

### Keywords

Wireless Sensor Network,  
MAC Protocol,  
Handover Management,  
Clustering,  
Energy Efficient

Received: March 5, 2016

Accepted: March 21, 2016

Published: May 6, 2016

# Improving and Minimizing Power Consumption

**Mohammed I. Alghamdi**

College of Computer Science and Information Technology, Computer Science department, Albaha University, Albaha, Kingdom of Saudi Arabia

### Email address

mialmushilah@bu.edu.sa

### Citation

Mohammed I. Alghamdi. Improving and Minimizing Power Consumption. *American Journal of Computation, Communication and Control*. Vol. 3, No. 1, 2016, pp. 1-9.

### Abstract

Energy conservation techniques play a vital role in mobile devices and sensors because their operational capability is limited by battery power. Over the past few years, energy and security have become very important in data-intensive applications running on high-performance computers. Web applications, file and media servers, transaction processing systems and search engines are a few examples of data-intensive applications. These data-intensive applications and services consume a huge amount of power for storing and retrieving data, since the applications and services are constantly servicing, processing, and supporting transactions throughout a day. Data intensive applications can result in high energy cost, which in turn is likely to become an increasing burden for companies maintaining the data-intensive computing servers.

## 1. Introduction

Previous studies have shown that data-intensive applications in data centers can consume several mega-watts of power [12] [13] [20]. A considerable amount of energy is expected to be used by large-scale storage systems, where a large number of disks are employed to handle numerous submitted disk requests [14] [15] [18] [19]. Traditional approaches to reducing energy dissipation in disk systems offer aggressively shutting down disks or placing disks into the standby mode when I/O loads are low. It is known that disk spindle motors exhibit noticeable energy overhead within hard disks. To address this issue, I apply in this study the dynamic speed control technique for power management in hard disks. The dynamic speed control technique performs very well under light I/O workload in disks with substantial idle time periods among disk requests. This approach relies on the new evolving technique - Dynamic Rotations per Minute, where one can dynamically choose the speed of a disk based on I/O load information. The dynamic speed control scheme is an efficient approach to conserving energy in disks, because slowing down disks can reduce energy consumption quadratically while degrading I/O response times linearly.

The increased number of devices that store, access, manipulate, or communicate security sensitive data has made security an important issue in the realm of storage systems. Security concerns range from user identification to secure information storage. Confidentiality services – one of the major security services provided in storage systems - are implemented using cryptographic technology, which are in most cases selected in accordance to security objectives and requirements of targeting systems.

I integrate the dynamic speed control technique with a wide variety of confidentiality security services to improve both energy efficiency and security of large-scale storage systems. I propose two approaches to integrating confidentiality services with the dynamic disk speed control technique. The first strategy, which is security aggressive in

nature, focuses on the improvement of storage system security followed by energy conservation. The second strategy gives higher priority to energy conservation as opposed to security optimization. Our approach is general in the sense that it can be readily applied to integrate a wide range of security services with various energy conservation techniques.

The rest of the paper is organized as follows. Section 2 gives a summary of the related work. Section 3 introduces the dynamic speed control technique for hard disks. A energy consumption model is also presented in Section 3. Section 4 describes the implementation issues of various security services in storage systems. Section 5 presents an approach to improve both energy efficiency and security in parallel disk systems. The experimental results and evaluations are presented in Section 6. Finally, Section 7 concludes the paper with discussions and future work directions.

## 2. Related Work

Abundant research has been done to improve energy efficiency of wireless devices like PDAs, cell phones, and sensors, thereby increasing battery life of the devices. Unlike mobile devices where battery life is critical, storage systems have to be energy efficient in order to lower electricity bills in data centres. In the past five years, much attention has been paid to energy conservation techniques for storage systems. For example, the conventional wisdom of saving energy in storage systems is to place idle disks into the standby mode or to shut down idle disks. Such dynamic power management schemes are effective, since significant energy in disks is consumed by spindle motors [2]. The dynamic power management techniques have been widely applied to reduce energy dissipation in disk drives in laptops as well as in high-performance computers [5] [6].

