

Sacrificing Reliability for Energy Saving: Is It Worthwhile for Disk Arrays?

Tao Xie

*Department of Computer Science
San Diego State University
xie@cs.sdsu.edu*

Yao Sun

*Department of Computer Science
San Diego State University
calvin@rohan.sdsu.edu*

Abstract

Mainstream energy conservation schemes for disk arrays inherently affect the reliability of disks. A thorough understanding of the relationship between energy saving techniques and disk reliability is still an open problem, which prevents effective design of new energy saving techniques and application of existing approaches in reliability-critical environments. As one step towards solving this problem, this paper presents an empirical reliability model, called Predictor of Reliability for Energy Saving Schemes (PRESS). Fed by three energy-saving-related reliability-affecting factors, operating temperature, utilization, and disk speed transition frequency, PRESS estimates the reliability of entire disk array. Further, a new energy saving strategy with reliability awareness called Reliability and Energy Aware Distribution (READ) is developed in the light of the insights provided by PRESS. Experimental results demonstrate that compared with existing energy saving schemes, MAID and PDC, READ consistently performs better in performance and reliability while achieving a comparable level of energy consumption.

1. Introduction

A hard disk drive (HDD) is a complex dynamic system that is made up of various electrical, electronic, and mechanical components. A malfunction of any of these components could lead to a complete failure of a hard disk drive. While the capacity, spindle speed, form factor, and performance of hard disk drives have been enhanced rapidly, the reliability of hard disk drives is improving slowly. The primary reasons are that the hard disk manufacturing technology is constantly changing, and that the performance envelope of hard disk drives is incessantly pushed. Although disk drive manufacturers claim that the MTBF (Mean Time Between Failure) of their

enterprise products is more than 1 million hours [25], storage system integrators and end users questioned the unrealistic reliability specification and usually found a much lower MTBF from their field data [8]. Since two decades ago, the need for large-scale storage systems has led to the introduction of configurations such as RAID (Redundant Array of Inexpensive Disks) disk arrays that provide efficient access to large volumes of data. To enhance system reliability, they mainly employ a variety of data redundancy mechanisms like data replication, parity-based protection, and Reed-Solomon erasure-correcting codes. Still, maintaining a high level of reliability for a large-scale storage system with hundreds of thousands of hard disk drives is a major challenge because the very large number of disks dramatically lowers down the overall MTBF of the entire system.

More recently, energy conservation for disk arrays has been an important research topic in storage systems as they can consume 27% of overall electricity in a data center [28]. A broad spectrum of technologies including power management [3][13][30], workload skew [4][21], caching [31], and data placement [28] have been utilized to save energy for disk arrays. While some of energy conservation schemes, such as caching based energy saving approaches, normally do not affect the disk reliability, power management based and workload skew based techniques, two mainstream categories of energy saving schemes for disk arrays, negatively affect the lifetime of disks. For example, power management based energy conservation schemes like Multi-speed [3], DRPM [13], and Hibernator [30] frequently spin up or spin down disk drives, which obviously affects disk drives' lifetime. Besides, workload skew oriented energy conservation techniques such as MAID [4] and PDC [23] utilize a subset of a disk array as workhorses to store popular data so that other disks could have opportunities to have a rest to save energy. Apparently, very high disk utilization is detrimental for the reliability of those overly used disks, whose high

failure rates in turn degrade the reliability of the entire disk array.

Unfortunately, although most of the researchers who proposed the energy-saving schemes above realized that their techniques could inherently and adversely affect the reliability of disks, only a few of them mentioned some intuitive ways such as limiting the power cycling of a disk to 10 times a day or rotating power-always-on disk role, which can alleviate the side-effects of their schemes on disk reliability [31]. Still, a deep understanding of the relationship between energy saving techniques and disk reliability is an open question. Consequently, it is risky and unwise to apply the energy saving schemes that are subject to potential reliability degradation in real storage systems before the following question can be answered: is it worthwhile for disk arrays to save energy at the price of a degraded reliability level?

To answer this question, a reliability model, which can quantify the impact of energy-saving-related reliability affecting factors (hereafter referred to as ESRRA factors) like operating temperature, speed transition frequency, and utilization on disk reliability, is fundamental. In this paper we present an empirical reliability model, called Predictor of Reliability for Energy-Saving Schemes (PRESS), which translates the ESRRA factors into AFR (Annualized Failure Rate). The PRESS model provides us a much needed understanding of the relationship between energy saving and disk reliability. With the PRESS model in hand, storage system administrators can evaluate existing energy-saving schemes' impacts on disk array reliability, and thus, choose the most appropriate one for their applications. Besides, energy-saving technique designers, assisted by PRESS, can devise new energy conservation schemes, which are able to achieve a good balance between energy saving and system reliability. To demonstrate how PRESS can be leveraged to guide the design of reliability-aware energy-saving strategies for disk arrays, we develop a new energy-saving technique called READ (Reliability and Energy Aware Distribution) in the light of the insights provided by PRESS.

The remainder of this paper is organized as follows. In the next section we discuss the related work and motivation. In Section 3, we describe the design of the PRESS model. The READ strategy is presented in Section 4. In Section 5, we evaluate the performance of READ based on real traces. Section 6 concludes the paper with summary and future directions.

2. Related work and motivation

Typical energy conservation techniques for parallel

disk arrays can be categorized into four broad groups: power management, workload skew, caching, and data placement. Power management based and workload skew based schemes are the two most popular categories of energy conservation techniques. Simply shutting down disks after a period of idle time to save energy is not feasible for parallel disk storage systems as they are normally used to serve server class workloads, where idle time slots are usually too small to justify the overhead caused by frequent spin up and spin down [3][13][30]. Therefore, power management mechanisms based on multi-speed disks like DRPM [13], Multi-speed [3], and Hibernator [30] have been proposed so that one can dynamically modulate disk speed to control energy consumption. Essentially, these techniques completely depend on the availability of the underlying hard disk multi-speed modes offered by disk manufacturers. Although real multi-speed (more than 2 speeds) DRPM disks are not widely available in the market yet [21], a few simple variations of DRPM disks, such as a two-speed Hitachi Deskstar 7K400 hard drive, have recently been produced [16].

The basic idea of workload skew based energy conservation techniques is to concentrate the majority of overall workload onto a subset of a disk array so that other disks can have chances to operate in low-power modes to save energy [4][23]. When multi-speed disks are employed, MAID and PDC can significantly save energy with only a small degradation in user response time [23]. The limitation of this type of techniques is that skewed workload adversely affects disk reliability due to the load concentration. Some of energy-saving strategies actually use a combination of several different techniques. For example, when utilizing multi-speed disks, MAID and PDC become hybrid techniques, which integrate disk power management into workload skew.

