

NAND 201: An Update on the Continued Evolution of NAND Flash

Jim Cooke

Sr. Technical Marketing Manager

Micron Technology, Inc.
September 6, 2011

A lot has changed with NAND Flash memory since my original [NAND 101 article](#) was published by EETimes in 2006. From the evolutionary changes of a continually shrinking NAND cell, to the performance-enhancing innovations that support increasingly advanced designs, this follow-on article chronicles the developments in NAND technology from 2006 through early 2011.

Market Changes

In 2006, single-level cell (SLC) NAND Flash devices were mainstream products that accounted for more than 80% of the devices on the market. At that time, many NAND Flash vendors were struggling with two bits per cell, known as multilevel cell (MLC), and SLC device densities were in the range of just a few gigabits. Today, Micron offers a range of NAND products with densities up to 512Gb in a single device.

Figure 1 shows the past, present, and projected output mix for the major NAND cell technologies. The high runners are all MLC devices, which have replaced SLC devices in holding approximately 80% of the total market. While 4-bit-per-cell technology (16 levels) grew to a few percent in 2010, it is expected to drop to 0% in 2011 due to the continued difficulty of reliably placing 16 discrete thresholds in a single cell.


Three-bit-per-cell devices (eight levels), are expected to grow from less than 10% in early 2009 to almost 30% by the end of 2011. These devices are used primarily in value-minded consumer products that can operate with lower NAND performance and fewer PROGRAM/ERASE cycles (also known as endurance).

Traditional MLC devices, which group two bits per cell (four levels), are ideal for applications that demand medium levels of higher performance and endurance; thus, MLC drives the majority of NAND output.

Lastly, SLC NAND is the technology of choice for high-performance, high-endurance, and high-reliability applications.

Other specialized NAND devices have been born out of necessity, and we will discuss those as well.






Source: Forward Insights Q4/10

Figure 1: NAND Production by Type

Figure 2 shows the expected application adoption of 3-bit-per-cell NAND devices. In addition to early drivers such as Flash cards and USB thumb drives, several other consumer application designs are expected to be based on this technology.



Source: Forward Insights Q4/10

Figure 2: 3-bit/cell Density Adoptions by Application



NAND Basic Operations

While the NAND cell itself has remained essentially the same over the last several years (albeit much smaller), almost everything else about NAND has changed (see Table 1).

Array Enhancements

As the NAND cell has continued to shrink, the array has become significantly more compact, allowing more cells to be integrated in a comparable die area.

The mainstream device in 2006 (Figure 3) was a 2Gb SLC device consisting of 2048 blocks; each block had 64 pages, and each page had 2112 bytes. In contrast, today's high-capacity devices (Figure 4) can include multiple die or logical unit numbers (LUNs). Each 32Gb monolithic SLC die has 4096 blocks, each block contains 128 pages, and each page has up to 8640 bytes—and it requires only a slightly larger silicon area than the 2Gb device of 2006. When the spare area needed to support the additional error correction code (ECC) is factored in, the result is more than 16 times the density.




Figure 3: 2Gb Array Organization of 90nm Micron SLC NAND Flash (2006)

