Power Saving Potential

http://www.eeclust.de

Timo Minartz
University of Hamburg
Research Group Scientific Computing

Motivation: Power-Proportionality

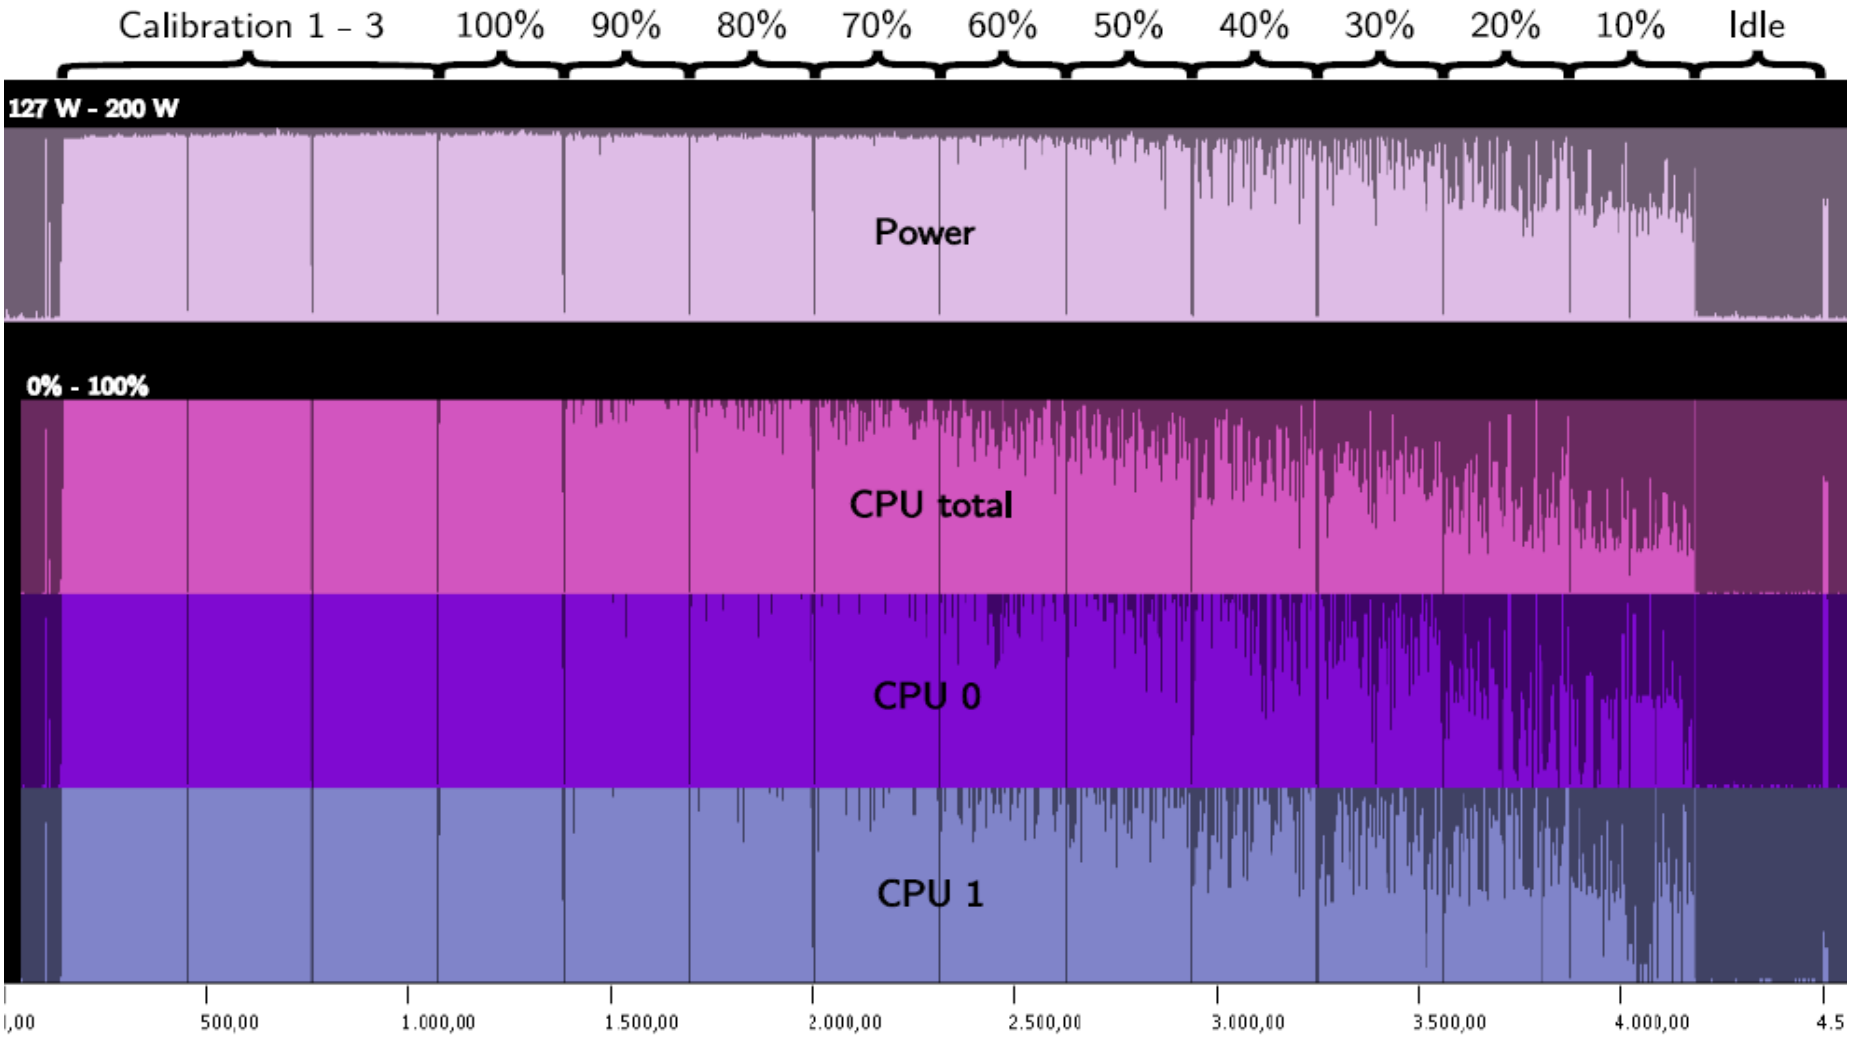

Poor Power-Proportionality

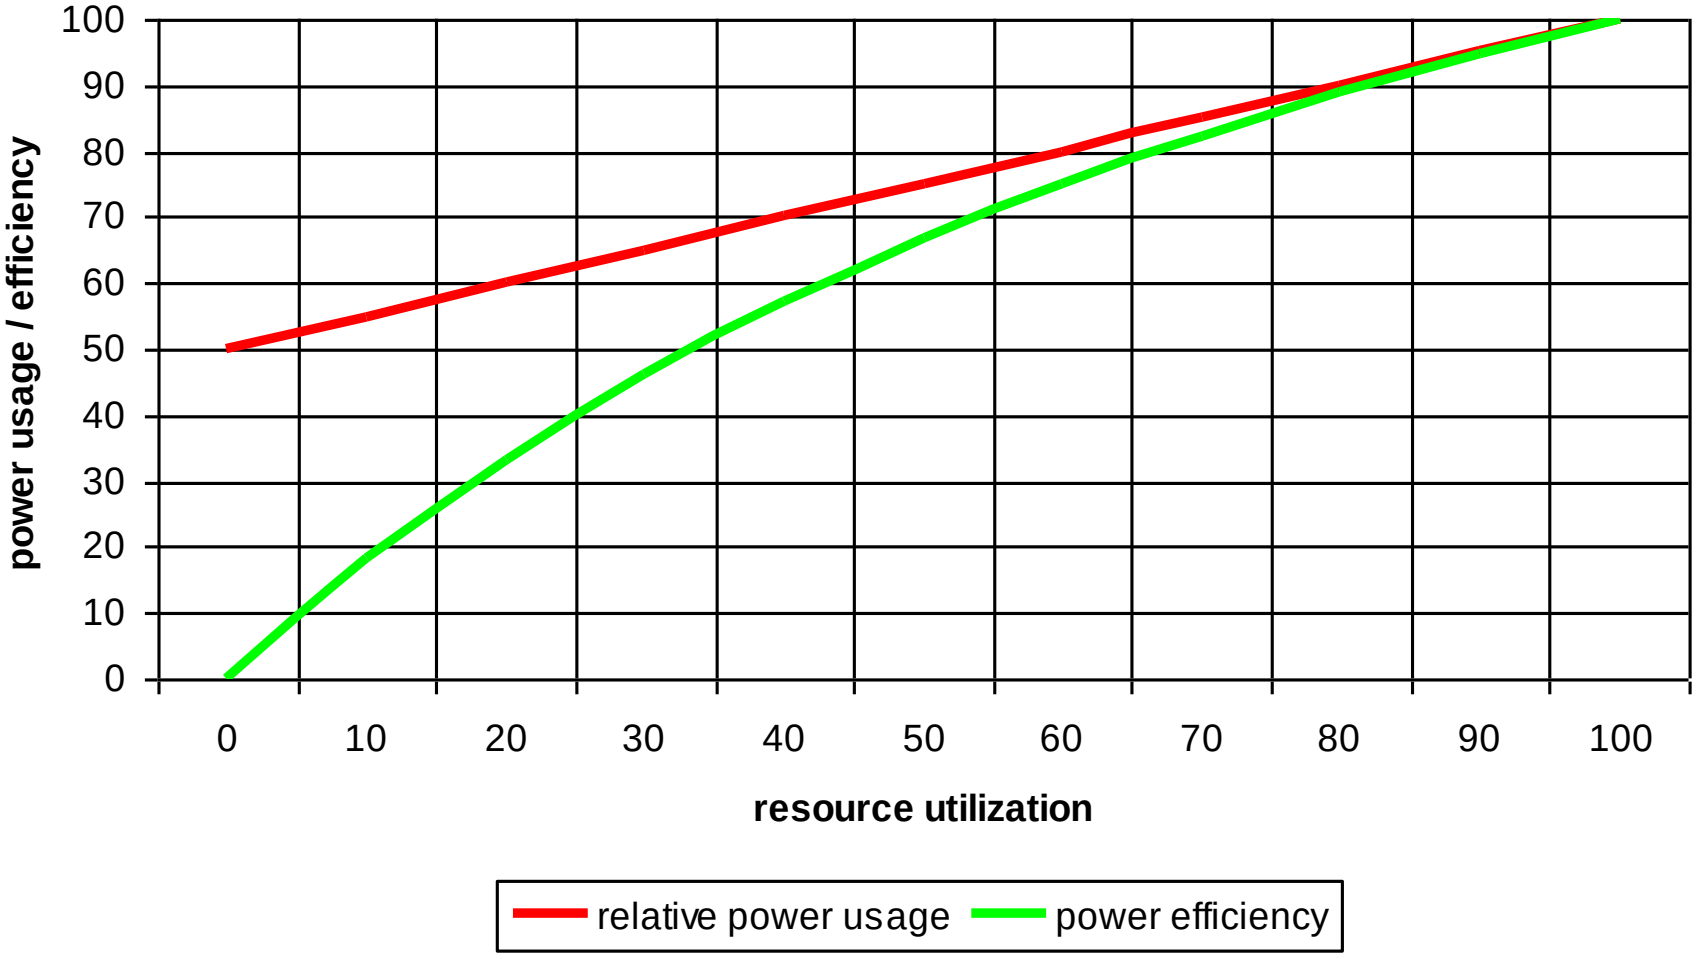

Resources

- Processors
  - Dynamic Voltage and Frequency Scaling (DVFS)
  - Power Gating
  - P-States and C-states
- Disks
  - Low rpm, head lift, cache flush
- Network
  - Frequency switching, port energy control
- Main memory
  - (Power control of memory banks)

eeClust Research Cluster

- 10 compute nodes (Dual Socket)
  - 5 Intel Xeon Nehalem X5560
  - 5 AMD Opteron Magny-Cours 6168
- 2 I/O nodes
  - 1 I/O node with 5 HDDS
  - 1 I/O node with 5 SSDs
- 3 ZES High Precision Power Meters
  - 4 channels each
  - Up to 20 Hz samples
  - 0.1 % precision

