Julius-Maximilians-
UNIVERSITÄT
WÜRZBURG

Common Errors and Assumptions in Energy Measurement and Management

Jóakim v. Kistowski
University of Würzburg

Symposium on Software Performance,
November 5th 2015, Munich, Germany

What is this Talk about?

- Measurement methodologies for energy efficiency
  - Focus on server systems
- Some pitfalls: Energy efficiency measurements can be unrepresentative or inaccurate if done incorrectly
- SPEC power methodology [1]: A methodology for standardized energy efficiency benchmarking
- Some results that challenge common implicit assumptions on energy efficiency of servers

Energy Efficiency of Servers

- Relationship of Performance and Power
- For transactional workloads:

$$\frac{transactions}{energy}\left[\frac{1}{J}\right] = \frac{throughput}{power}\left[\frac{1/s}{W}\right]$$

- Comparison of efficiency of different workload types is difficult
  - Different scales of transaction-counts / throughput
  - ➔ normalization

How to do it wrong…

PITFALLS IN POWER MEASUREMENT

Measuring at Maximum Load (1/2)

A typical server …

- has an average utilization between 10% and 50%,
- is provisioned with additional capacity (to deal with load spikes).

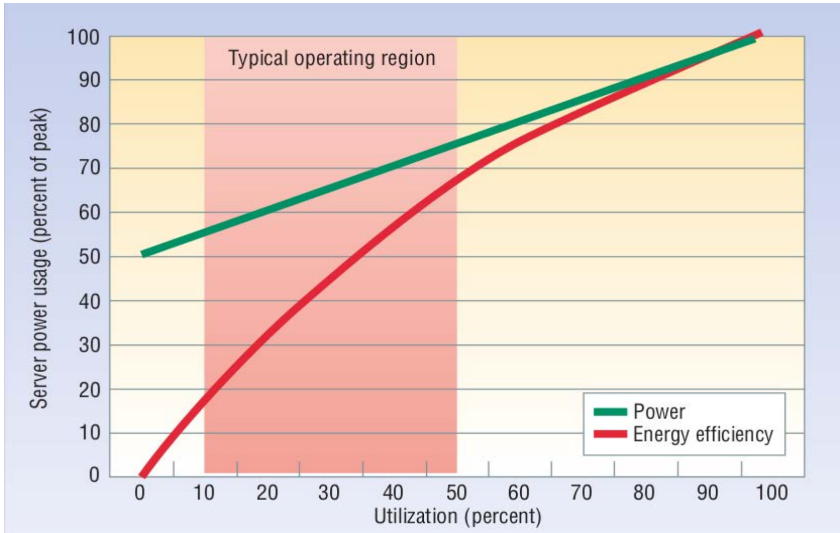

Energy Efficiency and Power Consumption of Servers [2]

- is not energy efficient at low utilization, more efficient at high utilization

Power consumption depends on server utilization.

Measuring at Maximum Load (2/2)

Bad Practice for…

- Full system power characterization
- Comparison of server systems intended for transactional workloads (most of them)

Good Practice for…

- HPC energy efficiency benchmarking

Varying Loads (1/2)

- Power meters have power measurement ranges
  - Lose measurement accuracy outside of range
  - Switching ranges takes time (~ 1 s)

