Wear Leveling

White Paper for UDM

June 2, 2009

Version 1.0

Apacer Technology Inc.

4th Fl., 75 Xintai 5th Rd., Sec.1, Hsichih, Taipei Hsien 221, Taiwan

Tel: +886-2-2698-2888 Fax: +886-2-2698-2889

Abstract

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. 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.

The Wear Leveling Algorithm

A wear leveling algorithm manages the uneven “wear” on the sectors of a flash media memory by distributing the writes through many sectors of the flash media. Integrated in to the firmware of the flash disk controller, the wear leveling algorithm is transparent to the overlying file system by keeping track of a map between the logical sectors and the physical sectors that mp on the flash media. In the ideal case, a wear-leveling algorithm will result in all the sectors of the flash media reaching their endurance limit nearly simultaneously, maximizing the usable lifetime of the flash media. With the use of aging mechanisms, it is possible to warn the user when endurance limits are reached and to proactively backup the contents before it is lost for good.

Wear leveling allows erase counts of blocks to be evenly distributed over the storage media so that write/erase endurance for the entire flash media could be increased. Wear leveling is an algorithm by which the controller recycles blocks with small erase counts in the flash array. Whenever a sector is written, the entire block must be moved since erase operations are at a block level. A spare block is identified in the spare sector pool; the old data is copied and appended with the new sector to be written. The mapping of logical to physical sector is updated and the old block is added to spare sector pool. If an error is detected during program or erases operation, the entire block is marked as bad on the bad block table and a spare block is used instead.

