

Heat-aware Loadbalancing - Is it a thing?

Lukas Iffländer, **Norbert Schmitt**, Andreas Knapp, and Samuel Kounev

Chair of Software Engineering – Universität Würzburg


November 13, 2020

11th Symposium on Software Performance 2020, 12.–13. November 2020, Leipzig, Germany

<https://se.informatik.uni-wuerzburg.de/>

Motivation


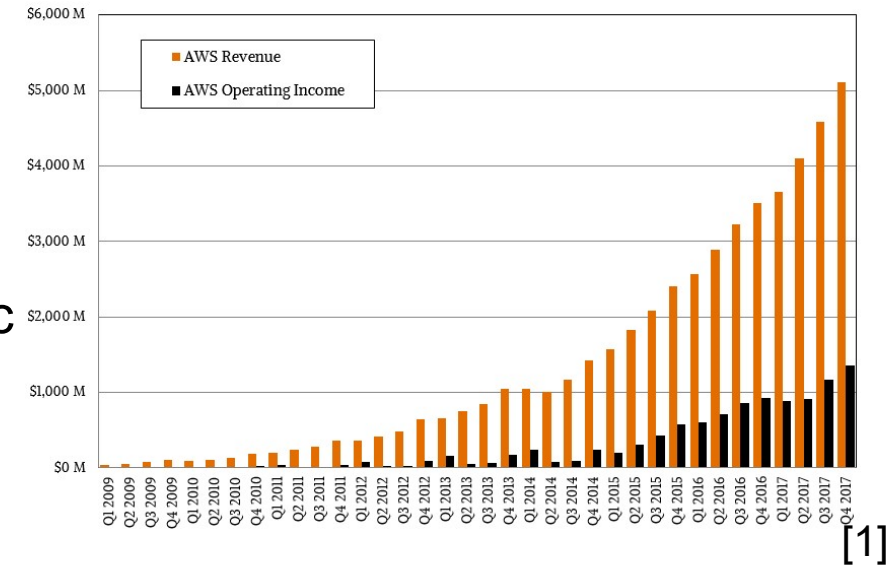
- Increasing amount of cloud users and services



[1]

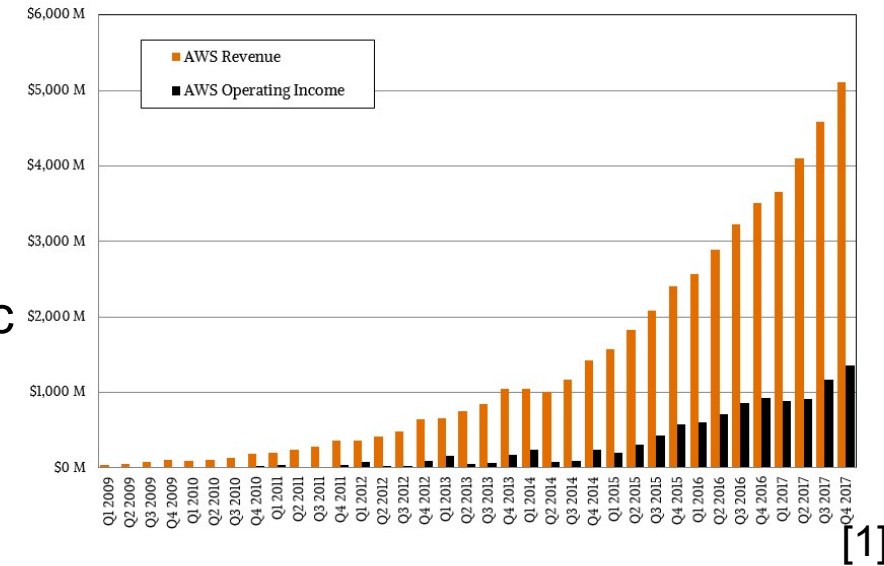
Motivation

- Increasing amount of cloud users and services
- Increasing requirements on computing power by more traffic
- ➔ Increasing computing power needed




Motivation

- Increasing amount of cloud users and services
- Increasing requirements on computing power by more traffic
- ➔ Increasing computing power needed




- Solution 1: More servers? ⚡




Motivation

- Increasing amount of cloud users and services
- Increasing requirements on computing power by more traffic
- ➔ Increasing computing power needed



- Solution 1: More servers? ⚡
- Solution 2: Better/Full usage of existing computing power ✓
 - ➔ make use of unallocated CPU boost power



