

## Wear Leveling

Author: Ken Perdue

AN98521 discusses the wear leveling approach, its advantages, and methods of implementation to extend the flash and product life expectancy.

### 1 Abstract

Today's embedded systems utilize flash devices for both code and data storage applications. Some of these applications use intensive programming and erasure operations which raises concerns about flash product life expectancy. Flash memory devices are limited to a finite number of Program Erase (PE) cycles. An embedded system's product life expectancy is tied directly to the flash usage model and the number of PE cycles guaranteed by the flash device. Utilization of appropriate systems tools like wear leveling can extend the flash and the product life expectancy. It is good practice to consider the flash usage model in a system to determine if wear leveling could extend flash life to meet the system requirements.

This document will highlight three areas:

- What is wear leveling?
- Advantages of using wear leveling.
- Wear Leveling methods.

### 2 Wear Leveling

Cypress's NOR flash endurance specification defines the conditions and the number of erase operations that can be successfully performed on a given flash erase unit (sector). Cypress offers two types of non-volatile NOR flash memory: Single-bit-per-cell floating-gate flash and Cypress's proprietary two-bit-per-cell MirrorBit® flash. These two technologies are designed to typically withstand 1,000,000 (floating gate) and 100,000 (MirrorBit) Program followed by Erase (PE) cycles.

There are cases where the flash endurance specification may not meet the projected cumulative erasures, for a given sector, in a particular system application (usage model). Flash File Systems (FFS) that employ wear-leveling algorithms provide a means to extend flash life expectancy.

Wear-leveling algorithms arrange or store data in a manner that sector erasures are distributed more evenly across the flash memory array or portions of the array. A host system using a flash file system performs its reads and writes to logical-sector addresses. A wear-leveling algorithm is typically part of a flash file system which remaps logical-sector addresses to different physical-sector addresses in the flash array. The wear-leveling algorithm tracks the least used sectors in the flash to identify the next area to write data. This process reduces or eliminates single sector failures resulting from a high concentration of erase cycles in a limited number of sectors during the expected system life time.

There is also an inverse relationship between the number of PE cycles performed on a sector and the duration of time the sector will reliably retain the data programmed into the sector. Typical data retention time is specified in the flash industry assuming no erase operations on a sector. As more PE cycles are applied to a sector the time that data will be reliably retained is reduced. Generally, this does not cause a problem for system operation because a sector that is frequently erased is likely to be erased again long before there is any danger of data loss. However, wear leveling reduces the average number of PE cycles on frequently cycled logical sectors, and thereby extends the data retention time on these sectors longer than it would be if all PE cycles were applied to the same physical sector.

Cypress's Application Note: [Practical Guide to Endurance and Data Retention](#) provides additional useful information to better understand Flash Endurance and Data Retention.

## 2.1 How Wear Leveling Affects Flash Life Expectancy

Below are example cases that highlight how flash (and the system) life expectancy is affected by wear leveling. These examples use a Cypress MirrorBit S29GL01GP flash device that has 1024 (128 Kbyte) sectors with a typical endurance rating of 100K PE cycles.

### 2.1.1 Example Application With No Wear Leveling

The first example assesses the flash device life expectancy in a system that does not utilize wear leveling. This system uses fixed addressing that maps each logical sector to the same physical sector, which results in all updates occurring at the same physical address. The sector associated with the subject physical addresses must be erased before new data can be programmed. In this example a 128 KByte data file will be updated 50 times per hour. Since there is no wear leveling the system reuses the same sector for each update.

The total flash life expectancy is:

$$= \frac{100,000 \text{ cycles} \times 1 \text{ sector used in the Flash}}{1 \text{ sector / file} \times 50 \text{ updates per hour} \times 24 \text{ hours per day}}$$

<84 days << 1 year

In this configuration where wear leveling is not utilized, the flash has a life expectancy of much less than one year.

