

AWARDS: An Adaptive Write Strategy for Secure Local Disk Systems

Mais Nijim, Xiao Qin^{*}, Tao Xie, and Mohamed ALghamidi

Department of Computer Science
New Mexico Institute of Mining and Technology
Socorro, New Mexico 87801
(mais,xqin,xietao,alghamdi)@cs.nmt.edu

Abstract

Since security is of critical importance for modern storage systems, it is imperative to protect stored data from being tampered or disclosed. Although an increasing number of secure storage systems have been developed, there is no way to dynamically choose security services to meet disk requests' flexible security requirements. Furthermore, existing security techniques for disk systems are not suitable to guarantee desired response times of disk requests. We remedy this situation by proposing an adaptive strategy (referred to as AWARDS) that can judiciously select the most appropriate security service for each write request while endeavoring to guarantee the desired response times of all disk requests. Experimental results show that AWARDS significantly improves security and overall performance over an existing scheme by up to 325.0% and 358.9% (with averages of 199.5% and 213.4%).

1 Introduction

In the past decade, storage systems have been an object of substantial interest because of an increasing number of emerging data-intensive applications like long running simulations[15], remote-sensing database systems[3], and out-of-core applications[11]. This trend can be attributed to the advances in computational power, disk performance, and high-speed networks. There are many cases where data-intensive applications require enriched security to protect data stored in storage systems from talented intruders. Further, a large number of data-intensive applications require guaranteed response times for interactive or high-priority data[6]. Therefore, storage systems are required to provide strong security and guaranteed response times for disk requests. This demanding re-

quirement is increasingly becoming a critical and challenging issue in the development of the next generation storage systems.

Although conventional storage systems are aimed at improving access times and storage space, many existing storage systems are vulnerable to a wide variety of potential threats. As such, the existing disk systems fail to meet the security requirements of modern data-intensive applications. To protect storage systems against all possible security threats, researchers have developed various ways of ensuring security of data in storage systems.

In this paper, we seek to present a novel adaptive write strategy for local disk systems providing a diversity of security services with various quality of security. The strategy can be seamlessly integrated into disk scheduling mechanisms in disk systems. The proposed strategy is conducive to the achievement of high security for the local disk systems while making the best effort to guarantee desired response times of requests. To prove the efficiency of the proposed approach, in our future study we will build an analytical model to measure the expected value of security levels and the probability that a disk request is completed before its desired response time.

The rest of the paper is organized as follows. In the next section we summarize related work. Section 3 describes the model of disk requests and the new architecture of storage systems. In section 4, we propose the adaptive write strategy for security-aware storage systems. Section 5 presents the experimental results. Finally, Section 6 concludes the paper with future directions.

2 Related Work

There is a large body of work in improving performance of disks, because disk I/O has become a serious performance bottleneck of computer systems. Previous techniques supporting high performance storage

^{*}Contact author. <http://www.cs.nmt.edu/~xqin>

[†]This paper appeared in the Proceedings of the 25th IEEE International Performance, Computing, and Communications Conference (IPCCC'06), Phoenix, Arizona, April 2006.

systems include disk striping[3] parallel file systems [7], load balancing[11], caching and buffering[7].

Disk scheduling algorithms also play an important role in reducing the performance gap between processors and disk I/O [4]. The shortest seek time first (SSTF) algorithm is efficient in minimizing seek times, but it is starvation-bound and unfair in nature [5]. The SCAN scheduling algorithm can solve the unfairness problem while optimizing seek times [5]. Reist and Daniel proposed a parameterized generalization of the SSTF and SCAN algorithms [12]. However, the above disk scheduling algorithms are unable to guarantee desired response times of disk requests.

Many data-intensive applications require that data is stored or retrieved before a desired response time [13]. Recently, many disk schedulers were implemented for a mixed-media data set, a mixture of data accessed by multimedia applications and best-effort applications [2]. Several disk-scheduling algorithms were proposed to provide quality of service guarantees to different classes of applications [2]. A salient difference between the proposed approach and the existing disk-scheduling algorithms in the literature is that our strategy is focused on maximizing security of a local disk. Moreover, our strategy is orthogonal to the existing disk scheduling policies in the sense that the novel strategy can be readily integrated into the existing disk schedulers to improve security of local disks.