Existing studies in disk failure analysis and reliability estimation can be generally divided into two camps: manufacturer technical papers and user empirical reports [22]. Disk manufacturers' investigations on disk failure and reliability estimation mainly employ two technologies, mathematical modeling and laboratory testing. Cole estimated the reliability of drives in desktop computers and consumer electronic by using Seagate laboratory test data and Weibull parameters [5]. Shah and Elerath from Network Appliance performed a series of reliability analyses based on field failure data of various drive models from different drive manufacturers [8][27]. The biggest problem in manufacturer papers is an overestimated MTBF of more than one million hours, which is unrealistic and

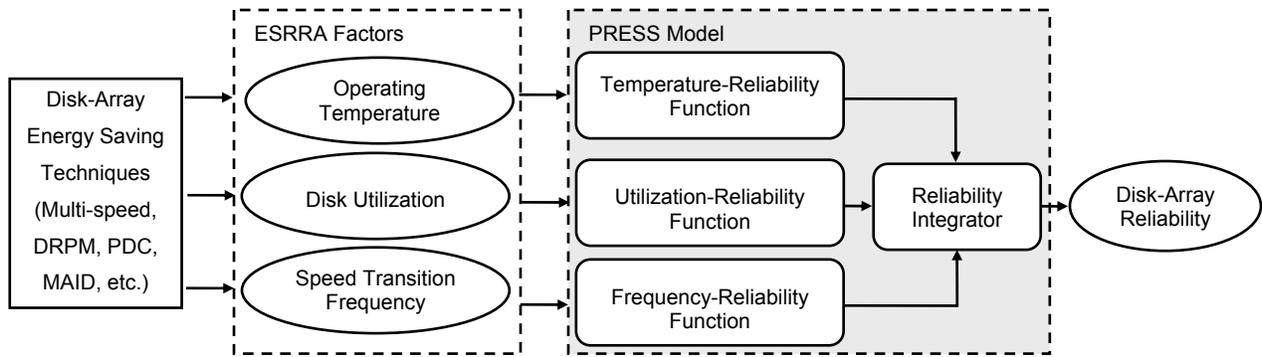


Figure 1. Overall architecture of the PRESS model

misleading. The cause of the problem can be attributed to the limitations of extrapolations from manufacturers' accelerated life experiments [22]. Compared with numerous vendor technical papers, there are only a very few number of user empirical reports. Schwarz et al. found a 2% disk failure rate from total 2,489 disks deployed at the Internet Archive in the complete 2005 year based on their Archive Observatory [26]. Very recently, two pioneer studies from Google [22] and CMU [25] opened up new perspectives for gaining a better understanding of disk failures in large-scale production systems. Schroeder and Gibson in [25] analyzed disk replacement data from several large deployments and observed a largely overstated datasheet MTBF specified by manufacturers. In fact, they found that the annual disk replacement rates in the field are usually in the range from 2% to 4%, which are much higher than manufactures' datasheet annual failure rate. Pinheiro et al. focused on finding how various factors such as temperature and activity level can affect disk drive lifetime [22]. Interestingly, they found a weak correlation between failure rate and either temperature or utilizations, which is against the results from many previous works.

Saving energy and maintaining system reliability, however, could be two conflicting goals. The side-effects of some energy-saving schemes on disk reliability may not be tolerated in reliability-critical applications like OLTP (online transaction processing) and Web services. Thus, a better understanding of the impacts of existing disk array energy-saving schemes on disk reliability is essential. Unfortunately, to the best of our knowledge, little investigation has been concentrated on this particular problem. Motivated by the importance of this largely ignored problem, in this work we study the effects of energy-saving schemes on disk reliability.

3. The PRESS model

3.1. An overview of the PRESS model

While a number of factors such as age, model, vintage, and altitude could affect disk reliability, operating temperature, disk utilization, and disk speed transition frequency are identified as major ESSRA factors based on our investigations. Thus, in this research we assume that all disks under investigation are the same in all non-ESSRA factors. Furthermore, we assume that all disks are older than 1 year, and therefore, the infant mortality phenomena will not be considered in this study.

Figure 1 depicts the overall architecture of the PRESS model. Energy-saving schemes such as DRPM [13], PDC [23], and MAID [4] inherently affect either part of the three ESSRA factors or all of them. Each of the three ESSRA factors is then fed into a corresponding reliability estimation function within the PRESS model. The PRESS model is composed of a reliability integrator module and three functions: temperature-reliability function, utilization-reliability function, and frequency-reliability function. While the former two functions are derived based on Google's results in [22], the last one is built from the spindle start/stop failure rate adder suggested by the IDEMA standards [17] and the modified Coffin-Manson model. Each of the three reliability functions individually outputs its estimated reliability values in AFR (Annualized Failure Rate), which then become the inputs of the reliability integrator module. The reliability integrator module translates the outputs of the three functions into a single reliability value for a disk array.

3.2. Operating temperature

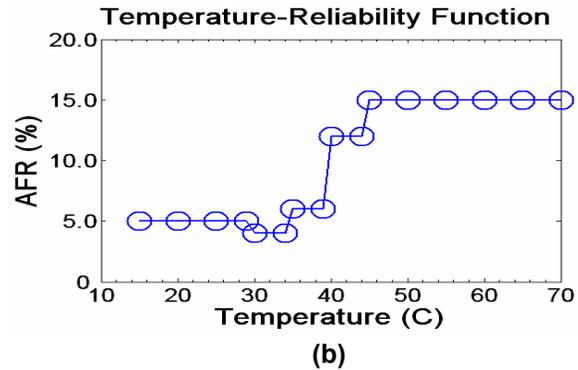
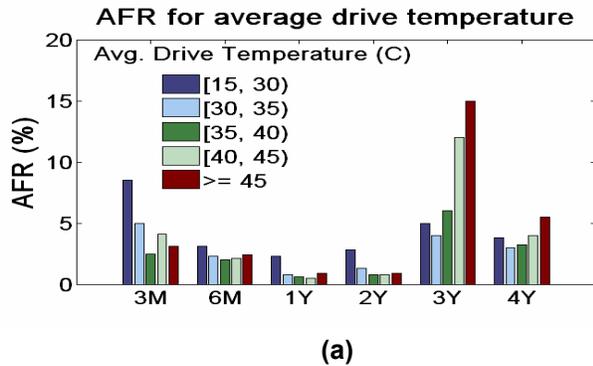


Figure 2. (a) Temperature impacts on AFR from [22]; (b) The temperature-reliability function

Operating temperature has long been believed as one of the most significant factors that affect disk reliability. This belief is supported by many previous investigations on the relationship between temperature and disk reliability [1][5][15][17]. High temperature was discovered as a major culprit for a number of disk reliability problems [15]. One such problem is off-track writes, which could corrupt data on adjacent cylinders. Even worse, spindle motor and voice coil motor running at high temperatures can lead to head crash [15]. Results from Seagate based on mathematical modeling and laboratory testing indicate that disk failure rate doubles when temperature increases 15 C [5]. More recently, research outcomes from Google using field collected data confirmed that disk operating temperature generally has observable effects on disk reliability, especially for older disks in high temperature ranges [22].

