AN98549

Endurance and Retention Management and Validation

AN98549 provides an introduction to the embedded program/erase operation mechanics and associated side effects. It provides a practical understanding of endurance and deliver guidance for developing appropriate flash subsystem test and validation routines, using Cypress MirrorBit NOR flash as the example.

1 Introduction

Flash memory devices are subject to application related physical degradation that can lead to device and system failure. Data retention and program/erase endurance are the terms used to describe the two principal end-of-life application dependent parameters for flash. Data Retention is the period of time a stored bit level can be reliably recovered. Endurance is the term used to describe the lifetime program/erase cycling capability of each sector or block within the flash array. Both Data Retention and Endurance are dynamic variables which are specified in terms of “typical” usage as well as qualified under worst case use conditions. These parameters are related in so much as the degradation associated with program/erase cycle accumulation impacts long term data retention.

The semiconductor technologies employed in modern NOR flash devices, e.g. Cypress MirrorBit®, do not support unlimited program/erase cycle accumulation or unrestrained same sector erase cycling rates. Excessive erase cycling rates will result in the permanent degradation of memory cells and a reduction of the program/erase endurance capability of the flash device. Excessive erase cycling rates are a common source of error during system level validation and reliability assessment, as well as premature system failure. It is important for the user to properly accommodate both endurance and erase cycling rate limitations when designing with and when creating validation strategies for flash subsystems.

This document provides an introduction to the embedded program/erase operation mechanics and associated side effects. It provides a practical understanding of endurance and deliver guidance for developing appropriate flash subsystem test and validation routines, using Cypress MirrorBit NOR flash as the example.

2 Program/Erase Operation Mechanics and Effects

Cypress's MirrorBit flash memory technology is based on a Silicon Nitride isolation barrier cell architecture that uses Channel Hot Electron (CHE) injection for programming and Hot Hole (HH) injection for erasure. These processes manipulate minute numbers of electrons which set charge levels underneath the cell Gate to represent persistent logical states within the flash cells and are similar to programming and erasing processes in Floating Gate technologies used in other NOR flash memory devices.

When programming via CHE injection, the cell transistor is biased with a moderate positive gate voltage, a moderately positive drain voltage, a grounded source and a grounded substrate. Under moderate voltage between the source and the drain, electrons in the channel will be accelerated, gain kinetic energy and become “hot.” The gate voltage, generating a vertical electric field, will then pull the hot electrons toward the silicon-oxide interface. Electrons that gained sufficient kinetic energy will overcome the silicon-oxide energy barrier and get injected into the charge-storage medium, i.e. silicon nitride layer. Hot electrons may also collide with channel atoms near the drain junction and cause secondary electron injection. Programming is complete when enough hot electrons have been injected into the charge-storage medium to raise the cell's threshold voltage to a predetermined level. The voltages of the source and drain terminals are switched to program the complementary bit of the MirrorBit cell. Not all hot electrons may have enough energy to get transported through the tunnel oxide and into the charge-storage medium. Most electrons will be generally collected by the drain terminal, while few may get trapped at the SiO2 layer (commonly referred to as the tunnel oxide).

