

NAND Flash Memory Reliability in Embedded Computer Systems

Ian Olson

INTRODUCTION

NAND flash memory – named after the NAND logic gates it is constructed from – is used for nonvolatile data storage in digital devices. The different types of NAND flash memory target different applications and have significantly different costs and longevity. This paper addresses the two main reliability concerns in NAND flash storage (data retention and endurance), what type of NAND flash memory is best suited for embedded computer systems, and how to optimize computer systems for maximum reliability of NAND flash storage.

History

In the late 1990s, NAND flash memory first began to be used in consumer products such as USB flash drives (also known as thumb drives) and digital cameras. Even though the cost was very high compared to other storage media, such as floppy disks, rotating hard drives, and CD-RWs, the size, durability, and power efficiency of flash memory opened a new world of portable data storage.

Due to its rapidly increasing capacity and decreasing price, flash memory became a viable primary storage medium by 2005 for embedded computer products like first-generation SEL computers (e.g., the SEL-3351 System Computing Platform). At that time, most flash memory was available in the form of memory cards such as CompactFlash[®], SD, and xD cards, which were typically used in digital cameras and portable MP3 music players. The 2.5-inch solid-state drives were also available and offered the greatest capacity, but cost prevented them from seeing widespread use until years later.

The SEL-3351 used the CompactFlash form factor as a primary storage device due to its large capacity and widespread availability. High-capacity for its time, the 8 GB industrial-grade CompactFlash card used in SEL computers in 2005 cost over \$1,000 compared to \$100 to \$150 for the same size consumer-grade CompactFlash. Industrial-grade cards were pricier for many reasons, including higher quality manufacturing processes, more rigorous testing and screening, temperature trials, and low market demand compared to consumer-grade devices.

Today, flash memory is used everywhere, as evidenced by the explosion of flash usage shown in Figure 1. Costs are still dropping, but the rate is slowing as flash manufacturers begin to hit the limits of physical NAND flash structure. Still, usage continues to rise at an amazing pace, with roughly 100 percent annual growth in the amount sold each year.

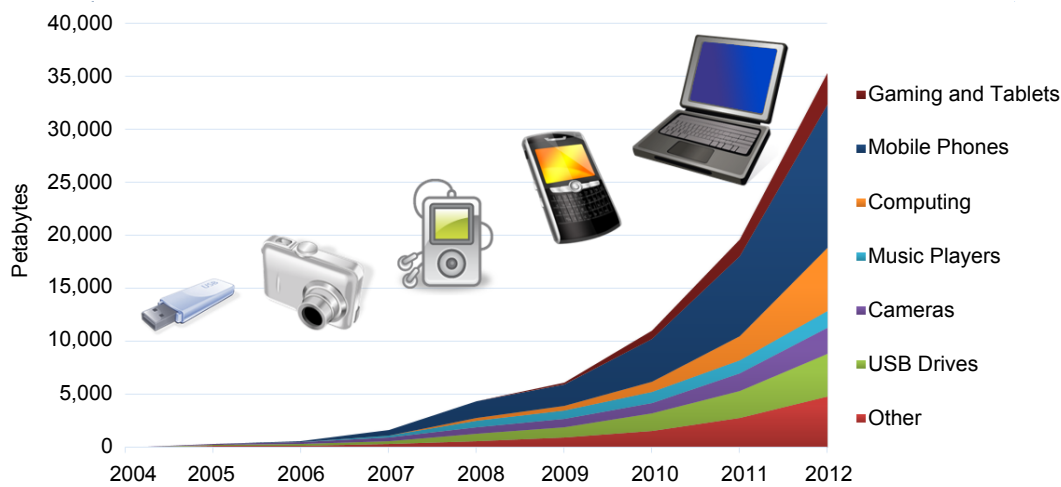


Figure 1 NAND Flash Memory Market Growth [1]

Though improvements in flash manufacturing technology continue to decrease cost and increase storage capacity, the industrial-grade flash memory storage devices used in SEL computers still cost roughly ten times more than similar consumer-grade devices. One of the biggest reasons is storage density, or how much data can be stored on a single flash chip.

