

this issue:

Nine Myths about
NAND and NOR,
Continued P.2

STI and LOCOS
Technologies React
Differently to ESD
Pulses P.3

Ask the Experts:
SWEAT EM Test
Structures P.3

Semitracks 2010
Tentative Course
Schedule P. 4



Nine **Myths** about NAND and NOR

[By Chris Henderson]
A Brief History of Non-Volatile
Memory, Part Three

Non-volatile memory technology, which can hold data with the power removed, has become critical to devices such as PDAs, cell phones, and digital cameras. In the final installment of our series on non-volatile memory, we'll examine some common myths about NAND and NOR flash memory. Although the article is dated by approximately 3 years, it will still give the reader an understanding of the differences.

Myth # 1: NAND Flash is slower than NOR.

The Facts: Although NOR Flash offers a slight advantage in random read access times, NAND offers significantly faster program and erase times. A 64KB erasable unit of NAND memory performs a typical program and erase sequence in an average of 17 milliseconds—significantly faster than the average time of 2.4 seconds for a comparable NOR unit. In a system application, the user can easily notice the difference. Additionally, NAND random read access performance supports system requirements without noticeably delaying the user.

The performance characteristics of NAND Flash—fast write (or program) speed, fast erase speed, and medium read speed—make NAND Flash ideal for low cost, high density, high speed program/erase applications. Today, many designers build upon the conventional cell phone memory architecture by increasing density of the NOR and PSRAM and adding NAND Flash to obtain greater performance and capacity for data storage. For high performance data storage requirements such as storing digital photos or downloading music, NAND write/erase speeds provide a distinct performance advantage.

Myth # 2: NAND is not reliable.

The Facts: Similar to the error correction process in a hard disk drive, NAND controller maps around bad memory areas and error correction code (ECC) correct bit errors. Historically, memory subsystems have used Hamming codes, while ECC has been common in hard drives and CDROMs. However, all controllers for NAND Flash have built-in ECC to automatically correct bit errors. The industry standard mandates a 10-14 correction standard, which allows only one bit uncorrectable error every 10^{14} bits (12.5 terabytes). The 10-14 correction standard minimizes bit errors to a level comparable to that of hard disk drives. Over the years, system designers have tested and proven the benefits of using ECC to detect and correct errors.

Myth # 3: NAND Flash is hard to integrate into a system.

The Facts: Today, a wide selection of NAND controllers and software drivers are available, making integration into a system relatively simple. NAND Flash has an indirect, I/O-like access. Therefore, data must be accessed through a command sequence instead of through the direct application of an address to the address lines. NAND Flash also has internal command, address, and data registers. Although the NAND interface may appear more cumbersome than the direct interface of NOR Flash, NAND Flash offers a significant advantage in that upgrading to a higher

density chip is relatively easy. Due to NAND's indirect interface, the external pinout (host connection) does not change with the density of the chip. Thus, different chips can use the same bus interface, similar to the way different densities of hard disk drives can use the same cable interface.

Myth # 4: MLC NOR is close to matching NAND capacities.

The Facts: The maximum available density currently available in MLC NOR Flash is 256 MB. The highest available capacity for MLC NAND Flash is currently 4 GB and the highest available capacity for SLC NAND is 2 GB. Multi-level Cell (or MLC) NOR, a common method to increase the capacity of NOR Flash, stores multiple charge levels (typically four), which enable the storage of 2 bits in a memory cell. However, implementing MLC architecture further reduces effective speed and write/erase endurance.

Myth # 5: MLC NAND does not have the performance or endurance to reliably store your digital photos.

The Facts: MLC NAND Flash allows each memory cell to store two bits of information, compared to one bit-per-cell for SLC NAND Flash, resulting in a larger capacity and lower bit cost. Currently, SLC Flash is rated to have approximately 100,000 cycles and MLC Flash is rated to have approximately 10,000 cycles. While SLC NAND may be more appropriate for some specific applications, the difference will not affect the many common consumer applications, including most digital camera users.

If a 256MB MLC card can typically store 250 pictures from a 4-megapixel camera (a conservative estimate), its 10,000 read/write cycles combined with wear-leveling algorithms in the controller will enable the user to store and/or view approximately 2.5 million pictures within the expected useful life of the card. That number is so far beyond the average number of photos taken by the typical user that the difference in endurance is not significant for this application. MLC NAND provides a very competitive level of performance and makes high density NAND cards more affordable, resulting in its growing popularity among consumers.