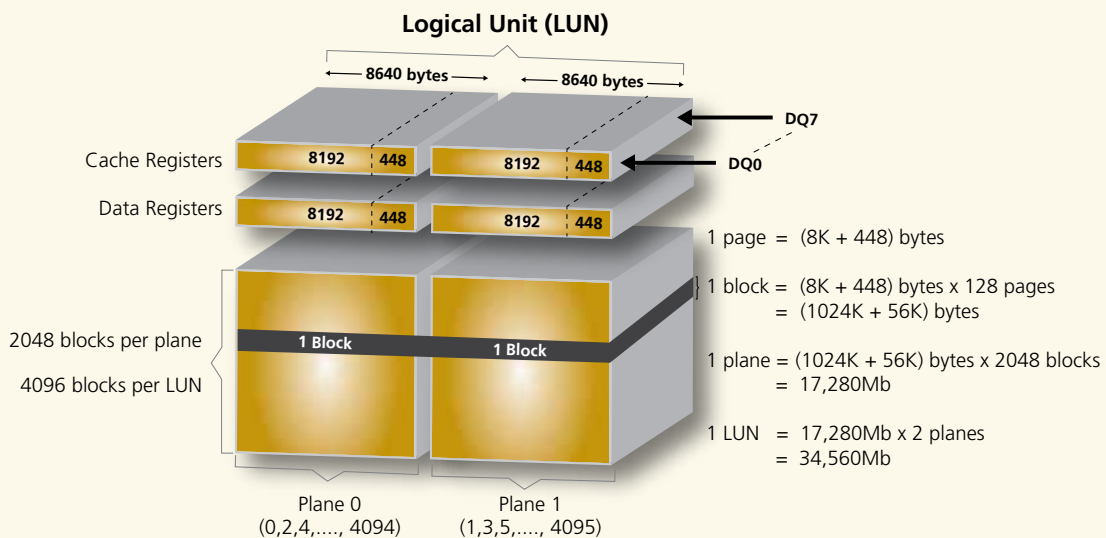


Figure 4: 32Gb Array Organization of 25nm Micron SLC NAND Flash (2010)



Table 1 shows the evolution of SLC and MLC devices. Focusing first on the SLC NAND, the programming time (t_{PROG}) has remained relatively constant at approximately 300 μ s, but the amount of data being programmed (the page size) has increased from 2112 bytes in 2006 to 8640 bytes for the latest 25nm production device. This increased page size provides a four-fold improvement in the overall array programming performance.

To improve performance even further, most NAND manufacturers have implemented what is referred to as multiplane technology, shown in Figure 4. Essentially, multiplane operations enable twice the number of array data operations in approximately the same amount of time. The details of these commands will be discussed later, but examples include multiplane ERASE or PROGRAM operations, where the array times are significant and can be amortized over double the amount of data, improving performance by almost two times. The SLC devices of 2006, by comparison, were single-plane, as shown in Figure 3.

	2006	2007	2008	2009	2010
Characteristics					
Monolithic Die Density	2Gb	8Gb	8Gb/ 16Gb	16Gb/ 32Gb	32Gb/ 64Gb
Geometry (nm)	90	72	50	34	25
t_{PROG} (μ s TYP)	~300	~700	~300/ ~900	~300/ ~900	~300/ ~1300
Page Size (bytes)	2112	2112	4314	4320	8640
Number of Planes	1	2	2	2	2
Number of Pages (SLC/MLC)	64	128	64/ 128	128/ 256	128/ 256
Block Size (KB)	128	256	256/ 512	512/ 1024	1024/ 2048
t_R (μ s)	25	50	25/ 50	25/ 50	35/ 75
ECC Required/ Code Word	1/528	4/528	4/539, 8/539	4/540, 12/540	8/540, 24/1080
NOP (Number of Partial-Page Operations)	8	1	4/ 1	4/ 1	4/ 1
Endurance	100,000	10,000	100,000/ 10,000	100,000/ 5000	60,000/ 3000
Interface (MTIs)	SDR 40	SDR 50	SDR 50	DDR 166	DDR 200
Packages	TSOP	TSOP	TSOP	TSOP, LGA, BGA*	TSOP, LGA, BGA*
Temperature Range	Commercial	Commercial	Commercial	Commercial, Industrial	Commercial, Industrial
Notes: Timeline years are approximate. Values specific to MLC devices are in bold . Shaded columns identify the NAND devices used for most of our comparisons. *Up to 8-die BGA packages offered.					

Table 1: Evolution of High-Density Micron SLC/MLC Devices



Doubling the number of SLC pages from 64 to 128 in 2009 enabled a more efficient array but produced the negative side effect of a larger block size, which can present management challenges for the Flash file system, especially when dealing with small data files.

The array read time (t_R) has remained very fast and relatively consistent, resulting in high performance for most READ operations.

The required ECC, NOP, and endurance specifications referenced in Table 1 are discussed as a group because they are all related to the robustness of the NAND cell. As NAND continues to scale, it becomes more challenging for cells to store data reliably. To compensate, the amount of ECC required continues to increase. Early SLC NAND devices were able to achieve an endurance rate of 100,000 with only one bit of ECC. These same devices were also able to support a NOP of eight, but that was reduced to four to match the requirement of the original 2112-byte, 64-page devices, which typically consisted of four 528-byte sectors. To remain competitive, NAND must shrink to smaller geometries while maintaining an acceptable system error rate, which requires a combination of higher levels of ECC, fewer NOPs, and/or lower endurance.

