A Power Consumption Analysis of Decision Support Systems

WOSP/SIPEW 2010
First Joint WOSP/SIPEW International Conference on Performance Engineering
San Jose, California, USA
January 28-30, 2010

Meikel Poess (Oracle Corporation), Raghunath Nambiar (Hewlett-Packard Company)

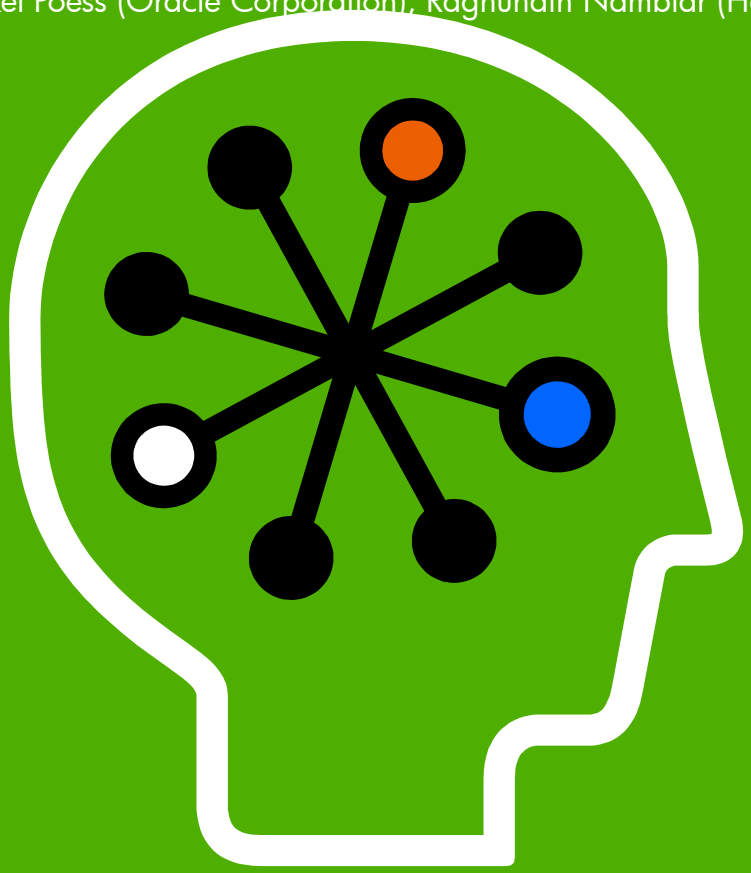

Motivation

- Energy efficiency is the top priority of today's data center mangers, because of
  - Increased demand for high performance and high capacity
    - Enterprise data warehouses double in size every three years, because business generate more data, need to store more data and retain data for a longer time
    - Petabyte data warehouses are increasingly common:
      - AT&T, ebay, Yahoo, Walmart are >5 petabytes
  - Increased energy cost
    - E.g. the disk drives alone for a 1 petabyte data warehouse would cost about $53,000 a year (assuming 300GB disk drives RAID10, 8W and $.11/kWh)

Motivation cont'

- Standard organizations such as SPEC, Neal Nelson and Associates, Green500 and SPC have responded to growing demand for energy benchmarks
- No energy benchmark for decision support systems
- TPC benchmarks do not include energy measurements yet

Motivation cont'

- Want to quantify energy consumption of large system configurations running complex workloads with a simple analytical model, which can
  - Utilize data from existing benchmarks
  - Show energy trends
  - Identify largest power consumers

Earlier Work

- *Energy cost, the key challenge of today's data centers: a power consumption analysis of TPC-C results*, VLDB2008
- ***A Power Consumption Analysis of Decision Support Systems*, WOSP/SIPEW 2010**
- *Tuning Servers, Storage and Database For Energy Efficient Data Warehouses*, to appear in ICDE 2010
- *Energy Benchmarks: A Detailed Analysis*, to appear in e-Energy 2010

