TECHNISCHE UNIVERSITÄT DRESDEN

Dresden Database Systems Group

Work-Energy Profiles: General Approach and In-Memory Database Application

Annett Ungethüm, Thomas Kissinger, Dirk Habich and Wolfgang Lehner

September 5th, 2016, TPCTC

Presenter: Tomas Karnagel

Hardware Properties

Dresden Database Systems Group

System Energy consumption

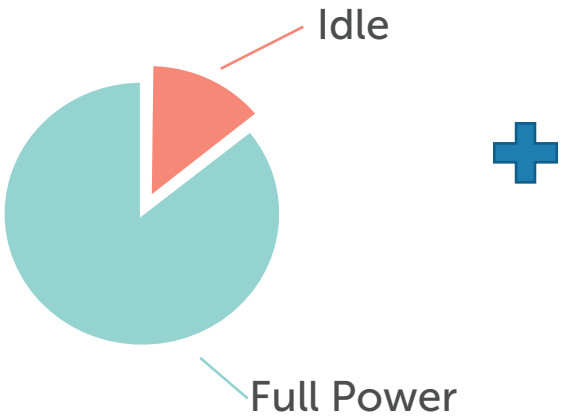

Lots of Tuning Knobs

- DVFS
- Sleep States
- Hyperthreads
- Core selection
- ...

