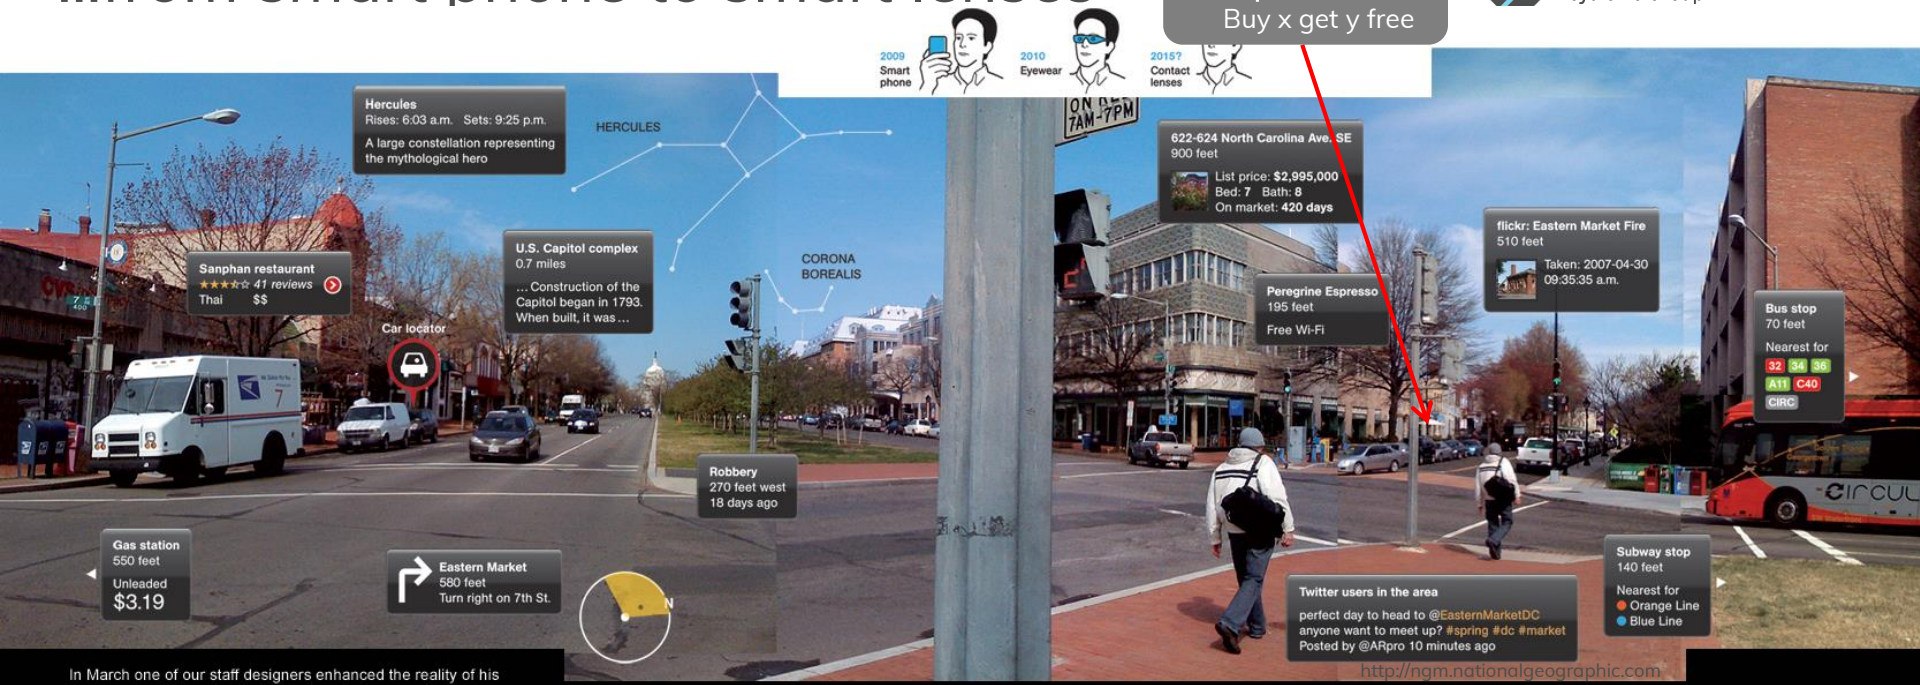

In March one of our staff designers enhanced the reality of his Washington, D.C., neighborhood. Smart phone applications (apps) added layers of information to what he saw—called out in this composite of five photos, each taken with his phone.

UP AND AWAY Point your phone at the sky and find stars hidden by daylight. Aim at a tourist spot and see its history plus info for visitors. For an augmented-reality check, tap into crime stats.

REAL DEALS Various apps can steer you to the cheapest gas around, mass-transit options, good food, and Wi-Fi spots. You can also learn the price of that town house that's up for sale.

STREET PALS The Tweeps Around app tells if tweeters are near. Flickr displays area photos by members (Eastern Market, above). In the works: an app to match faces to social-network profiles.

http://ngm.nationalgeographic.com

→ novel **Big Data Analytics** apps with ms-response time incorporating local context as well as global state