TPC-H Configuration trend

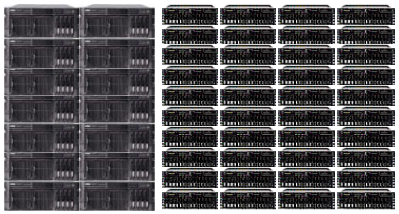

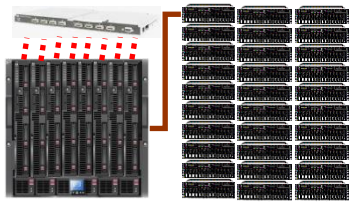

2000-2004 | 2004-2006 | 2006-2007 | 2007-2010

3K Queries/Rack | 7K Queries/Rack | 10K-40K Queries/Rack | 1Million Queries/rack

Core/Socket	1
Memory/4U	16GB
GB/disk	32GB (LFF)
Cores/Rack	40
Memory/Rack	320GB
Disk Capacity/Rack	7.5TB

Core/Socket	6
Memory/4U	1TB
GB/disk	2TB(LFF)/1TB (SFF)
Cores/Rack	1024
Memory/Rack	16TB
Disk Capacity/Rack	500TB

hp ORACLE

Typical TPC-H System

- Most common type (Type 1):
  - One or more servers with multiple CPUs
  - Small main memory to database size ration (<20%)
  - Servers are attached to an external storage subsystem
- Other types do not use external storage subsystems
  - Type 2 uses internal controllers and internal disks to host the database
  - Type 3 uses main memory to host the database

Sample TPC-H Systems

System Type	Type 1	Type 2	Type 3
Server	1 DL760G2	64 IBM x346	64 HP BL460c
QphH	4063.6	53,451.4	1,166,976
Price/Perf.	$43	$33	$5.42
Scale Factor	300	1000	1000
Processor	8x Intel Xeon MP 2.8 GHz with 2MB L2 cache	64 x 3.6 GHz Intel Xeon with 2MB L2 cache	128 x Quad-Core Intel Xeon X5450
Main Memory	16 GB	128 GB	2080 GB
Controller	5 Smart Array 5302	64 ServeRAID-7k Ultra320 SCSI controller	
External Drives	142 x 18.2 GB 15K RPM	N/A	256 146GB 10K SAS 2.5”
Internal Drives	2x 72GB 15K RPM	6x 73.4GB 15K RPM	72 x450GB 15K RPM

ORACLE

Workload Characterization and Performance Metric

- TPC-H's multi-user test
  - S users are emulated, each running 22 queries
  - Metric: $$TPC-H.Throughput@SF = \frac{S*22*3600*SF}{T_s}$$
- System vary widely in size:
  - E.g. at scale factor 100 the number of CPU cores vary between 2 and 96 and the number of disk drives varies between 4 and 344
- For this study we normalize the TPC-H Throughput metric by number of cores:

$$Normalizd.TPC-H.Throughput@SF = \frac{S*22*3600*SF}{T_s*C_C}$$

Power Consumption Model

- Power model estimates peak power consumption during steady state (multi-user test)
- Assumption: Total system power consumption equals the aggregate of the name plate power consumptions of key system components
- We consider only x86 based systems sold in the last ten years
- Power consumption model is applied to TPC-H benchmark results of the last seven years

Key Components Considered in our Power Consumption Model

Key Component	Power Consumption Range
Main Processor (CPU)	55 W to 165 W per CPU
External Disks	7.2 W to 19 W per disk
Internal Disks	7.2 W to 19 W per disk
Main System Memory	9 W per DIMM
Server Chassis	100 W + 30% of all components
Disk Enclosures	20% of all included disks

Power Consumption Model cont'

- Verified power consumption model against three reference configuration (see table before) for scale factors 100, 300 and 1000
- Model is simple but accurate enough to show power and performance per power trends
- Model is not accurate enough to size power supplies

