

03


NAND Basics

Understanding the Technology Behind Your SSD

Although it may all look the same, all NAND is not created equal: SLC, 2-bit MLC, 3-bit MLC (also called TLC), synchronous, asynchronous, ONFI 1.0, ONFI 2.0, Toggle 1.0, Toggle 2.0. To the uninitiated, all of this looks like nonsense. As the SSD market gets more complex, however, so does your buying decision. This paper will help to lay down some of the differences between various types of NAND technology, with the ultimate goal of making you a savvy SSD shopper.

A Brief Introduction to NAND Flash

NAND Flash memory stores data in an array of memory cells made from floating-gate transistors. Insulated by an oxide layer are two gates, the Control Gate (CG, top) and the Floating Gate (FG, bottom). Electrons flow freely between the CG and the Channel (see diagram to the right) when a voltage is applied to either entity, attracted in the direction to which the voltage is applied. To program a cell, a voltage is applied at the CG, attracting electrons upwards. The floating gate, which is electrically isolated by an insulating layer, traps electrons as they pass through on their way to the CG. They can remain there for up to years at a time under normal operating conditions. To erase a cell, a voltage is applied at the opposite side (the Channel) while the CG is grounded, attracting electrons away from the floating gate and into the Channel.



To check the status of a cell, a high voltage is applied to the CG. If the floating gate holds a charge (electrons are trapped there), the threshold voltage of the cell is altered, affecting the signal emanating from the CG as it travels through to the Channel. The precise amount of current required to complete the circuit determines the state of the cell. All of this electrical activity effectively wears out the physical structure of the cell over time. Thus, each cell has a finite lifetime, measured in terms of Program/Erase (P/E) cycles and affected by both process geometry (manufacturing technique) and the number of bits stored in each cell. The complexity of NAND storage necessitates some extra management processes, including bad block management, wear leveling, garbage collection (GC), and Error Correcting Code (ECC), all of which is managed by the device firmware through the SSD controller.

SLC vs. 2-bit MLC vs. 3-bit MLC NAND

NAND technology has been naturally progressing with the needs and expertise of the industry. In the simplest terms, the data stored in NAND flash is represented by electrical charges that are stored in each NAND cell. The difference between Single-Level Cell (SLC) and Multi-Level Cell (MLC) NAND is in how many bits each NAND cell can store at one time. SLC NAND stores only 1 bit of data per cell. As their names imply, 2-bit MLC NAND stores 2 bits of data per cell and 3-bit MLC NAND stores 3 bits of data per cell.

Advantages of MLC NAND

The more bits a cell stores at one time, the more capacity that fits in one place, thus reducing manufacturing costs and increasing NAND manufacturing capacity – a phenomenon called “bit growth.” This phenomenon has allowed NAND technology to penetrate a continually greater number of usage applications at increasingly higher capacities over the years.

NAND technology’s first home was in external storage devices (e.g. USB memory devices) at very modest capacities. As the technology matured, NAND found applications in a host of digital devices, including digital cameras, MP3 players, and mobile phones. Having proven itself a strong and durable performer, the technology made its way into consumer and finally enterprise solid state storage devices (SSDs). NAND’s rise in popularity and usefulness was directly a result of the

advances in semiconductor technology that allowed engineers to squeeze more bits into each cell. Capacities ballooned from megabytes (MBs) to gigabytes (GBs) as manufacturers were able to produce the same NAND bit capacity with less capital investment. A self-reinforcing cycle of increasing demand and decreasing pricing helped manufacturers to continue to meet industry needs without increasing supply costs – benefiting both consumers and device makers.

Cost vs. Capacity

	SLC	2-bit MLC	3-bit MLC
Price/GB	\$\$\$\$	\$\$	\$
External Storage	Mid 90's~	Mid 2000s~	'09~
USB Flash Drive*	4MB~	1GB~	2GB~
Memory card*	2MB~	1GB~	?
Digital Devices	Mid 90's~	Mid 2000's	'11~
Digital Camera	4MB~	128MB~	
MP3 Player	16MB~	512MB~	2GB~
Client SSD (2.5")	'06~	'08~	'12~
	32GB~	64GB~	120GB~
Enterprise SSD	'08~	'12~	'13~
	50GB~	128GB	?

* Less than 5% of Memory Cards and USB Flash Drives use SLC Memory as of 2012

Limitations of MLC NAND

Of course, adding more bits to each cell makes it more difficult to distinguish between states, reducing reliability, endurance, and performance. Indeed, determining whether a container is either full or empty (SLC) is much simpler than determining whether it is one quarter full, one half full, three-quarters full, or entirely full (MLC). This is why it can take up to 4 times as long to write and up to 2.5 times as long to read 3-bit MLC NAND than its SLC predecessor.

