2009 Sixth International Conference on Information Technology: New Generations

# An Energy-Efficient Reliability Model for Parallel Disk Systems

Fangyang Shen†, Xiao Qin, Andres Salazar†, Adam Manzanares, Kiranmai Bellam

†Department of CET
Northern New Mexico College
{fshen,asalazar}@nnmc.edu

Department of CSSE
Auburn University
{xqin, kzb000,acm0008}@auburn.edu

## Abstract

*In the last decade, parallel disk systems have increasingly become popular for data-intensive applications running on high-performance computing platforms. Conservation of energy in parallel disk systems has a strong impact on the cost of cooling equipment and backup power-generation. This is because a significant amount of energy is consumed by parallel disks in high-performance computing centers. Although a wide range of energy conservation techniques have been developed for disk systems, most energy saving schemes have adverse impacts on the reliability of parallel disk systems. To address this deficiency, we must focus on reliability analysis for energy-efficient parallel disk systems. In this paper, we make use of a Markov process to develop a quantitative reliability model for energy-efficient parallel disk systems using data mirroring. With the new model in place, a reliability analysis tool is developed to efficiently evaluate reliability of fault-tolerant parallel disk systems with two power modes.*

Key words: Parallel Disk Systems; Disk Reliability, Mean time to fail; Markov Chains; Mean time to recover

## 1. Introduction

In the last decade, parallel disk systems have been widely used to support a variety of data-intensive applications running on high-performance computing platforms. The reason behind the popularity of parallel disk systems is that parallel I/O is a promising avenue to bridge the performance gap between processors and I/O systems [30]. For example, redundant arrays of inexpensive disks or RAID can offer high I/O performance with large capacities [30]. The RAID systems improve reliability because they prevent data loss using disk redundancy.

Recent studies show that a significant amount of energy is consumed by parallel disks in high-performance computing centers [32, 33]. Energy conservation in parallel disk systems not only lowers electricity bills, but also leads to reductions of emissions of air pollutants from power generators. Moreover, the storage requirements of modern data-intensive applications and the emergence of disk systems with higher power needs make it desirable to design energy-efficient parallel disk systems. Energy conservation techniques developed for parallel disk systems include dynamic power management schemes [10, 11], power-aware cache management strategies [12], power-aware prefetching schemes [13], and multi-speed settings [16-18]. Many energy saving techniques significantly conserve energy in parallel disk systems at the apparent cost of reliability. Even worse, a parallel disk system consisting of multiple disks has a higher failure rate compared with a larger single disk system. This is due to the fact that there are more components in the parallel disk system. Due to hardware and software defects, failed components in a parallel disk system can ultimately cause failures in the system at run-time [29]. In addition to the energy efficiency issue, fault tolerance and reliability are major concerns in the design of modern parallel disk systems. These disk systems are expected to be the most stable part of high-performance computing systems. Therefore, it is not surprising that academic institutes, industry, and government agencies consider the reliability and energy-efficiency of parallel disks in their computing systems essential and critical to operations.

The long-term goal of this research is to develop novel energy conservation schemes with marginal adverse impacts on the reliability of parallel disk systems. To achieve this long-term objective, in this study we investigate fundamental theories to model the reliability of energy-efficient parallel disk systems using data mirroring.

In this paper we present a quantitative approach to addressing the issue of conserving energy without noticeably degrading reliability of parallel disk systems. We focus on parallel disk systems with data mirroring, which is a technique for maintaining two or more identical disk images on multiple disk devices (see, for example, [29]). The data mirroring technique leverages redundant data to significantly enhance reliability. Among the various energy conservation techniques, we focus on a widely adopted technique – dynamic power management, which dynamically changes the power state of disks. The dynamic power management technique turns disks into low-power states (e.g., sleep state) when the disk I/O workload is fairly light. In this study, we use a Markov process to develop a quantitative reliability model for energy-efficient parallel disk systems with data mirroring. Next, we evaluate the reliability and energy efficiency of parallel disk systems using the new reliability model.

With our novel reliability model in place, it becomes feasible to devise new approaches to seamlessly integrate energy efficiency and fault tolerance. The goal of these approaches is to provide significant energy savings, while ensuring high reliability for parallel disk systems without sacrificing performance. Our approach has several advantages over traditional qualitative approaches. This is because our model not only can quantify the reliability of energy-efficient parallel disks, but also can be used to balance energy efficiency and reliability.