Research Project: Energy Efficent Computing

Dresden Database Systems Group

„The Fog“: Core / Edge / Cloud-Computing

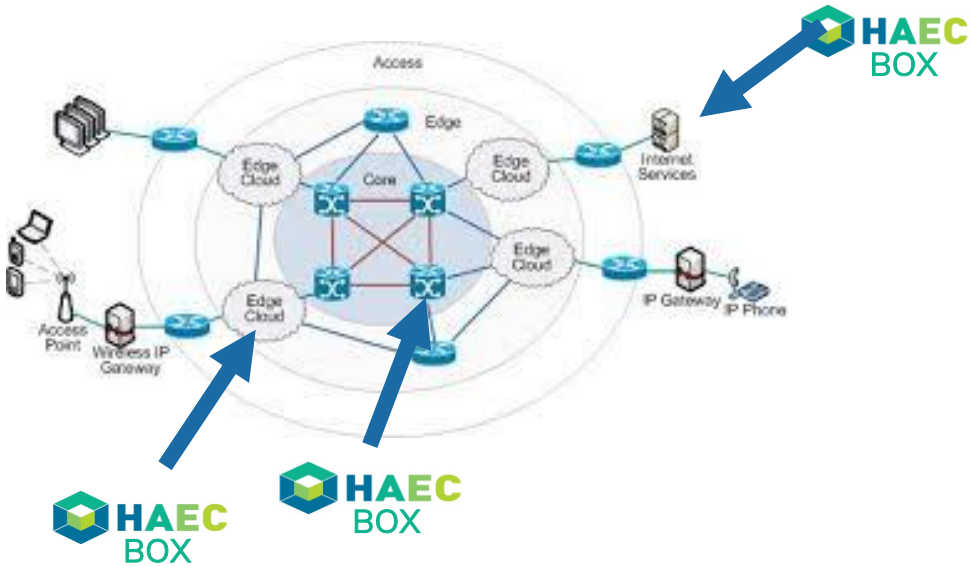

Information Management in
... Traffic
... Industrie / ...

„HPC in a Box“

... Simulation software
... Image processing
...

How to reduce energy ???

What at the main obstacles ???

Problem: Missing Energy Proportionality

Dresden Database Systems Group

Typical Workload Pattern

Center for Information Services and High Performance Computing (ZIH)