Another side effect of storing more bits per cell is an increase in the rate at which the NAND cells degrade. The state of a NAND cell is determined by the number of electrons present on the floating gate. The Oxide layers that trap electrons on the floating gate wear out with every program and erase operation. As they wear out, electrons become trapped, which affects the overall electrical properties of the cell and, consequentially, subsequent program and erase operations. With the oxide weakened, charges sometimes leak from the floating gate. While this is not a huge problem with SLC NAND because there are only two states to distinguish between, it can be a huge problem for 3-bit MLC NAND because there are 8 states to distinguish and very little room for differentiation – just a few electrons can make the difference between one state or another. Compounding matters is the fact that the oxide layer gets smaller with each advance in process geometry – as we shrink the size of the cell, the oxide becomes thinner. A thinner oxide layer will wear out faster, meaning the cell will have a lower lifespan.

MLC NAND Today

Manufacturers have made massive strides in managing the limitations of NAND technology. Sophisticated “bin sorting” algorithms allow 1st-tier manufacturers like Samsung to select only the highest-quality NAND chips for use in SSD devices. Advanced wear-leveling code ensures that NAND cells wear out evenly (to prevent early drive failure and maintain consistent performance), while garbage collection algorithms preemptively prepare fresh storage space (to improve write performance). Improved Error-Correcting Code (ECC) is able to detect and recover from errors at the bit level caused by the natural wear out of individual NAND cells. Supporting all of these maintenance features is over-provisioning, which ensures that the SSD controller always has swap space available to accomplish its tasks. As a precautionary measure, some vendors choose to pre-set a certain amount of mandatory over-provisioning at the factory, and there is always the option to manually set aside additional space for even further-improved performance (e.g. under demanding workloads).

NAND Characteristics


	SLC	2-bit MLC	3-bit MLC
Bits per Cell	1	2	3
Performance	★★★★	★★★	★★
Endurance	★★★★	★★★	★★
ECC complexity	★★	★★★	★★★★

* Today's SSD controllers are capable of minimizing the risks associated with 2-bit MLC and 3-bit MLC NAND through techniques like advanced Error Correcting Code (ECC).

On top of all of the above advancements, the modern trend towards increasing storage densities has a convenient side benefit – with higher capacities also comes higher performance, because increasing the number of NAND chips in an SSD allows for more parallelization, helping to overcome the inherently slow program times that MLC NAND suffers compared to its SLC predecessor (this same phenomenon also explains why the same or similar NAND flash can deliver very different performance, lifespan, and reliability among various NAND-based devices). The additional free space has the added benefit of providing the controller with unused capacity to use as a kind of non-official over-provisioning space. Together, all of the progress SSD vendors have made in NAND technology has led to lower cost drives with performance, endurance, and reliability that can satisfy the demands of an ever-increasing audience.

Industry Trends

History is set to repeat itself thanks to the above and other continuous advances in NAND technology. The entire Information Technology and Consumer Electronics industries, which voraciously consume NAND chips for devices like MP3 players, smart phones, memory cards, USB flash drives, car navigation systems, and the other devices that make up this digital age, have benefited greatly from increased capacity at more favorable pricing. It is now time for PC storage, which is rapidly moving away from traditional Hard Disk Drive (HDD) technology in favor of SSDs, to enjoy similar benefits. The “bit growth” phenomenon discussed above is pushing the industry towards wider



3-bit MLC NAND adoption. In fact, we may soon see 3-bit MLC NAND-based products begin to dominate the SSD industry. In reality, 2-bit MLC NAND’s performance and lifetime characteristics are such that they far exceed the requirements of most computing tasks. There are even fewer applications that require SLC-level endurance and reliability these days. Thus, the real-world benefits are well worth any tradeoffs in performance and endurance we must make as we move towards denser memory technology. Today’s NAND-based storage will still far outlive the useful lifespan of the devices it powers. Choosing to sacrifice a bit of excess lifetime and performance in favor of dramatic cost benefits will allow manufacturers to deliver NAND technology and its many advantages over traditional storage technology to a far wider audience.

Samsung has taken the first step by introducing 3-bit MLC NAND to the SSD market with its 840 Series SSD, made possible by its fully integrated design approach, proprietary controller architecture and firmware algorithms, and superior NAND quality. As mentioned previously, increased SSD capacity leads to increased performance and endurance, so SSDs using this technology will only improve as densities grow. Thus, 3-bit MLC may represent the beginning of the next personal storage revolution.