In recent years, the issue of security in storage systems has been addressed and reported in the literature. Riedel et al. developed a common framework of core functions required for any secure storage system [14]. To protect data in untrusted storage systems, researchers designed and implemented cryptographic file systems where data is stored in encrypted form [1]. Several key distribution schemes were proposed in SFS [8] and SNAD system [9]. Although a variety of secure storage systems were implemented, there is no adaptive way of choosing security services to meet disk requests' flexible security requirements. Furthermore, the above security techniques are not suitable for disk requests with desired response times. We remedy this situation by proposing an adaptive strategy that can judiciously choose the most appropriate security service for each write request while making the best effort to guarantee the desired response times of all disk requests.

In our previous work, we proposed a family of dynamic security-aware scheduling algorithms for clusters [16]. Unfortunately, these scheduling algorithms limit their applicability to computing resources and,

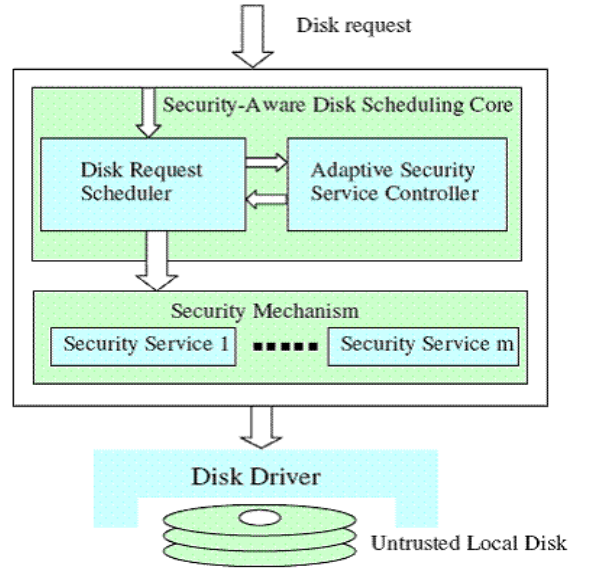


Figure 1: Architecture of a security-aware storage system

thus, our previous algorithms can not be employed to storage systems.

3 Architecture and Disk Requests with Security Requirements

3.1 Architecture of a Security-Aware Storage System

In this study we focus on local disk systems, and storage systems in parallel and distributed environments are out the scope of this paper. Our new algorithm is based on a security-aware storage system architecture. Figure 1. depicts the main components, i.e., a disk driver, a security mechanism, and a disk scheduling core, of this architecture.

The architecture is briefly overviewed as follows. The disk driver is responsible for controlling access to an untrusted local disk. The security mechanism provides an array of security services to guard data blocks residing in the local disk against unauthorized access and information theft. Without loss of generality, we consider in this study only confidentiality security services, where it is assumed that keys are known only to the owner, reader and writers [1]. The security mechanism can be readily extended to employ integrity and availability services. The disk scheduling core consists of two parts: a disk scheduler and an adaptive security service controller. While the scheduler implements generic logic and timing mechanisms for scheduling and waiting, the security service controller dynamically choose the most appropriate security

service for each disk request. Since the security service controller is independent of disk scheduling policies, the service controller is implemented separately for each disk scheduler. As such, it is easy to apply the security service controller to any disk scheduling policy implementation.

3.2 Modeling Disk Requests with Security Requirements

Each disk request submitted to a security-aware storage system specifies quality of service requirements, including security and performance requirements. A security requirement is defined as a lower bound security level, which is a value from 0.1 to 1.0. A performance requirement is posed as a desired response time. The quality of service requirements of disk requests can be derived from applications issuing these disk operations. Security quality of encryption services implemented in the security mechanism is measured by security levels. An encryption service with a high security level means the high quality of security provided by the service. For example, a disk request specifies a lower bound security level as 0.4. In this case, encryption services with security levels higher than or equal to 0.4 can successfully meet the disk request's security requirements.

A disk request r is characterized by five parameters, $r = (o, a, d, s, t)$, where o indicates that the request is a read or write, a is the disk address, d is the data size measured in KB, s is the lower security level bound, and t is the desired response time.