Types of NAND Flash Memory

Most NAND flash memory is produced in one of two forms: single-level cell (SLC) or multi-level cell (MLC). In general, both SLC and MLC are manufactured the same way. The major difference is how much data can be stored in a single NAND flash cell.

Flash memory stores data by charging flash cells to a specific level, and it reads data by checking the level of charge (the voltage) on the cells. SLC memory stores one bit per cell; a high voltage indicates a binary 0 while a low voltage indicates a binary 1. MLC stores two bits per cell, dividing the same voltage range into four regions, each corresponding to the binary pairs 11, 10, 01, and 00.

Almost all consumer-grade products use MLC flash memory because it costs roughly one-quarter that of consumer-grade SLC. MLC is more affordable because it has higher data density than SLC and, therefore, it is produced in greater volume. By the same token, the lower storage density of SLC flash memory impacts its popularity, resulting in low production volume and higher prices. When considering the cost differences of consumer-grade versus industrial-grade and MLC versus SLC memories, the combined effect leads to the large cost disparity between industrial-grade SLC flash devices and consumer-grade MLC flash devices.

FAILURE MECHANISMS IN NAND FLASH

The most commonly known failure mechanism in flash storage is related to the overall operational lifespan of the device. Writing and overwriting data to the flash memory result in program/erase (P/E) cycles. P/E cycles create a trapped charge in the NAND flash cells, which reduces the margin for bit errors to occur. As this trapped charge accumulates over time, the bit

error rate eventually becomes so high that the error-correcting code (ECC) cannot compensate, making the flash memory no longer useable.

The number of P/E cycles the flash device can endure over its lifespan while still providing reliable data storage is called endurance. Endurance is tested using a process defined by Joint Electron Devices Engineering Council (JEDEC) standards, and it is typically specified at a fixed temperature and number of P/E cycles (100,000 for SLC and 3,000 for MLC) [2]. While this is a very useful measurement, it only represents a single data point on what is actually a three-dimensional plane. Often overlooked, the two other dimensions are temperature and data retention, both of which are much more critical in embedded and industrial applications.

Data Retention

In flash storage, data retention is the measure of how long the integrity of data can be guaranteed after being written to the flash drive without suffering from data corruption. Once a flash cell is charged, the electrons stored in the cell leak across the NAND gate over time, causing the charge on the cell to decrease. With enough leakage, the voltage level on the cell will drift into the neighboring region, causing the incorrect binary value to be read.

Because SLC flash memory is only divided into two voltage regions, it has more margin for charge loss before a bit flip occurs (a 0 becomes a 1), as shown in Figure 2.




Figure 2 SLC Flash Data Storage

On the other hand, MLC can tolerate much less charge loss before data errors occur because it has a similar voltage range divided into four regions, as shown in Figure 3.




Figure 3 MLC Flash Data Storage

Data retention is a function of the rate of charge loss and the amount of margin between voltage regions. The charge loss is strongly driven by the temperature of the flash memory, and the voltage margin is affected by the number of P/E cycles the flash memory has endured.

A Complete View of Data Integrity

Assuming typical SLC NAND flash specifications of 1 year of data retention at 55 degrees Celsius after enduring 100,000 P/E cycles, and using the Arrhenius equation [3] to calculate the acceleration factor relative to temperature, we can estimate the effect of temperature on data retention. Additionally, using formulas from JEDEC and some empirical data, we can also estimate the increase in data retention when the flash memory is below the maximum number of P/E cycles. Finally, we can combine these two curves into a set of curves to estimate the data retention of an SLC flash drive given the average number of P/E cycles it has endured and the average operating temperature.




Figure 4 SLC Flash Data Retention Relative to Temperature (left) and P/E Cycles (right)

The graphs in Figure 4 are hypothetical. In practice, data retention at extremely low P/E cycles may not be as high as indicated because the mathematical function approaches infinity at zero P/E cycles.

