Energy Cost, The Key Challenge of Today's Data Centers: A Power Consumption Analysis of TPC-C Results

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

Historically, performance and price-performance of computer systems have been the key purchasing arguments for customers. With rising energy costs and increasing power use due to the ever-growing demand for computing power (servers, storage, networks), electricity bills have become a significant expense for today's data centers. In the very near future, energy efficiency is expected to be one of the key purchasing arguments. Some performance organizations, such as SPEC, have developed power benchmarks for single servers (SPECpower_ssj2008), but so far, no benchmark exists that measures the power consumption of transaction processing systems. In this paper, we develop a power consumption model based on data readily available in the TPC-C full disclosure report of published benchmarks. We verify our model with measurements taken from three fully scaled and optimized TPC-C configurations including client (middle-tier) systems, database server, and storage subsystem. By applying this model to a subset of 7 years of TPC-C results, we identify the most power-intensive components and demonstrate the existing power consumption trends over time. Assuming similar trends in the future, the hardware enhancements alone will not be able to satisfy the demand for energy efficiency. In its outlook, this paper looks at potential hardware and software enhancements to meet the energy efficiency demands of future systems. Realizing the importance of energy efficiency, the Transaction Processing Performance Council (TPC) has formed a working group to look into adding energy efficiency metrics to all its benchmarks. This paper is expected to complement this initiative.

1. INTRODUCTION

Historically, performance and price-performance of information systems have been the key purchasing arguments for customers. With rising electricity costs and steep increases in electricity use due to the ever-growing demand for computing power, electricity bills have become a significant expense for today's data centers. This opinion has been echoed by many independent IT analysts as well as sound industry research.

Gartner Group says energy costs may increase from 10% of the IT budget today to over 50% in the next few years [7]. Figure 1 compares the purchasing dollars spent on new servers with the power and cooling cost since 1996 and projects those numbers until 2010. The total power and cooling bill for servers in the US stands at a whopping $14 billion a year, and if current trends persist, that bill is going to rise to $50 billion by the end of the decade [10]. Forrester says servers would use about 30 percent of their peak electricity consumption while sitting idle 70 percent of the time. IDC says the cost to power servers will exceed the cost of the servers by next year [9]. The U.S. Department of Energy states that energy consumption for a data center can be 100 times higher than that of a typical commercial building.

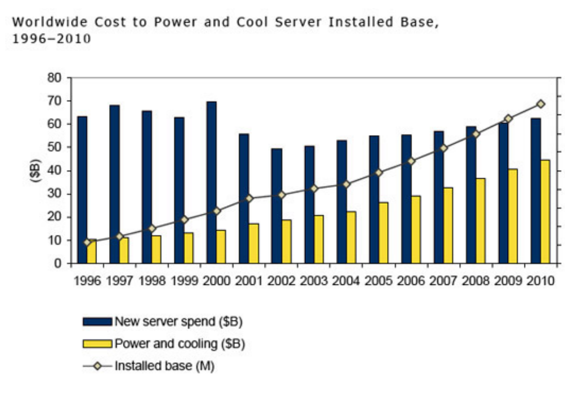

Figure 1: Worldwide Cost to Power and Cool Server Installed Base 1996-2010: Source IDC, 2007

Reducing electricity consumption in today's data centers is being addressed in initiatives from lawmakers, hardware vendors, and performance benchmark organizations. Lawmakers and benchmarks give incentives to direct system vendor's development efforts. On Thursday, July12 2006, the U.S. House of Representatives adopted bill H.R. 5646 [8], which calls for additional research to reduce energy costs and electricity consumption by computer servers and data centers. The bill directs the Environmental Protection Agency (EPA) to identify the potential energy and cost savings to the federal government and private business through the purchase of energy efficient servers, and includes an

Permission to make digital or hard copies of portions 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. Copyright for components of this work owned by others than VLDB Endowment must be honored.
Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists requires prior specific permission and/or a fee. Request permission to republish from: Publications Dept., ACM, Inc. Fax +1 (212)869-0481 or permissions@acm.org.

PVLDB '08, August 23-28, 2008, Auckland, New Zealand
Copyright 2008 VLDB Endowment, ACM 978-1-60558-306-8/08/08

indication from Congress that it is in the best interest of the United States that server buyers give high priority to energy efficiency.

Hardware vendors are addressing the energy crisis by investing in developing energy efficient components and various mechanisms to save energy. Currently, some processors are being equipped with demand-driven clock speed adjustments. Running a processor at high clock speeds results in better performance, but the same processor will consume less energy at lower frequency, which may be acceptable for the end user.

Software vendors have also started investing in power efficient coding methods. LessWatts.org is an open source initiative to create a community around saving power on Linux by bringing developers, users, and system administrators together to share power conserving innovations.

Over the past several years, major innovations have been gearing towards power conservation in consumer systems (personal desktop and notebook computers). Recently, the industry has seen initiatives from major vendors for extending these technologies from small to medium business to enterprise data centers. Some of the research and development can be applied across the board, but drastic enhancements are required for enterprise data centers. One initiative is the adoption of standard-based servers and blade servers because they are more cost effective and ahead in developing energy-efficient components than traditional large enterprise systems. While there is a huge demand for reducing energy consumption, another challenge exists in measuring it especially on large-scale systems. In this paper our focus is transactional systems: multi-tier architectures consisting of large numbers of servers and disk drives.

SPEC announced the industry's first power benchmark (SPECpower_ssj2008 [21]) to evaluate the power and performance characteristics of servers. The current version measures the performance of server side Java application, exercising the processors, caches, memory hierarchy as well as the implementations of the JVM (Java Virtual Machine), JIT (Just-In-Time) compiler, garbage collection, threads, and some aspects of the operating systems of standalone servers.

The Transaction Processing Performance Council (TPC) defines benchmarks (TPC-C [26] and TPC-E [31]) to measure the performance of large scale transactional systems, but lacks a power metric. Proposals are currently being discussed that are expected to lead to a TPC power metric in the future [18].

In this paper we develop a power consumption model based on data readily available in the TPC-C full disclosure reports of published benchmarks. We verify our model with measurements taken from three fully scaled, optimized and published TPC-C configurations, including client systems, database server, and storage subsystem. By applying this model to a subset of 7 years of TPC-C results, we identify the most power intensive components and demonstrate the existing power consumption trends over time. If these trends continue into the future, the hardware enhancements alone will not be able satisfy the demand for energy efficiency. Our first preference for this model was TPC-E, TPC's new transaction processing benchmark. However, we chose TPC-C due its large results set spanning an entire decade. This makes it the ideal candidate for a trend analysis. As of March, 2008, there are 237 TPC-C publications and 8 TPC-E publications. This paper looks at potential database and hardware enhancements to meet the energy efficiency demands of future systems.

