

# Energy efficiency in HPC

Sebastian Krey



# Outline

- 1 Overview
- 2 Infrastructure
  - Saving water
  - Data center temperature
  - First year of MDC operation
  - MDC Cooling optimization
- 3 HPC operation
  - Monitoring power consumption
  - IDLE nodes
  - Powercapping
  - Highspeed interconnect
- 4 Benchmarks
- 5 Summary

# Overview

- What is meant by energy efficiency?
- How can it be measured?
- Who is responsible for it?
- How can it be optimized?
- How can efforts be combined?
- Are there interactions between different efforts?

# What is energy efficiency?

- Electricity usage?
- Water usage?
- Space usage?

## How measure energy efficiency?

- What should be included?
- At which temporal scale should be measured?
- How can the data be accessed?
- Compute level measurements of processors or whole node?
- Absolute values or relative changes?
- Optimizie for energy to solution, peak consumption or power variation?

# Who is responsible?

- Data center operator?
- HPC Admins?
- HPC Users?
- All of them?
- How can they cooperate?

## How can it be optimized?

- Modern data center infrastructure
- Controller programming for cooling infrastructure
- Temperatures
- Sizing of infrastructure and IT systems
- Workflow tools
- General code optimization
- Optimizing code for specific hardware

# Interactions

- Can one optimization impair a different parameter?
- Users, admins and infrastructure operators often do not know each other
- Unified data collection from different monitoring systems often not possible
- Individual work often concentrated on the own domain.
- Static optimization or workload specific settings with automatic changes?

## Saving water

Reactions for planning a system with 32°C water supply temperature and dry air coolers in the middle of Germany:

- Insane, this does not work!
- You have guts!
- How often do you shutdown the system in summer?

Results after two summers


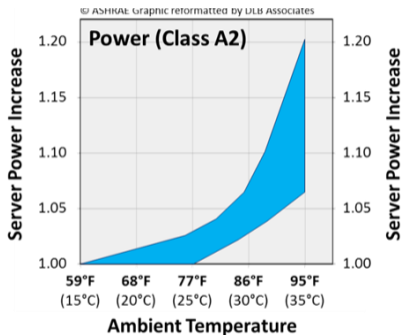
- It works!
- Water temperatures above 36°C (temporary full load limit) only few hours
  - ▶ Handling of these hours: Throttling to 60 kW per rack
- Reason: Drop in heat capture rate
  - ▶ Throttling to prevent overload on cooling for residual heat
- System shutdowns: Never

# Data center temperature

- What supply air temperatures are in use?
- What temperature limits does the equipment have?
- ASHRAE temperature limits
  - ▶ A1 up to 32°C
  - ▶ A2 up to 35°C
  - ▶ A3 up to 40°C
  - ▶ A4 up to 45°C
- ASHRAE Recommended up to 27°C
- Most servers are A2, some configurations A3
- Vendor surveys: permanent usage with 27°C no problem
- High efficiency data centers at Google 27°C, at Intel 32°C

# Data center temperature

Effect of supply air temperature on airflow requirements and power consumption



- Airflow and total power increase with temperature
- Fan power increases to the cube of the fan speed (RPM)

Source: ASHRAE\_TC0909\_Introduction\_and\_Overview\_09\_Jan\_2019.pdf

# Data center temperature

## Own experiments

- Air cooled compute nodes: up to 24°C flat fan curve, around 26°C high fan speeds → 24.5°C supply air
- DLC cooled compute nodes: up to 27°C flat fan curve → 27.5°C supply air
- Storage systems are more difficult, very rough fan control, reaction on cold aisle and hot aisle temperature changes → both temperatures have to be controlled.
- Storage supply air 24°C and return air not above 32°C

## MDC design targets

- capacity for 1.4 MW peak IT load
- space for 19 air cooled racks (11 compute, 5 infrastructure, 3 storage)
- air cooled capacity up to 300 kW
- space for 14 DLC racks (approx 78-80% heat capture rate)
- at least 1100 kW capacity for DLC systems with 32°C supply temperature (temporary increase up to 36°C possible)
- peak load operation at up to 19°C ambient, 75% at 32°C ambient
- full system PUE of at least 1.14

