

NAND

FLASH

Mr. NAND's Wild Ride:
Warning — Surprises Ahead

Robert Pierce
Senior Director, Flash Products
Denali Software, Inc.



On May 21, 2005 Six Flags Great Adventure introduced Kingda Ka—the tallest, fastest roller coaster on the planet—to the public. This \$25-million, Swiss-designed “rocket coaster” uses hydraulic motors to launch the trains along a horizontal section of track from zero to 128 mph in an impressive 3.5 seconds!

The train then begins a vertical ascent up a steel tower that peaks at 456 feet or 45 stories. Crossing over the apex the train enters a vertical descent plunging through a 270-degree spiral twist again reaching speeds in excess of 100 mph. Kingda Ka offers one final surprise before the brake run, a 129-foot tall camelback hill, that offers plenty of negative G’s also known fondly as “airtime”.¹

Just as the Kingda Ka roller coaster at Six Flags Great Adventure in Jackson, New Jersey started plunging coaster fans straight down that first big, spiraling, 418-foot drop, NAND Flash pricing started its own precipitous drop—a long, long drop that’s lasted nearly five years. It’s proven to be a wild, terrific ride for system and subsystem developers who have used the opportunity to harness cheap, non-volatile storage for several exciting new uses, including SSDs and NAND Flash caches that fully exploit the new low-cost-per-bit leader in semiconductor memory. But just as there are several surprises waiting at the bottom of the big Kingda Ka roller coaster drop, there are several big surprises waiting for NAND Flash users in the immediate future.

If you use or depend on NAND Flash memory, if you must stay on top of the changes happening to NAND Flash standards and conventions and the altered direction for NAND Flash pricing due to the technical consequences of lithographic advances and manufacturing realities, then this paper is written for you.

The End of One Ride; Starting the Next

NAND Flash manufacturing cost reductions of 60% per year sustained over nearly a decade have driven many technology changes, developments, compromises, and innovations. Prices have fallen even faster over the past five years, but the precipitous price decline could easily slow due to technical forces and in all likelihood, they will. Further, NAND Flash specifications are changing and will continue to change in predictable and unpredictable ways due to these forces. These changes are taking place throughout the industry, not with just one manufacturer’s NAND Flash devices. These changes will create new capabilities for NAND users, will impose extra performance burdens, and may ultimately limit the flexibility of NAND Flash in future device generations compared to what is available today.

This paper describes several of these trends in a series of warnings and enumerates the steps the semiconductor industry is taking to smooth product transition for NAND Flash users. The paper also highlights developments that may create problems for NAND Flash users in the immediate future. NAND Flash trends discussed in this paper include:

- Page-size trends and their impact on NAND devices and controllers
- Changes to the spare area on NAND Flash devices
- Changing page architectures and ways to use the new architectures
- Increasing bit error rates and the associated rise in error correction code (ECC) complexity

¹ http://www.ultimaterollercoaster.com/coasters/yellowpages/coasters/kingdaka_sfgadv.shtml

NAND Flash manufacturers employ a broad variety of technology roadmaps, semiconductor implementation methodologies, and cost/performance optimization-and-tradeoff strategies to develop and manufacture NAND Flash devices. Diverse applications for NAND Flash memories drive similarly diverse performance specifications and requirements. “One Technology Roadmap Does Not Fit All” is an important guiding principle here. Consequently, products from some NAND Flash vendors may take somewhat different approaches to those described in this paper even though the vendors are responding to the same technical and market pressures considered below.

NAND-Flash controller designers may find the warnings and comments in this paper useful as they analyze important trends in NAND Flash management. In particular, advanced NAND Flash devices need extensive software management and error-correcting methodologies to create fully-functional memory subsystems that provide error-free data storage and they’ll need even more management and better ECC methodologies in the near future.

The three basic NAND management methodologies currently applied to NAND Flash memory subsystems are:

- ECC circuitry
- Bad-block management, and
- Wear leveling

All three are implemented either as a dedicated IC or as an IP block in a system-on-chip (SoC), and often paired with NAND management software. No matter the approach used to solve the above mentioned items the architectural and technology changes to future NAND Flash devices require increasingly comprehensive changes to the error-correction techniques and sophistication of the NAND policies.