Energy dissipation in hard drives can be efficiently reduced by applying multi-speed disks because of very low power-state transition penalties. Song and Kinder developed novel energy-aware compilers for multi-speed disks [16]. The dynamic speed control scheme for power management in disks can conserve energy by adaptively choosing disk speed [2]. The speed of a disk is chosen to save energy without compromising I/O performance or disk response time. Our proposed approach relies on multi-speed disks, in which a dynamic speed control mechanism can judiciously reduce disk speed under light workload conditions. It is noteworthy that our approach is different from either multi-speed disks or the dynamic speed control technique, because ours seamlessly integrates energy-efficient disks with security services.

A wide range of security protocols and cryptographic algorithms have been thoroughly studied. Their vulnerabilities to different attacks and time to crack encrypted data have been revealed [7]. Comprehensive energy analysis of security protocols and cryptographic algorithms has been discussed by Potlapally *et al.* [4]. Their

results not only help in identifying energy bottlenecks, but also provide insight into the development energy-efficient security mechanisms and services. The energy consumption model used in this study is motivated by Potlapally *et al.*'s [4] energy analysis of cryptographic algorithms. In this research, I aim at developing a mechanism to assist improving both security and energy efficiency in large-scale storage systems. Our way of achieving high energy efficiency and security is orthogonal to energy-efficient security services in the sense that further energy savings can be achieved by integrating our approach with the existing security services with high energy efficiency.

## 3. Energy-Efficient Storage Systems

The existing energy conservation techniques, like the dynamic disk speed control technique, can be used in conjunction with other techniques that reducing disk accesses or data placement strategies to reduce head movements for power savings [8]. Spindle motors, which spin the disk platters, are the major power consumer even when disks are not serving a request as the disks are still spinning [9]. To minimize energy overhead incurred by spindle motors, I apply the dynamic disk speed control technique to conserve energy in secure storage systems.

*Table 1. Mechanical characteristic of motor:*

Parameter	Value	Units
Max. Permissible Speed	10,000	rpm
Rotor Inertia (J <sub>0</sub> )	3.84	gcm <sup>2</sup>
Torque Constant (K <sub>T</sub> )	9.1	mNm/A
Max. Continuous Current	0.708	A

### 3.1. Calculating RPM Transition Times

Since the dynamic disk speed control technique or DRPM dynamically changes the Rotations per Minute (RPM), there is time overhead that is required to change disk speed or RPM. In order to calculate the time required for a speed change, I will rely on physical data of spindle motors. I have obtained the necessary information from the datasheet of a Maxon EC-20, which is a 20 mm flat brushless permanent magnet DC motor [10], whose physical characteristics closely match those of a hard disk spindle motor. Table 1. summarizes the basic mechanical characteristics of this motor.

The motor specification is given by equation (1) for calculating the time  $\Delta t$  (in ms) required for a speed change of  $\Delta n$  RPM with a load inertia  $J_L$  [2] as

$$\Delta t = \frac{\pi}{300} \Delta n \frac{J_0 + J_L}{K_T I} \quad (1)$$

The mass of an individual platter 'm' was found to be 14.65 gm and the radius 'r' was 4.7498 cm. Using these values and the platter count as two, I calculate the moment of inertia of the load as

$$J_L = \frac{1}{2} n_p m r^2 = \frac{1}{2} \times 2 \times 14.65 \times (4.7498)^2 \quad (2)$$

$$J_L = 330.512 \text{ gcm}^2$$

Therefore, I have

$$\Delta t = 0.543 \times 10^{-4} \Delta n$$

where,  $\Delta t$  is the time overhead that is required for a speed change of  $\Delta n$  RPM. This shows that the time cost is linearly proportional to the change in amplitude of RPM [2].

### 3.2. Energy Consumption Model

Let us introduce a mathematical model used to quantify the energy consumption of a disk as a function of disk rotation speed. The relationship between motor voltage and rotation speed of disk spindle motors can be expressed by Eq. (3) below. Thus, the motor voltage  $V$  is related to angular velocity, i.e., rotation speed  $\omega$ . Hence, I have

$$V = K_E \omega \quad (3)$$

where  $K_E$  is a constant. More details on this energy consumption model can be found in [11]. The power of the motor can be calculated as

$$P = VI = \frac{V^2}{R} \quad (4)$$

where  $R$  is the resistance of the motor. The power of a disk can be derived from Eqs. (3) and (4) as follows

$$P = \frac{K_E^2 \omega^2}{R} \quad (5)$$

Eq. (5) indicates that there is a quadratic relation between power consumption and the rotation speed of a disk motor. A disk consumes a significant amount of energy when it is in the idle state i.e. when the disk is not transferring data. The idle power calculation is derived using the following formula (see also [2]), where  $rpm$  is spindle motor speed.

$$P_{idle} = 1.318 \times 10^{-7} rpm^2 - 4.439 \times 10^{-4} rpm + 8.643 \quad (6)$$

The above Eq. (6) shows that the idle power is quadratic with respect to the spindle motor speed.