NAND endurance has received some negative attention over the past few years. But endurance by itself doesn't tell the whole story. Effective wear leveling (spreading the erasing and programming [or cycling] evenly over the entire device), together with the higher densities of today's NAND, can greatly extend system endurance.

Consider MLC devices from 2008 and 2010 as examples. Density increased by 400%, while cell endurance only decreased 70% (from 10,000 to 3,000 cycles). The net result is actually better at the system level.

Note that the lower cycle rates associated with advanced NAND do not equate to low device endurance, but they do necessitate closer coupling between the NAND and controller to ensure optimized wear leveling algorithms.

Interface Enhancements

Another significant area of NAND Flash evolution is the interface. The Open NAND Flash Interface (ONFI) workgroup was founded in February 2006, and in December 2006 the group published its first specification, ONFI 1.0. Prior to ONFI, controllers or host firmware could issue a READ ID command to the NAND device, which would return manufacturer and design ID information, in addition to a few vendor-specific bytes. While better than nothing, the lack of device parameter information forced controller vendors to build internal firmware tables by manually entering data sheet parameters. These tables were then used by the firmware to interpret device capabilities.

One of the more significant outcomes of the ONFI 1.0 specification was the ability for controller firmware to ascertain the parameters and characteristics of the attached NAND Flash. ONFI 1.0 simplified support for the various timing specifications through timing modes. By reporting the supported timing modes, the NAND device could meet the minimum specifications, simplifying more than 30 various timing details into one of six preconfigured timing mode groupings.

In addition, ONFI 1.0 provides a parameter page data structure that controller firmware can easily read to request further details about the Flash, such as page size, number of blocks, ECC required, etc. (All released ONFI specifications are publically available at onfi.org.)

As the page size grows, the need for speed becomes more apparent. Using the original 2112-byte page device as an example, the fastest asynchronous clock specification of 20ns allows us to shift the entire 2112-byte page in less than 43 μ s. As the page size grows to 4320 or 8640 bytes, this time increases to 86 μ s and 172 μ s, respectively. With the SLC array read time (t_R) generally fixed at 25–35 μ s (depending on the geometry), the I/O bottleneck quickly becomes apparent. On a 25nm SLC device, it takes only 35 μ s to transfer all 8640 bytes from the array to the I/O register, but it takes an additional 172 μ s to get all 8640 bytes of data out of the register. This results in a total time of 207 μ s, or approximately 42 MB/s, not including the minimal command overhead.



Utilizing the source-synchronous ONFI 2.0 interface running a 10ns clock (DDR) would result in an I/O access rate of 200 MB/s. (ONFI 3.0 would result in 400 MB/s.) For this same example, the time to get 8640 bytes of data out of the register would be reduced from 172µs to 43µs. The array read time would remain 35µs. Adding the array read time of 35µs and the faster I/O time of 43µs results in a total time of 78µs, or approximately 110 MB/s.

Note that in order to keep these examples simple, we are considering only the single-plane read bandwidth from the NAND to the controller and are not using any cache or multiplane operations, which could provide a significant performance improvement. Also keep in mind that the controller typically has some inefficiencies, as well as ECC correction times that must be taken into account.

Command Enhancements


Table 2 shows the commands supported by Micron's 25nm NAND Flash. (Please note that this table is an example provided for general reference; check the device data sheet for specific implementation details. As this article is intended to discuss high-level evolutionary changes, commands will not be covered in detail.)

For easy comparison, the 13 basic commands supported by the 2006 2Gb device are shaded. These 13 commands have remained compatible with newer Flash devices.

The second column specifies whether the command is optional or mandatory (O/M) for NAND Flash, as required by ONFI.

Most high-density NAND packages consist of several die, often as many as eight. To minimize control signals, several die (usually two) typically share a single chip enable (CE). ONFI calls a single CE a target. Addressing several die on a single CE is accomplished via a logical unit number or LUN. The LUN is typically the next address bit adjacent to the most significant block address.