Figure 5 shows that a lightly used (1,000 P/E cycles) SLC flash memory operating at 55 degrees Celsius may have around 16 years of data retention. At higher temperatures, however, the data retention can be dangerously low: 6 months or less at 85 degrees Celsius.




Figure 5 SLC Flash Data Retention Relative to P/E Cycles at Different Temperatures

In comparison, MLC flash memory is typically only rated to endure 3,000 P/E cycles during its lifetime (over 30 times less than SLC), and storage devices that use MLC flash memory are typically only rated to operate at up to 55 or 60 degrees Celsius. Figure 6 compares SLC to the more limited operational range of MLC (3,000 cycles at 55 degrees Celsius).




Figure 6 Comparison of MLC and SLC Flash Data Retention at 55 Degrees Celsius

Although MLC flash memory has much lower endurance than SLC, it may still be considered acceptable for applications that are write-protected or that write very little data to the drive. However, when operating in a warm environment, the MLC flash drive system could fail due to data corruption within 1 to 2 years. In less severe conditions, the MLC flash drive provides moderately better reliability, but the point remains that for any given temperature and P/E cycle count, SLC flash memory provides almost 10 times the data retention and over 30 times the endurance of MLC.

Cold Operation

Flash memory is negatively affected by cold temperatures. While data retention is incredibly good at low temperatures, the ability of a flash cell to be charged accurately decreases rapidly as temperatures drop below 10 degrees Celsius. This means that data written while the flash cell is warm fit the data retention model discussed previously, but retention is reduced for data written while the cell is cold. This problem affects both types of flash, but it is more severe in MLC, making it less suitable than SLC in cold applications as well.

eMLC Flash

Enterprise MLC (eMLC), a variant of MLC, was recently introduced to the flash storage device market. Available at a slightly higher cost, eMLC offers significantly better endurance (typically 30,000 P/E cycles) than standard MLC. This makes eMLC appear to be an alternative to the more expensive SLC flash memory. However, the endurance improvement is at the expense of data retention. The original application for eMLC is in enterprise servers that use flash storage as a high-speed data cache, which requires high endurance and is not concerned with long-term data retention. The eMLC flash memory achieves higher endurance by overprovisioning (allocating a large amount of spare blocks) and allowing a relaxed P/E speed and charge level, which decreases

the wear from each P/E cycle. Unfortunately, the relaxed P/E process decreases data retention, making it worse than standard MLC at low P/E cycles.

IMPROVING RELIABILITY WITH SYSTEM DESIGN

While the problem of data retention can be a concern even for SLC flash memory, the demonstrated reliability of embedded and industrial computer systems over the last decade indicates that SLC is indeed an excellent storage medium for these applications. In addition, there are ways to design the embedded system to minimize P/E cycles, maintain data retention, and monitor flash health, further improving the reliability of flash storage.

Minimize P/E Cycles

The number of P/E cycles can be minimized by avoiding disk-write activity unless necessary for the application. For example, virtual memory in a Microsoft Windows® operating system (OS) can be disabled. Virtual memory creates a page file on disk and accesses it like random-access memory (RAM), with substantial amounts of read and write operations. Embedded systems typically have a static application load. If the system is designed with the appropriate amount of system RAM, the page file is completely unnecessary, and it can actually decrease performance and determinism if left enabled.

Most Windows Embedded operating systems include a feature called the enhanced write filter (EWF). The EWF caches all disk-write operations in RAM until a commit operation is executed. This operation writes the cached data all at once to the flash drive, eliminating the effect of periodic small disk-write operations that are unavoidable in many applications.

Maintain Data Retention

Although minimizing P/E cycles is the best first step, it is beneficial to not completely eliminate all P/E cycles. The controllers in flash drives have wear leveling algorithms, which spread the write operations across all flash cells evenly, to prevent any one flash cell from wearing out before all of the rest. Writing small amounts of data to the drive over time will cause the wear leveling feature to eventually rewrite all data on the drive, resetting the data retention clock and ultimately improving overall data retention (even though it is slowly increasing the P/E cycles).