Warning 1: Super-Sized Pages

Flash memory cards and USB drives are the volume leaders and users of NAND Flash, so their particular needs tend to drive many NAND Flash design decisions at the device level. For example, NAND Flash devices require the largest possible page size to achieve maximum read and write throughput. Files stored on USB sticks and SD cards are typically large files including: documents and presentations, still images, video, audio, and music files, and binary programs. All of these files tend to run at least 5 to 10 Mbytes and often considerably more. Because NAND Flash devices read and write data a page at a time, the trend to big files creates pressure on device designers to increase NAND Flash page sizes so that it takes fewer pages to store the files. Writing fewer pages boosts the NAND Flash throughput.

Consequently, NAND Flash page sizes have been steadily growing over the past five years as device designers strive to increase the block-transfer size and throughput rates while improving device layout. NAND Flash page size was just 256 bytes per page five years ago. Page sizes have ballooned to 8 Kbytes per page and 16-Kbyte page sizes will likely appear within another 2 to 3 years. **Table 1** shows Denali’s forecast for NAND Flash page sizes over the next few years, based on extensive conversations held with NAND Flash vendors.

Production Volume Level	2008	2009	2010	2011	2012 (projected)
Low	2KB	2KB	2KB	4KB	4KB
Medium	2KB	4KB	4KB	4KB	8KB

Table 1: Page Size production volume transition over time

In theory, larger NAND Flash pages increase the amount of data transferred per unit of time. The improved throughput is why USB sticks and SD cards drive this trend. End customers are impatient and want things done quickly, so fast NAND Flash drives sell somewhat better than slow ones.

However, certain applications buck this trend. For example, file and operating system (OS) transfer requests may not be aligned nicely on NAND Flash page boundaries resulting in incomplete transfers and causing the NAND Flash device to terminate the in-progress data transfer. This sort of problem is especially common for older file systems that don't understand NAND Flash devices and even occurs with some NAND Flash controllers.

SSDs and other mass-storage devices, which are tied closely to the OS, benefit from NAND Flash devices that retain the existing 2K or 4K page sizes because of these boundary-alignment problems. The smaller page sizes actually improve the I/O Operations per Second (IOPs) performance of such storage subsystems. For these mass-storage subsystems, a better way to achieve high throughput is to spread large file transfers across multiple NAND Flash devices with smaller pages. (Note: You can request a copy of Denali's MemCon 09 tutorial on 4K and 8K page size performance at www.denali.com).

Warning 2: Uneven Spare Areas

A NAND Flash device's spare area is primarily used to store the error-correction syndrome (the ECC bits) and redundancy check bits. The amount of spare area used directly depends on how the file system partitions the NAND Flash page. Dependencies include ECC sector size, page size, and system or file system usage.

The relatively small spare area associated with a page of NAND Flash can be consumed very quickly. Table 2 highlights the ECC requirements for the spare area by listing the number of bytes required based on the number of correction bits and the size of the ECC sector. Each row in **Table 2** corresponds to the ECC sector size (in bytes) and each column reflects the number of ECC bits required. The box at the intersection of each row and column in the table shows the number of spare area bytes required per sector. Multiplying the numbers in **Table 2** by the number of sectors per page gives the total number of spare-area bytes required for ECC for each page.

Page Size (Bytes)	# of Correction Bits															
	1	2	3	4	5	6	7	8	10	14	16	18	20	22	24	26
128	1	3	4	6	7	8	10	11	14	19	22	25	28	30	33	36
256	2	3	5	6	8	9	11	12	15	21	24	27	30	33	36	39
512	2	3	5	7	8	10	11	13	16	23	26	29	33	36	39	42
1024	2	4	5	7	9	11	12	14	18	25	28	32	35	39	42	46
2048	2	4	6	8	9	11	13	15	19	26	30	34	38	41	45	49
4096	2	4	6	8	10	12	14	16	20	28	32	36	40	44	48	52
8192	2	4	6	9	11	13	15	17	21	30	34	38	43	47	51	55

Table 2: ECC Spare Area requirement by page size

Table 3 displays a combination of correction bytes and the amount of spare area needed to support the different page and ECC-sector-size options. NAND Flash spare-area sizes vary both by manufacturer and among NAND Flash products from a single manufacturer. Depending on the device and manufacturer, a NAND Flash device with 4-Kbyte page size will have between 208 and 214 bytes of spare area at the 3X nm process node while a device with 8-Kbyte pages will have 420 to 520 bytes of spare area per page.