- Example

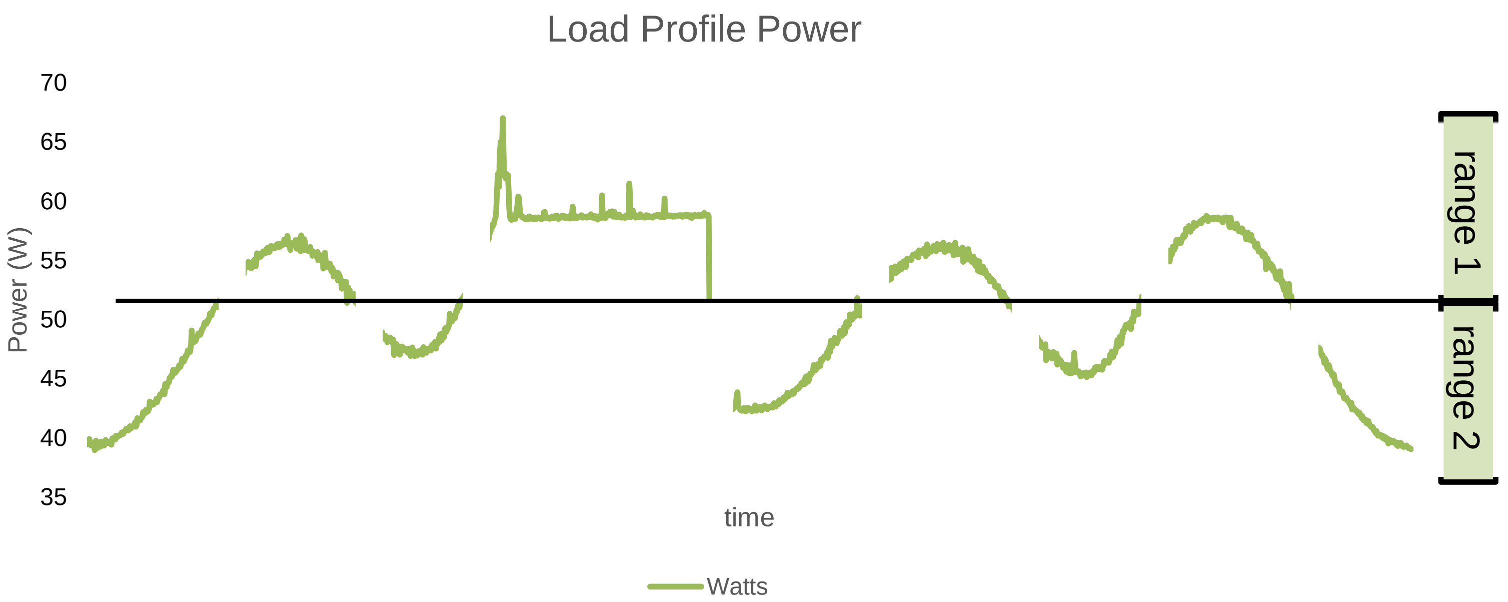

Varying Loads (2/2)

Lessons:

- Auto-Ranging is bad for varying loads
  - Lose measurements
- But:
  - Disabling auto-ranging decreases accuracy
- Measurement uncertainty depends on power meter
  - SPEC PTDaemon supported ➔ Less than 1% at optimal range
- Also:
  - Good load calibration is important

How to do it right…

SPEC POWER METHODOLOGY

SPEC Power Methodology

- Methodology for benchmarking of energy efficiency

- Goal:
  - Benchmarking at multiple load levels
  - Taking the quality criteria for benchmarks into account [3]:
    - Relevance
    - Reproducibility
    - Fairness
    - Verifiability
    - Usability

- Used in the following SPEC products:
  - SPECpower_ssj2008 [4]
  - SPEC SERT [5]
  - ChauffeurWDK

- Other Benchmarks that follow the methodology:
  - SAP Power Benchmark [6]
  - TPC Energy [7]

Load Levels

- Goal: For a given workload, achieve a load level of n% of system “utilization”.
- Utilization = $\frac{t_{busy}}{t_{busy} + t_{idle}}$
- DVFS increases CPU busy time at low load
  - ➔ increases utilization
  - Power over load measurements need to compensate

    How to compare?
- Our solution: Machine utilization
  - 100% utilization at calibrated maximum throughput
  - Load level = $\frac{current\ throughput}{max.\, throughput}$

SERT Architecture

- Controller System runs
  - SPEC Director: Chaffeur
  - Reporter
- PTDaemon
  - Network-capable power and temperature measurement interface
  - Can run on controller system or separate machine
- SUT runs
  - Host, which launches
  - Pinned SERT clients