Goal

Find Trade-off between Energy and Performance

Looking at Databases

Dresden Database Systems Group

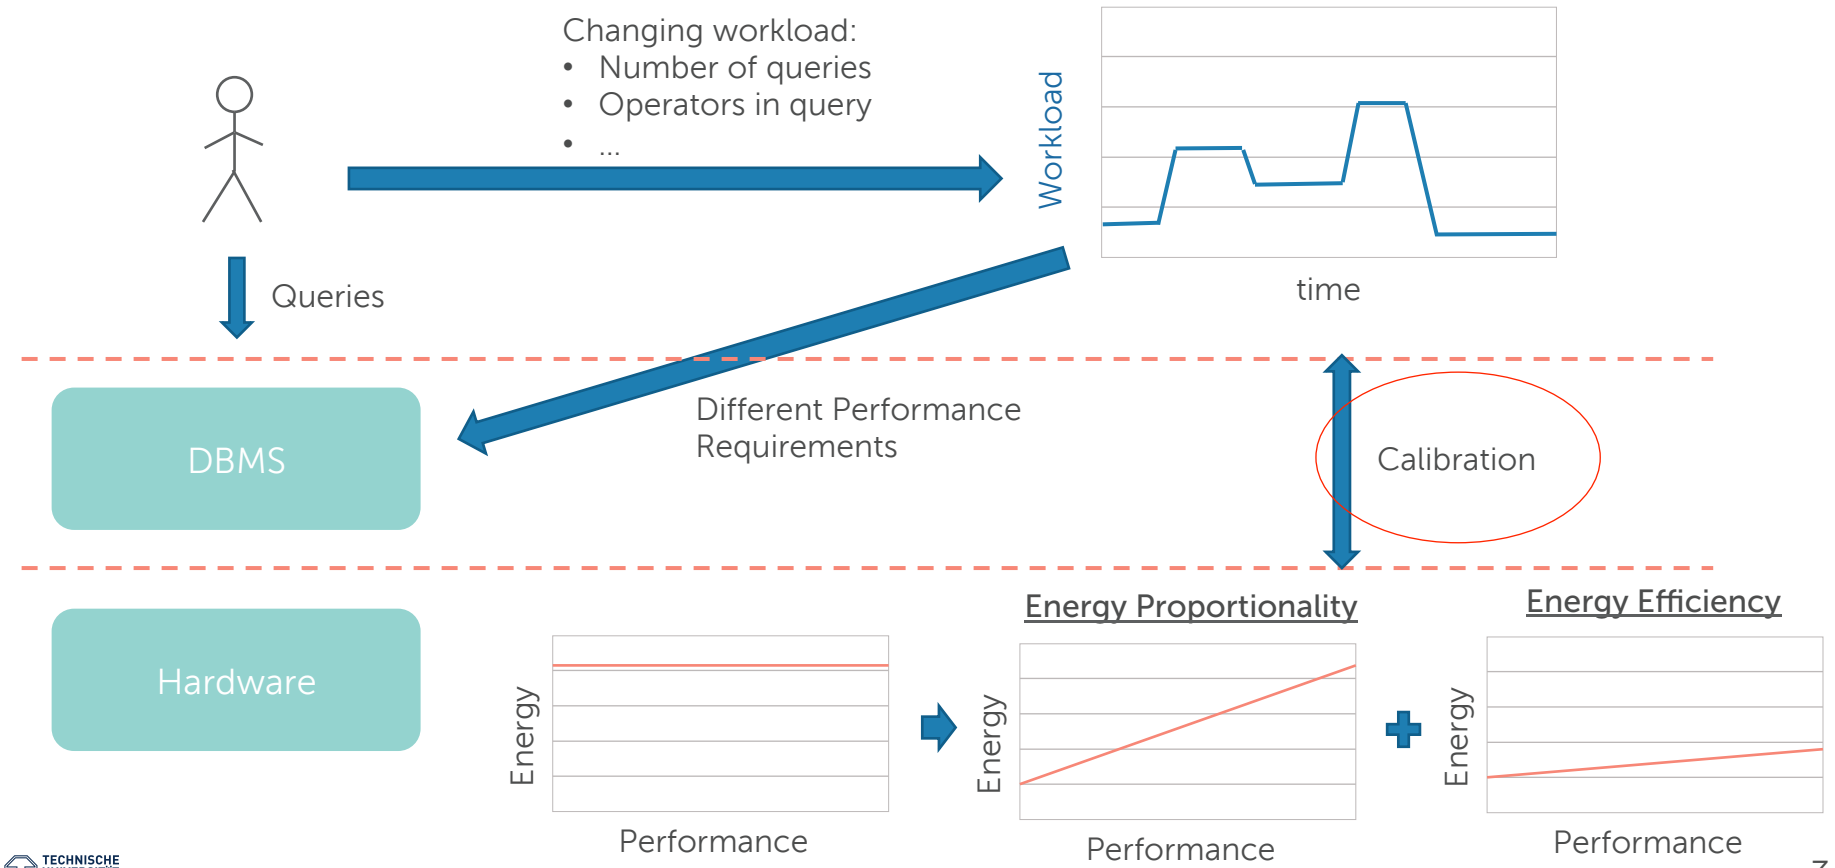

A Work-Energy Benchmark

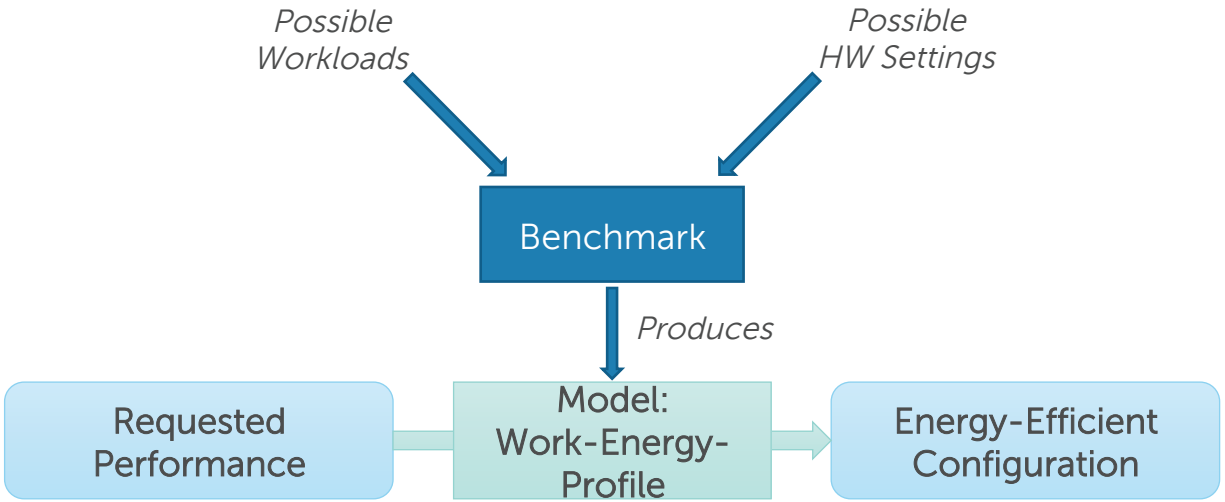

Overview: Work-Energy-Profiles