System and subsystem designers will find that it's very difficult to use NAND Flash devices from multiple suppliers if their file systems make use of the NAND Flash spare area for markers and or meta-data. While the combined capacity requirements for ECC and file-system data may be available in the spare area for a particular NAND Flash device from one vendor, there may not be sufficient capacity in a device from another vendor or even in a different NAND Flash device from the same vendor.

ECC Error Correction (Bits)	Number of Spare Bytes Per ECC Sector	Sector Size (Bytes)	Number of ECC Sectors per page	Page Size (Bytes)	Required Spare Area (Calc)	Manufacturer Spare Area in Bytes
8	13	512	4	2048	52	64
8	13	512	8	4096	104	218
8	13	512	8	4096	104	218
16	26	512	8	4096	208	218
24	39	512	8	4096	312	218
24	39	1024	8	8192	312	436
24	39	1024	8	8192	312	520

Table 3: Flash page structural design

Warning 3: Page and Spare Area Trade-offs

NAND Flash IC's page architecture can either be simple or very complex and confusing. **Figure 1** represents a typical 4-Kbyte page divided into eight 512-byte sectors for ECC purposes while **Figure 2** shows a typical 8-Kbyte page divided into 1-Kbyte ECC sectors.




Figure 1: NAND Flash with 4-Kbyte page divided into eight 512-byte pages





Figure 2: NAND Flash with 8-Kbyte page divided into eight 1-Kbyte pages

The 4K and 8K NAND Flash page sizes are currently the most common. For the NAND Flash and ECC controller to achieve the desired level of correction (Error Bit Threshold), each page is segmented into ECC sectors and the ECC sector size is typically specified in the datasheet by the NAND Flash manufacturer as is the recommended error correction use for the spare area. The fixed ECC sector size is required to allow the commonly used Bose, Ray-Chaudhuri, and Hocquenghem (BCH) error-correction code to reliably detect and correct the specified number of bits. Note that we're unlikely to see ECC sector sizes grow larger than 1 Kbyte because larger sector sizes make BCH coding inefficient, especially when dealing with ECC syndromes larger than 30 bits.

Subsystem designs that use NAND Flash may alter the bit allocation between the data area and the spare area to improve ECC effectiveness by reallocating bits from the data area to the spare area for additional ECC storage as shown in **Figure 3**. This bit reassignment increases storage reliability and improves ECC performance. Note that no changes are made to the NAND Flash device to accomplish this reallocation; it's done with special addressing routines typically written into the Flash Translation Layer (FTL) software. While this approach improves ECC effectiveness and therefore decreases the likelihood of a bit error, it also reduces device capacity.



- Change in Data as well as Spare area
- Smaller ECC sector size
- Higher Bit ECC correction possible

Figure 3: NAND Flash reduced page size, and increased spare area size to improve ECC reliability

Warning 4: More ECC (and Controller Cost) Ahead

Shrinking device lithographies and the trend towards multi-level cell (MLC) NAND devices and away from single-level cell (SLC) devices are driving a corresponding increase in the number of bits used for NAND Flash error correction (the Error Bit Threshold). Both trends are increasing the capacity per NAND Flash device to further reduce the per-bit storage cost. Several factors are responsible for an increase in NAND Flash bit errors, including:

- Charge injection and Trapped Charge are increasing the number of soft bit errors
- Increases in random hard errors (cell failures), and
- Write-cycle reductions

Figure 4 shows the increase in NAND Flash ECC bits needed to deal with the greater incidence of NAND Flash bit errors over several process nodes and extrapolates the trends to the 2X nm process node. Note that MLC devices currently attract the most focused development effort, as you might expect in an industry where cost/bit is of paramount importance, and MLC devices need far more error correction.