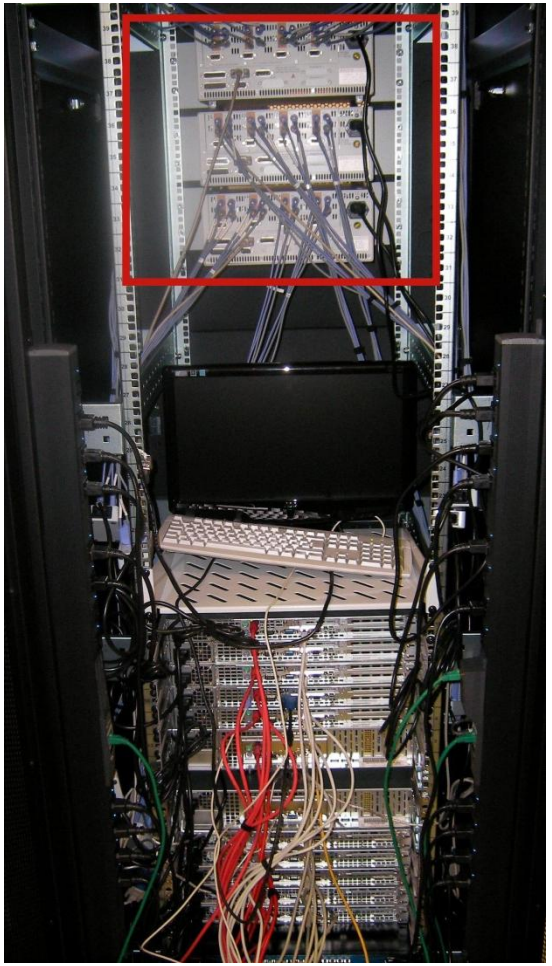

Power Saving Potential

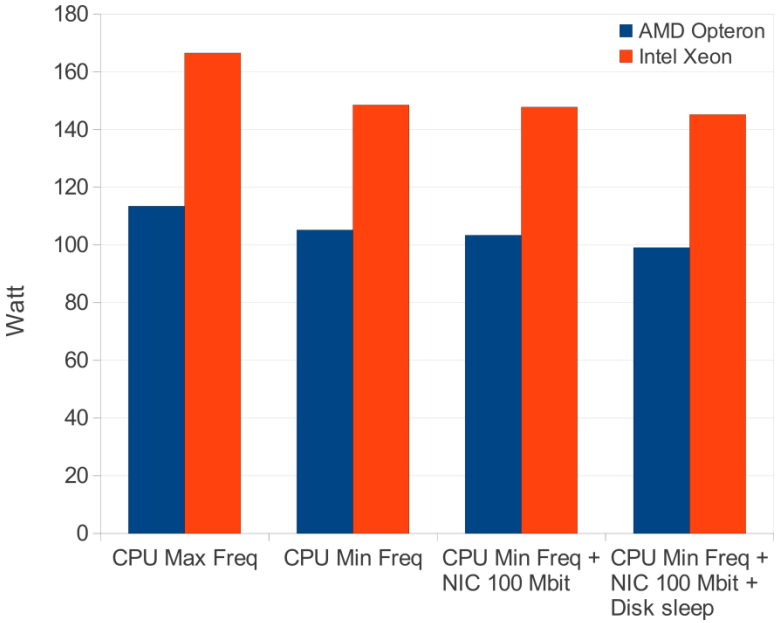

- Idle savings for AMD and Intel nodes
- C-States deactivated

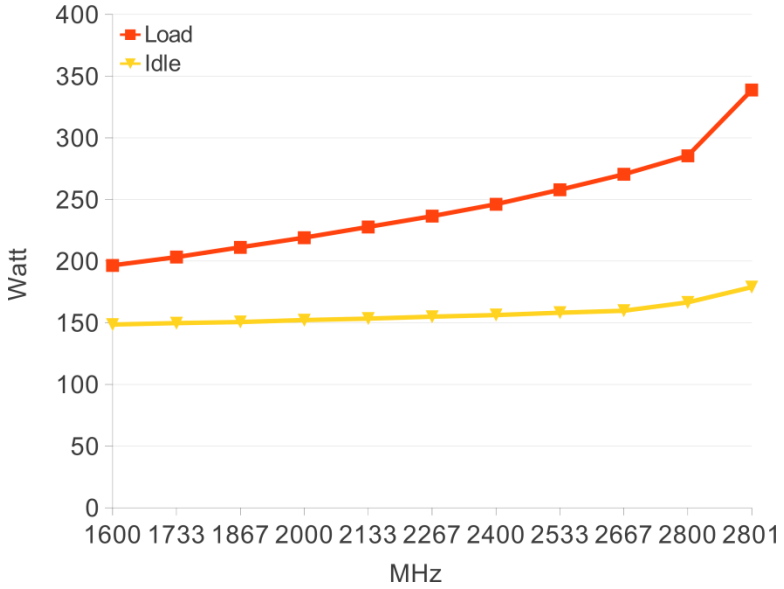

- Load savings for Intel Xeon nodes
- C-States deactivated

Spec Power (Intel)