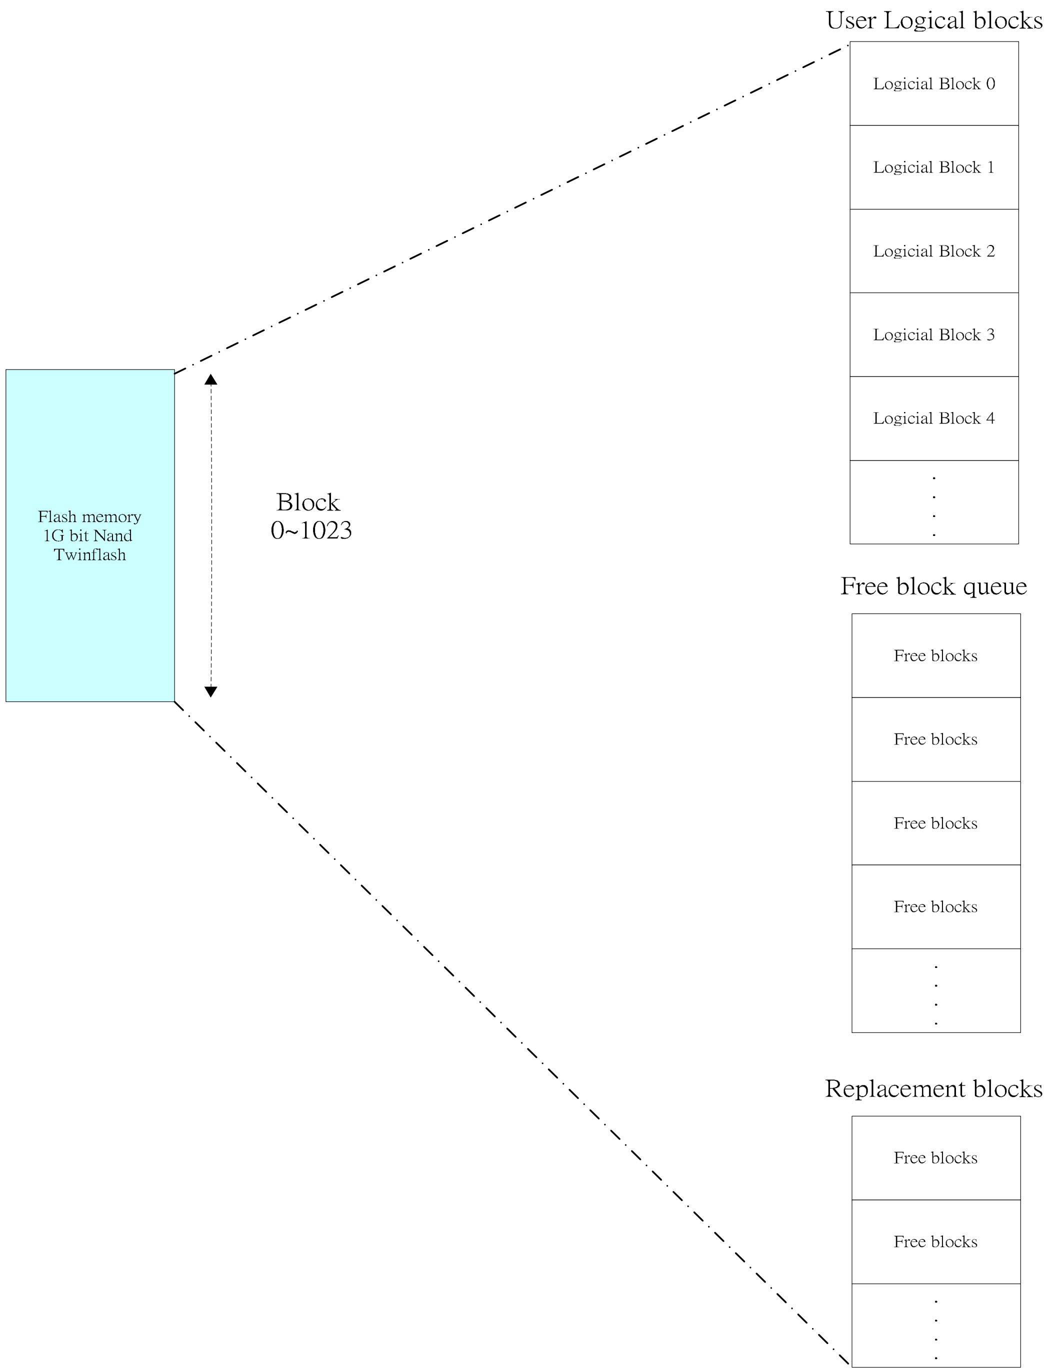

Figure 1 Wear Leveling Architecture

Test result of wear leveling

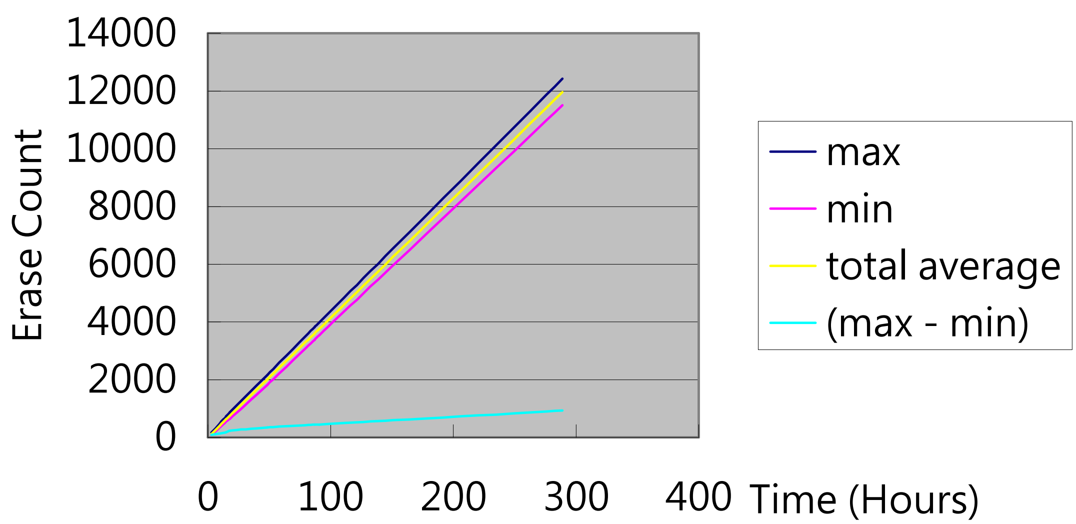

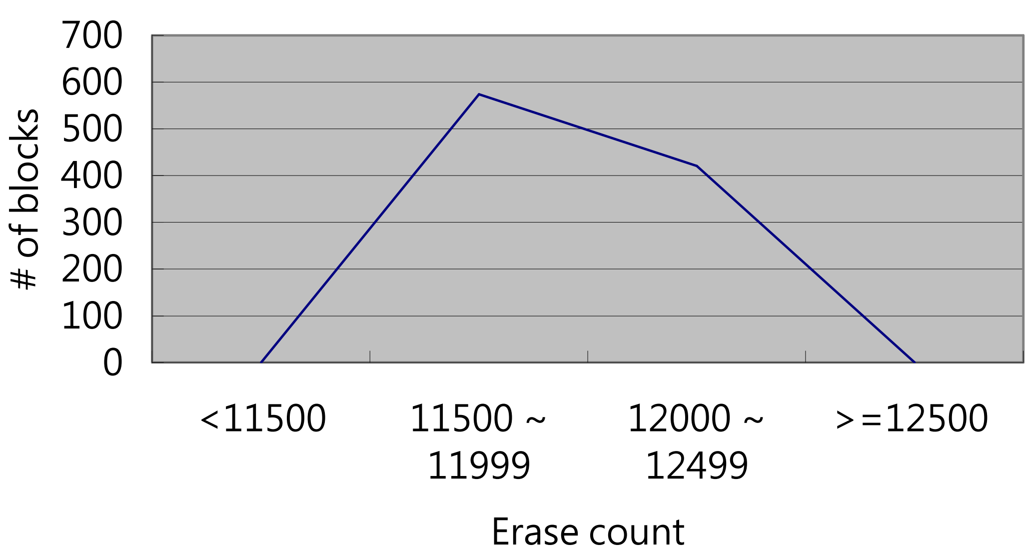

Bad Block management

NAND flash memory incorporates a certain percentage of the early bad blocks, and most NAND flash manufacturers indicate that less than 2% of the total blocks in the flash memory could be identified as bad. These blocks are taken as bad and should not be programmed or erased any further. Bad block management functions automatic identification on initialization where the bad blocks are located and map them out of the array. The maps are updated over time to incorporate any late defect blocks as the memory is used, and these specified blocks are instructed not to use for storage. In