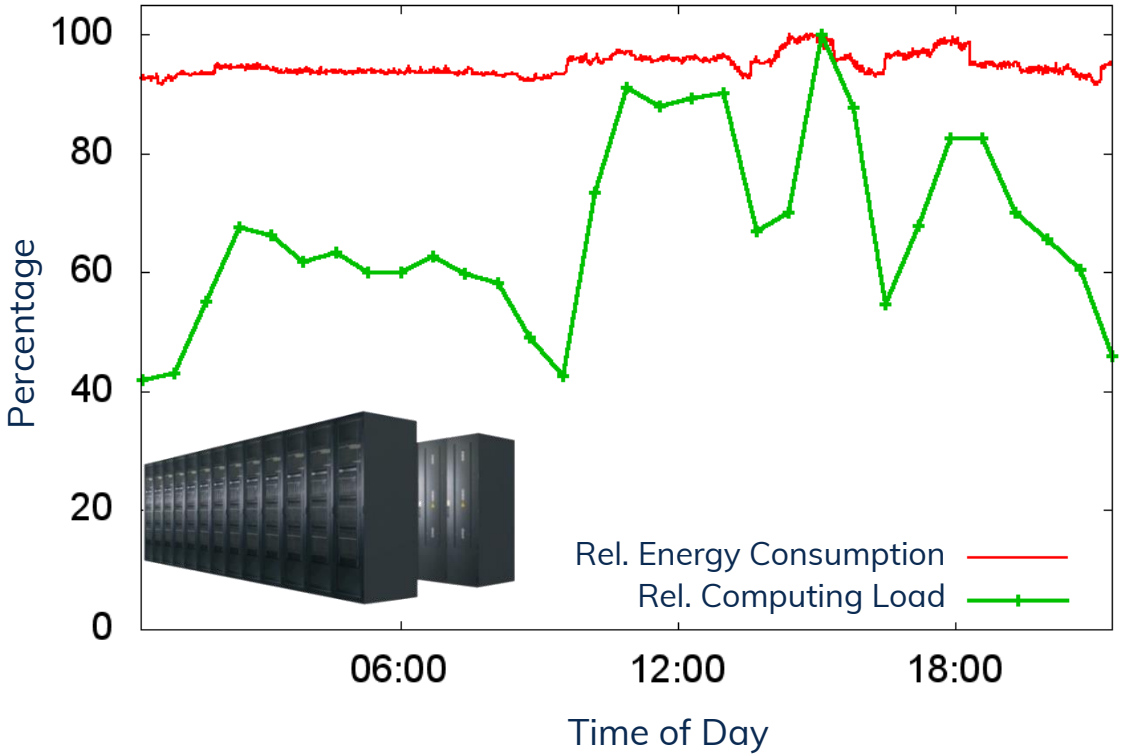

Missing Hardware Adaptivity

Dresden Database Systems Group

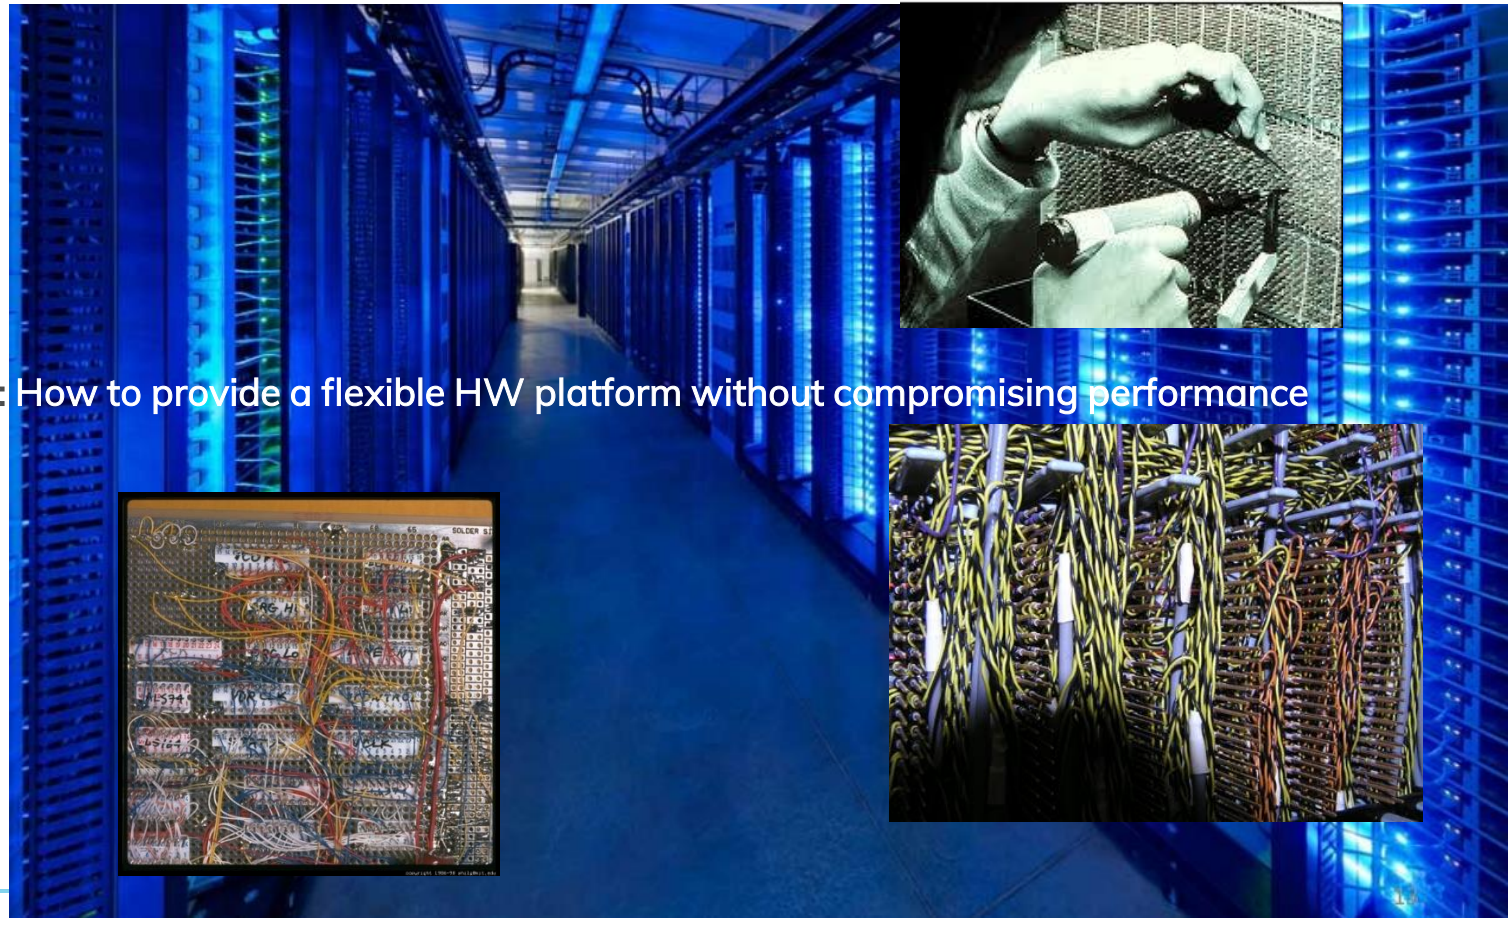

Key Question: How to provide a flexible HW platform without compromising performance

