Energy-Efficient Disk Replacement and File Placement Techniques for Mobile Systems with Hard Disks

Young-Jin Kim
School of Computer Science & Engineering
Seoul National University
Seoul 151-742, KOREA
youngjk@davinci.snu.ac.kr

Kwon-Taek Kwon
Computing Technology Lab
Samsung Advanced Institute of Technology
Gyeonggi-Do 446-712, KOREA
kwontaek.kwon@samsung.com

Jihong Kim
School of Computer Science & Engineering
Seoul National University
Seoul 151-742, KOREA
jihong@davinci.snu.ac.kr

ABSTRACT

Mobile systems have usually used hard disks as the secondary storage devices because of their high capacity per cost and high I/O throughput. However, their high power consumption is the main limiting factor for extending their adoptions in mobile systems. In this paper, we propose enhanced file placement techniques for mobile platforms with multiple smaller disks (instead of a single large disk). We investigate that how many smaller disks are necessary to obtain energy saving while maintaining the required performance using both a simplified energy model and a realistic trace-based simulator under the proposed multiple disk configurations. We also propose energy-efficient file placement techniques, which aggregate files with common attributes the same set of disks. By skewing I/O operations, the proposed techniques achieve additional energy saving. Experimental results show that the proposed techniques can reduce the energy consumption by up to 43% when eight 1″ disks are used instead of a single 2.5″ disk with an acceptable increase in the average response time.

Categories and Subject Descriptors

D.4.2 [**OPERATING SYSTEMS**]: Storage Management – *secondary storage.*

General Terms

Experimentation, measurement

Keywords

Energy saving, multiple disks, disk replacement, file placement, mobile systems.

1. INTRODUCTION

Despite close chase of flash memory, hard disk drives have been widely employed by the users for mobile and ubiquitous computing platforms. Therefore, the demand for disk drives with a small form-factor (2.5″ or less) is steadily rising in mobile devices such as PDAs, PMPs, MP3 players, and video camcorder [1]. This is because disks are still far more beneficial in capacity per cost and have high I/O bandwidth. However, since mobile systems operate based on batteries, significant power consumption of disk drives may shorten lifetimes of mobile systems critically, and so it is important to control the energy consumption of disk drives in mobile systems.

To date existing energy-efficient techniques mainly have been constrained to the mobile systems with a single disk drive and have not investigated the potential of energy saving when using more than one smaller disk instead of a larger one. In contrast to research on mobile storage systems, many recent studies of energy-efficient techniques based on arrays of multiple disks have been fulfilled in server storage systems [2, 3, 4, 5, 6]. Regardless of the scales of systems, techniques such as using multiple disks or distributing files in the viewpoint of energy efficiency can also be applied to mobile devices.

On the top of such background, we investigated that when we replace a larger high-power disk with two smaller low-power disks on a mobile system how much the energy saving of the system will be and how much the performance degradation will be [7]. We also proposed new file placement techniques which cluster related data into groups and migrate the correlated groups to one disk to obtain additional energy saving, compared with the previous data placement technique based on only the access frequencies of files. But in this work replacing a single 2.5″ disk with only two 1.8″ disks was considered and the proposed techniques were operated under a two disk configuration.

In this paper, our goal is to investigate how much the amount of energy savings can be when multiple (more than two) smaller form-factor disks are used for replacement and ascertain that how file placement techniques can be energy-efficient under generalized multiple disk configurations. First, we review performance and power characteristics of state-of-the-art small form-factor disks and examine how many smaller disks can be employed to save energy consumption while meeting tolerable performance degradation using a simplified disk energy model. Second, we propose energy-efficient file placement techniques with a generalized number of smaller disks, which aggregate files with common attributes to the same set of disks (not one disk) and skew I/O operations to save energy consumption, and evaluate combination of disk replacement and file placement techniques by simulation.

The rest of this paper is organized as follows. Section 2 presents characteristics of small form-factor disks and Section 3 describes how a 2.5″ disk can be replaced with multiple smaller disks to save energy consumption based on a simplified disk energy model.

SAC'07, March 11–15, 2007, Seoul, Korea.

