Technical Note

Wear Leveling in Micron® NAND Flash Memory

Introduction

This document describes the recommended wear leveling algorithm to be implemented in the flash translation layer (FTL) software for Micron® NAND Flash memory.

Wear Leveling and the Flash Translation Layer

In Micron single-level cell (SLC) and multilevel cell (MLC)NAND Flash memory, each physical block can be programmed and erased reliably up to 100,000 and 10,000 times, respectively. For write-intensive applications, it is recommended to implement a wear leveling algorithm to monitor and spread the number of write cycles per block. In memory devices that do not use a wear leveling algorithm not all blocks get used at the same rate. The wear leveling algorithm ensures that equal use is made of all the available write cycles for each block.

Wear leveling is implemented in the flash translation layer (FTL), which is the additional software layer between the file system and the NAND Flash memory. The FTL allows operating systems to read and write to NAND Flash memory devices in the same way as disk drives and provides the translation from virtual to physical addresses. Wear leveling can also be implemented by the file system directly on the NAND Flash (see Figure 1).

Refer to the data sheets for the full list of root part numbers and for further information on the devices (see “References” on page 5).

Figure 1: Software Tool Chain for an Embedded System Using NAND

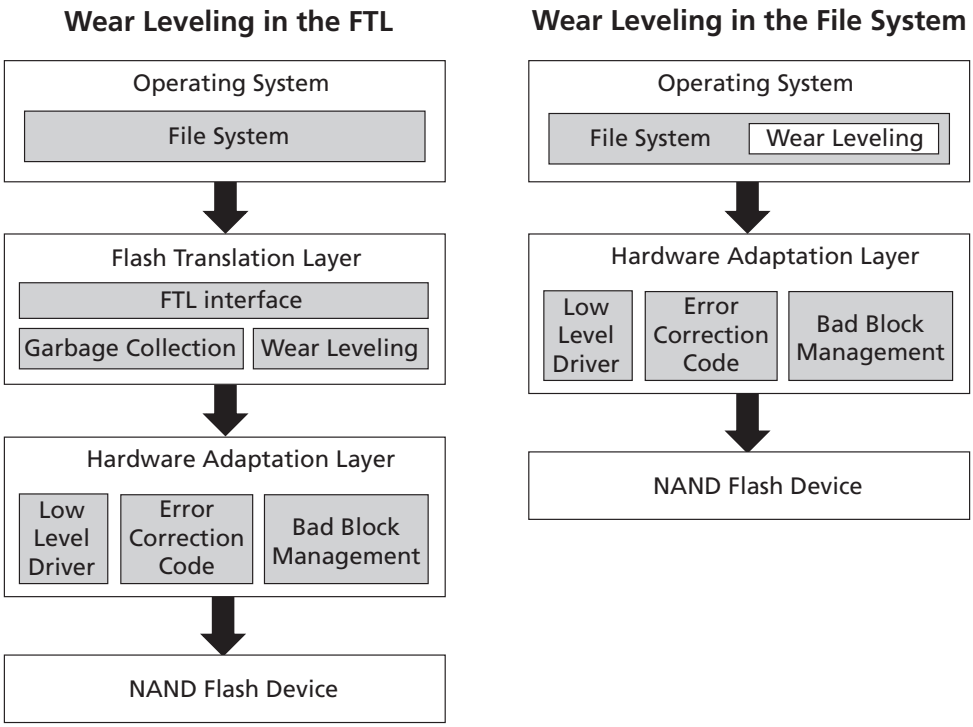

Increasing the Lifespan of a NAND Flash Memory Device

This section shows how using wear leveling can increase the lifespan of a NAND Flash device.

Lifetime Without Wear Leveling

For systems that have a file allocation table (FAT) based file system, the FAT table is always stored in the same virtual blocks. Frequent FAT table updates are required during data WRITE operations, which implies frequent erase cycles on the same physical blocks, hence a reduced NAND Flash lifetime.