Figure 1. CHE Injection Programming

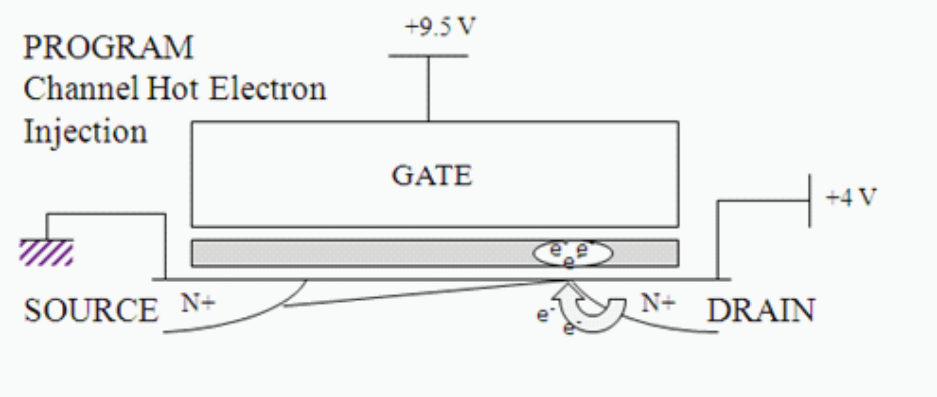

When erasing via HH injection, the cell transistor is biased with a high negative gate voltage, a moderately positive drain voltage, a floating source and a grounded substrate. The voltages of the source and drain terminals are switched to erase the complementary bit of the cell. During erase, electron/hole pairs are generated at the drain, the electrons are swept to the drain terminal, while the holes are swept in the strong junction field, getting hot. The negative bias of the control gate pulls the hot holes, injecting them across the oxide barrier and into the nitride storage layer. In the silicon nitride the holes may either recombine with the electrons injected during programming or get trapped nearby and compensate their field, thereby reducing the threshold voltage of the cell. The erase process continues with successive voltage pulses until the cell's threshold voltage has been reduced to the desired pre-determined level.

Figure 2. HH Injection Erasing

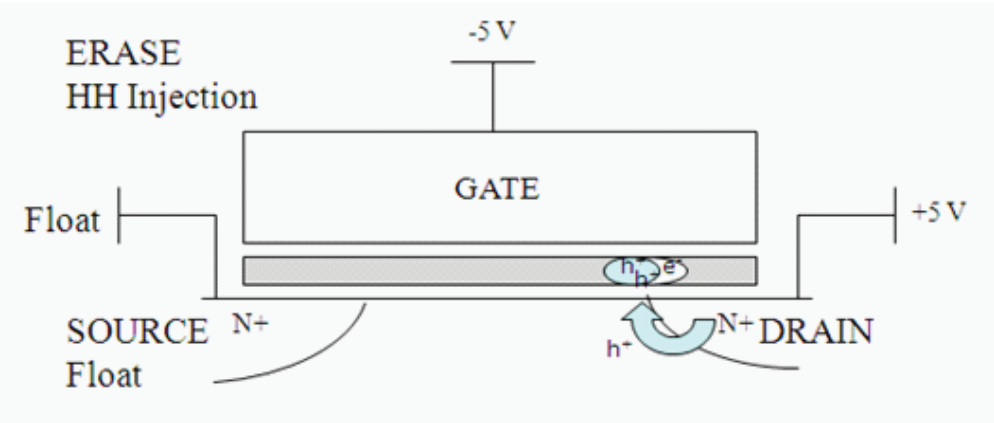

Most holes generated during the erase operation will recombine in the silicon nitride layer or be collected by the gate terminal, while few may get trapped at the SiO2 layer. The vast majority of these hot hole trap sites naturally dissipate via a self annealing process within a short period of time after the erase operation. However, if these charge trap sites are not given sufficient time to anneal, e.g. due to rapid erase cycling, they tend to become permanent. For this reason MirrorBit flash sectors should be erased at a rate that does not exceed 1000 times per day per sector, or once per ~90 seconds. Slower erase rates will further minimize damage, but may not be practical to implement in all cases.

In the event that a sector is repeatedly erased at a rate that prevents completion of the anneal process, trapped charge concentrations can be expected to build rapidly. It is possible that the trapped charge levels normally only seen in an end-of-life sector, i.e. a sector with 100,000 erase cycles, may be accumulated in as few as 1000 rapid erase cycles.