The security benefit gained by a disk request r_i can be measured by the security level σ_i of an encryption service facilitating confidentiality for the disk request. Likewise, the quality of security offered by a local storage system can be measured by a sum of security benefit of all incoming disk requests. Let R be a set of incoming disk requests. Our proposed AWARDS strategy strives to maximize the security benefit of the storage system. Thus, we can obtain the following non-linear optimization problem formulation to maximize the security benefit, where ρ_i is the real response time of the i^{th} disk request.

$$\text{Maximize } \sum_{r_i \in R} \sigma_i$$

$$\text{Subject to } \forall r_i \in R : s_i \leq \sigma_i \leq 1, \text{ and } \rho_i \leq t_i \quad (1)$$

3.3 Security Overhead Model

Now we consider security overhead incurred by confidentiality services. The security overhead model can be easily extended to incorporate other security services. Encryption is used to encrypt data blocks residing in local storage systems. There are nine encryption




Figure 2: Security overhead of encryption services

algorithms (see Table 1) implemented in the security mechanism. Based on the encryption algorithms' performance, each algorithm is assigned a security level. For example, level 0.9 implies that we use 3DES, which is the strongest yet slowest encryption function among the alternatives. Note that overhead of encryption depends on the chosen cryptographic algorithm and the size of data block. Figure 2. plots enciphering time in seconds as a function of the encryption algorithms and data size on a 175 MHz Dec Alpha600 machine [10][16]. Let σ_i be the security level of r_i , and the security overhead can be calculated using Eq. 3.2, where d_i is the data size and $P(\sigma_i)$ is a function used to map a security level i to the corresponding performance of encryption service listed in Table 1.

$$T_{security}(\sigma_i, d_i) = \frac{d_i}{P(\sigma_i)} \quad (2)$$

4 The Adaptive Write Strategy

This section presents the proposed adaptive write strategy, which is referred to as AWARDS throughout

Table 1: Cryptographic Algorithms Used for Encryption Service: Cryptographic Algorithms(CA), Security level(SL), Performance (P)

CA	SL, σ	P(KB/ms),P(σ)
SEAL	0.1	168.75
RC4	0.2	96.43
Blowfish	0.3	37.5
Knufu/Khafre	0.4	33.75
RC5	0.5	29.35
Rijndael	0.6	21.09
DES	0.7	15
IDEA	0.8	13.5
3DES	0.9	6.25




Figure 3: a-(top) Security levels and overhead of the requests. The system is running without AWARDS.
 b-(bottom) Security levels and response times of the requests. The system is running with AWARDS

this paper. We assume that the overhead of AWARDS is negligible when compared to the processing times of disk requests. In this study we consider a local disk system providing nine encryption services with different security levels (see Section 3.3). AWARDS aims at improving quality of security for local disk systems. To achieve this goal, AWARDS aggressively increases the security level of each incoming disk request under the condition that the request's response time does not exceed the desired response time. We make use of an example to elaborate the basic idea behind the AWARDS strategy. Suppose there are three write requests submitted to a local disk system at time 0. Table 2 shows the important parameters of the three write requests. We assume that the disk bandwidth is 30MB/Sec., and for each request, the sum of rotational latency and seek time is 8ms.

Prior to writing the data of a request to the local disk, the disk system enciphers the data using a selected encryption service. Similarly, the system decipheres the data of a read request after the data is retrieved from the disk. Hence, the processing time of each disk request logically consists of two parts: se-

Table 2: Important parameters of the three write requests. Requests(R), data size(d_i), Minimal security levels(s_i), Desired Response Time(t_i), Response time under AWARDS(T), Security Level under AWARDS(σ_i)

R	d_i	s_i	t_i	T	σ_i
r_1	90KB	0.2	18ms	17.7ms	0.8
r_2	150KB	0.1	41ms	40.7ms	0.7
r_3	30KB	0.3	55ms	54.5ms	0.9

curity overhead indicated by a shaded block and disk service time represented by an unshaded block (see Eq. (3) and Fig. 3). Figs. 3a illustrates the security levels and response times of the requests when the disk system is running without AWARDS. In this case, the minimal security requirement of each request is met at the minimal security overhead. In contrast, AWARDS can significantly increase the security levels provided that the desired response times can be guaranteed (see Fig. 3b). For example, the response times of the three requests are 17.7ms, 40.7ms, and 54.5ms, which are less than the respective desired response times (see Table 2). Specifically, with the AWARDS strategy in place, the security levels are improved by an average of 366.7%.