There are two different avenues to establishing a temperature-reliability relationship function. One is using mathematical modeling and laboratory testing techniques [5] and the other is employing user field data [22]. Although the two ways can provide us with temperature-reliability relationship functions with a similar trend, i.e., higher temperatures usually result in higher AFR, we selected the latter because field data mining is a more realistic, though not perfect, way to estimate disk reliability due to sufficient amount of failure statistics from real disk deployments.

In this study we only consider a simple type of multi-speed disks, namely, two-speed disks. We assume that the low speed mode is 3,600 RPM (revolutions per minute) and the high speed mode is 10,000 RPM. It is understood that operating temperature of a disk is affected by workload characteristics and several disk drive parameters like drive geometry, number of platters, RPM, and materials used for building the drive [18]. The change of RPM, however, becomes a primary influence on a disk's temperature when all other factors mentioned

above remain the same. This is because disk heat dissipation is proportional to nearly the cubic power of RPM [18]. Therefore, the increase of RPM results in excessive heat, which in turn leads to a higher temperature. Since there is no explicit information about the relationship between RPM and disk temperature, we derive temperatures of two-speed disks at 3,600 RPM and 10,000 RPM based on reported related work. The experimental results in [12] show that a Seagate Cheetah disk drive reaches a steady state of 55.22 C when running at 1,5000 RPM after 48 minutes. Considering that 10,000 RPM is only 2/3 of the disk's rotation speed, we argue that [45, 50] C is a reasonable temperature range for the high speed mode. Another experimental report [14] indicates that on average the temperature of a hard disk drive with 5400 RPM is 37.5 C. As a result, our assumption that the low speed mode 3,600 RPM is associated with a temperature falling in the range [35, 40] C is feasible.

Now we explain why we adopted the 3-year temperature-AFR statistics from [22] as our temperature-reliability function. One can easily makes the following two observations on Figure 2a, which is the Figure 5 in [22]. First, higher temperatures are not associated with higher failure rates when disks are less than 3 years old. Second, the temperature effects on disk failure rates are salient for the 3-year-old and the 4-year-old disks, especially when temperatures are higher than 35 C. The authors of [22] explain the reason of the first observation is that other effects may affect failure rates much more strongly than temperatures do when disks are still young. However, we have a different interpretation of this phenomenon. We believe that higher temperatures still have strong negative effects on younger disks as they do on older disks. The impacts of higher temperature on younger disks do not immediately turn out to be explicit disk failures just because the impacts-to-failure procedure is essentially an accumulation process and it takes some time. After all, higher temperatures make electronic

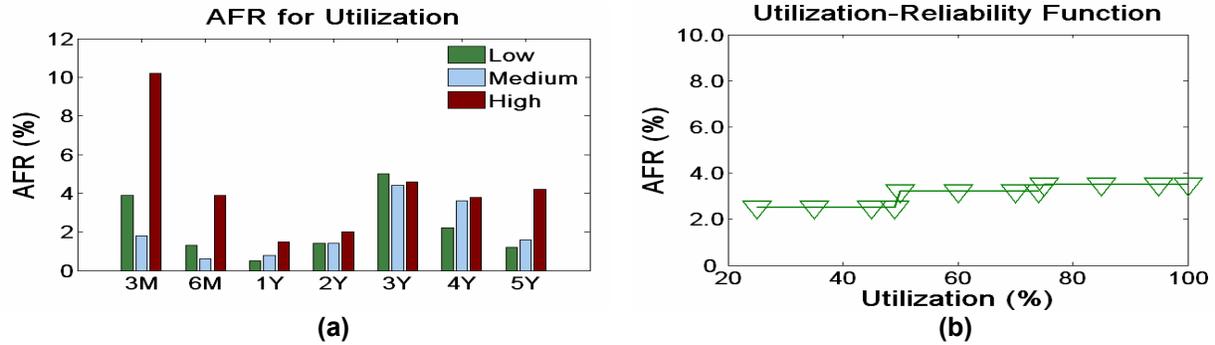


Figure 3. (a) Utilization impacts on AFR from [22]; (b) The utilization-reliability function

and mechanical components of disks more prone to fail prematurely [15]. The second observation, i.e., obvious higher failure rates associated with higher temperature ranges for 3-year-old disks, supports our explanation because earlier high temperature impacts on disks are eventually transformed into disk failures after one or two years. Therefore, we ignore the temperature-AFR results in [22] for disks younger than 3 years as they hide the temperature impacts on disk reliability. Although both 3-year-old disks and 4-year-old disks exhibit a high correlation between higher temperatures and higher failure rates, we finally decided to select 3-year-disk temperature-AFR data as the foundation of our temperature-reliability function. The primary reason is that the relationship between higher temperatures and AFR for 3-year-old disks fully demonstrates that higher temperatures have a prominent influence on disk failure rates because after 2-year higher temperature “torture” an observable number of disks fail in the third year. Apparently, these disk failures, which are originated in the first two years, should be included in the third year’s AFR. On the other hand, the 4-year-old disk results substantially lose the “hidden” disk failures, and therefore are not complete.

3.3. Disk utilization

Disk utilization is defined as the fraction of active time of a drive out of its total power-on-time. Since there is no enough detail in their measurements, researchers of [22] measured utilization in terms of weekly averages of read/write bandwidth for each drive and roughly divided them into three categories: low, medium, and high. Still, they found that using number of I/O operations and bytes transferred as utilization metrics provide very similar results [22]. Thus, we conclude that it is feasible to take the average bandwidth metric as the utilization metric because the number of I/O operations and bytes transferred of a disk are proportional to disk active time. Therefore, in

our utilization-reliability function we use the utilization metric in the range [25%, 100%] in stead of low, medium, and high employed in Figure 3 of [22]. We define low utilizations as utilizations in the range [25%, 50%). Similarly, a medium utilization is defined as a utilization within the scope [50%, 75%), whereas a high utilization falls in the range [75%, 100%).

