Microprocessors and Microsystems 37 (2013) 345–357

Contents lists available at SciVerse ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

A flash-aware write buffer scheme to enhance the performance of superblock-based NAND flash storage systems

Ning Lu, In-Sung Choi, Shin-Dug Kim *

Department of Computer Science, Yonsei University, 134 Shinchon-Dong Sudaemoon-Ku, Seoul 120-749, Republic of Korea

ARTICLE INFO

Article history:
Available online 25 July 2012

Keywords:
NAND flash
Storage system
Solid state disk
Flash memory
Write buffer

ABSTRACT

Most superblock-based NAND flash storage systems employ a high-speed write buffer to enhance their writing performance. The main objective is to bind data of adjacent addresses as much as possible in order to transform random data into sequential data, which then facilitates interleaving in the storage system. We have designed a new superblock-based buffer scheme for NAND flash storage systems that improves on traditional schemes. For buffer management, a series of lists need to be specified to monitor the data-flow changes in the current state of the buffered data and the NAND flash memory in order to maximize interleaving during the flush operation. Experimental results show that the proposed scheme achieves higher write speed performance in almost all configurations, with greater than 50% speedup in some cases. Our proposed flash-aware write buffer (FAWB) scheme achieves this higher write performance with a required buffer space of only 1/4th–1/8th that of other schemes, resulting in higher efficiency.

1. Introduction

Many mobile devices, such as MP3 players, digital cameras, video recorders, smart phones, and personal digital assistants (PDAs), use NAND flash memory for their storage system due to its characteristics of non-volatility, low power consumption, and high random access speed. Until recently, production was relatively expensive for mass capacity. However, as manufacturing technology has advanced, mass capacity appliances have gradually become popular, and now many mass-storage-demanding handheld devices use NAND flash memory.

Along with the benefits, NAND flash memory has some inherent defects. First, read and write operations must be implemented in page units rather than byte units. Second, the write operation is much slower than the read operation. Third, an erase operation is required before a write operation, instead of using direct overwrite. And finally, the erase operation must be implemented in a block unit, which is considerably larger than a page unit. Some main characteristics of these NAND flash memory specifications have been shown in [1,2]. In order to overcome these drawbacks, the storage system usually employs a high-speed write buffer. This write buffer can be made of relatively expensive non-volatile memories such as PRAM, FRAM, MRAM [1], and power backup DRAM. In traditional schemes, the main objective of the buffer is to bind random data in a group in order to transform it into sequential data, which results in maximum access speed close to that of sequential writing.

In this paper, we propose a new write buffer scheme, called a flash-aware write buffer (FAWB) that enhances the write performance of previous NAND storage systems. In the previous NAND storage systems, the write buffer usually packs the interrelated data merely and ignores the running state and detail of the underlying components, e.g. flash translation layer. Thus, the write buffer is always blind to the underlying components, and eventually the synergy is declined. One significant change from the FAWB scheme is a communication channel to be built between the write buffer and the underlying flash translation layer, and thus the FAWB scheme monitors the current running information of the NAND storage system to manage buffer coordination with the flush operation and obtain optimized writing performance, as shown in performance evaluation. From the overall running information, an optimum scheduling method is determined and used to conduct the buffer flush operation. Thereby all sub-controllers for a group of NAND flash memory are never idle, increasing the interleaving level. Additionally, the least recently used (LRU) method is employed to manage the gathered running information and to avoid frequent merge operations. We showed by simulation that the FAWB achieves higher write performance, and the required buffer space is reduced to only 1/4th to 1/8th that of former schemes.

The remainder of this paper is organized as follows. The next section briefly introduces the basic mechanisms of the mass NAND flash memory storage system. Previous write buffer schemes are compared in Section 3. Section 4 discusses the FAWB scheme and its detailed operation. Evaluation and comparison with previous

* Corresponding author.
E-mail addresses: luning@yonsei.ac.kr (N. Lu), cis53@yonsei.ac.kr (I.-S. Choi), sdkim@yonsei.ac.kr (S.-D. Kim).

0141-9331/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.micpro.2012.07.005

schemes are presented in Section 5. Finally, we draw conclusions in the final section of this paper.

2. Background and related work

In this section, we briefly describe the characteristics of NAND flash memory and its usage in mass storage systems. We also describe a previous superblock management method called a hybrid mapping FTL, a related buffering method known as the flash-aware buffer (FAB) [3], and the cold and largest cluster (CLC) [4].

A NAND flash memory chip is composed of a fixed number of blocks, and each block holds a fixed number of pages, which is typically 32, 64, or 128. Each page has a data area ranging from 512 bytes to 4 Kbytes, as well as a spare area that is typically 16 bytes. The spare area is often used to store management information or error correction codes.