## 2. Related Work

High reliability is one of the key design goals of modern parallel disk systems. In the past decade, a variety of practical and fundamental reliability models have been constructed for storage systems. For example, conventional reliability models include Markov Chain models [2-5], analytic queuing models [6, 7], queuing networks [8], and iterative models [9]. Although the existing models can accurately measure the reliability of disk systems, all of the models are inadequate to quantify the

978-0-7695-3596-8/09 $25.00 © 2009 IEEE
DOI 10.1109/ITNG.2009.295



IEEE computer society

Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 22:54:40 UTC from IEEE Xplore. Restrictions apply.

reliability of parallel disk systems, in which energy saving techniques are implemented.

The issue of energy efficiency in storage systems has received increasing attention. Energy conservation techniques proposed in previous studies include dynamic power management schemes[10,11], power-aware cache management strategies[12], power-aware prefetching schemes[13], software-directed power management techniques[14], redundancy techniques [15], multi-speed settings[16-18], and dynamic voltage scaling[19]. However, the previous research did not investigate the impact of energy saving techniques on the reliability of storage systems. Our work represents the first attempt to build reliability models for the design process of fault-tolerant and energy-efficient parallel disk systems.

A number of studies on energy conservation in parallel disk systems showed that energy-saving techniques should not sacrifice system reliability [20-24]. However, little attention has been paid to the integration of reliability into energy-conservation techniques for disk arrays in general, and for mobile disk arrays in particular. It is imperative to build an energy-dependent reliability model to reveal relationships between energy-saving techniques and the reliability of parallel disk systems. Therefore, in this study we develop a new energy-aware reliability model. This model is accompanied by a reliability analysis tool able to systematically quantify the reliability of a particular type of energy-efficient and fault-tolerant parallel disk system.

Geist and Trivedi proposed an analytic model to measure the performance and reliability of mirrored disk systems [5]. Their model shows that the disk mirroring technique can substantially improve both reliability and performance of disk systems. Our model is radically different from theirs in the fact that mirrored disk systems considered in our model have high energy-efficiency by the virtue of dynamic power management.

Unsal et al. investigated the relationship between fault tolerance techniques and energy consumption in real-time systems. This was achieved by establishing the energy efficiency of application level fault tolerance over other software-based fault tolerance methods [25]. Although this study integrated fault tolerance and energy-saving techniques in real-time systems, the above techniques are inadequate to address the reliability issues in energy-efficient parallel disk systems.