Performance Trends SF=100

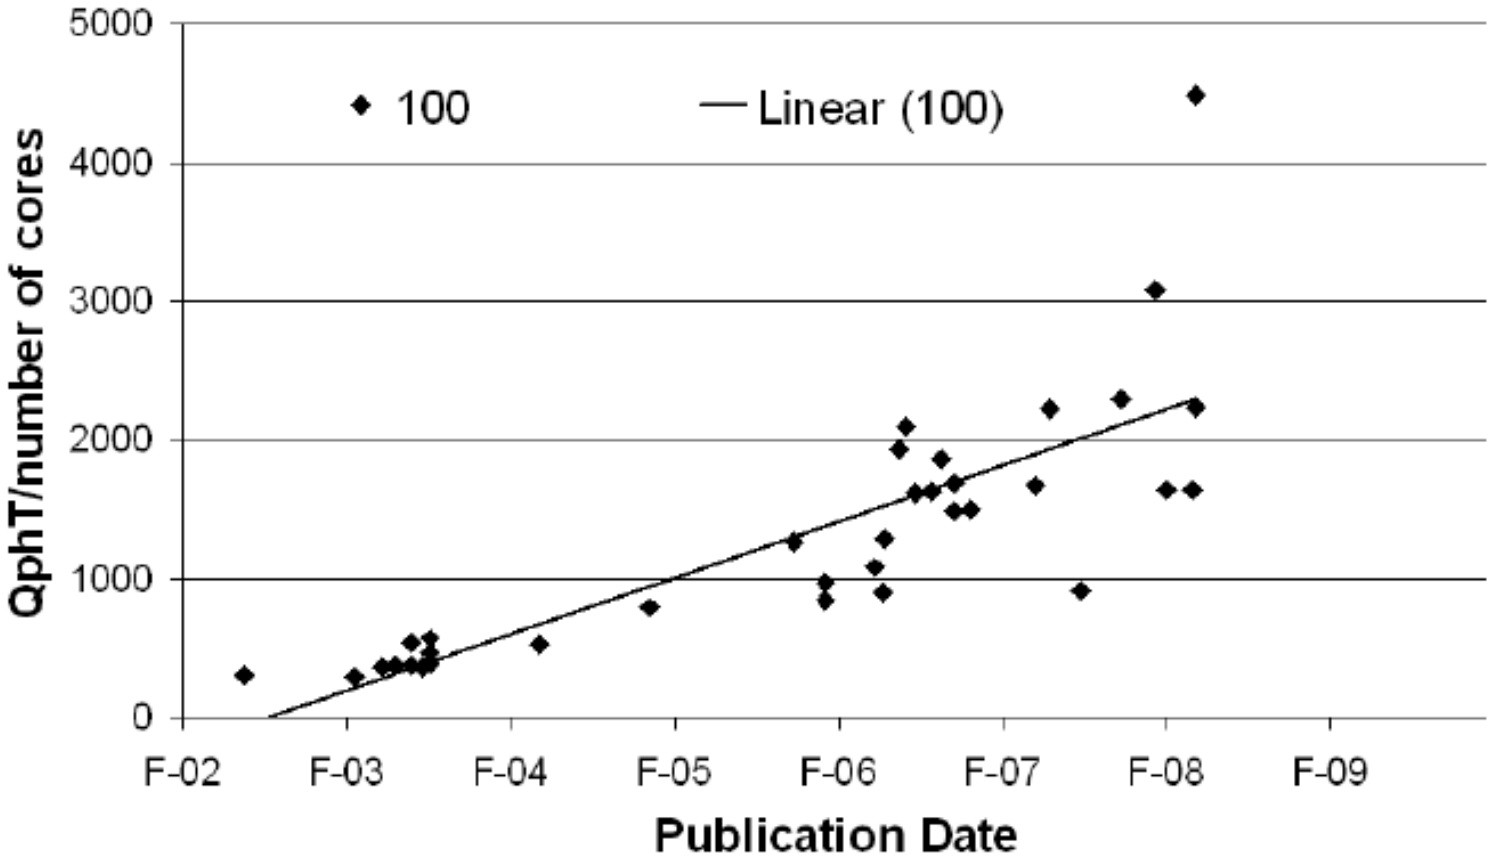

Performance Trends SF=300