Dresden Database Systems Group

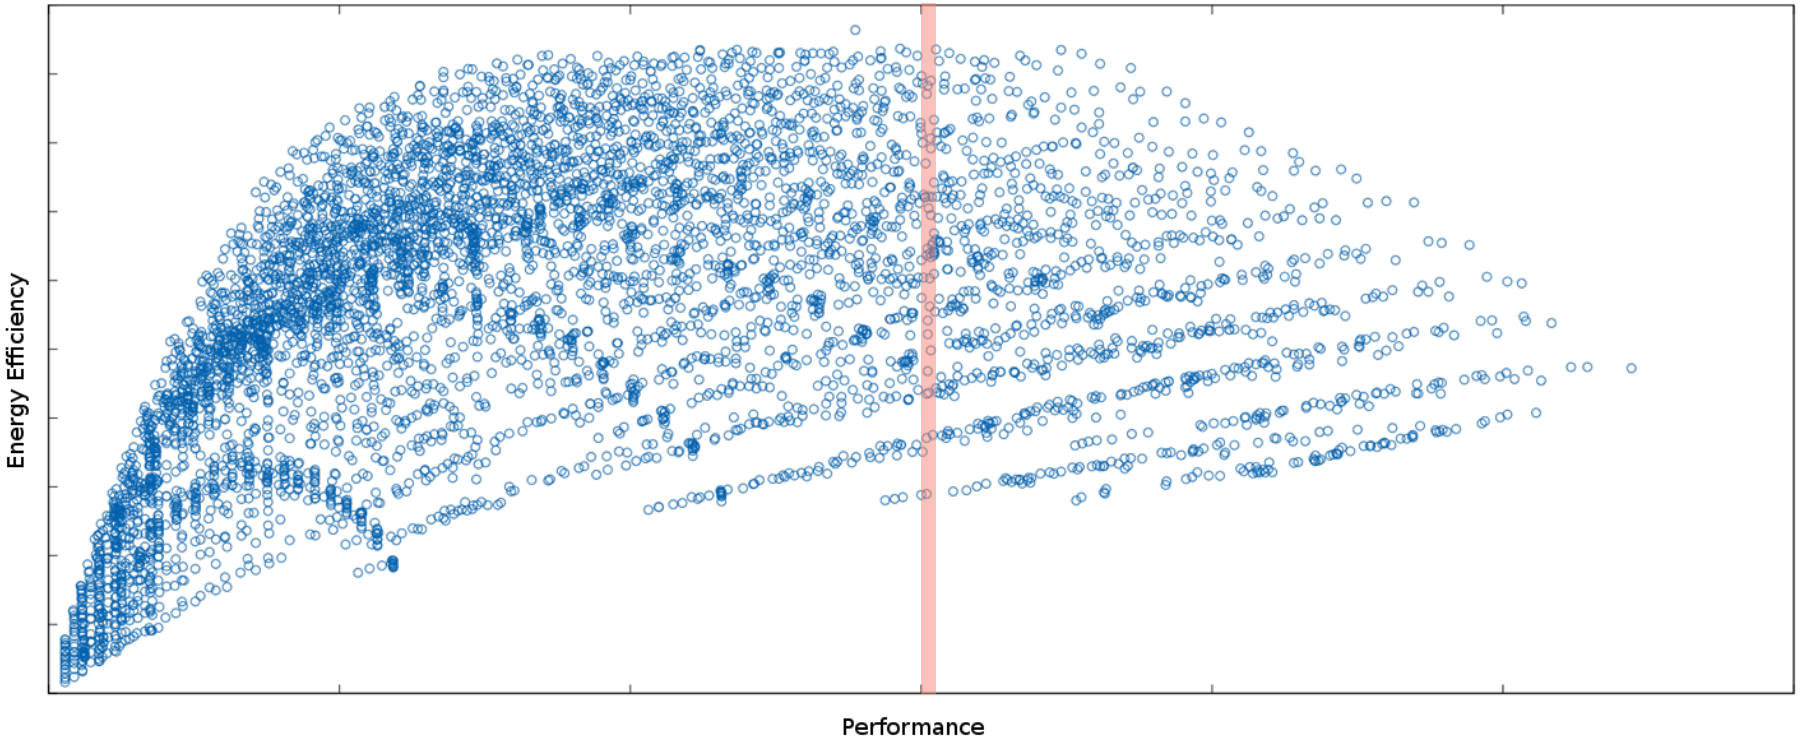

Benchmark Concept

- Concept
- Performance
- Energy Efficiency
- Hardware Configurations

Benchmark Overview

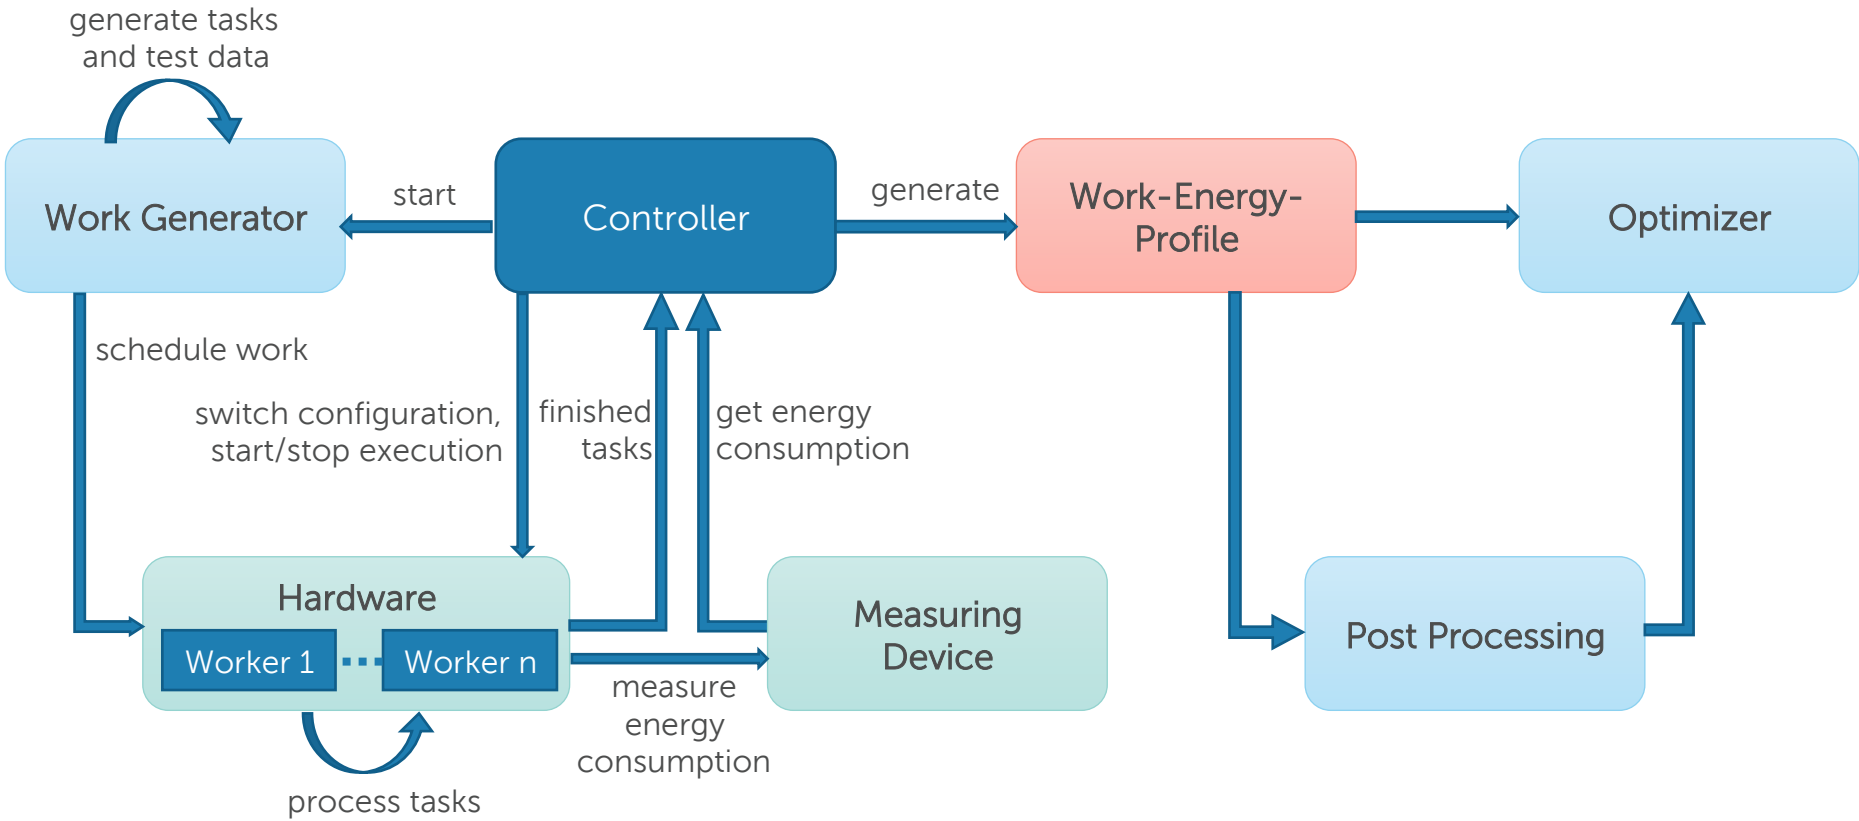

What do we need?