Asynchronous vs. Synchronous NAND

At the heart of all electronic devices is the concept of a regulatory signal, which coordinates the rate at which instructions are executed and acts as a kind of conductor to keep everything in sync. Historically, NAND devices used two regulatory signals: the “RE signal” (for Read Events) and the “WE signal” (for Write Events). As NAND evolved and increased in speed, however, it became necessary to introduce a new signal, called a “strobe.” Present in all modern DDR NAND implementations, the strobe helps the controller to handle read and write data at high speed. It starts after receiving a signal from the host indicating a read or write event – think of a stopwatch, which can start and stop on demand. Depending on its current task, the strobe plays slightly different roles. During a write event, it is directly managed by the controller, whereas it plays a more supporting role during a read event.

With ONFI 2.0, synchronous logic was introduced to the ONFI NAND specification. Synchronous devices operate a “free-running” clock, which means that, as long they are receiving power, the clock will continuously run – much like a standard wall clock. This clock is used as a reference for the strobe. Modern ONFI NAND implementations eschew the use of a clock once again in favor of using only the strobe. Every generation of Toggle NAND was asynchronous.

Asynchronous




Synchronous




Traditional vs. DDR NAND

The advent of DDR NAND was a key breakthrough in increasing NAND speed. Traditional NAND, also known as “Single Data Rate (SDR)” NAND, was only capable of processing data on “one edge” of its regulatory signal. ONFI 1.0 and pre-Toggle Samsung and Toshiba NAND belong to this category. The figure below shows a Read operation for illustrative purposes.



In contrast, modern NAND can handle data on “both edges” of its regulatory signal (or strobe, as it is known). This concept is similar to how DDR RAM works (DDR stands for “Double Data Rate”).




Toggle DDR NAND works much the same way as the DDR NAND diagram above. The strobe that is used to coordinate the NAND operations has rhythmic rises and falls, and Toggle DDR NAND is capable of processing data on both the rise and fall.

Processing data on “both edges” of the strobe obviously results in significant speed and efficiency gains. The different generations of Toggle DDR and ONFI NAND are distinguished by how fast the strobe (or clock in the case of Synchronous ONFI NAND) itself runs (how fast it can complete one cycle).

Toggle NAND vs. ONFI NAND

In the NAND flash industry, there is some variation in implementation depending on manufacturer. Samsung and Toshiba make Toggle NAND, whereas everyone else makes what is called ONFI NAND.

Each generation of NAND, whether Toggle or ONFI, differs in the speeds it is capable of delivering: ONFI 1.0 and pre-Toggle NAND from Samsung and Toshiba (both asynchronous) are capable of speeds up to 50MB/s, ONFI 2.x (synchronous) and Toggle 1.0 NAND (asynchronous) are capable of delivering up to 200MB/s and 133MB/s respectively, and ONFI 3.0 and Toggle 2.0 (both asynchronous) are capable of delivering up to 400MB/s. Today’s newest NAND chips are either Toggle 2.0 or ONFI 3.x.



So why choose one standard over the other? The latest ONFI NAND chips, because ONFI has alternated between synchronous and asynchronous methods, must be built with backwards compatibility for both synchronous and asynchronous logic. Because Toggle NAND has always been asynchronous, Toggle 2.0 NAND benefits from reduced complexity and, in turn, increased flexibility when it comes to designing smaller chip sizes and bringing next-generation products to market faster. Furthermore, because of Samsung and Toshiba's positions as market leaders in the NAND industry, a majority of the market has standardized on Toggle NAND.

NAND Type & Performance

How do speed differences among the various generations of individual NAND chips affect the overall speed of an SSD? And what does this mean for you in terms of real-world impact?

In practice, the use of synchronous versus asynchronous NAND has no impact on SSD performance. What actually affects performance is the generation of the NAND (e.g. SDR/ONFI 1.0, Toggle 1.0/ONFI 2.0, etc.). When discussing NAND type, therefore, it is inaccurate to refer to it simply as either "synchronous" or "asynchronous" when referring to performance.

Why Samsung?

As the #1 player in the worldwide memory market for over 20 years, no one has more experience manufacturing NAND than Samsung. Furthermore, Samsung's fully-integrated manufacturing approach means that it has intimate knowledge of every nuance of each component of your SSD. While generic controller manufacturers must optimize their chips to work with both ONFI and Toggle NAND, Samsung can focus all of its design efforts on making Toggle NAND work perfectly with its proprietary controller technology. The end result is product characterized by awesome performance and unrivaled reliability – a product only Samsung could build.