Now we present the AWARDS strategy, which adaptively adjusts security levels of write requests (see Fig. 4). It is worth noting that AWARDS is unable to adjust security levels of read requests, because the local disk system has to use a corresponding encryption service to decipher the data of a read request after the cipher is read from the disk. Before increasing the security level of a write request r_i , AWARDS must ensure that check that r_i and those write requests with earlier desired response times can be finished before their desired response times (see condition (2) in the following property). Therefore, the following property needs to be satisfied in the AWARDS strategy.

Property 1. If the security level of a write request r_i is increased by 0.1, the following conditions must hold:

(1) The current security level of r_i is less than

$$0.9, \text{ i.e., } \sigma_i < 0.9 \quad (3)$$

$$(2) \quad \forall r_k \in Q, t_k \geq t_i : es(r_k) + T(r_k, \sigma_k) \leq t_k; \quad (4)$$

where Q is the waiting queue, $es(r_k)$ is the start time of request r_k , and $T(r_k, \sigma_k)$ is the processing time of $r_k \in Q$. The start time of $es(r_k)$ can be expressed by

$$es(r_k) = \sum_{r_l \in Q, t_l \leq t_k} T(r_l, \sigma_l) \quad (5)$$

where the expression on the right side of Eq. (5) is the total processing time of disk requests whose desired response time is earlier than that of r_k .

The processing time $T(r_k, \sigma_k)$ in condition (4) can be computed by Eq. (6).

$$T(r_i, \sigma_i) = T_{seek}(a_i) + T_{rot}(a_i) + \frac{d_i}{B_{disk}} + T_{security}(d_i, \sigma_i) \quad (6)$$

```

1. Input:  $r_i$ , a disk request  $r_i$  submitted to the local disk system;  $Q$ , a waiting queue
2. Insert  $r_i$  into  $Q$  based on the earliest desired response time first policy;
3. for each write request  $r_j$  in the waiting queue  $Q$  do
4.    $\sigma_j \leftarrow s$ ;  $i^0$  Initialize the security levels  $\sigma_j$ 
5.   for each write request  $r_j$  in the waiting queue  $Q$  do
6.     if  $\sigma_j < 0.9$  then  $i^0$   $\sigma_j$  can be further increased  $\sigma_j$  (see property 1(1))
7.       Increase security level  $\sigma_j$  by 0.1;
8.       for each request  $r_j$  with longer desired response time do (see property 1(2))
9.         if  $r_i$  can not be completed before its desired response time then
10.           Decrease security level  $\sigma_j$  by 0.1; break;
11.       end for
12.     if  $\sigma_j$  is not increased in Step 7 then break;
13.   end for

```

Figure 4: The adaptive write strategy for local disk systems

where T_{seek} and T_{rot} are the seek time and rotational latency, $\frac{d_i}{B_{disk}}$ is the data transfer time that largely depends on the data size d_i and the disk bandwidth B_{disk} , and $T_{security}(\sigma_i, d_i)$ is the security overhead that lies in the security level and security-critical data size (see Eq. (2)).

The adaptive write strategy for secure local disk systems, or AWARDS, is described in Fig. 4. AWARDS aims at improving quality of security and guaranteeing disk requests' desired response times. To achieve high security, the AWARDS strategy optimizes the security levels of write requests (see Step 7).

Upon the arrival of a disk request, AWARDS insert the request into the waiting queue based on the earliest desired response time first policy, meaning that disk requests with earlier response time are processed first. Before proceeding to the optimization of security levels of write requests in the queue, AWARDS first initialize the security levels of all the write requests to the minimal levels (see Step 4). Step 7 then gradually enhance the security level of each request r_j under the conditions that (1) the current security level of r_j does not exceed 0.9 (see Step 6); and (2) the desired response times of the requests being processed later than r_j can be achieved (see Steps 8-10). The process of optimizing security levels repeatedly performs (see Step 5); and stops when a request's desired response time can not be met (see Step 12). In doing so, the AWARDS strategy can maximize the security levels of write requests (see Step 7) while guaranteeing the desired response times of all the disk requests in the queue (see Steps 9 and 10). The time complexity of AWARDS is evaluated as follows.