The relationship between utilization and disk reliability has been investigated previously [2][5][22]. A conclusion that higher utilizations in most cases affect disk reliability negatively has been generally confirmed by two widely recognized studies. One is a classical work from Seagate, which utilized laboratory testing and mathematical modeling techniques [5]. The other is a new breakthrough, which analyzes the utilization impacts on disk reliability based on field data from Google [22]. Authors of [22] measured 7 age groups of disks (3-month, 6-month, 1-year, 2-year, 3-year, 4-year, 5-year, see Figure 3 in [22]) and found that only 3-year-old group exhibits an unexpected result, i.e., low utilizations result in a slightly higher AFR than higher utilizations do. The two explanations for this “bizarre” behavior provided by [22] are not convincing in our views. Their first explanation is the *survival of the fittest* theory. They speculate that the drives that survive the infant mortality phase are the least susceptible to the failures caused by higher utilizations, and result in a population that is more robust with respect to variations in utilization levels [22]. If this is the case, they cannot explain why the results from the 4-year-old disk group and the 5-year-old disk group immediately restore the “wired” behavior to a “normal” one, i.e., higher utilizations correlate to higher AFR. The second explanation they made is that previous results such as [5] can only better model early life failure characteristics, and thus, it is possible that longer term population studies could discover a less significant effect later in a disk’s lifetime. Again, if this is true, they cannot explain why we still see a noticeable higher utilization with higher AFR behavior for disks in their age 4 and 5. In fact,

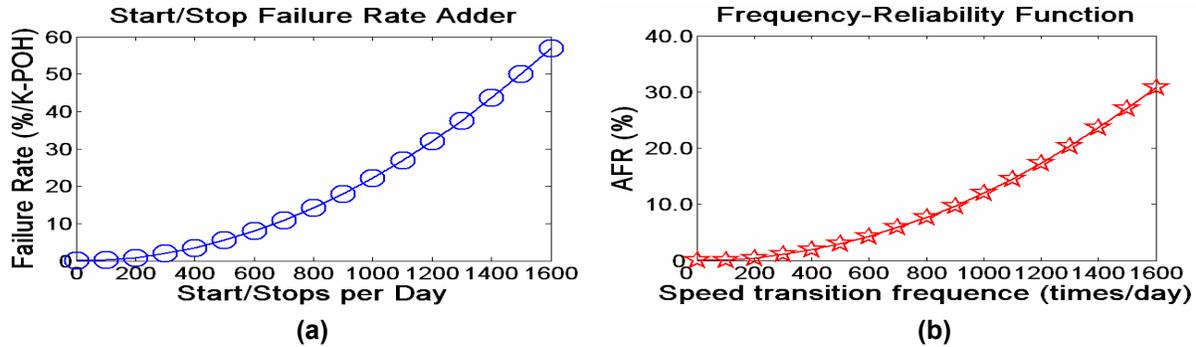


Figure 4. (a) Start/Stop failure rate adder [17]; (b) The frequency-reliability function

their second explanation conflicts with their observation that only very young and very old age groups show the expected behavior. Based on our observations on Figure 3a, we argue that a reasonable explanation for this unexpected behavior is that disk drives in their middle ages (2 or 3 years) are strong enough in both electronic and mechanical parts to resist the effects of higher utilizations. Therefore, AFR of disks in these two age groups has little correlation with utilization. Our speculation is supported by the evidence that failure rates of different utilization levels are very close to each other for disks in these two age groups and failure rate distribution exhibits some randomness. We selected the results from 4-year-old disk group as our utilization-reliability function mainly because that (1) we only consider disks older than 1 year; (2) results from 2-year and 3-year groups cannot provide any explicit utilization impacts on disk reliability although many previous research confirms that these impacts do exist; (3) 5-year results are less useful because disks normally only have five year warranty; and (4) the results from 4-year disks match the reliability versus duty cycle outcomes of [5].

3.4. Disk speed transition frequency

The disk speed transition frequency (hereafter called frequency) is defined as the number of disk speed transitions in one day. Establishing a frequency-reliability function is the most difficult task in this research primarily because multi-speed disks have not been largely manufactured and deployed. Thus, no result about the impacts of frequency on disk reliability has been reported so far. Although the applications of multi-speed disks are still in their infancy, we believe that they will no doubt have a huge impact on energy-saving for disk-based storage systems in the not-so-distant future. Therefore, now it is the time to obtain a basic understanding of the relationship between frequency and reliability. Our frequency-reliability function is built on a combination of the spindle

start/stop failure rate adder suggested by IDEMA [17] and the modified Coffin-Manson model.

We start our investigation on this challenging issue from a relevant disk usage pattern parameter, namely, spindle start/stop rate (SSSR), defined as the times of spindle start/stop per month [8][17]. The rationale behind is that disk speed transitions and spindle start/stops essentially generate the same type of disk failure mode, spindle motor failure, though with different extents. A disk reliability report discovered that each spindle start-and-stop event causes some amount of fatigue to occur at the heads and the spindle motor [24]. In fact, spindle motor failure is one of the most common disk drive failure modes [19]. That is why disk drive manufacturers normally set 50,000 as the start/stop cycle limit and suggest no more than 25 power cycles per day to guarantee specified performance. A disk speed transition event could cause a similar reliability issue as a spindle start/stop occurrence does because speed transitions incur some amount of fatigue, noise, heat dissipation and vibration as well [19]. We believe, however, the degree of reliability impacts caused by speed transitions is relatively lower than that of caused by spindle start/stops. The reason is two-fold. First, during a start up process, a spindle has to increase its speed from zero to maximum. However, a speed transition event, e.g., from a low speed to a high speed, only needs to promote spindle's speed from its current value to an immediate higher value. Therefore, the costs of a speed transition between two contiguous speed levels in terms of energy consumption and time are less than that of a spindle start/stop, which in turn brings a disk drive less heat dissipation, a main reason for fatigue. Second, there is no salient peak power issue associated with speed transitions. It is understood that peak power within a short period of time is detrimental to disk reliability [19].

Both start/stop events and disk speed transitions incur temperature cycling, the main cause of fatigue failures [7]. The damage caused by temperature

cycling accumulates each time a hard disk drive undergoes a power cycle or a speed transition. Such cycles induce a cyclical stress, which weakens materials and eventually makes the disk fail [7]. We utilize the modified Coffin-Manson model (Equation 1) because it is a widely-used model, which works very well for failures caused by material fatigues due to cyclical stress [9]. It evaluates the reliability effects of cycles of stress or frequency of change in temperatures. The Arrhenius equation involved describes the relationship between failure rate and temperature for electronic components (Equation 2).

The spindle start/stop failure rate adder curve presented by IDEMA is re-plotted as Figure 4a. It indicates, for example, a start/stop rate of 10 per day would add 0.15 to the AFR for disks older than one year. Since IDEMA only gives the curve in a start/stop frequency range [0, 350] per month, we extend it to [0, 1600] per day using quadratic curve fitting technique. We derive our frequency-function based on Figure 4a and the modified Coffin-Manson model, which is listed as Equation 1 as below:

$$N_f = A_0 f^{-\alpha} \Delta T^{-\beta} G(T_{\max}), \quad (1)$$

where N_f is the number of cycles to failure, A_0 is a material constant, f is the cycling frequency, ΔT is the temperature range during a cycle, and $G(T_{\max})$ is an Arrhenius term evaluated at the maximum temperature reached in each cycle. Typical values for the cycling frequency exponent α and the temperature range exponent β are around -1/3 and 2, respectively [9]. The term $G(T_{\max})$ can be calculated using the following Arrhenius equation [9]:

$$G(T) = A e^{(-E_a / KT)}, \quad (2)$$

where A is a constant scaling factor, E_a is the activation energy, K is the Boltzmann's constant (i.e., 8.617×10^{-5}), and T is the temperature measured in degrees Kelvin (i.e., $273.16 +$ degrees in Celsius) at the point when the failure process takes place.

