

# A Power Consumption Analysis of Decision Support Systems

Meikel Poess  
Oracle Corporation  
500 Oracle Parkway  
Redwood Shores, CA-94065, USA  
650-633-8012  
Meikel.Poess@oracle.com

Raghunath Othayoth Nambiar  
Hewlett-Packard Company  
11445 Compaq Center Dr. W  
Houston, TX-77070, USA  
281-518-2748  
Raghu.Nambiar@hp.com

## ABSTRACT

Enterprise data warehouses have been doubling every three years, demanding high compute power and storage capacities. The industry is expected to meet such compute demands, but dealing with the dramatic increase in energy requirements will be challenging. Energy efficiency has already become the top priority for system developers and data center managers. While system vendors focus on developing energy efficient systems there is a huge demand for industry-standard workloads and processes to measure and analyze energy consumption for enterprise data warehouses. SPEC has developed a power benchmark for single servers (SPECpower\_ssj2008), but so far, no benchmark exists that measures the power consumption of large, complex systems. In this paper, we present a simple power consumption model for enterprise data warehouses based on the industry standard TPC-H benchmark. By applying our model to a subset of 7 years of TPC-H publications, we identify the most power-intensive components where research and development should focus and also analyze existing power consumption trends over time. This paper complements a similar study conducted for enterprise OLTP systems published by the same authors at VLDB 2008 and the Transaction Processing Performance Council's initiative of energy metric to its benchmarks.

## Categories and Subject Descriptors

B.8.2 [Hardware]: Performance and Reliability – Performance Analysis and Design Aids

## General Terms

Measurement, Performance, Experimentation, Standardization

## Keywords

Power and performance, energy efficiency, benchmarking, software performance testing, use of benchmarks in industry and academia, performance tuning and optimization

## 1. INTRODUCTION

Companies have realized that despite the recent drop in oil prices, developing economies will inevitably reach a level of sustained growth that will eventually cause energy prices to rise and to reach never-imagined heights in the near future. In the last decades, performance and purchase price of hardware and software

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WOSP/SIPEW'10, January 28–30, 2010, San Jose, California, USA.  
Copyright 2010 ACM 978-1-60558-563-5/10/01...\$10.00.

were the dominant concerns of data center managers. Consequently, market forces have driven down the price of hardware and software while performance has increased substantially, which in many cases led to the purchase of more systems if performance bottlenecks occurred. However, over the last few years, the cost of owning and maintaining large data centers has become a serious concern for data center managers, especially with the increase in energy cost.

Reducing power consumption is also at the top of the priority list for government agencies as they challenge data center managers and system developers to reduce power consumption. The U.S. Environmental Protection Agency (EPA) has been working with various organizations to identify ways in which energy efficiency can be measured, documented, and implemented not only in data centers but also in the equipment they house.

Other organizations that have responded to the growing demand for an energy benchmark are Standard Performance Evaluation Corporation (SPEC), Neal Nelson and Associates, Green500, and Storage Performance Council (SPC). SPEC, a non-profit group of computer vendors, system integrators, universities, research organizations, publishers, and consultants, developed SPECpower\_ssj2008.

The benchmarks developed by the above consortia are very significant milestones towards designing and deploying standards to measure energy efficiency in servers and storage subsystems. However, they do not address the need for measuring power consumption of complex data warehouse (DW) systems that can be characterized by a combination of compute and I/O intensive operations. According to the Winter Corporation Survey of 2008, the largest data warehouses triple every three years. There is no reason to believe that this trend is going to change any time soon, challenging the experts to develop methodologies to measure the energy consumption of very large configurations. Market surveys show that large, commercial data warehouses have crossed the petabyte boundary. At such scales, energy consumption will be significant. For instance, in order to keep all data of a 1-petabyte DW in a RAID10 protected storage system requires 3744 300GB disk drives. Assuming an 8-watt average power consumption per disk (using energy efficient disks and estimating idle times) the storage subsystem alone (no cooling cost included) would consume about 262800kWh a year.