2. THE TPC-C BENCHMARK

TPC Benchmark C (TPC-C) [26] is an On Line Transaction Processing (OLTP) workload. Since its establishment in 1992 the TPC-C benchmark has been the industry benchmark to measure the performance of complex OLTP systems. It is a mixture of read-only and update-intensive transactions that simulate the activities found in complex OLTP application environments. It does so by exercising a breadth of system components associated with such environments, which are characterized by several conditions:

- Simultaneous execution of multiple transaction types that span a breadth of complexity
- Multiple on-line terminal sessions
- Moderate system and application execution time
- Significant disk input/output
- ACID properties
- Non-uniform distribution of data access through primary and secondary keys
- Databases consisting of many tables with a wide variety of sizes, attributes, and relationships
- Contention of data access and update

TPC-C is modeled after actual production applications and environments. It evaluates key performance factors such as user interface, communications, disk I/Os, data storage, and backup and recovery. The difficulty in designing TPC benchmarks lies in reducing the diversity of operations found in a production application while retaining its essential performance characteristics, namely, the level of system utilization and the complexity of its operations.

TPC-C benchmark is accepted in the industry as the most credible transaction processing benchmark with a large body of results across all major hardware and database platforms. The highly tuned and optimized nature of the TPC-C configurations makes it the best candidate for power modeling for transactions processing systems.

2.1 A Typical TPC-C System

The typical TPC-C system is designed in 3 tiers:

1. Tier: Driver System
2. Tier: Client
3. Tier: Database Server

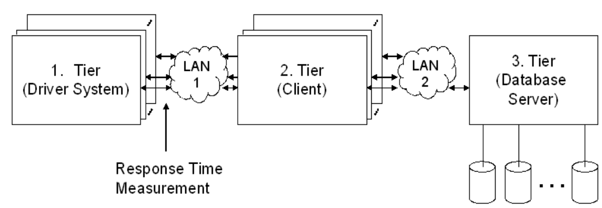

Figure 2: Typical TPC-C System Setup (conceptual)

The Driver System emulates the user load. It represents users generating the TPC transactions using a remote terminal emulator (RTE). The Clients run the TPC-C application and the transaction monitor or database RPC library, for instance, Tuxedo, or ODBC. The Database Server runs the database management systems (DBMS), transaction monitor, and TPC-C stored procedures, for instance, DB2, Oracle, SQL Server, or Tuxedo. The tiers are connected through a local area network (LAN). The transaction response time is measured on the driver system: the start time is when the transaction is generated by the Driver System and the end time is when the commit is received by the Driver System.

2.2 The TPC-C Workload

TPC-C simulates operators performing various transactions which are used typically in wholesale company against a transactional system,. It is not our intent to give a detailed overview of TPC-C, but rather to focus on the elements that impact the way we measure power consumption of the TPC-C system, that is transaction workload, execution rules, space requirements, and metrics.

A detailed overview of TPC-C is available at the websites referenced in notes [16] and [26]. Figure 3 details TPC-C's execution model.

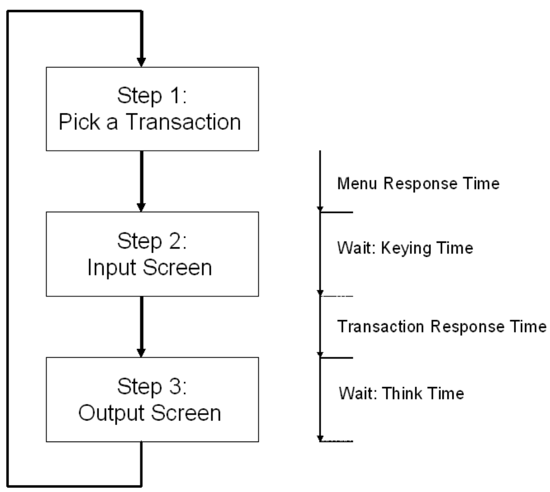

Figure 3: TPC-C Execution Model

Step 1 selects a transaction type from the menu according to a weighted distribution: approximately one "payment" transaction for each "new order" transaction and approximately one "order status" transaction, one "delivery" transaction, and one "stock level" transaction for every 10 "new order" transactions. This mix results in the complete business processing for each order. Then the system waits for the input/output screen to be displayed. The menu response time is measured as Menu RT. Step 2 emulates the user entering the required number of input fields (keying time is defined in the specification). The transaction response time is measured between finishing entering the keys of Step 2 and receiving the answer from the server. Then the emulator waits for the defined minimum think time while the input/output screen remains displayed. At the end of the think time the emulated user loops back to select a transaction type from the menu. A certain percentage of each transaction must exercise the rolled back function of the database. For instance, a fixed 1 percent of the new order transactions are chosen at random to simulate user data entry errors and demonstrate rolling back update transactions.

The TPC-C benchmark throughput is driven by the activity of the emulated terminals connected to each modeled warehouse. To increase the throughput, more warehouses and their associated terminals must be configured. Each warehouse requires a number of rows to populate the database along with some storage space to maintain the data generated during a defined period of activity called a 60-day period. These requirements define how storage space and database population scale with throughput. The intent of the scaling requirements is to maintain the ratio between the transaction load presented to the system under test, the cardinality of the tables accessed by the transactions, the required space for storage, and the number of terminals generating the transaction load. For each active warehouse in the database, the SUT must accept requests for transactions from a population of 10 terminals. The WAREHOUSE table is used as the base unit of scaling. The cardinality of all other tables (except for ITEM) is a function of the number of configured warehouses (cardinality of the WAREHOUSE table). This number, in turn, determines the load applied to the system under test which results in a reported throughput.

The configured disk space must fulfill the 60 day requirement. That is, the test database must be built including the initial database population and all indices present during the test. It must sustain the reported throughput during an eight hour period. The total storage space allocated for the test database must be broken down into the following:

- **Free-Space**: any space allocated to the test database and which is available for future use. It is comprised of all database storage space not used to store a database entity or not used as formatting overhead by the data manager.
- **Dynamic-Space**: any space used to store existing rows from the dynamic tables. It is comprised of all database storage space used to store rows and row storage overhead for the dynamic tables. It includes any data that is added to the database as a result of inserting a new row independently of all indices. It does not include index data or other overhead such as index overhead, page overhead, block overhead, and table overhead.
- **Static-Space**: any space used to store static information and indices. It is comprised of all space allocated to the test database and which does not qualify as either Free-Space or Dynamic-Space.

The TPC-C performance reported in a benchmark publication is the transaction throughput during steady state condition. The performance is measured during the measurement interval. The measurement interval must begin after the system reaches steady state, be long enough to generate reproducible throughput results that would be representative of the performance that would be achieved during a sustained eight hour period and extend uninterrupted for a minimum of 120 minutes. Although the measurement interval may be as short as 120 minutes, the system under test must be configured to run the test at the reported tpmC for a continuous period of at least eight hours without operator intervention, maintaining full ACID properties. For example, the media used to store at least 8 hours of log data must be configured if required to recover from any single point of failure. A graph of the throughput of the new order transaction versus elapsed time (for example wall clock) must be reported for both ramp-up time and measurement interval. At least 240 different intervals should be used with a maximum interval size of 30 seconds. The opening and the closing of the measurement interval must also be reported and shown on the graph. The start time for each of the checkpoints must be indicated on the graph. This is important to note since our power consumption model uses peak power consumption, which is equivalent to the power consumption during the steady state of the TPC-C run.

