플래시 메모리를 사용한 쓰기 캐시 정책 연구

김영진*, 알드히노*, 이정배*, 임기욱*
*선문대학교 컴퓨터공학부
e-mail: youngkim@sunmoon.ac.kr, nindja@gmail.com, {jblee, rim}@sunmoon.ac.kr

A Study on Write Cache Policy using a Flash Memory

Young-Jin Kim*, Aldhino Anggorosesar*, Jeong-Bae Lee*, Kee-Wook Rim*
*Dept. of Computer Science & Engineering, SunMoon University

Abstract

In this paper, we study a pattern-aware write cache policy using a NAND flash memory in disk-based mobile storage systems. Our work is designed to face a mix of a number of sequential accesses and fewer non-sequential ones in mobile storage systems by redirecting the latter to a NAND flash memory and the former to a disk. Experimental results show that our policy improves the overall I/O performance by reducing the overhead significantly from a non-volatile cache over a traditional one.

1. Introduction

Hard disks have poor I/O performance for random I/O requests due to a seek time depending on the position variation of a disk head between successive I/O requests. Thus, there have existed a number of works for improvement of the I/O performance in disk-based storage systems. The representative is using a cache beside a disk to keep the requested data with the expectation of their being accessed again in the near future.

Extensive studies on read caches are reported in the literature to achieve significant I/O performance enhancement. However, studies of write caches have been shown relatively less.

In this paper, we propose a simple, but practical write cache policy for a NAND flash memory in disk-based mobile storage systems. The proposed policy redirects write requests to a non-volatile cache of a NAND flash memory or a disk according to access patterns in order to optimize the overall I/O performance in a disk-based mobile storage system.

2. Motivation

A NAND flash memory has many advantages over hard disks such as fast I/O access time, lower-power consumption, and higher shock resistance although its cost per unit capacity is still high [1]. Thus, using a NAND flash memory as a write cache can be attractive because its non-volatility is useful in write caching and the cost may be acceptable if the size of a NAND flash write cache is properly small over a disk.

Gill *et al.* proposed a performance-enhancing technique based on wise ordering for writes called WOW. WOW is a write request ordering technique proposed for a non-volatile cache (hereafter, NVC) [2]. However, WOW doesn't consider the real I/O latency of an NVC itself since it assumes a volatile DDR memory to be an NVC.

3. Proposed Policy

We propose a simple, but practical pattern-aware write cache policy, which considers the access patterns of the load as well as the I/O features of a NAND flash memory in a

disk-based mobile storage system to optimize the overall I/O performance.

If a write access is random, it will be inserted into the MRU location of an NVC. Otherwise, the write will be forwarded to a disk. Therefore, redirecting data requests toward the NVC is likely to absorb a large number of I/O accesses, which otherwise will be directed to the disk, and powers down the disk, which can obtain performance improvement as well as energy saving concurrently. When a dirty block is evicted from the NVC, it will be flushed to the disk. Otherwise, it will be discarded.

4. Experiments and Results

Figure 1 shows the average I/O response times of LRU and our proposed policy for a generic workload trace [3] obtained in Windows XP when the size of an NVC varies. We can notice that the proposed policy has the faster I/O response time as the cache size grows than LRU. The proposed pattern-aware policy is shown to have shorter response time than LRU by up to 52%. Figure 2 shows that the flash write count of the proposed policy is quite smaller than that of LRU. This is because the flash write count is reduced by the proposed policy since it redirects the random writes to the NVC instead of the disk.

5. Conclusions

We proposed a pattern-aware write cache policy using a NAND flash memory in disk-based mobile storage systems. Our work aims at facing a mix of a number of sequential accesses and fewer non-sequential ones in mobile storage systems by redirecting the latter to a NAND flash memory and the former to a disk. Experimental results show that our policy improves the overall I/O performance by reducing the overhead significantly from a non-volatile cache over a traditional one. As future work, we will investigate how queuing method can have synergistic influence in combination with the proposed policy.

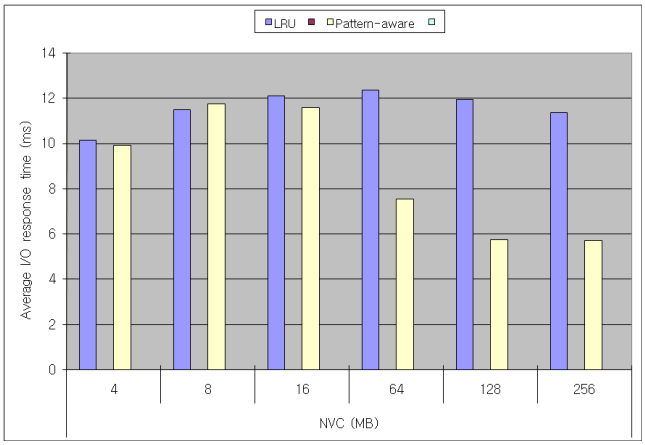

(Figure 1) Average I/O response times

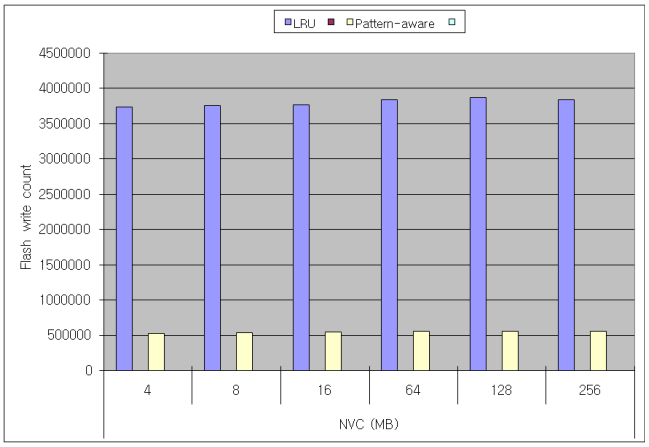

(Figure 2) Flash write counts

Acknowledgment

This research was supported by the MKE (Ministry of Knowledge Economy), Korea, under the ITRC (Information Technology Research Center) Support program supervised by the IITA (Institute of Information Technology Advancement) (IITA-2008-C1090-0801-0020).

References

[1] G. Lawton, "Improved flash memory grows in popularity," IEEE Computer, vol. 39, no. 1, pp. 16-18, Jan. 2006.

[2] B. Gill and D. Modha, "WOW: Wise ordering for writes-combining spatial and temporal locality in non-volatile caches," Proc. the 4th USENIX Conference on FAST, San Francisco, 2005.

[3] Y.-J. Kim, S.-J. Lee, K. Zhang, and J. Kim, "I/O performance optimization technique for hybrid hard disk-based mobile consumer devices," IEEE Transactions on Consumer Electronics, vol. 53, issue 4, pp. 1469-1476, Nov. 2007.