Even though there are benchmarks to measure system level power consumption, none exists to provide power consumption on large DWs. With this paper we attempt to fill this gap by defining a power consumption model for enterprise data warehouses based on industry standard TPC-H [12] benchmark. The motivation for choosing the TPC-H benchmark are numerous: TPC-H is a widely adapted benchmark standard; TPC-H publications have been tracking technology changes and enhancements in the data ware-

house space; there is a large body of published results; data is available on all major hardware and database platforms; most system vendors, database vendors, and end users can relate the TPC-H workload to their setups.

The power consumption model introduced in this paper is based on data that is readily available in the TPC-H full disclosure reports of published benchmarks. The model was verified by measuring the power consumption of three fully scaled TPC-H systems, including servers, storage, and network. By applying this model to a subset of 7 years of TPC-H results (78 published results) that used x86 processors at scale factors 100, 300, and 1000, the paper identifies the most power intensive components and demonstrates existing power consumption trends over time.

The remainder of this paper is organized as follows: Section 2 characterizes those parts of the TPC-H workload that are pertinent in understanding our power consumption model. Section 3 presents the power consumption model and its verification. Section 4 applies the power consumption model to a subset of the TPC-H result set to perform various trend analyses.

## 2. WORKLOAD CHARACTERIZATION AND PERFORMANCE METRIC

This section serves as a general introduction to the workload characteristics of decision support systems pertinent to the subsequent power consumption model. It continues with the definition of the exact workload and performance metric that is assumed by the power consumption model.

Generally, decision support workloads can be divided into three distinct types of typically parallel operations: Initial load, incremental load, and queries. These types of workloads can be run in single- and multi-user modes. The single-user mode stresses a system’s ability to parallelize operations to answer a given request in the least amount of time, such as overnight batch job processing. The multi-user mode stresses the system’s ability to schedule concurrent requests from multiple users to increase overall system throughput. Because of these differences, systems are usually tuned differently for single and multi-users modes.

Each TPC-H result is obtained on a database with a specific size indicated by the scale factor (SF). The scale factor in GB equals the raw data outside the database, for example, SF=100 means that the sum of all base tables equals 100 GB. The TPC rules prohibit comparing TPC-H results between scale factors. The primary performance metric in TPC-H is the composite performance metric (QphH). It equally weights the contribution of the single-user and the multi-user. In addition, it defines multiple numerical quantities, such as TPC-H Throughput, a quantity to measure the multi-user performance of a system when running S concurrent users, each executing 22 queries. It is defined as the ratio of the total number of queries (S\*22) executed over the length of the multi-user interval  $T_s$ , and must be computed as:

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

Within a particular scale factor, TPC-H results vary widely because of the different system sizes. For instance, at scale factor 100 the number of processor cores installed varies between 2 and 96 while the number of disk drives varies between 4 and 344. This makes analyzing performance trends impossible. Therefore, for this study we use a TPC-H throughput numerical quantity normalized by the number of cores ( $C_c$ ):

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

Note that this quantity normalizes the multi-user test to scale factor and number of concurrent users. It is not a TPC-H metric.

### 2.1 Typical TPC-H Systems

Following the strict rules of the TPC-H specification, a benchmarked system is composed of a driver that submits queries to a system under test (SUT). The SUT executes these queries and replies to the driver, which resides on the SUT’s hardware and software.

The most traditional type of SUT comprises of one or more servers with multiple CPUs, small main memory to database size ratio, typically less than 20 percent, and one or more HBAs (controllers) (Type1). These systems are either directly attached to a storage subsystem or, in case of clustered systems, funnelled through one or multiple switches. The disk subsystem is comprised of multiple disk enclosures and usually many hard disk drives. Seventy-one percent of our sample set uses this type of configuration.