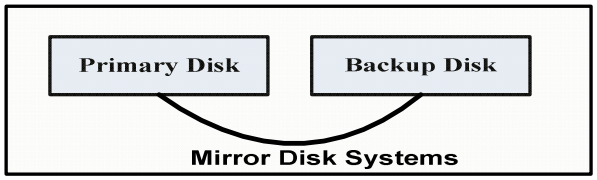



**Fig. 1. A parallel disk systems with data mirroring.**

## 3. Modeling Reliability of Mirroring Disks

### 3.1. Reliability models

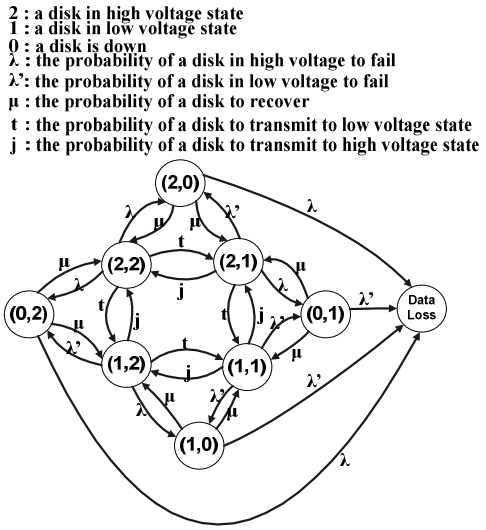



**Fig. 2 Mirroring disks with two supply voltage levels**

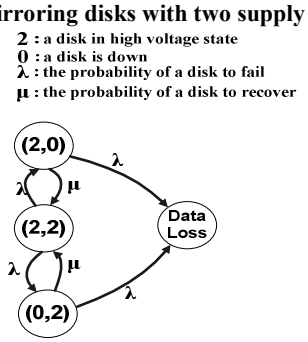



**Fig. 3 Mirroring disks with one voltage level**

In this study, we focus on mirroring disk systems (see Fig. 1). Let us suppose disk failure and repair rates are modeled using the exponential distribution, we build a reliability model based on the Markov process to estimate the reliability and energy savings in parallel disk systems with data mirroring. We aim at utilizing the reliability model to quantify the mean-time-to-data-loss or MTTDL of a fault-tolerant parallel disk system with the data mirroring technique. Without loss of generality, we consider disks with two power states, i.e., a high-voltage state and a low voltage state. In addition, we focus on disk systems with one primary disk and one backup disk. The Markov reliability models of mirroring disks are diagramed in Fig. 2 and 3. Note that Fig. 2 plots the model for mirroring disks with two supply voltage levels; Fig. 3 outlines the model for mirroring disks with a single voltage level.

**Table 1. Parameters used to characterize reliability of mirroring disks with two voltage levels**

| |
|---|
| Sp = the state of a primary disk |
| Sp or Sb = 2: disks are in the high voltage state |
| Sp or Sb = 1: disks are in the low voltage state |
| λ = failure rate of a disk in the high voltage state |
| t = probability of transition from a high voltage state to a low voltage state |
| Sb = the state of a backup disk |
| Sp or Sb = 0: disks failed |



Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 22:54:40 UTC from IEEE Xplore. Restrictions apply.

| |
|---|
| μ = recovery rate |
| λ' = failure rate of a disk in the low voltage state |
| j = probability of transition from a low voltage state to a high voltage state |

The reliability of mirroring disks with two voltage states can be quantitatively characterized by parameters summarized in Table 1. The parameters (λ, μ, t, j) capture the reliability of mirroring disks with two supply voltage levels. Let (Sp, Sb) denote the system state, where Sp, Sb are the states of primary and backup disks, respectively. Si = 2 (i= b and p) signifies that the disk is at the high voltage level; Si = 1 (i= b and p) indicates that the disk is at the low voltage level; and Si = 0 (i= b and p) means that the disk is in the failure state. For example, an initial state (2, 2) indicates that both primary and backup disks are at the high voltage supply level. The state (0, 0) means that the primary and backup disks have unrecoverable failures, (i.e., data loss). Once the system state transits to one of the four states ((0, 2), (2, 0), (1, 0), (0, 1)), further failure can cause the system state to transition to the data loss state.

We model the failure rate of disks as a function of voltage. Thus, the failure rate function is:

$$\lambda(v) = v \cdot 10^{\frac{d(1-v)}{1-v_{low}}}, \quad (1)$$

where $v$ is the normalized supply voltage, $v_{low}$ is the normalized minimum voltage, and $d$ is a constant. A larger $d$ indicates that the failure rate is more sensitive to the discrepancy between the high and low voltages. In our experiments, $d$ is set to 0, 2, 4, 6, respectively. A similar failure rate model for processors was introduced by Zhu et al. [26].

### 3.2. Reliability estimation

The mean-time-to-data-loss (MTTDL) can be calculated using the fundamental matrix M defined as $M = [m_{ij}] = [I - Q]^{-1}$, where Q is the truncated matrix given below [18], and matrix I is the identity matrix corresponding to the dimension of matrix Q, N is the number of total disks.

$$Q = \begin{bmatrix} 1 - N\lambda & N\lambda \\ \mu & 1 - (N-1)\lambda - \mu \end{bmatrix} \quad (2)$$

The truncated stochastic transitional probability matrix for mirroring disks with two voltages is given as follows:

$$Q = \begin{pmatrix} 1-2t-2\lambda & t & t & 0 & \lambda & \lambda & 0 & 0 \\ j & 1-t-j-\lambda-\lambda' & 0 & t & 0 & \lambda' & \lambda & 0 \\ j & 0 & 1-t-j-\lambda-\lambda' & t & \lambda' & 0 & 0 & \lambda \\ 0 & j & j & 1-2j-2\lambda' & 0 & 0 & \lambda' & \lambda' \\ \mu & 0 & \mu & 0 & 1-\lambda-2\mu-t & 0 & t & 0 \\ \mu & \mu & 0 & 0 & 0 & 1-\lambda-2\mu-t & 0 & t \\ 0 & \mu & 0 & \mu & j & 0 & 1-\lambda'-2\mu-j & 0 \\ 0 & 0 & \mu & \mu & 0 & j & 0 & 1-\lambda'-2\mu-j \end{pmatrix} \quad (3)$$

