

Flash Translation layer used for mapping schemes like BAST and FAST

Gayke Pratibha S.

Assistant Professor ,

Department of Information Technology,

Dr.Vithalrao Vikhe Patil College of Engineering, Ahmednagar (MS), India

Abstract : Flash is a type of electronically erasable programmable read- only memory (EEPROM). NAND Flash Memory is used in handheld electronic devices like mobile, cameras, iPods, music players, is also used as an alternative storage medium for Hard Disk Drives (HDD) in PCs and Laptops. Flash memory is important as nonvolatile storage for mobile consumer electronics due to its low power consumption and shock resistance. An intermediate software layer called Flash Translation Layer (FTL) is used to overcome these obstacles.

Many efforts for optimizing the working of address mapping schemes have been done by different research workers. Though various schemes are designed and proposed but there is no literature available providing mathematical computations comparing the performance of the various mapping schemes in the form of time complexity. In this paper we have tried to find out the comparative cost of block merge operation required during garbage collection for some representative mapping schemes like BAST and FAST .

Keywords: Flash Translation layer, Mapping, BAST and FAST

1. Introduction:

Flash memory which is an electronic non volatile memory can be electrically erased and reprogrammed. It was developed from EEPROM (electrically erasable programmable read-only memory). While a flash device can read any of its pages, it may only write to one that is in a special state called erased. This is called as 'erase before update' characteristics of Flash. Due to which flash do not support overwrite i.e. in-place update operation, instead it supports out of-place-update. Flash Translation Layer maintains a mapping table of virtual address to physical address.

1.1 Characteristics of NAND flash memory

1) The previous data should be erased before a new data can be written in the same place. This is usually called erase- before-write characteristic.

2) Normal read/write operations are performed on a per- *page* basis, while erase operations on a per- *block* basis. The erase block size is larger than the page size by 64~128 times. In MLC (Multi-Level Cell) NAND flash memory, the typical page size is 4KB and each block consists of 128 pages.

3) Flash memory has limited lifetime; MLC NAND flash memory wears out after 5K to 10K write/erase cycles. There are three basic operations in NAND flash memory: read, write (or program), and erase. The read operation fetches data from a target page, while the write operation writes data to a page. The erase operation resets all values of a target block to 1. NAND flash memory does not support in-place update.

FTLs can be categorized into two classes according to their mapping granularities: Sector-mapped FTLs and block-mapped FTLs. A Sector-mapped FTL literally maps a logical address into a physical address in a page unit. It is highly flexible as a logical page can be written to any physical page in NAND flash memory. On the other hand, the mapping unit of a block-mapped FTL is a block. And later on a hybrid mapping scheme proposed by D. Park came into existence.

A NAND flash memory chip is composed of a fixed number of *blocks*, where each block typically has 32 *pages*. Each page in turn consists of 512 bytes of the main data area and 16 bytes of the spare area. The page is the basic unit of read and write operations in NAND flash memory.

1.2 REVIEW OF PREVIOUS FTL SCHEMES

A. Page Level Mapping:

In page mapping scheme the smallest logical unit that FTL uses for address translation is a page. A page is made up of certain number of smallest units called as sector. This is most efficient scheme in that a logical page can be mapped into any physical page in flash memory. This mapping information is kept in the form of tables called as mapping tables.

B. Block- mapping FTL

The pure block-mapping FTL is another classic FTL scheme. Block-mapping table is used to store and manage the mapping information between LBN and Physical Block Number (PBN). If there are m pages in a block, the size of the block-mapping table is m times smaller than its page-mapping counterpart. In a block-mapping FTL, one LPN must be mapped to a fixed page offset in any physical block (i.e., direct mapping). If this page offset has been written before, the LPN cannot be written to any other page in this block even if there are free pages in the same physical block. In this case, all existing valid data in the block as well as the data to be written must be copied to a new clean block, and the old block is marked for erase, incurring one erase and a number of read/write operations. Compared with the page-mapping FTL the block-mapping FTL requires extra operations to serve a request, adversely affecting the performance. Since both the block-mapping and page-mapping FTLs have their aforementioned disadvantages, they are rarely used in SSD commercial products in their pure forms.

