Three-Bit Hot/Cold Page Clustering Mechanism for Improving Endurance of Flash Memory

Taehwa Lee[1], Jaehyuk Cha[2], Youjip Won[2], Suyong Kang[2]
Jongmoo Choi[3], Sungroh Yoon[4]

{[1]Department of Electronics and Computer Engineering,
[2]Department of Computer Science and Engineering}, Hanyang University, Seoul 133-791, Korea
[3]Department of Computer Science and Engineering, Dankook University, Yongin, Gyeonggi 448-701, Korea
[4]Department of Electrical and Computer Engineering, Seoul National University, Seoul 151-744, Korea
{alghost, chajh, yjwon, sykang}@hanyang.com
choijm@dankook.ac.kr, sryoon@snu.ac.kr

Abstract. In order to reuse a specific page, as in the case of a rewrite operation, a flash memory has to erase the block in which the page is included; this is a drawback of the flash memory. A block, which is a unit of the erase operation for the flash memory, has limited number of erase operations permitted, and if this count exceeds the threshold, the block can become a bad block and cannot be used in any future operations. Further, if SSD(Solid State Drive) exceeds a certain number of bad blocks, it cannot be used. SSD, which is a flash-memory-based storage device, has FTL to overcome the above, just as other block devices such as HDD do. In this paper, three-bit hot/cold page clustering and victim block selection based on MIN-MAX GAP are proposed to extend the endurance of SSD by improving the main policies of FTL, such as victim block selection and block allocation. The proposed method decreases the erase count and improves the wear leveling. The performance evaluation reveals that the proposed method has comparable result with the existing methods and shows approximately 77% improvement in terms of the wear leveling.

Keywords: Flash Memory, Wear Leveling, Garbage Collection, Endurance

1 Introduction

Recently, a considerable amount of attention has been paid to NAND-flash-memory-based SSD(Solid State Drive) as an alternative to the traditional HDD(Hard Disk Drive) storage. This can be attributed to the strengths of SSD, such as fast speed, low power consumption, and durability. However, in a flash memory, an internal storage element of SSD, a read/write operation is performed on the basis of a page unit, and the overwrite operation is not possible, thereby requiring the erase operation performed on the basis of a block unit, which is a set of pages. Further, each block has a limited number of erase operations permitted, and the SSD cannot be used if the

number of bad blocks, which are blocks that have exceeded the limited number of erase operations, is more than a certain percentage of the total number of blocks. SSD has a s/w module called the flash translation layer (FTL), which complements the physical characteristics of the flash memory; this module can be used as a block device such as HDD. The garbage collection, which is one of main features of FTL, erases blocks. Wear leveling is a method to extend the endurance of SSD by leveling the erase count of each block evenly. In this paper, a mechanism that introduces two policies into a general page mapping technique is proposed. First, three-bit hot/cold page clustering is proposed to have many invalid pages in order to find a victim block. Second, a block consisting of valid pages cannot become a victim block in order to ensure that the erase count does not increase. In order to utilize blocks uniformly, victim block selection based on MIN-MAX GAP is proposed to disperse the pages in the abovementioned block to the other blocks. The proposed mechanism is successful in decreasing the erase count and improving the wear leveling.

The rest of this paper is organized as follows: In Section 2, a three-bit hot/cold page clustering is proposed. In Section 3, victim block selection based on MIN-MAX GAP is proposed. In Section 4, evaluation experiments of a comparison between the proposed method, one-bit clustering such as FeGC[1] and Dual-Pool[2], and a method without clustering are discussed. The results of the proposed method reveal comparable performance with respect to the erase count and approximately 77% improvement in terms of the wear leveling. We have considered the total erase count, and the dispersion of the block erase count of each block, which verifies wear leveling, as the performance evaluation indices for checking the endurance expectancy. Finally, the conclusion is presented in Section 5.

2 Three-Bit Hot/Cold Page Clustering

When receiving a write request from a host, FTL attempts to allocate free pages. If free pages cannot be allocated, a new free block is allocated. It is required to have a policy that all the pages in one block become invalid at the same time. In general, pages have different frequencies of use. Thus, if pages that have a similar frequency of use are gathered in one block with the clustering technique, the valid pages in the same block are more likely to become invalid at the same time because of the pages that are reused in a similar period. Therefore, when a victim block is selected during garbage collection, blocks that have relatively few valid pages are present, thereby reducing the degradation of the product endurance because of relatively few copy operations of the valid pages. To this end, in this study, three-bit state information of a page was used for defining HOT for the most frequently referred-to pages, and COLD for the least referred to pages, and the other six states were categorized into states between HOT and COLD according to their frequency of reference.

When free pages corresponding to a certain state are all filled in a block, a new free block should be allocated. The past state of the requested page is utilized at this time. For a page that is closer to the HOT state, which indicates a high erase count, a block that has a low erase count is allocated. Through a page's state definition and clustering of blocks by state, the erase count can be reduced by a reduction of the

copies of the valid pages during garbage collection. Further, wear leveling can be achieved by the block allocation policy using the page's state.