We first demonstrate how we derive the value of $G(T_{\max})$ using Equation 2. As we discussed in Section 3.2, the maximum disk operating temperature is set to 50 C when a disk is running at its high speed at 10,000 RPM. Thus, T_{\max} is equal to $273.16+50=323.16$ Kelvin. Also, E_a is suggested to be 1.25 [9]. Therefore, $G(T_{\max})=A*3.2275 \times 10^{-20}$. Next, we show how we obtain $A*A_0$, the product of the two constants A and A_0 using Equation 1. Since the suggested daily power cycle limit is 25, we set f equal to 25. Also, the temperature gap from an ambient temperature 28 C to the maximum operating temperature 50 C is 22 C, which means that ΔT is equal to 22. Besides, we know that the maximum number of power cycles specified in

a disk datasheet is normally 50,000. We let N_f be 50,000. Consequently, based on Equation 1, we obtain $A*A_0 = 2.564317 \times 10^{26}$. Now we calculate N'_f , the number of speed transitions to failure assuming that the number of speed transitions per day is 25, the same as daily power cycle limit. Here, the temperature T_{\max} is set to 45 C, the midway value of the low temperature 40 C and the high temperature 50 C (see Section 3.2). The reason is that speed transition is bi-directional in the sense that a speed transition could either increase or decrease disk temperature. Now ΔT in Equation 1 is equal to 10 because this is the gap between the low temperature range and the high temperature range (see Section 3.2). Based on Equation 1 and the calculated value of $A*A_0$, we conclude that N'_f is equal to 118529, the number of disk speed transitions to failure, which is roughly twice of N_f , the maximum number of power cycles suggested. We view this as strong evidence that a disk speed transition can cause about 50% effects on reliability as that of incurred by a spindle start/stop. Therefore, we scale down the spindle start/stop failure rate adder curve (Figure 4a) by half and change the unit of the X axis to times per day to obtain our frequency-reliability function (Figure 4b). The expression based on quadratic curve fitting for the reliability-frequency function is Equation 3, where R is the reliability in AFR and f is the disk speed transition frequency.

$$R(f) = 1.51e^{-5}f^2 - 1.09e^{-4}f + 1.39e^{-4}, f \in [0, 1600] \quad (3)$$

3.5. PRESS it all together

The reliability integrator module in Figure 1 has two functions. First, it combines the three reliability functions together to establish the PRESS model, which can then be used to predict the reliability for each single disk in a disk array. More specifically, the PRESS model estimates an AFR value for a single disk after its values of the three ESRR factors have been provided. Since we 3-dimensional people have no 4-dimensional perspective, we present two 3-dimensional figures to represent the PRESS model at operating temperature 40 C (Figure 5a) and 50 C (Figure 5b), respectively. In Section 3.2, we justified why it is reasonable to set the temperature range [35, 40] C for disk speed 3,600 RPM and the temperature range [45, 50] C for disk speed 10,000 RPM. Within these feasible temperature range settings, we suppose that disks in low speed have operating temperature 40 C, whereas disks in high speed are at 50 C. Second, after obtaining AFR for each disk in a disk array, the reliability integrator module outputs the AFR of the least reliable disk as the overall reliability for the entire

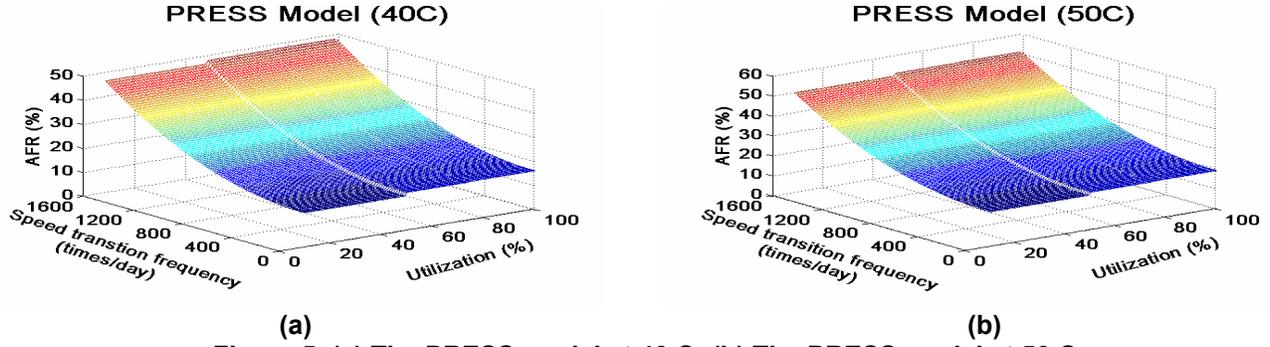


Figure 5. (a) The PRESS model at 40 C; (b) The PRESS model at 50 C

disk array. We argue that the reliability level of a disk array is only as high as the lowest level of reliability possessed by a single disk in the array.

The PRESS model yields several important insights on how to make trade-offs between energy-saving and reliability when developing energy conservation techniques for disk array systems. First, disk speed transition frequency is the most significant reliability-affecting factor among the three ESSRA factors. Based on our estimation in Section 3.4, the number of disk speed transitions should be limited to less than 65 ($118529/5/365 \approx 65$) per day in order to guarantee a 5-year performance warranty. Thus, it is not wise to aggressively switch disk speed to save some amount of energy. We argue that the high AFR caused by a high speed transition frequency would cost much more than the energy-saving gained. Normally, the value of lost data plus the price of failed disks substantially outweigh the energy-saving gained. Thus, it is not worthwhile for disk arrays to save energy by frequently switching disk speed. Next, operating temperature is the second most significant reliability-affecting factor. A high temperature can be caused by a long time running at high speed. Hence, workload-skew based energy-saving schemes need to rotate the role of workhorse disks regularly so that the scenario that a particular subset of disks is always running at high temperature can be prevented. Finally, since the differences in AFR between high utilizations and medium utilizations are slim, an uneven utilization distribution in an array should not be overly concerned.

