

NAND Flash Management

White Paper

February 18, 2009

Version 1.2

NAND Flash Management

Overview

Flash memory has many benefits over conventional rotating media including compact size, low power consumption, zero seek time, and better shock resistance. These properties make flash memory storage ideal for the mobile environment including digital cameras, portable devices, and other embedded systems. However, NAND flash has a limited number of write/erase cycles. There are many sophisticated algorithms for evening out the usage of flash blocks across the whole address space (wear leveling) to improve the expected life of a product. In addition, NAND devices contain a certain number of bad blocks and require extensive error correction to be performed on the read data. A host system needs to consider handling and managing bad blocks, not just the ones identified at initialization, but also the ones that go bad with usage. Next, most NAND devices require error-correction code (ECC) for the same data block size. Implementing an error-correction algorithm in firmware will consume lots of precious arithmetical computations, and if the error frequency is great, it will slow the system down significantly, as well as lower the overall performance. Moreover, the feature of power loss recovery demonstrates greatly helpful during sudden power interruptions or brown-out conditions and where the inadvertent removal of a battery could threaten the integrity of data. Overall, how a storage system manages the memory is the key to understanding the extended reliability of the host that relies on these storage systems.

NAND Flash memory

NAND Flash memory allows only two states: erased and non-erased. In the erase state, a byte may be either all ones (0xFF) or all zeroes (0x00) depending on the flash device. A given bit of data may only be written when the media is in an erase state. After being written, the bit is considered unusable. In order to return the bit to its erase state, a significantly larger size of a flash must be erased.

There are two type NAND flash memory technologies available today: Single-Level Cell (SLC) and Multi-Level Cell (MLC). SLC technology stores one bit per cell, and MLC stores two bits per cell. In general, SLC NAND flash memory is specified at 100k endurance cycle, and MLC is specified at 10k endurance cycles.

In fact, NAND flash memory requires accesses in units of pages and blocks. Pages are the smallest unit for read and write operations of flash memory. Typical page size is 2K or 4K bytes depending on flash memory. Each page will store 4 or 8 sectors of 512 bytes each depending on page size. There are, typically, 16 additional bytes associated with each sector that may be used to store ECC codes and wear leveling data structures. In contrast, a block is the smallest unit for erasure of flash memory. The number of pages in a block depends on the specific flash memory. In general, SLC flash memory has 64 pages in a block but MLC flash memory has 128 pages.




Figure 1: NAND Flash chip (2K page) MLC Architecture

The Need for Wear-Leveling

All types of flash memory technology are subject to write endurance limitations. After repeated program and erase cycles, some flash memory bits will no longer retain data. All flash memory manufacturers provide a specification with their devices indicating how many write cycles it can endure before anticipated failure. To compensate for endurance limitation and maximize the lifecycle of the flash memory storage devices, the latest technologies in flash memory management incorporates global wear-leveling algorithms. Global Wear-leveling algorithms are transparent to the overlying files system by keeping track of a map between the logical blocks and the physical blocks that map onto the flash memory and automatically even out flash usage across the whole address space by directing the controller to rotate memory writes to blocks that have lowest usage. In the ideal case, a wear-leveling algorithm will result in all the sectors of the flash memory reaching their endurance limit approximately the same time. Thus the flash memory life is maximized extensively with expectancy beyond the endurance specification.

The Global Wear Leveling Algorithm

The global wear leveling algorithm evens out the “wear” of sectors by writing to a new physical sector every time a logical sector is written to prevent any frequently updated data from staying at the static area so that wear leveling could be evenly applied to all blocks. Static areas contain any data that does not change, and are ignored by dynamic wear leveling. Such static data may include operating system files, table look-ups, executable files, and etc. Global wear leveling frequently replaces blocks in this area with block in the hot area, and thus each block in all areas has the same probability to be used.




Figure 2: Wear Leveling for Flash Memory

Revision History

Revision	Date	Description	Remark
1.0	April 30, 2008	Official release	
1.1	January 20, 2009	Context revised	
1.2	February 18, 2009	Passages trimmed	

Apacer Technology Inc.

9/F, No. 100, Hsin Tai Wu Rd.
Hsichih, Taipei County 221, Taiwan
Tel: +886-2-2696-1666 Fax: +886-2-2696-1668
www.apacer.com

Copyright © 2009 Apacer Technology Inc. All Rights Reserved.
Information in this document is subject to change without prior notice.
Apacer and the Apacer logo are trademarks or registered trademarks of Apacer Technology Inc.
Other brands, names, trademarks or registered trademarks may be claimed as the property of
their respective owners.