C. Hybrid Level Mapping

A family of the hybrid mapping schemes is introduced to address the shortcomings of the Sector-mapping and block-mapping FTLs. In a typical hybrid FTL, physical blocks are logically partitioned into two groups: data blocks and log blocks. When a write request arrives, the hybrid FTL first writes the new data in a log block and invalidates the data in the corresponding target data block.

Block-mapping information for data blocks and page-mapping information for log blocks are kept in a small RAM for performance purposes. When all the

log blocks are full, their data are flushed into the data blocks immediately and they are then erased to generate new free log blocks. More specifically, the valid data in data blocks and the valid data in the corresponding log block must be merged and written to a new clean data block. This process is called a merge operation. Further, merge operations can be classified into three types depending on their overhead. Full merge occurs, when the log block is selected as a victim block and not written sequentially from the first page to the last page, and all the valid data in it and in its corresponding data block are copied to a new clean block.

This process requires m read operations, m write operations and two erase operations, where m is the number of pages in a block. When the log block is written sequentially from the first page to the last page of a logical block, this log block can replace the corresponding data block, a merge operation called switch merge. This type of merge requires only one erase operation. Partial merge takes place when the log block is written sequentially from the first page to a middle page in a block, and the last part of data will be copied from the corresponding data block. Partial merge requires several read and write operations and one erase operation. A number of variations of the hybrid FTL schemes have been proposed recently, including BAST, FAST, LAST, Superblock Reconfigurable FTL. More recently, Demand-based FTL (DFTL) was proposed to address the RAM-capacity problem of the page-mapping FTL by storing only the “hot” mapping information in RAM based on temporal locality of workloads.

DFTL is shown to significantly outperform hybrid FTLs. A hybrid technique, as its name suggests, first uses a block mapping technique to get the corresponding physical block, and then, uses a sector mapping technique to find an available empty sector within the physical block.

1.3 COMPARATIVE SCHEME

A. Block Associative Sector Translation (BAST) exclusively associates a log block with a data block. In presence of small random writes, this scheme suffers from **log block thrashing** that results in increased full merge cost due to inefficiently utilized log blocks.

Table 1. Measures of BAST scheme

Garbage collection cost (worst and best case) : considering number of random requests = K* (number of log Blocks)	$((3T/100)*K*N)$ Read + $((3T/100) * K*N)$ Write + $(2*3T/100)$ Erase T: total blocks in flash N: number of ages/block Log block: 3% of T
RAM requirement	Less, Proportional to number of log blocks
Search time (worst case)	Time to search the page map table of a single log blocks.
Usefulness	In case of sequential read write and update pattern

B. Super Block FTL scheme utilizes existence of *block level* spatial locality in workloads by combining consecutive logical blocks into a super block. It maintains page- level mappings within the superblock to exploit temporal locality in the request streams by separating hot and cold data within the superblock. However, the three-level address translation mechanism employed by this scheme causes multiple OOB area reads and writes for servicing the requests. More importantly, it utilizes a fixed superblock size which needs to be explicitly tuned to adapt to changing workload requirements.

C. Fully Associative Sector Translation (FAST) allows log blocks to be shared by all data blocks. This improves the utilization of log blocks as compared to BAST. FAST keeps a single sequential log block dedicated for sequential updates while other log blocks are used for performing random writes. Thus, it cannot accommodate multiple sequential streams and does not provide any special mechanism to handle temporal locality in random streams.