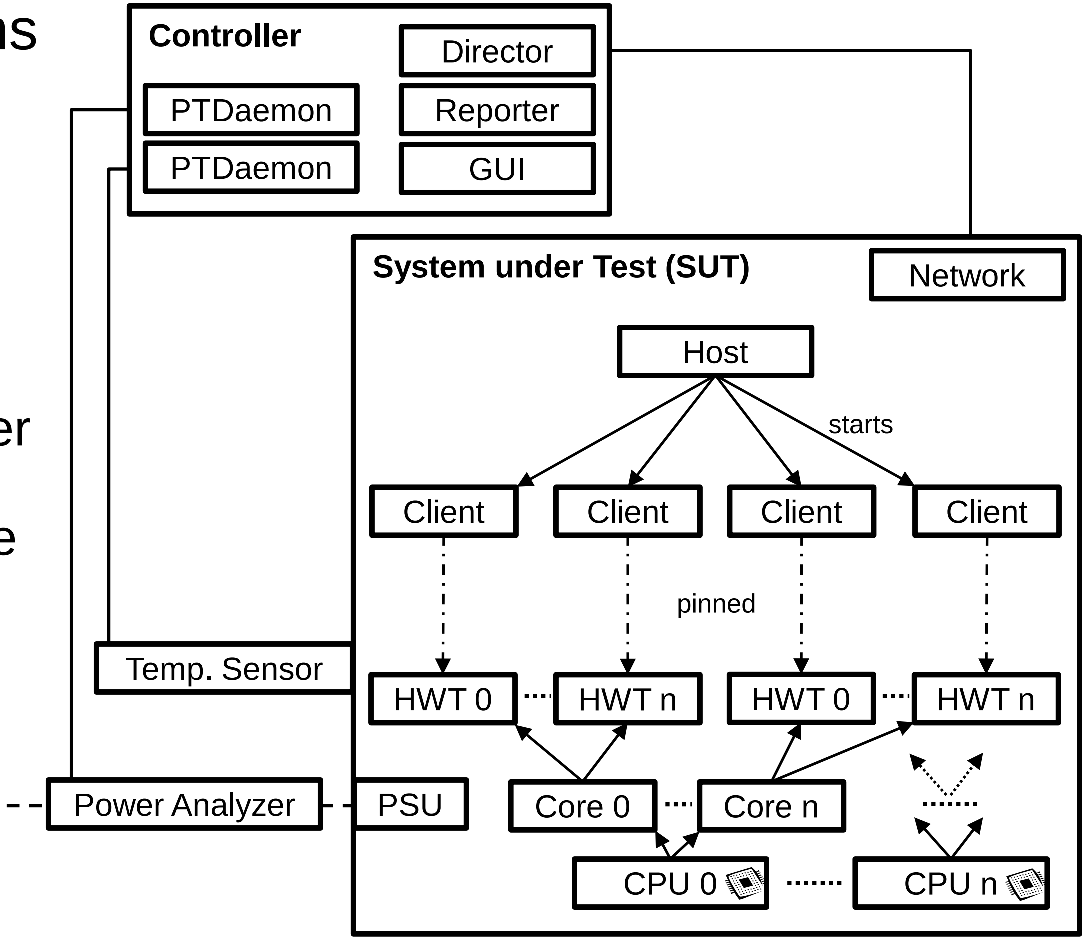

SERT Measurement (1/2)

- Transactional workloads are dispatched in “Intervals”:
  - Warmup
  - Calibration
    - Multiple intervals
    - Maximum transaction rate
  - Graduated Measurement Series
    - Multiple intervals at decreasing transaction rate
    - Target transaction rate is percentage of calibration result
    - Exponentially distributed wait times between transactions