The following example calculates how many times a FAT table (FAT32 and a cluster size of 2KB) is updated when writing a 10MB file to a NAND Flash memory with a physical erase unit of 16KB (NAND small page device).

To write a file of 10MB, 5KB entries in FAT and 5KB clusters in the file system are required. This corresponds to 640 physical NAND Flash blocks.

This means that the file can be written at the same location 20 times:

$$20 \times 5120 = 102400$$

This is greater than the maximum number of program/erase cycles.

The expected NAND Flash lifetime can be calculated as follows:

$$\text{Expected lifetime} = \frac{\text{Size of NAND flash} \times n\text{umber of erase cycles} \times \text{FAT overhead}}{b\text{ytes written per day}}$$

This means that if the application writes at 3KB/s, the expected lifetime of the NAND blocks is:

$$\text{Expected lifetime} = \frac{10Mbyte \times 20 \times 0.7}{(3Kbyte/s) \times 24 \times 60 \times 60} = 0.55 \text{ days}$$

In a NAND Flash, when virtual blocks are mapped to the same physical blocks, the lifetime of the device is significantly reduced, independently of its size.

Lifetime with Wear Leveling

Wear leveling extends the lifetime of NAND Flash devices because it ensures that even if an application writes to the same virtual blocks over and over again, the PROGRAM/ERASE cycles will be distributed evenly over the NAND Flash memory.

For example, the expected lifetime of a 64MB (512Mb) NAND Flash device can be calculated as follows:

$$\text{Expected lifetime} = \frac{64\text{Mbyte} \times 100Kcycles \times 0.7}{(3KBbyte/s) \times 24 \times 60 \times 60} = 18{,}124 \text{ days (about 49.7 years)}$$

In this example, 0.7 is the file system overhead.

Wear Leveling Algorithms

Wear leveling is associated with a block aging table (BAT) to store information about which blocks have been erased in a selected period of time.

There are two kinds of wear leveling that can be implemented in the FTL:

- Dynamic wear leveling
- Static wear leveling

Dynamic Wear Leveling

When applying the dynamic wear leveling, new data is programmed to the free blocks (among blocks used to store user data) that have had the fewest WRITE/ERASE cycles.

Static Wear Leveling

With static wear leveling, the content of blocks storing static data (such as code) is copied to another block so that the original block can be used for data that is changed more frequently.

Static wear leveling is triggered when the difference between the maximum and the minimum number of WRITE/ERASE cycles per block reaches a specific threshold. With this particular technique, the mean age of physical NAND blocks is maintained constant.

References

The following documents related to NAND Flash memory are available on *www.micron.com*

- NANDxxx-A, single-level cell, small page, 528-byte/264-word page, 3 V supply voltage, NAND Flash memory data sheets
- NAND01G-B2B_NAND02G-B2C, single level cell, large page, 2112-byte/1056-word page, 1.8 V/3 V supply voltage, NAND Flash memory data sheets
- NANDxxGW3C2B, multilevel cell, large page, 2112-byte page, 3 V supply voltage, NAND Flash memory data sheets
- Garbage collection in NAND Flash memory
- Error Correction Code in Single-Level Cell NAND Flash Memory
- Bad Block Management in NAND Flash Memory

Conclusion

Implementing wear leveling is recommended as part of the software tool chain (either in the FTL or file system) to increase the lifetime of NAND Flash in an embedded system.

In addition, it is recommended that you implement garbage collection and bad block management algorithms. It is mandatory that you implement error correction code algorithms.

To help integrate NAND Flash memory in applications, Micron can provide a full range of software solutions, such as a file system, sector manager, drivers, and code management. Contact your Micron sales representative or visit *www.micron.com* for more details.

8000 S. Federal Way, P.O. Box 6, Boise, ID 83707-0006, Tel: 208-368-3900
www.micron.com/productsupport Customer Comment Line: 800-932-4992