Table 1. Characteristics of small form-factor disks

Parameters \ Form factor / Model		2.5″ Travelstar 80GN (Hitachi GST)	1.8″ MK4004GAH (Toshiba)	1″ ST1.3 (Seagate)	1″ Microdrive 3K8 (Hitachi GST)	0.85″ MK4001MTD (Toshiba)
Capacity (GB)		80	40	12	6, 8	4
Rotational speed (rpm)		4,200	4,200	N/A*	3,600	3,600
Avg. rotation time (ms)		7.1	7.1	N/A*	8.3	8.3
Avg. seek time (ms)		12	15	N/A*	12	16
Power (W)	Active	2.3	1.4	0.792	0.627	0.6
	Idle	0.95	0.4	0.254	0.264	0.45
	Standby	0.25	0.2	0.0429	0.066	0.12
Physical size	Weight (g)	99	62	14	13	8.5
	Area (cm^2)	70	42.39	12	12	7.68

(*: the average access time is shown as 21ms [10])

Section 4 explains energy-efficient file placement techniques that use file migration across multiple disks and Section 5 describes our simulation environment and presents simulation results. Related work and conclusions are given in Sections 6 and 7, respectively.

2. CHARACTERISTICS OF SMALLER DISKS

Table 1 shows parameters of five state-of-the-art small form-factor disks from manufacturers' documents: the Travelstar 80GN [8], MK4004GAH [9], ST1.3 [10], Microdrive 3K8 [11], and MK4001MTD [12]. The Travelstar 80GN is usually used in laptop computers and the other disks are used in handheld devices such as PDAs and PMPs.

From Table 1, we see that smaller disks tend to consume less power but have worse performance because their platters rotate more slowly. We also notice that the degree of performance degradation is different from that of power consumption improvement when comparing smaller form-factor disks with a 2.5″ one. For instance, the MK4004GAH consumes 1.4W in its active state and 0.2W in the standby state, which are about 1.6 times less and 1.25 times higher than the Travelstar 80GN. But, in terms of performance the access time of the MK4004GAH is 1.16 times longer than that of the Travelstar 80GN. The access time is calculated with a sum of an average rotation time and an average seek time. The MK4001MTD has low active power consumption 3.8 times and low standby power consumption 2.1 times as much as the Travelstar 80GN, but its access time is 1.27 times longer than that of the Travelstar 80GN.

Noticeable observations are two-fold. First, the inverse of the ratio of each smaller disk's power consumption over that of the 2.5″ disk tends to be higher than the ratio of each smaller disk's access time over that of the of the 2.5″ disk. This indicates that a considerable amount of energy saving is possible with a small delay if a smaller disk would be used instead of a 2.5″ disk. Second, the inverse of the ratio of each smaller disk's active power consumption over that of the 2.5″ disk is different from the ratio of each smaller disk's standby power consumption over that of the 2.5″ disk according to the type of a smaller disk. The MK4004GAH and MK4001MTD have higher inversed ratios of active power than those of standby power over the Travelstar 80GN. On the contrary, with the ST 1.3 and Microdrive 3K8, the inversed ratios of standby power are higher than those of active power. This characteristic should be considered importantly because the standby power consumption of a disk usually has much influence on the overall energy consumption of a disk-based storage system.

When we consider replacing a larger high-power disk with multiples of smaller lower-power disks, the differences between the ratios of each parameter such as access time, power consumption, capacity, and size should be investigated. We will focus on the access time and power consumption of each disk. This is because the trade-off between performance and power consumption influences the energy efficiency of the whole system directly, deciding the number of smaller disks for replacement. Then, we will also take into account the remaining parameters for practical replacement.