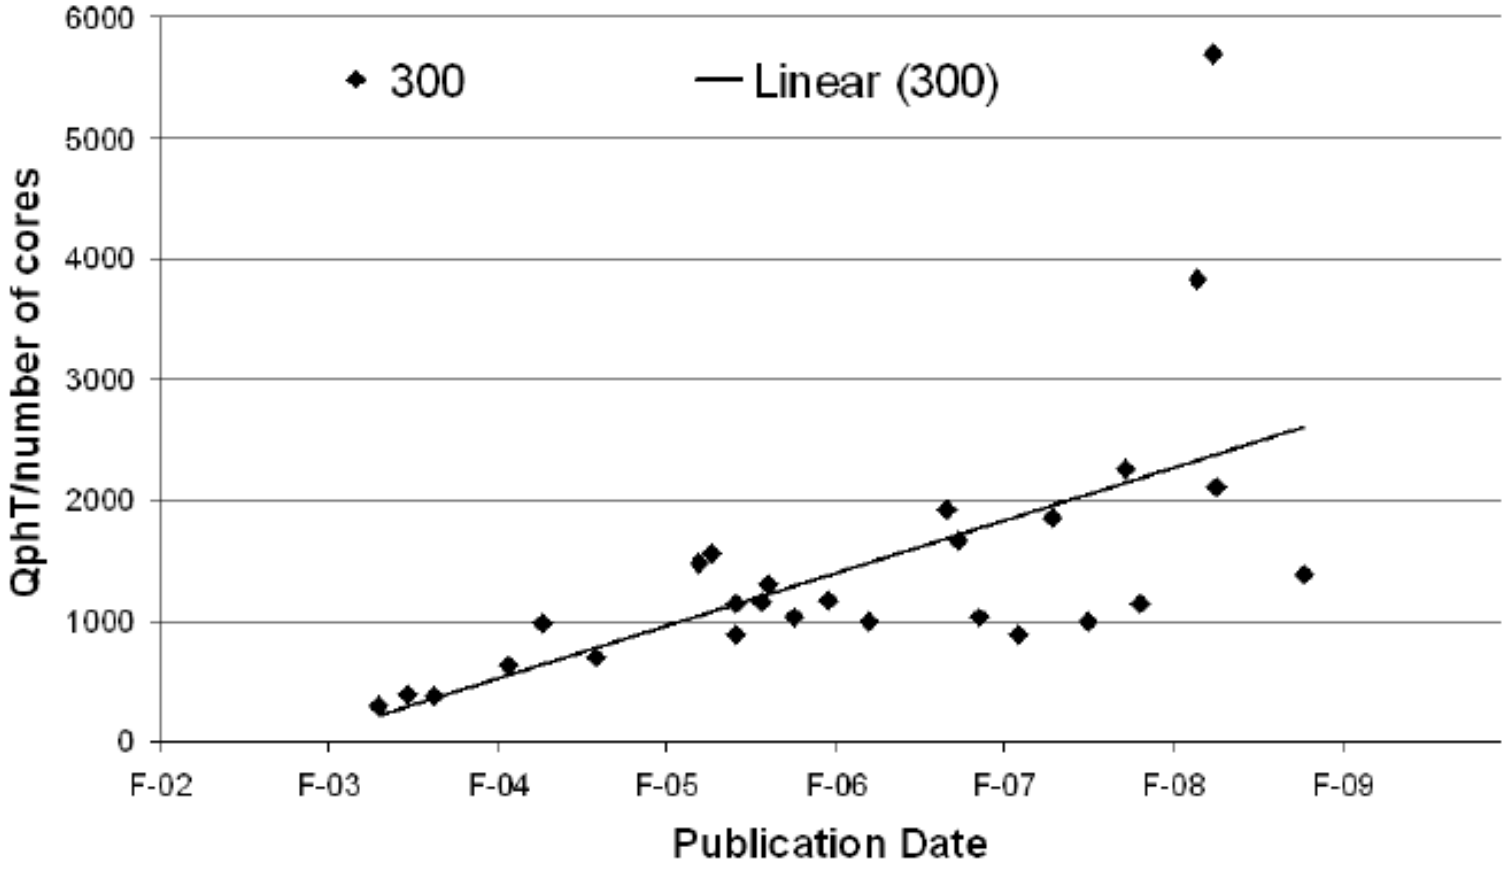

Power Consumption per Performance Trends SF=300

