

Interplay of Energy and Performance for Disk Arrays Running Transaction Processing Workloads

Sudhanva Gurumurthi[†] Jianyong Zhang[†] Anand Sivasubramaniam[†] Mahmut Kandemir[†]
Hubertus Franke[‡] N. Vijaykrishnan[†] Mary Jane Irwin[†]

[†] Dept. of Computer Science and Engineering, The Pennsylvania State University, University Park, PA 16802
{gurumurt,jzhang,anand,kandemir,vijay,mji}@cse.psu.edu

[‡] IBM Research Division, Thomas J. Watson Research Center, Yorktown Heights, NY 10598
frankeh@us.ibm.com

Abstract

The growth of business enterprises and the emergence of the Internet as a medium for data processing has led to a proliferation of applications that are server-centric. The power dissipation of such servers has a major consequence not only on the costs and environmental concerns of power generation and delivery, but also on their reliability and on the design of cooling and packaging mechanisms for these systems. This paper examines the energy and performance ramifications in the design of disk arrays which consume a major portion of the power in transaction processing environments.

Using traces of TPC-C and TPC-H running on commercial servers, we conduct in-depth simulations of energy and performance behavior of disk arrays with different RAID configurations. Our results demonstrate that conventional disk power optimizations that have been previously proposed and evaluated for single disk systems (laptops/workstations) are not very effective in server environments, even if we can design disks than have extremely fast spinup/spindown latencies and predict the idle periods accurately. On the other hand, tuning RAID parameters (RAID type, number of disks, stripe size etc.) has more impact on the power and performance behavior of these systems, sometimes having opposite effects on these two criteria.

1 Introduction

The growth of business enterprises and the emergence of the Internet as a medium for data processing has resulted in a proliferation of applications that are server-centric. Both small-scale businesses and large corporations employ servers for managing inventory and accounts, performing decision-support operations, and hosting web-sites. Traditionally, the main targets for optimization in these environments have been performance (either response time to a single user or improving overall throughput), reliability, and availability. However, power consumption is increasingly becoming a major concern in these environments [6, 4, 10]. Optimizing for power has been understood to be important

for extending battery life in embedded/mobile systems. It is only recently that the importance of power optimization in server environments has gained interest because of the cost of power delivery, cost of cooling the system components, and the impact of high operating temperatures on the stability and reliability of the components.

It has been observed that data-centers can consume several Mega-watts of power [20, 6, 27, 7]. This power is not only essential to keep the processors up and running, but is also drained by the I/O peripherals - particularly the disk subsystem - which in these server environments are employed in abundance to provide high storage capacities and bandwidth. Specifically, most of these servers use some type of RAID disk configuration to provide I/O parallelism. While this parallelism benefits performance, the additional disks drain power at the same time, and it is not clear whether the performance benefits justify the power that they consume. It is important to study these trade-offs when designing I/O subsystems for the data centers, and to our knowledge this is the first paper to do so.

Disk power management has been considered in previous research [9, 21, 16, 15, 23, 13] as an important issue in single disk environments (primarily laptops), where the goal has been to conserve battery power by spinning down the disks during periods of idleness. There are several differences between those studies and the issues under consideration in this paper. First, we are examining systems that employ a large number of disks (RAID configurations) for server systems, rather than single disk systems. Second, the workloads in server environments are significantly different from those on traditional workstation/laptop environments. We have to deal with several users/transactions at the same time, with workloads being much more I/O intensive. The server workloads typically have a continuous request-stream that needs to be serviced, instead of the relatively intermittent activity that is characteristic of the more interactive desktop and laptop environments. This makes it difficult to obtain sufficiently long idle-periods to profitably employ conventional disk power-management schemes which use mode control. Further, response time is a much more critical issue in web-servers and database-servers in the corporate world (compared to desktop workloads), and we need to be careful when there is the issue of degrading performance to gain power savings. In addition, the I/O subsystems in server environments offer sev-

eral more parameters for tuning (RAID configuration, number of disks, striping unit, etc.) as opposed to single disk systems. Typically, these have been tuned for performance, and it is not clear whether those values are power efficient as well. Finally, server disks are also physically different from their laptop and desktop counterparts. They are typically heavier and have much larger spinup and spindown times and are more prone to breakdown when subjected to these mechanical stresses [3] (reliability and availability are of paramount importance in these environments). All these differences in system environments and workload behavior warrant a rethinking of the way power-management needs to be done for these systems.

Transaction processing workloads are amongst the most common and I/O intensive, of the commercial applications. We specifically focus on the TPC-C and TPC-H workloads in this paper [31]. These workloads are extensively used in the commercial world to benchmark hardware and software systems. TPC-C is an On-Line Transaction Processing (OLTP) benchmark, that uses queries to update and lookup data warehouses. TPC-H, in contrast, involves longer-lived queries that analyze the data for decision-making (On-Line Analysis Processing - OLAP).

We carry out a detailed and systematic study of the performance and energy consumption across the RAID design-space (particularly RAID-4, RAID-5 and RAID-10), and examine the interplay between power and performance for these workloads. This is conducted with a trace-driven simulation using the DiskSim [12] simulation infrastructure. DiskSim provides a detailed disk-timing model that has shown to be accurate [11], and we have augmented this infrastructure for power measurements.

We first show that traditional disk power management schemes proposed in desktop/laptop environments are not very successful in these server workloads, even if we can design the disks to spinup and spindown very fast and predict the idle periods accurately. Consequently, the options for disk power management require the investigation and tuning of the different parameters, which this paper undertakes. We demonstrate that tuning the RAID configuration, number of disks, and stripe size has more impact from the power optimization angle. The values of system parameters for best performance are not necessarily those that consume the least power, and vice-versa.

It should be noted that it is possible to have large idle periods during intervals of *light load* (at nights, weekends, etc.). During those times, any power saving technique (even simple heuristics that wait a few minutes before powering down the disks/system) would suffice. Our focus in this work is more at periods of heavier load which is what TPC-C and TPC-H are intended to capture (and which is when more power would be dissipated and cooling becomes more important).

The rest of the paper is organized as follows. Section 2 presents the related work. Section 3 provides a brief overview of the RAID configurations used in this paper. Section 4 describes the workloads and metrics used in the evaluation, along with details of the simulated hardware. Section 5 presents the results of the study and Section 6 summarizes the contributions of this work.