Figure 3. Trapped Charge Build Up

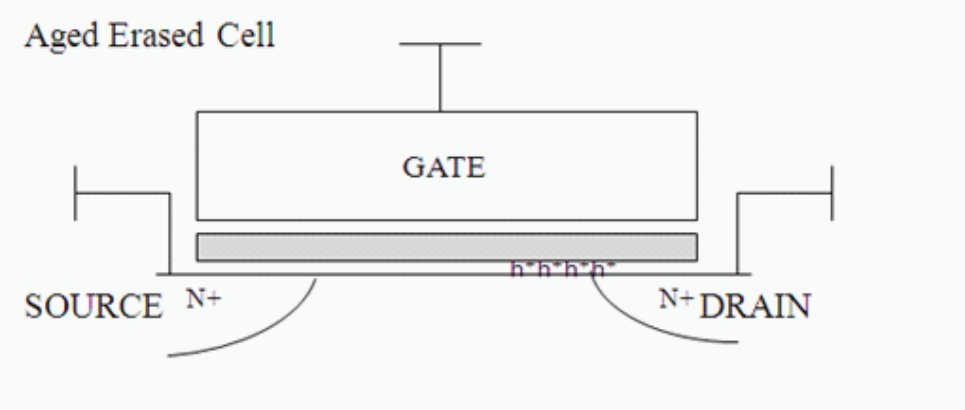

Persistent charge accumulation at the SiO2 layer tends to slow down programming and erase operations, and when excessive will result in the inability of the cell to function properly. This can result in read, programming and/or erase errors, with both immediate and latent characteristics. It is important to note, however, that any potential lasting side effects of an erase operation are localized only to those cells within the sector that is being erased, and do not affect the long-term health of any other sector in the array.

3 In-System Endurance Management

MirrorBit flash sectors are designed to withstand 100,000 program/erase cycles under typical conditions, which are defined by data sheet environmental limits and recommended same sector cycling rates, e.g. operation at nominal voltages within the -40/+85°C temperature range with same sector erases occurring at least 90 seconds apart. Given that most systems are designed to operate for more than the 100 days necessary to reach end-of-life endurance at the maximum 1,000 cycles per day cycling rate, the slower erase cycling rates mandated by realistic system lifetimes will further ensure that the typical 100,000 program/erase endurance cycles will be realized.

For reliable system operation, the endurance capabilities of the flash must be aligned with the system data capture requirements. Sectors used for code storage generally are updated rarely so no special considerations related to endurance are required. Sectors used for data logging and other short term data storage applications may rapidly accumulate erase cycles depending on the frequency and size of the data set(s) being captured. In these applications, the user needs to accurately assess the lifetime data capture requirements of the system and create a data capture strategy that does not exceed the endurance capability of any sector and does not exceed the same sector erase rate guidance. This often warrants the use of wear leveling methodologies to spread the sector erases over a large number of sectors. Wear leveling works via creation of a partition within the flash array to allow a file system to evenly distribute data captures over a group of sectors sufficiently large to assure even sector cycling that does not exceed lifetime per sector cycle accumulation, and also provides adequate anneal time between same sector erases.

3.1 Example: Basic System Endurance Assessment

A radar system has to capture 57 kB of parameter data whenever auto calibration is performed. When operational, auto calibration is performed once per minute. The system will be operated twelve hours per day, every day, for twenty years, at up to 85°C. The system requires 24 MB of flash for boot, operating system and application code. What S29GL-S will best support this system's code storage and lifetime parameter data logging requirements?

First, it is useful to estimate how the data is to be captured in the target flash device. The S29GL256S has 128 kB sectors. As such, the two auto calibration data sets can be captured in a sector while leaving sufficient room for file system overhead.

Next, the sector fill rate and lifetime accumulation should be determined. One capture per minute and two data captures per sector results in a sector fill rate of 30 per hour. With 12 operational hours per day, the resulting sector fill rate is 360 per day. Over the system lifespan of 20 years, that amounts to 2,629,800 sector fills. S29GL-S flash provides 100,000 program/erase cycle endurance, so the flash will support the lifetime auto calibration data capture requirements of this system by wear leveling the data parameter storage over at least 27 sectors.

Use of more sectors for data storage will provide endurance margin by lowering the lifetime accumulation per sector below 100,000.

Next, verify that the sector cycling rate will not be exceeded. Since two data sets will be stored per sector at 60 second intervals, rapid sector cycling will be avoided as same sector erases will necessarily be separated by more than 90 seconds.

Since the application requires 192 sectors (24 MB) of code storage and at least 27 sectors (3.3 MB) for data storage, the 256 sector (32 MB) S29GL256S will support this application. When implementing the application with this device, wear leveling the calibration data over all 64 sectors (8 MB) not used for code storage would be desirable as it would result in less stress on the sectors used for data capture.