3. DISK REPLACEMENT

Motivational example in [7] shows that perfect file placement can derive a large amount of energy saving using two 1.8″ disks instead of a single 2.5″ disk. Our aim of this paper is to investigate that it is feasible to replace a 2.5″ disk with more than two smaller form-factor (maybe less than 1.8″) disks to save energy consumption while maintaining a required performance level. To do this, we take into account a simplified disk energy model based on disk replacement and perfectly distributed file placement as shown in Figure 1. Suppose that there are four files (A, B, C, and D) on a 2.5″ disk and each file is accessed sequentially and steadily during the same period T such that during each T the disk stays in the active state without entering a lower-power mode. Figure 1 (a) shows this behavior of the 2.5″ disk while files are accessed. If we assume the 2.5″ disk to be the Travelstar 80GN in Table 1 total energy consumption can be calculated as eq. (1) with the total elapsed time being $4T$ (the unit is assumed to be second).

$$E_{single} = 2.3 \times 4T \tag{1}$$

Let us assume that we replace this 2.5″ disk with four smaller disks. And we assume that files are equally located on the disks and I/O accesses are uniformly distributed across them. Then, each disk stays in the active state only during each period of file accesses as shown in Figure 1 (b). Since smaller disks have longer access time than the 2.5″ disk a period of file accesses can be represented by mT, where m is larger than 1. And, for the time, we assume that the inverse of the ratio of each smaller disk's power consumption over that of the 2.5″ disk is equal to the ratio of each smaller disk's access time over that of the of the 2.5″ disk (i.e. m). Then, total energy consumption of the disks can be calculated as

$$E_{multiple} = \left(\frac{2.3}{m}\right) \times 4 \times (mT) + \left(\frac{0.25}{m}\right) \times 3 \times 4 \times (mT) \tag{2}$$

where 2.3 is the active power and 0.25 is the standby power value of the Travelstar 80GN and smaller disks are assumed to enter the standby state immediately when there are no accesses. Since there

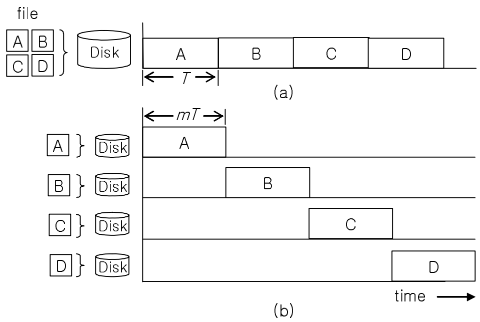

Figure 1. File access behaviors of a 2.5″ disk and multiple smaller disks under the perfect file placement

is no overlapped period during accesses the total elapsed time is $4mT$.

If we expect energy saving from disk replacement $E_{multiple}$ in eq. (2) should be less than E_{single} in eq. (1). Since $E_{single} = 9.2T$ in eq. (1) and eq. (2) is abbreviated into $E_{multiple} = 12.2T$ and $E_{multiple}$ is larger than E_{single}, we cannot acquire energy saving in this replacement.

Now we introduce the observations described in Section 2. Since the inverse of the ratio of power consumption is higher than the ratio of access time, eq. (2) is changed into eq. (3).

$$E_{multiple} = \left(\frac{2.3}{m_1}\right) \times 4 \times (m_2 T) + \left(\frac{0.25}{m_1}\right) \times 3 \times 4 \times (m_2 T) \tag{3}$$

where $1 < m_2 < m_1$. If we apply the other observation to eq. (3) it will be transformed into eq. (3′).

$$E_{multiple} = \left(\frac{2.3}{m_1}\right) \times 4 \times (m_2 T) + \left(\frac{0.25}{m_3}\right) \times 3 \times 4 \times (m_2 T) \tag{3′}$$

If we select the MK4001MTD in Table 1 for disk replacement the values of m1, m2, and m3 in eq. (3′) will be 3.83, 1.27, and 2.08, respectively. Since $E_{multiple} = 4.97T$ and $E_{multiple}$ is less than E_{single} we can acquire 46% energy saving. Here, if we take five MK4001MTDs $E_{multiple}$ in eq. (3′) becomes $7T$ and we still have the potential of energy saving.

