



Towards an Enhanced Benchmark Advocating Energy-Efficient Systems



Daniel Schall, Volker Hudlet,

Manuel Kern



Overview

- TPC-* Benchmarks & TPC-Energy
- Energy Consumption
& Usage Patterns of DB servers
- Benchmark Proposal
- Future Work

TPC-* Benchmarks

- TPC-C, TPC-H, TPC-E
- performance-centric
- measure throughput at peak
- pricing specifications

→ tpmC, tpsE,
price / performance

TPC-C spec.

TPC-H spec.


TPC-E spec.

TPC-Energy

- add-on to any of the TPC benchmarks
- power consumption during benchmark run
- idle power consumption


→ tpmC, tpsE,
price / performance

→ Watt / performance




TPC-* & TPC-Energy

- limited to peak throughput
- sufficient for performance-oriented customers
- What about „average“ customers with „average“ servers?

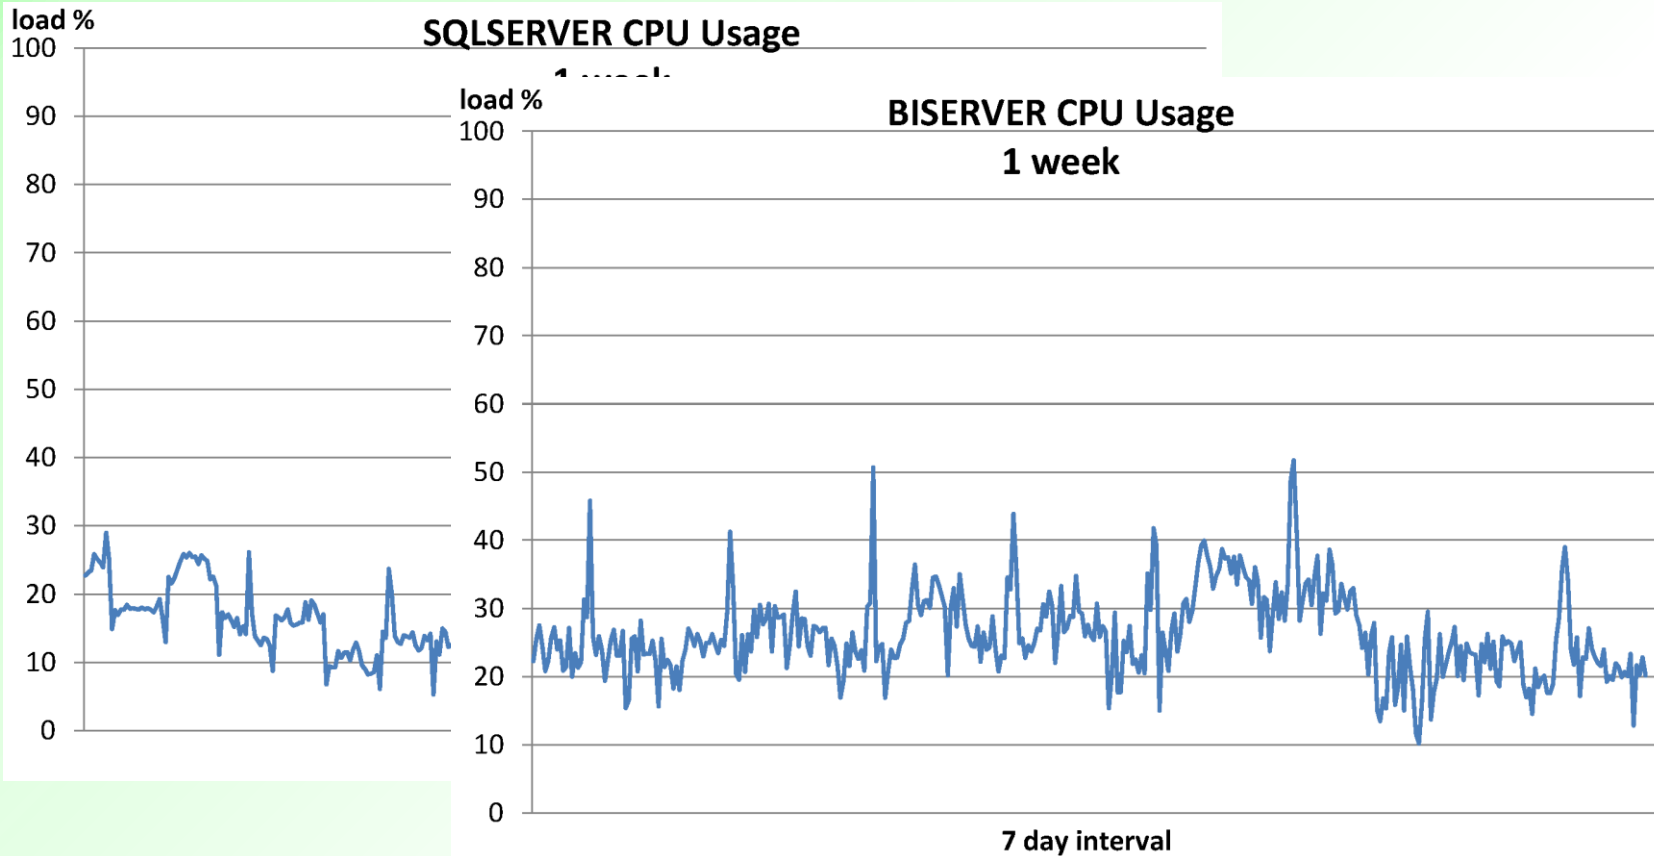


What does your server do all day?



Average CPU utilization of more than 5,000 servers,
see A. Barroso and U. Hölzle: The Case for Energy-Proportional Computing

What does your server do all day?



Study by SPH AG, Stuttgart

Monitoring for 1 week, ERP backend & analysis servers

What does your server do all day?

- > 50 % more or less idle




- 20 – 30 % doing „something“



- < 10 % under stress




D`Oh?



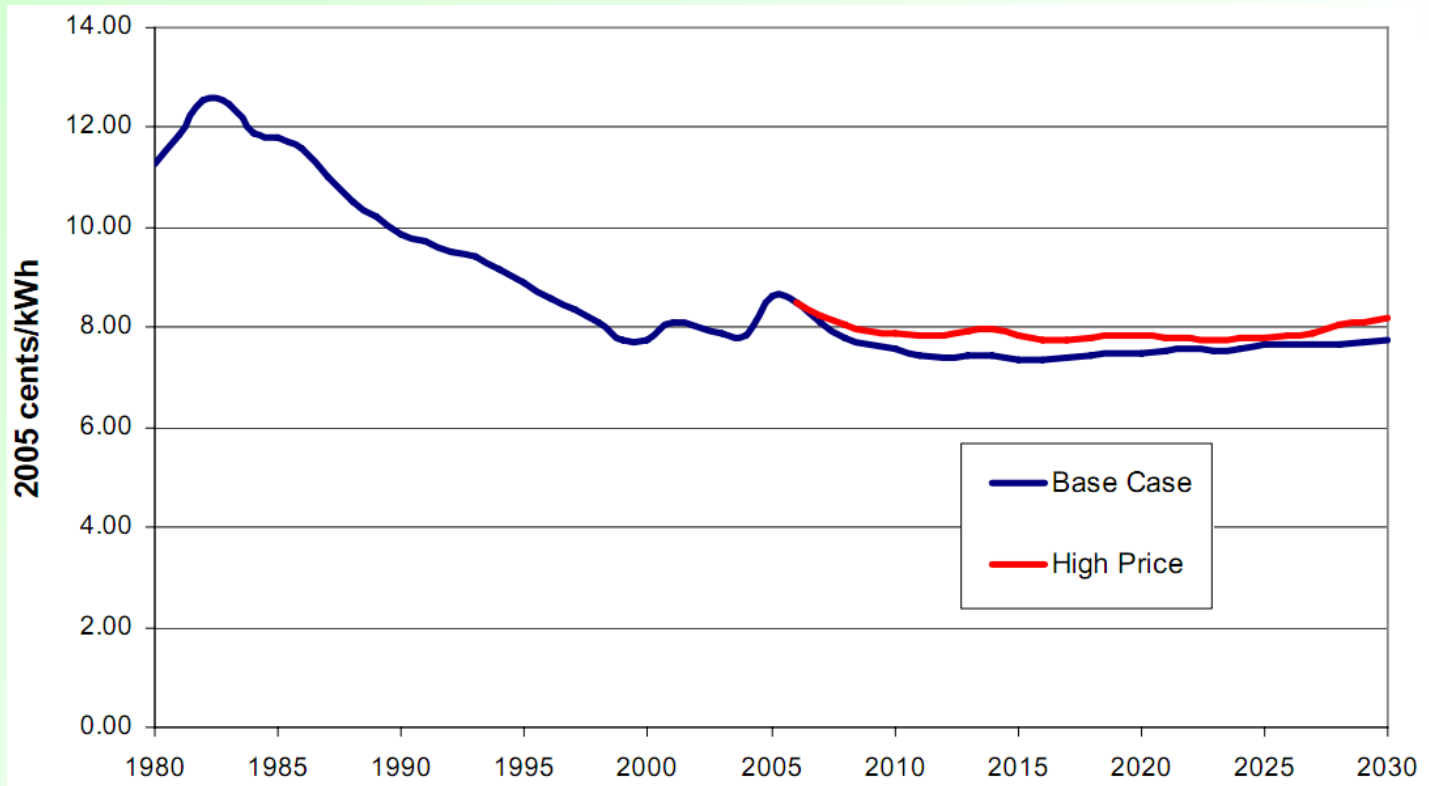
Power by component at different activity levels,
see A. Z. Spector: Distributed Computing at Multi-dimensional Scale, MIDDLEWARE 2008

D`Oh!