Myth # 6: MLC NAND does not have high enough performance for streaming video.

The Facts: The performance of MLC NAND is sufficient to support the 6 to 8 Mbits/second transfer rate needed to store MPEG2 compressed video on a memory card. Many manufacturers large block NAND manufactured on 90nm process technology can write approximately 2.5 MB/second.

Myth # 7: SLC NAND is a generation ahead of MLC NAND.

The Facts: On several companies' roadmaps, SLC development leads MLC by only two to three months. Presently, for each new generation, SLC chips are designed with MLC requirements in mind, so little lag-time exists between the two types of NAND. The real

issue is market acceptance, not actual time-to-market for the next generation. However, MLC development is currently well-timed to match market acceptance, with 512MB and 1GB cards widely available to meet market demand.

Myth # 8: The additional circuitry needed for MLC NAND takes up a significant amount of real estate.

The Facts: MLC NAND requires relatively minimal circuitry. A 4 GB MLC NAND Flash chip provides approximately 1.95 times greater density than a 2 GB SLC NAND chip.

We believe that the more important question to the user is "What density can you get in a chip today?" Presently, the highest density MLC NAND Flash in production is 4 GB, whereas the highest density SLC NAND in mass production is 2 GB. The market demand for ever-higher densities of removable storage makes the lower-cost, higher density MLC card attractive to users and continues to enable new applications to emerge. The rated storage capacity of 2 GB SLC NAND is 271MB, compared to 529 MB for a 4 GB MLC NAND Flash chip, for a density increase of approximately 1.95 times.

Myth # 9: NAND Flash is a slow storage technology.

The Facts: NAND Flash offers excellent performance for data storage. As a point of comparison, NAND can offer significantly faster performance and reliability than a hard disk drive, depending on the number and size of files transferred. For a random access of a 2 KB file, a typical hard disk drive might take approximately 10ms to retrieve a file, while NAND Flash would take about 0.13ms to retrieve a similar size file. For a comparable write function with the 2kB file, NAND can perform up to 20 times faster than a typical hard drive. Because it is a solid state memory with no moving parts, NAND flash features a significantly shorter random access time compared to a mechanical hard disk drive.

<u>Parameter</u>	<u>NOR Flash</u>	<u>NAND Flash</u>
Density	Up to 128 Mbits	Up to 2 Gbits
Read Speed	55 MB/sec	27 MB/sec
Write Speed	0.15 MB/sec	8.3 MB/sec
Erase Speed	32 kB/sec	85 MB/sec
Rewrite Speed	26 kB/sec	7.5 MB/sec
Interface	Random Access	I/O – Indirect Access
Applications	Execute in Place	Program/Data Mass Storage

STI and LOCOS Technologies React Differently to ESD Pulses

[Technical Tidbit]

Shallow Trench Isolation (STIs) and Local Oxidation of Silicon (LOCOS) technologies each react uniquely to an ESD event. The illustrations in Figure 1 help to explain the electric field increase associated with scaling a LOCOS isolation. The junction at the top represents an implanted junction in a LOCOS process, and the junction at the bottom represents an implanted junction in an STI process.

As junctions are scaled in depth in a LOCOS process, the curvature on the edge of the metallurgical junction leads to high electric fields with scaling. The increase in curvature can lead to electrical breakdown, hot electron effects, and non-uniform current distribution. However, in STI isolation, the junction terminates at the edge of the STI trench, yielding a planar junction edge. Therefore, in STI implanted regions, the curvature problem does not produce the adverse affects seen in the LOCOS technology.

The graph on the following page shows how LOCOS and STI isolation affects the temperature rise associated with an ESD pulse. Because the deep trench associated with STI causes the heat to be confined, the temperature rises significantly in the STI sample. The temperature rise in the LOCOS-isolated device is much lower.

For further information on ESD and Latchup, please consider attending our ESD/Latchup Design Short Course.





Figure 1. Top: Implanted LOCOS junction; Bottom: Implanted STI junction.