Similarly, we define the transitional probability matrix for mirroring disks with a single voltage as

$$Q_2 = \begin{pmatrix} 1-\mu-\lambda & \mu & 0 \\ \lambda & 1-2\lambda & \lambda \\ 0 & \mu & 1-\mu-\lambda \end{pmatrix} \quad (4)$$

We assume that the mirroring disk system starts in state (2, 2); thus, MTTDL can be expressed as $MTTDL = \sum_{j=1}^{d+1} m_{1j}$.

### 3.3. Energy consumption model

The following parameters are introduced to model energy dissipation rates in mirroring disks.

$T_{total}$ : the total operation time of a disk system.

$E_{total}$ : the total energy consumption (measured in Joule) in a total time of $T_{total}$ .

$P_{ij}$ : the probability that a disk system is in state (i, j). Thus, we have $\sum_{j=0}^{2}\sum_{i=0}^{2} P_{ij} = 1$ .

$V_L$ : the low voltage of a disk; $V_H$ : the high voltage of a disk .

The energy consumption of a mirroring disk system with two voltage levels can be written as:

$$E_{total} = 3600 \cdot \left( \left( \sum_{i=0}^{2} P_{i1} + \sum_{j=0}^{2} P_{1j} \right) \cdot V_L + \left( \left( \sum_{i=0}^{2} P_{i2} + \sum_{j=0}^{2} P_{2j} \right) \cdot V_H \right) \right) \cdot T_{total} \quad (5)$$

The energy consumption of a mirroring disk system with a single high voltage within 24 hours is

$$E_{total} = 3600 \cdot (P_{22} + P_{02} + P_{20}) \cdot V_H \cdot T_{total} . \quad (6)$$

**Table 2. Parameters of energy-efficient mirroring disks. MTTR: Mean-time-to-recover**

| Parameters | Default Values | Values I | Values II | Values III | Values IV |
|---|---|---|---|---|---|
| **MTTF(year)** | 5—30 | 5,10,15,30 | 5,10,15,30 | 5--30 | 30 |
| **MTTR(day)** | 1,3,6,7 | 1-7 | 1,3,7 | 1 | 1 |
| **T** | 0.1,0.5,0.9 | 0.1,0.5,0.9 | 0.1,0.5,0.9 | 0.1,0.5,0.9 | variable |
| **D** | 2 | 2 | 2 | 2 | 2 |
| **High Voltage(W)** | 13.5 | 13.5 | 13.5 | 13.5 | 13.5 |
| **Low Voltage(W)** | 10.2 | 10.2 | 10.2 | 10.2 | 10.2 |
| **Time(Hour)** | 24 | 24 | 24 | 24 | 24 |
| **Fmin** | 0.46 | 0.46 | 0.46 | 0.46 | 0.46 |

## 4. Performance Evaluation

To determine the strength of our reliability model, we make use of the model to measure the reliability and energy dissipation of two mirroring disk systems. The first mirroring disk system is energy efficient by the virtue of two supply voltages; the second mirroring disk system is a traditional system without dynamic power management. Table 2 summarizes the key configuration parameters of a mirroring disk system. The parameters are chosen to resemble real-world disks like the Seagate ST-34501W.

### 4.1 Impact of disk failure rate

The relationship between disk failure rate λ, and MTTDL is shown in Fig. 4. Similarly the relationship between λ and energy consumption is shown in Fig. 5. The results plotted in Figs. 4 and 5 are obtained by using default values taken from Table 2. It is evident from Fig. 4 that MTTDL significantly increases with the increase of λ. This is because MTTDL of a system is directly related to MTTF. We observe from Fig. 5 that



Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 22:54:40 UTC from IEEE Xplore. Restrictions apply.

the energy consumption of a disk system increases gradually with the increase of λ. The reason is that more energy is consumed by the disk system when each disk is active for a longer period of time. Figs. 4 and 5 suggest that gradually increasing λ allows us to find a compromise between MTTDL and energy consumption.