## 4. Security Implementation

Security services have increasingly become critical for data-intensive computing systems, where storage systems are major computing resources. Although a variety of security mechanisms have been developed to address the security issues in storage systems, there is a lack of adaptive way to dynamically choose security services to meet the ever-changing needs of secure storage systems. In other words, existing security mechanisms for storage systems generally do not factor in dynamically changing security requirements when choosing security services.

In many cases, security mechanisms require significant computational overheads. It is important to make best tradeoff between security and performance, which in turn affect energy efficiency of storage systems. Dynamic security needs can be straightforwardly met with an array of security mechanisms with various security strengths.

It is mandatory to deploy security services to protect security-critical applications that are also data-intensive. Since snooping and alteration are two common attacks in storage systems, let us address two security services (confidentiality and integrity services [4] [12]) to guard against the common threats to storage systems. Snooping, an unauthorized interception of information can be countered by confidentiality services. Alteration is an unauthorized change of information. Integrity services can be used to cope with threats of alteration. With the two types of security services in place, users can flexibly select from a wide range of security services to form an integrated security protection against all sorts of threats and attacks in a large-scale storage systems. Therefore, in this study I focus on confidentiality security services – one of the most popular security services that have been commonly implemented and adopted. Our approach is general in the sense that it can be extended to incorporate other security services depending upon the security requirements of data-intensive applications or of data sets handled by the applications.

Among many cryptographic algorithms, I decided to investigate IDEA (International Digital Encryption Algorithm) and have it implemented in our testbed. IDEA can provide security protection at different levels depending on the *number of rounds*. When the number of rounds in IDEA is changing, the performance of encryption varies accordingly. For example, the throughput of the 8-round IDEA is 13.5 KB/ms. If one intends to improve the security quality of security, the number of rounds must be increased from 8 to 24. However, 24-round transformation consumes roughly 3 times the energy of the 8-round IDEA encryption. In addition, the throughput of 24-round transformations much lower than that of the 8-round, as the performance of 24-round IDEA is as low as 5.4 KB/ms. Detailed information concerning the performance of IDEA can be found in [12]. In the next section, I discuss how to dynamically choose the number of rounds for IDEA while meeting energy-efficiency and security needs.

## 5. Improving Security and Energy Efficiency

Improving energy efficiency of security-aware storage systems is challenging, because security and energy efficiency are often two conflicting goals. There are two general ways of potentially achieving high energy efficiency and security. In the first approach, one may gradually improve security subject to power constraints (see, for example, [17]). In the second approach, energy efficiency can be boosted, subject to security requirements. To implement

these two general approaches, I will have to rely on energy efficiency requirements and security requirements. As such, I assume that both energy efficiency and security requirements of disk requests are derived from data-intensive applications issuing the requests. This assumption is practical, because I can extend the disk I/O interface by incorporating two parameters specifying energy efficiency and security requirements. Application programmers are responsible to specify the requirement parameters in accordance with the energy efficiency and security constraints of applications issuing disk requests.

Now I am positioned to propose two efficient ways of integrating confidentiality services with the dynamic disk speed control technique. Although I pay particular attention to a specific security service (i.e., confidentiality) and a widely investigated energy conservation technique (i.e., dynamic speed control), our approach is general in the sense that it can be readily applied to seamlessly integrate a wide range of security services with various energy conservation techniques.

#### A. Security Aggressive Approach (SAA)

The first security aggressive approach (SAA) is focusing on the improvement of storage system security under specified energy conservation constraints. The second strategy to be presented in the next sub-section gives higher priority to energy conservation as opposed to security optimization.

This security-aggressive approach can be employed in large scale storage systems where security is critical and a security breach is unacceptable. The security-aggressive approach or SAA systematically increases security levels of the confidentiality service for data stored in a storage system, subject to power constraints. The goal of this approach is to deliver high quality of security while maintaining energy efficiency requirements for storage systems. Before further optimizing the energy efficiency of a storage system, SAA attempts to aggressively increase the security level of each submitted disk request. Such a security level increasing process is suspended if the energy efficiency requirements of the storage system can no longer be met.