2 Related Work

Disk power management has been extensively studied in the context of single disk systems, particularly for the mobile/laptop environment. Many current disks offer different power modes of operation (active - when the disk is servicing a request, idle - when it is spinning and not serving a request, and one or more low power modes which con-

sume less energy than idle where the disk may not spin). Managing the energy consumption of the disk consists of two steps, namely, detecting suitable idle periods and then spinning down the disk to a low power mode whenever it is predicted that the action would save energy. Detection of idle periods usually involves tracking some kind of history to make predictions on how long the next idle period would last. If this period is long enough (to outweigh spin-down/spinup costs), the disk is explicitly spun down to the low power mode. When an I/O request comes to a disk in the spindown state, the disk first needs to be spun up to service this request (incurring additional exit latencies and power costs in the process). One could pro-actively spin up the disk ahead of the next request if predictions can be made accurately, but many prior studies have not done this. Many idle time predictors use a time-threshold to find out the duration of the next idle period. A fixed threshold is used in [21], wherein if the idle period lasts over 2 seconds, the disk is spun down, and spun back up only when the next request arrives. The threshold could itself be varied adaptively over the execution of the program [9, 16]. A detailed study of idle-time predictors and their effectiveness in disk power management has been conducted in [13]. Lu et al. [22] provide an experimental comparison of several disk power management schemes proposed in literature on a single disk platform.

If we move to high-end server class systems, previous work on power management has mainly focussed on clusters, employing techniques such as shutting off entire server nodes [6], dynamic voltage scaling [4], or a combination of both [10]. There has also been work to reduce the power consumption by balancing the load between different nodes of a cluster [27]. Shutting off nodes is possible if the computational load can be re-distributed or there exist mirrors for the data. The application of voltage scaling can reduce the CPU energy consumption, and has indeed been shown to reduce energy up to 36% for web-server workloads [4]. Investigation of power optimization for SMP servers has looked into optimizing cache and bus energy [24].

Our focus in this paper is on power optimizations for the I/O (disk) subsystem, particularly in the context of transaction processing workloads. In web server workloads (where there are fewer writes), duplication of files on different nodes offers the ability to shut down entire nodes of the web serving cluster (both CPU and its disks) completely. Further, a web page (file) can be serviced by a single node (disks at that node), and the parallelism on a cluster is more to handle several requests at the same time rather than parallelizing each request. On the other hand, database engines use disks for parallelism by striping the data across them, and mirroring is done (if at all) on a much smaller scale. Each database query typically involves several disks for its data processing. Even in data centers which use a cluster for transaction processing, there is usually a Storage Area Network that is accessible by all the nodes, on which the disks reside. In such cases, our work can be applied to manage the disks, while the previous cluster management techniques can be applied to the CPUs. Finally, transaction processing workloads can have very different characteristics from web server workloads [25] - locality between requests is expected to be higher in web servers compared to transaction processing workloads, potentially returning pages from the cache.

There have been some studies that have attempted power management in the context of disk arrays. [32] looked into minimizing the disk energy consumption of a laptop disk by replacing it with an array of smaller form-factor disks. The argument was that the power consumed by a disk is a direct function of its size. However, this study does not look at striping the data across these disks for I/O parallelism, and is thus not applicable to server environments. In [8], the

authors have proposed replacing a tape-backup system with an array of disks that are kept in the spindown state as long as possible. This work targets archival and backup systems, where idleness of the disks is much easier to exploit, and writes overwhelm the reads.

To our knowledge, power management for high performance I/O subsystems with database workloads is largely unexplored. Our work can be applied either to SMP systems (which constitute the bulk of transaction processing servers today) that are directly attached to RAIDs, or to cluster environments which interface to these parallel disks via a Storage Area Network.

3 RAID Overview

Redundant Array of Independent/Inexpensive Disks (RAID) [26] employs a bunch of disks to serve a request in parallel, while providing the view of a single device to the request. If there are n disks, with each disk having a capacity of B blocks, then the RAID address space can be visualized as a linear address-space from 0 to $nB - 1$. The unit of data distribution across the disks is called the *striping-unit* or just *stripe*. A stripe consists of a set of consecutive blocks of user data. Since there are multiple disks present, the reliability of the array can go down. Consequently, RAID configurations use either parity or mirroring for error detection and recovery. There are several RAID configurations based on how the data is striped and how the redundancy is maintained. We consider RAID levels 4, 5, and 10 in this paper (the latter two are among the more popular ones in use today). A more detailed exposition of these RAID levels and our RAID-10 model can be found in [14].

4 Experimental Setup

Before we get to the results of this study, we go over the transaction processing workloads that we use and give details on the platform that is used for the simulation.

4.1 Workloads

As explained earlier, this paper focuses mainly on transaction processing workloads. These workloads use database engines to store, process and analyze large volumes of data that are critical in several commercial environments. Many of these are also back-end servers for a web-based interface that is used to cater to the needs of several hundreds/thousands of users, who need low response times while sustaining high system throughput.

We specifically use two important transaction processing workloads identified by the Transaction Processing Council (TPC) [31]. While ideally one would like to run simulations with the workloads and database engines in their entirety in direct-execution mode, this would take an inordinate amount of time to collect data points with the numerous parameters that we vary in this study. Consequently, we have used device level traces from executions on actual server platforms to drive our simulations.

- **TPC-C Benchmark:** TPC-C is an On-Line Transaction Processing (OLTP) benchmark. It simulates a set of users who perform transactions such as placing orders, checking the status of an order etc. Transactions in this benchmark are typically short, and involve both read and update operations. For more details on this benchmark, the reader is directed to [28]. The tracing was performed for a 20-warehouse configuration with 8 clients and consists of 6.15

million I/O references. The traced system was a 2-way Dell PowerEdge SMP machine with Pentium-III 1.13 GHz processors with 4 10K rpm disks running IBM's EEE DB-2 [17] on the Linux operating system.

- **TPC-H Benchmark:** This is an On-Line Analytical Processing (OLAP) benchmark and is used to capture decision-support transactions on a database [29]. There are 22 queries in this workload, and these queries typically read the relational tables to perform analysis for decision-support. The trace that is used in this study was collected on an IBM Netfinity SMP server with 8 700 Mhz Pentium III processors and 15 IBM Ultrastar 10K rpm disks, also running EEE DB-2 on Linux, and consists of 18 million I/O references.

The above traces are those used in our default configuration, and we have also studied the impact of the dataset size in our experiments. We would like to mention that the overall trends hold across datasets, and the detailed results are omitted here in the interest of space. The traces have been collected at the device level and give the timestamp, type of request (read/write), logical block number and number of blocks. We map the logical block numbers to the disk parameters based on the RAID configuration.

4.2 Simulation Environment

In this study, we use the DiskSim simulator [12], augmented with a disk power model, to study the performance and power implications of RAID configurations on the transaction processing workloads. DiskSim provides a large number of timing and configuration parameters for specifying disks and the controllers/buses for the I/O interface. The default parameters that we use in this study are given in Table 1. The RPM and disk cache size have been chosen to reflect what is popular today for servers. The power values are taken from the data sheets of the IBM Ultrastar 36ZX [19] disk, which is used in several servers. The reader should note that spinup/spindown operations are quite expensive and these values have also been obtained from the data sheets for the IBM Ultrastar 36ZX.

Parameter	Value
Number of Disks: 32	
Stripe Size: 16 KB	
Capacity	33.6 GB
Rotation Speed	12000 rpm
Disk Cache Size	4 MB
Idle Power	22.3 W
Active (Read/Write) Power	39 W
Seek Power	39 W
Standby Power	12.72 W
Spinup Power	34.8 W
Spinup Time	26 secs.
Spindown Time	15 secs.
Disk-Arm Scheduling Algorithm	Elevator
Bus Type	Ultra-3 SCSI

Table 1. Default Disk Configuration Parameters. Many of these have been varied in our experiments.