## MDC cooling infrastructure

- 2 dry air coolers, each providing 660 kW cooling capacity at 32°C ambient for 36°C primary supply temperature
- 3 downflow CRAC units, each has 105 kW capacity, two compressor circuits and a free cooling circuit allowing mixed mode operation
- 6 in-row coolers, each has 37.5 kW capacity, a speed controlled compressor and a free cooling circuit allowing mixed mode operation
- 3 CoolIT end of row CDUs to create 3 building blocks consisting of 4 compute racks, 2 in-row coolers and one CDU.

## First year of MDC operation

- Air cooled compute racks peak power consumption 23.5 kW per rack
- Warm water cooled racks peak power consumption 96 kW per rack
- Long term average power consumption in normal operation 80% of peak (70kW per rack)
- Typical full load power consumption 88-92% of peak (85-88kW per rack)
- Powerusers with highly optimized code constant 95-97% of peak (91-93kW per rack)
- Regular powerswings of 400 kW (start of two 256 node jobs) within less than 3 mins (up to 120 kW for air cooled systems).
- Power swings up to 600 kW happen from time to time

## First year of operation

- Full data center PUE 1.13 +0.07/- 0.03  
(includes cooling, UPS for storage and infrastructure, environmental monitoring, video surveillance, access control, fire extinguishing and alert systems, lights, cleaner sockets,  $1.13=1.045 \times 1.081$ )
- PUE for air cooled systems 1.24 + 0.15/-0.07  
(includes UPS,  $1.24=1.045 \times 1.187$ )
- PUE for DLC cooled systems 1.07 +0.02/- 0.03  
(DLC with in-row coolers for residual heat,  $1.07=1.045 \times 1.024$ )
- Average primary supply temperature 17.8°C, 33.4 hours above 32°C, 6.8 hours above 36°C (load reduction necessary)
- 98.7% of the year primary supply above 8°C, 70% above 14°C

# Observations during operation and cooling experiments

## Air cooled systems

- Large power consumption increase of air cooled systems at cold aisle temperatures  $>26^{\circ}\text{C}$  (fans)
- Very little differences at lower temperatures
- PUE plateau at primary temperatures between  $10^{\circ}\text{C}$  and  $14^{\circ}\text{C}$  under high load (up to  $15^{\circ}\text{C}$  under light load)
- PUE improvement at primary temperatures  $<10^{\circ}\text{C}$  very small (less than 0.03)
- At low loads 78% CRAC fan speed sufficient
- At high load cooling performance increases up to 87% CRAC fan speed
- Temperature oscillations due to power cycling of compressor circuits at the same temperatures

## Research questions

- How exact are the CDU cooling duty estimates and the heat capture rate calculations based on these
- Effect of the DLC supply temperature on the heat capture rate
- Effect of the rack temperature on the compute node power consumption and heat capture rate
- Interactions between rack temperature, DLC supply temperature on the heat capture rate and total power consumption
- Finding the best operating point, allowing as much free cooling as possible without increasing the compute node power consumption

## Test conditions for cooling optimization of DLC systems

- Tests performed at 89kW (at initial cooling settings 27°C/27°C) per Rack (2 exclusively reserved racks with Firestarter and fixed clock speeds), measured at PDU
- Response surface split plot experimental design with additional star points for rack and secondary fluid circle supply temperatures (rack temperature is difficult to change parameter)
- Temperature range: DLC supply 22-32°C (star at 18°C, 34°C and 36°C) rack 24-30°C (star at 22°C, 32°C skipped as 30°C already too high for sensible operation)
- Heat capture rate of DLC based on CDU cooling duty estimates
- Heat capture rate of in-row cooler based on primary flowrate and dT through measured at balancing valve
- 10 to 20 min. waiting time before measurements to allow stabilization of values