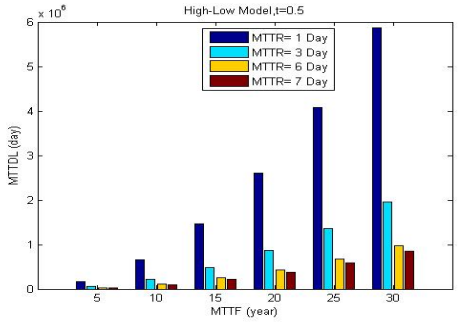



**Fig. 4. Impact of disk failure rate λ on MTTDL.**

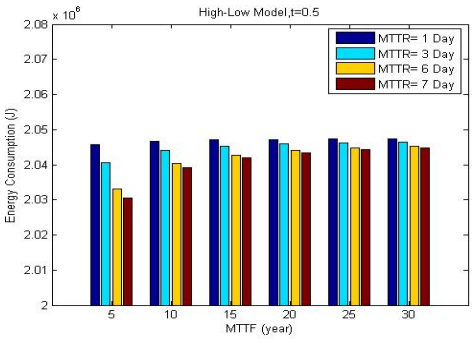



**Fig. 5. Impact of disk failure rate λ on energy.**

## 4.2 Impact of mean-time-to-recover

Fig. 6 reveals the correlation between mean-time-to-recover (MTTR) and MTTDL, whereas Fig. 7 shows the relationship between $\mu$ and energy consumption. The simulation settings for these figures are using values I in Table 2. We observe that MTTDL and MTTR are inversely related. This result is expected because a large MTTR leads to a low recover rate $\mu$, which is also inversely proportional to MTTR. As the recovery rate is lowered, the reliability of mirrored disks also decreases. Fig. 7 clearly illustrates that the energy consumption for a disk system decreases slowly with the increase of MTTR. This is mainly because a large value of MTTR implies a smaller likelihood of having a failed disk repaired.

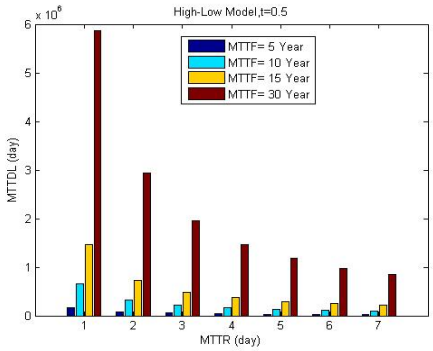



**Fig. 6. Impact of recovery rate on MTTDL.**

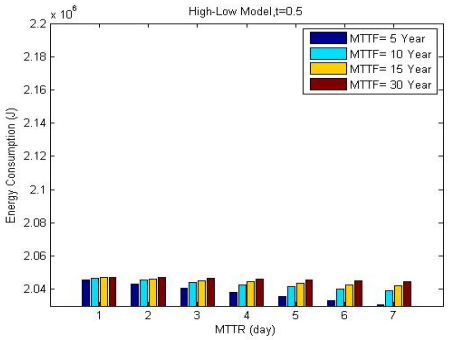



**Fig. 7. Impact of recovery rate on energy consumption.**

## 4.3 Probability of power transitions

Recall that parameter *t* represent the probability of a transition between the high voltage and low voltage states. Fig. 8 shows the relationship between *t* and MTTDL; Fig. 9 plots the correlation between t and energy consumption. The simulation settings can be found in the third column (values II) in Table 2. Fig. 8 reveals that MTTDL decreases rapidly with the increase of t. This is because the MTTDL of a system is related to the amount of time each disk stays in the high voltage state. A smaller t means a disk stays longer periods of time in the high voltage state, thereby enhancing the reliability of the disk system. This implies that we can increase the reliability of a disk system by increasing the supply voltage level of each disk in the system.

Fig. 9 shows that the energy consumption for a disk system gradually decreases with the increase of t. This is because less energy is consumed by the disk system when each disk stays in the low voltage state for a longer period of time.

## 4.4 Impact of parameter *d*

Recall that parameter *d* in Eq. (1) represents the sensitivity of the failure rate on disk voltage supplies. Figs. 10 and 11 plot the impacts of the parameter *d* on MTTDL and energy consumption respectively. The settings of this experiment are



Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 22:54:40 UTC from IEEE Xplore. Restrictions apply.