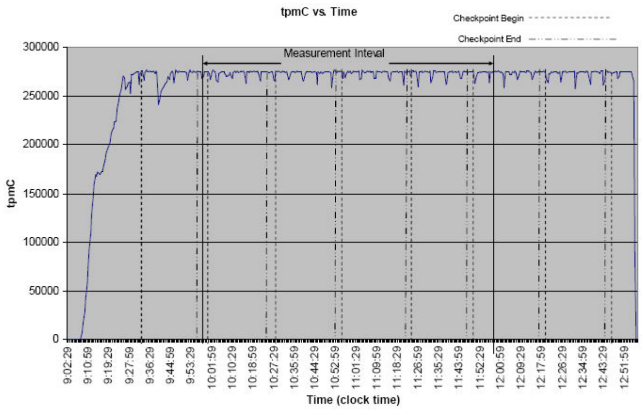

Figure 4: Throughput versus Time: TPC-C Publication 107111201

Figure 4 graphs the number of new order transaction (tpmC) during the measurement interval as achieved in benchmark publication 107111201 (see Section 2.3 for a detailed description). The figure also includes the start and the end of taking checkpoints in regular intervals as required by the benchmark specification.

Another important metric of TPC-C benchmarks is price-performance. The price-performance metric [$/tpmC] is calculated by dividing the three-year cost of ownership of all components by the tpmC. See TPC's pricing specification [21] for how the three-year TCO is calculated.

2.3 Example Publication 107111201

As an example for a real TPC-C system setup, we describe the current[1] price-performance record TPC-C results in the category of two-processor systems [25]. The system achieved 273,666 tpmC at a price-performance of $1.38 per tpmC. The benchmark configuration is depicted in Figure 5. The client systems consisted of five DL360 Generation 5 servers, each equipped with two Intel dual core X5130 2.0GHz with 4MB L2 cache, one dual-port integrated Giga-bit network controller and integrated SMART array controller holding a 36GB 10K SCSI Attached Storage (SAS) small form factor (SFF) disk drive. Microsoft COM+ on each client system served as the queuing mechanism to the database. Each delivery request was submitted to Microsoft COM+ asynchronously with control being returned to the client process immediately and the deferred delivery part completing asynchronously.

The database server consisted of one HP ProLiant ML370 Generation 5 system equipped with two Intel quad core processors X5460 3.16GHz with 12M L2 cache, 64GB of main memory, dual-port integrated Giga-bit network controller and nine storage controllers. Two of the storage controllers were Smart Array Controller P600 that were connected to disk trays internal to the server holding a total of 13 SAS SFF disk drives, which were powered by the server power supply. One of the disks was used for the operating system and database software. The remaining 12 disks are connected to one of the two Smart Array Controller P600 controllers holding the database redo log files.

The storage subsystem consisted of 28 modular disk array enclosures of type MSA70. Each MSA70 had 25 36GB 15K 2.5 SAS SFF disk drives holding the database tables and indexes.

[1] As of February 4th, 2008

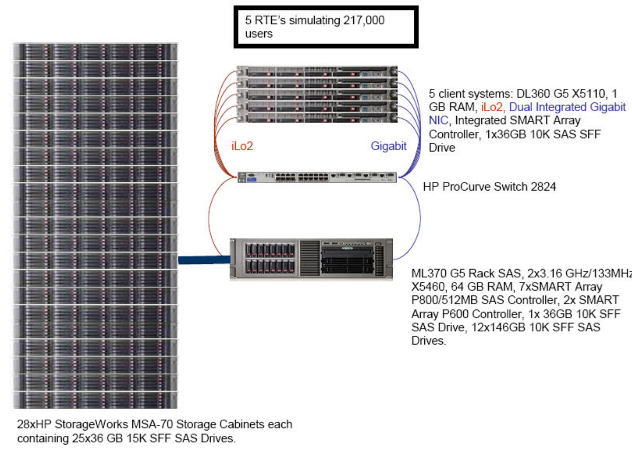

Figure 5: System Overview Publication 107111201

3. HOW POWER CONSUMPTION OF TPC-C SYSTEMS CAN BE ESTIMATED

In the previous section, we have shown how intricate TPC-C systems are. Measuring or even estimating their power consumption is difficult because of the assortment of components that are involved in their three-tier architecture and the lack of power measurements of its individual components from large-scale deployments. In this section we develop a simplified power consumption model that can be applied to any published TPC-C result and representative transaction processing systems.

3.1 System Boundaries

The three-tier architecture, conceptualized in Figure 2, includes the driver systems (Tier 1), middle-tier (Tier 2) a.k.a clients and database server (Tier 3). Today's transaction system users are most often connected through the Internet rather than through closed circuit systems. Hence, in most cases, the deployment of a transaction system does not include the driver systems. In order to estimate the power consumption cost of TPC-C transactions, we only include Tier 2 and Tier 3 systems into our model. In TPC terms, these systems are referred to as the System Under Test (SUT). The database server (Tier 3) is typically comprised of one or more compute systems and a storage subsystem, usually comprising of one or more RAID devices (Redundant Arrays of Independent Disks). We refer to the "container" of the RAID devices as disk enclosures. Note that the TPC-C three-year cost of ownership includes both hardware and software costs for all systems in the second and third tiers.

Even reducing the scope to the second and third tiers, we are faced with very complex systems. Power consumption estimates are very difficult. As described in endnote [6], we circumvent this problem by utilizing an indirect power estimation model that is simple yet accurate in predicting power consumption during the measurement interval of a TPC-C benchmark run. The components we include in our model are listed in Figure 8. They represent the major power consumers of a typical TPC-C system: CPU, memory, disks, server chassis and disk enclosure. The first column lists a description of the components. The second column lists the estimated power consumption of these components obtained from the specification (peak power consumption/nameplate power consumption) of the components used in the published benchmark 107111201 [25]. We will describe how we obtain these numbers in detail in Section 3.2. The following columns in

Figure 8 note how many components are included in the clients, database server, and storage systems of this particular benchmark publication. We will describe this configuration in detail since this is one of the systems that we will use to validate our model.

3.2 Power Consumption Model

In our power consumption model we assume that the peak power consumption of an entire system during the measurement interval is identical to the aggregate of the individual nameplate power consumptions. The components that we include in our model are CPU, memory, disks, server chassis, and disks enclosures. We differentiate between database server and client CPUs since the client CPUs are typically less powerful than the database server CPUs. This is because a TPC system is usually sized around the database server, that is, the number of clients and their CPU choice is a function of how fast the database can drive those systems. We also differentiate between internal and external disks for the same reason.

rocessor Description	DP [W]
AMD 8220SE 2.8 GHz	93
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 2.2GHz Dual Core - 2.2 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 Itanium 2 Processor 6M - 1.5 GHz	107
Intel Itanium2 - 1 GHz	100
Intel Itanium2 - 1.6 GHz	130
Intel Pentium III Xeon - 900 MHz	50
Intel Pentium Xeon MP - 1.6 GHz	55
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

Figure 6: Processor Peak Power Consumption ([11] [2])

For each of the components listed above we determine its peak power consumption. We obtain the peak power consumption of CPUs from their manufacturer's specification [2] [11]. The peak power consumption for CPUs is depicted as Thermal Design Power (TDP). For instance, the TDP of the client CPUs in [25] our sample configuration (Intel X5130) is 65 watts per processor, while that of the database server CPUs (Intel X5460) is 120 watts. For a selection of processor TDP see Figure 6. As for main memory we approximate the power consumption of main memory by assuming 9 watts per memory DIMM[2] [6]. The peak power consumption levels of the disk drives are obtained from the manufacturers' web sites [19] [20]. Our example system uses 36G 15K RPM SFF SAS as external disk drives which have a 9.2 watt peak power consumption and 36G 10K RPM SFF SAS as internal disk drives, which have a 7.2 watt peak power consumption. For a selection of disk drive peak power consumption, see Figure 7.