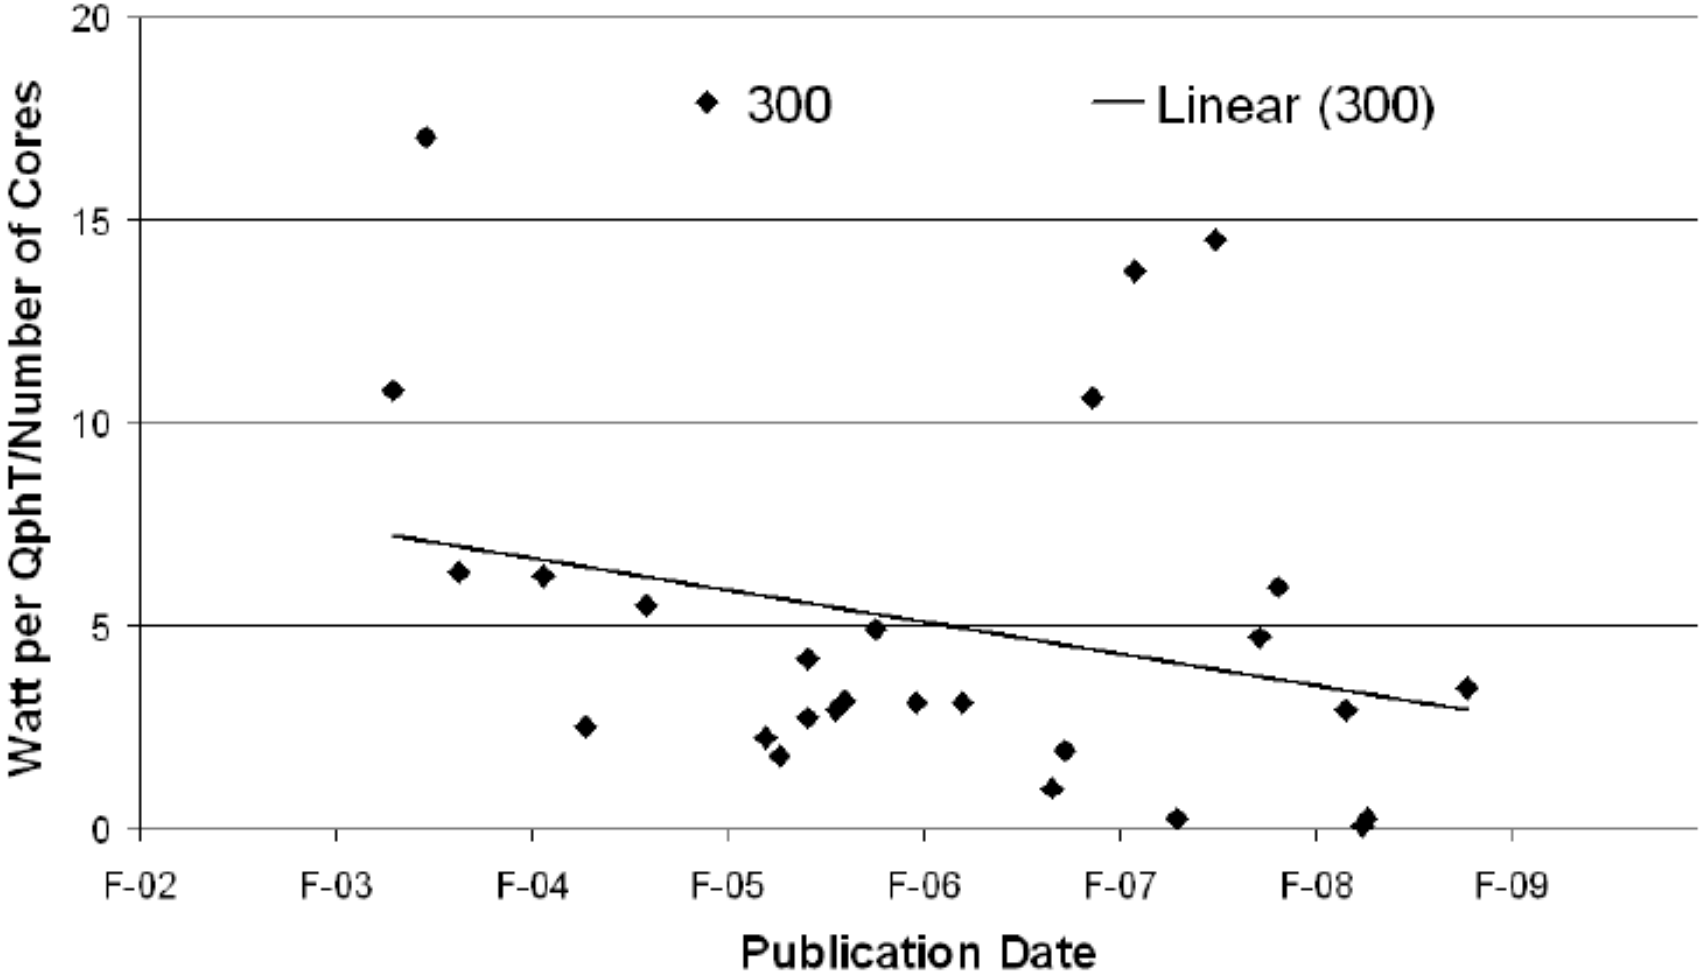

Power Consumption per Performance