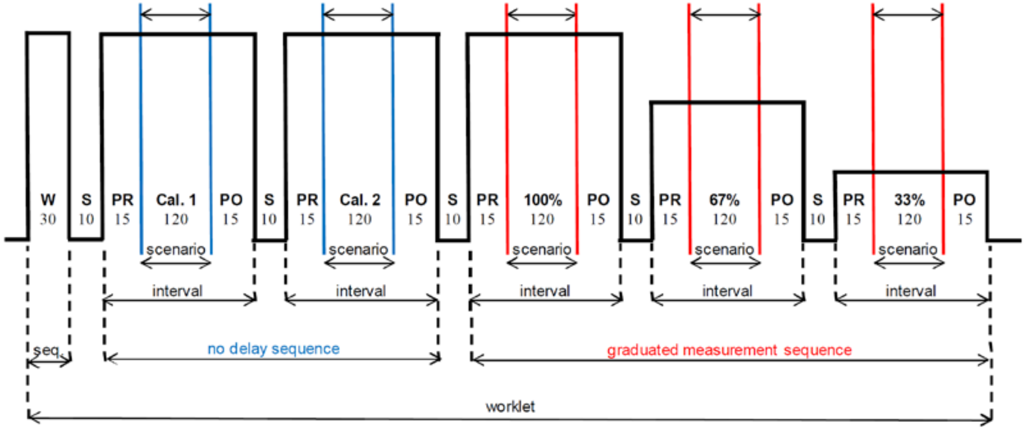

SERT Measurement (2/2)

- Separate measurement intervals at stable states
  - 10 second sleep between intervals
  - 15 second pre-measurement run
  - 15 second post-measurement run
  - 120 second measurement

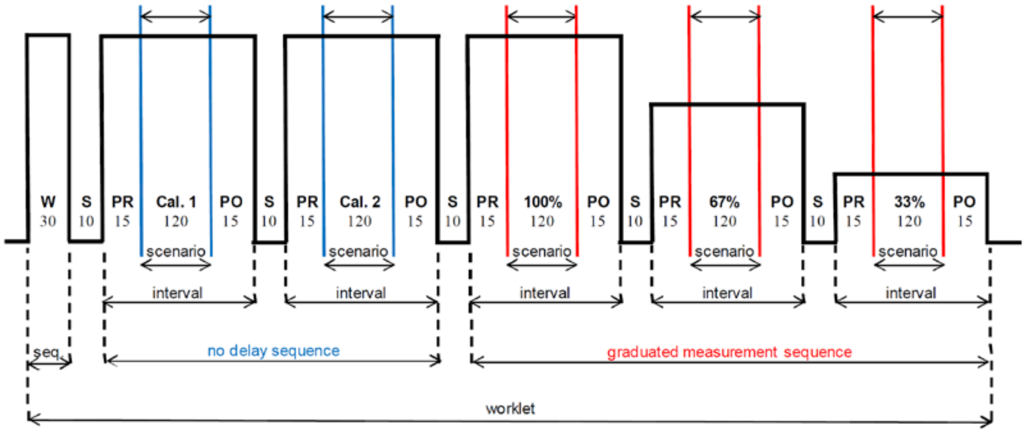

- Temperature analyzer for comparable ambient temperature
- Power Measurements: AC Wall Power

Performance and Power Variation

- Throughput results from load level definition
  - Throughput variation is measure of benchmark driver stability
  - Throughput coefficient of variation > 5% ➔ invalid interval

- Power consumption results from SUT response to load
  - Power variation is measure of SUT stability
  - CVs often < 1% on state-of-the-art x86 systems

Workloads

- Workloads can be anything, as long as…
- … they have a measurable throughput
- … allow for result validation

- Common Workloads:
  - SPEC SERT: “Worklets”
    - 7 CPU Workets
    - 2 HDD Worklets
    - 2 Memory Worklets
    - 1 Hybrid Worklet (SSJ)
  - SPECpower_ssj2008: Buisiness Transactions
  - TPC Energy
  - ChauffeurWDK: Allows custom workload creation

Motivating future work…

SOME MEASUREMENT RESULTS

The Software Stack Matters! (1/2)

(With differing extent)