Dresden Database Systems Group

Work

- A task which is repeated for every configuration
- Same test data and break conditions for every configuration

Performance

Performance = work done / time

Power and Energy

Power and Energy

Dresden Database Systems Group

- Finished operations
- Power curve
- Energy

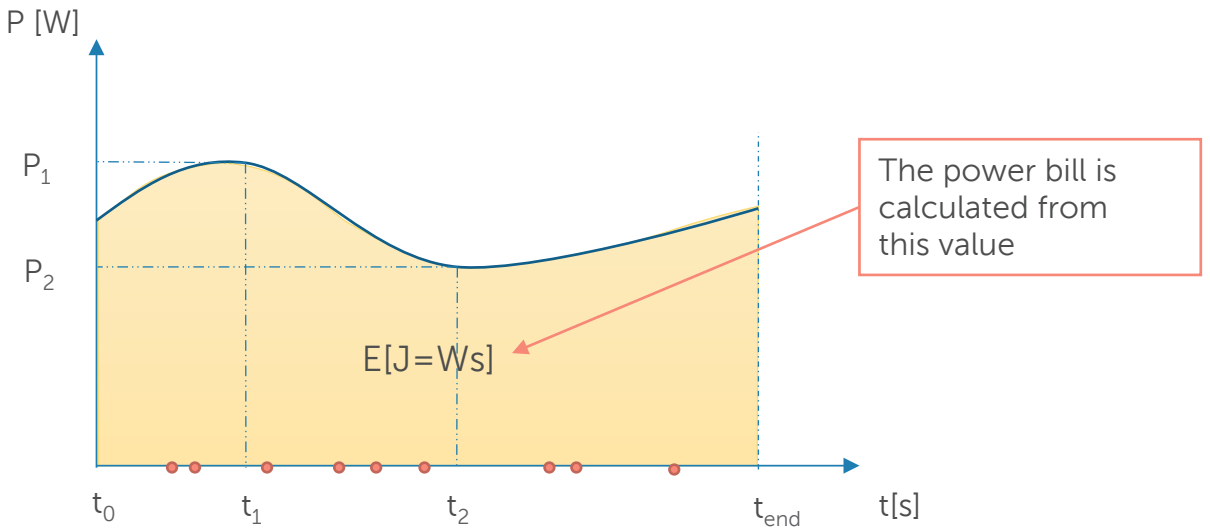

Minimize the energy, not the power

What do we need?

Work

- A task which is repeated for every configuration
- Same test data and break conditions for every configuration

Performance

Performance = work done / time

Power and Energy

$$P(t) = v(t) \cdot i(t)$$

$$E = \int_{t_0}^{t_{\text{end}}} P(t)\,\mathrm{d}t = \int_{t_0}^{t_{\text{end}}} v(t) \cdot i(t)\,\mathrm{d}t$$

Energy Efficiency

- Work-Energy Quotient (WEQ)
- WEQ = work done / energy

Hardware Configurations

Dresden Database Systems Group

Depend on hardware and operating system

Common configuration parameters:

1. Frequency of physical cores (f_i)
2. Defining active or idle workers

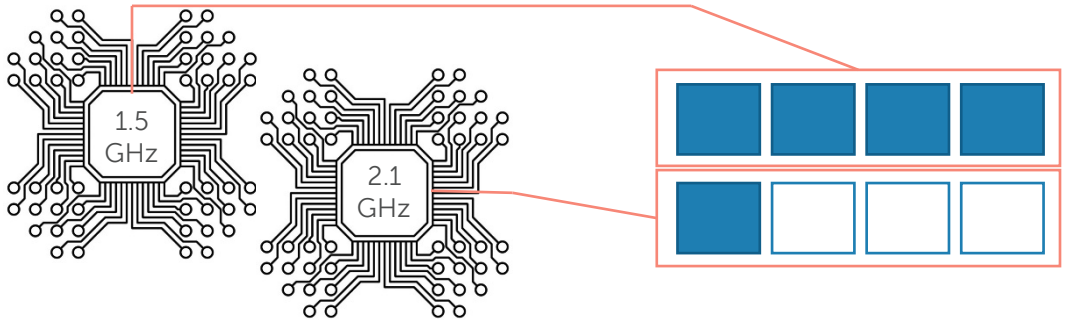

f_1=1.5 GHz, f_2=2.1 GHz

Benchmark Application

- Workload
- Hardware Setup
- Initial Energy Profiles
- Performance Effects

Typical database memory access patterns

Dresden Database Systems Group

Main memory is the bottleneck for in-memory database systems

=> Low-level work operations with significantly different memory access patterns

- E.g., compute, scan, lookup, copy

Overview: Heterogeneous Test Hardware

ODROID-XU3 - ARM® big.LITTLE. TECHNOLOGY

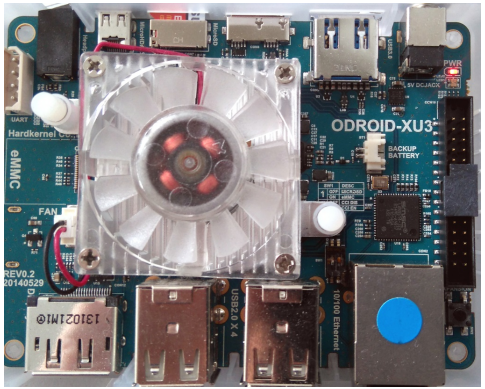

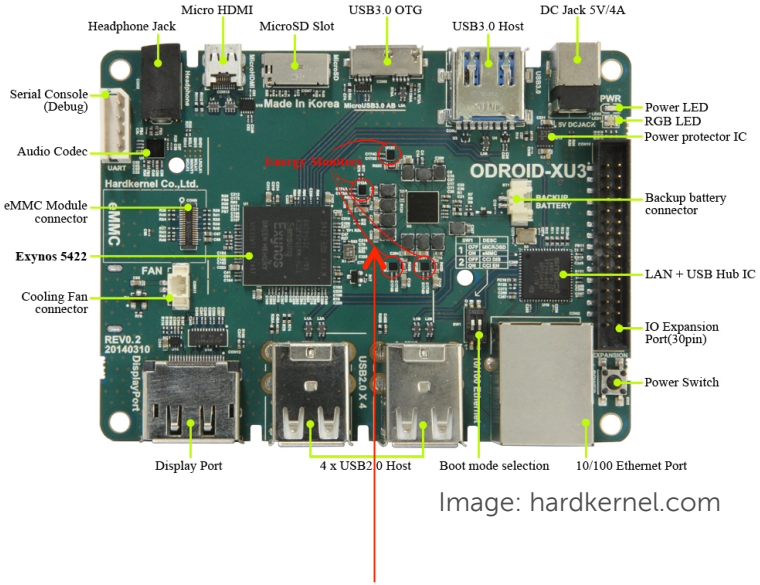

Image: hardkernel.com

	LITTLE Cluster	Big Cluster
Core Description	ARM-Cortex A7	ARM-Cortex A15
Number of Cores	4	4
Frequency Range	0.2 GHz – 1.4 GHz	0.2 GHz – 2.0 GHz
Frequency Step Range	100 MHz	
Number of Freq. Steps	13	19
Pipelines	1	3
Execution	In-order	out-of-order

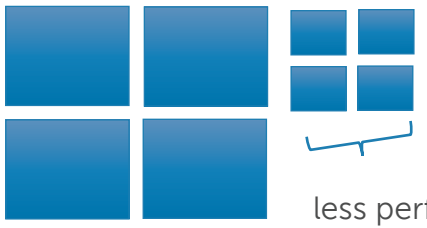

less performance
But much less power

4 Hardware Power Sensors

Configuration Search Space:
13 * 19 (frequencies) *
5 * 5 (core options)
= 6175 options