A computer running a typical Windows or Linux® OS will have enough disk activity from background tasks and logging to not require this kind of upkeep. However, an embedded OS with a write filter could benefit from a background task writing small amounts of data to a partition that is not protected by the write filter.

In the future, as more advanced features are incorporated into flash drive controllers, background integrity scans will be built into the controller and this type of upkeep will be automatic and completely transparent to the OS and application software. An example of this already in use is read-disturb monitoring. Whenever a flash cell is read, the charge of nearby flash cells in the same block or page can be disturbed, eventually causing data corruption. When the flash controller detects this corruption during a read command, its read-disturb feature will automatically rewrite that block of data. In this case, all that could be required for periodic data refreshing would be reading the entire logical contents of the drive on a slow periodic interval.

Monitor Status

Most flash drives that operate on fixed-disk interfaces (such as CompactFlash, CFast, and 2.5-inch solid-state drives) provide status information through the Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) interface. Although much of the information available through S.M.A.R.T. is vendor-specific, the interface usually provides the operating temperature of the drive and often includes data for the average number of P/E cycles. Some drives even provide a percent estimate of the remaining drive life. Also useful in predicting imminent drive failure is the spare block count, which typically starts to decrease rapidly as the drive nears the end of its useful life.

Status monitoring could be incorporated into the background task mentioned previously. The task could monitor the operating temperature and P/E cycle count on the drive to dynamically adjust the rate that the flash drive is written to, optimizing P/E cycles versus data retention at any temperature.

Overprovision

Probably the easiest way to maximize data integrity on any system using flash storage is to purchase a large-capacity flash drive. Due to the wear leveling feature built into the flash drive controller, a larger drive takes longer to wear out because the P/E cycles are spread across more flash cells. This results in both longer drive life and longer data retention. While buying the largest drive available is often not economically viable, buying a drive that is at least one size larger than necessary for the application is money well spent.

CONCLUSION

While flash memory storage has known weaknesses, it is still far superior in reliability and performance to what is currently the only alternative storage technology: rotating hard drives. Even though the initial cost of SLC flash memory is considerably higher than MLC, the long-term cost is significantly less. MLC flash-based drives will corrupt and lose data within the operational lifetime of an embedded computer system, even if the computer is not run in extreme temperatures. In the event of a drive failure, the cost to replace the drive, recover lost data, and reconfigure the computer typically exceeds the cost of an SLC flash-based drive. This is why SEL computers use SLC flash memory, as SEL will always strive to use the most reliable storage technologies available in order to minimize maintenance and total cost of ownership and to maximize reliability and availability.

REFERENCES

- [1] X-bit, "Flash and Solid State Drives: Facts and Forecasts," November 2010. Available: <http://www.xbitlabs.com/articles/storage/display/flash-ssd-forecast.html>.
- [2] JEDEC standard JESD218A, Solid State Drive (SSD) Requirements and Endurance Test Method.
- [3] JEDEC standard JEP122G, Failure Mechanisms and Models for Semiconductor Devices.

BIOGRAPHY

Ian Olson works as a lead application engineer for Schweitzer Engineering Laboratories, Inc. (SEL) in the computing systems group. Ian joined SEL in 2005 to support customers with communications, protocols, and system integration. Now, as a lead application engineer, Ian helps manage the growth, direction, and technology of computer products. Ian graduated from the University of Idaho in 2004 with a bachelor's degree in computer engineering.

© 2014 by Schweitzer Engineering Laboratories, Inc. All rights reserved.

All brand or product names appearing in this document are the trademark or registered trademark of their respective holders. No SEL trademarks may be used without written permission.

SEL products appearing in this document may be covered by US and Foreign patents.

SCHWEITZER ENGINEERING LABORATORIES, INC.

2350 NE Hopkins Court • Pullman, WA 99163-5603 USA
Tel: +1.509.332.1890 • Fax: +1.509.332.7990
www.selinc.com • info@selinc.com