# Observations during operation and cooling experiments

## DLC systems

- Reduction of DLC (secondary fluid circle) temperature increases heat capture rate (up to 90% at 18°C)
- CDU controller oscillations at  $dT > 8K$  between primary and secondary circuit
- $dT \leq 4K$  between primary and secondary supply temperature results in 100% CDU valve opening most of the time
- DLC return temperatures  $< 40^\circ C$  result in drop of CDU heat exchanger efficiency  $\rightarrow$  lower primary  $dT \rightarrow$  reduced dry air cooler efficiency
- No cooling gain by increasing the in-row fan speed  $> 90\%$
- Less than 65% in-row fan speed increases compute node temps
- Full free cooling (in-row) needs  $dT = 9-12K$  between rack and primary supply
- Compute chassis fan control has an exponential curve, rack temp  $> 28^\circ C$  large fan speed increase
- Secondary supply temperatures  $> 34^\circ C$  trigger high fan speeds

# Logical consequences

## Primary circuit and air cooled systems

- Dry air coolers have now 14°C minimum allowing power saving for more than 30% of the year.
- Each CRAC unit has individual start temperatures for its compressor circuits and zone control, resulting in a 7 step power control
- Cold aisle target temperature increased to 23.5°C (highest possible temperature to ensure cold aisle  $\leq 26^\circ\text{C}$  without unnecessary compressor cycling)
- Dynamic CRAC fan speed control from 78% up to 87%

# Logical consequences

## DLC systems

- Rack temperature target for DLC systems increased to 27.5°C
- Fan speed limits for in-row cooler fans now 65% and 90%
- Fine tuning of hydraulic balancing under partial load to achieve same free cooling capacity for all in-row coolers under all load conditions  
→ full free cooling at up to 20°C primary supply (5K higher than before)
- Dynamic control of DLC supply temperature following the primary fluid circle temperature with  $dT=5K$ , minimal temperature of 22°C and maximum of 34°C (resulting in 79-85% heat capture rate)
- Heavy tuning of CDU PID parameters to allow stable control with primary to secondary  $dT>8K$  (standard parameter range was not sufficient so cooperation with Luc from CoolIT to provide patched firmware)

## Monitoring power consumption

- IPMI sensors (often only PSU)
- IPMI DCMI (not on all systems available)
- Intel Node Manager (only via FreeIPMI and Intel Data Center Manager)
- Metered PDUs
- Official metering from infrastructure group (often difficult to integrate in HPC system monitoring)
- RAPL (not always interpretable results)
- Integration in Slurm to provide users Wh per job (easy integration for DCMI)

## IDLE nodes

Easiest power save method is shutdown of IDLE nodes:

Emmy: 10 air cooled racks, 11 water cooled racks

- One air cooled rack of IDLE nodes: 1.9% powersave
- One water cooled rack of IDLE nodes: 7.4% powersave
- Power save potential:
  - 80% system load : about 5.5%
  - 60% system load : about 15%

Enabling C-states C1E and lower for powered IDLE nodes has a major effect on IDLE power consumption (up to 40%).

C1E and lower effect memory and network latencies → Disable in job prolog.

# Powercapping

- Emergency plans for long term power supply problems
- How much electricity is needed for an operational ready system without compute (users can login, access their data and compute could be started)
- Powersave potential of different performance/power limit methods:
  - ▶ Turbo off
  - ▶ Frequency limit
  - ▶ Custom TDP
  - ▶ Platform based power capping
- How do these different methods effect different workloads (compute, memory, communication intensive)
- TCO optimized procurements in the last years limit possible perf/watt gains

# Minimal Powerconsumption

Minimal energy for an operational ready cluster system:

Can be reached within a few minutes in case of power supply problems.