For each disk request issued to the storage system, the SAA scheme escalates the security level of the request by increasing the number of rounds in the IDEA while satisfying the following three conditions: (1) increasing the security level will not exceed the specified total energy consumption budget (i.e., energy efficiency requirement); (2) the increment of the security level must guarantee the desired I/O response time (i.e., performance requirement) of the disk request; (3) increasing the security level of the current request still can satisfy the desired response time of any previously admitted disk requests.

After the security levels of all the pending disk requests have been optimized, the SAA strategy can further enhance energy efficiency of the storage systems. Thus, once the security has been maximized using the 24-round IDEA, the security-aggressive approach makes an effort to save energy using the dynamic disk speed control technique. In the SAA

scheme, I save energy by dynamically changing the disk spindle motor speed based upon desired I/O response times of disk requests. Thus, the following two conditions must be met when the disk speed is adjusted: (1) the decreasing disk speed must guarantee the desired response times of the current disk request; (2) decreasing disk speed can satisfy the desired response time of previously admitted and pending requests.

To meet the above two conditions, I track the response time of  $n$  pending disk requests in a window called the  $n$ -request window. At the end of each  $n$ -request window, I calculate the percentage change in the response time compared to using the dynamic speed control technique.

When the percentage change is larger than an upper threshold, one can conclude that the desired response times of disk requests are unlikely to be guaranteed. Therefore, it is imperative to increase the disk spindle motor speed to serve disk requests faster. In contrast, if the percentage change is between the upper threshold and the lower threshold, then the disks maintain the same disk speed to serve the next  $n$  requests. If the percentage change is less than the lower threshold, it indicates that the desired response times are more likely to be met. Thus, one can reduce the speed of the disk spindle motors to conserve energy provided that the requests can meet their corresponding desired response times.

In our experiments I set the upper threshold and lower threshold as 20% and 10%, respectively. It does not imply by any means that the upper and lower thresholds must be fixed at 20% and 10%. The configuration of thresholds largely depends on data-intensive applications' needs, meaning that the thresholds may vary from application to application. Ideally, a mechanism should be implementing to dynamically adjust the upper and lower thresholds in accordance with changing workload environments.

#### B. Energy Saving Aggressive Approach (ESA)

Unlike the first SAA scheme, the second strategy – called energy-saving aggressive scheme (ESA) – gives higher priority to energy conservation as opposed to security and performance optimization. The energy-saving aggressive approach concentrates on providing significant energy savings while achieving minimum levels of security and I/O performance. Obviously, the second scheme can be employed if there is a pressing need to conserve energy in large-scale storage systems savings and to maintain a certain level of data security.

To create opportunities to reduce energy dissipation in disks, the energy aggressive approach initially and intentionally chooses to provide the minimum security by using the 8-round IDEA. Compared with 16-round IDEA with a higher security level, 8-round IDEA with a low security level can encrypt and decrypt data sets faster, leaving more opportunities to conserve energy by reducing the disk speed. Again, the energy-aggressive approach is capable of offering significant energy savings by the virtue of the dynamic disk speed control technique.

Given a set of newly issued disk requests, the disk speed is reduced to save energy while satisfying the following two


conditions: (1) the decreasing disk speed must guarantee the desired response times of all the pending disk requests; (2) decreasing disk speed must meet the minimum security requirements the disk request.

The implemented ESA scheme does not further enhance security for two reasons. First, the ESA approach is energy-saving aggressive. Second, high security levels come at the cost of energy efficiency. In the future, I will extend the ESA strategy by considering security optimization after the energy saving concerns have been addressed. In this extended ESA algorithm, the energy consumption of the disk system is minimized first. Then, the security levels of all the pending disk requests may be increased as long as their desired I/O response times are met.

## 6. Evaluation

In this section, I analyze experimental results obtained from extensive simulations. I generated disk traces to represent a wide range of disk workload conditions. I assume that disk requests arrive in a disk system according to a *Poisson* process. In our experiments, I varied both request inter-arrival times and average data size to study their impacts on the energy efficiency and performance of storage systems.