### 2.1.2 Extending Flash Life Expectancy Via Wear Leveling

In this second example, wear leveling is utilized. This means the data file or logical sector(s) are no longer tied to a single physical sector, allowing the file programming and erasure to be distributed across a number of sectors. After a limited number of PE cycles on a physical sector, a new physical sector with less PE cycles on it is selected as the new location of data in the frequently updated logical sector. The logical to physical address map is updated to point to the new physical location for the logical sector.

Even though the logical sector may receive many PE cycles, those PE cycles are performed only a few times on each of the physical sectors that hold the logical sector during a small portion of system operational life. In this way the PE cycle wear is spread across multiple physical sectors. Sometimes, the lowest cycled sector needed for the new location of a frequently cycled logical sector already contains data. In this case the data in the low cycled sector is first copied to another unused sector before erasing the low cycled sector.

This example will use 96 of the total 1024 available sectors in the wear-leveling process. Note the 96 sectors will be programmed and erased equally; the 128 Kbyte file is programmed at a rate of 50 file updates/ hour. Distributing the Erasure cycles across the 96 sectors significantly extends the useful flash life.

The total flash life expectancy using wear leveling is:

$$= \frac{100,000 \text{ cycles} \times 96 \text{ sectors}}{1 \text{ sector / file} \times 50 \text{ updates per hour} \times 24 \text{ hours per day}}$$

~8000 days ~ 21.9 years

This example demonstrates that wear leveling increases the flash life expectancy to ~22 years.

Most file updates do not require changing large amounts of data such as the entire contents of a 128 KB sector. In the last simple example, even a small update of 512 bytes within the 128 KB sector did require copying the entire sector to RAM, updating the desired 512 bytes, erasing the physical sector, then writing back all 128 KB from RAM. This requires a great deal of copying and programming overhead for each small update and cycles the entire sector for every small update.

A more efficient FFS will provide logical to physical location mapping for much smaller blocks of data such as 512 bytes or 4 KB blocks. With this finer granularity of mapping it is only necessary to program an updated version of the modified block to an erased portion of a sector and mark the old copy of the block in the map as deleted. In this way each small update uses a much smaller portion of a sector and only when all of a sector has been used for these small updates is it necessary to copy the remaining valid blocks in the sector to a new sector and erase the sector to make room for more updates. This approach greatly increases the number of times data may be

updated without requiring a PE cycle on a sector and thus further extending the update capacity or operational life of the flash.

This third example will assume the same 96 sectors of available space but updates will be performed in 512 byte blocks. There are 256 blocks of 512 bytes in each 128 KB sector.

The total flash life expectancy using 512 byte blocks with wear leveling is:

$$= \frac{100,000 \text{ cycles} \times 96 \text{ sectors} \times 256 \text{ blocks}}{50 \text{ block updates per hour} \times 24 \text{ hours per day}}$$
$$\cong 2048000 \text{ days} \cong 5610 \text{ years}$$

Or the block updates could be increased to 28000 updates per hour:

$$= \frac{100,000 \text{ cycles} \times 96 \text{ sectors} \times 256 \text{ blocks}}{28000 \text{ block updates per hour} \times 24 \text{ hours per day}}$$
$$\cong 3657 \text{ days} \cong 10 \text{ years}$$

Clearly, a modern FFS with wear leveling can dramatically increase the useful life time of the flash device to meet the needs of nearly any application.

## 2.2 Wear Leveling Methods

Wear leveling is a system level tool to address flash endurance limitations by arranging data so that required erasures and re-writes are distributed more evenly across the flash storage medium. In this way, no single flash sector fails due to a high concentration of erase cycles during the desired system life time. A designer needs to understand the flash usage model for a given design and its life expectancy. This information can be used to determine the appropriate wear leveling method to be implemented. Below are three simple ways wear leveling could be utilized:

### 2.2.1 A System with No Wear Leveling

Some applications do not need frequent updates and can meet the required system life time requirement without wear leveling. But, the system designer must recognize that the system life expectancy can be limited by frequently erasing and exhausting a single or a few sectors in a flash device. A system not utilizing wear leveling may see flash sector failures due to a high concentration of erase cycles in a limited number of sectors. Note in many cases, there may still be areas in the flash array that are either unused or under used. Effectively, the system life ends much sooner than most of the flash array because of the dependence on frequent cycling of only a few physical sectors.