Predicted US electricity use for data centers,
see J. Koomey: GROWTH IN DATA CENTER ELECTRICITY USE 2005 TO 2010

D`Oh!



U.S. Average electricity price by year,

see DOE EIA Annual Energy Outlook 2006

Conclusions

- typical server load is far from peak
- TPC-benchmarks limited
- energy spendings are steadily rising

- need for a new benchmarking paradigm
- energy proportionality

What is Energy Proportionality?

$$EP(x) = \frac{PC_{ideal}(x)}{PC_{act}(x)} = \frac{x}{PC_{act}(x)}$$

power consumption linear to load x


dependent on load level x

real power consumption

IDEAL SYSTEM

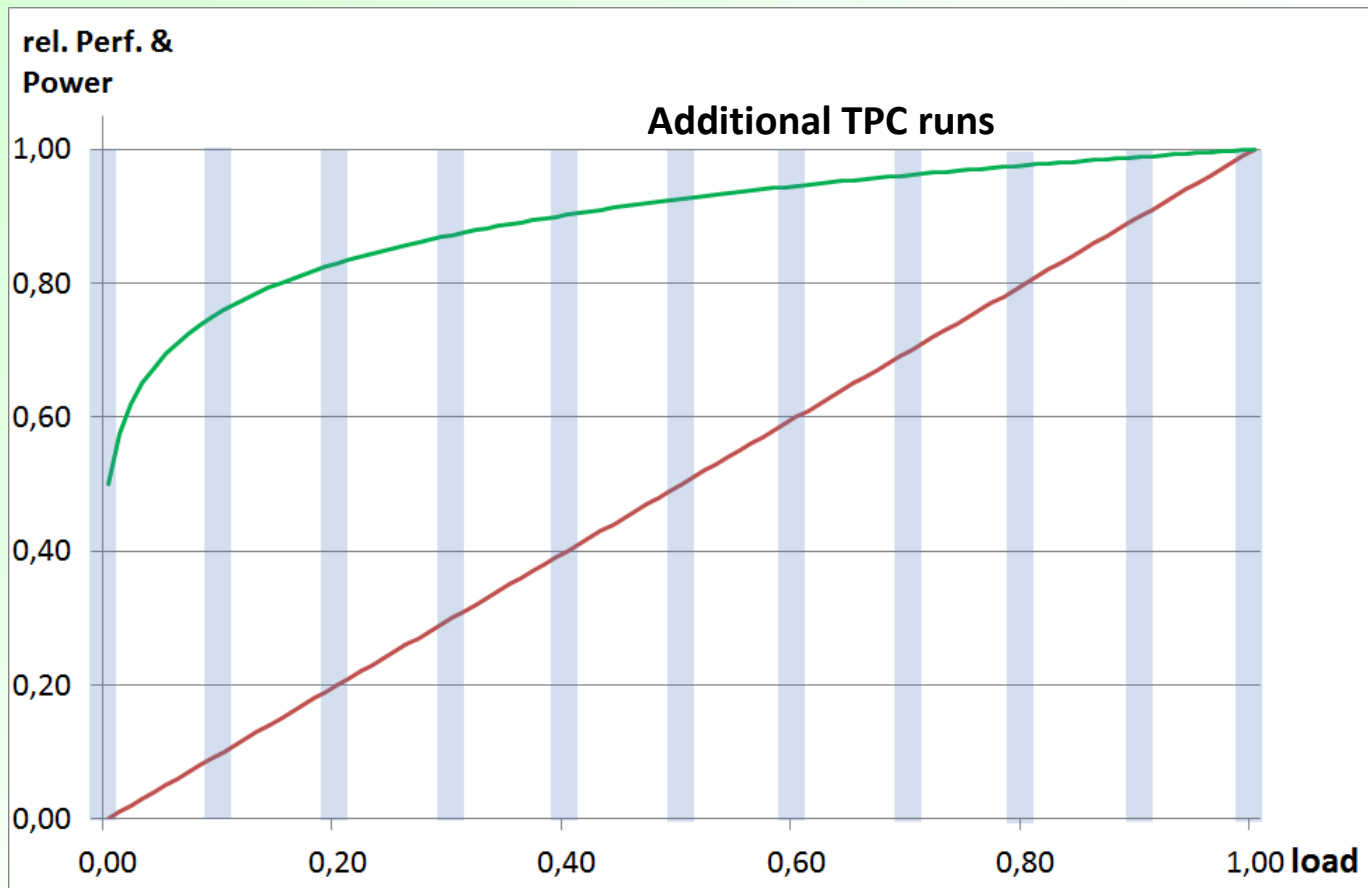
DB server running at 90% load → energy consumption: 90% of peak