Goals

Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion

1. Is it possible to exceed permanently and feasibly the clock rate with host relay?



Goals

Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion

1. Is it possible to exceed permanently and feasibly the clock rate with host relay?



2. What is the effect on the computing power compared to normal behavior?

Goals

Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion

1. Is it possible to exceed permanently and feasibly the clock rate with host relay?



2. What is the effect on the computing power compared to normal behavior?

3. What is the effect on the energy consumption compared to normal behavior?



Contributions

Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion

1. A first approach to heat-aware load-balancing

Contributions

Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion

1. A first approach to heat-aware load-balancing

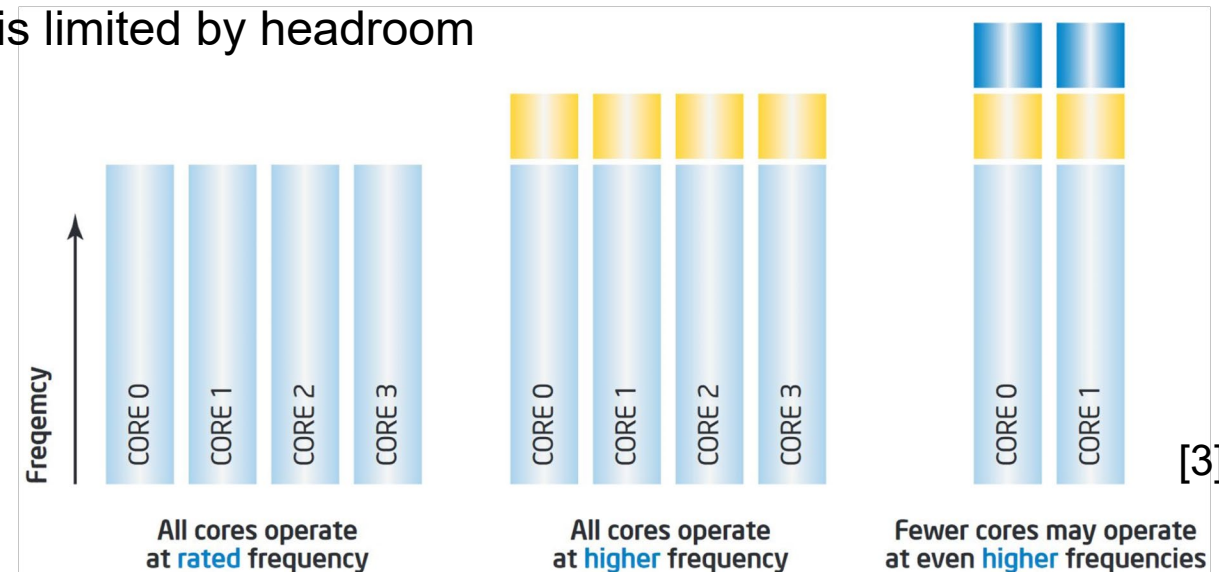
2. A prototype implementation with SDN and an initial evaluation

Approach - Boosting

Foundations ➤ Reasearch Question ➤ Approach ➤ Evaluation ➤ Conclusion

➤ Intel Turbo Boost:

- CPU-Boost ≠ Overclocking in general
- CPU Boost respects specifications of hardware
- Max. n° of boosted cores and max. frequency is limited by headroom
- OS requests more computing power → Boost
- Example Intel:
 - Stepwise increasing clock with 133,33 MHz
 - Time in boost state depends on:
 - Temperature
 - Energy



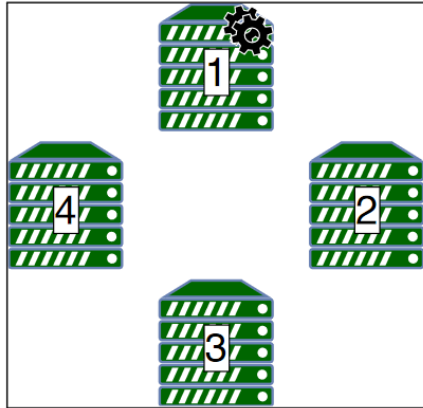
➤ Other common used boost: AMD Turbo Core / AMD Precision Boost

➤ Boosts are usually disabled in related work and HPC environments due to unpredictability

Motivation ➤ Goals & Contributions ➤ Approach ➤ Evaluation ➤ Conclusion

Approach

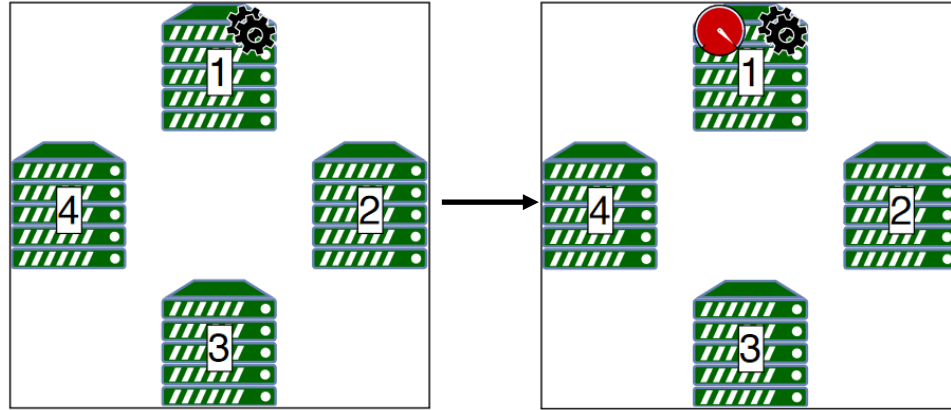
Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach


Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach


Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach

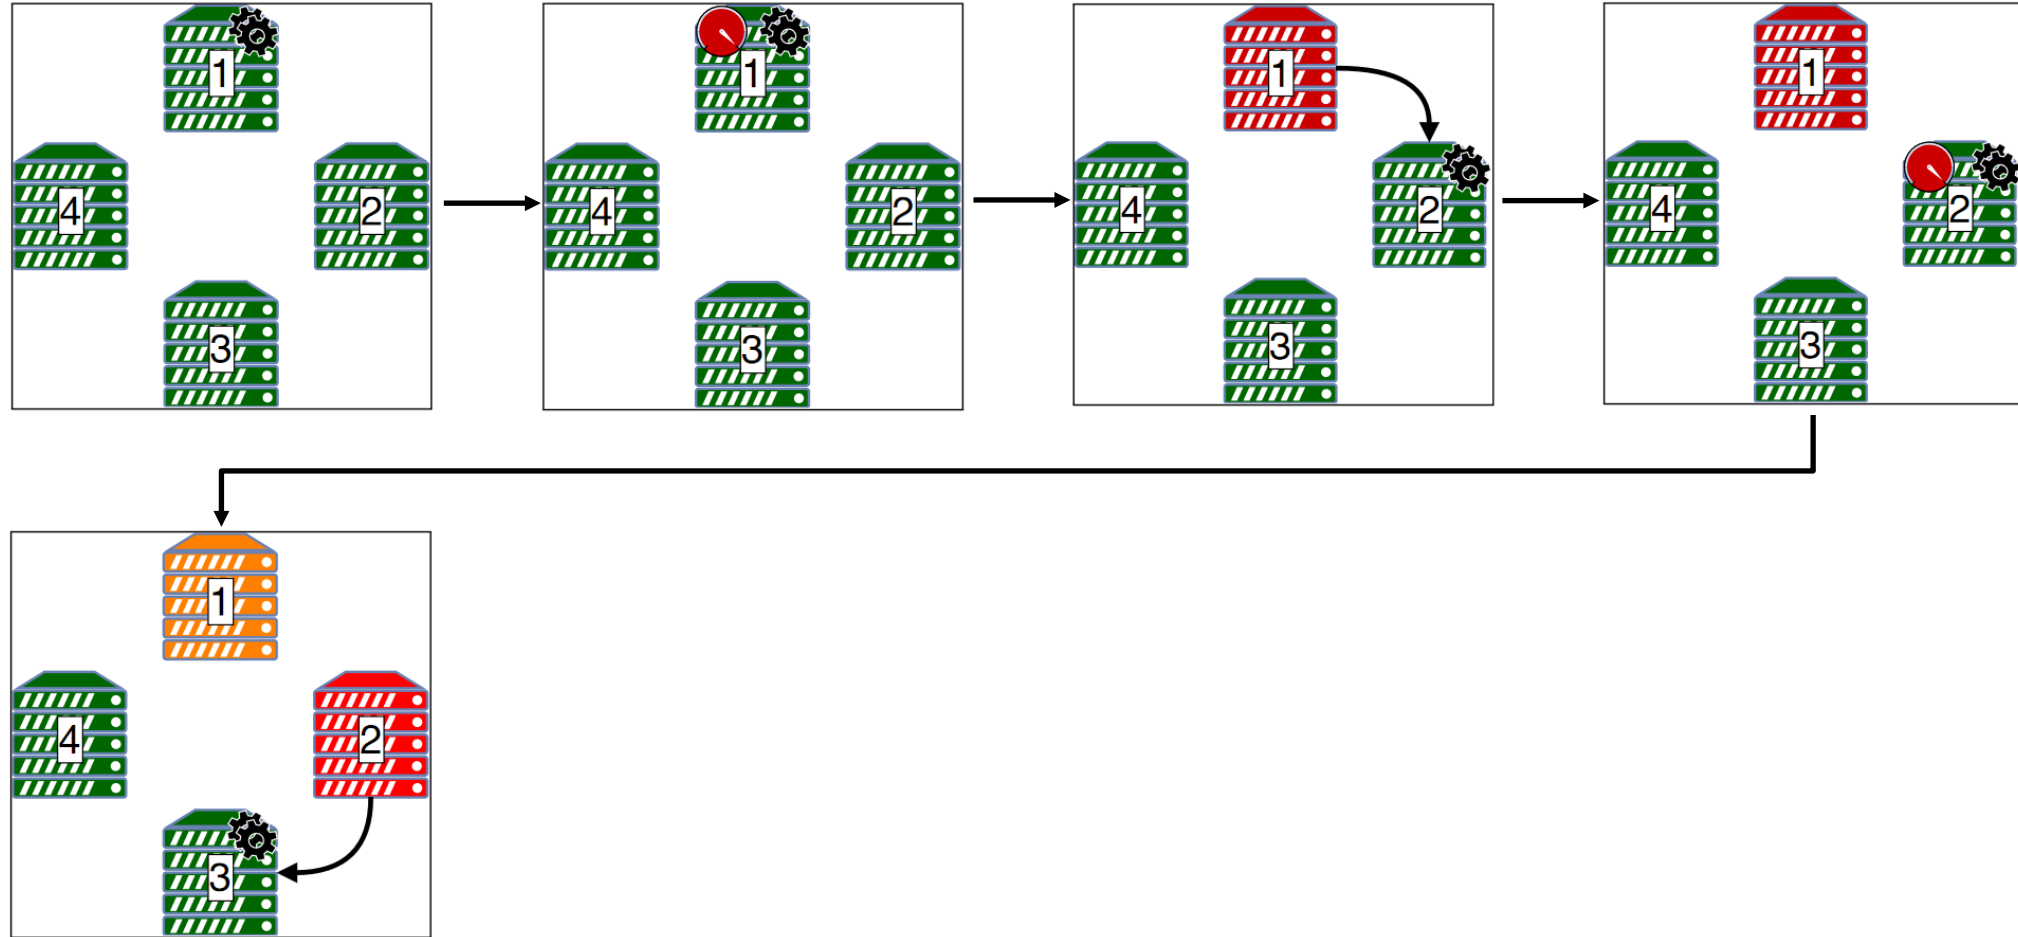
Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach


Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach


Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach

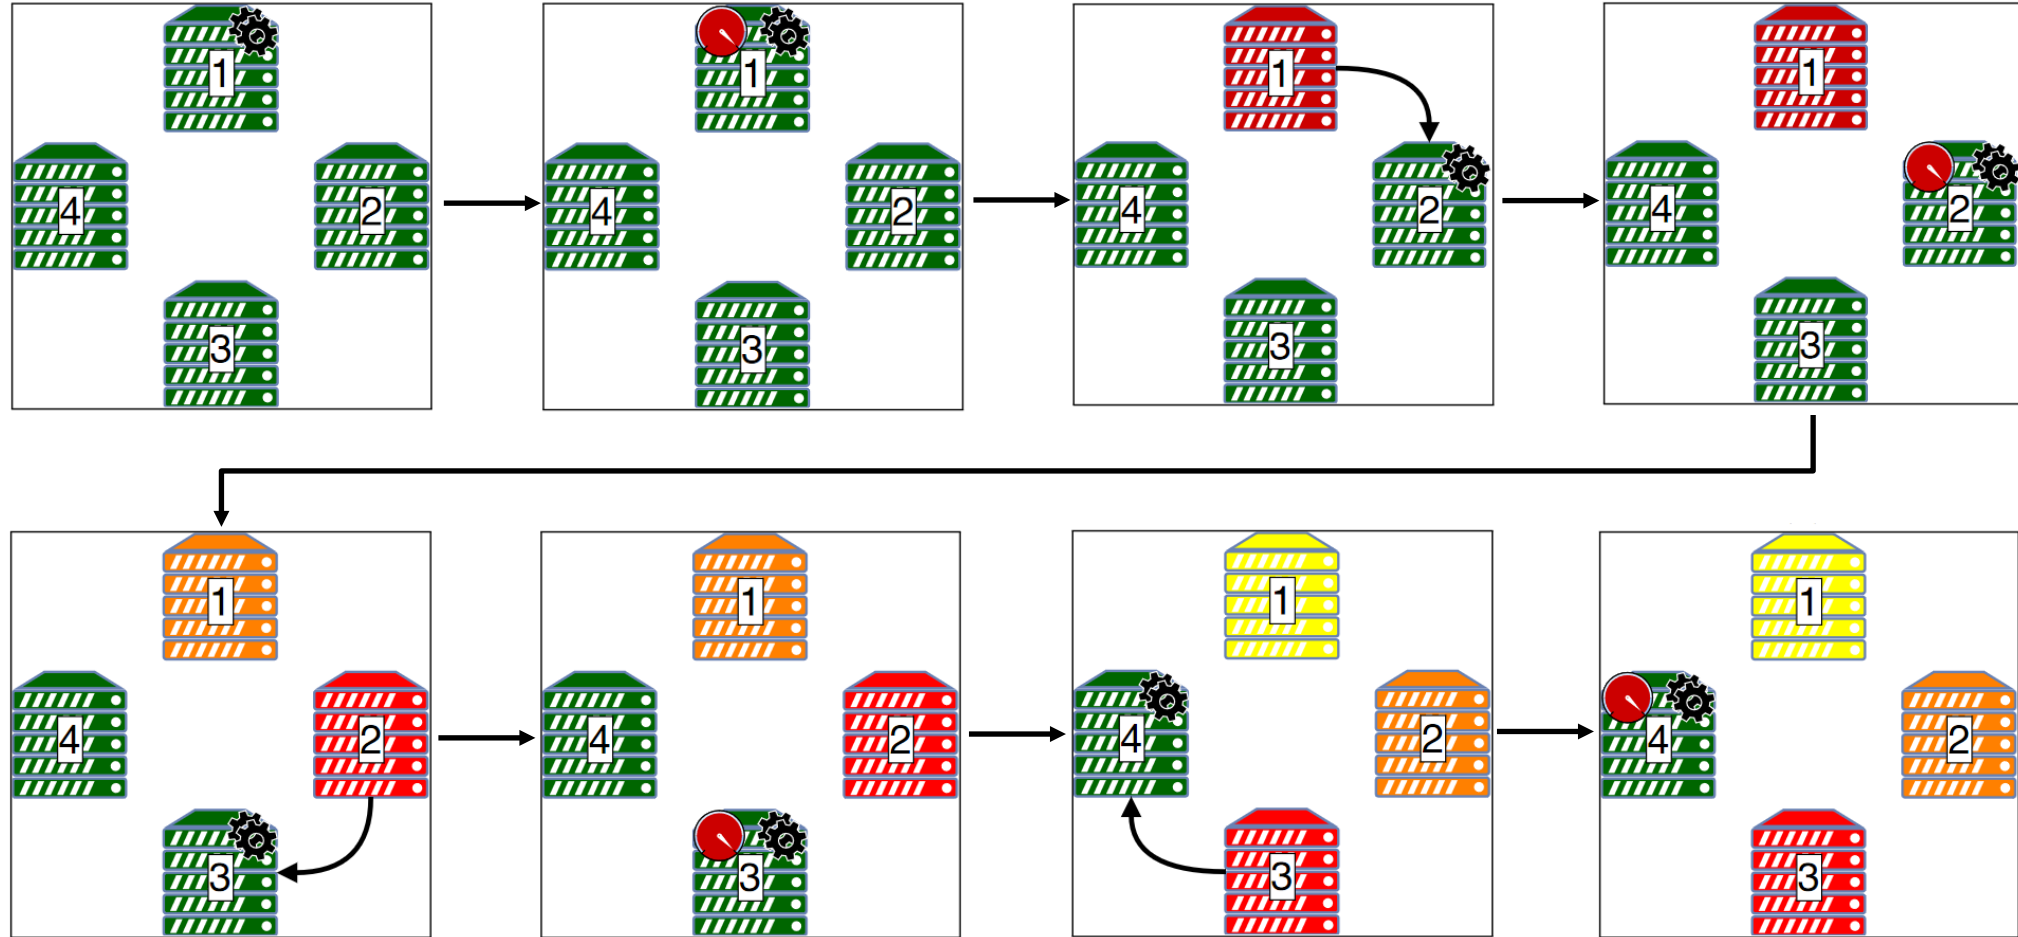
Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach


Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach

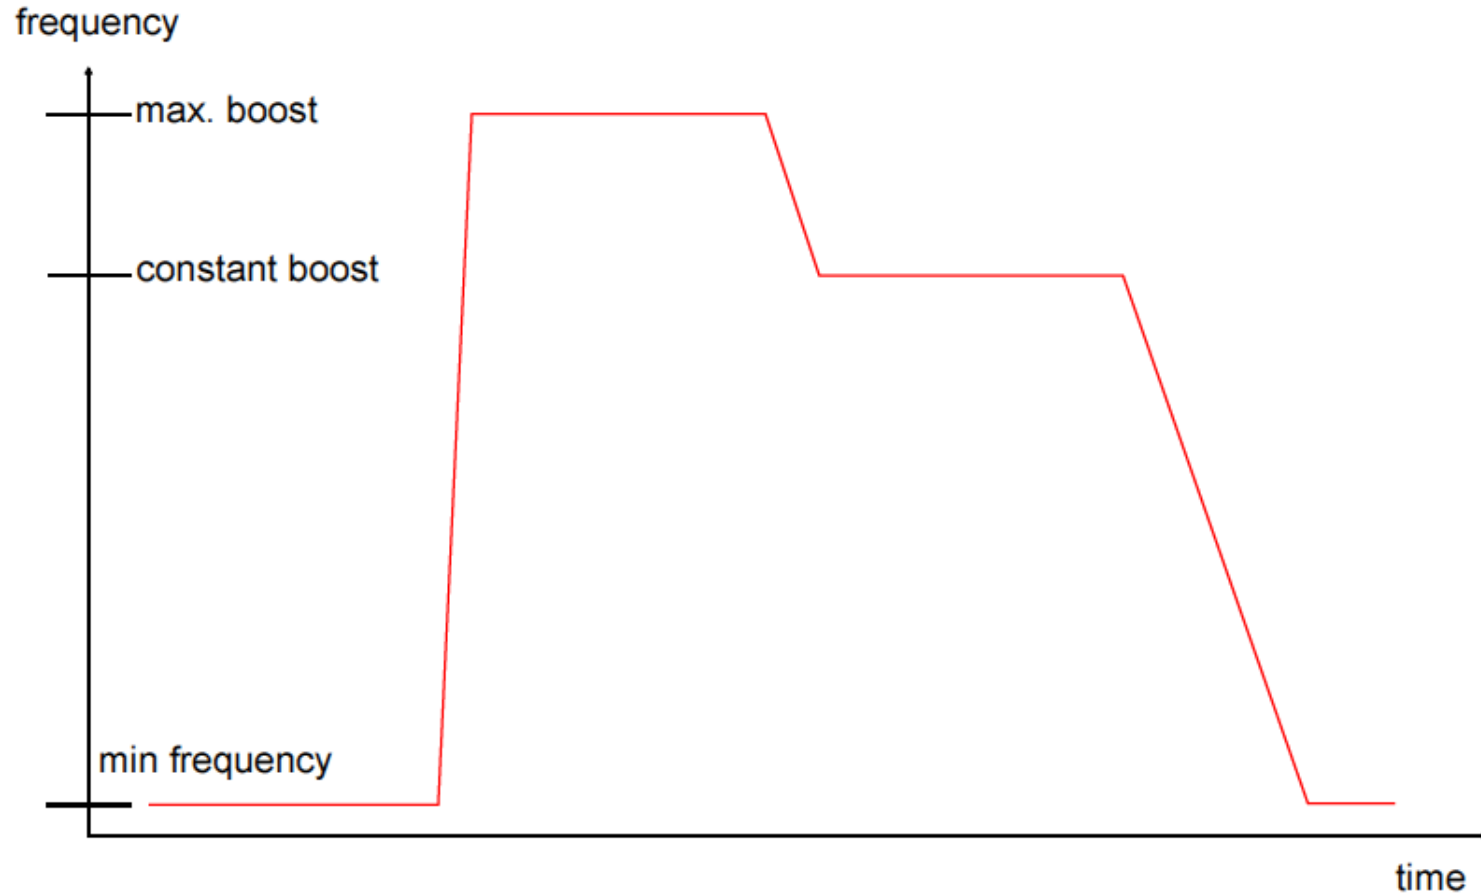
Foundations ➤ **Research Question** ➤ Approach ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach

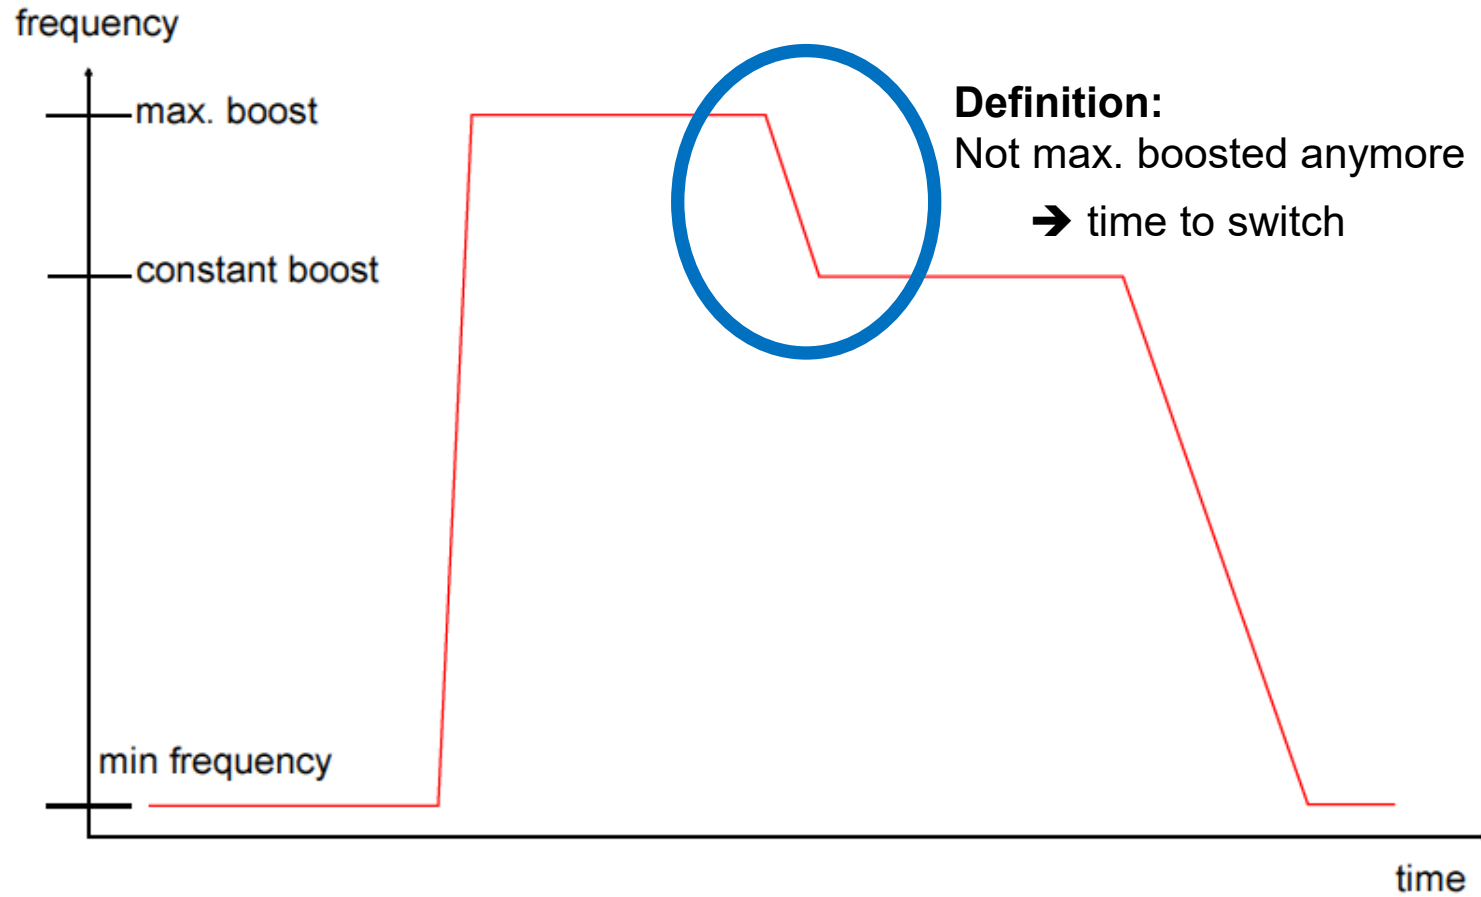
Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach

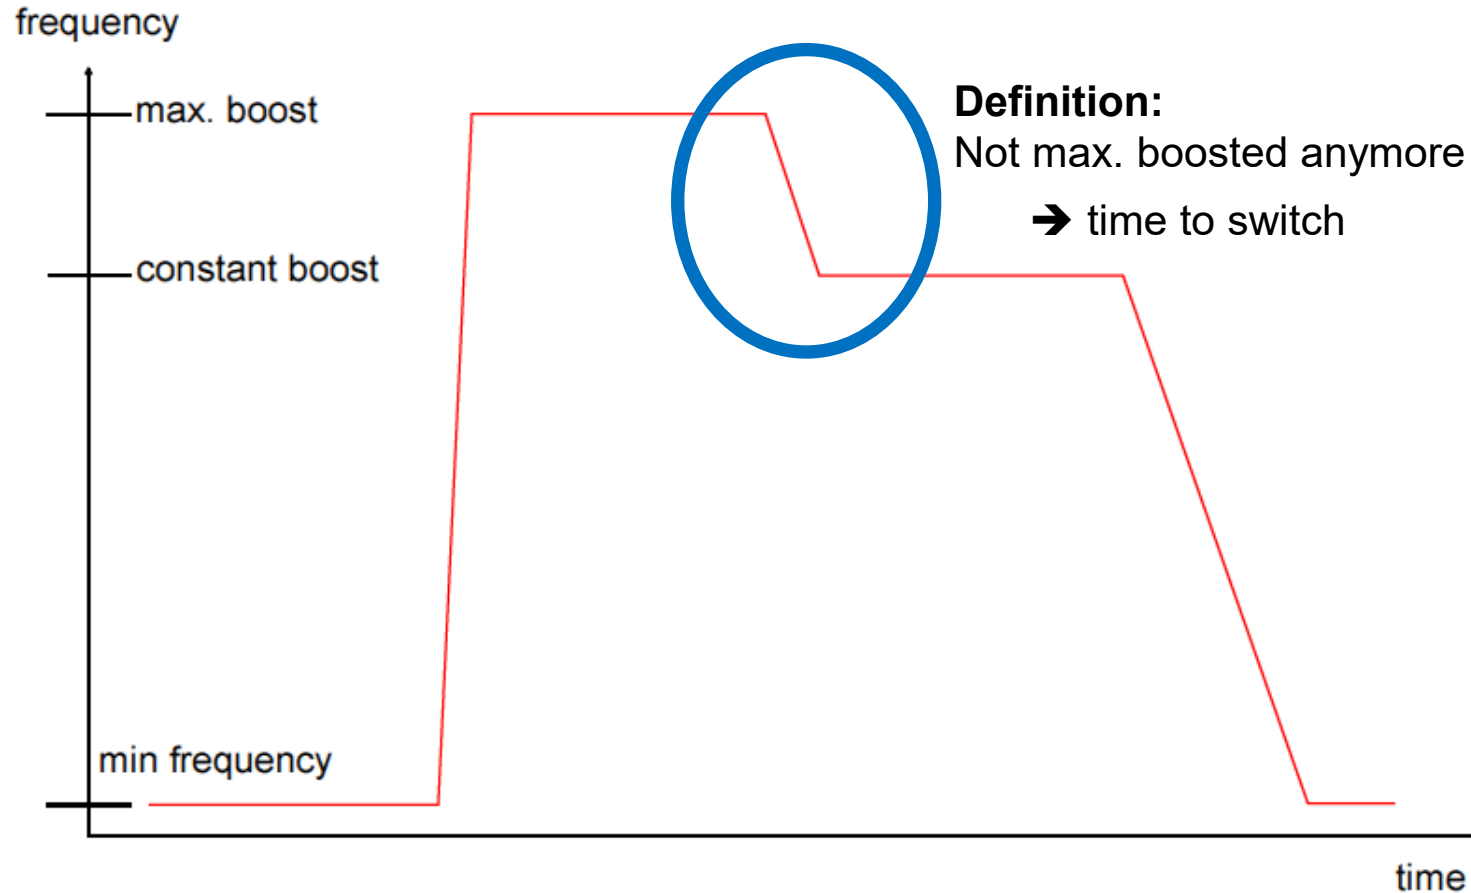
Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion



Motivation ➤ Goals & Contributions ➤ **Approach** ➤ Evaluation ➤ Conclusion

Approach

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion



➤ Assumptions

- The workload puts enough stress on a server that it will boost
- The server is under stress that the migrated workload will only be adequately serviced when boosted

Prototype Implementation

Foundations ➤ Reasearch Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

- Client: **HTTP Load Generator**



Prototype Implementation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

➤ Client: **HTTP Load Generator**

➤ Worker:

- Workload: **Bungee-LU-Servelet**

- Data collector: **Telegraf**



Prototype Implementation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

➤ Client: **HTTP Load Generator**



➤ Worker:

- Workload: **Bungee-LU-Servelet**



- Data collector: **Telegraf**



➤ Monitor

- Database: **InfluxDB**



- Dashboard: **Chronograf**



- SDN-Controller framework: **Ryu**



Prototype Implementation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

➤ Client: **HTTP Load Generator**



➤ Worker:

- Workload: **Bungee-LU-Servelet**



- Data collector: **Telegraf**  *telegraf* [6]

➤ Monitor

- Database: **InfluxDB**  *influxdb* [6]

- Dashboard: **Chronograf**  *chronograf* [6]

- SDN-Controller framework: **Ryu**  [7]

➤ Switch: Software-Switch with **Open vSwitch**



Prototype Implementation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

➤ Client: **HTTP Load Generator**

➤ Worker:

- Workload: **Bungee-LU-Servelet**

- Data collector: **Telegraf**



➤ Monitor

- Database: **InfluxDB**



- Dashboard: **Chronograf**




- SDN-Controller framework: **Ryu**



➤ Switch: Software-Switch with **Open vSwitch**




➤ Testbed with three servers, SDN, and controller



Evaluation

➤ Three evaluated scenarios:

- A. Highly overbooked
- B. Medium overbooked
- C. Not overbooked




A


Evaluation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

- Three evaluated scenarios:
 - A. Highly overbooked
 - B. Medium overbooked
 - C. Not overbooked



A




B

Motivation ➤ Goals & Contributions ➤ Approach ➤ **Evaluation** ➤ Conclusion


Evaluation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

- Three evaluated scenarios:
 - A. Highly overbooked
 - B. Medium overbooked
 - C. Not overbooked