[2] DIMM: Dual In-Line Memory Module

nterface	apacity GB]	PM 1000]	orm Factor[3] nches]	eak Power onsumption [W]
	36	10	3.5	12.5
	36	15		14.5
	72	10		12.6
	72	15		13.2
	146	15		14.9
	300	10		16
	300	15		18.1
Serial Attached SCSI	73	15		13.7
	146	15		15.4
	300	15		18.5
	36	10	2.5	7.2
	36	15		9.2
	72	10		8.4
	72	15		9.2
	146	10		9

Figure 7: Disk Peak Power Consumption

Since the server chassis and its infrastructure (fan, power supply, and so on) are sized according to its components (CPU, memory), we express its power consumption as 30 percent of the power consumption of its components [4][23] plus a fixed overhead of 100 watts. This applies to both client and database servers. For our example, this means that the overhead for all five client server chassis is 635 watts and 261 watts for the database server chassis. Similarly we account for the power consumption of the disk enclosure (disk array). Since usually all slots are being utilized in a given enclosure, we approximate the power consumption of the disks enclosures with 20% power overhead of the aggregate power consumption of all external disks.

As a result, the output of our power consumption model is an estimate for the peak power consumption of an entire system under full load. The estimated power consumption of our sample configuration is summarized in Figure 8.

Description	Power Consumption [W]		
	Clients	DB Server	Storage Subsystem
CPU	5x65=325	2x120=240	n.a
Memory	10x9=90	16x9=144	n.a
Disk Drives	5x9=36	151	700x9.2=6440
Server Chassis	635	261	n.a.
Disk enclosures	n.a.	n.a.	1288
Total	**1086**	**796**	**7728**

Figure 8: Components of the Published TPC-C Benchmark 107111201 [25]

3.3 Verifying the Power Consumption Model To The TPC-C Workload

The power consumption model introduced in the previous section is just an estimate of how power is consumed in real systems. The real power consumption of a specific system depends on many factors, such as the workload that is run, how balanced the system is, and the environmental parameters of the system under test. In this section we verify the power model using the TPC-C workload. We compare the measurements of three systems used in TPC-C benchmark publications. We first outline the boundaries in which our power consumption model is valid in the context of TPC-C. Then we define the power measurement methodology, and finally we compare the two numbers.

[3] The form factor indicates the size of the disk drive. The two most common sizes are 3.5 and 2.5 inch.

3.3.1 Power Model Boundaries

The workload plays a major role in the power consumption of any system. Workloads that utilize a system 100 percent for the entire duration of the measurement interval are suitable for this verification. Workloads that impose oscillating system utilization are not suitable for verification. For instance, consider the power test of a TPC-H [13] system. Being a single user test it measures how well a system can parallelize a given query in order to deliver the result in the least amount of time. The characteristics of most of the TPC-H queries are such that not all queries utilize the entire system the entire time. Not all resources (e.g. IO, CPU, Memory, network) are fully used during the execution of the query. For instance, a hash join is CPU bound during the build phase of its hash table and, usually, IO bound during its probe phase. TPC-H's 22 queries are so diverse that it is impossible to setup a system that assures 100 percent utilization 100 percent of the time during the power test. Consequently, the system consumes more power in the IO subsystem during some time of the power test and more CPU power during other times of the power test. Contrarily, the TPC-C benchmark is constructed such that the performance numbers are obtained during the steady state of the system (see Figure 4), during which all components are fully utilized.

Another important factor in evaluating our power consumption model is system balance. Depending on the application and system, an optimal component ratio has to be maintained to keep all the components (CPU, disks, controllers) utilized during the measurement interval. If a system does not have the optimal ratio between these components, our power consumption models will not produce accurate estimates. This is because it assumes that all components are used during the duration of the measurement. Due to the nature of the TPC-C benchmark workload and the typical business objective of demonstrating performance and price-performance, all publications have been maintaining such component ratios. No vendor can afford to over-configure one part of the system because all parts that are used in a benchmark need to be disclosed and priced. And price-performance is widely being used by system vendors to showcase their advantages over those of their competitors. For instance, if a vendor over-configures a database server with 50% more CPUs, those CPUs need to be priced, and, since the number of CPUs is disclosed, the result will be used by competitors to show that they can achieve the same performance with fewer CPUs. Lastly, some database vendors tie their pricing model to the number of CPUs, while some tie it to the number of disks. This inconsistency makes it even more unattractive to publish unbalanced TPC-C performance results.

Environmental parameters certainly play a big role in how much power is consumed by a certain system. Power measurements taken in a server room at 0 degrees Celsius are certainly lower than the same measurements in a server room at 30 degrees Celsius. This is, however, irrelevant for our power consumption model since we can assume the same temperature for all benchmarks.

3.3.2 Power Measurement Methodology

For each of our test systems we use the same power measurement methodology, similar to that used in "JouleSort: A Balanced Energy-Efficiency Benchmark" [22]. During the steady state of the TPC-C workload we use five digital powermeters to simultaneously capture the power consumption of two randomly chosen client systems, the entire database server and two randomly chosen disk enclosures. For each of the five measurement sets we compute the average power consumption during the measurement interval. The two power measurements for the identical components of the client system and disk enclosure show a difference of less than 0.5% indicating that it is safe to assume that all instances of the same system are being utilized equally. Hence, the power consumption for all client systems and the entire storage subsystems can be extrapolated by multiplying the measured numbers by the actual number of systems used.

3.3.3 Power Model Verification

We use three different systems to verify our power consumption model. Each of the following three systems published TPC-C benchmark results. System A has been described in Section 2.3. System B is similar to System A in respect to the database server, but different in terms of the client systems and storage subsystem. Like System A, it uses a two-processor system for the database server with 64 Gigabytes of memory. However, instead of an ML360, it uses an HP ML370. The same client systems are used in System A and System B. System B runs different software. For an operating system, it runs Windows, and for DBMS, it runs SQL-Server. The third system (System C) uses two more CPUs for the database server for a total of four. It runs a DL580 as the database server. Instead of dual core processors it uses quad core processors. It uses a larger number of client systems, which are different to the above client systems.

