Solid State Disk (SSD) Management for Reducing Disk Energy Consumption in Video Servers

Minseok Song and Manjong Kim
School of Computer and Information Engineering, Inha University, Incheon, Korea

***Abstract*—A promising approach to reducing disk energy consumption is to use multi-speed disks with lower rotational speeds, and allowing disks to run slowly when workloads are light can reduce their large contribution to the power used by video servers. We propose an SSD cache management scheme for video servers which use multi-speed disks. We formulate an integer linear problem (ILP) that determines videos cached on the SSD with the aim of minimizing overall disk energy consumption. Simulations show that our caching scheme allows disks to run at lower speeds, which saves disk energy consumption greatly.**

I. Introduction

The increasing number of video application services such as digital libraries, UCC (user created contents) and movie-on-demand increases the storage demand of the data centers greatly, which makes the energy consumption of video servers a significant problem. Servers are typically built on a redundant array of independent disk (RAID) which are one of the biggest energy consumers among the components of a server.

Reducing the speed at which a disk spins reduces its power consumption. Gurumurthi et al. [1] have suggested that multi-speed disks could save a lot of power consumption in servers. Based on this, Zhu et al. [2] proposed a flexible disk placement architecture called Hibernator to reduce disk speeds as often as possible. Multi-speed disks are now on the market (e.g. Hitachi drive [3] and the SONY drive [4]), and are being used to build energy-efficient storage systems [5].

Flash-based solid state disks (SSDs) provide many technical merits such as low-power consumption and shock resistance [6]. But due to their high cost, it is almost impossible to use SSDs solely for video servers that typically require large storage space.

We propose an SSD cache management scheme for video servers which use multi-speed disks. We formulate an integer linear problem (ILP) that determines videos cached in the SSD with the aim of minimizing disk energy consumption, and explore how cache allocation affects the overall energy consumption.

II. System Model and Problem Formulation

To support the periodic retrieval of a video server, we use round-based scheduling in which time is divided into equal-sized periods, called rounds, and each client is served once in each round [7]. For example, to serve a stream at 1.5Mbps with a round length of 2 seconds, the server needs to read 3Mbits of data during every round.

A multi-speed disk can service some level of requests at reduced speeds without transitioning to a full speed, thus saving energy. Suppose that each disk has NS speed levels. We use a typical seek time model in which a constant seek time T_s is required for one read of contiguous data [7]. The rotational delay and the data transfer rate vary with the speed level. Let $r_d(l)$ be the data transfer rate and $T_d(l)$ the rotational latency of a disk running at speed level l $(l = 1, ..., NS)$. Let ND be the number of disks in the array.

Each video file i has a size of S_i and a bit-rate of b_i, $(i = 1, ..., NV)$, where NV is the number of videos. Disk bandwidth utilization is defined to be the ratio of total service time to round length [7]. Let R be the length of round. The amount of data that must be read during R for video i is b_iR to keep up with playback rate, and reading a video file i incurs a seek and rotational delay overhead of $T_s + T_d(l)$ and a reading time of $\frac{b_iR}{r(l)}$; therefore, reading video i increases disk bandwidth utilization by $\frac{T_s+T(l)+\frac{b_iR}{r(l)}}{R}$.

We calculate disk bandwidth per unit storage of each video i when the lowest speed level is selected, wv_i which is $\frac{NC_i(T_s+T_d(1)+\frac{b_iR}{r(1)})}{S_iR}$, where NC_i is the number of concurrent clients requesting video i. We introduce the arrays of videos, where each element $V_{k,m}$ represents a video with the m^{th} highest value of wv_i among videos stored on disk k $(k = 1, ..., ND)$. We cache videos with higher wv_i values onto SSD first, because serving these videos from SSD reduces disk bandwidth utilization most. We thus have the following combinations of videos cached on SSD for disk k as follows: $CV_k = \{\phi, \{V_{k,1}\}, \{V_{k,1}, V_{k,2}\}, ..., \{V_{k,1}, ..., V_{k,NE_k-1}\}\}$, where NE_k represents the number of elements in CV_k. Let $SR_{k,n}$ be the storage requirement needed to store videos in the n^{th} element in CV_k. We introduce a binary variable $x_{k,n}$ as follows: if only videos of the n^{th} element in CV_k among videos stored on disk k are cached on SSD, then $x_{k,n} = 1$; otherwise $x_{k,n} = 0$.

Let $C_{i,j}$ denote the j^{th} client requesting video i, $(j = 1, ..., NC_i)$. Let G_k be a set of clients requesting videos stored on disk k, $(k = 1, ..., ND)$. Let $Q_{k,n}$ be a set of clients requesting videos in the n^{th} element in CV_k. We determine disk bandwidth utilization, $D_{k,n}(l)$ of disk k for the speed level of l when $x_{k,n} = 1$ as follows:

$$D_{k,n}(l) = \sum_{C_{i,j} \in G_k - Q_{k,n}} \frac{(T_s + T_d(l) + \frac{b_iR}{r(l)})}{R}.$$

This research was supported by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education, Science and Technology(2009-0065248)

TABLE I
PERFORMANCE AND POWER CONSUMPTION OF EACH RPM LEVEL.

speed level	1	2	3	4	5
RPM	2880	3960	5040	6120	7200
$r(l)$	24.56MB/s	33.77MB/s	42.98MB/s	52.19MB/s	61.4MB/s
$T_d(l)$	5ms	3.6ms	2.86ms	2.35ms	2ms
$P_i(l)$	4.08W	5.11W	6.48W	8.17W	10.2W
$P_a(l)$	7.38W	8.41W	9.78W	11.47W	13.5W
$P_s(l)$	7.38W	8.41W	9.78W	11.47W	13.5W