In our simulated I/O subsystem, the constituent disks of the RAID array are attached to an array-controller using 2 Ultra-3 SCSI buses. In our experiments, half the disks are on each bus (typically, each bus can sustain the bandwidth needs of up to 16 disks). The array controller stripes the data (as per the RAID configuration) across all these disks. It also has an on-board cache, which can potentially avoid some of the disk accesses. The array controller is in turn

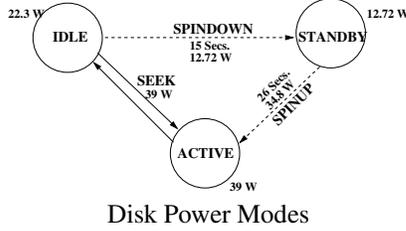


Figure 1. RAID Configuration and Power Modes

interfaced to the main system via the I/O bus, and the power issues for those parts are not studied here.

Figure 1 shows the power mode transitions for each disk. In the experiments where no power management is undertaken, *Active* and *Idle* are the only modes of operation, with the former being used during actual media accesses, and the latter when not performing any operation (disk is only spinning). In the experiments with explicit mode control, we use an additional low power state - *Standby* - where the power consumption is much lower (see Table 1), and a cost is expended to bring the disk to Active state before servicing a request (transitions to and from this state are shown as dashed lines).

The parameters that we vary include the RAID configuration (4, 5 and 10), the number of disks, and the stripe size.

4.3 Metrics

In our evaluation, we use four metrics, namely, *total energy consumption over all the requests* (E_{tot}), *average energy consumption per I/O request* (E), *response-time per I/O request* (T), and *energy-response-time product* ($E \times T$). These can be defined as follows:

- The total energy consumption (E_{tot}) is the energy consumed by all the disks in the array from the beginning to the end of the trace. We monitor all the disk activity (states) and their duration from the start to the end of the simulation, and use this to calculate the overall energy consumption by the disks (integral of the power in each state over the duration in that state).
- The energy consumption per I/O request (E) is E_{tot} divided by the number of I/O requests. A previous study on energy management of server clusters also uses a similar metric (Joules per Operation) for capturing the impact of energy/power optimization for a given throughput [7].
- The response-time (T) is the average time between the request submission and the request completion. This directly has a bearing on the delivered system throughput.
- The product of the previous two ($E \times T$) measures the amount of energy or performance we can tradeoff for the other to have an overall beneficial effect. For instance, if we increase the number of disks in the array, and, get much more improvement in response time than the additional energy consumption, then we can consider this optimization to be a net winner and the product would quantitatively capture this effect. We understand that the energy-delay product requires a complete system characterization to really quantify how the savings of response time in one hardware component can affect the energy of other hardware components (and vice-versa). In this paper, we use this term more to qualify the *relative importance* of energy and response time.

5 Results

In the following subsections, we explore the possible benefits of traditional disk power management using spin-downs for these server workloads and show that there is not much scope for energy savings with this approach. Subsequently, we investigate different hardware parameters that can be tuned for power savings.

5.1 Effectiveness of Conventional Disk Power Management

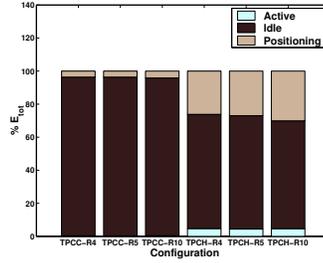


Figure 2. Breakdown of Total Energy Consumption in Different Disk Modes (R4, R5 and R10 refer to RAID-4, RAID-5 and RAID-10 respectively).

In order to optimize disk energy, we need to first examine the energy profile of an actual execution. Figure 2 shows the energy consumption breakdown for the three RAID configurations and two workloads, in terms of the energy consumed when in active mode, idle mode, and during head movements (positioning). It can be seen that contrary to expectations, the least amount of energy is actually spent in the active mode. Most of the energy is expended when the disk is idle, and to a lesser extent for positioning the head. This suggests that one should optimize the idle mode for energy and one possibility is to draw from the ideas of power mode transitions which tries to put the disk in standby mode when it is not performing any operation. This is explored in the next few experiments.

Optimizing disk power for prolonging battery life has been explored in earlier research with smaller form factor disks, by exploiting idle times (by being able to predict them) to spindown the disk. In the following results, we examine how applicable those techniques can be for the server environments under investigation. We first examine the predictability of idle times between disk activities. Next, we examine the duration of idle times to see how much scope is there for employing these techniques. Finally, we use an oracle predictor (that is accurate when predicting idle times both in terms of detecting when an idle period starts and what its duration would be) to see what is the maximum we can hope to gain by these techniques.

5.1.1 The Predictability of Idle-Periods

Prediction of disk requests based on prior history has been a topic of previous research [30]. One commonly used technique is “autocorrelation analysis”, wherein data with good correlations are conducive for fitting with ARIMA models [5] as a time-series. Essentially, an autocorrelation at lag i

is computed between the observation pairs (idle time periods) $y(t)$ and $y(t + i)$ [13]. The resulting values are plotted as a graph for different lags. A sequence that lends itself to easier prediction models is characterized by a graph which has a very high value for small lags and then steeply falls to low values for larger lags (i.e. recent history has more say on the prediction making it easier to construct a model based on these). Note that observations can be negatively correlated as well. Further, there could be some repetitive sequences which can cause spikes in the graph, resulting in deviations from monotonic behavior. The reader is referred to [5] for further explanation on such time-series models.

We have conducted such an autocorrelation analysis of the idle periods of the disks for 50 lags. The resulting lag-plots are found in [14].

Overall, we observe that we do not have good correlation of idle periods. Either the values degrade slowly or degrade sharply but the absolute values are quite low. We show in Table 2, the mean (μ) and standard deviation (σ) across all the disks for the first five lags. As we can observe, except in a few cases, the mean does not cross 0.12 and the standard deviation is not very high either. All these results suggest that it is difficult to get good predictions of idle times based on previous (recent) history. These results are in contrast with those for normal workstation workloads, which have been shown to have higher correlations [13].

Note that, though it is difficult to obtain good predictability of the idle periods using time-series analysis, which relies on the recent past for making predictions, it is possible that good prediction accuracy could be obtained by other means. For example, if the idle periods could be fitted to a probability distribution, it may be possible to predict the duration of an idle period with higher probability. However, as we shall show in section 5.1.3, even with perfect prediction, conventional disk power management does not provide much savings in the energy consumption.

5.1.2 The Duration of Idle-Periods

One could argue that even if we are not able to accurately predict the duration of the next idle period, it would suffice if we can estimate that it is larger than a certain value as far as power management is concerned. In order to ascertain the number of idle-periods that could potentially be exploited for power management, we plotted the idle periods of the disks as a Cumulative Density Function (CDF). These plots can be found in [14].