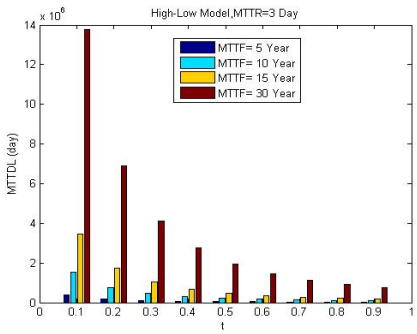



**Fig. 8. Impact of t on MTTDL. *t*: Probability of transition from high voltage to low voltage**

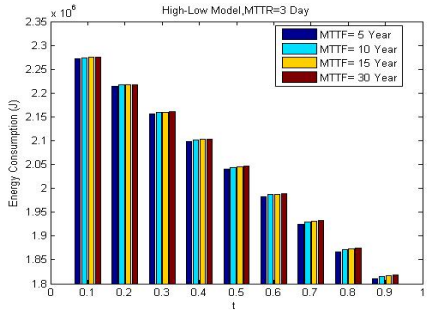



**Fig. 9. Impact of t on energy. *t*: Probability of transition from high voltage to low voltage**

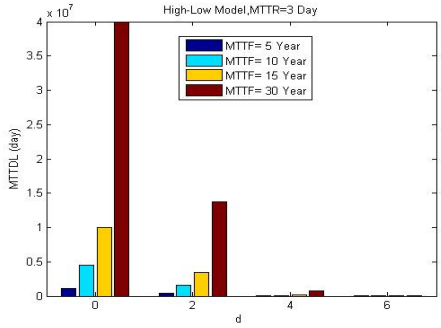



**Fig. 10. Impact of parameter d on MTTDL.**

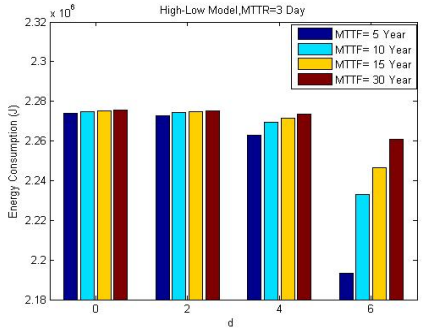



**Fig. 11. Impact of parameter d on energy consumption.**

identical to the previous experiments except that we now vary d from 0 to 6 by 2. Fig. 10 shows that increasing the value of d results in a rapid decrease in MTTDL. The rationale behind this result is that MTTDL relies on failure rate, which is directly related to d. We can increase the reliability of a disk system by setting a small value of d. In contrast, when d is smaller than 6, energy consumption is less sensitive to d. The energy consumption of the disk system will now increase slowly with the increase of d (see Fig. 11).

### 4.5 Impact of high voltage level

Fig. 12 shows the relationship between high voltage and MTTDL, whereas Fig. 13 illustrates the effect high voltage has on energy consumption. Fig. 12 shows that MTTDL decreases with an increasing value of the high voltage. This trend can be explained by the fact that MTTDL is likely to decrease if there are a large number of transitions between the high and low voltages. As the gap increases between the high and low voltages, the disk system reliability decreases. We observe from Fig. 13 that the energy consumption of the disk system increases with increasing values of high voltage. This is because more energy dissipation is caused by disks when they stay in the high voltage state for longer periods of time.