addition, a reliable system should manage bad blocks over the lifetime of a NAND Flash device since bad blocks may be triggered during the run-time of the device. When getting a run-time bad block, it may contain previously programmed valid data and may be recoverable. It is a risk losing information by storing it in a bad block if bad blocks are not handled very well. Also, this mechanism should scan out the bad blocks despite power loss.

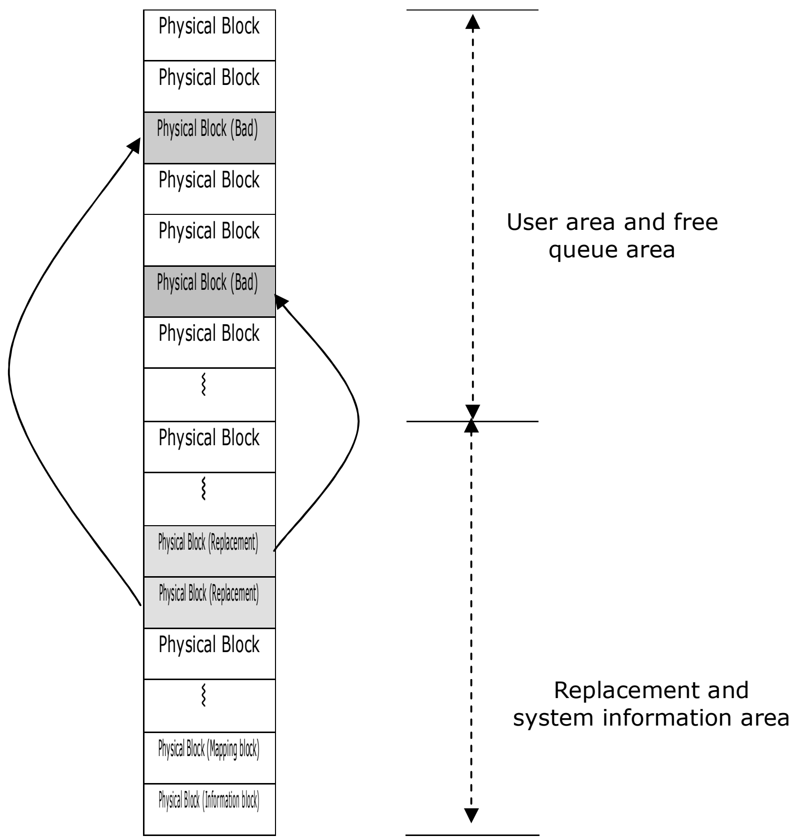

Figure 2 Bad Block Management Diagram

Revision History

Revision	Date	Description	Remark
1.0	June 2, 2009	Official release	

Apacer Technology Inc.

4th Fl., 75 Xintai 5th Rd. Sec.1
Hsichih, Taipei County 221, Taiwan
Tel: +886-2-2698-2888   Fax: +886-2-2698-2889
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.