We observe that, whether it be the overall results or that for an individual disk, idle times are extremely short. In fact, the configurations for TPC-H do not show any visible idle times greater than even 1 second. TPC-C, on the other hand, shows some idle times larger than 2 seconds, but this fraction is still quite small (less than 1% in most cases).

These results indicate that there is not much to be gained with traditional power management techniques, regardless of the predictability of the idle times, if spinup/spin-down times are in the ranges indicated in Table 1.

5.1.3 Limits of Traditional Disk Power Management

Figure 3 (a) shows the maximum energy savings that we can hope to get for these workloads and RAID configurations without any degradation in response times for the spin-down/spin-up values shown in Table 1 for the server disk under consideration. We are assuming an oracle-predictor which has perfect knowledge of the next idle period for this calculation. For TPC-C, performing traditional disk power management actually hurts the overall energy consumption, as the durations of the long idle periods are not sufficient

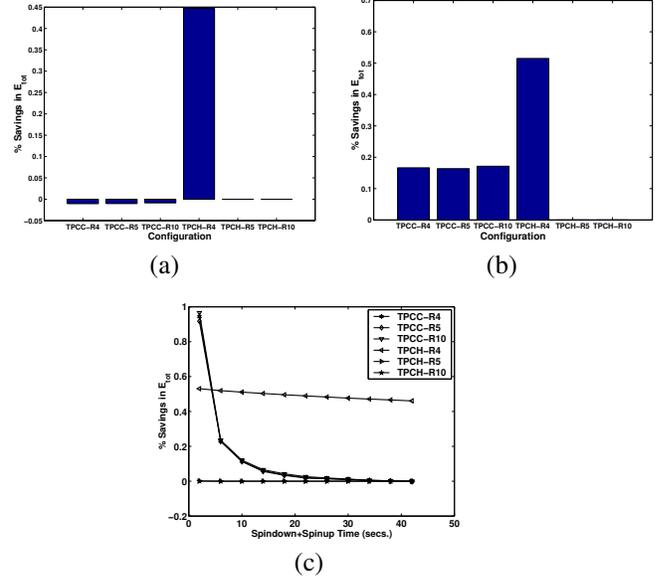


Figure 3. Percentage Savings in the Total Energy Consumption with Disk Spindowns using a perfect idle time prediction oracle. The savings are given for (a) current server class disk (spindown+spinup = 41 secs), (b) an aggressive spinup + spindown value (9 secs, for a state-of-the-art IBM Travelstar disk used in Laptops), and (c) for different spinup + spindown values.

enough to overcome the energy cost of spinning up the disk. Even in the best case (RAID-4 for TPC-H), the percentage energy savings is quite negligible (less than 0.5%).

We have also investigated how this situation would change if we had a laptop-type disk, whose spin-down/spin-up times are typically much lesser than those of server disks. Even when we assume values of 4.5 seconds for spin-down and spin-up (which is in the range of state-of-the-art laptop disks [18]), the energy savings for these workloads are still quite small as is shown in Figure 3 (b). Figure 3 (c) plots the energy savings that can possibly be obtained for different values of spinup+spindown latencies. As can be seen, even if these latencies become smaller than 2 seconds, we get less than 1% improvement in the energy consumption. It is not clear if we can get these numbers down to those values for server class disks without any degradation in performance. Even if we do, these may need very powerful spindle motors, which can in-turn increase the energy consumption.

Despite the high idle power that was shown in Figure 2, we find that idle periods themselves are not very long. The contribution to the idle power is more due to the number of idle periods than the duration of each. This also suggests that it would be fruitful if we can develop techniques to coalesce the idle periods somehow (batching requests, etc.) to better exploit power mode control techniques.

Benchmark	RAID Level	Lag 1		Lag 2		Lag 3		Lag 4		Lag 5	
		μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
TPC-C	RAID-4	-0.064	0.013	0.070	0.020	0.037	0.014	0.045	0.016	0.031	0.009
	RAID-5	-0.044	0.016	0.087	0.019	0.058	0.015	0.057	0.013	0.050	0.016
	RAID-10	0.014	0.020	0.076	0.019	0.057	0.015	0.043	0.014	0.049	0.015
TPC-H	RAID-4	0.066	0.012	0.101	0.017	0.083	0.015	0.090	0.014	0.085	0.015
	RAID-5	0.085	0.011	0.130	0.009	0.115	0.011	0.120	0.010	0.116	0.010
	RAID-10	0.092	0.005	0.139	0.005	0.118	0.005	0.125	0.005	0.121	0.005

Table 2. Autocorrelation Statistics of All Disks Over 5 Lags. For each lag, μ and σ denote the mean and standard-deviation of the autocorrelation at the given lag respectively.

5.2 Tuning the RAID Array for Power and Performance

Since power mode control does not appear very productive when we do not have the flexibility of extending response times, it is then interesting to examine what factors within the I/O architecture can influence its design for power-performance trade-offs. In the following discussion, we look at three important parameters - the RAID configuration (RAID 4, 5 and 10), the number of disks across which the data is striped, and the stripe size. Traditional studies have looked at these parameters only from the performance angle.

5.2.1 Impact of Varying the Number of Disks

Figure 4 shows the T , E , and $E \times T$ (as defined in section 4.3) as a function of the number of disks that are used in the three RAID configurations for the two workloads. Please note that the third graph is normalized with respect to the leftmost point for each line. The energy-response time product has been normalized this way since we are more interested in the trend of a single line than a comparison across the lines. Figure 5 shows the total energy consumption (across all disks) broken down into the active, idle and positioning components. Due the size of the I/O space addressed by the workloads, we chose configurations for RAID 10 starting from 30 disks. We make the following observations from these graphs.

When we examine the performance results, we notice little difference between the three RAID configurations for the TPC-C workload. Though the TPC-C workload has a high amount of write-traffic (14.56% of the total number of requests issued), even for the RAID-4 configuration, beyond 26 disks, there was little variation in the response time. On the other hand in TPC-H, that has a lesser percentage of write requests (8.76%), the RAID-4 configuration showed greater performance sensitivity when the number of disks were increased compared to the other two RAID configurations. This is due to a combination of two factors, namely, the inter-arrival time of the requests and the size of the data accessed per write request. It was found that the average inter-arrival time between requests for RAID-4 TPC-C was 119.78 ms whereas it was 59.01 ms for TPC-H. Further, for TPC-C, 98.17% of the writes spanned atmost one stripe-unit whereas in TPC-H, 99.97% of the writes were over 2 stripe-units. These two factors caused a greater amount of pressure to be put on the parity disk. When the number of disks increases, there is some improvement across the three configurations (due to increase in parallelism), but this improvement is marginal, except for RAID-4 TPC-H. It should be noted that there are some variations when increasing the disks because of several performance trade-offs. The benefit is the improvement in bandwidth with parallelism, and

the downside is the additional overheads that are involved (latency for requests to more disks, and the SCSI bus contention). But these variations are not very significant across the range of disks studied here, and the main point to note from these results is that there is not much improvement in response time beyond a certain number of disks.