In the first experiment, I focus on workload conditions in terms of the disk request inter-arrival times. Fig. 1 plots the energy savings provided by the security-aggressive approach or SAA when the average I/O inter-arrival time is 100, 500, and 1000 ms. Note that the larger the inter-arrival times, the lighter the I/O workload condition. Thus, a 1000 ms inter-arrival time represents a light workload, whereas a 100 ms inter-arrival time represents a relatively high workload. In addition, I evaluated if changing the window size in SAA has an impact on the energy efficiency of the security-aggressive scheme. I achieve this goal by setting the window size to 100, 250, and 500. To calculate energy savings, I chose to compare a disk system employing SAA against the same disk system without deploying the SAA scheme.




**Figure 1.** The energy efficiency of the security-aggressive approach (SAA). The impact of request inter-arrival time on energy savings of SAA.

From Figure 1 I observe that with increasing window size, the energy savings gained by the security-aggressive scheme

become more pronounced. For example, the energy saving achieved by SAA under a window size of 500 is more than that of SAA under a window size of 100. A large window size leads to higher energy efficiency because large window sizes allow the security-aggressive scheme to modulate the speed of the disk spindle motor speed less frequently. The larger window sizes help in reducing the overhead incurred by the dynamic speed control mechanism.

Fig. 1 also shows that regardless of the window size, large inter-arrival times lead to large energy savings. For example, when the window size is 100, the energy savings are 35% and 50% as the inter-arrival time is set to 100 and 1000 ms, respectively. This result indicates that light I/O loads offer great opportunities for SAA to conserve energy by reducing the disk speed. Given a light workload, the disk has a strong likelihood to operate at lower speed for longer periods. The implication is that although SAA gives higher priority to security than energy efficiency, SAA is still able to significantly reduce energy dissipation in disk systems with lighter workloads.



**Figure 2.** I/O response time of the security aggressive scheme (SAA) using different types of disk systems.


Fig. 2 shows the degradation in response time for the SAA approach when the inter-arrival time is varied from 100 to 1000 ms. It is evident from Fig. 2 that the improvement in security and energy efficiency comes at the cost of performance (i.e., I/O response time). Results plotted in Fig. 2 show that there is a marginal increase in I/O response times of the disk requests.

One way to alleviate performance degradation due to energy savings is to increase the number of disks in a storage system. To evaluate the impact of the number of disks in a parallel disk system, I tested the performance of a parallel disk system with 8 disks. Fig. 2 indicates that performance degradation in a multiple disk system is smaller than that of a single disk system. This result suggests that the performance degradation problem can be partially solved by applying multiple disks.

From Figs. 1 and 2 I conclude that the SAA scheme can significantly reduce the energy dissipation of disk systems with marginal degradation in I/O response time. It is worthwhile to trade marginal performance degradation for substantial energy savings.

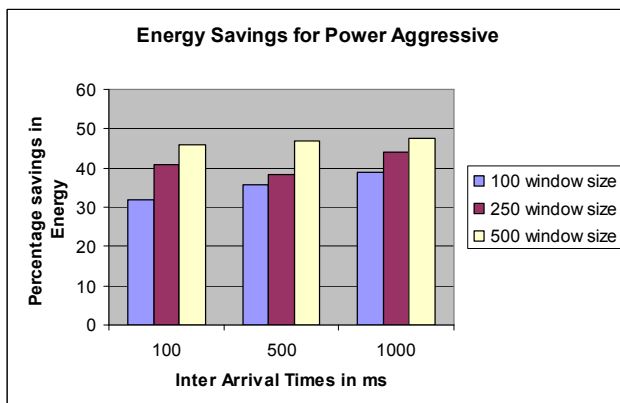
Now I analyze the impacts of I/O arrival rates and window

size on I/O performance degradation in response time. Fig. 3 shows the response time comparisons between two multiple-disk systems with and without applying the SAA scheme. In this set of experiments, the simulated disk systems contain eight disks. The results plotted in Fig. 3 are consistent with those illustrated in Fig. 2.



**Figure 3.** The performance evaluation of the security-aggressive approach. The impact of request inter-arrival time on response time.

Fig. 3 clearly shows that parallel storage systems with multiple disks effectively lower the performance degradation incurred by the dynamic speed control technique. In this specific case, the parallel disk system not only noticeably shortens the response time, but also successfully alleviates the performance degradations. I can attribute the results shown in Fig. 3 to the fact that there are multiple disks serving an array of disk requests in parallel. It is to be noted here that to make the comparison fair, I ensure two evaluated disk systems have identical energy consumption when measuring the response time.




**Figure 4.** The performance evaluation of the energy-saving aggressive approach. The impact of request inter-arrival time on I/O response time.