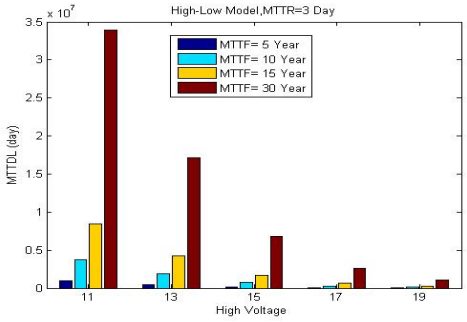



**Fig. 12. Impact of high voltage level on MTTDL.**

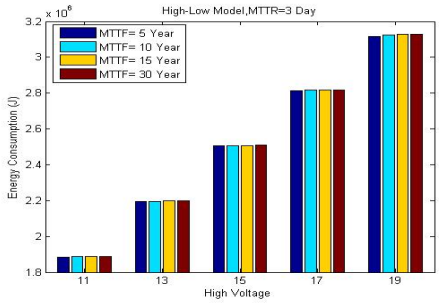



**Fig. 13. Impact of high voltage level on energy consumption**

## 5. Conclusion

In this paper we focused on energy-efficient parallel disk systems with data mirroring. We developed a reliability model for mirroring disk systems where the dynamic power management technique is employed to significantly reduce energy dissipation in disks. Our reliability model, which is based on a Markov process, makes it possible to estimate reliability of mirroring disks with two supply voltage levels. With the reliability model in place, we developed a tool to evaluate the reliability and energy efficiency of parallel disk systems with data mirroring. In future studies, we first will extend our model by incorporating the effects and cost of state transitions between high voltage and low voltage. Then, we will develop a reliability model for energy efficient RAID systems, which are widely used in data-intensive computing systems.



Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 22:54:40 UTC from IEEE Xplore. Restrictions apply.

## Acknowlegements

Thanks for David F. Trujillo for the support on this work. This research was supported in part by a grant from the U.S. Department of Education.

In addition, this work was supported in part by the US National Science Foundation under Grants No. CCF-0742187, No. CNS-0757778, No. CNS-0831502, No. OCI-0753305, No. DUE-0621307, and No. DUE-0830831, and Auburn University under a startup grant.

## References

[1] A.B. Smith, C.D. Jones, and E.F. Roberts, "Article Title", Journal, Publisher, Location, Date, pp. 1-10.

[2] S.H. Baek, B.W. Kim, E.J. Joung, and C.W. Park, "Reliability and Performance of Hierarchical RAID with Multiple Controllers," Proc. ACM Int'l Symp. Principles of Distri. Comp., pp. 246–254, 2001.

[3] G.A. Gibson and D.A. Patterson, "Designing Disk Arrays for High Reliability," J. Parallel and Distri. Computing, vol. 17, no. 1-2, pp. 4–27, 1993.

[4] R.Y. Hou and Y.N. Patt, "Using Non-volatile Storage to Improve the Reliability of RAID5 Disk Arrays," Proc. Int'l Symp. Fault-Tolerant Comp., pp. 206–215, 1997.

[5] R. Geist and K. Trivedi, "An Analytic Treatment of the Reliability and Performance of Mirrored Disk Subsystems," Proc. Int'l Symp. Fault-Tolerant Comp., pp. 442-450, 1993.

[6] S. Chen and D. Towsley, "The Design and Evaluation of RAID 5 and Parity Striping Disk Array Architectures," J. Parallel and Distributed Comp., vol. 17, pp. 58-74,1993.

[7] S. Chen and D. Towsley, "A Performance Evaluation of RAID Architectures," IEEE Trans. Computers, vol. 45, no. 10, pp. 1116-1130, 1996.

[8] E. Drakopoulos and M.J. Merges, "Performance Analysis of Client-Server Storage Systems," IEEE Trans. Computers, vol. 41, no. 11, pp. 1442-1452, 1992.

[9] E. Geist, D. Finkel, and S.K. Tripathi, "Availability of A Distributed Computer Systems with Failures," Acta Informatica. vol. 23, no. 6, pp. 643-655,1986.

[10] F. Douglis, P. Krishnan, and B. Marsh, "Thwarting the Power-Hunger Disk," Proc. USENIX Conf., 1994.

[11] K. Li, R. Kumpf, P. Horton, and T.E. Anderson, "A Quantitative Analysis of Disk Drive Power Management in Portable Computers," Proc. Winter USENIX Conf., pp. 279-292, 1994.

[12] Q. Zhu, F.M. David, C.F. Devaaraj, Z. Li, Y. Zhou, and P. Cao, "Reducing Energy Consumption of Disk Storage Using Power-Aware Cache Management," Proc. High-Performance Computer Architecture, pp. 118-129, 2004.

[13] S. W. Son and M. Kandemir, "Energy-aware Data Prefetching for Multi-speed Disks," Proc. ACM Int'l Conf. Comp. Frontiers, pp. 105-114, 2006.

[14] S.W. Son, M. Kandemir, and A. Choudhary, "Software-directed disk Power Management for Scientific Applications," Proc. Int'l Symp. Parallel and Dist. Processing, pp. 4-14, 2005.

[15] E. Pinheiro, R. Bianchini, and C. Dubnicki, "Exploiting Redundancy to Conserve Energy in Storage Systems," Proc. Sigmetrics and Performance, pp. 15-26, 2006.