On the other hand, the energy consumption keeps rising with the number of disks that are used for the striping. If we look at the detailed energy profile in Figure 5, we observe that most of the energy is in the idle component (as mentioned earlier). When the number of disks is increased, the positioning component does not change much, but the idle component keeps increasing linearly, impacting the overall energy consumption trend as well.

Comparing the two workloads, we find that the active energy (though a small component in both) is more noticeable in TPC-H compared to TPC-C. This is because the former does much more data transfers than the latter. In terms of head positioning power, again TPC-H has a larger fraction of the overall budget, because the number of seeks and average seek distances are higher in this workload [14]. We feel this happens because TPC-H queries can manipulate several tables at the same time, while TPC-C queries are more localized.

One interesting observation that can be seen in both the energy results in Figure 4 and Figure 5 is that the total and the breakdown are comparable across the three RAID configurations for a given number of disks. To investigate this further, we studied the average queue length to different groups of disks for each RAID configuration: (i) for the normal data disks and the parity disk separately in RAID-4, (ii) average over all the disks for RAID-5, and (iii) for each of the two mirrors in RAID-10. The interested reader is referred to [14] for the queue length values.

We saw that the load on the disks (except for the parity disk in RAID-4 which is known to be a bottleneck) across the configurations is comparable, regardless of the number of disks in the range chosen. Since the idle energy is directly related to the load on the disks, and the loads are comparable, the overall energy (which is dominated by the idle energy) and its breakdown are more or less similar across the configurations.

If we are only interested in performance, one can keep increasing the number of disks. This is a common trend in the commercial world where vendors publish TPC results with large disk configurations (even though the improvements may be marginal). On the other hand, power dissipation gets worse with the number of disks. The relative influence of the two factors depends on the nature of the workload. The energy growth is in fact much more influential of these two factors for TPC-C and for RAID-10 in TPC-H, and is the determining factor in the energy-response time product graph. However, the performance gains of using more disks plays a more dominant role in the product for RAID-4 and RAID-5 TPC-H.