Missing Software Adaptivity

Dresden Database Systems Group

- For Software Components

→ WHAT to execute

- For System Scheduling

→ WHEN to execute

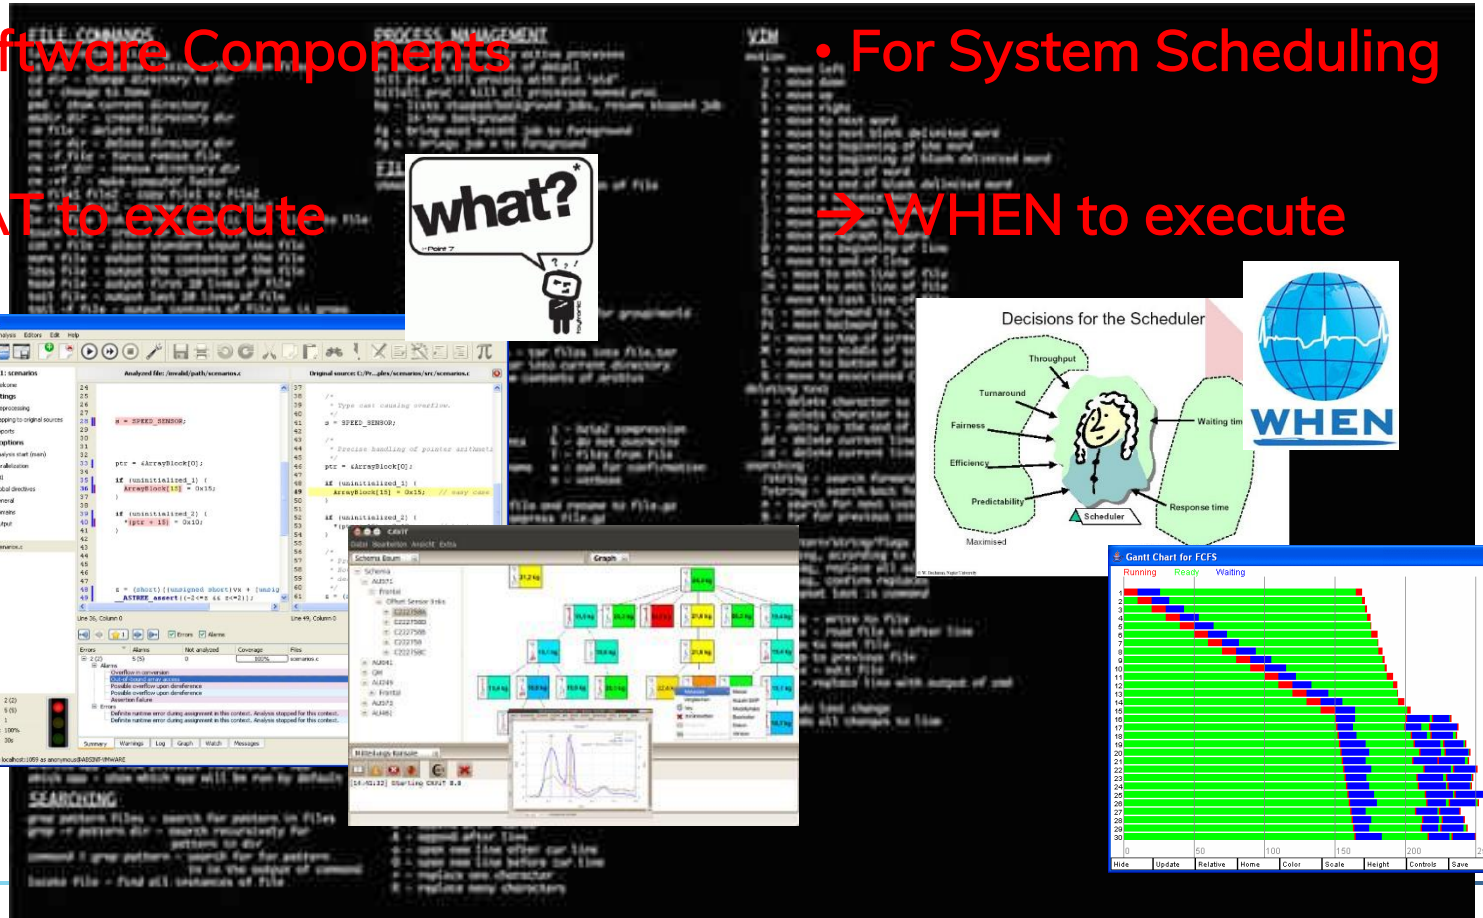

Highly Adaptive Energy-Efficient Computing

Dresden Database Systems Group

Simulation and Prototyping

Information Processing

- Application
- Middleware
- Runtime
- Operating System

Devices & Circuits

- Processor

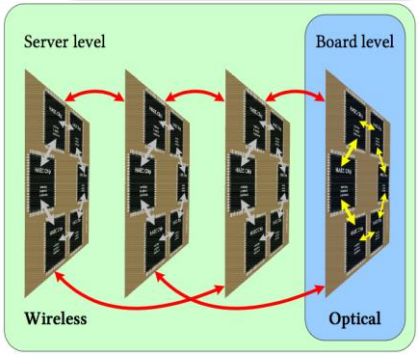

Flexible Software

Energy Control Loop

Flexible Hardware Interconnects

Energy-Adaptive High-Speed Computing Platform

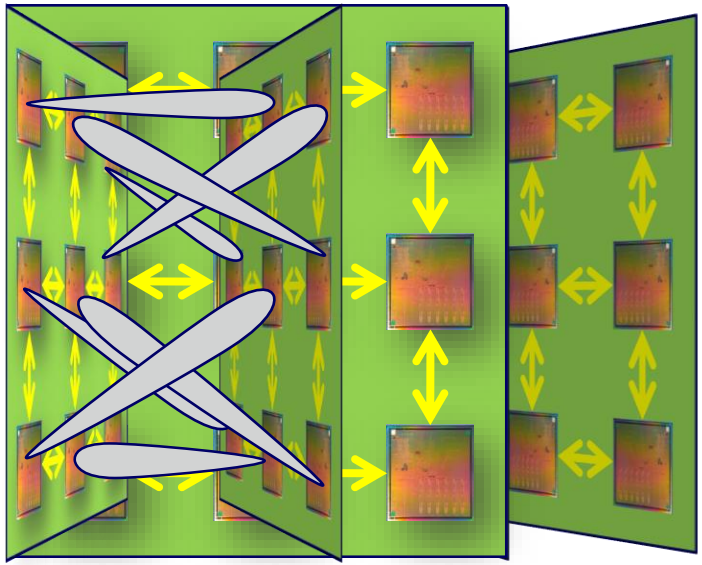

How realistic is this?

Optical Interconnect

- adaptive analog/digital circuits for e/o transceiver
- embedded polymer waveguide
- packaging technologies (e.g. 3D stacking of Si/III-V hybrids)
- 90° coupling of laser

Radio Interconnect

- on-chip/on-package antenna arrays
- analog/digital beamsteering and interference minimization
- 100Gb/s
- 100-300GHz channel
- 3D routing & flow management

Deployment and Increase of Compute Power

https://www.nvidia.com/en-us/data-center/tesla-v100/

https://www.top500.org/featured/systems/asci-white-lawrence-livermore-national-laboratory/

	Nvidia V100 (2017)	IBM ASCI White (2000)
Number of Processor Cores	3584	8192 (512 nodes x 16 IBM Power3)
Double-Precision Performance	7.5 TeraFLOPS	7.2 TeraFLOPS
NVIDIA NVLink™ v2 Interconnect Bandwidth	2x150 GB/s	N/A
PCIe x16 Interconnect Bandwidth	2x16 GB/s	N/A
Memory Capacity	16 GB	**6 TB DRAM** (Power 3 w/ **16 MB L2 cache**)
Max. overall data transfer speed	900 GB/s	?
Weight	450 gramm	106 tons
Energy consumption	**300W**	**3 MW**

Customizable Processor Model

Dresden Database Systems Group

	Intel i7-920		DBA_2LSU_EIS
Throughput (elements/s)	1,100 mio		1,203 mio
Clock frequency	2.67 GHz		0.41 GHz
Max. TDP	130 W	>>	0.135 W
Cores/Threads	4/8		1/1
Feature size	45 nm		65 nm
Area (logic & memory)	263 mm^2	>>	1.5 mm^2

~ ±x%

Scalable Data Management Runtime

Dresden Database Systems Group

ERIS

... We have it under control!!!
... It's built for adaptivity!!!