Now we generalize the number of disks while obtaining energy savings. By augmenting 4 disks to n disks eq. (3′) is transformed into eq. (4) with the total elapsed time being $4m_2T$. Since we assume that only one file is ideally located on each disk, the number of files is equal to n, which is the number of disks.

$$E_{multiple} = \left(\frac{2.3}{m_1}\right) \times n \times (m_2 T) + \left(\frac{0.25}{m_3}\right) \times n(n-1) \times (m_2 T) \tag{4}$$

Table 2 summarizes the values of m_1, m_2, and m_3 in eq. (4), which can be easily calculated from Table 1. And Figure 2 shows energy consumption of each smaller disk with the number of disks varying using eq. (4). The energy consumption values are normalized on the basis of that of the Travelstar 80GN. We can notice that when the number of disks are set as 2 all the disk type can derive significant energy saving. But, we notice that as the number of disks increases only two 1″ disks (i.e. the ST1.3 and Microdrive 3K8) still save energy by up to at least 25%. For the MK4004GAH, when only two disks are used the energy can be saved. This observation corresponds to the motivational example of disk replacement in [7], which revealed that only two MK4004GAHs may replace a Travelstar 80GN in terms of energy saving with tolerable performance degradation.

Table 2. Ratios of active power, standby power, and access time of smaller disks over the Travelstar 80GN

Parameters	MK4004 GAH	ST1.3	Microdrive 3K8	MK4004 MTD
m_1	1.64	2.9	3.67	3.83
m_2	1.16	1.1	1.06	1.27
m_3	1.25	5.83	3.79	2.08

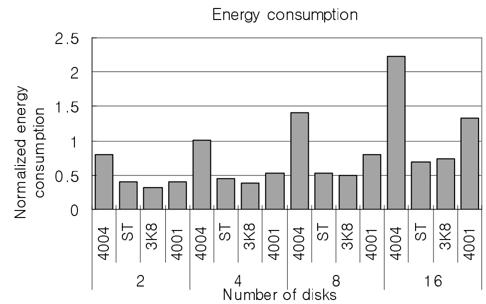

Figure 2. Normalized energy consumption of smaller disks varying the number of disks

One remarkable point is that multiples of the MK4001MTD, which has the lowest active power consumption and the smallest form factor among the disks, may save a small amount of energy while the ST1.3 and Microdrive 3K8 save a large amount of energy. This may be because 1) the standby power of the MK4001MTD is much higher than that of the ST1.3 (or Microdrive 3K8). Since the remaining disks except active one can be put into the standby state due to perfect file placement during the whole execution time, higher standby power will make a greater impact on the total energy consumption of the system; 2) although the MK4001MTD has less active power consumption than the ST1.3 (or Microdrive 3K8) it has longer access time than the ST1.3 (As is shown in Table 2, the access time delay of the MK4001MTD is 13% and 16.5% longer than that of the ST1.3 and Microdrive 3K8, respectively.). In short, in case of the MK4001MTD the saved energy due to relatively lower active power is offset by the increment due to higher standby power and relatively longer access time, compared with the ST1.3 or Microdrive 3K8. And this offset goes further as the number of disks grows and energy saving becomes unavailable relatively fast.

With the ST1.3 and Microdrive 3K8, the maximum numbers of disks which can save energy ideally are 31 and 24, which are invisible in Figure 2. If we should consider capacity, weight, and area in Table 1 simultaneously when replacing a Travelstar 80GN with multiple ST1.3s or Microdrive 3K8s, the area parameter restricts the feasible number of disks to 5, which is much less than the maximum number. And we notice that if we have space enough for accommodating more than 5 disks, the ST1.3 will be more qualified for practical replacement than the Microdrive 3K8 because 7 ST1.3s have capacity comparable to that of a Travelstar 80GN but 7 Microdrive 3K8s don't.

4. ENERGY-EFFICIENT FILE PLACEMENT

In this section, we describe energy-efficient file placement techniques, which aggregate files with common attributes (e.g., degree of correlation or access frequency) to the same set of disks and skew I/O operations to save energy consumption. We review the PDC technique which concentrates frequently-accessed file onto the same set of disks [5]. Next, we describe our energy-efficient file placement techniques simply, which are COR and

COMBINED. The COR technique concentrates correlated file requests onto the same set of disks and lets other disks have more idle times and the COMBINED technique combines the PDC and COR techniques simply (For details, refer to [7]).