In what follows, I evaluate the performance and energy efficiency of the energy-saving aggressive scheme or ESA. Fig. 4 shows the energy savings achieved by the energy-saving aggressive approach. Compared to the results plotted in Fig. 1, the results depicted in Fig. 4 show that the ESA scheme is more energy efficient than the security-aggressive approach or SAA, as expected. For example, when the

window size is 100 and the average inter-arrival time is 100 ms, ESA reduced the energy consumption by 32% whereas the SAA scheme reduces energy consumption by 25%. The energy efficiency of ESA is higher than that of SAA, because ESA focuses on energy conservation rather than the security optimization by simply providing minimum security using the 8-round IDEA.

I also observe from Fig. 4 that energy savings gained by the energy-saving aggressive approach increases with increasing inter-arrival times. This energy-efficiency trend is the same as that of the security-aggressive approach (see Fig. 1). The results suggest that I can apply the ESA scheme to significantly improve the energy efficiency of a disk system under light I/O workload conditions. Higher energy efficiency offered by ESA becomes diminished when the I/O load is increasing.

The energy dissipation in disk systems can be further reduced by the energy-saving aggressive scheme when the window size is increased, because a large window size decreases the number of disk speed transitions that cause extra energy overhead incurred by disk speed transitions. Interestingly, I observe by comparing Figs. 1 and 4 that the energy-saving aggressive scheme is less sensitive to the window size than the security-aggressive scheme.



**Figure 5.** Impact of data size on energy efficiency of the security-aggressive and energy-aggressive approaches.

Now I measure performance degradation introduced by the ESA scheme. Again, I tested the baseline disk system without applying ESA to evaluate response-time degradation. Fig. 5 shows performance in I/O response times of disk requests when we varied the window size and inter-arrival time. The results indicate that compared with the baseline disk system, the average response time of the disk system with ESA is only increased by a few milliseconds. These results confirm that it is worth trading a few milliseconds performance degradation for reducing energy consumption in the disk system by an average of 42%.

After comparing the results in Fig. 2 with those in Fig. 5, I observe that the ESA noticeably shortens the I/O response time as compared to the security-aggressive scheme by an average of 12 ms. The main reason behind this I/O performance improvement is that the overhead incurred by the security service is more significant than the overhead introduced by the

dynamic speed control technique. These results suggest that the I/O performance can be substantially improved if I reduce the quality of security in disk systems while maintaining the minimal security requirements. It makes sense to improve I/O performance at the cost of security provided that no benefit is gained by increasing security levels after the minimal security requirements are met.

From Figs. 4 and 5 I conclude that it is possible to significantly conserve energy in disk systems with a slight degradation in I/O response time. Figs 4 and 5 indicate that the energy-saving aggressive approach performs better than the security aggressive approach. However, the security aggressive scheme achieve higher quality of security compared with the energy-saving aggressive approach.




Figure 6. I/O response time of the energy-saving aggressive scheme (ESA) using different types of disk systems.

Fig. 6 shows the I/O performance with respect to response time for the energy-aggressive approach or ESA when it is employed in parallel disk systems. For comparison purpose, I also plot the I/O response time in a single disk system. Similar to the case of SAA, the I/O performance of the ESA strategy can be significantly improved using parallel disks. For example, a parallel disk system helps in reducing the average I/O response time by more than 50% since multiple disk requests can be served by parallel disks simultaneously. Again, to make the comparison fair, I configure the disks in a way that the energy consumed by the parallel disk system is the same as the case of a single disk system.




Figure 7. Energy consumption of the energy-saving aggressive (ESA) and security aggressive (SAA) approaches when data size is increased.

Now I am positioned to quantitatively compare the energy efficiency and performance of the two approaches by varying the average data size of disk requests. As such, in this experiment the data size was increased from 100 KB to 800 KB. The energy efficiency of each approach is measured in terms of energy savings; performance is evaluated in terms of response times.

Fig. 7 shows that the energy savings yielded by the two approaches becomes more prominent when the data size is increased. The result suggest that disk systems serving large disk requests can benefit more in terms of energy savings from both the security-aggressive and energy-saving aggressive schemes. A second observation drawn from Fig. 7 is that the energy-saving aggressive scheme is more energy-efficient than the security-aggressive approach. This is because the energy consumed by encrypting data in the security-aggressive approach significantly increases when the data size becomes larger, indicating that there is a major difference in the energy consumption of the two approaches when the data size is increased.