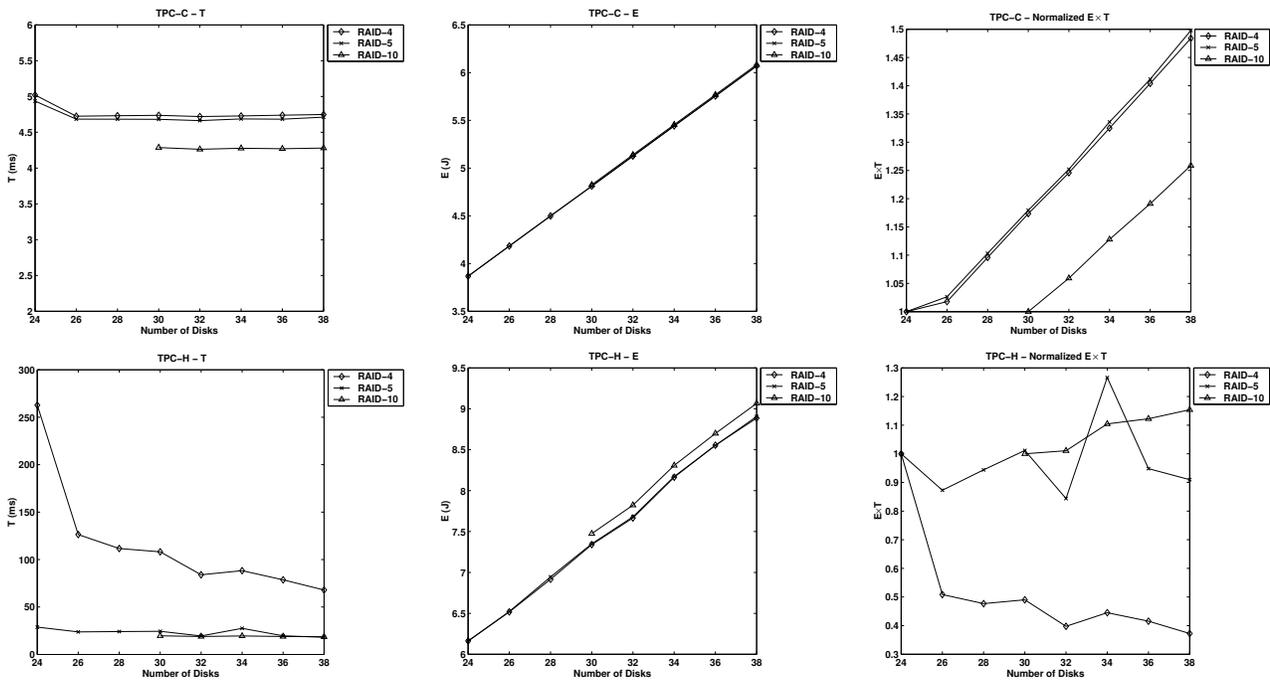


Figure 4. Impact of the Number of Disks in the Array

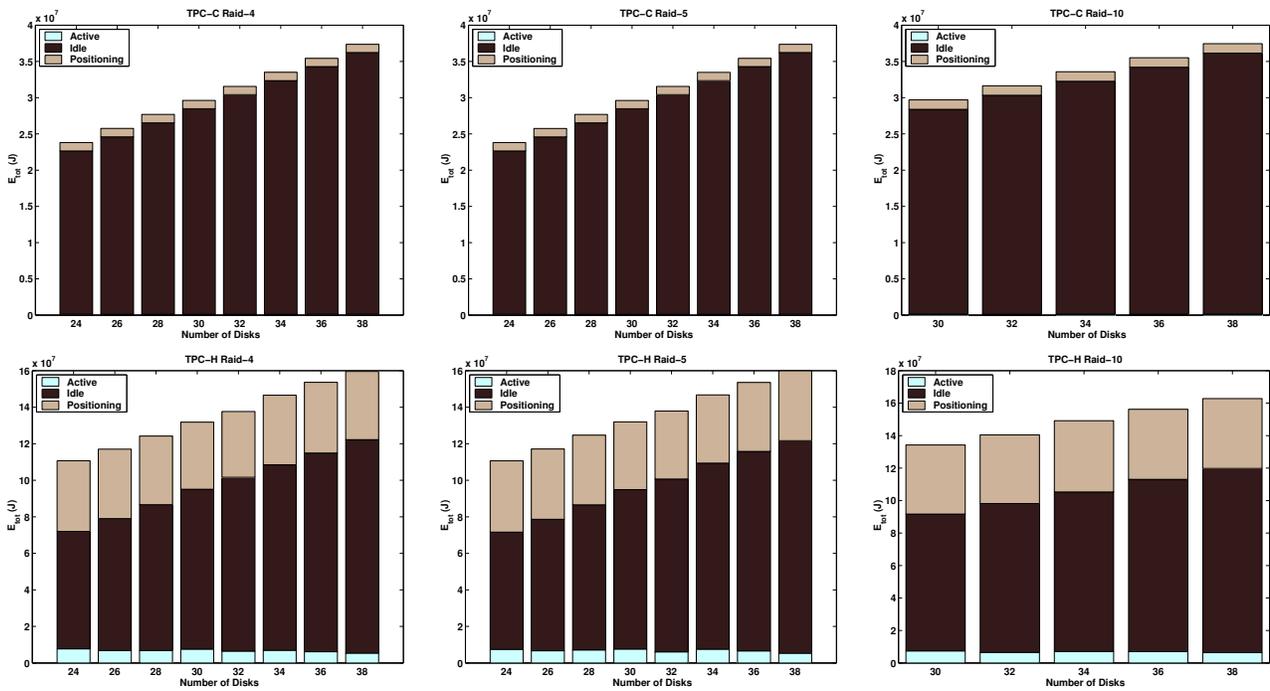


Figure 5. Impact of the Number of Disks in the Array - Breakdown of Total Energy Consumption (E_{tot})

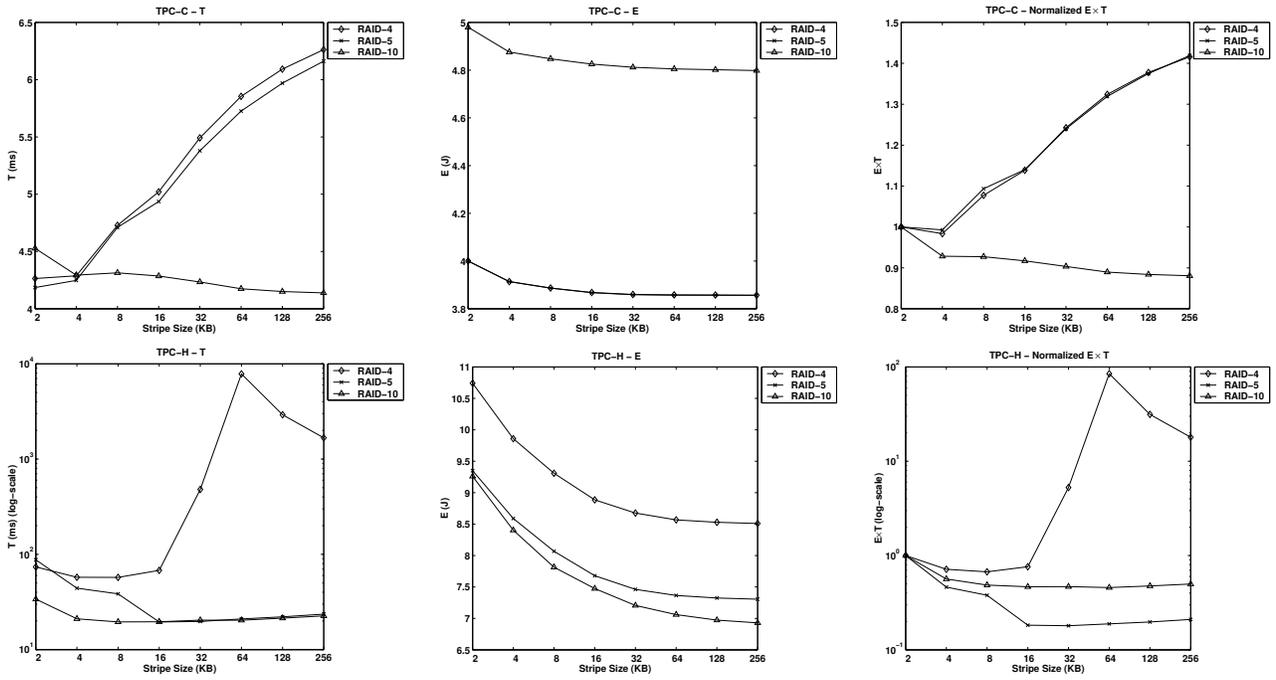


Figure 6. Impact of the Stripe Size

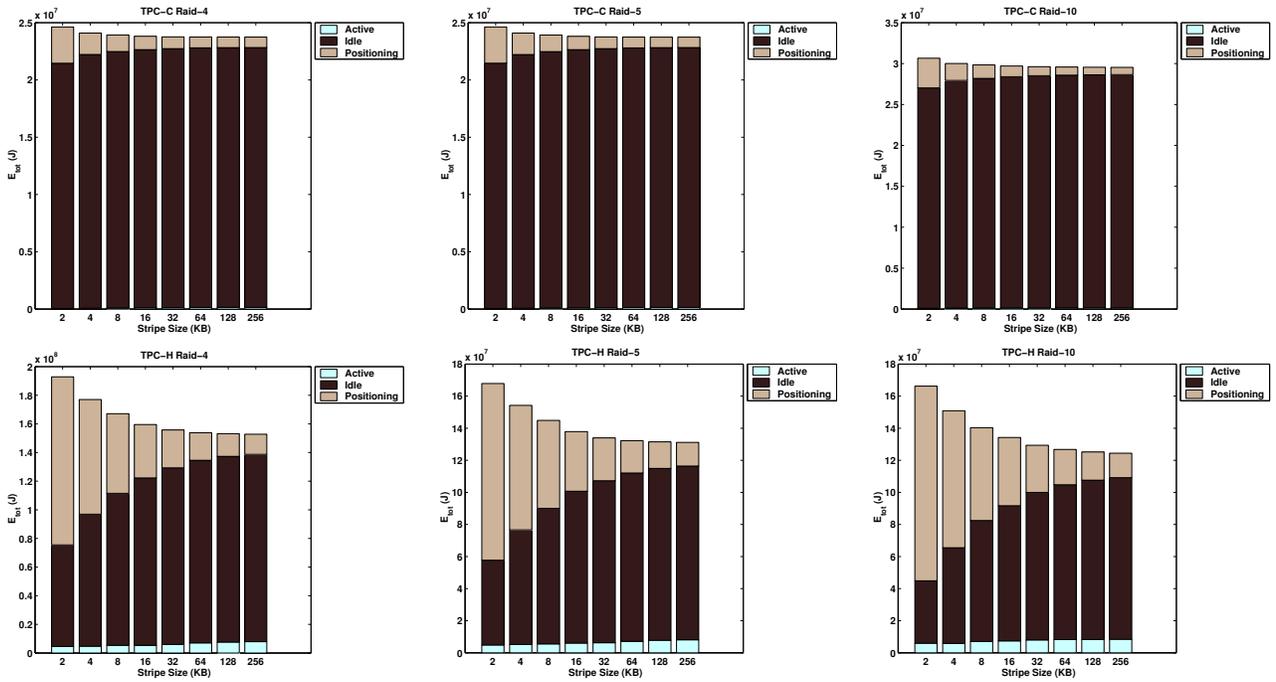


Figure 7. Impact of the Stripe Size - Breakdown of Total Energy Consumption (E_{tot})

5.2.2 Impact of Stripe Size

Figure 6 shows the impact of varying stripe sizes on the three RAID configurations for the two workloads. In these experiments, the number of disks used for each configuration was chosen based on what gave the best energy-response time product in Figure 4 (24 for RAID levels 4 and 5 in TPC-C and 30 for RAID-10; 38,32, and 30 for RAID levels 4,5, and 10 of TPC-H respectively). Figure 7 breaks down the energy consumption in these executions into the idle, active and positioning components.