Typical database memory access patterns on the Odroid-XU3

Dresden Database Systems Group

Compute-Intensive

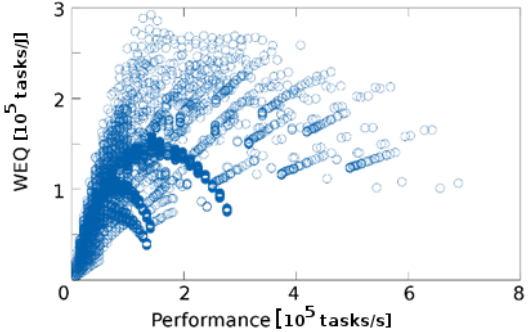

Scan-Operation

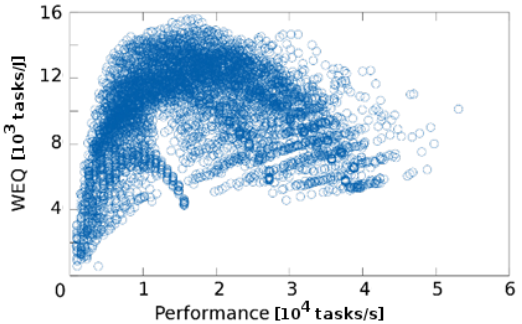

Lookup-Operation

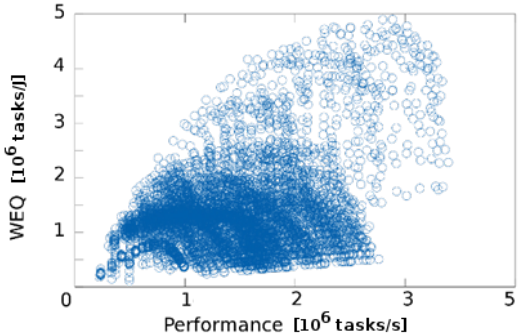

Copy-Operation

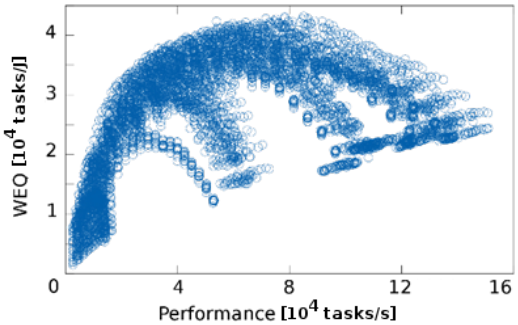

Varying frequency on active cluster

Dresden Database Systems Group

compute

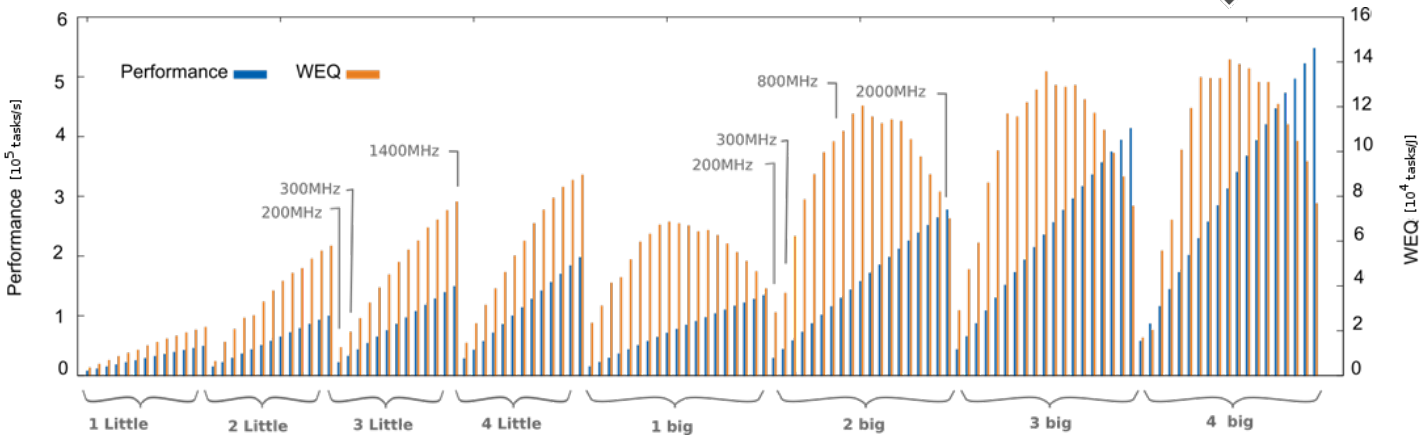

Copy

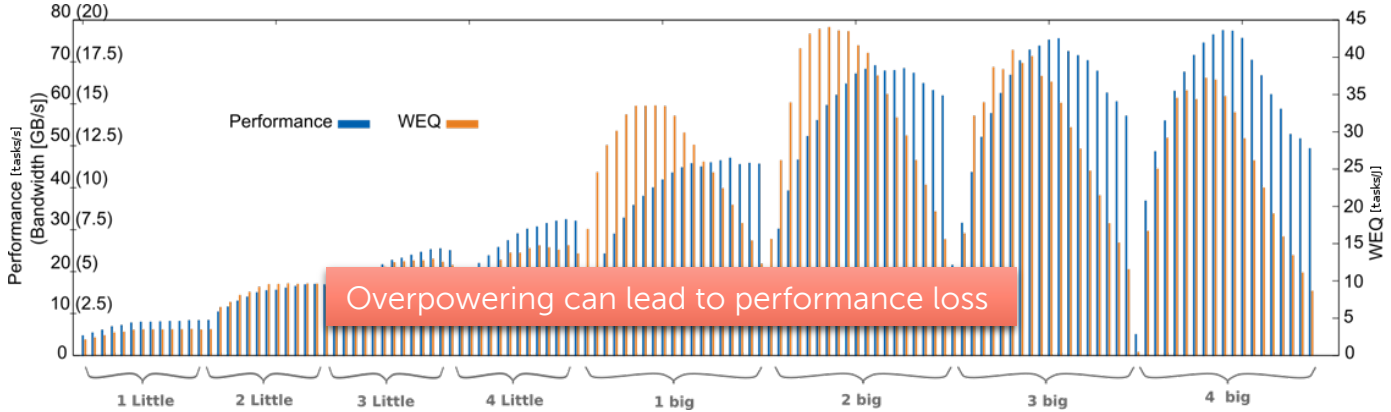

Varying frequency on idle cluster

Dresden Database Systems Group

compute