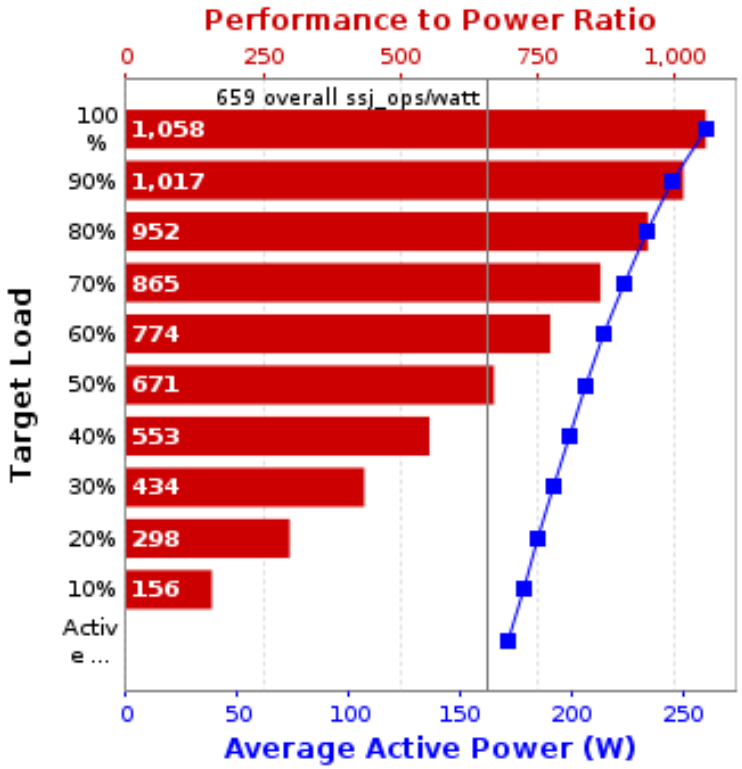

P-States deactivated
C-States deactivated
TurboBoost deactivated
SMT activated

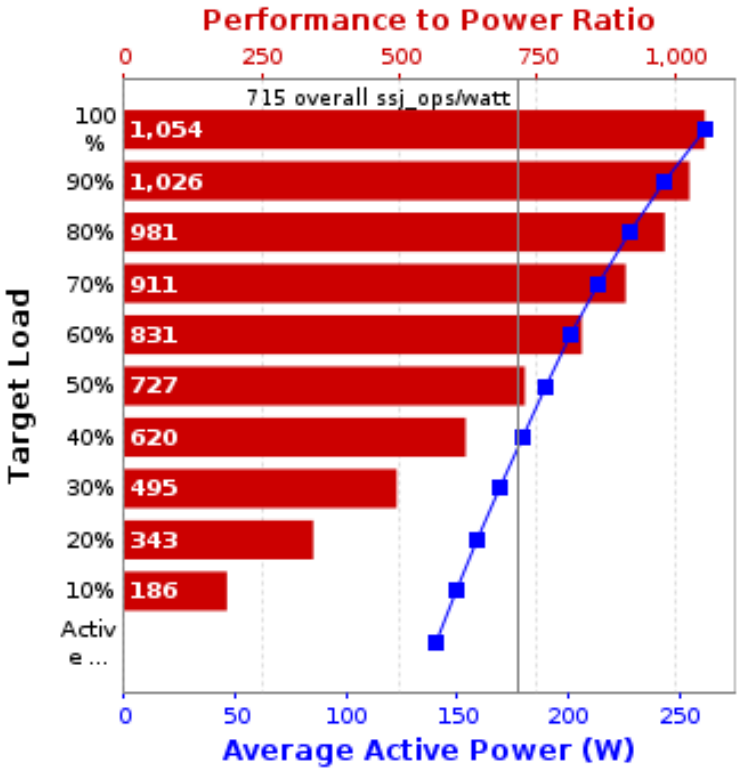

P-States activated
C-States activated
TurboBoost deactivated
SMT activated

Spec Power (AMD)

