

Wear Leveling in NAND Flash Memory

Introduction

NAND flash memory is widely used in today's embedded systems for code and data storage applications. Some applications are required to perform numerous Program and Erase operations on the NAND flash memory after system boot-up. NAND flash memories are generally specified with a limited number of Program/Erase (P/E) cycles per block. If the P/E cycles are not evenly distributed across the memory, individual memory blocks can exceed their endurance specification limit. To prevent this scenario from happening, a technique known as "wear leveling" is widely used in NAND flash memory management. Essentially, wear leveling is a way to average out the number of P/E cycles across all usable blocks in a NAND flash device so that the number of bad blocks created over time as a result of frequent PE cycling is minimized. Today, wear leveling has become a critical part of NAND flash management in embedded systems that require frequent Program and Erase operations.

What is Wear Leveling

All flash memory devices suffer wear as they are erased and reprogrammed. They all have finite P/E cycle endurance limits. The actual number of usable P/E cycles will vary from manufacturer to manufacturer, but they all have wear out mechanisms that will eventually render the device unreliable and unusable. When using flash memory in an application, it is important to evaluate how frequently data within the flash will be updated and make sure the application does not exceed the cycle endurance specified by the manufacturer. The wear out mechanism occurs at the cell level and not at the chip level. It is possible to exceed the endurance specification of cells within a single block of memory while the rest of the chip remains intact and fully usable. This can occur if a single block is repeatedly erased and reprogrammed while the remaining blocks are either unused or programmed with static (unchanging) data.

To maximize NAND flash lifetimes, techniques have been developed to evenly distribute the burden of repeated P/E cycles over a larger set of blocks. The name given to these techniques is wear leveling and they can greatly enhance and extend the apparent endurance of the flash memory.

The following example shows the endurance-related lifespan of a NAND flash block with and without wear leveling.

First assume no wear-leveling is used in an application where a specific memory block is to be updated 100 times per hour. Using Macronix SLC NAND MX30LF1G08AA as an example, we can expect each physical block to be able to endure 100,000 P/E cycles (please refer to the latest datasheet). Since the embedded system does not implement any wear leveling, we can expect the 100,000-cycle threshold to be reached in a short period of time.

$$(100,000 \text{ cycles} \times 1 \text{ fixed block}) / (100 \text{ updated} \times 24 \text{ hours}) = 41.7 \text{ days}$$

As the example shows, with no wear leveling, a memory block will need to be taken out of service after roughly 1.5 months of usage.

Let's take a look what happens when wear leveling is implemented and P/E cycles are distributed over 10 different blocks.

$$(100,000 \text{ cycles} \times 10 \text{ fixed block}) / (100 \text{ updated} \times 24 \text{ hours}) = 417 \text{ days} \geq 1 \text{ year}$$

By spreading the P/E cycles over 10 different blocks, we are able to extend the lifespan of the specific blocks by 10x.

In a real embedded system, wear leveling is done through the help of an address translation table which serves as a mapping between the host address (logical address) and the NAND flash address (physical address). The table must be continuously maintained. The wear leveling algorithm is responsible for translating logical addresses received from the application into physical addresses within the NAND device and updating the translation tables whenever the contents of any block is relocated. The block containing the table needs to be managed through wear leveling as well. Otherwise, the block may become worn out and the transfer table corrupted, resulting in a system failure. In general, the spare area of each individual usable block is used to record Erase counts to keep track of P/E cycles.

Wear Leveling Techniques

Dynamic Wear Leveling and Static Wear Leveling are two primary ways to implement wear leveling.

Dynamic Wear Leveling

Dynamic Wear Leveling implements wear leveling only on blocks that contain dynamic data. The blocks containing static data remain untouched. For example, if 25% of NAND capacity is for code storage that is seldom updated during its life time and 75% of NAND capacity is used for storing dynamic data that is updated frequently under embedded system operation, the wear leveling would be applied to the 75% of the memory used for dynamic data, while the remaining 25% of memory would be excluded from wear leveling. The advantage of this kind of implementation is that it ensures critical code will be kept in the most reliable portion of the NAND and not be subject to P/E cycling degradation.

Static Wear Leveling

Static Wear Leveling involves the wear leveling of all usable blocks, including the blocks containing static data. In the example above, all of the blocks reserved for code storage would also be made available to the wear leveling algorithm. The code stored there would be moved to free blocks having higher wear. Blocks previously occupied by code would in theory have less wear and would be freed for dynamic data storage. This extends the apparent wear life of the device by utilizing the wear available in code blocks too. To ensure good code data retention, care should be taken when moving code to blocks previously used for dynamic data. The target blocks for code storage should not be near the end of their useful wear life. This scheme increases the overall complexity of the wear leveling algorithm due to the need to track the usage and wear of all blocks (static code and dynamic data; not just dynamic data).



Summary

With the implementation of wear leveling for NAND flash memory, the lifespan of NAND flash memory can be extended to its maximum endurance limit. Macronix provides SLC NAND flash memory for customers who desire both reliability and lower cost data storage in their embedded systems.

References

Revision No.	Description	Page	Date
REV. 1	Initial Release	ALL	MAR. 10, 2014



Except for customized products which have been expressly identified in the applicable agreement, Macronix's products are designed, developed, and/or manufactured for ordinary business, industrial, personal, and/or household applications only, and not for use in any applications which may, directly or indirectly, cause death, personal injury, or severe property damages. In the event Macronix products are used in contradicted to their target usage above, the buyer shall take any and all actions to ensure said Macronix's product qualified for its actual use in accordance with the applicable laws and regulations; and Macronix as well as it's suppliers and/or distributors shall be released from any and all liability arisen therefrom.

Copyright© Macronix International Co., Ltd. 2014. All rights reserved, including the trademarks and tradename thereof, such as Macronix, MXIC, MXIC Logo, MX Logo, Integrated Solutions Provider, NBit, Nbit, NBiit, Macronix NBit, eLite-Flash, HybridNVM, HybridFlash, XtraROM, Phines, KH Logo, BE-SONOS, KSMC, Kingtech, MXSMIO, Macronix vEE, Macronix MAP, Rich Audio, Rich Book, Rich TV, and FitCAM. The names and brands of third party referred thereto (if any) are for identification purposes only.

For the contact and order information, please visit Macronix's Web site at: <http://www.macronix.com>