Table 2. Measures of FAST mapping scheme

Garbage collection cost (worst and best case) : considering number of random requests = K* (number of log Blocks)	$((3T/100)*K*N)$ Read + $((3T/100) * K*N)$ Write + $(2*3T/100)$ Erase T: total blocks in flash N: number of ages/block Log block: 3% of T
RAM requirement	Less, Proportional to number of random log blocks
Search time (worst case)	Time to search the page map table of all log blocks.
Usefulness	

D. Locality-Aware Sector Translation (LAST)

Scheme tries to alleviate the shortcomings of FAST by

providing multiple sequential log blocks to exploit spatial locality in workloads. It further separates random log blocks into hot and cold regions to reduce full merge cost. In order to provide this dynamic separation, LAST depends on an external locality detection mechanism. However, Lee et al. themselves realize that the proposed locality detector cannot efficiently identify sequential writes when the small-sized write has sequential locality. Moreover, maintaining sequential log blocks using a block-based mapping table requires the sequential streams to be aligned with the starting page offset of the log block in order to perform switch-merge. Dynamically changing request streams may impose severe restrictions on the utility of this scheme to efficiently adapt to the workload patterns.

1.4 DFTL Architecture

DFTL makes use of the presence of temporal locality in workloads to judiciously utilize the small on-flash SRAM. Instead of the traditional approach of storing all the address translation entries in the SRAM, it dynamically loads and unloads the page-level mappings depending on the work- load access patterns. Furthermore, it maintains the complete image of the page-based mapping table on the flash device itself. There are two options for storing the image: (i) The OOB area or (ii) the data area of the physical pages. We choose to store the mappings in the data area instead of OOB area because it enables us to group a larger number of mappings into a single page as compared to storing in the OOB area. For example, if 4 Bytes are needed to represent the physical page address in flash, then we can group 512 logically consecutive mappings in the data area of a single page whereas only 16 such mappings would fit an OOB area. Moreover, the additional space overhead incurred is negligible as compared to the total flash size. A 1GB flash device re- quires only about 2MB (approximately 0.2% of 1GB) space for storing all the mappings.

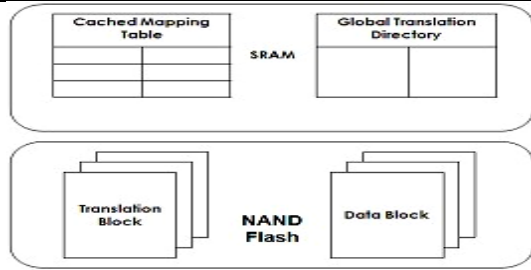


Figure 5 DFTL Architecture

1.5 Comparison of Existing State-of-the-art FTLs with DFTL

A. Full Merge - Existing hybrid FTL schemes try to reduce the number of full merge operations to improve their performance. DFTL, on the other hand, completely does away with full merges. This is made possible by page-level mappings which enable relocation of any logical page to any physical page on flash while other hybrid FTLs have to merge page-mapped log blocks with block-mapped data blocks.

B. Partial Merge - DFTL utilizes page-level temporal locality to store pages which are accessed together within same physical blocks. This implicitly separates hot and cold blocks as compared to LAST and Superblock schemes which require special external mechanisms to achieve the segregation. Thus, DFTL adapts more efficiently to changing workload environment as compared with existing hybrid FTL schemes.

C. Random Write Performance - As is clearly evident, it is not necessarily the random writes which cause poor flash device performance but the intrinsic shortcomings in the design of hybrid FTLs which cause costly merges (full) on log blocks during garbage collection. Since DFTL does not require these expensive full-merges, it is able to improve random write performance.

D. Block Utilization - In hybrid FTLs, only log blocks are available for servicing update requests. This can lead to low block utilization for workloads whose working-set size is smaller than the flash size. Many data blocks will remain un-utilized (hybrid FTLs have block-based mappings for data blocks) and unnecessary garbage collection will be performed. DFTL solves this problem since updates can be performed on any of the DATA blocks.