System A: 2 CPUs, published result 107111201[25]
System B: 2 CPUs, published result 108010701 [27]
System C: 4 CPUs, published result 107090502 [24]

Tier	Power Consumption [W]								
	System A			System B			System C		
	Power Model	Meas-ureme	Diff [%]	Powe Mode	Measur ment	Dif [%	Powe Mode	Measur ment	Diff [%]
Storag	7728	6720	13	6973	6240	11	11631	9600	17
Client	1086	845	22	1813	1352	25	2369	2028	14
DB-Serve	796	705	11	618	510	17	909	820	10
Total	**9610**	**8270**	**14**	**9404**	**8102**	**14**	**1490**	**12448**	**17**

Figure 9: Comparison Power Consumption Model and Power Measurements of Three TPC-C Systems

Using the measurement methodology presented in the previous section, Figure 9 shows the power consumption of the client systems, database server, and the storage subsystem of systems A, B, and C, both estimated by the power model and measured with power meters. The paper, "Power Provisioning for a Warehouse-sized Computer" [6] refers to this as the difference between the nameplate value and actual peak power. For each of the three layers, Figure 9 shows the difference in percent between the estimate and the measurement. The power model over-estimates the system's power consumption. The difference varies between 10 and 25 percent. Overall the difference between the three systems is between 14 and 17 percent. The paper [6], which applied a different workload on smaller systems, shows that the difference between the nameplate model and actual peak power consumption is 30 percent. The 15 percent difference in the power consumption estimation with our power model and the actual measurements is smaller than the difference found in the paper [6]. The difference between the modelled number of a server (251W) and its measured power consumption referenced in the paper [6] is about 40 percent. (In order to calibrate our power model to the TPC-C workload, we abate our power model number by 15 percent.

4. HISTORIC TREND ANALYSIS OF PUBLISHED TPC-C RESULTS

The power consumption model developed in the previous section allows us to estimate the power consumption of any published TPC-C result because all information necessary for the power consumption model is readily available in the TPC-C Full Disclosure Report (FDR). In this section, we will apply the power consumption model to a subset of all available results (defined in Section 4.1). We then analyze the following trends: performance, price/performance, total system power consumption and transaction/power trends. We finish this section with a analysis of components that consume the most power in TPC-C systems.

4.1 Four Processor TPC-C Results

As of March 2nd, 2008, a total of 237 TPC-C Version 5 results have been published. They differ in data volume, systems configurations, hardware architectures, and database management systems. Fourteen hardware vendors have published results on five different database management systems using scale-up and scale-out solutions. The systems range from 2 processors/2 cores to 64 processors/128 cores.

In this section, we conduct a historic trend analysis of Version 5 benchmark results in respect to performance, price, and power consumption. In order to be able to compare systems realistically, we only consider results using a single system with 4 processors and up to 16 cores for the database server. Since the number and size of the client systems is a function of the size of the database server, they vary from 2 to 48 processors. As of March 2nd, 2008, 64 TPC-C Version 5 results used four processor systems for the database server. Seven different hardware vendors published with four different database vendors. The results can be downloaded from the TPC website [1].

4.2 Performance Trends

Before analyzing the power consumption trend, we study the performance and price/performance trends. Figure 10 shows the transaction-performance [tpmC] of our sample data set.

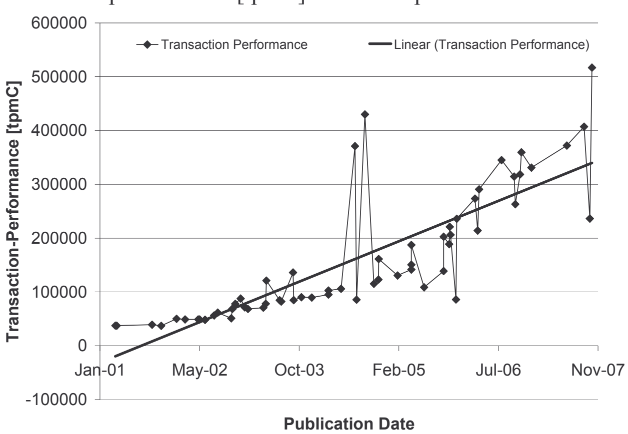

Figure 10: Performance Trend of TPC-C Results

The diamond shaped graph indicates the actual tpmC numbers, while the solid line shows the trend of time. Four-processor systems achieved about 40,000 tpmC in early 2001. The current performance leader in four-processor systems achieved about 520,000 tpmC. This is a 13x increase in a period of seven years. The solid line is the linear trend line of all four processor results. There are two results noticeably higher (both on IBM's Power 5) but the overall trend of transaction-performance of system is about 1.28x increase per year (9x increase over a seven year period). This is roughly in line with Moore's law [11].

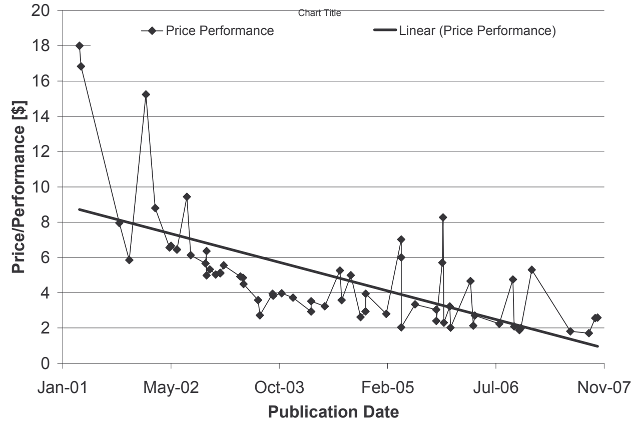

Figure 11: Price-Performance Trend of TPC-C Results

In the same period, the price-performance has dropped from \$18.00 per tpmC to \$1.71 per tpmC. This is a 10.5x improvement. The diamond shaped graph in Figure 11 shows the price-performance of all 64 four-processor results. The solid line indicates the linear trend line. The trend line polishes the picture by ignoring the very expensive early results. According to its trend line the price-performance increases at a rate of 9x in the seven year period (1.28x per year). This is roughly in line with the transaction performance improvement.

4.3 Power Consumption Trends

Now we analyse the power trend of our sample TPC-C result set. For each result we compute the total peak power consumption according to the calibrated power consumption model we developed in Section 3.2. We are interested in two trends: a) how the system power evolved over time; and b) how the transaction performance per system power evolved over time.

Using our power model Figure 12 graphs the system power consumption for the 64 TPC-C results. The lowest reported system power consumption is about 2000W, while the highest system power consumption is about 15,000W. This is a difference of 7.5x. The trend line (solid line) reduces this difference to about 6x over the seven year period (0.86x per year). This is slightly lower than the 9x-performance and price-performance increase we saw in Section 4.2.

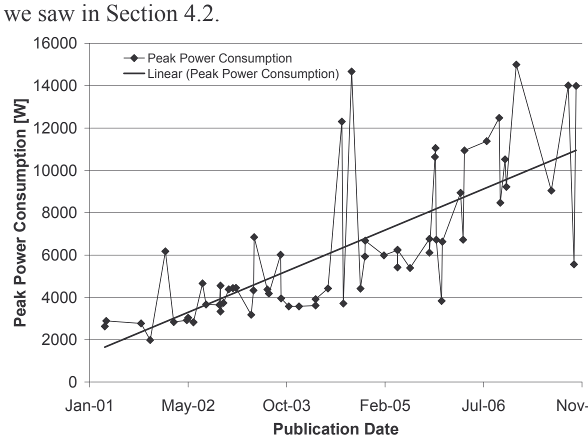

Figure 12: System Power Trend TPC-C Results

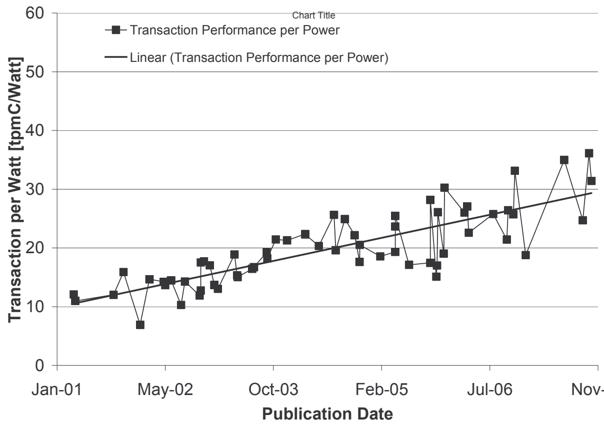

Figure 13: Performance per Power Trend TPC-C Results

Figure 13 shows the power used per tpmC for all 64 results. The lowest transaction-per-watt is ca. eight, while the highest is ca. 37. This is an increase of only 4.6x over seven years. The trend line further decreases the increase to below three (0.4x per year).