- Massively Parallel
- In-Memory
- Communication *w/ varying patterns*
- Adaptivity *Multiple dimensions*

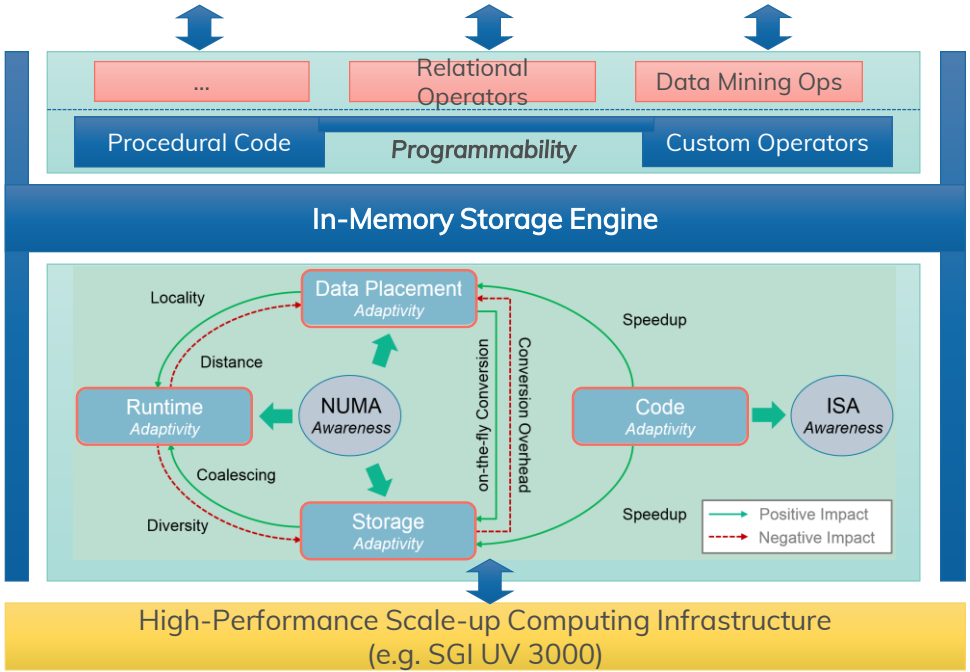

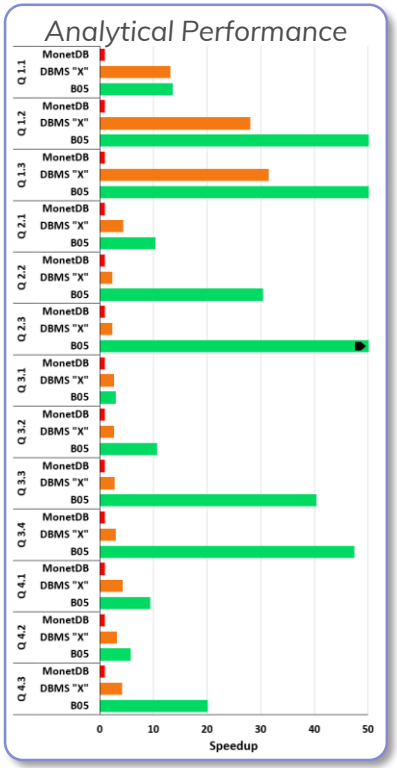

Energy Awareness

Dresden Database Systems Group

Power Breakdown Haswell-EP

- 19% static
- 81% dynamic

≅ load dependent

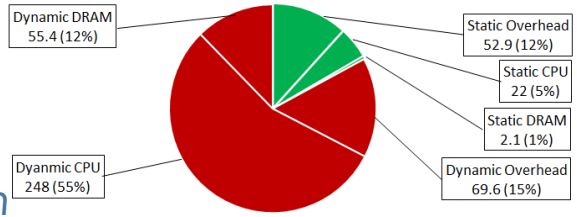

Initial Evaluation

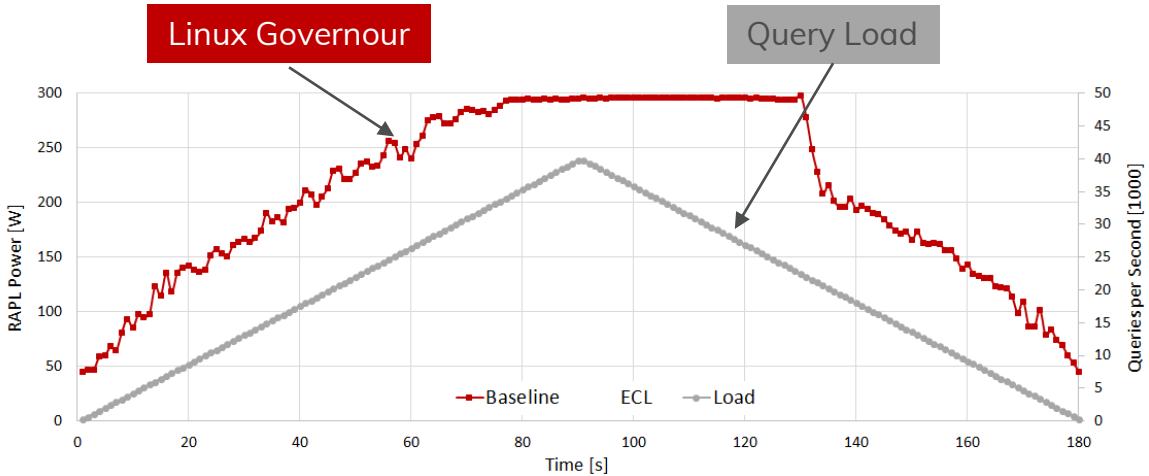

HARDWARE CONFIGURATION KNOBS

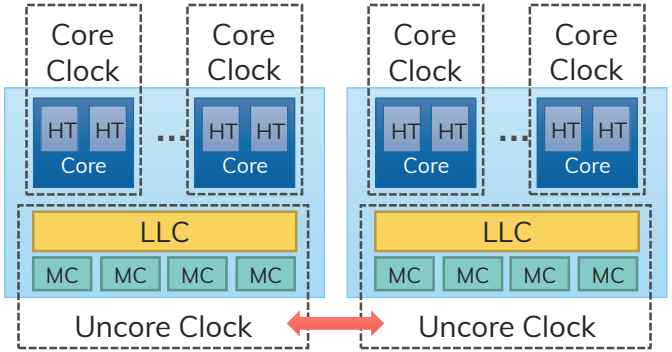

OBSERVATIONS:

1) There are opportunities
2) There are many knobs to tune

...but: workload knowledge makes a difference

Dresden Database Systems Group

same performance /
most energy efficient configuration

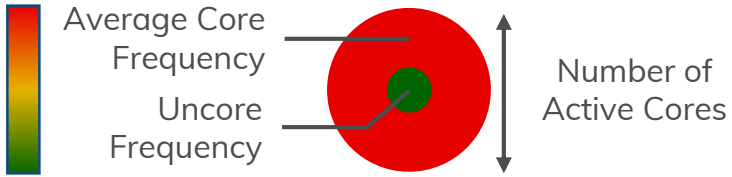

Lookup-Operation

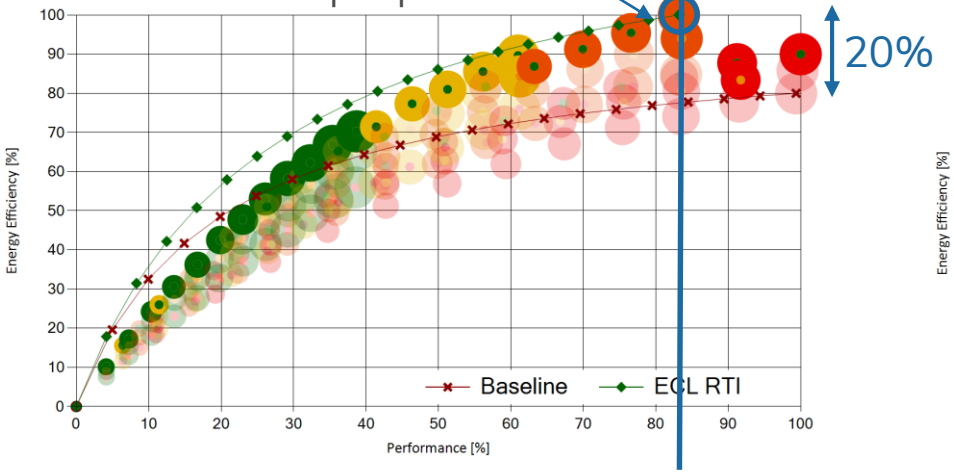

Scan-Operation

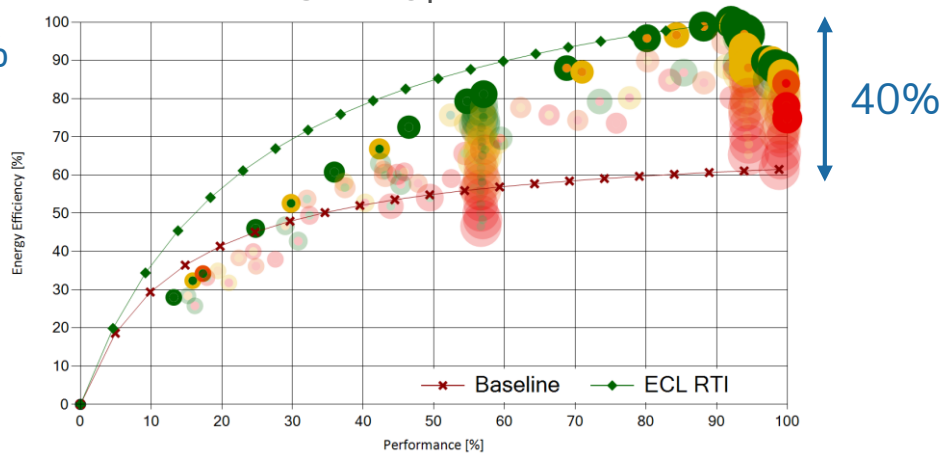

Typical Database Memory Access Patterns (Odroid-XU3) – dependency on operator

Dresden Database Systems Group

Compute-Intensive

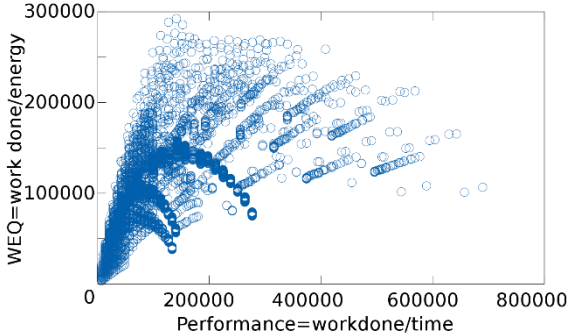

Scan-Operation

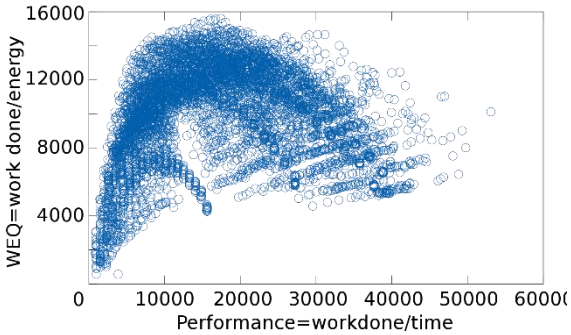

Lookup-Operation

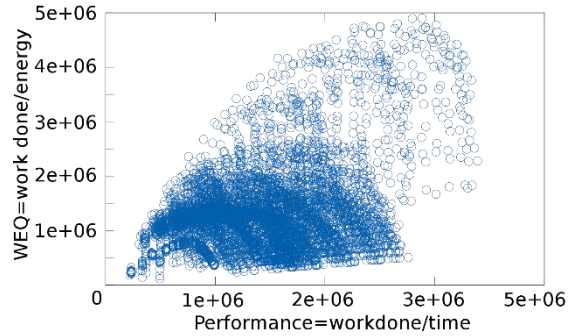

Copy-Operation

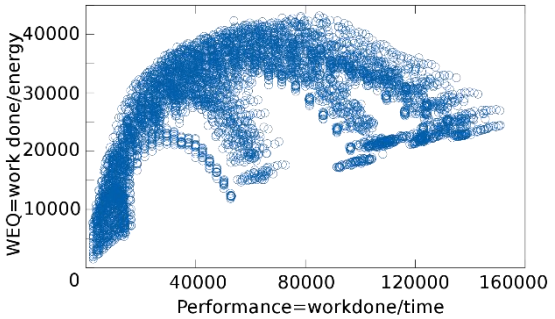

Let the DB do the job (ERIS ECL)