4.1 PDC: Popular Data Concentration

The idea of PDC is to concentrate the most popular (i.e. frequently-accessed) data by migrating it to a subset of the disks, so that the other disks can be sent to a lower-power mode to save energy. PDC redistributes data across the disk array according to its popularity, in an orderly fashion. The first disk then stores the most popular data, the second disk stores the next most popular data, and so on. The least popular data and data that is apparently never accessed will be stored on the last few disks. Files are migrated to the target disk until it is full or the expected load approaches its maximum bandwidth.

However, if the frequency of file access varies significantly with time, PDC may cause a lot of file migration, which itself uses energy and also limits the possibility of energy saving by idle disks. Furthermore, when new files are created they will be stored on the disk with the least popular disk data, interrupting the sleep of that disk. Therefore, if PDC is applied to mobile workloads with varying file popularity, the energy saving may be limited.

4.2 COR: Load Skewing of Correlated Data

The idea of COR is to concentrate most of the correlated data onto one set of disks. *Correlated data* is data which is repeatedly occurred in the same order. This occurs when data requests are generated in the same order by an application (for instance, requests for libraries or configuration files) or files are sent to the disk system so that they can be shared between applications (for instance, when the output of one application is input to another). For example, if files are accessed in the order A, B, C, D, E, B, C and D, then B, C and D will be identified as a group of correlated data. If B, C and D are accessed frequently and all reside on one disk, the other disks are likely to enjoy long idle times while these files are accessed, and can enter lower-power states. Since many such groups of correlated data are found in mobile workloads, this method of load skewing can lead to significant energy saving. The goal of COR is to distribute data across more than two sets of disks so that one set of disks stores most of the correlated data.

COR operates in two phases: first, data accesses that take place in the same order are identified and classified into groups; second, data is moved between disks. To identify and classify groups of correlated data, the system creates a pointer from each file accessed to the one to be accessed next. Although files may be accessed in different sequences, we only keep information for the sequence following the next recent access to determine whether the further accesses take place in the same order based on [13].

4.3 COMBINED: Combining COR and PDC

PDC moves popular data to a fixed subset of disks and this may involve an unnecessarily large number of migrations. Compared with PDC, COR has no way of identifying groups of correlated data on pre-assigned disks and may therefore miss additional chances of concentrating disk I/Os. Furthermore, if the recently accessed data exhibits a low correlation, even though each file is frequently accessed, COR may lose the chance of skewing the load to save energy because it has no knowledge of the popularity of each file.

The COMBINED technique combines COR and PDC to overcome these disadvantages [7]. COMBINED is a simple combination of techniques: PDC lets file migration occur based on the frequency with which each file is accessed. At the same time, COR repeatedly tries to move groups of correlated data to the same set of disks. If a file migration request by PDC conflicts with one by COR, COR has priority. Thus, when file migrations result from popularity, only the files which are not registered in the current groups of correlated data will be moved.

5. EXPERIMENTAL RESULTS

In Section 3, we investigated how many smaller disks can be used to save energy consumption using a simplified disk energy model based on perfect file placement. In this section, we verify if the result of investigation on multi-disk replacement is feasible and evaluate how the file placement techniques can affect energy efficiency of mobile platforms with multiple small disks through simulations.

5.1 Simulation Setup

We implemented a multi-disk power and performance simulator by extending the trace-based simulator of [7], which models power and performance of only two disks and applies various file placement techniques along with data migration. The performance and power parameters of all disks are the same as those given in Table 1, but the capacities of the 2.5″ disk and the other disks are bounded to 800MB and 400MB, respectively.

And we built a new synthetic trace generator to produce various and realistic mobile workloads. Our trace generator can control request rate, read/write ratio, maximum size of each file, and total size of files. It maintains 1000 file identifiers and creates 100 pools from them. Each of 99 pools has a sequence with a limited number of file identifiers, which is generated randomly in order to represent an individual file access pattern. The last (that is, 100th) pool has all the file identifiers. According to the given request rate and maximum file size, the trace generator selects one from the 100th pool or the remaining pools randomly and with equal probability, and it issues I/O requests for the file identifiers of the selected pool randomly until the elapsed time approaches the given trace length. Default parameters are as follows: mean request inter-arrival time = 70 (ms), trace time = 80 (min), maximum file size = 20 (MB), total file size = 300 (MB), write ratio=0.4.