4.4 Power Consumption Distribution

We will now look at the power distribution of the different components of a typical TPC-C system. We compute the percentage of the power consumed by each component listed in our power model (see Section 3.2) by averaging the power consumption numbers of all 64 results. Figure 14 shows the power consumption of each component as a percentage of the overall power consumption. Note, that the numbers are rounded. Grouping the components in client systems, database server, and storage subsystem, the figure shows that the power consumption of an entire TPC-C system is dominated by the storage subsystem. 79.1 percent of all power is consumed by the storage subsystem. Eleven percent is consumed by the database server and 9.9 percent is consumed by the client systems. Within the storage subsystem the largest power consumer are the disks (63 percent of the total power), followed by the disk enclosures (12.6 percent of the total power). Seven percent of the power is consumed by the database server's memory and 1.6 percent is consumed by the database server's CPUs. The power consumption of the disk in the database server is negligible. The clients' CPUs consume 4.2 percent; 2.1 percent is consumed by memory; 1.3 percent is consumed by the disks; and 3.9 percent is consumed by the chassis.

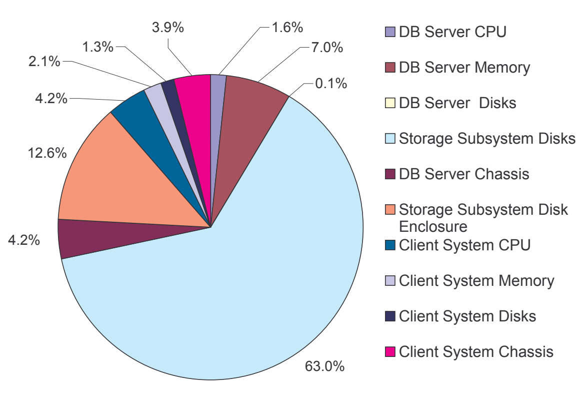

Figure 14: Average Power Consumption of Major Parts Used in TPC-C Benchmarks

5. CURRENT APPROACHES TO SOLVE THE POWER CONSUMPTION PROBLEM

The previous section clearly identifies the storage subsystem as being the largest of the power consumers. In this section we summarize the approaches employed by hardware and software vendors to address the problem.

5.1 Small Form Factor Disk Drives

Enterprise data access and transfer demands are no longer driven by advances in CPU processing alone. Performance, reliability, capacity, space, and power consumption of storage subsystems have been critical factors for mission-critical applications. Small Form Factor (SFF) 2.5" Serial Attached SCSI (SAS) disk drives are 70% smaller and use half the power of traditional 3.5" U320 SCSI. They also permit better airflow. In addition, SAS technology has 1.75 million hours MTBF (Mean time between failure), compared to U320 with 1.5 million hours MTBF.

5.2 Solid State Disk Drives

In the past few years, we have seen power saving technologies like spinning down the disks on notebook computers and differing raid rebuild and other activities intelligently. All have been effective, but are still limited to the baseline power needed to spin disks and move their heads. NAND-based solid state disks, which have no moving parts, are expected to have an active power consumption during read or write operations in the range of 0.03-0.06 watts while idle power is minimal. Compared to an average of 12 watts for a LFF drive and about 8 watts for a SFF disk drive, the power consumption of solid state disks is much lower. Although they are not commercially available yet, they are expected to be widely accepted by the industry in the near future.

5.3 Large Memory Configurations

Transaction processing performance can be improved significantly by using large main memory. Accessing data from main memory is an order of magnitude faster than accessing data from disk, which dramatically increases performance. Most DBMS use in-memory structures (buffer cache) to store database pages. Increasing the buffer cache increases the buffer cache hit ratio so that I/O operations issued by transactions do not need to access data from the storage subsystem.

As demonstrated in the trend analysis in section 4.4, the storage subsystem is the largest power consumer. Reducing the number of disk drives and disk enclosures in the storage subsystem can significantly decrease total energy use. The industry trend of gradually increasing memory density while decreasing memory cost makes very large memory configurations very attractive to customers. HP's current generation of 2- and 4-processor industry-standard servers supports 128GB and 256GB of main memory, respectively. Additionally, the cost-per-gigabyte of main memory is now under $200, compared to $1000 five years ago; this trend is more than faithful to Moore's law.

The above approach does not unfold its full potential for the TPC workload because of TPC-C's continuous scaling model. Being praised as one of the major strengths of the TPC-C benchmark, the continuous scaling model is also one if its weakness from the viewpoint of large memory systems. As explained in detail in Section 2.2, the TPC-C specification mandates a fixed ratio between the transaction load, the cardinality of the tables accessed by the transactions, the required space for storage, and the number of users generating the transaction load. Due to this ratio, the higher the performance of a system, the larger the database needs to be. Increasing the ratio of buffer cache to database size (by adding more memory) results in higher performance, which results in a large database, which reduces the buffer cache to database size ratio. This design characteristic of TPC-C makes it unattractive to use to reduce the power consumption of a TPC-C result by adding memory.

5.4 Data Compression

The result of our detailed power analysis shows that on average 70 percent of all power of a typical on-line application is consumed in the storage subsystem (controllers, disks, and disk enclosures). Data compression can significantly reduce the number of disks necessary to run such an application. Oracle introduced index compression in Oracle 8 and table compression in Oracle 9. Oracle 11 takes the compression concept a step further by offering an advanced compression option, which presents customers with a comprehensive set of compression capabilities to help maximize resource utilization and reduce costs. It allows IT administrators to significantly reduce their overall database storage footprint by enabling compression for all types of data, that is regular structured data (numbers, characters), unstructured data (documents, spreadsheets and XML), or backup data.

The new compression features in Oracle Database 11g are not limited to read-only applications. Online transaction processing (OLTP) Table Compression allows structured or relational data to be compressed during all types of DML operations, such as INSERT, UPDATE and DELETE. This new feature enables compression for any application and leverages a sophisticated and intelligent algorithm that minimizes the compression overhead during write operation, thereby making it viable for all application workloads. Additionally, compressed data uses less space on disk and utilizes memory more efficiently so it significantly improves performance of queries by reducing disk I/Os and improving memory efficiency.

Another area of data compression is performing regular database backups. The storage requirements for maintaining database backups and backup performance are directly impacted by the size of the database.

Apart from the obvious power and subsequent operational cost reductions possible with compression, there are other benefits, such as capital investment in storage subsystems, and reduction in overall resource requirements including memory, IO, and network bandwidths. However, this is commercially practical only if performance does not suffer at the same rate as data is compressed.