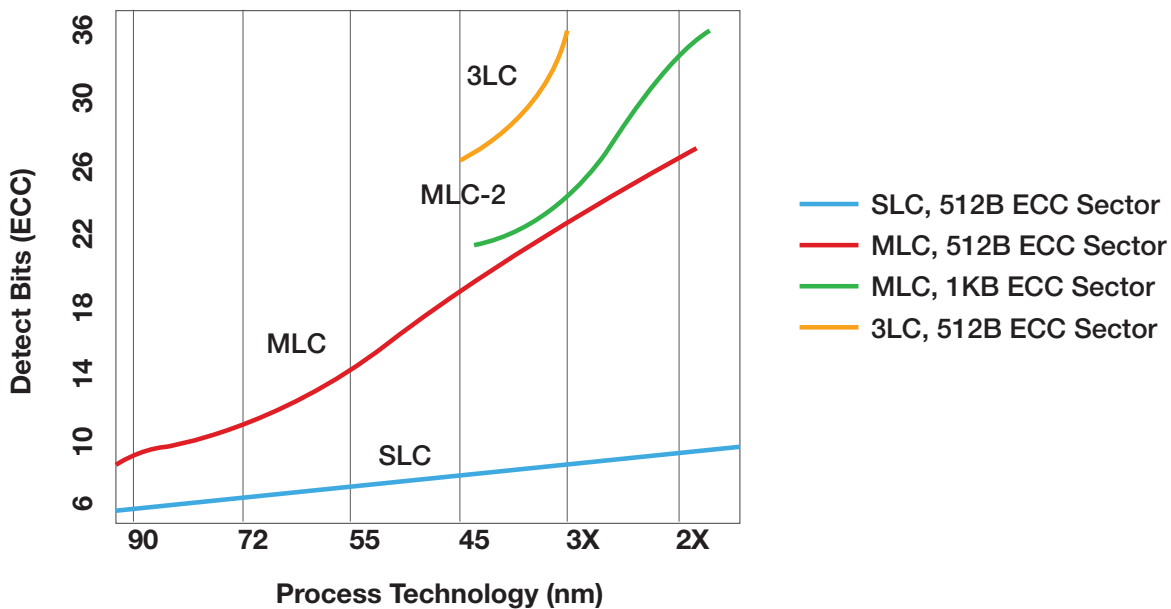


Figure 4: Error correction threshold trend by flash technology

As mentioned already, the current trend is to use 1-Kbyte ECC sector sizes. Although not really required, the industry has naturally converged and is moving from 512-byte to 1-Kbyte ECC sector sizes to accommodate larger ECC codes. This trend directly affects the area and therefore the cost of the hardware implementations for NAND Flash controllers. An 8-bit ECC generator/checker consumes about 15K gates while a 26-bit ECC generator/checker requires about 100K gates. This change disproportionately affects production cost in low-end application markets because low-end products are far more sensitive to the controller's silicon area and cost.

Beyond 32 bits of ECC, new and more advanced ECC approaches include mixed ECC solutions, new algorithms, DSP solutions, and hybrid devices that perform advanced ECC. Perhaps some completely different form of error correction will be needed. For example, one possibility is the use of Cyclical Redundancy Check(CRC)—which is a common error-detection-and-correction scheme used for hard disk drives that performs very fast “Go/No Go” comparisons for sub-sector accesses and reduces the read latency for good data blocks. Stay tuned as there are more advanced approaches to error correction in various stages of research and development beyond CRC.

Conclusion

NAND Flash devices continue to evolve rapidly, complicating their use. There are numerous interrelated design factors and dimensions. Optimum use of any specific NAND Flash device is highly design-dependent and will continue to be so. New applications must employ better NAND-management solutions to improve reliability and increase system performance. These improvements will be incorporated into software, hardware, and system interfaces.

Finding the optimum mix of these design details requires a tremendous amount of device knowledge, industry trend data, solid relationships with the device manufacturers, and related system-design experience. For example, newer high-speed NAND Flash interfaces such as Toggle and ONFi 2 will simply add to the long list of complexities and designers need to understand the trade-offs involved. Denali has the knowledge, experience, and the relationships in the NAND Flash arena. Contact Denali's NAND Flash experts for help with your next NAND Flash design. We can help turn the wild ride into a profitable one.

Author

Robert Pierce is Senior Director of Flash Products for Denali Software. He brings more than twenty years of experience in the semiconductor and IP industry to bear on NAND Flash analysis. He has held several senior management positions in startup and large IC companies, including eSilicon, Infineon, and OKI Semiconductor. Mr. Pierce was one of the co-developers of the RLDRAM architecture, and pioneered the embedded DRAM market for Infineon.

For More Information

www.denali.com/flash

www.denali.com/dmr