When this concept of multiple die (LUNs) sharing a single CE was introduced, it was generally understood that additional information would be needed for the system to take full advantage of it. The second-to-the-last column in Table 2 indicates whether the command on each row can be issued when other LUNs are busy—information that is crucial in order to fully leverage the NAND's ability to interleave or pipeline operations among several die in a single package. A typical example would be to pipeline the data loading and programming operations among all die in a particular package because effective pipelining enables maximum performance. Figures 5A and 5B show the maximum sequential performance that can be achieved using the ONFI 2.2 source-synchronous interface, assuming the timing parameters of Table 1 for the 2010 25nm device.



*Maximum sequential performance assuming no controller overhead.

Figure 5A: 25nm Micron MLC Performance




Figure 5B: 25nm Micron SLC Performance



NAND 201: An Update on the Continued Evolution of NAND Flash

Command	O/M	1st Cycle	# of Valid Address Cycles	Data Input Cycles	2nd Cycle	Acceptable While Accessed LUN Is Busy	Acceptable While Other LUNs Are Busy	Target-Level Commands
READ	M	00h	5	–	30h	–	Y	–
Multiplane	O	00h	5	–	32h	–	Y	–
COPYBACK READ	O	00h	5	–	35h	–	Y	–
CHANGE READ COLUMN	M	05h	2	–	E0h	–	Y	–
CHANGE READ COLUMN ENHANCED	O	06h	5	–	E0h	–	Y	–
READ CACHE RANDOM	O	00h	5	–	31h	–	Y	–
READ CACHE SEQUENTIAL	O	31h	0	–	–	–	Y	–
READ CACHE END	O	3Fh	0	–	–	–	Y	–
BLOCK ERASE	M	60h	3	–	D0h	–	Y	–
Multiplane	O	60h	3	–	D1h	–	Y	–
READ STATUS	M	70h	0	–	–	Y	Y	–
READ STATUS ENHANCED	O	78h	3	–	–	Y	Y	–
PAGE PROGRAM	M	80h	5	Yes	10h	–	Y	–
Multiplane	O	80h	5	Yes	11h	–	Y	–
PAGE CACHE PROGRAM	O	80h	5	Yes	15h	–	Y	–
COPYBACK PROGRAM	O	85h	5	Optional	10h	–	Y	–
Multiplane	O	85h	5	Optional	11h	–	Y	–
SMALL DATA MOVE ¹	O	85h	5	Optional	11h	–	Y	–
CHANGE WRITE COLUMN ²	M	85h	2	Optional	–	–	Y	–
CHANGE ROW ADDRESS ²	O	85h	5	Optional	–	–	Y	–
READ ID	M	90h	1	–	–	–	–	Y
READ PARAMETER PAGE	M	ECh	1	–	–	–	–	Y
READ UNIQUE ID	O	EDh	1	–	–	–	–	Y
GET FEATURES	O	EEh	1	–	–	–	–	Y
SET FEATURES	O	EFh	1	4	–	–	–	Y
RESET LUN	O	FAh	3	–	–	Y	Y	–
SYNCHRONOUS RESET	O	FCh	1	–	–	Y	Y	Y
RESET	M	FFh	0	–	–	Y	Y	Y

Notes:
1. The first opcode for SMALL DATA MOVE may be 80h if the operation is a program only, with no data output. The last second cycle of a SMALL DATA MOVE is a 10h command that confirms the PROGRAM or COPYBACK operation.
2. CHANGE WRITE COLUMN specifies the column address only. CHANGE ROW ADDRESS specifies the row address and the column address. Refer to the specific command definitions.

Table 2: Historical Comparison of NAND Commands



Multiplane Commands

Many of the new commands in Table 2 are multiplane commands, which make better use of the physical NAND array by providing concurrent ERASE, PROGRAM, or READ operations across multiple planes. These concurrent operations must be of the same type; for example, you cannot mix an ERASE operation on one plane with a PROGRAM operation to another plane.