It is to be expected that a higher stripe size will lead to less head positioning latencies/overhead, providing better scope for sequential accesses and possibly involve fewer disks to satisfy a request. However, this can adversely affect the degree of parallelism, which can hurt performance. We find the latter effect to be more significant in determining the performance. Between the two workloads, TPC-H requests are larger, and hence the point where the adverse effects become more significant tend to shift to the right for TPC-H. Of the RAID configurations, RAID-4 is more susceptible to stripe size changes because the sequentiality problem is higher there, i.e. one disk (the parity) can turn out to become a bottleneck. This becomes worse for TPC-H, which exercises the parity disk to a larger extent due to reasons explained in the previous section, making RAID-4 performance much worse. RAID-5 and RAID-10 for TPC-H are much better (note that the y -axis for TPC-H response time and energy-response time product graphs are in log-scale to enhance the readability of the RAID-5 and RAID-10 lines) though the overall above explanations with regard to stripe size still hold.

Increasing stripe size can lead to fewer disks being involved per request, and higher sequential accesses per disk (reducing seek overheads). Consequently, the head positioning overheads drop, having a consequence on its energy decrease as is shown in the energy profile graphs of Figure 7. This drop in positioning energy causes the overall energy to decrease as well. For both the workloads, the decrease in the energy consumption is not significant beyond a certain point. This is because of two other issues: (i) the idle component for the disks not involved in the transfer goes up (as can be seen in the increase in idle energy), and (ii) the active component for the disks involved in the transfer goes up (the number of accesses per disk does not linearly drop with the increase in stripe size, making this number degrade slower than ideal, while the transfer energy per request grows with the stripe size). These two offset the drop in the positioning component. This effect is much more pronounced for TPC-C compared to TPC-H, since in the latter the positioning overhead is much higher, as mentioned in section 5.2.1. Finally, the reader should note that despite these overall energy changes, the percentage variations are in fact quite small since the scale of the energy graphs in Figure 6 is quite magnified.

The energy-response time product indicates that response time is a much more significant factor in determining stripe size than energy variations components when stripe size is increased.

Different criteria thus warrant a different stripe size. If performance is the only goal, a smaller stripe size of around 4KB appears to be a good choice. If energy is the only goal, then wider stripes of around 256K seem better (though the energy savings may not be significantly better than a smaller stripe). Overall, a stripe size of 4-16K seems a good choice from the energy-response time product perspective.

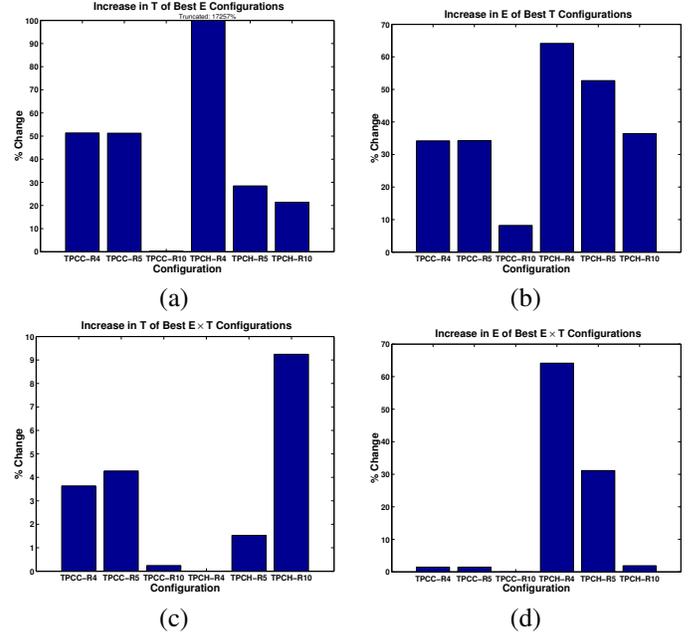


Figure 8. The Effect of Tuning with different Performance and Energy Criteria

Benchmark	RAID Level	Best T	Best E	Best E x T
TPC-C	RAID-4	32/4 KB	24/256 KB	24/4 KB
	RAID-5	32/4 KB	24/256 KB	24/4 KB
	RAID-10	32/4 KB	30/256 KB	30/256 KB
TPC-H	RAID-4	38/8 KB	24/256 KB	38/8 KB
	RAID-5	38/32 KB	24/256 KB	32/32 KB
	RAID-10	38/8 KB	30/256 KB	30/64 KB

Table 3. Optimal Configurations for the Workloads. For each configuration, the pair of values indicated give the number of disks used and the stripe-size employed.

5.2.3 Implications

Having conducted an investigation of the different parameters, we put these results in perspective in Figure 8 which shows the trade-offs between performance tuning and energy tuning. We show four graphs in this figure: (a) the percentage increase (over the best-performing version) in response time for the best-energy (per I/O request) version; (b) the percentage increase (over the best-energy version) in energy consumption for the best-performing version; (c) the percentage increase (over the best-performing version) in response time for the best energy-response time product version; and (d) the percentage increase (over the best-energy version) in energy consumption for the best energy-response time product version. Table 3 gives the configurations that generate the best E , T , and $E \times T$ values. Overall, we observe that performance and energy optimizations can lead to very different choices of system configurations. We would like to mention that overall trends presented in this paper are not very different even when we go for different dataset sizes.

6 Conclusions and Future Work

This paper has conducted an in-depth examination of the power and performance implications of disk arrays (RAIDs) that are used for transaction processing workloads. It has used real traces of TPC-C and TPC-H, and has simulated their execution on three different RAID configurations (RAID-4, RAID-5, and RAID-10) using DiskSim [12] which has been extended with power models. From this detailed examination, this paper makes the following contributions:

- We show that conventional power mode control schemes which have been extensively used in laptop/workstation environments do not show much benefit for these workloads if we do not have the luxury of stretching response times. Even though the idle power is quite high and so is the number of idle periods, the duration of each is rather small. This makes it difficult to offset the high spindown/spinup costs associated with server disks (even if we assume very optimistic costs for these). This is true even if we had a perfect oracle predictor of idle periods. Another problem with frequent spinup/spindown operations is the decrease in mean time between failures, which is an important consideration for server environments.

- On the other hand, with the current state of technology, tuning of RAID parameters has more to gain, and allows the scope for different optimizations - whether power or performance. Increasing the number of disks, though adversely impacts the energy consumption, may buy significant performance benefits depending on the nature of the workload. The choice of stripe size, on the other hand, is more determined by the performance angle than the energy consumption. We found that these parameters had a more determining impact on energy and response time than the choice of RAID configuration itself (particularly RAID-5 and RAID-10) since the load on the disks are comparable.