- CARO:**
- Login-, admin nodes, network, BMC: 54 kW
  - Storage: 10 kW
  - plus cooling

**Emmy:** Total 120 kW

- Computeracks (edge switches, BMCs): 31 kW
- Service nodes and director switches 22 kW
- Storage 25 kW
- Cooling 42 kW

## Turbo off/Frequency limit

### Turbo off

- Very easy to implement
- Performance degradation workload dependent
- For Emmy power consumption reduction around 11% (8% air cooled, 14% water cooled nodes)
- Affects also the high speed interconnect

### Frequency limit

- Above base clock only possible for Intel CPUs
- Fine granular adjustment (200MHz)
- Less throttling for AVX2 and AVX512 codes
- Good balancing of performance variation between nodes
- For Emmy from 1% at 3.0GHz to 11% at 2.4GHz power consumption reduction

## Custom TDP

- Linear power/perf relation up to certain point, then stronger power increase
- Power capping can move performance optimized CPUs into linear range
- Custom TDP provides a few (2-3) vendor defined levels
- Custom TDP for Intel Skylake and Cascadelake CPUs only effective for Non-Turbo operation
- Affects only CPU, not mainboard, network, storage etc. → Strong limit on CPU reduces performance per watt
- For CARO (AMD Rome 7702) 165 W and 200 W result in 5.0 GFlops/W HPL
- Emmy: combination of Turbo off and cTDP reduces perf/watt by 8-15% in Gromacs

## Power capping

Intel Node Manager allows global or event based platform power limiting. Similar features for AMD based systems exist probably but are unknown to me.

- Power limit affects all components of system
- Fine granular limiting possible
- Control via BIOS, FreeIPMI or Intel Data Center Manager
- Main control based on dynamic frequency limits for the CPU
- AVX2 and AVX512 clocks (lower) will also be limited
- Problem: Same power limits, different CPU clocks (production variance)
- HPC requires same performance on all nodes to minimize load imbalances
- Sometimes hiccups of the controller resulting in bad performance

## Highspeed interconnect

Reduction on CPU clock has a direct effect on the number of interrupts, that can be processed, which results in performance degradaton of the highspeed interconnect.

- Omni-Path:**
- Very sensitive to changes in CPU clock
  - Turbo off increases latencies by 10%
  - IPoIB high interrupt load, Turbo off reduces bandwidth by 20%
  - RDMA bandwidth less affected (1-8%)
  - Fine granular limiting of CPU clock beneficial for communication sensitive workloads
- Infiniband:**
- Less CPU sensitive, Benchmarks are still mssing
  - IPoIB probably similar results to Omni-Path

# Benchmarks

Ongoing benchmarks with different tools from our benchmark suite

- Gromacs (highly optimized compute intensive)
- OpenFOAM (memory bandwidth intensive)
- BQCD (communication latency sensitive)

## Results: Gromacs

- Turbo off**
- Air cooled nodes: No performance impact, 3% power save
  - Water cooled nodes: 3% performance loss, 8% power save
  - Multinode jobs: no performance loss, 11% power save

- Power cap.**
- Air cooled nodes: performance/watt increase for 500W and 450W limit
  - Water cooled nodes: performance/watt increase for 800W, 750W reduces perf/watt

**cTDP/Turbo off** Not beneficial, up to 15% perf/watt loss

## Results

- OpenFOAM** Performance behaviour with Turbo off similar to Gromacs, but even more power save (18%)
- BQCD** Very latency sensitive but better performance with Turbo off due to more balanced performance of compute nodes outweigh the higher network latency performance loss due to reduced clock rate
- FESOM** Measureable performance decrease with Turbo off, but smaller than powersave so better perf/watts

Software which is not using all cores can have up to 20% performance loss with Turbo boost due to lower clock rate, so we allow to enable Turbo boost on a small subset of our nodes on demand.

# GPU Results

- Until now only HPL measurements
- Grete with unthrottled A100 69.12 kW for 1.83 PFlop/s
- Grete efficiency optimized with throttling to 1020 MHz 48.71 kW for 1.56 PFlop/s

# Summary

- Energy efficiency is a complex topic
- For good results different domains have to be integrated
- Interactions between domains have to be considered
- Contact between domain specialists have to be established
- Common language and understanding is critical