In fact, how to quantitatively and accurately measure reliability impacts caused by various factors is still an open question [5][8][22][25]. The PRESS model is only a step towards finding a way to quantitatively approximating the reliability effects imposed by the three ESSRA factors. We believe that our model is reasonable due to the following two reasons. First, the foundation of our PRESS model is solid. Our temperature-reliability function and

utilization-reliability function come from a state-of-the-art work [22], which studies the impacts of the two factors on disk reliability based on field data from a large disk population over 5 years. In addition, the modified Coffin-Manson model, which captures the relationship between failure and cyclic stress, has been used successfully to model materials fatigue failures due to repeated temperature cycling as device is turned on and off [7][9]. Second, although the measurements of reliability in terms of AFR are not completely objective in the PRESS model, the improvements of our READ algorithm compared with two existing energy-saving approaches in terms of reliability are still valid because all algorithms are evaluated by using the same set of reliability functions under the same conditions.

4. The READ strategy

Several previous studies [6][11] show that the distribution of web page requests generally follows a Zipf distribution [20] where the relative probability of a request for the i 'th most popular page is proportional to $1/i^\alpha$, with α typically varying between 0 and 1. Inspired by the observations of this highly skewed data popularity distribution, two traditional energy-saving techniques, MAID [4] and PDC [23], concentrate the majority of workload onto a subset of a disk array so that other disks can have chances to operate in low-power modes to save energy. PDC dynamically migrate popular data to a subset of the disks so that the load becomes skewed towards a few of the disks and others can be sent to low-power modes [23]. Since only a small portion of data would be accessed at a given time, the idea of Massive Array of Idle Disks (MAID) [4] is to copy the required data to a set of "cache disks" and put all the other disks in low-power mode. Later accesses to the data may then hit the data on the cache disk(s). A common goal of both PDC and MAID is to increase idle times by rearranging data among the disk array and lower the disks' speed down

Input: A disk array D with n 2-speed disks, a collection of m files in the set F , an epoch P , idleness threshold H , a disk maximum allowed times of speed transitions per day S , speed transition times for each disk $T(n)$, and the skew parameter θ

Output: A file allocation scheme $X(m)$ for each epoch P

1. Use Eq. 4 to compute the number of popular files and the number of unpopular files
2. Use Eq. 5 to compute γ , the ratio between the number of hot disks and the number of cold disks
3. Hot disk number $HD = \frac{\gamma * n}{\gamma + 1}$, cold disk number $CD = n - HD$, $d_h=1$, $d_c=1$
4. Configure HD of n disks to high speed mode and set CD of n disks to low speed mode
5. Sort all files in file size in a non-decreasing order
6. Assign all popular files onto the hot disk zone in a round-robin manner
7. Assign all unpopular files onto the cold disk zone in a round-robin manner
8. **for** each epoch P **do**
9. Keep tracking number of accesses for each file
10. Re-sort all files in number of accesses during the current epoch
11. Re-calculate the skew parameter θ and re-categorize popular and unpopular for each file
12. **for** each previously hot file that becomes unpopular **do**
13. Migrate it to the cold disk zone
14. Update its record in the allocation scheme X
15. **end for**
16. **for** each previously cold file that becomes popular **do**
17. Migrate it to the hot disk zone
18. Update its record in the allocation scheme X
19. **end for**
20. **for** each disk $d_i \in D$ **do**
21. **if** $S/2 \leq T(d_i)$ // Still has room in terms of disk speed transitions to spin down
22. $H=2H$; // Double the idleness threshold H to reduce future disk speed transitions
23. **end if**
24. **end for**
25. **end for**

Figure 6. The READ strategy

[23]. None of the two algorithms applied any mechanisms to limit the reliability impacts introduced by them.

Our READ strategy is motivated by data popularity locality as well and it employs data redistribution and multi-speed disks. We adopted several similar assumptions that PDC used. We also assume that each request accesses an entire file, which is a typical scenario for Web, proxy, ftp, and email server workloads [23]. In addition, the distribution of requests generally follows a Zipf-like distribution with α in the range $[0, 1]$. Also, each file is permanently stored on one disk and there is no stripping or mirroring is used [23]. We decide not to use stripping for two reasons. One is that we want to make the comparisons between READ and the two conventional algorithms in a fair manner as they didn't employ stripping. The other is that the average file sizes in the real web workload are much smaller than a normal stripping block size 512 KB. Further, no requests can be served when a disk is switching its speed. The general idea of READ is to control disk speed transition frequency based on the statistics of the workload so that disk array reliability can be guaranteed. Also, READ employs a dynamic file redistribution scheme to periodically redistribute files across a disk array in an even manner to generate a more uniform disk utilization distribution. A low

disk speed transition frequency and an even distribution of disk utilizations imply a lower AFR based on our PRESS model.

The set of files is represented as $F = \{f_1, \dots, f_u, f_v, \dots, f_m\}$. A file f_i ($f_i \in F$) is modeled as a set of rational parameters, e.g., $f_i = (s_i, \lambda_i)$, where s_i , λ_i are the file's size in Mbyte and its access rate. In the original round of file distribution, READ orders the files in terms of file size because we assume that the popularity in terms of access rate of a file is inversely correlated to its size. And then READ splits the file set into two subsets: *popular file set* $F_p = \{f_1, \dots, f_h, \dots, f_u\}$ and *unpopular file set* $F_u = \{f_v, \dots, f_c, \dots, f_m\}$ ($F = F_p \cup F_u$ and $F_p \cap F_u = \emptyset$). Next, a disk array storage system consists of a linked group $D = \{d_1, \dots, d_e, d_f, \dots, d_n\}$ of n independent 2-speed disk drives, which can be divided into a *hot disk zone* $D_h = \{d_1, \dots, d_h, \dots, d_e\}$ and a *cold disk zone* $D_c = \{d_f, \dots, d_c, \dots, d_n\}$ ($D = D_h \cup D_c$ and $D_h \cap D_c = \emptyset$). Disks in the hot zone are all configured to their high speed modes, which always run in the high transfer rate t^h (Mbyte/second) with the high active energy consumption rate p^h (Joule/Mbyte) and the high idle energy consumption rate i^h (Joule/second). Similarly, disks in cold zone are set to their low speed modes, which continuously operate in the low transfer rate t^l (Mbyte/second) with the low active energy consumption rate p^l (Joule/Mbyte) and

the low idle energy consumption rate i^l (Joule/second). All disks have the same capacity c .