5.5 Grid Computing

The conventional approach to application deployment, where each application is dedicated one physical server, inevitably leads to over-provisioning and under-utilization of server hardware assets. In this approach, IT organizations provision at least one server for every application or service. With high availability (HA) requirements some applications require at least two servers. Unlike in a TPC-C run the workload on these servers is not constantly in a steady state but fluctuates over time. In extreme cases servers are idle during those times of the night when no data maintenance tasks are conducted. Consequently, on average, servers are utilized only at a fraction of their total load capacities. While these servers are not utilized, costs associated with power, cooling, network infrastructure, storage infrastructure, administrative overhead, and floor space continue to accrue. The result is a convoluted data center that uses more electricity than it should.

Database grid computing and operating system virtualization have the potential to drastically reduce the power consumption of today's data centers by viewing vitalizing compute power.

The goals of grid computing are closely aligned with capabilities and technologies that Oracle and HP have been developing for years. Oracle provides substantial grid computing technology since Oracle Database 10*g*. Oracle and HP envision grid computing by orchestrating many small servers and storage subsystems into one virtual computer [14]. There are three levels of abstraction: the first level contains server nodes; the second level contains database applications; and the third level contains storage subsystems. This three-level architecture, which allows for a very flexible grid implementation, requires a shared-disk implementation. Shared-disk architectures provide dynamic resource sharing between applications of one database and between databases and provisioning by virtualization of resources. This enables enterprises to dynamically allocate resources for various enterprise tasks based on changing business priorities, reducing underutilized resources and decreasing overcapacities. Applications can easily share compute and data resources by migrating between servers on the grid to leverage available resources. Schedulers on the grid track resource availability, and assign resources accordingly.

5.6 Virtualization

Virtualization is the pooling and sharing of resources, including servers, storage and networking. In a virtualized environment, the logical functions of computing, storage and network elements are separated from their physical functions. Elements from these pools can then be manually or dynamically allocated to meet the changing needs and priorities of a business. This leads to a drastic reduction in the volume of servers in data centers and a dramatic savings on power and cooling costs. It can also reduce real estate cost by reducing the footprint of the data center [35] [36].

5.7 Beyond Storage

Once the improvements to reduce power consumed by disk drives has become main stream the focus will be on other components of a transaction processing system, such as CPU, memory, and so forth. HP is investing heavily in these areas.

5.7.1 Low Power Processors

Examples for low power processors are Intel LV – Low Voltage and AMD HE – High Efficiency processors. The HP Server Based Computing (HP SBC) test environment in September 2007 [34] highlights the power/performance advantages of lower wattage Quad-Core Intel® Xeon® processors in HP industry standard servers. Test results indicate that for a specific HP industry standard server configuration, a lower wattage processor reduced power consumption by 21%, incurring a performance penalty of only 8%. Thus, if system power consumption is a major concern, a lower wattage processor may deliver more than acceptable performance at a reduced cost.

The Dual-Core Intel® Xeon® processor LV and Dual-Core Intel® Xeon® processor ULV are members of Intel's growing product line of multi-core processors. Each dual-core processor combines the benefits of two high-performance execution cores with intelligent power management features to deliver significantly greater performance-per-watt over previous single-core Intel Xeon processor-based platforms. The dual-core/dual-processor capabilities are ideal for a wide range of low-power communications and embedded applications.

AMD has expanded the breadth of its low-power solutions with AMD Opteron processor Models 1218 HE, 2218 HE and 8218 HE. Designed to offer industry-leading performance-per-watt at only 68-watt maximum thermal design power, these processors are ideal for energy-conscious customers looking to reduce power and cooling bills and to achieve greater density in the data center. AMD Opteron HE processor models now include three 1000 Series models, bringing the benefits of reduced thermals over previous AMD Opteron 1000 Series processors to entry-level server customers while preserving the enterprise reliability they value. All of the new processors feature AMD PowerNow!™ technology which is designed for reduced system level energy consumption, with multiple levels of lower clock speed and voltage states that can reduce processor power consumption by as much as 75 percent during idle times.

5.7.2 Low Power Memory

Memory is one of the most critical components in transaction processing servers, and one of the easiest to upgrade quickly to enhance and improve overall system performance and to reduce power consumption. The dependence on industry standard servers to run memory-intensive applications is pushing the memory capacity of servers to new levels. HP constantly evaluates new energy technologies to insure our product offerings provide the most efficient, reliable, and effective results with maximum performance from our servers while still reducing power consumption. HP continues to review new technologies as they become available and will offer them for use with HP servers as they meet the requirements of our customers.

DDR-2 SDRAM is the second generation of DDR SDRAM. It offers data rates of up to 6.4 GB/s, lower power consumption, and improvements in packaging. At 400 MHz and 800 Mb/s, DDR-2 increases memory bandwidth to 6.4 GB/s – 800 percent more than original SDRAM. DDR-2 SDRAM achieves this higher level of performance and lower power consumption through faster clocks, 1.8-V operation and signaling, and simplification of the command set.

DDR-3, the third-generation of DDR SDRAM technology, will make further improvements in bandwidth and power consumption. Manufacturers of DDR-3 will initially use 90 nm fabrication technology and move toward 70 nm as production volumes increase. DDR-3 will operate at clock rates from 400 MHz to 800 MHz with theoretical peak bandwidths ranging from 6.40 GB/s to 12.8 GB/s. DDR-3 uses 1.5-V signaling (compared to 1.8 V for DDR-2) for lower power consumption. A thermal sensor integrated on the DIMM module signals the chipset to throttle memory traffic to the DIMM if its temperature exceeds a programmable critical trip point.

5.7.3 Energy Efficient Power Supplies

Power supplies convert high-voltage alternating current (AC) into low-voltage direct current (DC) for use by the electronic circuits in office equipment, telecommunications, and consumer electronics. Over 2.5 billion AC/DC power supplies are currently in use in the United States alone. About 6 to 10 billion are in use worldwide. Most likely, the computer you're using to view this article wastes 30-40% of the electrical power it consumes because it is using an inefficient power supply. While the best power supplies are more than 90% efficient, some are only 20 to 40% efficient, wasting the majority of the electricity that passes through them. As a result, today's power supplies consume at least 2% of all U.S. electricity production. More efficient power supply designs could cut that usage in half, saving nearly $3 billion and about 24 million tons of carbon dioxide emissions per year [5].

HP industry standard server power supplies lead the industry globally in efficiency, reliability, power density, size, flexibility, cost, and commonality. As founding members of key power efficiency consortiums like the Green Grid organization for EnergyStar, and Climate Savers Computing, HP continues to drive the standards for power optimization on industry standard servers. HP HE power supplies exceed 90% efficiency under standard.

5.7.4 Blade Systems and Thermal Logic Technology

Blade Systems are self-contained infrastructure designed to address management, utilization, and power and cooling. Blades range from servers and storage devices to workstations and virtual desktops.

HP and IDC forecast [37] a 69% reduction in energy consumption over a three-year period for IT organizations that migrate to blade architectures. HP Thermal Logic Power and Cooling Technology for the HP c-Class blade systems [38] offers built-in thermal instrumentation and controls to adjust and shift power load and thermal control automatically, based on changes in workload demand and environment. This gives the ability to raise system performance without exceeding the power and cooling capacity.