[16] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Fanke, "DRPM: Dynamic Speed Control for Power Management in Server Class Disks," Proc. Int'l Symp. Computer Arch., pp. 169-179, 2003.

[17] D.P. Helmbold, D.E. Long, T.L. Sconyers, and B. Sherrod, "Adaptive Disk Spin-Down for Mobile Computers," Mobile Networks and Applications, vol. 5, no. 4, pp. 285-297, 2000.

[18] P. Krishnan, P. Long, and J. Vitter, "Adaptive Disk Spindown via Optimal Rent-to-buy in Probabilistic Environments," Proc. Int'l Conf. Machine Learning, pp. 322-330, 1995.

[19] C.M. Krishna and Y.H. Lee, "Voltage-clock-scaling Techniques for Low Power in Hard Real Time Systems," IEEE Real-Time Technology and Appl. Symp., pp. 156–165, 2000.

[20] P. Pillai and K.G. Shin, "Real-time Dynamic Voltage Scaling for Low-Power Embedded Operating Systems," Proc. ACM Symp. Operating Sys. Principles, 2001.

[21] S. Gurumurthi, J. Zhang, A. Sivasubramaniam, M. Kandemir, H. Fanke, N. Vijaykrishnan, and M. Irwin, "Interplay of Energy and Performance for Disk Arrays Running Transaction Processing Workloads," IEEE Int'l Symp. Perf. Analy. Sys. and Software, pp. 123-132, 2003.

[22] E. Pinheiro, R. Bianchini, and C. Dubnicki, "Exploiting Redundancy to Conserve Energy in Storage Systems," Proc. Sigmetrics and Performance, pp. 15-26, 2006.

[23] C. Weddle, M. Oldham, J. Qian, and A. Wang, "PARAID: The Gear-Shifting Power-Aware RAID," Technical Report 060323, Florida State University, 2006.

[24] Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton, and J. Wilkes, "Hibernator: Helping Disk Arrays Sleep Through the Winter," Proc. 12th ACM Symp. Operating Sys. Principles, pp. 177-190, 2005.

[25] S. Unsal, I. Koren, and C.M. Krishna, "Towards Energy-Aware Software-Based Fault Tolerance in Real Time Systems," Proc. Intl' Symp. Low Power Electronics and Design, pp. 124-129, 2002.

[26] D.-K. Zhu, R. Melhem and D. Mosse, "The Effects of Energy Management on Reliability in Real-Time Embedded Systems," Proc. IEEE/ACM Int'l Conf. Computer-aided Design, pp. 528-534, 2004.

[27] B. Schroeder and G.A. Gibson, "Disk Failures in The Real World: What does an MTTF of 1,000,000 hours mean to you?," 5th USENIX Conf. File and Storage Tech., pp. 1-16, 2007.

[28] M. Xie, Z.L. Yang, and O. Gaudong, "More on The Mis-Specification of the Shape Parameter with Beibull-to-Exponential Transformation," Proc. Int'l Quality and Reliability Engineering, vol.16, no. 4, pp. 281-290, 2001.

[29] D. Bitton, "Disk Shadowing," Proc. Int'l Conf. Very Large Data Bases, pp. 331-338, 1998.

[30] Y.B. Lee and C.S. Lui, "Automatic Recovery from Disk Failure in Continuous-Media Servers," IEEE Trans. Parallel and Distr. Sys., vol.13, no. 5, pp. 499-515, 2002.

[31] D.A. Patterson, G.A. Gibson, and R.H. Katz, "A Case for Redundant Arrays of Inexpensive Disk," Int'l Conf. Management of Data, Proc. the ACM SIGMOD Int'l Conf. Management of Data, pp. 109-116, 1988.

[32] X. Qin, "Design and Analysis of a Load Balancing Strategy in Data Grids," Future Generation Comp. Sys.: The Int'l J. Grid Comp., vol. 23, no.1, pp. 132-137, 2007.

[33] Z.L. Zong, M.E. Briggs, N.W. O'Connor, and X. Qin, "An Energy-Efficient Framework for Large-Scale Parallel Storage Systems," Proc. Int'l Parallel and Distributed Processing Symp., CA, 2007.



Authorized licensed use limited to: Auburn University. Downloaded on November 17,2026 at 22:54:40 UTC from IEEE Xplore. Restrictions apply.