The second most common type of system used in TPC-H publications uses only internal controllers and disks to host the database (Type 2) [18]. The servers look similar to the first configuration type. Twenty-three percent of our sample set uses this type of configuration.

In the last few years, benchmark publications emerged that use servers with a main memory to database size ratio of more than 50 percent. These systems use no or very small storage systems. Six percent of our sample set fall in this category (Type 3) [19].

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

Figure 1: Sample TPC-H benchmark publications

## 3. POWER CONSUMPTION MODEL

The previous section demonstrated how typical TPC-H systems are configured. Estimating their power consumption is difficult because of the assortment of components that are involved and the lack of power measurements of individual components in large-scale deployments. This section develops a simplified power consumption model that can be applied to any published TPC-H result and representative data warehouse systems. The model complements recent research in this area [13][4]. It is meant to estimate peak power consumption of the above systems. It is not suitable for calculating an accurate power consumption number for system provisioning purposes. But it is sufficient to analyze historic results for trend analysis and to identify most power-intensive components of a decision support system.

The proposed power consumption model is based on the assumption that the peak power consumption of an entire system during a multi-user run can be derived from the aggregate of the nameplate

power consumptions of individual components. We include the following key components into our model:

1. Main processor (CPU)
2. External disks
3. Internal disks
4. Main system memory
5. Server chassis
6. Disk enclosures

We differentiate between internal and external disks because internal disks typically have lower performance characteristics, therefore, lower power consumption compared to external disks. For each of the components listed above we determine its peak power consumption as illustrated in the next sections.

Table 1 shows the peak power consumption of selected CPUs as obtained their manufacturer’s specifications [1][8]. The peak power consumption is depicted as Thermal Design Power (TDP). Please note that the list of CPUs in Table 1 is only a subset of the CPUs that have been sold over the last 10 years. The subset includes data for x86 based CPU types considered in this study.

Processor Description	TDP [W]
AMD 8220SE- 2.8 GHz	93
AMD 8220SE 2.8 GHz	93
AMD DC 8220SE- 2.8 GHz	93
AMD Opteron - 2.0 GHz	89
AMD Opteron - 2.2 GHz	85
AMD Opteron - 2.4 GHz	85
AMD Opteron - 2.6 GHz	93
AMD Opteron - 2.8 GHz	93
AMD Opteron - 3.0 GHz	92.6
AMD Opteron - 3.20 GHz	119
AMD Opteron 2.2GHz Dual Core - 2.2 GHz	93
AMD Opteron 252 - 2.6 GHz	93
AMD Opteron Dual Core 1 MB L2-2.4 GHz	95
Intel DC Itanium2 Processor 9050-1.6 GHz	130
Intel Dual-Core Itanium2 1.6GHz	130
Intel Pentium III Xeon - 900 MHz	50
Intel Pentium Xeon MP - 1.6 GHz	55
Intel Xeon 7041 - 3.0 GHz	165
Intel Xeon 7140 3.4GHz	150
Intel Xeon 7350 2.93GHz	130
Intel Xeon MP - 1.6 GHz	55
Intel Xeon MP - 2.0 GHz	57
Intel Xeon MP - 2.7 GHz	80
Intel Xeon MP - 2.8 GHz	72
Intel Xeon MP - 3.0 GHz	85

Table 1: Thermal Design Power (TDP) of considered CPUs

We approximate the peak power consumption of main memory with 9 watts per DIMM as was done in other studies, see [4] [9].