READ places popular files onto the hot disk zone and unpopular files onto the cold disk zone. The ratio between hot disk number and cold disk number in a disk array is decided by the load percentages of popular files and unpopular files in the whole file set. The load of a file f_i is defined as $h_i = \lambda_i \cdot sv_i$, where sv_i , λ_i are the file's service time and its access rate. Since we assume that each request sequentially scans a file from the beginning to the end, sv_i is proportional to s_i , the size of file f_i . Thus, the load of file f_i can also be expressed as $h_i = \lambda_i \cdot s_i$. Besides, we assume that the distribution of file access requests is a Zipf-like distribution with a skew parameter $\theta = \log \frac{A}{100} / \log \frac{B}{100}$, where A percent of all accesses are directed to B percent of file [20]. The number of popular files in F is defined as $|F_p| = (1-\theta) * m$, where m is the total number of files in F . Similarly, the number of unpopular files is $|F_u| = \theta * m$. Thus, the ratio between the number of popular files and the number of unpopular files in F is defined as δ

$$\delta = \frac{(1-\theta)}{\theta}. \quad (4)$$

The ratio between the number of hot disks and the number of cold disks is defined as γ , which is decided by the ratio between the total load of popular files and the total load of unpopular files:

$$\gamma = \frac{\sum_{i=1, f_i \in F_p}^{(1-\theta)*m} h_i}{\sum_{j=1, f_j \in F_u}^{\theta*m} h_j}. \quad (5)$$

Figure 6 depicts the READ algorithm. READ assigns sorted popular files in F_p onto the hot disk zone in a round-robin manner with the first file (supposed most popular one) onto the first disk, the second file onto the second disk, and so on. Similar file assignment strategy is applied for sorted unpopular file in F_u onto the code disk zone. After all files in F have been allocated, READ launches an Access Tracking Manager (ATM) process, which records each file's popularity in terms of number of accesses within one epoch in a table called File Popularity Table (FPT). The FPT table with the latest popularity information for each file will be used later by the File Redistribution Daemon (FRD). At the end of each epoch, FRD re-orders all files based on their access times recorded during the current epoch in the FPT table and then redefine popular file set F_p and unpopular file set F_u accordingly. A hot file will be migrated to the cold disk zone if its new position in the entire re-sorted file set is out of the newly defined hot file set range. It will stay in the hot zone, otherwise. Similarly, a previous cold file will be migrated to the hot disk zone if its new ranking is within the new hot

file set scope.

5. Performance evaluation

5.1. Experimental setup

We developed an execution-driven simulator that models an array of 2-speed disks. The same strategy used in [23] to derive corresponding low speed mode disk statistics from parameters of a conventional Cheetah disk is adopted in our study. The main characteristics of the 2-speed disk and major system parameters are similar as [23]. The number of disks in the simulated disk array varies from 6 to 16. The performance metrics by which we evaluate system performance include *mean response time* (average response time of all file access requests submitted to the simulated 2-speed parallel disk storage system), *energy consumption* (energy consumed by the disk systems during the process of serving the entire request set, and *AFR* (Annualized failure rate of a disk array). Each disk has an AFR calculated based on the PRESS model. The highest one is used to designate the AFR of the entire disk array.

We evaluate the three algorithms by running trace-driven simulations over a web I/O trace (WorldCup98-05-09 [2]), which has been widely used in the literature. Since the simulation times in our experiments are much shorter compared with the time span of the three traces, we only choose one day data from each trace. The WorldCup98 trace includes 4079 files with total 1,480,081 requests and the average request arrival interval is 58.4 ms.

5.2. Experimental results

We conduct our performance evaluation of the three energy-saving algorithms on a simulated platform of a disk array consisting of 6 to 16 disks. The READ algorithm consistently outperforms MAID and PDC algorithms in reliability by up to 39.7% and 57.5%, respectively. READ constrains each disk's number of speed transitions so that it cannot be larger than S , which is set to 40 in our study. READ accomplishes this by gradually enlarging the idleness threshold value. In our implementation, we simply double the idleness threshold value once READ finds that a disk's current number of speed transitions reaches half of S .

In terms of energy conservation, READ performs obviously better than the two baseline algorithms in a light workload condition (Figure 7a). On average READ results in 4.8% and 12.6% less energy consumption compared with MAID and PDC,

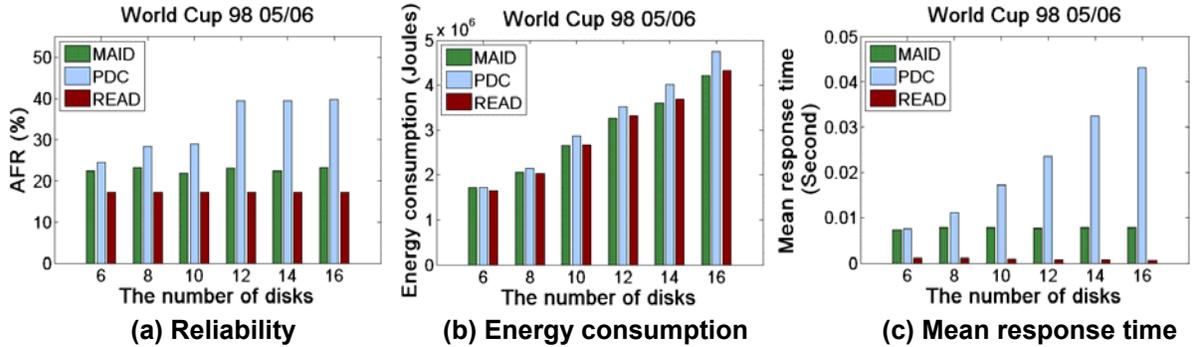


Figure 7. An overall comparison of the three algorithms

respectively. One important observation is that a large number of disk spin downs does not necessarily bring us more energy savings. On the contrary, a disk spin down can cause more energy consumption if the idle time is not longer enough to compensate the energy cost during disk spin down and spin up. This conjecture is demonstrated by the high energy consumption of MAID and PDC in Figure 7. We notice that our READ performs slightly worse than MAID in energy consumption in a heavy workload condition (Figure 7b) when the number of disks changes from 12 to 16. The reason is that MAID still have disk spin downs when the disk number increases and these disk spin down in deed bring energy conservation because in most cases the idle times are longer enough to compensate disk transition energy cost. On the other hand, our READ algorithm has no disk spin downs, and thus disks are always running at high speed. The READ algorithm delivers much shorter mean response times in all cases (Figure 7c) primarily due to its very few number of disk transitions.

6. Conclusions

In this paper, we establish an empirical reliability model PRESS, which can be utilized to estimate reliability impacts caused by the three ESRR factors. The PRESS model is built on a state-of-the-art work [22] and our own investigation on the relationship between disk speed transition frequency and reliability. In particular, our frequency-reliability function reveals that it is not a good idea to save energy if disk speed transition frequency is always higher than 65 times per day. Further, with the light shed by the PRESS model, we develop and evaluate a novel energy saving strategy with reliability awareness called READ (Reliability and Energy Aware Distribution). The READ strategy exploits popularity locality of I/O workload characteristics and an adaptive idleness threshold to limit each disk's speed transition times per day to provide a good reliability. Besides, it generates

a more even load distribution to further alleviate reliability side-effect. Our trace-driven experimental results show that when workload is not extremely heavy the READ strategy results in an average 24.9% and 50.8% reliability improvement compared with MAID and PDC, respectively. Meanwhile, in terms of energy consumption, READ in most cases still outperforms the two traditional approaches.