Dresden Database Systems Group

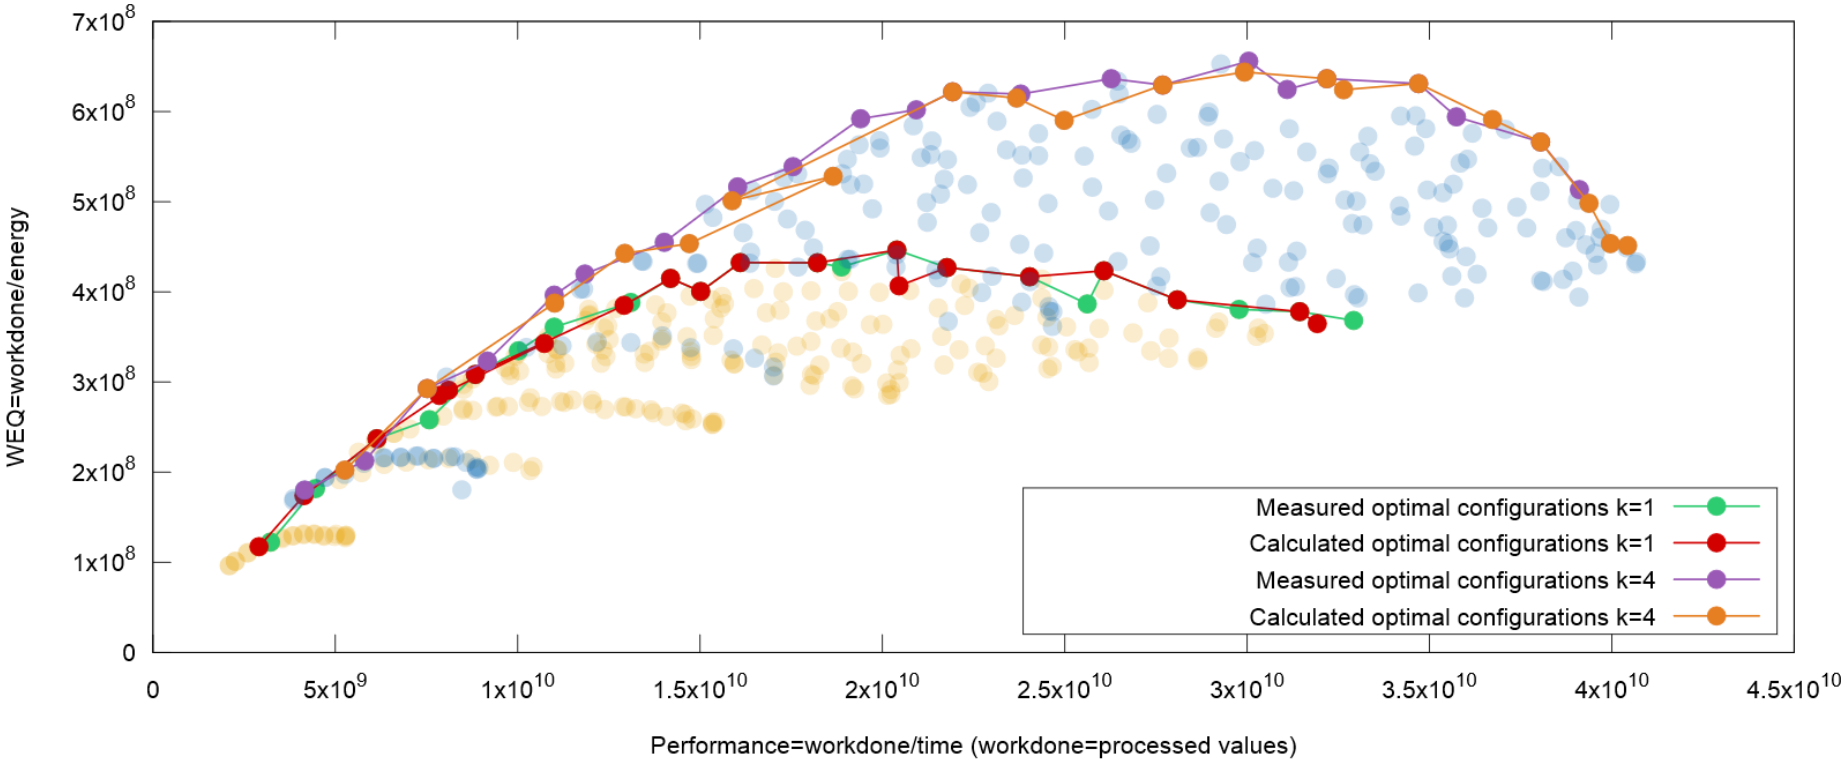

Benchmarking Energy Savings

Dresden Database Systems Group

Query Load

Linux Governour

DB-controlled

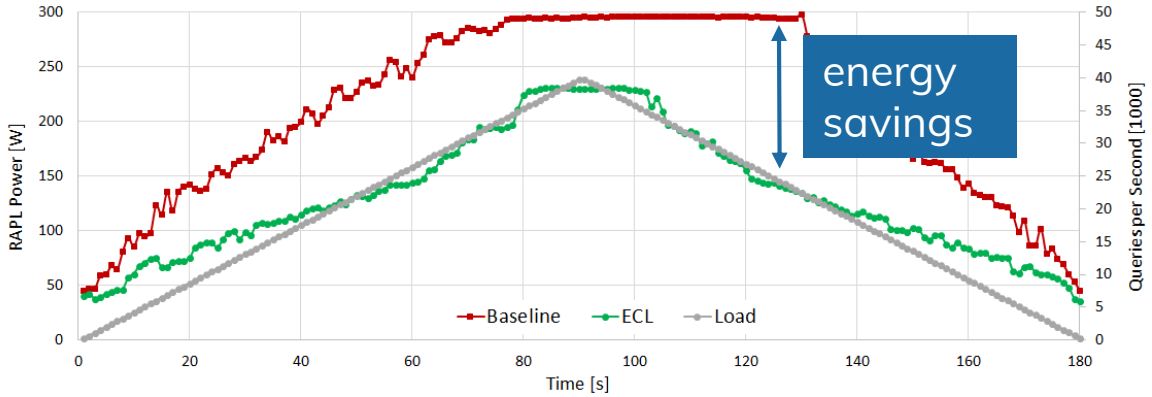

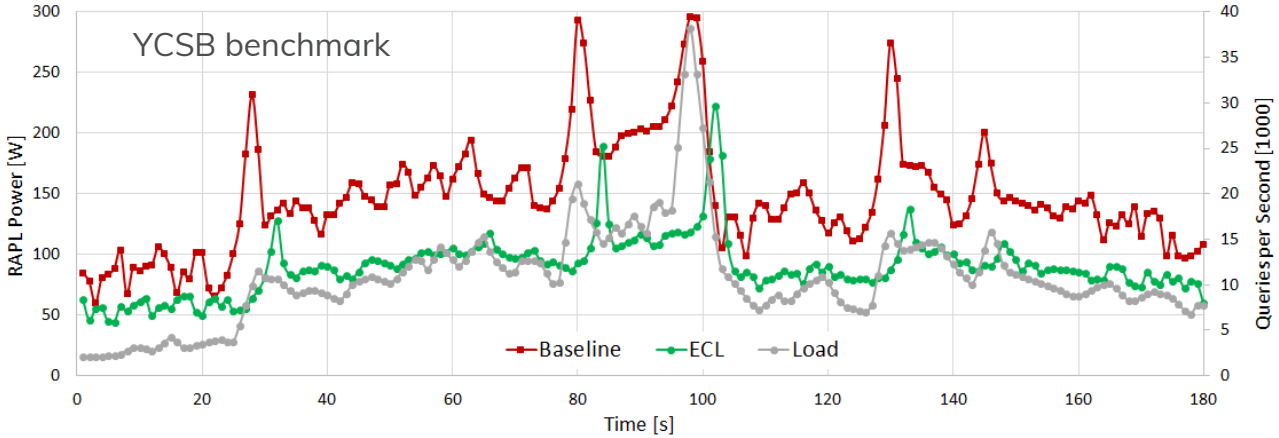

SSB + TATP

Dresden Database Systems Group

	Workload	Spike	Twitter	Most Energy-Efficient Configuration
Indexed	Key-Value *50M Pairs*	20.4 %	22.7 %	 Cores: 24@2.6GHz Uncore: 1.2GHz
	TATP *SF 200*	21.6 %	23.4 %	 Cores: 24@1.9GHz Uncore: 1.2GHz
	SSB *SF 5*	15.8-20.7 %	19.8-22.5 %	 Cores: 24@1.9GHz Uncore: Ø 1.9GHz