Trends SF=300

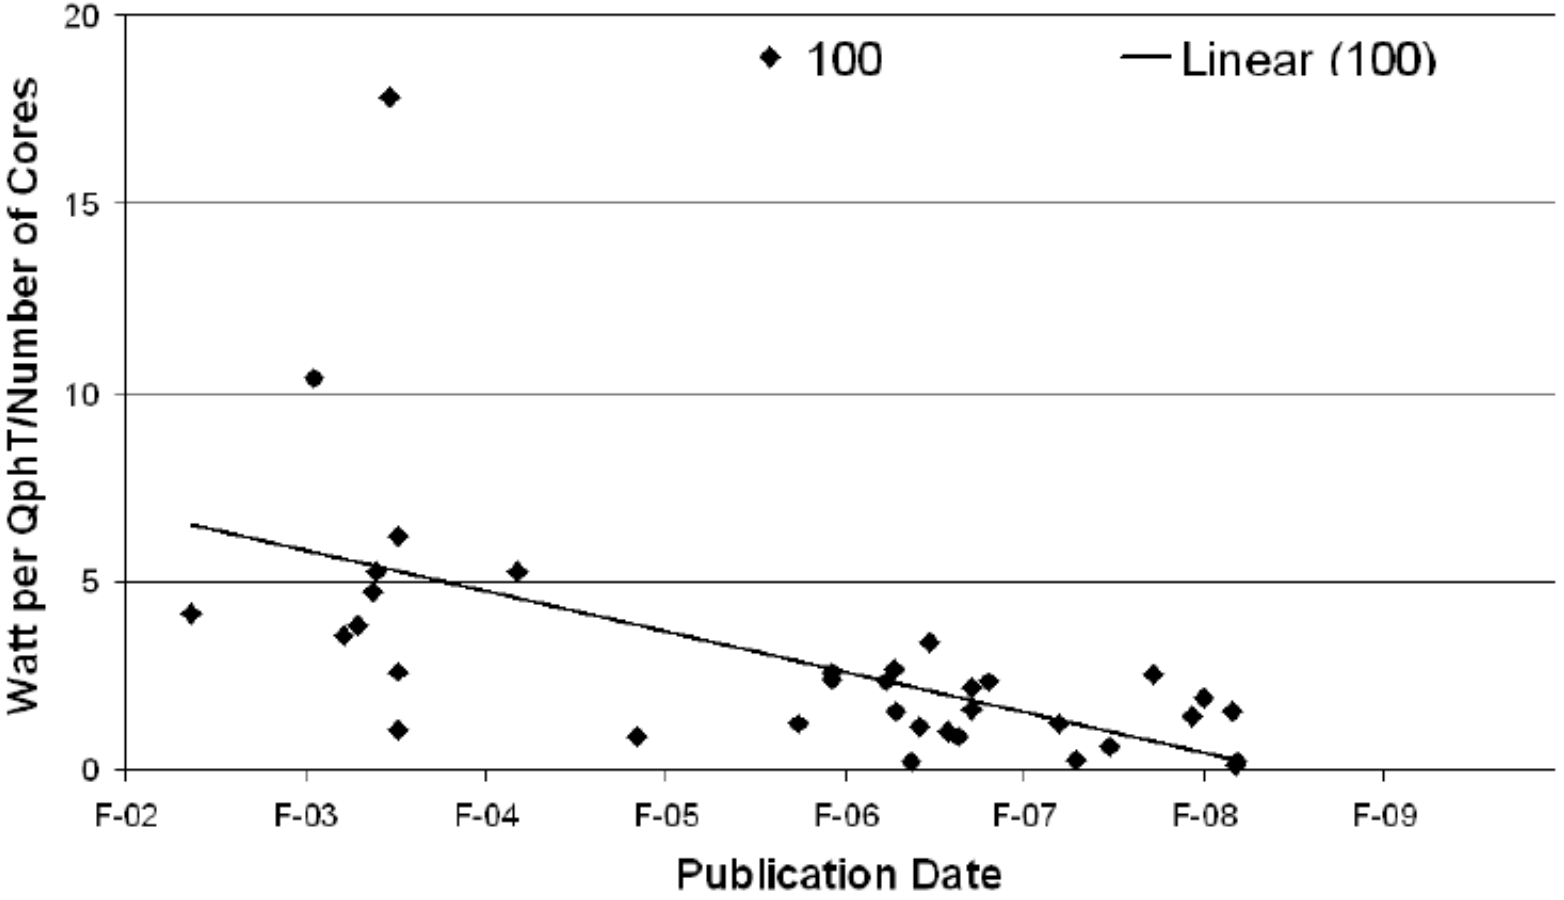

Power Consumption Distribution