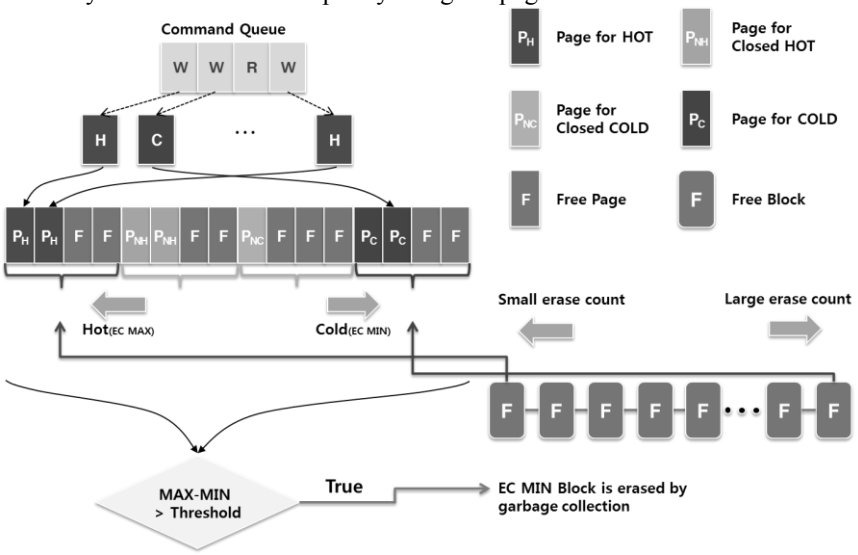

Fig. 1 Main features of FTL used in three-bit hot/cold page clustering and victim block selection based on MIN-MAX GAP to ensure wear leveling

For the victim block selection policy required for garbage collection, the widely used greedy technique was improved and used. The greedy technique is a method to minimize the valid page copies by selecting a block that has the least number of valid pages within the block. However, if many COLD pages are present in a block, it is difficult for the block to become a victim block because there is no change caused by the greedy technique. In this study, when pages' states of a victim block candidate are COLD, a weight is assigned so that a block that has many COLD pages has more chances of selection as a victim block.

3 Victim Block Selection based on MIN-MAX GAP

Although during three-bit hot/cold page clustering, we can expect wear leveling in general, we cannot guarantee it depending on the use pattern. To guarantee this, a policy of victim block selection based on MIN-MAX GAP was designed. In this policy, when the difference between the maximum and the minimum values of the block erase count for each block exceeds the threshold, a block that has the minimum erase count is forced to become a victim block during garbage collection. The COLD pages in the victim block are allocated to a block that has a high erase count, and HOT pages are allocated to the victim block because of the low erase count, thereby achieving wear leveling ultimately. This can ensure wear leveling because the threshold for the maximum and minimum erase counts was fixed.

4 Performance Evaluation

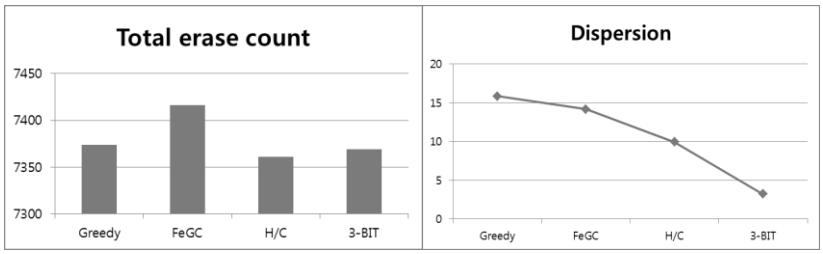

Fig. 2 Comparison of total erase count and dispersion between the algorithms

The trace used for this experiment was trace extracted by running TPC-B, which was performed in PostgreSQL 9 installed in Solaris 10. For the flash memory used for the simulation, the physical page size is 8k, the number of pages per block is 128, and the total number of blocks was set as 512, which are the same conditions as those for a virtualized 512MB flash memory. This experiment compared FTL (greedy algorithm) to which the wear-leveling technique is not applied, and the previously proposed FeGC and H/C. H/C is an algorithm for clustering as two states using one-bit-like Dual-Pool. The experiment result showed that the mechanism proposed in this paper had comparable performance in terms of the erase count but approximately 77% improvement in terms of the dispersion.

5 Conclusion

In this paper, three-bit hot/cold page clustering and victim block selection based on MIN-MAX GAP were proposed in order to improve the endurance of flash-memory-based SSD. According to the experimental result, the proposed mechanism had almost the same erase count as the greedy algorithm and the dispersion of erase counts was improved by approximately 77%.

Acknowledgment: This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MEST) (No. 2010-0013076) and sponsored by IT R&D program MKE/KEIT [No. 10035202, Large-Scale Hyper-MLC SSD Technology Development].

References

1. O. Kwon, K. Koh, J. Lee, H. Bahn, “FeGC: An efficient garbage collection scheme for flash memory based storage systems,” The Journal of Systems and Software, vol.84, pp.1507-1523, 2011.
2. L.-P. Chang, “On efficient wear-leveling for large-scale flash-memory storage systems,” Proc. of the 2007 ACM Symposium on Applied Computing, pp.1126-1130, 2007. Checklist of Items to be Sent to Volume Editors