Let $P_s(l)$ be the power required during the seek phase, while $P_a(l)$ is the power required during the active phase, and $P_d(l)$ is the power consumed during the idle phase when the speed level is l. We calculate the energy consumption during R in each phase as follows:

1) The total seek time is $\sum_{C_{i,j} \in G_k - Q_{k,n}} T_s$. When the speed level is l, the energy required to perform seeks during R denoted by $ES_{k,n}(l)$, is $\sum_{C_{i,j} \in G_k - Q_{k,n}} T_s P_s(l)$.
2) The energy required for reading data during R, denoted by $EA_{k,n}(l)$, is calculated as $\sum_{C_{i,j} \in G_k - Q_{k,n}} \frac{b_i R}{r(l)} P_a(l)$.
3) The energy required when no disk activity is taking place or a disk is waiting for a sector to arrive underneath a head, denoted by $EI_{k,n}(l)$ as follows: $\sum_{C_{i,j} \in G_k - Q_{k,n}} (R - \frac{b_i R}{r(l)}) P_d(l)$.

We therefore determine the power consumption of disk k when speed level is l and $x_{k,n} = 1$, $P_{k,n}(l)$ as follows:

$$P_{k,n}(l) = \frac{ES_{k,n}(l) + EA_{k,n}(l) + EI_{k,n}(l)}{R}.$$

Disk bandwidth utilization must be smaller than or equal to 1; based on this, we define a set of feasible speed levels of disk k when $x_{k,n} = 1$, $W_{k,n} = \{l \mid D_{k,n}(l) \leq 1 - \alpha\}$, where α represents additional disk bandwidth reserved to prepare for the overload condition. To reduce the energy consumption, the lowest speed level in $W_{k,n}$ needs to be selected; we refer to this speed level as $L_{k,n}$. A total size of the cached videos must not exceed the SSD size, TS. We can then formulate the cache selection problem ($\mathcal{CSP}$) that determines $x_{k,n}$ as follows:

$$\text{Minimize} \quad \sum_{k=1}^{ND} \sum_{n=1}^{NE_k} x_{k,n} P_{k,n}(L_{k,n})$$

$$\begin{aligned} \text{subject to} \quad & \sum_{k=1}^{ND} \sum_{n=1}^{NE_k} x_{k,n} SR_{k,n} \leq TS, \\ & \sum_{n=1}^{NE_k} x_{k,n} = 1, (k = 1, ..., ND), \\ & x_{k,n} \in 0, 1. \end{aligned}$$

III. EXPERIMENTAL RESULTS

We evaluate the effectiveness of our scheme through simulations of an array of 20 five-speed disks in Table I. The arrival of client requests follow a poisson distribution where the inter-arrival time is 4 seconds, and the access probability follows a Zipf distribution with parameter $\theta = 0.271$. We considered an SSD with 240 MB/s, which is measured for real sequential workloads [6]. The SSD consumes 1.3W in active mode and 0.7W in idle mode [8]. We consider a workload consisting of 800 video files that support 9Mb/s. α is assumed to be 0.1, and the round length is 2 seconds. To find the optimal solution to $\mathcal{CSP}$, we ran an lp_solve program [9]. We profile the energy consumption of the disk array over 2 hours.

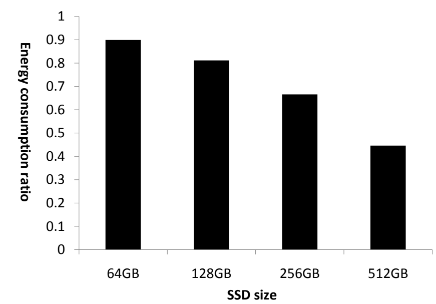

Fig. 1. Energy consumption ratio against the SSD size.

We assess how energy consumption depends on the size of the SSD, and Fig. 1 shows the energy consumption ratio relative to the case where the SSD cache is not used. From the figure, we observe that increasing the cache size gives disks more opportunities to stay at lower speeds, which decreases disk energy consumption.

IV. CONCLUSIONS

We have formulated an integer linear problem that determines videos cached on the SSD to minimize disk energy consumption. Simulations show that using a 256GB SSD can save up to 33% disk energy consumption.

As our future works, we are going to handle several important issues such as SSD bandwidth limitation and replacement overheads. We also plan to apply our scheme to real video servers.

REFERENCES

[1] S. Gurumurthi., "Power management of enterprise storage systems," Ph.D. dissertation, Pennsylvania State University, 2005.
[2] Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton, and J. Wilkes., "Hibernator: helping disk arrays sleep through the winter," *ACM Operating Systems Review*, vol. 39, no. 5, pp. 177–190, 2005.
[3] "Hitachi power and acoustic management - quietly cool," *Hitachi White Paper*, 2004.
[4] H. Yada, H. Ishoioka, T. Yamakoshi, Y. Onuki, Y. Shimano, M. Uchida, H. Kanno, and N. Hayashi., "Head positioning servo and data channel for hdds with multiple spindle speeds," *IEEE Transactions on Magnetics*, vol. 36, no. 5, pp. 2213–2215, 2000.
[5] "http://www.nexsan.com/sataboy/automaid.php."
[6] F. Chen, D. Koufaty, and X. Zhang., "Understanding intrinsic characteristics and systems implications of flash memory based solid state drives." in *Proceedings of the ACM SIGMETRICS*, 2009, pp. 181–192.
[7] E. Chang., "Storage and retrieval of compressed video," Ph.D. dissertation, University of California at Berkeley, 1996.
[8] "Samsung green ssd," *Samsung White Paper*, 2010.
[9] "http://lpsolve.sourceforge.net/5.5/."