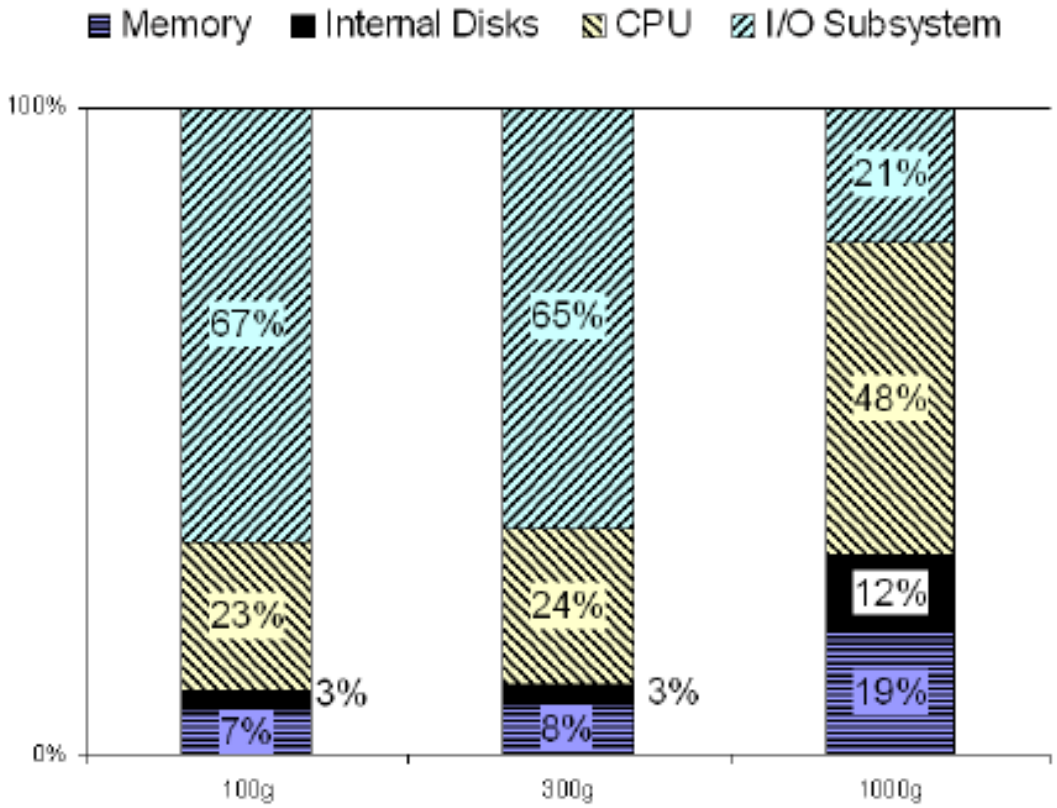

Figure 11: Average Power Consumption of