- Operating System [8]
  - Impact on base consumption and power scaling behavior

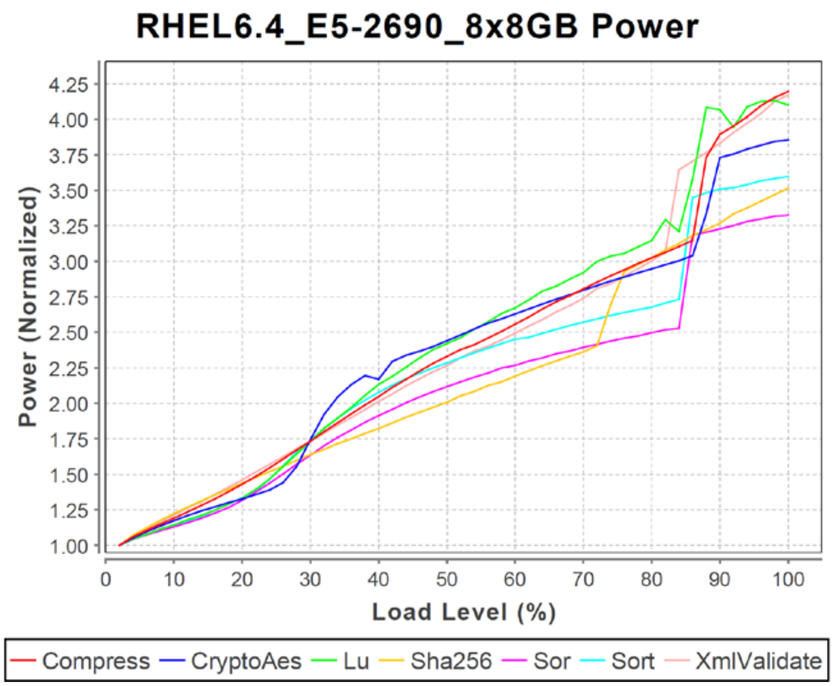

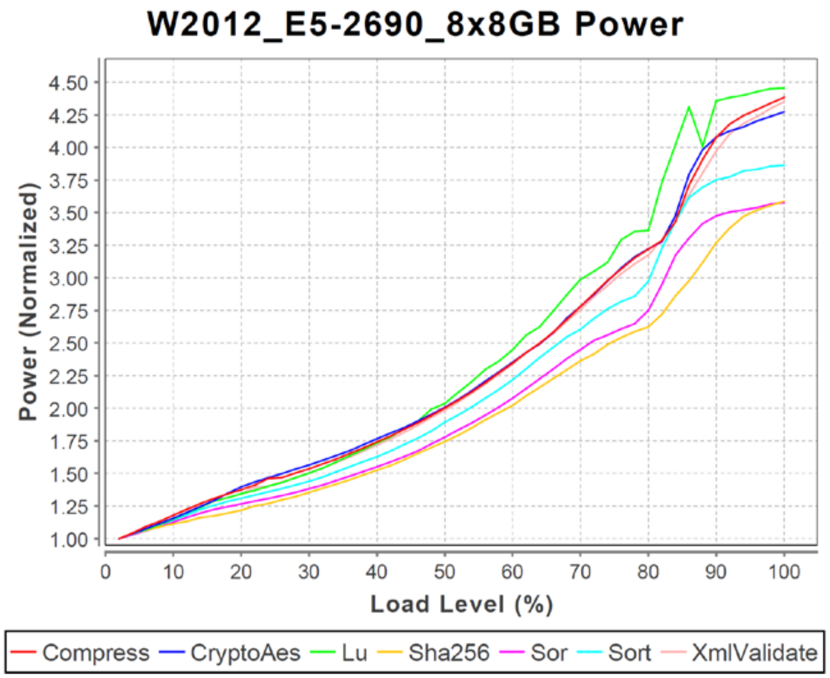

The Software Stack Matters! (2/2)

(With differing extent)