DB server running at 20% load → energy consumption: 20% of peak



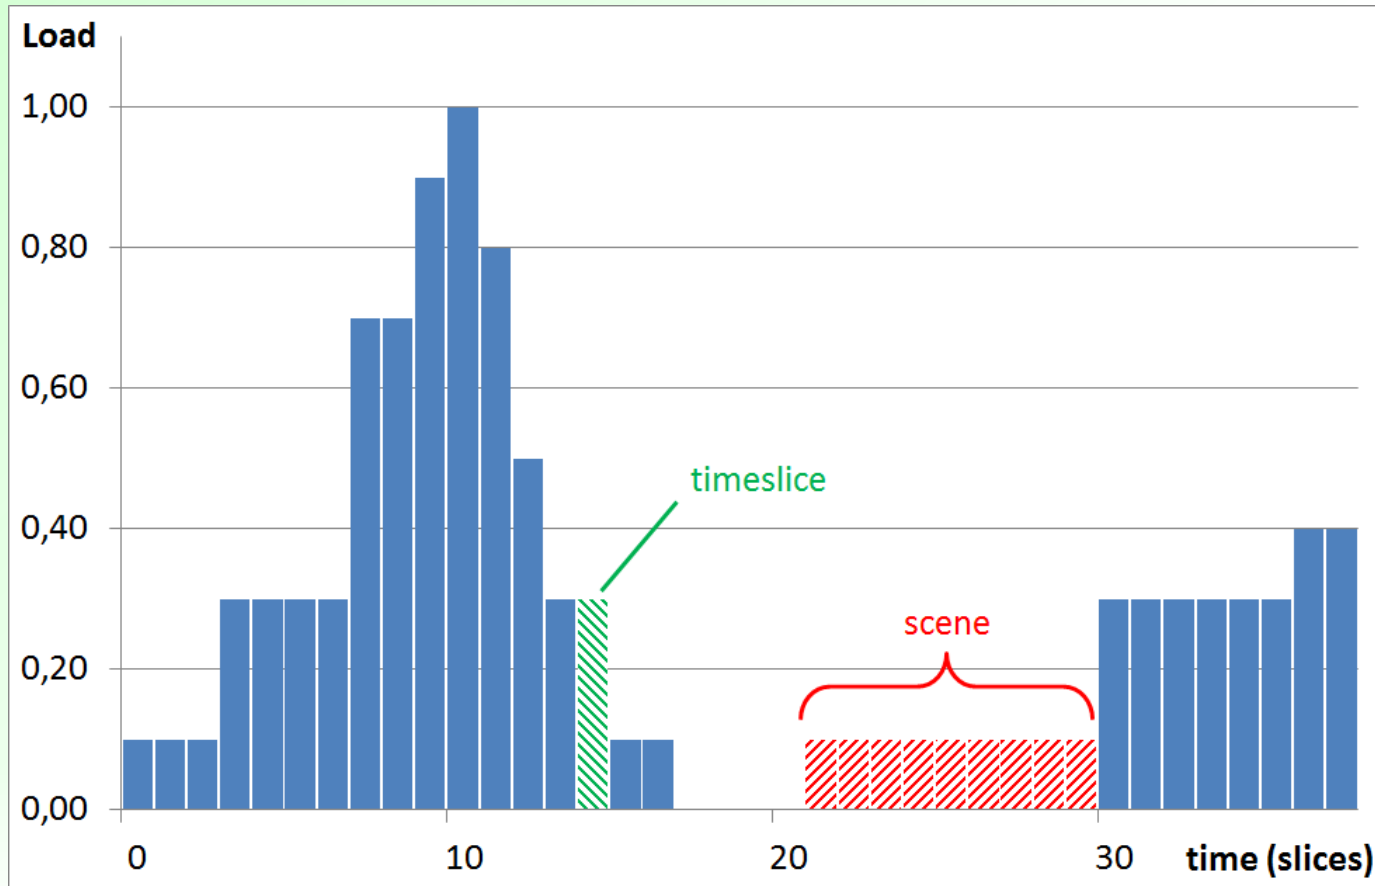
Measuring Energy Proportionality

- measure energy consumption at different load levels



Dynamic Energy Proportionality

- dynamically changing load over time



Summary

- Static Energy Proportionality
 - fixed to one load level per run
 - measure energy proportionality per level
- Dynamic Energy Proportionality
 - varying load during run
 - measure ability to adapt to load

Future work

- implement proposal
- benchmark throttling
 - TPC-C : increase think time?
 - TPC-H : decrease number of concurrent streams?
- predefined load profiles
 - for a „typical“ database server