In comparing PDC, COR, and COMBINED file placement techniques, we assume that PDC migrates files every 5 minutes and COR does when the number of correlated data groups is beyond a given threshold. We compared the energy consumption and average request response time of the disk-based storage systems with 2, 4, 8, and 16 disks against using a single 2.5″ disk and used a threshold-based power control policy for power management of the disk(s).

5.2 Simulation Results

Figure 3 (a) and (b) show the energy consumption and average request response time of four different disks with the disk number varying when COR is applied to the trace generated with default parameters (We omitted the results of PDC and COMBINED as they are rather similar to that of COR.). And Figure 3 (c) and (d) show the energy distribution and disk time breakdown over four disks when the number of disks is four for each disk type and COR is applied. We notice that Figure 3 (b) is very similar to

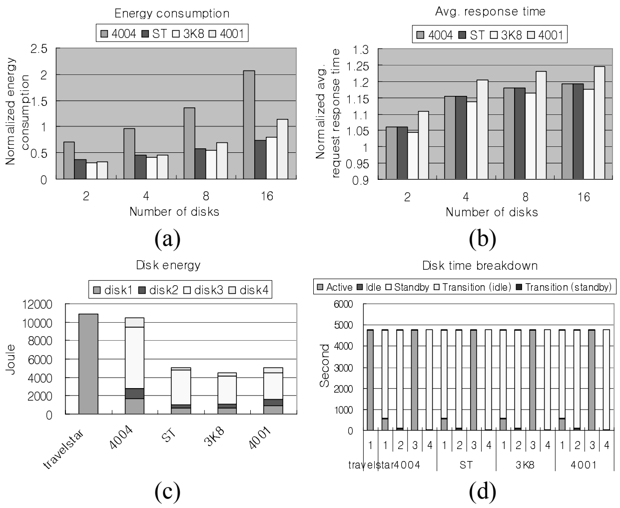

Figure 3. (a) energy consumption and (b) avg. request response time of smaller disks varying the number of disks. (c) energy distribution and (d) disk time breakdown when four disks are used for each disk type.

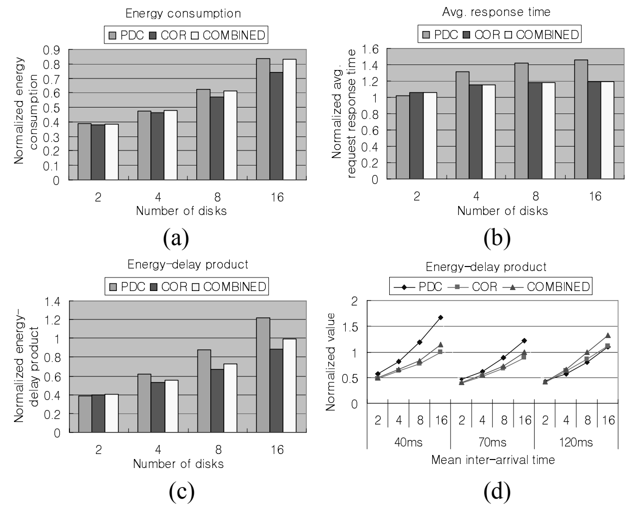

Figure 4. (a) energy consumption, (b) avg. request response time, and (c) energy-delay product of each file placement technique varying the number of disks. (d) enery-delay product according to a different mean inter-arrival time.

Figure 2 in Section 3 and this means that the simplified disk model with perfect file placement corresponds well to the realistic one with dynamic file placement according to realistic workloads in terms of energy consumption and performance.

The Microdrive 3K8 consumes the least energy for all the number except sixteen of disks. This is because since initial distribution of files is uniform over the disks COR involves almost all the disks in file migration, and the lowest active power of the Microdrive 3K8 affects the total energy consumption. As shown Figure 3 (c), comparing the results of the ST1.3 and Microdrive 3K8, the energy consumption of the third disk of the Microdrive 3K8 is less than that of the third disk of the ST1.3. According Figure 3 (d) the third disk for each disk type is active almost all the time.