Peak power consumption levels of disk drives vary widely with the disk’s form factor, size, and rotational speed. Table 2 summarizes the peak power consumption of those disk drives that are considered in this study. They are obtained from manufacturers web sites [22][23]. (FF=Form Factor)

FF=2.5 RPM=10K		FF=2.5 RPM=15K		FF=3.5 RPM=7.2K		FF=3.5 RPM=10K		FF=3.5 RPM=15K	
[GB]	[W]	[GB]	[W]	[GB]	[W]	[GB]	[W]	[GB]	[W]
36	17	36	10.0	240	11.35	9	9.7	18	13.2
36	7.2	36	9.2	465	13	9	10.0	18	10.0
72	8.4	72	14.2			36	12.5	18.2	9.7
73	10.5	72	9.2			36	10.0	32	10.0
146	10.0	146	14.0			72	12.6	36	14.5
146	9.0					73	11.0	36	15
						146	14.2	72	13.2
						146	11.4	73	16.2
						160	12.8	146	14.2
						250	11.35	146	19.0
						300	16.4	300	17.6
								300	14.4

Table 2: Disk Peak Power Consumption

### 3.1 Server Chassis Peak Power Consumption

Recent studies [3][24] suggest that the power consumption of the server chassis and its infrastructure (fan, power supply, and so forth) can be expressed as a percentage of the nameplate power consumption of its main components. In respect to TPC-H systems, the main components are CPU, main memory, and internal disks. Since the server chassis and its infrastructure (fan, power supply, and so forth) are sized according to its components (CPU, memory), we express its power consumption as 30 percent of the power consumption of its components plus a fixed overhead of 100 watts.

### 3.2 Disk Enclosure Peak Power Consumption

As in the server chassis case, we approximate the power consumption of the disks enclosures with 20 percent power overhead of the aggregate power consumption of all external disks.

### 3.3 Power Consumption Estimates

The TPC-H server systems considered in this study can be conceptualized with the following components:

Component	Count	Power Consumption [W]
CPU per server	$C_C$	$P_C \in [55,165]$ see Table 1
Memory DIMM per server	$C_M$	9
Internal Disks per sever	$C_{DI}$	$P_{DI} \in [7.2,19]$ see Table 2
Number of servers	$C_S$	$P_S$ (calculated below)

Figure 2: Number of Database Server Components and their Power Consumptions

The power consumption of a database server ( $P_{DS}$ ) can be broken down into the above components plus the overhead of the chassis (30 percent plus 100 Watts). The power consumption of the components can be estimated by aggregating the peak power consumption of all components:

$$C_C * P_C + 9 * C_M + C_{DI} * P_D \quad E_1$$

The power consumption of the chassis is:

$$(C_C * P_C + 9 * C_M + C_{DI} * P_D) * 0.3 + 100 \quad E_2$$

Hence, power consumption of an entire server can be estimated as

$$P_S = (C_C * P_C + 9 * C_M + C_{DI} * P_D) * 1.3 + 100 \quad E_3$$

Each I/O subsystem consists of the following components:

Component	Count	Power Consumption [W]
Number of Enclosures	$C_E$	$P_E$ (calculated below)
External disks per enclosure	$C_{DE}$	$P_{DE} \in [7.2,19]$ see Table 2

Figure 3: Number of IO Subsystem Components and their Power Consumptions

The power consumption of the I/O subsystem ( $P_{IO}$ ) can be broken down into the above components plus the overhead of the enclosure (20 percent).

$$P_{IO} = C_E * C_{DE} * P_{DE} * 1.2 \quad E_4$$

Hence, the total power consumption of the entire system can be estimated as the sum of the server and I/O subsystems.

$$P = P_S + P_{IO} \quad E_5$$

### 3.4 Verifying the Power Consumption Model with Sample Configurations

Three different systems were used to verify the power consumption model. System A follows the traditional approach of a strong I/O subsystem. It consists of one HP ProLiant ML370 G5 server [6] as the database server and four HP StorageWorks 70 Modular