Some new commands in Table 2 include:

- **CHANGE READ COLUMN ENHANCED:** Provides the ability to specify the complete 5-byte address, which is useful when interleaving between multiple die or multiple planes.
- **READ CACHE RANDOM:** Expands the sequential cache capabilities by adding the ability to select the next nonsequential page from either plane by specifying a 5-byte address.
- **CHANGE ROW ADDRESS:** Extremely useful for some applications, like Secure Digital (SD) or MultiMedia Card (MMC), in which the host can simply terminate a data transfer. Issuing a CHANGE ROW ADDRESS command after data has already been loaded into the NAND data register allows storage of the data fragment into a block reserved for these fragments, providing more efficient handling.
- **READ PARAMETER PAGE:** Introduced in the ONFI 1.0 specification, the parameter page is a convenient way to ascertain all the capabilities and characteristics of the particular NAND device. This command is accepted by the target only when all die (LUNs) on the target are idle.
- **READ UNIQUE ID:** Returns a unique 16-byte serial number, which can be useful for implementing security features.
- **GET FEATURES/SET FEATURES:** Allows you to read the current setting via GET FEATURES or to change the setting via SET FEATURES. Examples include timing modes, output drive strength, and array operations, such as OTP (one-time program).
- **RESET LUN:** Used to put a particular LUN (or die) on a target (CE) into a known condition and to abort command sequences in progress. This command is accepted only by the LUN addressed by the RESET LUN (FAh) command, even when that LUN is busy.

- **SYNCHRONOUS RESET:** When the synchronous interface is active, the SYNCHRONOUS RESET (FCh) command is used to put a target into a known condition and to abort command sequences in progress. This command is accepted by all die (LUNs), even when they are BUSY.
- **SMALL DATA MOVE:** Useful for small controllers that do not have enough RAM to buffer the entire page. This command allows the controller to access the data in pieces that are smaller than the full page size.

Addressing Enhancements

Like the NAND array and interface, addressing has also followed an evolutionary migration, as shown in Tables 3 and 4. The first two bytes (the column address) have grown from 12 bits to 14 bits to accommodate the larger 8640-byte page size. Another obvious difference in the two tables is the bit naming in the third address byte. While the six least significant bits (LSBs) of Table 3 specify which of the 64 pages are being addressed, they were renamed for clarity in the transition to the newer devices. Address bytes 3–5 specify the page and block addresses, which have grown to 7 and 12, respectively, providing the ability to address the 128 pages and the 4096 blocks of the 25nm SLC NAND device. As mentioned previously, the LSB address (BA7) is used to determine which plane is selected.

Addressing Differences Between SLC and MLC

Essentially, an MLC device has the same page layout as its SLC counterpart, except with twice the number of pages per block. If Table 4 were describing an MLC device storing two bits per cell, it would have an additional page address bit (PA7) as the most significant bit (MSB) of the third address byte. The LSB of the block address would move to the fourth address byte, and all remaining block address bits would be shifted left accordingly.

The remaining address bit (LA0) in the fifth address byte is used to select which die (LUN) is enabled when two die share a single CE. Future NAND devices may have more than two die on a single CE. To specify the die address, additional LA (LUN addresses) would be required.

The bold specifications in Table 1 show the major timing differences for MLC, which results in the lower performance numbers of these higher-density devices.



Cycle	I/O7	I/O6	I/O5	I/O4	I/O3	I/O2	I/O1	I/O0
First	CA7	CA6	CA5	CA4	CA3	CA2	CA1	CA0
Second	LOW	LOW	LOW	LOW	CA11	CA10	CA9	CA8
Third	RA19	RA18	RA17	RA16	RA15	RA14	RA13	RA12
Fourth	RA27	RA26	RA25	RA24	RA23	RA22	RA21	RA20
Fifth	LOW	LOW	LOW	LOW	LOW	LOW	LOW	RA28
Notes: CAx = column address; RAx = row address								

Table 3: Addressing for 90nm Micron 2Gb SLC NAND Flash (2006)

Cycle	DQ7	DQ6	DQ5	DQ4	DQ3	DQ2	DQ1	DQ0
First	CA7	CA6	CA5	CA4	CA3	CA2	CA1	CA0
Second	LOW	LOW	CA13	CA12	CA11	CA10	CA9	CA8
Third	BA7	PA6	PA5	PA4	PA3	PA2	PA1	PA0
Fourth	BA15	BA14	BA13	BA12	BA11	BA10	BA9	BA8
Fifth	LOW	LOW	LOW	LOW	LA0	BA18	BA17	BA16

Table 4: Addressing for 25nm Micron 32Gb SLC NAND Flash (2010)

Package Enhancements

The predominant package in 2006 was the 48-lead TSOP (12 x 20mm), but demand for more efficient packages has since led to the adoption of the land grid array (LGA) package. The LGA package is similar to a traditional ball grid array (BGA) package but without the balls. The 14 x 18mm, 52-land LGA package has become commonplace in many consumer products. For improved package and solder joint reliability, many designers prefer the 100-ball BGA package.