Theorem 1. The time complexity of AWARDS is $O(n^2)$, where n is the number of disk requests in the

Table 3: Disk Parameters (IBM Ultrastar 36Z15)

Size	18.4 GB
RPM	15000
Seek Time, T_{seek}	7.18 ms
Rotational Time, T_{rot}	4.02 ms
Disk Bandwidth, B_{disk}	30 MB/Sec

waiting queue.

Proof. To increase the security level of a request, it takes $O(n)$ time check condition (4) (see Step 8). Since there are $O(n)$ number of write requests in the waiting queue, the time complexity of optimizing security levels of write requests is: $O(n)O(n) = O(n^2)$.

5 Evaluation

To quantitatively evaluate the performance of the AWARDS strategy, we simulate a complete local disk system. We also simulated nine encryption services (see Table 2) on top of the local disk system. Disk parameters, which are shown in Table 3, are similar to those of the IBM Ultrastar 36Z15. Performance comparisons between a disk system with AWARDS and another system without employing AWARDS (referred to as the Original strategy) are discussed in Section 5.1. Section 5.2 studies the impact of data size on system performance. Section 5.3 examines the performance sensitivities to disk bandwidth. Finally, Performance impact of security requirements is evaluated in Section 5.4.

The following four important performance metrics are used to evaluate the proposed AWARDS strategy. (1) Satisfied ratio is defined as a fraction of total arrived disk requests that are found to be finished before their desired response time. (2) Average security level is measured by the average value of all the disk requests issued to a local disk system. (3) Average security overhead is measured in seconds. (4) Overall performance is defined as a product of the satisfied ratio and average security level. Our empirical results were obtained based on a wide range of random workload conditions. Table 4 outlines the important workload configuration parameters for the simulated disk system used in our experiments.

5.1 Overall Performance Comparisons

This experiment is aimed at comparing AWARDS against Original, which is a strategy without making use of AWARDS. To stringently evaluate the performance of AWARDS and its competitive strategy, in this experiment we set write ratio to 100%. We

Table 4: Workload Configuration

parameter	Value(Fixed)-(Varied)
Disk Bandwidth	30KB/s
Request Arrival Rate	(0.1,0.2,0.3,0.4,0.5) No./Sec.
Desired Response Time	10 Sec.
Security Level	(0.5)-(0.1,0.2,0.3,0.4,0.5,0.6,0.7,0.8,0.9)
Write Ratio	(100%)-(0%, 10%, 20%, 30%, 100%)
Data Size	(500 KB)-(300,400,500,600,700) KB

increased disk request arrival rate from 0.1 to 0.5 No./Sec. Other workload parameters were fixed to the same values as those listed in Table 4. Figure 5 plots the four performances metrics for the AWARDS and Original strategies. Fig. 5a reveals that AWARDS maintains a very close performance in satisfied ratio to the Original strategy. Fig. 5b shows that AWARDS significantly outperforms Original in average security level by an average of 138.2%. As the request arrival rate increases, the average security levels of the two strategies decrease. However, AWARDS always achieves higher average security levels compared with those of the Original strategy. This result can be explained in part by Fig. 5c, which shows that the average security overhead of AWARDS is constantly higher than that of the alternative. In other words, high security levels of AWARDS are achieved at the cost of high security overheads. Fig. 5d clearly reveals that AWARDS noticeably outperforms the Original one in terms of overall performance. Specifically, AWARDS obtains an improvement in overall performance over the Original strategy by an average of 125.6%. The performance improvement can be attributed to the fact that AWARDS adaptively enhance security levels of each write requests under the condition that all requests in the disk can achieve their desired response times.

