Technical white paper

Solid State Drive Technology

Differences between SLC, MLC and TLC NAND

Table of contents

Executive summary

This document outlines the differences between three different NAND technologies: SLC, MLC, and TLC.

- SLC - Single level cell provides high performance, low power consumption; faster write speeds, and higher program/erase cycles per cell, but at a higher cost. These devices are mainly used in high end server applications.
- MLC - Multi level cell provides a lower cost than SLC; however, MLC has a lower endurance limit and lower program/erase cycles than SLC, and is a good fit for commercial/workstations platforms and applications.
- TLC - Three level cell provides higher density, but has lower endurance limits, slower write read speeds, and lower program/erase cycles than SLC and MLC. To this point, TLC NAND has been used primarily in flash devices; however, technological advancements are making TLC an option for normal applications.

SLC vs MLC vs TLC

NAND cell technology

A single NAND flash die is typically subdivided into 512 MB blocks, and each block is typically subdivided into 4 KB pages. You can read and write to individual pages as long as they are empty; however, once a page is written, it can't be overwritten - it must be erased first before it can be written to again. This sequence of writing, erasing, and then writing again is known as the Program/Erase cycle. The P/E cycle can serve as a criterion for quantifying the endurance of a flash storage device.

Physically, all three NAND technologies consist of similar transistors; the only difference is that they store different amounts of charge. All three work similarily, applying voltage to a cell to go from an "off" state to an "on" state. SLC uses two distinct voltage states representing one bit per cell and two logic levels (0 and 1). MLC uses four distinct voltage states representing four logic states (00, 01, 10, 11) or two bits. TLC uses eight distinct voltage states representing eight logic states (000, 001, 010, 011, 100, 101, 110, and 111), or three bits.

Since SLC uses two voltage states, these states can be further apart, reducing the potential to incorrectly interpret the state of the cell and allowing for standard NAND error conditions, like read disturbs. The possibility for errors increases with TLC NAND, so it requires more ECC correction as the NAND wears out, because there are three bits to correct instead of one for SLC and two for MLC.

Write amplification

In write amplification, the actual amount of data written is a multiple of the logical amount intended to be written. Because flash memory must be erased before it can be rewritten, the process to perform these operations results in moving user data and metadata more than once. This multiplying effect increases the number of writes required over the life of the NAND, which shortens the time it can reliably operate. This increase in write consumes bandwidth and reduces random write performance.

Write endurance

All SSDs specify an endurance number that relates to the number of Program/Erase cycles that the underlying NAND array can accept. This endurance, defined in Total Bytes Written (TBW), specifies the number of host writes that can be written to the drive. In this endurance number, SSD vendors assume a workload, write amplification, and wear level efficiency that determine the actual P/E cycles allowed by the NAND array.

Table 1. Write endurance

	SLC	MLC	TLC
Bits per cell	1	2	3
P/E Cycles	100,000	3,000	1,000
Read Time	25us	50us	~75us
Program Time	200-300us	600-900us	~900-1350us
Erase Time	1.5-2ms	3ms	~4.5ms

SLC NAND flash typically rates endurance at about 100K cycles. MLC NAND flash is typically rated about 3K-5K cycles. TLC NAND flash is typically rated about 1K-3K cycles. Fewer available PE cycles make it important that TLC maintain low write amplification, because high write amplification will more quickly use up PE cycles.

Wear leveling algorithms and write amplification are never perfect. Providing a specific number for endurance is very difficult because every drive behaves differently and user workload varies. DPS (digital signal processing) solutions can help with lower endurance NANDs. Every time you program or erase the cell, you degrade it, and the DSP reads the voltage change and adapts to it. While a DSP won't make the NAND last indefinitely, it does prevent stress to the NAND by allowing it to last for more P/E cycles. If it can extend NAND life, it helps suppliers use smaller process nodes and possibly more bits per cell.