### 2.2.2 Dynamic Wear Leveling

Dynamic wear leveling spreads the PE cycles across sectors in a portion of the total memory space where updates are done on a regular basis. Systems using dynamic wear leveling do not touch static data. In a flash usage model where 75% of the array is used for code storage or static data, only 25% is available for wear leveling. This model offers increased flash endurance compared to no wear leveling, but all updates are concentrated within 25% of the available space. The advantage is some wear leveling without requiring a means to remap all accesses to memory. The static areas can be read directly without software or hardware overhead to locate the physical sectors where code or static data is placed.

### 2.2.3 Static Wear Leveling

A static wear-leveling algorithm tracks the least used physical sectors across the entire flash address space and uses all sectors to more evenly distribute PE cycles over the entire flash array. If a low cycle sector is unused, update programming is simply done in the erased sector. If a low cycle sector contains static data, the data is copied to a more heavily used location before erasing and programming the low cycle sector with an update. The static wear leveling further optimizes the flash endurance compared to no wear leveling or dynamic wear leveling.

However, it requires that all read accesses to the flash go through some type of logical to physical mapping that can slow down the reading process for static as well as dynamic logical sectors. Often this mapping for code and static data is handled by a hardware Memory Management Unit (MMU) in the system microprocessor. When a static area of the flash is first accessed an interrupt handler is called to access the FFS logical to physical map and move the mapping into the MMU. The MMU then performs the logical to physical mapping at high speed (needing no additional access time) as code or static data is needed by the microprocessor.

### 3 Conclusion

Today's embedded systems have progressed beyond using NOR flash exclusively for code storage to more feature rich data storage applications requiring periodic data updates. The examples highlight how the flash's useful life expectancy can be dramatically improved when wear leveling is employed. It is recommended to understand the flash usage model of the intended application and utilize system tools like a FFS with wear leveling as required to meet product life expectations.

Cypress makes available flash driver and FFS software. The flash file system source code provided by Cypress delivers a production-grade block driver and stand-alone file system that are specifically designed for Cypress flash memory. This file system significantly reduces the amount of software development time, enabling faster time to market by simplifying flash memory integration.

### 4 References

- Cypress Application Note  
[Practical Guide to Endurance and Data Retention](#)
- Cypress software support  
<http://www.spansion.com/Support/Pages/DriversSoftware.aspx>

## Document History Page

Document Title: AN98521 - Wear Leveling Document Number: 001-98521				
Rev.	ECN No.	Orig. of Change	Submission Date	Description of Change
**	–	–	04/30/2008	Initial version
*A	–	–	06/15/2010	Global updates
*B	4958448	MSWI	10/12/2015	Updated in Cypress template
*C	5823035	AESATMP8	07/18/2017	Updated logo and Copyright.
*D	6163315	PRIT	05/02/2018	Sunset review.

## Worldwide Sales and Design Support

Cypress maintains a worldwide network of offices, solution centers, manufacturer's representatives, and distributors. To find the office closest to you, visit us at [Cypress Locations](#).

### Products

Arm® Cortex® Microcontrollers	<a href="http://cypress.com/arm">cypress.com/arm</a>
Automotive	<a href="http://cypress.com/automotive">cypress.com/automotive</a>
Clocks & Buffers	<a href="http://cypress.com/clocks">cypress.com/clocks</a>
Interface	<a href="http://cypress.com/interface">cypress.com/interface</a>
Internet of Things	<a href="http://cypress.com/iot">cypress.com/iot</a>
Memory	<a href="http://cypress.com/memory">cypress.com/memory</a>
Microcontrollers	<a href="http://cypress.com/mcu">cypress.com/mcu</a>
PSoC	<a href="http://cypress.com/psoc">cypress.com/psoc</a>
Power Management ICs	<a href="http://cypress.com/pmic">cypress.com/pmic</a>
Touch Sensing	<a href="http://cypress.com/touch">cypress.com/touch</a>
USB Controllers	<a href="http://cypress.com/usb">cypress.com/usb</a>
Wireless Connectivity	<a href="http://cypress.com/wireless">cypress.com/wireless</a>