Smart Array enclosures [7] as the storage subsystem. The database server is equipped with two quad core CPUs (Xeon® CPU X5450 @ 3.00GHz) and 16 Gigabytes (GB) of main memory (4 DIMMS of 4GB each). The operating system is stored on two internal 146GB 10K rpm 2.5 inch SAS disk drives. Each of the four HP StorageWorks 70 Modular Smart Array enclosures is equipped with 25 36GB 15K rpm 2.5 inch SAS disk drives. System B uses internal storage only instead of external enclosures to host the database. The server consists of one HP ProLiant ML370 G5 server [7] with two internal I/O controllers, each attached to eight 146GB 10K rpm 2.5 inch SAS disk drives. The server is equipped with two quad core CPUs (Intel® Xeon® CPU X5450 @ 3.00GHz) with 16 Gigabytes (GB) of main memory (4 DIMMS of 4GB each), Two internal disks are used for the operating system. System C consists of 64 HP ProLiant BL 460c server blades, each with two quad-core Intel® Xeon® X5450 Processors @ 3.00 GHz, 32GB of main memory (8 DIMMS of 4GB each), and two internal disk drives (146GB 10K rpm SAS 2.5 inch). For each of our test systems we use the same power measurement methodology, similar to that used in [13]

	Power Consumption [W]								
	System A			System B			System C		
	Power Model	Measurement	Diff [%]	Power Model	Measurement	Diff [%]	Power Model	Measurement	Diff [%]
DB-Server	433	405	7	511	395	29	27366	23320	17
I/O Sub-system	1180	891	32	N/A	N/A	N/A	N/A	N/A	N/A
<b>Total</b>	<b>1613</b>	<b>1296</b>	<b>24</b>	<b>511</b>	<b>395</b>	<b>29</b>	<b>27366</b>	<b>23320</b>	<b>17</b>

Figure 4: Comparison of the Power Consumption Model and Power Measurements of Three TPC-H Systems

Using the measurement methodology presented in the previous section, Figure 4 shows the power consumption of the database server and the storage subsystem of systems A, B, and C, both estimated by the power model and measured with power meters. For the database server and the storage subsystem, Figure 4 shows the percentage difference between the estimate and the measurement. The paper, “Power Provisioning for a Warehouse-sized Computer” [4] refers to this as the difference between the nameplate value and actual peak power. The power consumption difference between the model and the measurements of the three systems is between 17 and 29 percent – 23 percent on average. The paper [5], which applied a different workload on smaller systems, shows that the difference between the nameplate model and actual peak power consumption can be up to 30 percent. Using the TPC-C workload, we calculated the difference between modelled and measured power consumption of database clients and servers to be between 10 and 26 percent (15 percent on average) [13]. The slightly higher average difference between measured and estimated power consumption, when compared to [5] and [13], can be explained with the oscillating nature of the TPC-H workload (see **Error! Reference source not found.**). The oscillating behaviour leads to periods in which resources are less utilized and therefore draw less power, which led to a larger gap between nameplate power consumption and estimated power consumption. We therefore calibrate our power model to the TPC-H workload by abating our power model numbers by 15 percent.

## 4. HISTORIC TREND ANALYSIS

The power consumption model developed in the previous section allows us to estimate the power consumption of any published TPC-H result because all information necessary for the power consumption model is readily available in the TPC-H Full Disclosure Report (FDR). In this section, we apply the power consump-

tion model to a subset of all available results. We then analyze the performance and total system power consumption trends. We finish this section with an analysis of components that consume the most power in TPC-H systems.

### 4.1 Performance Trends

In this section we study the multi-user performance trends for each of the three scale factors: 100, 300, and 1000. We use the normalized TPC-H Throughput@SF metric developed in Section 2 to compare results within identical scale factors. Figure 5 shows the performance trend for scale factor 100, Figure 6 shows the performance trend for scale factor 300, and Figure 7 shows the performance trend for scale factor 1000. For each scale factor we show the actual normalized TPC-H Throughput numbers (diamond shaped graph) and the linear trend line (solid line). In the beginning of 2002, systems achieved about 300 QphT per processor core at scale factor 100. The current performance leader achieved about 4500 QphT per core in 2008 at scale factor 100. This is a 15x improvement in performance over 6 years. This positive trend can also be observed by ignoring the extreme cases. The trend line shows an increase of about 388 QphT per core per year, essentially doubling performance every year.