Future directions of this research can be performed in the following directions. First, we will extend our scheme to a fully dynamic environment, where file access patterns can dramatically change in a short period of time. As a result, a high file redistribution cost may arise as the number of file migrations increases substantially. One possible solution is to use file replication technique. Second, we intend to enable the READ scheme to cooperate with the RAID architecture, where files are usually striped across disks in order to further reduce the service time of a single request. For the web server environment, files are usually very small, and thus stripping is not crucial. However, for large files such as video clips, audio segments, and office documents, stripping is needed.

Acknowledgement

This work was supported by the US National Science Foundation under Grant No. CCF-0742187.

References

- [1] D. Anderson, J. Dykes, and E. Riedel, "More Than an Interface - SCSI vs. ATA," *Proc. 2nd USENIX Conf. File and Storage Technologies*, pp. 245-257, 2003.
- [2] M. Arlitt and T. Jin, *1998 World Cup Web Site Access Logs*, August 1998, Available at <http://www.acm.org/sigcomm/ITA/>.
- [3] E. V. Carrera, E. Pinheiro, and R. Bianchini, "Conserving Disk Energy in Network Servers," *Proc. Int'l Conf. Supercomputing*, pp. 86-97, 2003.

- [4] D. Colarelli and D. Grunwald, "Massive Arrays of Idle Disks for Storage Achieve," *Proc. of Supercomputing*, November 2002.
- [5] G. Cole, "Estimating Drive Reliability in Desktop Computers and Consumer Electronics Systems," *Technology Paper TP-338.1*, Seagate Technology, November 2000.
- [6] C. Cunha, A. Bestavros and M. Crovella, "Characteristics of WWW Client-based Traces," *Technical Report*, 1995-010, Boston University, 1995.
- [7] N. Durrant and R. Blish, "Semiconductor Device Reliability Failure Models," ismi.sematech.org/docubase/document/3955axfr.pdf, 2000.
- [8] J.G. Elerath, "Specifying Reliability in the Disk Drive Industry: No More MTBF's", *Proc. IEEE Annual Reliability and Maintainability Symp.*, pp.194-199, January 2000.
- [9] Engineering Statistics Handbook, the National Institute of Standards and Technology (NIST) website, <http://www.sbtionline.com/nist/index.html>, 2007.
- [10] G.R. Ganger, B.L. Worthington, R.Y. Hou and Y.N. Patt, "Disk arrays: high-performance, high-reliability storage subsystems," *IEEE Computer*, Vol. 27, No. 3, pp. 30-36, March 1994.
- [11] S. Glassman, "A caching relay for the Web," *Conf. World-Wide Web*, pp.165-173, 1994.
- [12] S. Gurumurthi, A. Sivasubramaniam, and V.K. Natarajan, "Disk Drive Roadmap from the Thermal Perspective: A Case for Dynamic Thermal Management," *Proc. 32 Int'l Symp. Computer Architecture*, pp.38-49, 2005.
- [13] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Fanke, "DRPM: Dynamic Speed Control for Power Management in Server Class Disks," *Proc. Int'l Symp. of Computer Architecture*, pp. 169-179, June 2003.
- [14] Hard Drive Info, GEN-X-PC company white paper, http://www.gen-x-pc.com/harddrive_info.htm.
- [15] G. Herbst, IBM's Drive Temperature Indicator Processor (Drive-TIP) Helps Ensure High Drive Reliability, *In IBM Whitepaper*, October 1997.
- [16] Hitachi Power & Acoustic Management—quietly cool, *White paper*, Hitachi Corp., 2004.
- [17] IDEMA Standards, "Specification of Hard Disk Drive Reliability", document number R2-98.
- [18] Y. Kim, S. Gurumurthi and A. Sivasubramaniam, "Understanding the Performance-Temperature Interactions in Disk I/O of Server Workloads," *Proc. 12th Int'l Symp. High-Performance Computer Architecture*, pp.176-186, 2006.
- [19] C.M. Kozierok, Hard Disk Spindle Motor, <http://www.pcguide.com/ref/hdd/op/spin.htm>.
- [20] L.W. Lee, P. Scheuermann, and R. Vingralek, "File assignment in parallel I/O systems with minimal variance of service time," *IEEE Transactions on Computers*, Vol. 49, Issue 2, pp. 127-140, 2000.
- [21] A.E. Papatthasiou and M.L. Scott, "Power-efficient Server-class Performance from Arrays of Laptop Disks," *WIP USENIX 2004 Annual Technical Conference*, Boston, 2004.
- [22] E. Pinheiro, W. Weber, and L. Barroso, "Failure Trends in a Large Disk Drive Population," *Proc. 5th USENIX conference on File and Storage Technologies*, San Jose, 2007.
- [23] E. Pinheiro and R. Bianchini, "Energy Conservation Techniques for Disk Array-Based Servers," *Proc. ACM Int'l Conf. on Supercomputing*, June 2004, St. Malo, France.
- [24] Reliability of Hard Disk Drives (HDD), technical paper, University of Maryland, 2003.
- [25] B. Schroeder and G.A. Gibson, "Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to You?" *USENIX FAST*, 2007.
- [26] T. Schwartz, M. Baker, S. Bassi, B. Baumgart, W. Flagg, C. Ingen, K. Joste, M. Manasse, and M. Shah, "Disk failure investigations at the internet archive," *Proc. 14th NASA Goddard, 23rd IEEE Conf. Mass Storage Systems and Technologies*, May 2006.
- [27] S. Shah and J.G. Elerath, "Reliability analysis of disk drive failure mechanisms," *Proc. IEEE Reliability and Maintainability Symp.*, pp. 226-231, 2005.
- [28] T. Xie and Y. Sun, "No More Energy-Performance Trade-Off: A New Data Placement Strategy for RAID-Structured Storage Systems," *The 14th Annual IEEE International Conference on High Performance Computing (HiPC 2007), Lecture Notes in Computer Science (LNCS 3834)*, pp.35-46, Goa, India, December 18-21, 2007.
- [29] Q. Xin, E.L. Miller, T. Schwarz, D.D.E. Long, S.A. Brandt, and W. Litwin, "Reliability mechanisms for very large storage systems," *Proc. 20th IEEE Conf. Mass Storage Systems and Technologies*, pp. 146-156, 2003.
- [30] Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton, and J. Wilkes, "Hibernator: helping disk arrays sleep through the winter," *Proc. 12th ACM Symposium on Operating Systems Principles*, pp. 177-190. 2005.
- [31] Q. Zhu and Y. Zhou, "Power-aware storage cache management," *IEEE Transactions on Computers*, Vol. 54, Issue 5, pp. 587-602, 2005.