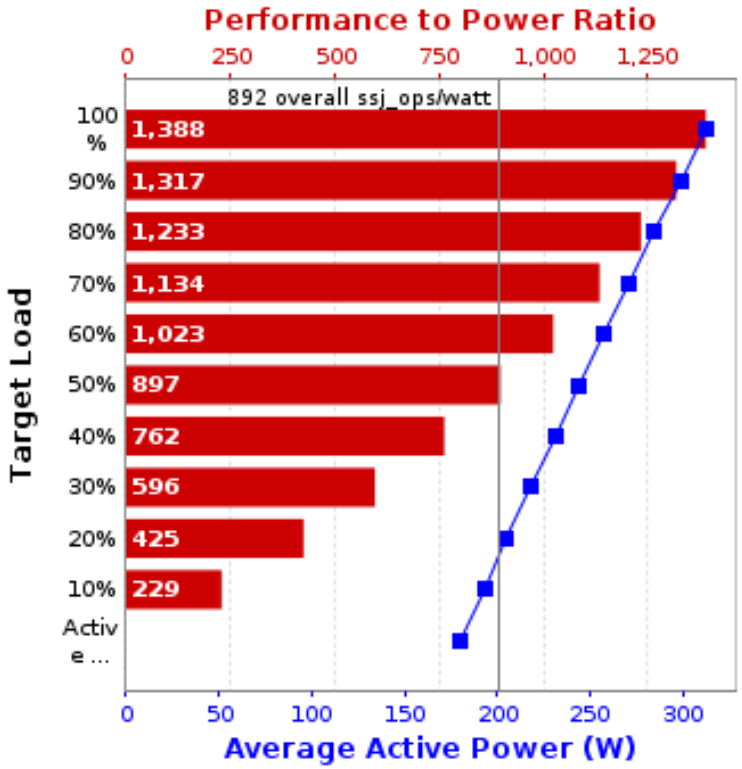

P-States deactivated
C1E deactivated

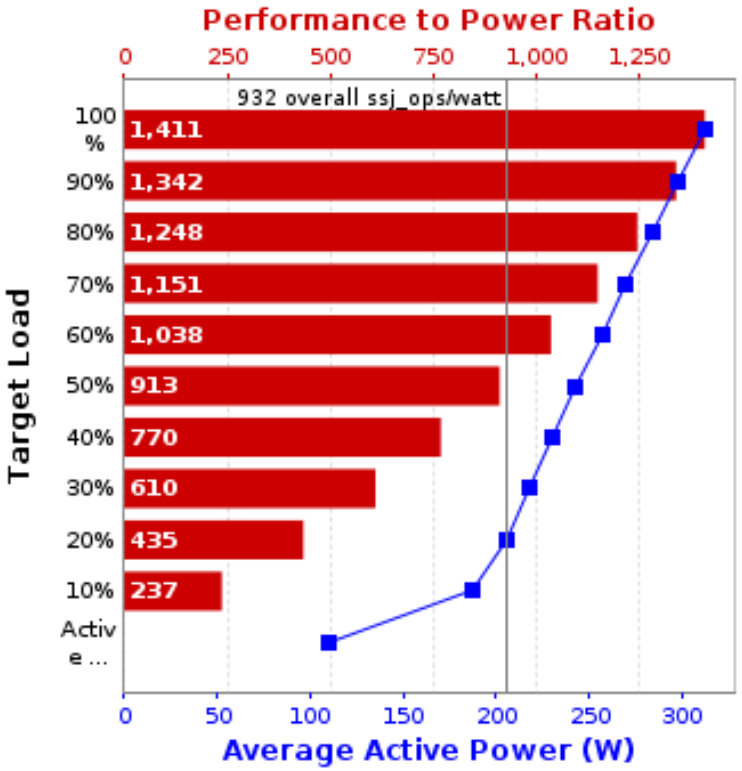

P-States activated
C1E activated

Conclusion

- 10 node test cluster with power saving mechanisms
- Identified C-State / C1E power saving potential for idle nodes
  - AMD: 40 %
  - Intel: 30 %
- Identified P-State power saving potential (idle and load) for nodes
  - AMD: 30 % Load, 8 % Idle
  - Intel: 30 % Load, <1 % Idle