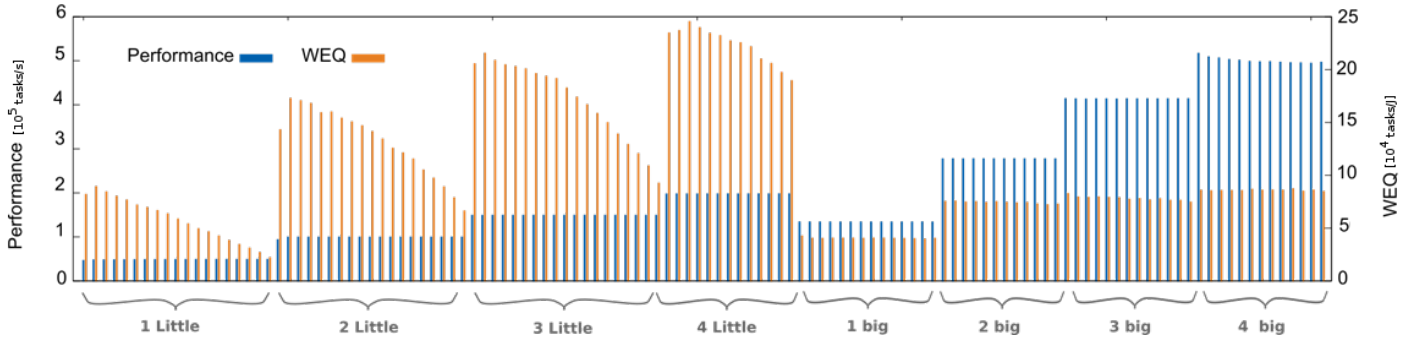

Copy

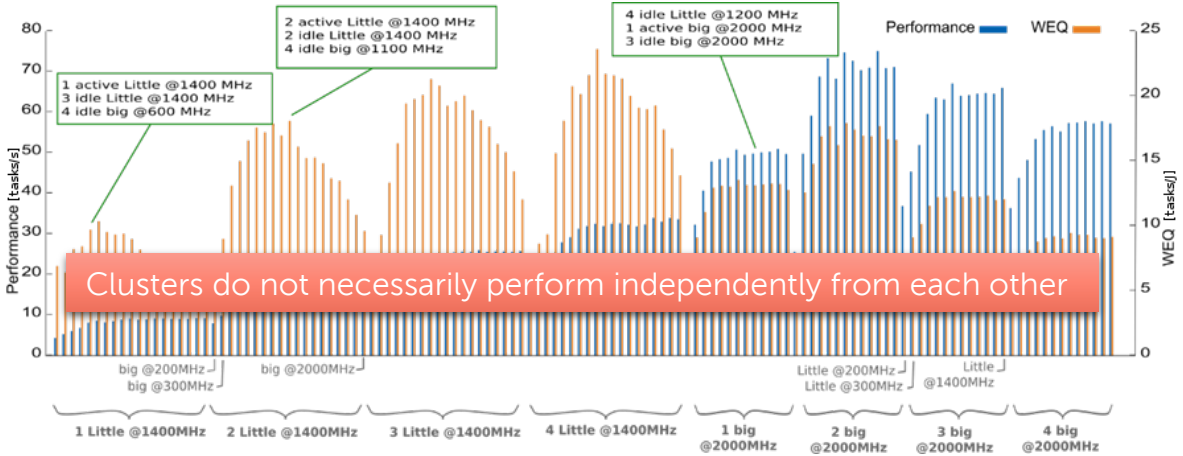

Clusters do not necessarily perform independently from each other

Optimal Configurations

Dresden Database Systems Group

Compute-Intensive

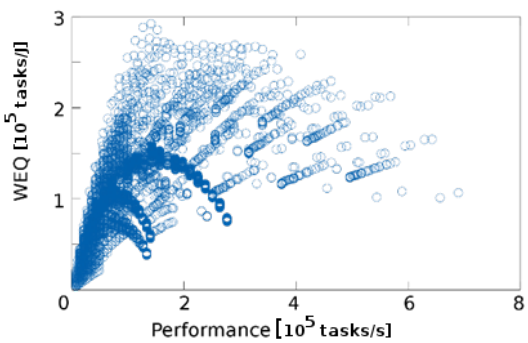

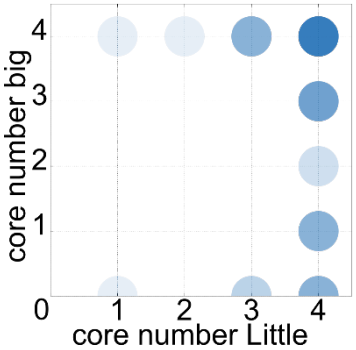

Scan-Operation

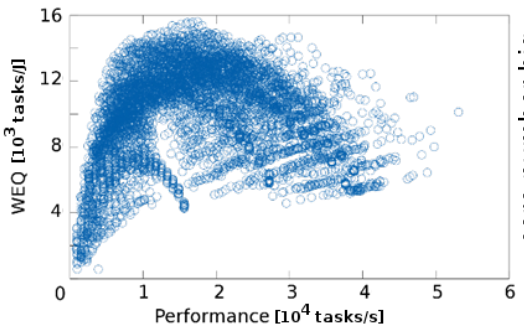

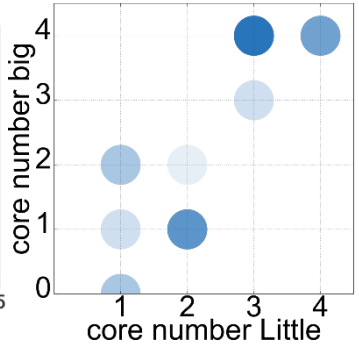

Lookup-Operation

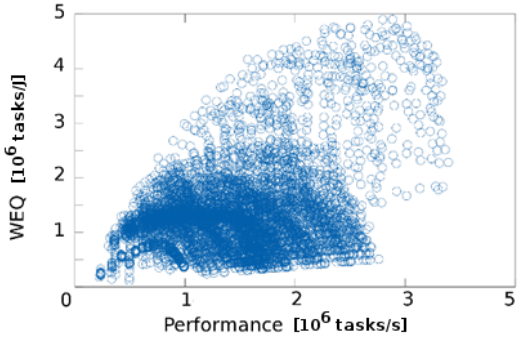

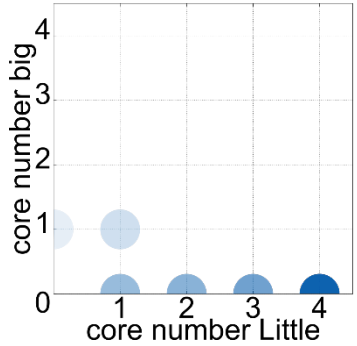

Copy-Operation

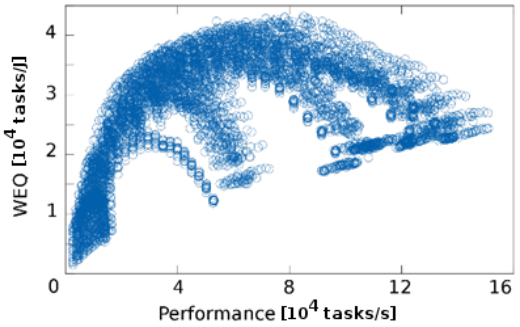

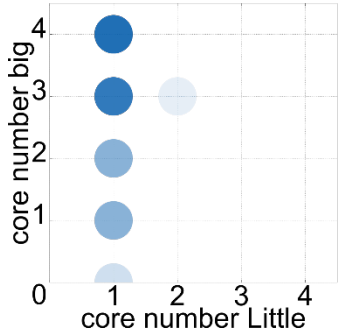