A




C

Motivation ➤ Goals & Contributions ➤ Approach ➤ **Evaluation** ➤ Conclusion

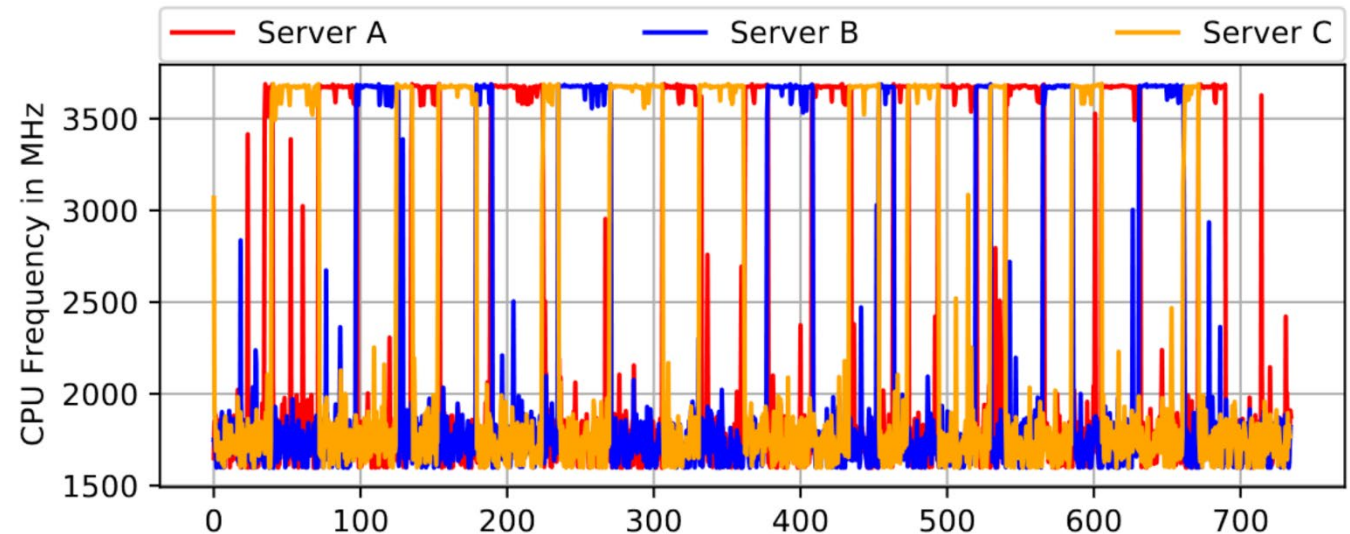
Evaluation

Foundations ➤ Research Question ➤ **Approach** ➤ Evaluation ➤ Conclusion

- Three evaluated scenarios:
 - A. Highly overbooked
 - B. Medium overbooked
 - C. Not overbooked



B



C

Motivation ➤ Goals & Contributions ➤ Approach ➤ **Evaluation** ➤ Conclusion

Preliminary Measurements

Foundations ➤ Reasearch Question ➤ Approach ➤ **Evaluation** ➤ Conclusion

	Switched	Unswitched
Average response time in seconds	0.85	1.83
Number of requests sent/reveiced/lost	1310 / 1310 / 0	1310 / 1304 / 6
Temperature max. r04/r05/r06 in °C	61.99 / 61.13 / 61.35	69.76 / * / *
Temperature average r04/r05/r06 in °C	59.29 / 58.81 / 59.13	65.96 / * / *
Average CPU frequency in MHz r04/r05/r06	2459.86 / 2147.86 / 2333.61	3518.44 / * / *
Average CPU maximum in MHz	3670.24	3580.32
Average power consumption (idle: 98.19W) in Watt	135.54	147.98
Additional power consumption in total and percent	37.35 / 38.03%	49.79 / 50.71%

Conclusion

Basics ➤ Reasearch Question ➤ Approach ➤ Evaluation ➤ **Conclusion**

Achievements

- Working heat-aware and boost-oriented load rotation
- Higher average of max. CPU frequency

Conclusion

Basics ➤ Reasearch Question ➤ Approach ➤ Evaluation ➤ **Conclusion**

Achievements

- Working heat-aware and boost-oriented load rotation
- Higher average of max. CPU frequency

Open Tasks

- Include performance metrics
 - Throughput
 - Latency
- Eliminate simplifications
- Stateful migration
- Tests with additional workloads
- Extend to heterogenous systems
- Power and temperature evaluation

Thank You!

<https://se.informatik.uni-wuerzburg.de/>

References

- [1] <https://3s81si1s5ygj3mzby34dq6qf-wpengine.netdna-ssl.com/wp-content/uploads/2018/02/aws-q4-2017-revenue-income.jpg>
- [2] <https://www.grandviewresearch.com/static/img/research/iot-market.png>
- [3] <https://www.intel.com/content/dam/doc/white-paper/intel-microarchitecture-white-paper.pdf>
- [4] https://se.informatik.uni-wuerzburg.de/fileadmin/10030200/user_upload/logos/HTTPlogo.png
- [5] https://se.informatik.uni-wuerzburg.de/fileadmin/_processed_/9/9/csm_csm_BUNGEE_02_837dac3f07_dddb202d42.png
- [6] <http://influxdata.github.io/design.influxdata.com/branding-docs/img/research/tick-stack-1.jpg>
- [7] <https://osrg.github.io/ryu/css/images/LogoSet02.png>
- [8] <https://raw.githubusercontent.com/openvswitch/openvswitch.github.io/master/assets/vswitch.png>