Similarly, at scale factor 300, systems achieved about 288 QphT per core in 2003, while the current leader in this scale factor achieved 5691 QphT per core in 2008. This is a 20x performance improvement over 5 years. As in the scale factor 100 case, the trend line shows that performance doubles every year.




Figure 5: Performance Trend SF 100

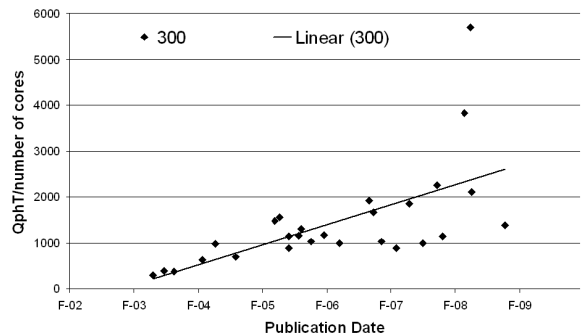


Figure 6: Performance Trend SF 300

At scale factor 1000, systems achieved 289 QphT per core in 2003, while the current leader achieved 3793 QphT per core in 2008. This is a 13x improvement in performance over 5 years. Comparing the trend line of the 1000 scale factor case with the 100 and 300 scale factor cases shows a stronger tendency towards higher performance. The trend line at scale factor 1000 indicates that performance doubles every six months.




Figure 7: Performance Trend SF 1000

These results are quite different from the results obtained in [13]. This study of TPC-C results found that performance over the last 8 years is roughly in line with Moore’s law [10]. As Figure 5, Figure 6 and Figure 7 indicate, TPC-H results undoubtedly outperformed the hardware performance improvement predicted by Moore’s Law.

### 4.2 Power Consumption Trends

Now we analyze the power trend of our sample TPC-H result set. For each result we compute the total peak power consumption according to the calibrated power consumption model we developed in Section 3.3. We are interested in how the normalized Throughput performance per system and how power consumption evolved over time.

The following figures (Figure 8, Figure 9, and Figure 10) show the system power consumption per normalized Throughput performance (QphT/number of cores) for scale factors 100, 300, and 1000. At scale factor 100, power consumption decreases from a peak of 10.4 W to 0.13W, an 80x reduction. However, the 10.4 W power consumption looks like an extreme data point. The trend line shows that power consumption per normalized Throughput performance decreases at a rate of 1.4W per year.

Figure 9 shows the power consumption over the last six years of scale factor 300 results. The trend line shows that performance decreased at a rate of about 0.8W per year. This is less than the 1.4 rate drop in power consumption at scale factor 100. Please note, however, that the 300 scale factor case has more outliers compared to the scale factor 100 case.




Figure 8: Power Consumption Trend SF 100

Figure 10 shows the power consumption per normalized QphT of 1000 scale factor results. Power consumption dropped from about

15W to 0.5W per normalized QphT from 2004 to 2009. This is a 30x decrease, the largest we have seen in this study.




Figure 9: Power Consumption Trend SF 300

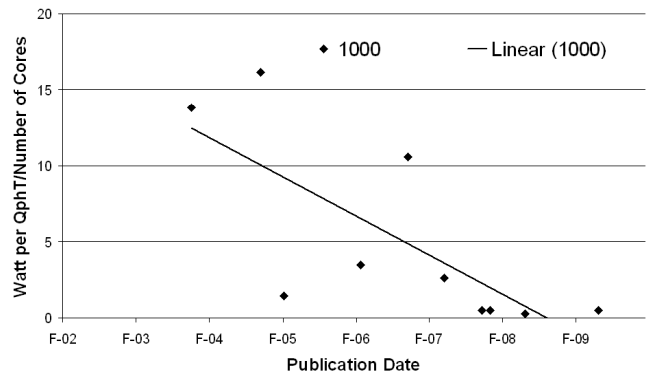


Figure 10: Power Consumption Trend SF 1000

### 4.3 Power Consumption Distribution

We will now look at the power distribution of the key components of a typical TPC-H system. We compute the percentage of the power consumed by each component listed above by averaging the power consumption numbers of all 78 results. As in our previous analysis, we distinguish between scale factors 100, 300, and 1000. Figure 11 shows the power consumption of each component as a percentage of the overall power consumption.