6. CONCLUSION

Energy cost is already the number one challenge for today's data centers. TPC-C is known as the most credible transaction processing performance benchmark, as can be seen by the large number of benchmark results and due to the fact that it is the most requested test for server requests for proposals (RFP). This paper introduced a power consumption estimation model for TPC-C benchmarks. The accuracy of this model is verified by measuring power consumption of recently published TPC-C benchmarks. The model was applied to a large subset of TPC-C results to show performance and power performance trends. They show that performance is increasing and price-performance is decreasing – at nearly the rate Moore predicted 43 years ago. The model also shows that the performance-per-watt is not increasing at the same rate. It is increasing at a much lower rate (approximately 0.4x per year).

The paper further identified the components that consume the most power in TPC-C systems. Based on list of top power consuming components the paper discussed database and hardware enhancements that can help alleviate the power crisis.

Having realized the importance of energy efficiency, the TPC is working on developing energy metrics for all its benchmark. The analysis done in this paper will contribute directly to the TPC's effort to develop such metrics. Once defined, the power metric will have a ground-breaking impact on customer purchase decisions as TPC's performance and price-performance metrics have done in the past. These energy metrics will further enhance the relevance of TPC benchmarks for the industry.

7. FUTURE WORK

The power consumption model developed in this paper enables the estimation of peak power consumption during the TPC-C workload and similarly sized transaction processing systems. We are currently extending this model to include the 1st Tier systems. The resulting model will extend power estimates beyond Internet applications to closed circuit systems and systems that are deployed within the Intranet of a large corporation. This is particularly important for power estimates of entire new systems.

We are also investigating how a power performance metric can be applied to other TPC benchmarks, such as TPC-H and TPC-E. Ultimately, customers are interested in energy costs over the lifetime of a system. The simplest approach is to assume 24 hour peak performance and a nationwide average price per KWh. However, for most customers this is not applicable. A detailed analysis of how actual systems are being utilized is required.

Another area of interest is to analyze the impact of power preserving techniques that are mentioned in Section 5. Early measurements on real systems that have some of the hardware improvements enabled show very promising results. Once we have studied these measurements, we will try to incorporate the power improvement recommendations into our power consumption model.

Furthermore, we are interested in analyzing how a transactional system can be optimized for power with database techniques. To what extent can the number of disks be reduced without drastically decreasing performance? Where is the cross-over between performance and number of disks, CPU, memory, and number of client systems?

8. ACKNOWLEDGEMENTS

The authors would like to acknowledge Mike Nikolaiev for his guidance; Bryon Georgson and David Adams for TPC configurations and power measurements; Partha Ranganathan and Mehul Shah for their input on the energy model; Wade Vinson, David Adams, Klaus-Dieter Lange, Paul Cao and David Miller for their comments and feedback; Ray Glassstone and Tracey Stewart for their support.

9. REFERENCES

[1] 2 Processor TPC-C results as of 1-Feb-2008: http://www.tpc.org/tpcc/results/tpcc_advanced_sort.asp?PRINTVER=false&FLTCOL1=tpcc.c_server_procs&FLTCOLOPR1=%3D&FLTCHO1=2&ADDFILTERROW=&filterRowCount=1&SRTCOL1=tpcc.c_tpmc&SRTDIR1=ASC&ADDSORTROW=&sortRowCount=1&DISPRES=100+PERCENT&include_withdrawn_results=none

[2] AMD Opteron family of processors: http://www.amd.com

[3] Barroso, Luiz André. 2005. *The Price of Performance: An Economic Case for Chip Multiprocessing*. *ACM Queue*, September. http://labs.google.com/papers/priceofperformance.html

[4] Dimitris Economous, 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

[5] Efficiency of Power Supplies in the Active Mode: http://www.efficientpowersupplies.org/

[6] 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

[7] Gartner Urges IT and Business Leaders to Wake up to IT's Energy Crisis: http://www.gartner.com/it/page.jsp?id=496819

[8] H.R. 5646 [109th]: To study and promote the use of energy efficient computer servers in the United States: http://www.govtrack.us/congress/billtext.xpd?bill=h109-5646

[9] IDC, Enterprise Class Virtualization 2.0—ApplicationMobility, Recovery, and Management, Document #DR2007_5MEW, February 2007

[10] IDC, Solutions for Datacenters's Thermal Challenges,http://www.blade.org/docs/wp/idc_cool_blue_whitepaper.pdf

[11] Intel family of processors: http://www.intel.com

[12] J. Mitchell-Jackson, J. G. Koomey, B. Nordman, and M. Blazek: *Data center power requirements : measurements from silicon valley*. Energy (Energy) ISSN 0360-5442, 28(4):837 – 850, 2003

[13] Meikel Poess, Chris Floyd: *New TPC Benchmarks for Decision Support and Web Commerce*. SIGMOD Record 29(4): 64-71 (2000)

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

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

[16] Overview of the TPC Benchmark C: The Order-Entry Benchmark: http://www.tpc.org/tpcc/detail.asp

[17] TPC pricing specification http://www.tpc.org/pricing/spec/Price_V1.0.1.pdf

[18] Raghunath Othayoth Nambiar: *3-Year Energy Cost in TPC Benchmarks*, http://library.hp.com/techcon07/75.pdf

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

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

[21] SPECpower_ssj2008: http://www.spec.org/power_ssj2008/

[22] Suzanne Rivoire, Mehul Shah, Parthasarathy Ranganathan, Christos Kozyrakis, JouleSort: *A Balanced Energy-Efficiency Benchmark*, Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data (SIGMOD), June 2007

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

[24] TPC-C Benchmark Result 107090502: http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=107090502

[25] TPC-C Benchmark Result 107111201: http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=107111201

[26] TPC-C Benchmark Revision 5.9: http://www.tpc.org/tpcc/default.asp

[27] TPC-C Result 108010701: http://www.tpc.org/tpcc/results/tpcc_result_detail.asp?id= 108010701

[28] TPC-C results: http://www.tpc.org/tpcc/results/tpcc_advanced_sort.asp

[29] TPC-D Version 2.1: http://www.tpc.org/tpcd/default.asp

[30] TPC-DS Draft Version: http://www.tpc.org/tpcds/tpcds.asp

[31] TPC-E Benchmark Revision 1.4.0: http://www.tpc.org/tpce/default.asp
[32] TPC-H Version 2.6.0: http://www.tpc.org/tpch/default.asp
[33] US Census Bureau, *Unadjusted and Adjusted Estimates of Monthly Retail and Food Services Sales by Kinds of Business:2001*, Department stores (excl.L.D) 4521.
[34] Using lower wattage processors to achieve acceptable performance in an HP Server Based Computing environment September 2007: http://h71028.www7.hp.com/ERC/downloads/4AA1-7946ENW.pdf?jumpid=reg_R1002_USEN
[35] Virtualization- The Most Impactful Solution to Data Center Power Crisis: www.vmware.com/files/pdf/Energy_Efficiency_WP.pdf
[36] Xen: http://www.xen.org/
[37] Forecasting Total Cost of Ownership for Initial Deployments of Server Blades. IDC White Paper: ftp://ftp.compaq.com/pub/products/servers/blades/idc-tco-deployment.pdf
[38] HP BladeSystem c-Class portfolio: http://h18004.www1.hp.com/products/blades/components/c-class-components.html