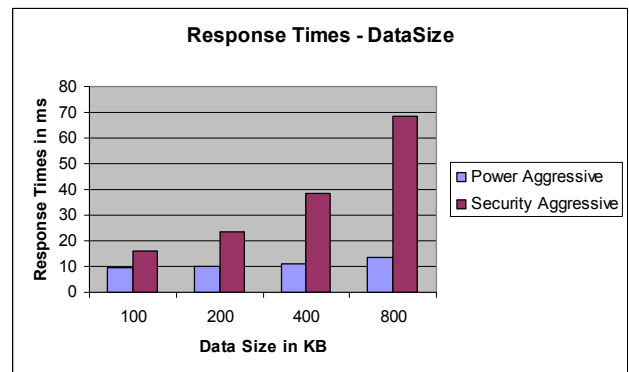


Figure 8. Impact of data size on I/O performance of the security-aggressive (SAA) and energy-saving aggressive (ESA) approaches.

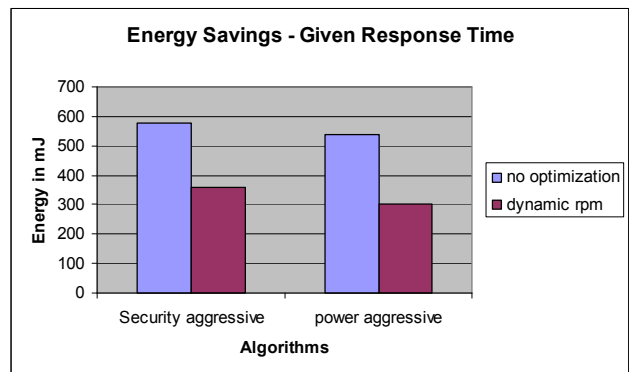


Figure 9. Energy consumption comparison of the security aggressive (SAA) approach and the energy-saving aggressive (ESA) approach.

Fig. 8 shows the impacts of data size on the I/O performance in terms of response time of the two approaches. As the data size increases, the I/O response times of the security-aggressive and energy-aggressive approaches both go up. Fig. 8 demonstrates that the performance of the energy-saving aggressive scheme is higher than that of the

security-aggressive scheme. For example, when the data size is set to 800KB, the response time of the energy aggressive approach is 14 ms, whereas the response time of the security-aggressive approach is as high as 69 ms.

Fig. 9 shows the energy consumption of the SAA and ESA approaches when the desired I/O response times are specified. Fig. 9 indicate that both approaches can conserve energy in disk systems while guaranteeing the specified desired response times. Energy savings are achieved by SAA and ESA by dynamically reducing disk speed provided that the desired response times can be met.

## 7. Summary and Future Work

In this study I designed and implemented two approaches that aim to integrate security services with the dynamic disk speed control technique. The first strategy, which is security aggressive in nature, is focused on the improvement of security followed by conserving energy in disk systems. The second strategy is an energy-saving aggressive approach that gives higher priority to energy conservation as opposed to security optimization.

The security-aggressive approach or SAA systematically increases security levels of the confidentiality service for data stored in a storage system, subject to power constraints. The goal of this approach is to deliver high quality of security while maintaining energy efficiency requirements for storage systems. Before further optimizing the energy efficiency of a storage system, SAA attempts to aggressively increase the security level of each submitted disk request. The energy-saving aggressive scheme or ESA makes an effort to conserve energy provided that the minimum security and performance requirements can be met. The energy-aggressive approach concentrates on providing significant energy savings while achieving minimum levels of security and meeting performance constraints.

Our experimental results show that the energy-saving aggressive approach has higher energy efficiency than the security-aggressive approach. However, the quality of security achieved by the security-aggressive scheme is higher than that of the energy-saving aggressive approach. Additionally, I evaluated the impacts of data size on energy efficiency and I/O performance of the two approaches. Numerical results show that the energy savings yielded by the two approaches become more pronounced when the data size is increased. The results also suggest that disk systems serving large disk requests can benefit more in terms of energy savings from both the security-aggressive and energy-saving aggressive schemes.

As for our future work, I will further extend the SAA and ESA strategies by considering disk requests with deadlines. The extended scheme is expected to conserve energy of real-time disk systems while maintaining certain levels of security. Moreover, I will consider additional security services like intrusion detections, authentication services, and authorization services.