- JVM [8]
  - Little impact through secondary effects

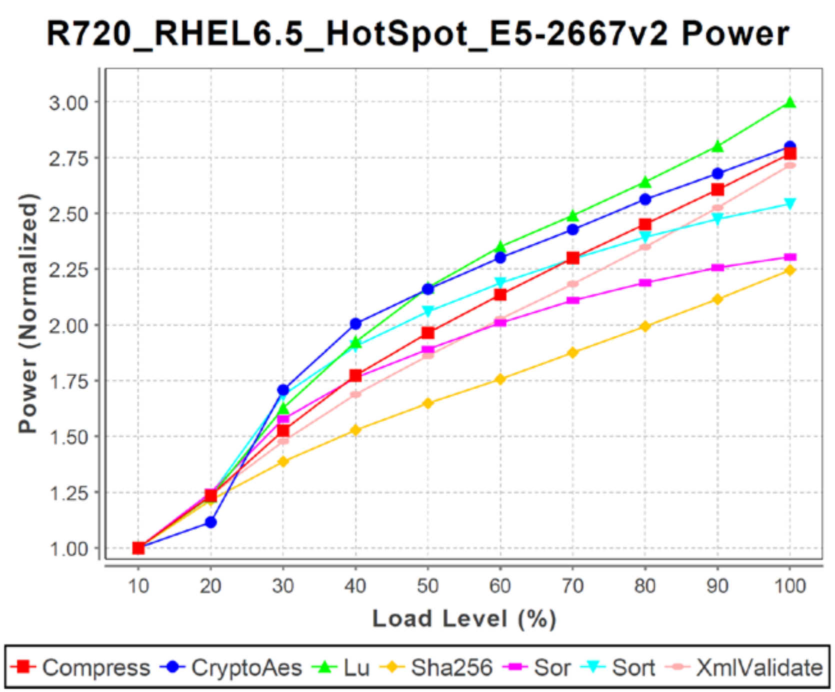

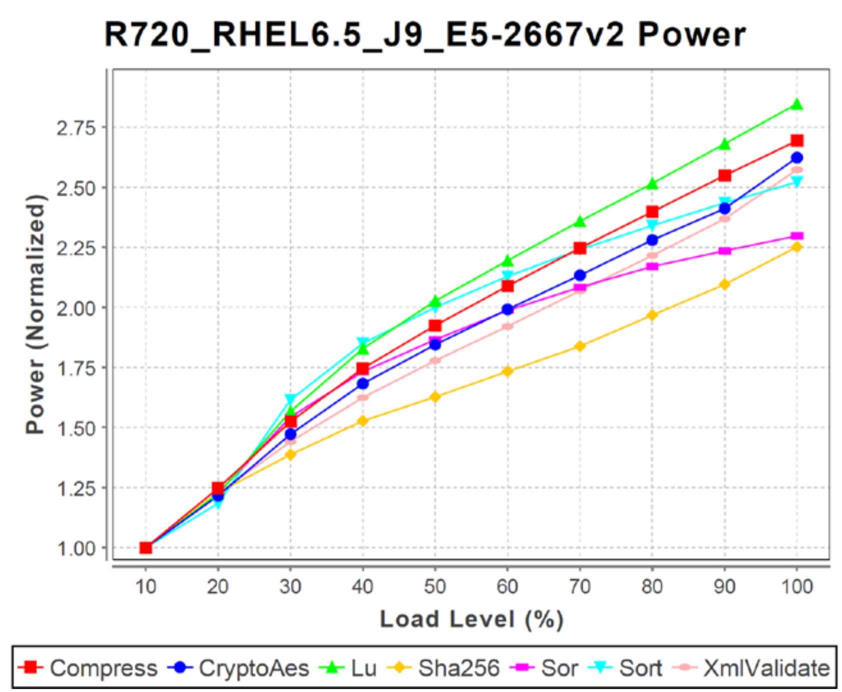

Maximum Energy Efficiency

- Energy Efficiency depends on multiple factors
  - Hardware
  - Software Stack
  - Workload
  - Load Distribution