And the MK4001MTD is shown to consume more energy than the ST1.3 or Microdrive 3K8. As described in Section 3, this can be mainly contributed to considerable standby power consumption and longer access time. We notice that in Figure 3 (c) the third disk's energy of the MK4001MTD is less than that of each 1″ disk but all the remaining disks' energy consumptions of the MK4001MTD are more than those of each 1″ disk. This is from the fact that the MK4001MTD consumes a little less active energy but much more standby energy than either of the 1″ disks. But, the effect of longer access time was mitigated for the average request response time, as shown in Figure 3 (b). We found that each actual seek time was shorter than the nominal seek time in Table 1 and for the given trace the seek time had a rather low portion out of the total response time, which consists of a seek time, a spin-up delay, a queue delay, and so on. And due to the uniform initial distribution of files over disks, the average response time increases as the number disks increases.

Since when the number of disks is two and four the energy savings of the MK4004GAH are 29% and 4%, respectively, we can replace a single Travelstar 80GN with two or four MK4004GAHs. But we also should consider the aspect of response time, and if we consider an energy-delay product metric more than two MK4004GAHs are not possible because the value goes beyond 1 (Due to the space limit a figure on energy-delay product was omitted, but we may figure out the values of each disk from Figure 3 (a) and (b).). We found that all disk types except the MK4004GAH have still lower energy-delay product values than 1 with eight disks. Moreover, this was found to be true for sixteen ST1.3s or Microdrive 3K8s. Therefore, if we only take into account the two aspects of energy consumption and response time replacing a 2.5″ disk with multiples of smaller disks may be said to be feasible. Specifically, eight ST1.3s and Microdrive 3K8s saved energy by 43% and 45%. But, if we consider a practical replacement based on capacity, weight, and area in addition to the two parameters, the number of possible disks will be limited below eight.

Figure 4 (a) and (b) show the energy consumption and average response time of the ST 1.3 with the number of disks varying when PDC, COR, and COMBINED are applied to the same trace. The ST 1.3 was selected because it is the second best energy saver but has the largest capacity except the MK4004GAH. COR exhibits the least energy consumption among three file placement techniques and saves more energy than PDC by up to 12% and has an 18% improved response time over PDC at best. Figure 4 (c) shows that COR and COMBINED have better performance than PDC and 16 ST 1.3s along with these two techniques can replace a single 2.5″ disk in terms of energy-delay product. And we found that these techniques save 48% and 43% energy when five and seven ST 1.3s are used under the constraints of area, weight, and capacity for feasible replacement, respectively. Figure 4 (d) shows the energy-delay product of three techniques when a mean inter-arrival time between the requests is 40, 70, and 120ms (70ms is a default parameter.). We notice that as the load goes higher PDC has a larger energy-delay product value due to the increased migration overhead while COR and COMBINED show better performance. This result may be said to correspond to the behaviors of PDC, which was shown in [7].

The main conclusions from the simulation results can be summarized as follows. First, replacing a 2.5″ disk with multiple smaller disks is feasible and can be beneficial in saving energy with negligible performance degradation. Second, the enhanced file placement techniques can save more energy under more than two smaller disk configurations.

6. RELATED WORK

Carrera *et al*. [2] and Papathanasiou *et al*. [3] investigated the possibility of saving energy by replacing high-speed server disks with arrays of smaller form-factor disks with almost the same aggregate I/O throughput. But Carrera *et al*. do not look at replacing high-performance disks with a set of lower-power disks deeply, and they focus on using two-speed SCSI disks to save energy. Papathanasiou *et al*. simply assumed that the original contents of a server disk are already mirrored on three laptop disks, and do not consider data migration.

Gurumurthi *et al*. [4] proposed a multiple-speed disk technique called dynamic rotations per minute (DRPM) for disk array based servers, which exploits access patterns that exhibit short idle intervals. But, DRPM has no migration of data between disks.

Pinheiro *et al*. [5] proposed a technique called popular data concentration (PDC) that dynamically migrates popular disk data to a subset of the disks in an array for network servers. The energy efficiency and performance delay of PDC is shown to vary considerably, depending on parameters such as request rate and migration period.