Future Work

The Energy-Control Loop

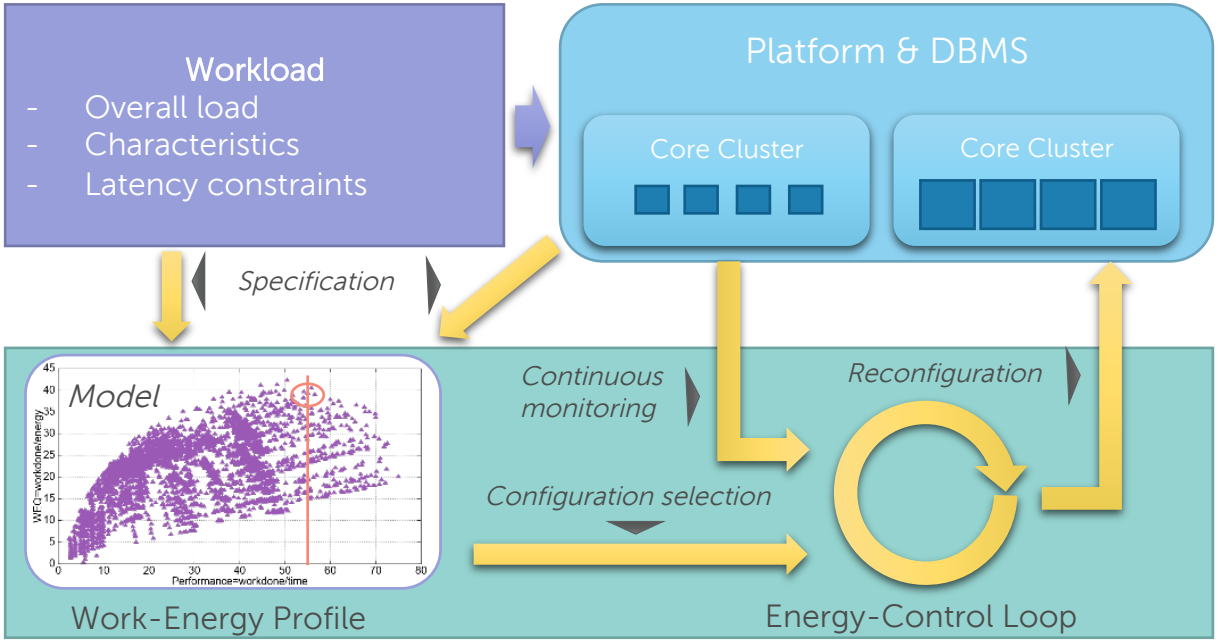

The Energy-Control Loop

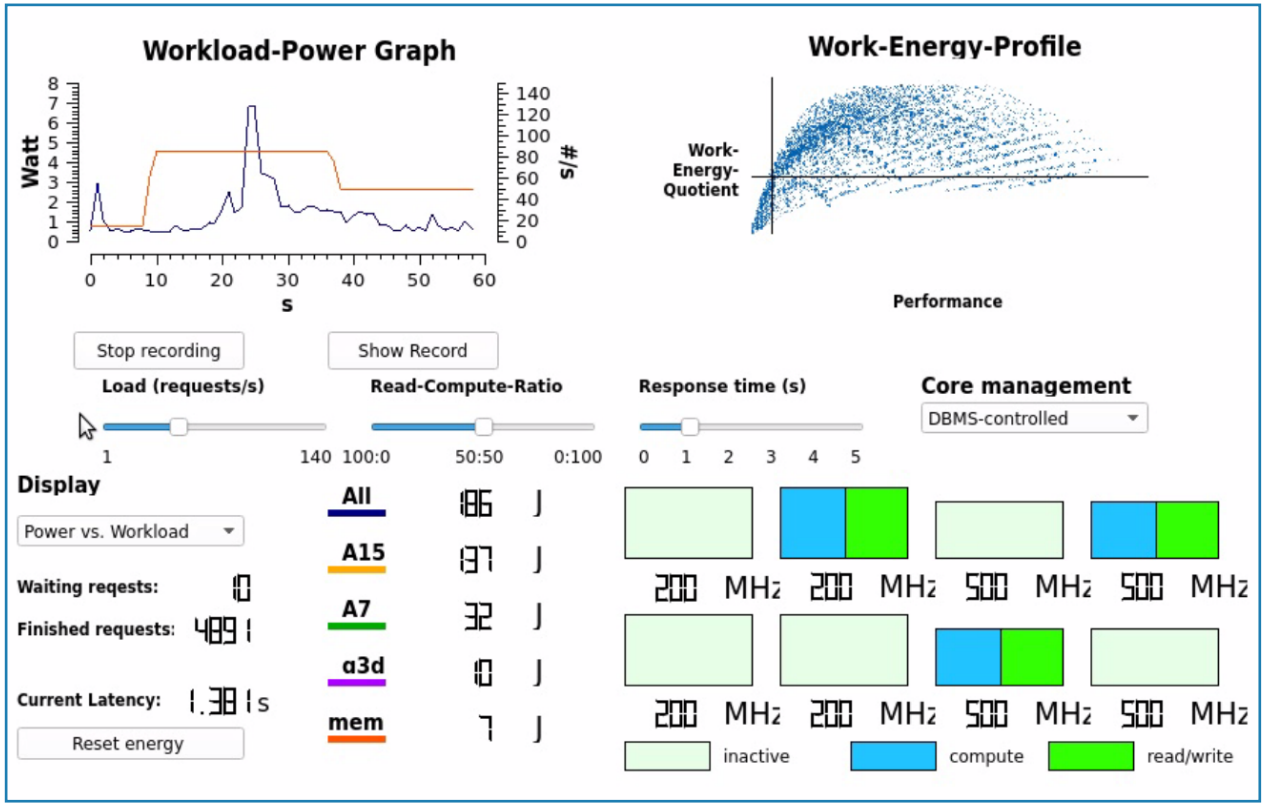

Energy Elasticity on Heterogeneous Hardware using Adaptive Resource Reconfiguration LIVE in Sigmod '16

TECHNISCHE UNIVERSITÄT DRESDEN

Dresden Database Systems Group

Work-Energy Profiles: General Approach and In-Memory Database Application

Annett Ungethüm, Thomas Kissinger, Dirk Habich and Wolfgang Lehner

September 5th, 2016, TPCTC

Follow-up questions:
annett.ungethuem@tu-dresden.de

Presenter: Tomas Karnagel