Non-Indexed	Key-Value *50M Pairs*	37.3 %	38.6 %	 Cores: 16@1.2GHz Uncore: 2.1GHz
	TATP *SF 200*	28.7 %	29.3 %	 Cores: 21@1.9GHz Uncore: 2.1GHz
	SSB *SF 5*	27.3-29.8 %	28.7-29.8 %	 Cores: Ø 21@1.9GHz Uncore: 2.1GHz

Summary

Energy Consumption within ICT

- More and more pressing issue!

Different Approaches

- Green Data Center
- Energy-Efficient Hardware
- Adaptive Software

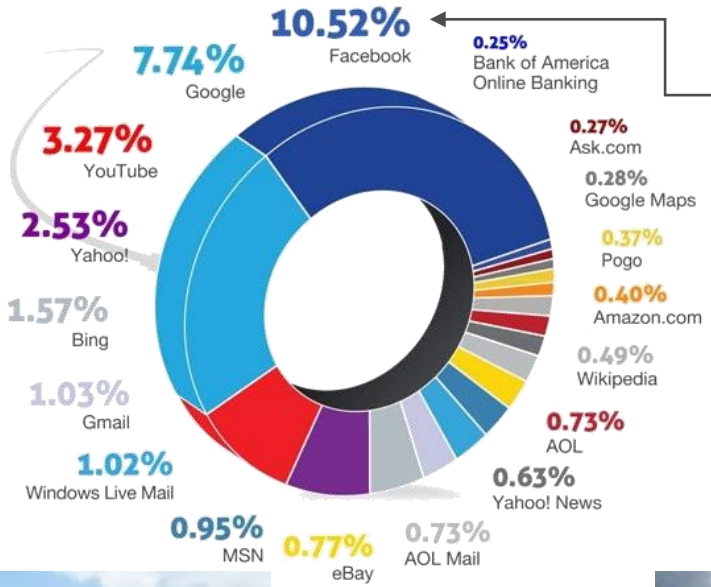

Dresden Database Systems Group

≈60.000 servers
9 data centers

each data center to power
1,730 to 4.615 homes

ERIS