Zhu *et al*. [6] proposed a combinational technique called Hibernator, which combines intelligent speed setting and data migration to save energy on multi-speed disk arrays, with response time guaranteed by a service-level agreement (SLA). Hibernator exploits a RAID5-like striping scheme to achieve redundancy, and its migration techniques are largely specific to database servers.

The above techniques are all targeted to server workloads and our method aims at mobile workloads. In the meantime, in [7] novel file placement techniques under a two disk configuration were proposed, which cluster related data into groups and migrate the correlated groups to one disk while the other disk is put into a lower-power mode. But, This work used only two 1.8″ disks instead of a single 2.5″ disks and did not investigate generalized disk replacement and file placement problems for mobile systems with more than two smaller (equal to or less than 1.8″) disks.

7. CONCLUSIONS

In recent years, for small form-factor disks much effort has been made not only to improve the demerits of hard disks but also to enlarge capacity while making size smaller. We expect that in the near future the prices of normal small form-factor disks would go down due to technical innovation and using multiple disks might be sufficiently feasible.

We have proposed enhanced file placement techniques for mobile platforms with multiple smaller disks instead of a single large disk. We investigated that how many smaller disks are necessary to obtain energy saving while maintaining the required performance using both a simplified energy model and a realistic trace-based simulator under multiple disk configurations. We also proposed energy-efficient file placement, which aggregate files with common attributes the same set of disks and skew I/O operations to put the other sets of disks into lower-power modes. Trace-based simulation results showed that the proposed techniques can reduce the energy consumption by up to 43% when eight 1″ disks are used instead of a single 2.5″ disk with an acceptable increase in the average response time.

In the proposed techniques, we assumed that there is at most one active disk at a given time. That is, we did not consider parallel disk accesses to multiple disks which may improve the average response time as well as the energy efficiency of the mobile storage system. Our future work focuses on efficiently exploiting these overlapped accesses for a higher energy efficiency.

8. ACKNOWLEDGMENTS

This research was supported by the MIC (Ministry of Information and Communication), Korea, under the ITRC (Information Technology Research Center) support program supervised by the IITA (Institute of Information Technology Assessment). The ICT at Seoul National University provided research facilities for this study.

9. REFERENCES

[1] L. D. Paulson, "Will hard drives finally stop shrinking?," *IEEE Computer*, vol. 38, no. 5, pp.14-16, 2005.

[2] E. V. Carrera, E. Pinheiro, and R. Bianchini, "Conserving disk energy in network servers," in *Proc. of the 17th International Conference on Supercomputing*, June 2003.

[3] A. Papathanasiou and M. Scott, "Power-efficient server-class performance from arrays of laptop disks," Technical Report 837, Department of Computer Science, University of Rochester, May 2004.

[4] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, "DRPM: Dynamic speed control for power management in server class disks," in *Proc. of the International Symposium on Computer Architecture*, June 2003.

[5] E. Pinheiro and R. Bianchini, "Energy conservation techniques for disk array-based servers," in *Proc. of the 18th International Conference on Supercomputing (ICS'04)*, June 2004.

[6] Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton, and J. Wilkes, "Hibernator: helping disk arrays sleep through the winter," in *Proc. of the 20th ACM Symposium on Operating Systems Principles*, Oct. 2005.

[7] K.-T. Kwon, "Low-power file placement techniques for embedded systems with multiple disks," Master's thesis, Seoul National University, Aug. 2006.

[8] Hitachi GST, Travelstar 80GN. http://www.hitachigst.com/tech/techlib.nsf/products/Travelstar_80GN.

[9] Toshiba, MK4004GAH. http://www3.toshiba.co.jp/storage/ english/spec/hdd/mk4004gs.htm.

[10] Seagate, ST1.3. http://www.seagate.com/products/consumer_electronics/st1series.html.

[11] Hitachi GST, Microdrive 3K8. http://www.hitachigst.com/tech/techlib.nsf/products/Microdrive_3K8.

[12] Toshiba, MK4001MTD. http://www3.toshiba.co.jp/storage/english/spec/hdd/mk4001.htm.

[13] T. M. Kroeger and D. D. E. Long, "The case for efficient file access pattern modeling", in *Proc. of the Seventh Workshop on Hot Topics in Operating Systems*, March 1999.