This research takes a step towards the eventual goal of being able to make good a priori, power-aware RAID design decisions in an automated manner, as suggested in [2, 1]. Our ongoing work is examining issues about extending idle times by possibly batching requests (studying the trade-offs between extending response times and saving energy), and investigating other server workloads (web servers and hosting centers). Another avenue for research that we plan to explore is to design disk arrays that use a combination of disks with different performance and energy consumptions. There are several interesting issues related to reducing the energy consumption without significantly affecting performance by appropriately directing the requests to the different disks, ensuring good load-balance etc.

Acknowledgements

This research has been supported in part by NSF grants: 0103583, 0097998, 9988164, 0130143, 0093082, and 0103583, 0082064, NSF CAREER Awards 0093082 and 0093085, and MARCO 98-DF-600 GSRC.

References

- [1] E. Anderson, M. Kallahalla, S. Spence, and R. Swaminathan. Ergastulum: an approach to solving the workload and device configuration problem. Technical Report HPL-SSP-2001-05, HP Laboratories Storage Systems Program, 2001.
- [2] E. Anderson, R. Swaminathan, A. Veitch, G. Alvarez, and J. Wilkes. Selecting RAID Levels for Disk Arrays. In *Proceedings of the Conference on File and Storage Technology (FAST)*, pages 189–201, January 2002.

- [3] P. Bohrer, D. Cohn, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, R. Rajamony, F. Rawson, and E. Hensbergen. Energy Conservation for Servers. In *IEEE Workshop on Power Management for Real-Time and Embedded Systems*, May 2001.
- [4] P. Bohrer, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, C. McDowell, and R. Rajamony. *The Case for Power Management in Web Servers*, chapter 1. Kluwer Academic Publications, 2002.
- [5] G. Box and G. Jenkins. *Time Series Analysis Forecasting and Control*. Holden-Day, 2nd edition, 1976.
- [6] J. Chase, D. Anderson, P. Thakur, A. Vahdat, and R. Doyle. Managing Energy and Server Resources in Hosting Centers. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP'01)*, pages 103–116, October 2001.
- [7] J. Chase and R. Doyle. Balance of Power: Energy Management for Server Clusters. In *Proceedings of the 8th Workshop on Hot Topics in Operating Systems (HotOS)*, May 2001.
- [8] D. Colarelli and D. Grunwald. Massive Arrays of Idle Disks for Storage Archives. In *Proceedings of Supercomputing*, November 2002.
- [9] F. Douglass and P. Krishnan. Adaptive Disk Spin-Down Policies for Mobile Computers. *Computing Systems*, 8(4):381–413, 1995.
- [10] E. Elnozahy, M. Kistler, and R. Rajamony. Energy-Efficient Server Clusters. In *Proceedings of the Workshop on Power-Aware Computer Systems (PACS'02)*, pages 124–133, February 2002.
- [11] G. Ganger. *System-Oriented Evaluation of I/O Subsystem Performance*. PhD thesis, The University of Michigan, June 1995.
- [12] G. Ganger, B. Worthington, and Y. Patt. *The DiskSim Simulation Environment Version 2.0 Reference Manual*. <http://www.ece.cmu.edu/ganger/disksim/>.
- [13] R. Golding, P. Bosch, and J. Wilkes. Idleness is not sloth. Technical Report HPL-96-140, HP Laboratories, October 1996.
- [14] S. Gurumurthi, J. Zhang, A. Sivasubramaniam, M. Kandemir, H. Franke, N. Vijaykrishnan, and M. Irwin. Interplay of Energy and Performance for Disk Arrays Running Transaction Processing Workloads. Technical Report CSE-02-014, The Pennsylvania State University, October 2002.
- [15] T. Heath, E. Pinheiro, and R. Bianchini. Application-Supported Device Management for Energy and Performance. In *Proceedings of the Workshop on Power-Aware Computer Systems (PACS'02)*, pages 114–123, February 2002.
- [16] D. Helmbold, D. Long, T. Sconyers, and B. Sherrod. Adaptive Disk Spin-Down for Mobile Computers. *ACM/Baltzer Mobile Networks and Applications (MONET) Journal*, 5(4):285–297, December 2000.
- [17] IBM DB2. <http://www-3.ibm.com/software/data/db2/>.
- [18] IBM Hard Disk Drive - Travelstar 40GNX. <http://www.storage.ibm.com/hdd/travel/tr40gnx.htm>.
- [19] IBM Hard Disk Drive - Ultrastar 36ZX. <http://www.storage.ibm.com/hdd/ultra/ul36zx.htm>.
- [20] J. Jones and B. Fonseca. Energy Crisis Pinches Hosting Vendors. <http://iwsun4.infoworld.com/articles/hn/xml/01/01/08/010108hnpower.xml>.
- [21] K. Li, R. Kumpf, P. Horton, and T. E. Anderson. Quantitative Analysis of Disk Drive Power Management in Portable Computers. In *Proceedings of the USENIX Winter Conference*, pages 279–291, 1994.
- [22] Y.-H. Lu, E.-Y. Chung, T. Simunic, L. Benini, and G. Micheli. Quantitative Comparison of Power Management Algorithms. In *Proceedings of the Design Automation and Test in Europe (DATE)*, March 2000.
- [23] Y.-H. Lu and G. Micheli. Adaptive Hard Disk Power Management on Personal Computers. In *Proceedings of the IEEE Great Lakes Symposium*, March 1999.
- [24] A. Moshovos, G. Memik, B. Falsafi, and A. Choudhary. JETTY: Filtering Snoops for Reduced Energy Consumption in SMP Servers. In *Proceedings of the 7th International Symposium on High-Performance Computer Architecture (HPCA)*, January 2001.
- [25] V. Pai. *Cache Management in Scalable Network Servers*. PhD thesis, Rice University, November 1999.
- [26] D. Patterson, G. Gibson, and R. Katz. A Case for Redundant Arrays of Inexpensive Disks (RAID). In *Proceedings of ACM SIGMOD Conference on the Management of Data*, pages 109–116, 1988.
- [27] E. Pinheiro, R. Bianchini, E. V. Carrera, and T. Heath. Load Balancing and Unbalancing for Power and Performance in Cluster-Based Systems. In *Proceedings of the Workshop on Compilers and Operating Systems for Low Power*, September 2001.
- [28] TPC-C Benchmark V5. <http://www.tpc.org/tpcc/>.
- [29] TPC-H Benchmark. <http://www.tpc.org/tpch/>.
- [30] N. Tran. *Automatic ARIMA Time Series Modeling and Forecasting for Adaptive Input/Output Prefetching*. PhD thesis, University of Illinois at Urbana-Champaign, 2002.
- [31] Transaction Processing Performance Council. <http://www.tpc.org/>.
- [32] R. Youssef. RAID for Mobile Computers. Master's thesis, Carnegie Mellon University Information Networking Institute, August 1995.