All three packages are offered with up to eight die, in a variety of configurations. Both the LGA and BGA multidiel packages can be offered with two completely independent interfaces. Having separate interfaces allows users to concatenate the outputs, yielding a 16-bit-wide interface that is attractive for applications requiring higher performance. If performance is not critical, users typically tie the two 8-bit buses together to save I/O pins on the controller.

The Future of NAND Devices

Evolution of Managed Devices

In order to maintain system data reliability levels, higher levels of ECC and advanced signal-processing algorithms will be needed in the near future. These requirements will push a larger and larger share of the market toward the use of partially and fully managed devices.

Fully Managed Devices

Near the end of 2006, the MultiMedia Card Association (MMCA) agreed to work with JEDEC to create an embedded version of their card specification. The result was the initial publication of the e•MMC™ specification version 4.1 in June 2007; the specification has undergone several revisions since the initial publication.



Basically, e•MMC devices are similar to many other fully managed devices in that they provide a simple, block-oriented interface to the user. While this simplicity may be a good solution for many implementations that are trying to get products to market quickly, e•MMC's slower-performing, single-threaded interface isn't necessarily the best choice for high-performance applications. In addition, e•MMC devices are not as deterministic as raw NAND devices because the controller may be busy with garbage collection or other time-consuming operations. If the controller is busy with these operations when the request for data is issued, it can have a big impact on performance.

Partially Managed Devices

The issue most customers object to when interfacing directly to NAND Flash is the ever-increasing amount of ECC required. Supporting the additional ECC generally requires a new version of the customer's controller chip, a migration that is both expensive and time-consuming. Many customers prefer to handle the required block management, as it helps them differentiate their products and allows them full control over the NAND, providing higher levels of performance through the use of multithreaded techniques.

In response to the needs of these customers, Micron is offering a new family of products called ClearNAND™ devices. ClearNAND devices are offered in similar LGA and BGA packages, with signals that are compatible to the footprint of traditional NAND. ClearNAND devices utilize the same ONFI asynchronous and synchronous interfaces that users are accustomed to and include a thin controller to manage the ECC and data integrity challenges of NAND. More details on these ClearNAND devices can be found at micron.com/cleannand.

For several years, pundits have been projecting that NAND will be unable to scale reliably, but so far, NAND vendors have proven them wrong. However, it is clear that even more ECC and other advanced algorithms will be required to continue the migration to smaller and smaller cell geometries.

In the near future, interfacing to raw NAND Flash will be an option for only a minority of the market, the very large customers who are almost as familiar with NAND Flash as the manufacturers are. ClearNAND technology will become the solution for the majority of the market, for those customers who cannot or choose not to deal with the ever-increasing challenges of NAND, yet still want to interface with NAND directly to achieve the higher performance that the architecture can provide. ClearNAND Flash devices will enable these mainstream customers to take advantage of the high densities and attractive costs of finer geometries, while leaving the complexities (such as ECC) to the NAND manufacturers.

Growth of Tailored and Specialty Devices

In 2011, the worldwide NAND market is expected to grow to more than \$20 billion and continue to drive the development of NAND Flash tailored for specific markets. Perhaps the best examples of a tailored NAND device are Enterprise NAND™ devices, unique products that increase endurance by up to six times that of standard MLC NAND. Many enterprise customers welcome the trade-off of higher endurance versus slower erase and programming timing characteristics.

Conclusion

With the continued growth of the NAND market and the remaining potential for cell shrinks, it is clear that NAND Flash will persist in its evolution by offering ways to reduce controller signals, increase speed, and provide better management of the NAND cell. The best way to stay informed about these technical proposals is to join onfi.org.

micron.com

Products are warranted only to meet Micron's production data sheet specifications. Products and specifications are subject to change without notice.

Micron, the Micron logo, and ClearNAND™ are trademarks of Micron Technology, Inc. e-MMC™ is a trademark of the JEDEC Solid State Technology Association. All other trademarks are the property of their respective owners. ©2011 Micron Technology, Inc. All rights reserved. 09/06/11 EN.L