5.2 Impact of Data Size

In this section, we varied data size from 300 to 700 KB to examine the performance impact of data size on the local disk system. Again, other workload parameters were kept unchanged (see Table 4). Fig. 6a shows that when the data size increases from 300 to 700 KB the AWARDS strategy delivers similar satisfied ratios as those of Original. This result, which is consistent with the result presented in Fig. 5a, demonstrates that AWARDS achieves a good performance in satisfied ratio. Like Fig. 5b, Fig. 6b shows a significant improvement of AWARDS in security level over




Figure 5: Performance impact of arrival rate

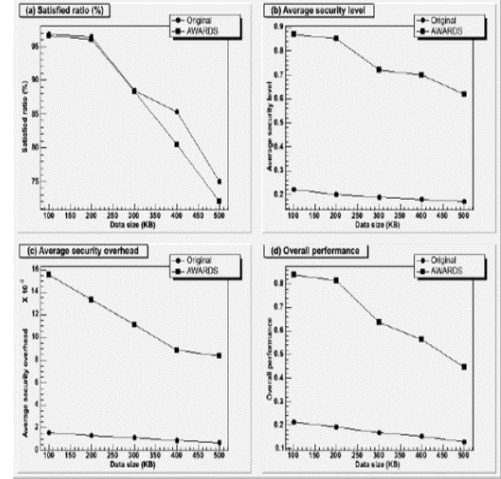


Figure 6: Performance impact of data size

Original. Interestingly, it is observed from Fig. 6b that the average security level gradually drops with the increasing value of data size. This is because the increasing data size results in a overloaded condition, which in turn leads to decreasing security levels of disk requests in the queue in order to finish most of the requests before their desired response time. Further, we observe from Figs. 6c and 6d that when the data size goes up, the average security overhead and overall performance of AWARDS decreases, because the security levels of disk requests are lowered down due to the high workload. By contrast, the average security level and overhead of the Original strategy only slight reduce with the increasing value of data size. These results indicates that Original is insensitive to data size.

5.3 Impact of Disk Bandwidth

In this, experiment we investigated the performance of AWARDS and Original when the disk bandwidth varies from 10 MB/sec to 50 MB/sec. An important observation drawn from Fig. 7a is that as the bandwidth increases, the satisfied ratios of the two strategies rise gently. The result can be explained by the fact that the high disk bandwidth leads to short transfer times, which in turn result in short processing times of disk requests. Consequently, more disk requests can be finished before respective desired response times. Interestingly, the satisfied ratio curves of the two alternatives begin to merge when the bandwidth is larger than 40 MB/sec. Fig. 7b shows that the average security level increases as the disk bandwidth is increased, because processing times of disk requests become smaller in the light of high disk bandwidth. The shortened processing times enable AWARDS to further increase security levels at the expense of higher security overheads (see Fig. 7c). Thanks to the increasing satisfied ratio and average security level, the overall performance of AWARDS is substantially boosted in case of a disk system providing high disk bandwidth (see Fig. 7d).

5.4 Security Level Range

In this group of experiments, we studied performance impact of the security requirements of disk requests. The maximal required security level is varied from 0.5 to 0.9, whereas the minimal required security level is fixed at a value of 0.1.

It is observed from Fig. 8a that when the maximal required security levels increase, the satisfied ratios of the two strategies decrease. The main reason for this result is that when disk requests require higher security levels, the security overhead is inevitably grows (see Fig. 8c). The growing security overhead in turn causes a significant drop in the satisfied ratio. We

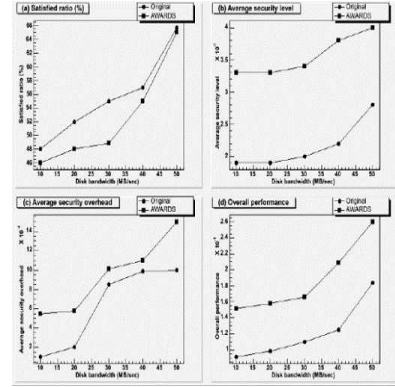


Figure 7: Performance impact of disk bandwidth




Figure 8: Performance impact of security level range

observe from Fig 8b that the amount of improvement in average security level becomes more prominent with the increasing value of maximal required security level. This performance trend can be explained by the fact that the larger the maximal required security level, the more opportunities for AWARDS to dynamically increase security level of each write request. Because the average security level rapidly increases, the overall performance of AWARDS also goes up as the maximal required security levels is increased (see Fig. 8d).