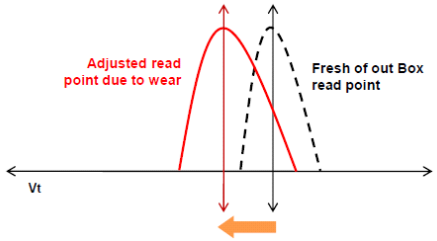

As higher density devices come online, a new coding technique known as Adaptive Endurance Code (AEC) can help increase the number of program/erase cycles that a flash device can endure. Data-dependant wear implies that flash device lifetime can significantly be increased by converting data into bit patterns prior to programming, which cause minimal wear. AEC can be adapted to data compressibility to maximize endurance gains with low system overhead cost.

Block copy time is the largest limiting factor for random write performance. As block copy times increases, random performance decreases.

Key factors that impact NAND Flash random write performance are as follows:

- Number of pages per blocks
- Increase of tPROG
- Increase in I/O transfer time

Figure 1: Block copy time

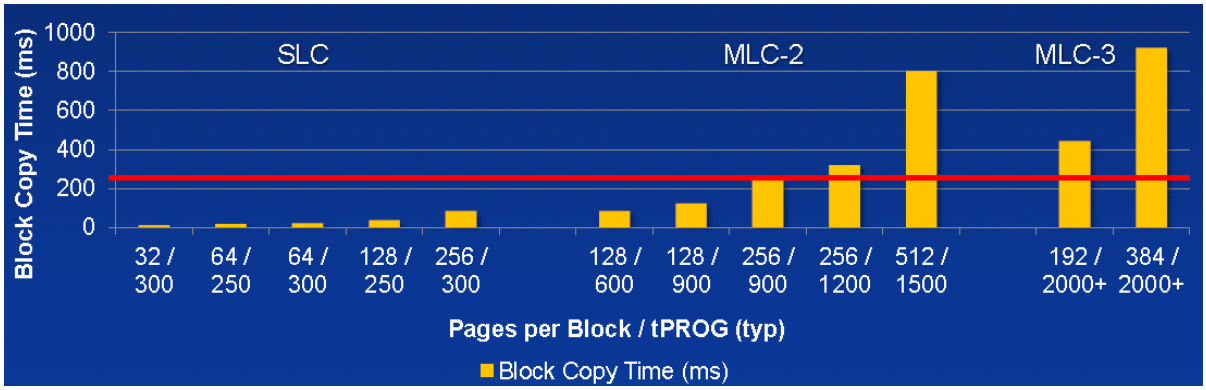

*note: tProg = program time

SMART technology

S.M.A.R.T (Self-Monitoring, Analysis and Reporting Technology) is standard in all storage devices. It monitors fault conditions and decreases the chance of data loss by giving adequate warning to these conditions. Some of the many attributes that are reported include re-allocated NAND block count and wear leveling count. However, SMART is not a replacement for good disaster recovery planning.

Cost per gigabyte

The cost per gigabyte of NAND varies greatly for a given lithography. The more the bits per cell, the cheaper the cost. SLC is the most expensive, while TLC is the least expensive.

Conclusion

Cost is always an important consideration, and TLC with its more complex voltage sensing capabilities helps to reduce the cost per gigabyte. This cost saving comes with some limitations, such as lower endurance and slower write performance. However, with the introduction of technologies such as DSP and AEC, device P/E cycle and endurance can be extended from original design capabilities.

References

- Adaptive Endurance coding for NAND Flash; {ashishja, franceschini, lastrasl,karidis}@us.ibm.com
- NAND Flash Architecture and Specification Trends; mabram@micron.com
- Understanding the Robustness of SSD under Power Faults; Mai Zang, Joseph Tucek, Feng Qin, Mark Lillibridge, The Ohio State Unitversity and HP Labs

For more information

For current HP driver, support, and security alerts delivered directly to your desktop, go to hp.com/go/getconnected.

To help us improve our documents, please provide feedback at hp.com/solutions/feedback.

Sign up for updates
hp.com/go/getupdated

736112-001, May 2013, Rev.1