Key Components Used in TPC-H Benchmarks

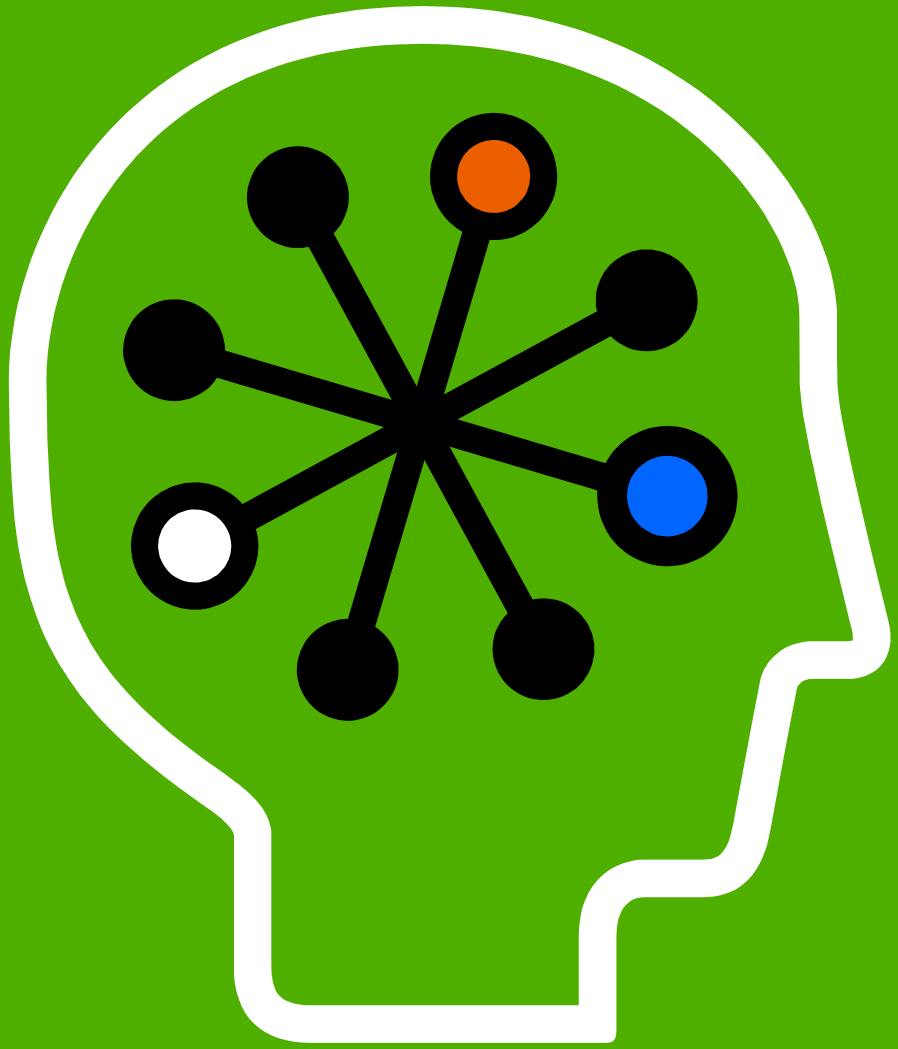

Q & A