6 Summary

In this paper, we considered the flexible security requirements of disk write requests in the context of a local disk system. To protect stored data from being tampered or disclosed, we proposed a security-aware storage system architecture. Next, we developed an adaptive write strategy for secure disk systems (AWARDS for short). The AWARDS strategy can adaptively choose an appropriate security service

for each write request in a way to maximize the security of the local disk system while making efforts to achieve desired response times of all incoming disk requests. Experimental results demonstratively show that our strategy outperforms an existing scheme in security and overall performance by up to 325.0% and 358.9% (with averages of 199.5% and 213.4%). In our future study we will build an analytical model to measure the expected value of security levels and the probability that a disk request is completed before the desired response time, and the expected value of disk requests' security levels.

Acknowledgments The work reported in this paper was supported in part by the New Mexico Institute of Mining and Technology under Grant 103295 and by Intel Corporation under Grant 2005-04-070.

References

- [1] M. Blaze, "A Cryptographic File System for UNIX," *Proc. ACM Conf. Communications and Computing Security*, 1993.
- [2] P. Bosch and S. J. Mullender, "Real-Time Disk Scheduling in a Mixed-Media File System," *Proc. Real-Time Technology and Applications Symp.*, pp. 23-33, 2000.
- [3] C. Chang, B. Moon, A. Acharya, C. Shock, A. Sussman, and J. Saltz. "Titan: a High-Performance Remote-Sensing Database," *Proc. the 13th Int'l Conf. Data Engineering*, Apr 1997.
- [4] Jr. Coffman and M. Hofri, "Queueing Models of Secondary Storage Devices," *Stochastic Analysis of Computer and Communication Systems*, Ed. Hideaki Takagi, North-Holland, 1990.
- [5] P. J. Denning, "Effects of Scheduling on File Memory Operations," *Proc. AFIPS Conf.*, 1967.
- [6] Z. Dimitrijevic and R. Rangaswami, "Quality of Service Support for Real-time Storage Systems," *Proc. Int'l Conf. IPSI*, Sv. Stefan, Montenegro, October 2003.
- [7] B. Forney, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau, "Storage-Aware Caching: Revisiting Caching for Heterogeneous Storage Systems," *Proc. Int'l Symp. File and Storage Tech.*, 2001.
- [8] D. Mazieres, M. Kaminsky, M. Kaashoek and E. Witchel, "Separating key management from file system security," *Proc. ACM Symp. Operating System Principles*, December 1999.
- [9] E. Miller, D. Long, W. Freeman and B. Reed, "Strong Security for Distributed File Systems," *Proc. Symp. File Systems and Storage Technologies*, January 2002.
- [10] E. Nahum, S. O'Malley, H. Orman, R. Schroepel, "Towards High Performance Cryptographic Software," *Proc. IEEE Workshop Arch. and Impl. of High Perf. Comm. Subsystems*, August 1995.
- [11] X. Qin, H. Jiang, Y. Zhu, and D. R. Swanson. "Improving the Performance of I/O-Intensive Applications on Clusters of Workstations" *Cluster Computing: The Journal of Networks, Software Tools and Applications*, Vol. 8, No. 4, Oct. 2005.
- [12] R. Reist and S. Daniel, "A Continuum of Disk Scheduling Algorithms," *ACM Trans. on Computer Sys.*, pp.77-92, Feb. 1987.
- [13] L. Reuther, M. Pohlack, "Rotational-Position-Aware Real-Time Disk Scheduling Using a Dynamic Active Subset," *Proc. IEEE Real-Time System Symp*, 2003.
- [14] E. Riedel, M. Kallahalla, and R. Swaminathan, "A Framework for Evaluating Storage System Security," *Proc. the 1st Conf. File and Storage Technologies*, Monterey, CA, Jan. 2002.
- [15] T. Tanaka, "Configurations of the Solar Wind Flow and Magnetic Field around the Planets with no Magnetic field: Calculation by a new MHD," *J. Geophysical Research*, pp.17251-17262, Oct. 1993.
- [16] T. Xie and X.Qin, "A New Allocation Scheme for Parallel Applications with Deadline and Security Constraints on Clusters," *Proceedings of the 7th IEEE International Conference on Cluster Computing (Cluster'05)*, Boston, Sept. 2005.