4 Endurance Ramifications on Test and Validation

Systems using flash memory for specialized applications commonly need to evaluate the flash operational performance and reliability by simulating the final product's durability under worst case conditions. When attempting to understand the long term behavior of flash memory subsystems, use of accelerated laboratory test conditions are often desired as they enable extrapolation of short term test results to validate that long term flash behaviors meet system lifetime requirements. Understanding flash technology, as it relates to program/erase cycling endurance, is necessary to properly develop flash subsystem validation test routines.

While MirrorBit flash erase performance specifications support sector erase rates in excess of 1 per second, flash performance will degrade, and loss of functionality may occur, if testing accelerates same sector erase cycling rates in excess of 1 per 90 seconds. As such, accelerated test result extrapolation may not provide an accurate prediction of long term flash performance if structured improperly.

Erase time, for some applications, is a primary concern. Flash erase time will exhibit some variability during system test and this variability can be significant if the test program performs multiple same sector erase operations in an unacceptably compressed period of time. Rapid same sector erase cycling will result in the lengthening of subsequent program and erase operations which can be misinterpreted as a prediction of improper operation or flash failure in a real system environment. When validating erase time, it is necessary to ensure adequate anneal time between same sector erases in the test flow to avoid extrapolated erase times from being artificially lengthened. In this case, any assessment of sector erase times must avoid erasure of the same sector within 90 seconds. This can be accomplished by erasing a block of sectors instead of one sector repeatedly and averaging the sector group erase time over the quantity of sectors in the group. Alternatively, sector erase time can be assessed (although less accurately) by erasing the same sector repeatedly provided a 90 second delay is inserted between sector erase commands.

Program/erase endurance, for some applications, is a primary concern. When performing program/erase endurance validation, it is necessary to structure the test such that no sector is erased more than once per 90 seconds. Program/erase endurance validation tests should consecutively erase sectors in a large group followed by at least a 90 second delay, and looping for the desired endurance validation point (not to exceed 100,000 cycles). This method will ensure that no sector is erased at a rate exceeding 1 per 90 seconds and will provide validation that all sectors typically will meet the minimum endurance requirements of the system. Pre-programming all bits in a sector to 1 and verifying all bits are subsequently erased to 0 are part of the embedded erase operation algorithm so no additional programming operation or validation steps are necessary. Testing should monitor the embedded erase operation status using data polling or status register accesses (after the required 90 second delay) to capture erase failures if they occur as well as to ensure maximum test performance (versus only waiting a set period of time between issuance of sector erase commands). Monitoring of the RY/BY# output is also an option and will also deliver maximum test performance and allow errors to be detected by placing a suitable watchdog on the RY/BY# assertion period.

5 Summary

System usage of flash and test and validation program development must take into consideration the properties and characteristics of the flash technology in order to produce reliable test results. In particular, program/erase endurance capabilities of the flash must be properly accommodated to ensure System requirements are met and test and validation test results are properly representative of flash performance. Erasing rates of individual sectors must be limited to no more than 1000 evenly spaced program/erase cycles per day to insure that specified flash lifetimes are achieved.

Document History Page

Document Title: AN98549 - Endurance and Retention Management and Validation Document Number: 001-98549				
Rev.	**ECN No.**	**Orig. of Change**	**Submission Date**	**Description of Change**
**	–	–	09/27/2011	Initial version
*A	4958550	MSWI	10/12/2015	Updated in Cypress template
*B	5814936	AESATMP8	07/13/2017	Updated logo and Copyright.

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	cypress.com/arm
Automotive	cypress.com/automotive
Clocks & Buffers	cypress.com/clocks
Interface	cypress.com/interface
Internet of Things	cypress.com/iot
Memory	cypress.com/memory
Microcontrollers	cypress.com/mcu
PSoC	cypress.com/psoc
Power Management ICs	cypress.com/pmic
Touch Sensing	cypress.com/touch
USB Controllers	cypress.com/usb
Wireless Connectivity	cypress.com/wireless

PSoC® Solutions

PSoC 1 | PSoC 3 | PSoC 4 | PSoC 5LP | PSoC 6

Cypress Developer Community

Forums | WICED IOT Forums | Projects | Video | Blogs | Training | Components

Technical Support

cypress.com/support

Cypress Semiconductor
198 Champion Court
San Jose, CA 95134-1709