Table 3. Comparative analysis of various FTL schemes

FTL Scheme	Merge cost of block during GC	Lookup Performance	RAM Req.	Mapping granularity
Pure Page level	N/A	N lookup c	Much more	Page
Pure block level	Much More	N lookup c	Very less	Block
BAST	More	Less	less	Page for log Blocks
FAST	Less than BAST	Much more	Less	Page for log Blocks Block for data blocks
LAST	Less than FAST	Much more	Less	Page & blocks for log blocks. Block for data blocks
Demand Paged	N/A	Lesser	Lesser	Page

1.6 Conclusion:

From the response time got from the various workload it is observed that the proposed innovative algorithm outperforms the FAST scheme. FAST is taken as a representative scheme for comparison as it has been considered an optimal one among the various FTL schemes that are available. We proposed a complete paradigm shift in the design of the FTL with our Demand-based Flash Translation Layer (DFTL) that selectively caches page-level address mappings. Our experimental evaluation using Disk Sim with realistic enterprise-scale workloads endorsed DFTL's efficacy for enterprise systems by demonstrating that DFTL offered (i) Improved performance, (ii) reduced garbage collection overhead (iii) Improved overload behavior and (iv) Most importantly unlike existing hybrid FTLs is free from any tunable parameters.

REFERENCES

1. Liu D., Wang Y., Qin Z., Shao Z., Guan Y.: A Space Reuse Strategy for Flash Translation Layers in SLC NAND Flash Memory Storage Systems, IEEE

- Transactions on Very Large scale Integration (VLSI) Systems, pages 1-14, May – 2011, ISSN: 1063- 8210 Volume: PP Issue:99
2. Shin I.: Light weight sector mapping scheme for NAND-based block devices, IEEE Transactions on Consumer Electronics, pages: 651 – 656, May 2010, ISSN: 0098-3063 Volume: 56 Issue:2
 3. “Flash Sim: A Simulator for NAND Flash-Based Solid-State Drives” in Advances in System Simulation, 2009. SIMUL '09. First International Conference on pages 125-131 ISBN: 978-1-4244- 4863-0
 4. Aayush Gupta Youngjae Kim Bhuvan Urngaonkar “DFTL: A Flash Translation Layer Employing Demand-based Selective Caching of Page- level Address Mappings” Computer Systems Laboratory, department of Computer Science & Engineering. The Pennsylvania State University, Univesity Park, PA 16802, Technical Report CSE-08-012 August 2008.
 5. Dongchul Park, Biplob Debnath, and David Du “CFTL: A Convertible Flash Translation Layer with Consideration of Data Access Patterns”, Technical Report Department of Computer Science and Engineering University of Minnesota September 14, 2009.
 6. S. Lee, D. Shin, Y. Kim, and J. Kim. LAST: Locality- Aware Sector Translation for NAND Flash Memory- Based Storage Systems, in Proceedings of the International Workshop on Storage and I/O Virtualization, Performance, Energy, Evaluation and Dependability (SPEED2008), February 2008.
 7. Chung, D. Park, S . Park, D. Lee, S. Lee, and H. Song. System Software for Flash Memory: A Survey. In Proceedings of the International Conference on Embedded and Ubiquitous Computing, pages 394–404, August 2006.
 8. J. Kang, H. Jo, J. Kim, and J. Lee. A Superblock-based Flash Translation Layer for NAND Flash Memory. In Proceedings of the International Conference on Embedded Software (EM-SOFT) , pages 161–170, October 2006. ISBN 1-59593-542-8.
 9. S. Lee, D. Park, T. Chung, D. Lee, S . Park, and H. Song. A Log Buffer based Flash Translation Layer Using Fully Associative Sector Translation. IEEE Transactions on Embedded Computing Systems, 6(3):18, 2007. ISSN 1539–9087.
 10. T.-S. Chung, D.-J. Park, S. Park, D.-H. Lee, S.-W. Lee, and H.-J. Song, “A survey of Flash Translation Layer,” J. Syst. Archit., vol. 55, no. 5-6, 2009.
 11. Shinde /Gayke Pratibha, Mrs. Suvarna “Efficient Flash Translation layer for Flash Memory” International Journal of Scientific and Research Publications, Volume 3, Issue 4, April 2013 ISSN 2250-3153