### PSoC® Solutions

[PSoC 1](#) | [PSoC 3](#) | [PSoC 4](#) | [PSoC 5LP](#) | [PSoC 6 MCU](#)

### Cypress Developer Community

[Forums](#) | [WICED IOT Forums](#) | [Projects](#) | [Video](#) | [Blogs](#) | [Training](#) | [Components](#)

### Technical Support

[cypress.com/support](http://cypress.com/support)

All other trademarks or registered trademarks referenced herein are the property of their respective owners.



© Cypress Semiconductor Corporation, 2008-2018. This document is the property of Cypress Semiconductor Corporation and its subsidiaries, including Spansion LLC ("Cypress"). This document, including any software or firmware included or referenced in this document ("Software"), is owned by Cypress under the intellectual property laws and treaties of the United States and other countries worldwide. Cypress reserves all rights under such laws and treaties and does not, except as specifically stated in this paragraph, grant any license under its patents, copyrights, trademarks, or other intellectual property rights. If the Software is not accompanied by a license agreement and you do not otherwise have a written agreement with Cypress governing the use of the Software, then Cypress hereby grants you a personal, non-exclusive, nontransferable license (without the right to sublicense) (1s) under its copyright rights in the Software (a) for Software provided in source code form, to modify and reproduce the Software solely for use with Cypress hardware products, only internally within your organization, and (b) to distribute the Software in binary code form externally to end users (either directly or indirectly through resellers and distributors), solely for use on Cypress hardware product units, and (2) under those claims of Cypress's patents that are infringed by the Software (as provided by Cypress, unmodified) to make, use, distribute, and import the Software solely for use with Cypress hardware products. Any other use, reproduction, modification, translation, or compilation of the Software is prohibited.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CYPRESS MAKES NO WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, WITH REGARD TO THIS DOCUMENT OR ANY SOFTWARE OR ACCOMPANYING HARDWARE, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. No computing device can be absolutely secure. Therefore, despite security measures implemented in Cypress hardware or software products, Cypress does not assume any liability arising out of any security breach, such as unauthorized access to or use of a Cypress product. In addition, the products described in these materials may contain design defects or errors known as errata which may cause the product to deviate from published specifications. To the extent permitted by applicable law, Cypress reserves the right to make changes to this document without further notice. Cypress does not assume any liability arising out of the application or use of any product or circuit described in this document. Any information provided in this document, including any sample design information or programming code, is provided only for reference purposes. It is the responsibility of the user of this document to properly design, program, and test the functionality and safety of any application made of this information and any resulting product. Cypress products are not designed, intended, or authorized for use as critical components in systems designed or intended for the operation of weapons, weapons systems, nuclear installations, life-support devices or systems, other medical devices or systems (including resuscitation equipment and surgical implants), pollution control or hazardous substances management, or other uses where the failure of the device or system could cause personal injury, death, or property damage ("Unintended Uses"). A critical component is any component of a device or system whose failure to perform can be reasonably expected to cause the failure of the device or system, or to affect its safety or effectiveness. Cypress is not liable, in whole or in part, and you shall and hereby do release Cypress from any claim, damage, or other liability arising from or related to all Unintended Uses of Cypress products. You shall indemnify and hold Cypress harmless from and against all claims, costs, damages, and other liabilities, including claims for personal injury or death, arising from or related to any Unintended Uses of Cypress products.

Cypress, the Cypress logo, Spansion, the Spansion logo, and combinations thereof, WICED, PSoC, CapSense, EZ-USB, F-RAM, and Traveo are trademarks or registered trademarks of Cypress in the United States and other countries. For a more complete list of Cypress trademarks, visit [cypress.com](http://cypress.com). Other names and brands may be claimed as property of their respective owners.