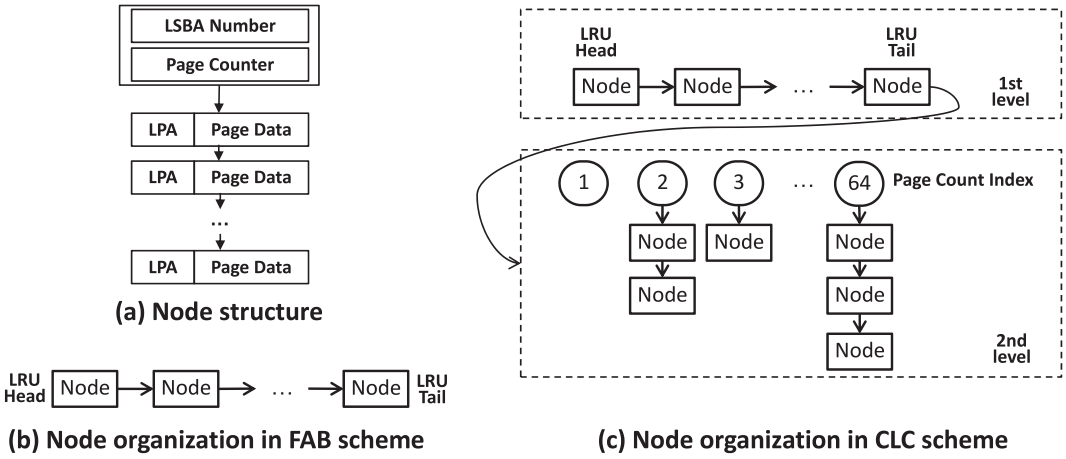

Fig. 1. Buffer technique of the FAB and CLC schemes.

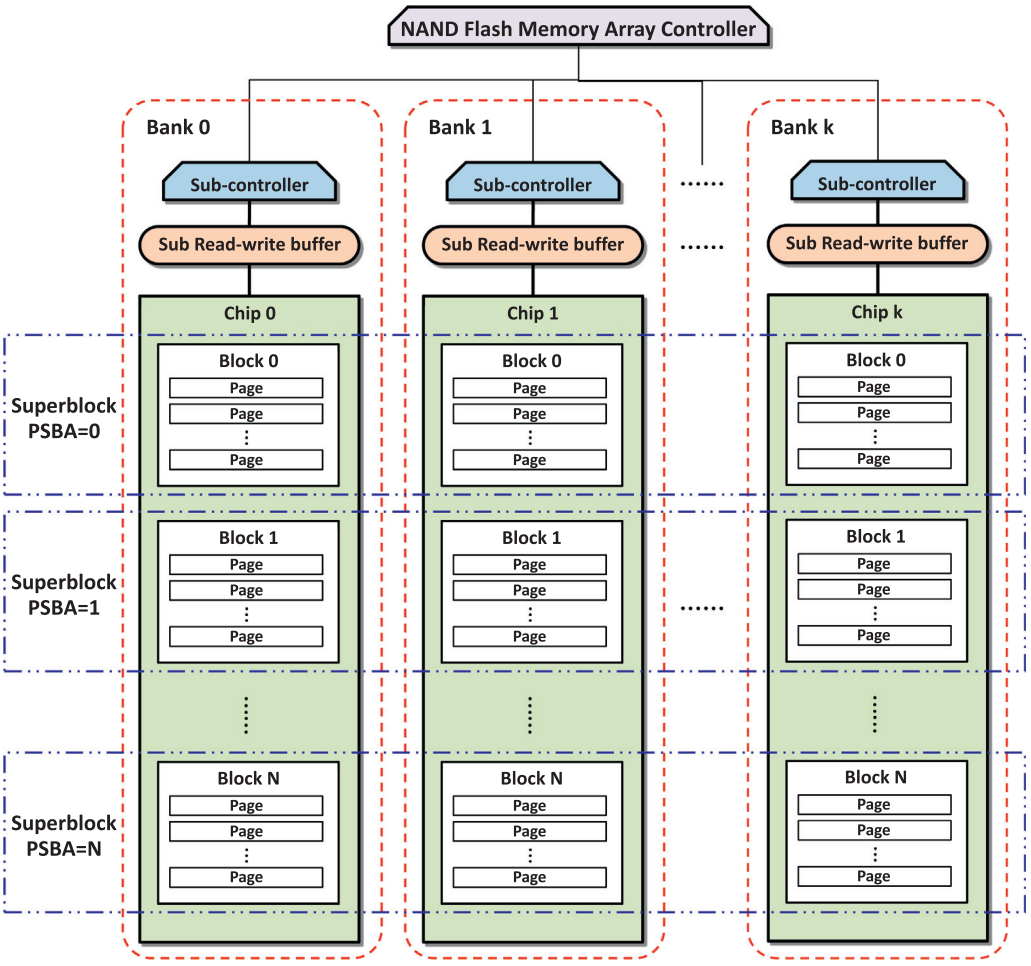

Fig. 2. Superblock-based NAND memory organization.