The power consumption distribution across the different components is very similar for scale factors 100 and 300. Power consumption of systems used to publish results for scale factors 100 and 300 is dominated by the I/O subsystem. For both scale factors, about 66 percent of all power is consumed in the I/O subsystem. The second-largest power consumers are the CPUs. As indicated in Table 1, power consumption (TDP) per CPU varies from 50W to 165W. Scale factor 100 and 300 results indicate that about 23% of all power is consumed by the CPUs. The third largest power consumer is the memory. About 7% of all power for scale factor 100 and 300 results is consumed by memory. Internal disks play a smaller role for scale factors 100 and 300. Only 3% of all power is consumed by internal disks. Since they are quite different from the smaller scale factors, we discuss the results for scale factor 1000 separately. The governing power consumer in the 1000 scale factor category is the CPU rather than the I/O subsystem. Forty-eight percent of all power is consumed by the CPUs, while 21% of power is consumed in the I/O subsystem. Memory takes the third spot with 19% and internal disks take the last spot with 12%.

The difference between the power distribution of the smaller scale factors of 100 and 300 compared to the larger scale factor of 1000 seems to be related to a recent development using large memory systems rather than large I/O systems. This development has two significant effects: higher performance and lower power consumption. Both effects are beneficial for customers as they can achieve their performance goal and still spend less on system power.

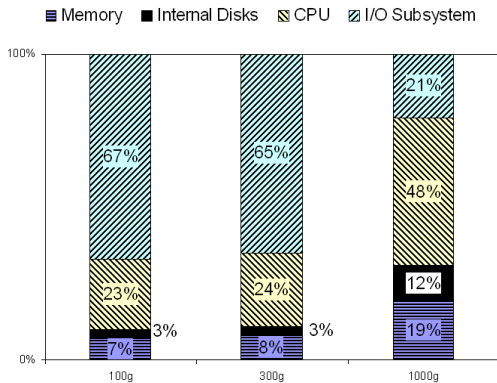


Figure 11: Average Power Consumption of Key Components Used in TPC-H Benchmarks

## 5. CONCLUSION

In this paper we are building upon the research that we have published at VLDB 2009 [13] for TPC-C systems (OLTP) to develop a power consumption model for decision support systems. Specifically, this paper introduced a power consumption estimation model for TPC-H benchmarks. The accuracy of this model was verified by measuring power consumption of systems tuned for the TPC-H workload. The model was applied to a large subset of TPC-H results to show performance and power performance trends.

The paper further identifies the components that consume the most power in TPC-H systems depending on the database size. We hope that the model we have developed in this paper will help customers evaluate systems for which no power consumption tests have been conducted.

## 6. Acknowledgement

The authors would like to acknowledge Mike Nikolaiev and Ray Glasstone for their guidance, Paul Cao for power measurement methodologies, and David Miller for his comments and feedback.

## 7. REFERENCES

[1] AMD Opteron family of processors: <http://www.amd.com>

[2] Benoît Dageville, Mohamed Zaït: SQL Memory Management in Oracle9i. VLDB 2002: 962-973

[3] Dimitris Economou, Suzanne Rivoire, Christos Kozyrakis, and Parthasarathy Ranganathan: Full-system Power Analysis and Modeling for Server Environments, Workshop on Modeling, Benchmarking, and Simulation (MoBS), June 2006

[4] Fan, Xiaobo, Wolf-Dietrich Weber, and Luiz André Barroso. 2007. Power Provisioning for a Warehouse-sized Computer. Proceedings of the 34th International Symposium on Computer Architecture in San Diego, CA. Association for Computing Machinery, ISCA '07.