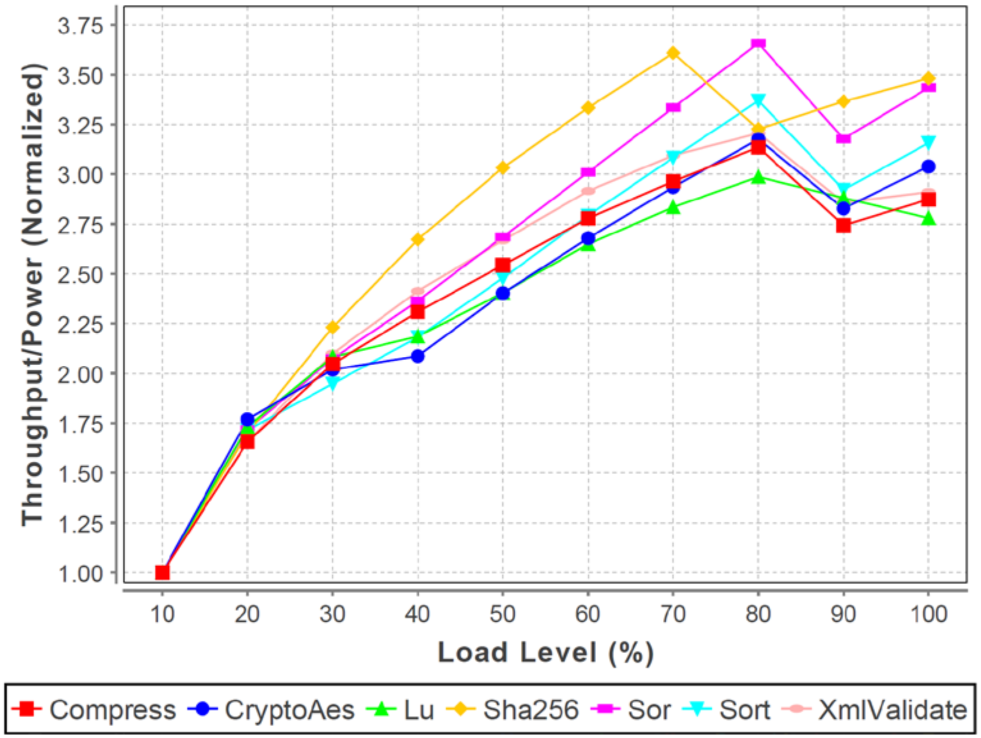

- Maximum Energy Efficiency is often reached at < 100% load

- Result: Load Consolidation is not most efficient load distribution strategy [9]

Conclusions

- Power and energy efficiency measurements has many pitfalls
  - Can lead to inaccurate or missing results

- SPEC power methodology is an established standard to avoid errors in energy efficiency benchmarking
  - Goal: Energy efficiency characterization at multiple load levels

- Results demonstrate that energy efficiency and energy efficiency scaling depend on many factors, including hardware, software stack, workload, etc.

Thanks for listening!

joakim.kistowski@uni-wuerzburg.de
http://se.informatik.uni-wuerzburg.de

Trademark and Disclaimers

The SPEC logo, SPEC, and the benchmark and tool names, SPECpower_ssj, SERT, PTDaemon are registered trademarks of the Standard Performance Evaluation Corporation. Reprint with permission, see spec.org.

The opinions expressed in this tutorial are those of the author and do not represent official views of either the Standard Performance Evaluation Corporation, Transaction Processing Performance Council or author's company affiliation.

References

[1] SPEC Power and Performance Benchmark Methodology. http://spec.org/power/docs/SPEC-Power_and_Performance_Methodology.pdf .

[2] L. Barroso and U. Holzle. The Case for Energy Proportional Computing. *Computer*, 40(12):33-37, Dec 2007.

[3] J. von Kistowski, J. A. Arnold, K. Huppler, K.-D. Lange, J. L. Henning, and P. Cao. How to Build a Benchmark. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering (ICPE 2015)*, New York, NY, USA, February 2015. ACM.

[4] K.-D. Lange. Identifying Shades of Green: The SPECpower Benchmarks. *Computer*, March 2009.

[5] K.-D. Lange and M. G. Tricker. The Design and Development of the Server Efficiency Rating Tool (SERT). In *Proceedings of the 2nd ACM/SPEC International Conference on Performance Engineering*, ICPE'11, New York, NY, USA, 2011. ACM.

[6] SAP Power Benchmarks Specification. http://global.sap.com/solutions/benchmark/pdf/Specification_SAP_Power_Benchmarks_V12.pdf .

[7] M. Poess, R. O. Nambiar, K. Vaid, J. M. Stephens Jr, K. Huppler, and E. Haines. Energy benchmarks: a detailed analysis. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, 2010. ACM.

[8] J. von Kistowski, H. Block, J. Beckett, K.-D. Lange, J. A. Arnold, and S. Kounev. Analysis of the Influences on Server Power Consumption and Energy Efficiency for CPU-Intensive Workloads. In *Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering (ICPE 2015)*, Austin, TX, USA, February 2015. ACM.

[9] J. von Kistowski, J. Beckett, K.-D. Lange, H. Block, J. A. Arnold, and S. Kounev. Energy Efficiency of Hierarchical Server Load Distribution Strategies. In *Proceedings of the IEEE 23nd International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2015)*, Atlanta, GA, USA, October 5-7, 2015. IEEE.