## References

- [1] E. V. Carrera, E. Pinheiro, and R. Bianchini, "Conserving Disk Energy in Network Servers," *Proc. Int'l Conf. Supercomputing*, 2003.
- [2] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, "Dynamic Speed Control for Power Management in Server Class Disks," *Proc. 30<sup>th</sup> Annual Int'l Symp. Computer Architecture*, 2003.
- [3] J. S. Chase, and R. P. Doyle, "Balance of Power: Energy Management for Server Clusters," *Proc. 8<sup>th</sup> Hot Topics in Operating Systems Workshop*, 2001.
- [4] N. R. Potlapally, S. Ravi, A. Rag, and N. K. Jha, "A study of Energy Consumption Characteristics of Cryptographic Algorithms and Security Protocols," *IEEE Trans. Mobile Computing*, 2006.
- [5] F. Douglass, P. Krishnan, and B. Bershad, "Adaptive Disk Spin-Down Policies for Mobile Computers," *Proc. 2<sup>nd</sup> USENIX Symp. Mobile and Location-Independent Computing*, 1995.
- [6] K. Li, R. Kumpf, P. Horton, and T. Anderson, "Quantitative Analysis of Disk Drive Power Management in Portable Computers," *Proc. USENIX Winter Conf.*, 1994.
- [7] W. Freeman, and E. Miller, "An Experimental Analysis of Cryptographic Overhead in Performance-Critical Systems," *Proc. 7<sup>th</sup> Int'l Symp. MASCOT*, 1999.
- [8] I. Hong, and M. Potkonjak, "Power Optimization in Disk-Based Real-Time Application Specific Systems," *Proc. IEEE/ACM Int'l Conf. ICCAD*, 1996.
- [9] E. P. Harris, S. W. Depp, W. E. Pence, S. Kirkpatrick, M. Sri-Jayantha, and R. R. Troutman, "Technology Directions for Portable Computers," *Proceedings of IEEE*, 1995.
- [10] Maxon motor. <ftp://ftp.maxonmotor.com/public/Download>
- [11] T. Kenjo, *Electronic Motors and Their Controls*, Oxford University Press, 1993.
- [12] T. Xie, and X. Qin, "Scheduling Security-Critical Real-Time Applications on Clusters," *IEEE Trans. Computers*, vol. 55, no. 7, pp. 864-879, 2006.
- [13] T. Xie and X. Qin, "An Energy-Delay Tunable Task Allocation Strategy for Collaborative Applications in Networked Embedded Systems," *IEEE Trans. Computers*, vol. 57, no. 3, pp. 329-343, March 2008.
- [14] A. Manzanares, K. Bellam, and X. Qin, "A Prefetching Scheme for Energy Conservation in Parallel Disk Systems," *Proc. NSF Next Generation Software Program Workshop*, April 2008.
- [15] Z.-L. Zong, X. Qin, M. Nijim, X.-J. Ruan, K. Bellam, and M. Alghamdi, "Energy-Efficient Scheduling for Parallel Applications Running on Heterogeneous Clusters," *Proc. 36<sup>th</sup> Int'l Conf. Parallel Processing*, Sept. 2007.
- [16] S. W. Song, M. Kandemir, and A. Choudhary, "Software-directed disk power management for scientific applications," *Proc. Int'l Symp. Parallel and Distributed Processing*, April 2005.
- [17] R. Chandramouli, S. Bapatla, K. P. Subbalakshmi, and R. N. Uma, "Battery Power-Aware Encryption," *ACM Trans. Information and System Security*, pp. 162-180, May 2006.



- [18] X.-J. Ruan, A. Manzanares, K. Bellam, X. Qin, "DARAW: A New Write Buffer to Improve Parallel I/O Energy-Efficiency," *Proc. the 24th Annual ACM Symposium on Applied Computing*, March 2009.
- [19] T. Xie and H. Wang. "MICRO: A Multi-level Caching-based Reconstruction Optimization for Mobile Storage Systems," *IEEE Trans. Computers*, vol. 57, no. 10, p.1386-1398, 2008.
- [20] M. Nijim, A. Manzanares, and X. Qin, "An Adaptive Energy-Conserving Strategy for Parallel Disk Systems," *Proc. the 12th IEEE Int'l Symp. Distributed Simulation and Real Time Applications (DS-RT)*, Oct. 2008.