[5] Fan, Xiaobo, Wolf-Dietrich Weber, and Luiz André Barroso. 2007. Power Provisioning for a Warehouse-sized Computer. Proceedings of the 34th International Symposium on Computer Architecture in San Diego, CA. Association for Computing Machinery, ISCA '07. [http://labs.google.com/papers/power\\_provisioning.pdf](http://labs.google.com/papers/power_provisioning.pdf)

[6] HP ProLiant ML370 G5 Server series specifications: <http://h10010.www1.hp.com/wwpc/us/en/sm/WF06a/15351-15351-241434-241646-241477-1121474.html>

[7] HP StorageWorks 70 Modular Smart Array Specifications: [http://h18004.www1.hp.com/storage/disk\\_storage/msa\\_disksarrays/drive\\_enclosures/msa70/specs.html](http://h18004.www1.hp.com/storage/disk_storage/msa_disksarrays/drive_enclosures/msa70/specs.html)

[8] Intel family of processors: <http://www.intel.com>

[9] Micron Technology, Inc. Calculating Memory System Power for DDR. <http://download.micron.com/pdf/technotes/DDR/TN4603.pdf>

[10] Moore, G.E.: *Cramming More Components Onto Integrated Circuits*, Proceedings of the IEEE, Volume 86, Issue 1, Jan 1998 Page(s):82 - 85

[11] NEAL NELSON BENCHMARK LABORATORY <http://www.worlds-fastest.com/wf3010.html>

[12] Poess M, Floyd C.: New TPC Benchmarks for Decision Support and Web Commerce. SIGMOD 29(4) 64-71 (2000)

[13] Poess, M, Nambiar, R.O., Energy cost, the key challenge of today's data centers: a power consumption analysis of TPC-C results. PVLDB 1(2): 1229-1240 (2008)

[14] Poess, M., Nambiar, R.O.: *Large Scale Data Warehouses on Grid: Oracle Database 10g and HP ProLiant Systems*. VLDB 2005: 1055-1066

[15] Poess, Nambiar, R. O., Walrath D.: Why You Should Run TPC-DS: A Workload Analysis. VLDB 2007: 1138-1149

[16] Poess, Smith B., Kollár L, Larson P.A.: TPC-DS, taking decision support benchmarking to the next level. SIGMOD Conference 2002: 582-587

[17] Power Measurement Tutorial for the Green500 List <http://www.green500.org/docs/tutorials/tutorial.pdf>

[18] Publication 105021401 IBM eServer xSeries 346 [http://www.tpc.org/tpch/results/tpch\\_result\\_detail.asp?id=105021401](http://www.tpc.org/tpch/results/tpch_result_detail.asp?id=105021401)

[19] Publication 109060301 HP BladeSystem c-Class 128P RAC [http://www.tpc.org/tpch/results/tpch\\_result\\_detail.asp?id=109060301](http://www.tpc.org/tpch/results/tpch_result_detail.asp?id=109060301)

[20] Stephens J.M., Poess, M.: MUDD: a multi-dimensional data generator. WOSP 2004: 104-109

[21] SPEC – Power and Performance. [http://www.spec.org/power\\_ssj2008/docs/SPECpower\\_ssj2008-Design\\_overview.pdf](http://www.spec.org/power_ssj2008/docs/SPECpower_ssj2008-Design_overview.pdf)

[22] Specifications of disk drive from Hitachi: <http://www.hitachigst.com/tech/techlib.nsf/techdocs/>

[23] Specifications of disk drives form Seagate <http://www.seagate.com/staticfiles/support/disc/manuals/>, <http://www.seagate.com/docs/pdf/datasheet/>

[24] Suzanne Rivoire, Mehul Shah, Parthasarathy Ranganathan, Christos Kozyrakis, Justin Meza: *Modeling and Metrology Challenges for Enterprise Power Management*, IEEE Computer, December 2007