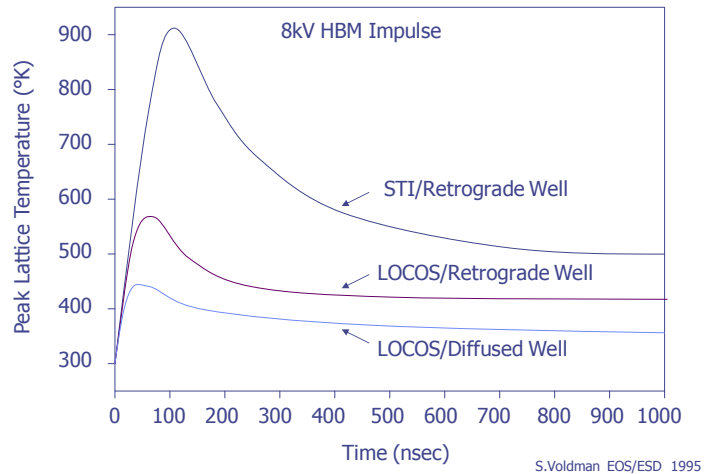


Figure 2. Temperature rise associated with an ESD pulse.



Ask the Experts

Q: How do you calculate the EM lifetime using the SWEAT test structures when the metal width is not uniform, alternating between narrow and wide segments? What is the normal EM failure location at SWEAT test structures?

A: The EM lifetime is normally calculated as a time to a percent increase in resistance in the structure

(maybe 10%, or some other number). Given the temperature that you tested and the current density, you can then use Black's equation to project down to use conditions. The big problem with the SWEAT structure is that it does not replicate the metal layout on an IC, so most people use ASTM structures or other via-blocked structures that replicate the metal layout on a chip more effectively. The normal failure location will be at the beginning of a narrow segment based on the directions the electrons flow. Thus, if the electrons flow left to right in the structure on the slide, then the EM failure location is usually near the left end of a narrow segment.

To post, read, or answer a question, visit <http://forums.semitracks.com/>. We look forward to hearing from you!

2010 Semitracks Course Schedule

Mark Your Calendars and
Register Today!



Customer Feedback

If you have a suggestion or a comment regarding our courses, online training, discussion forums, or reference materials or if you wish to suggest a new course or location, please feel free to call us at 1-505-858-0454 or e-mail us at info@semitracks.com.

To submit questions to the Q&A section, inquire about an article, or suggest a topic you would like to see covered in the December newsletter, please contact Alicia Constant by email at alicia.constant@semitracks.com.

We are always looking for ways to enhance our courses and educational materials.

~

For more information on Semitracks online training or public courses, visit our website!
www.semitracks.com

[Failure and Yield Analysis](#)

Malaysia
January 19-22, 2010

[Semiconductor Reliability](#)

Malaysia
January 25-28, 2010

[Photovoltaics Technology and Manufacturing](#)

San Francisco, CA
February 1, 2010

[CDMA and EV-DO Technology](#)

San Jose, CA
February 8-10, 2010

[Semiconductor Reliability](#)

San Jose, CA
February 8-10, 2010

[Failure and Yield Analysis](#)

San Jose, CA
February 15-18, 2010

[MEMS Technology](#)

Austin, TX
March 1-2, 2010

[Optimizing Factory Performance](#)

Austin, TX
March 3-5, 2010

[Defect-Based Testing](#)

Austin, TX
March 8-9, 2010

[Design for Reliability](#)

Austin, TX
March 10-11, 2010

[MEMS Technology](#)

Cambridge, UK
March 15-16, 2010

[Defect-Based Testing](#)

Melaka, Malaysia
March 24-26, 2010

[Failure and Yield Analysis](#)

Shanghai, China
March 29-April 1, 2010

[Photovoltaics Technology and Manufacturing](#)

Malaysia
April 19-21, 2010

[Packaging Technology and Metallurgy](#)

Malaysia
May 26-28, 2010

[MEMS Technology](#)

Malaysia
May 31-June 2, 2010

[ESD/Latchup Design](#)

Boston, MA
June 14-15, 2010

[Bipolar and BiCMOS Process Integration](#)

San Jose, CA
June 21-22, 2010

[CMOS Process Integration](#)

San Jose, CA
June 23-24, 2010

[Interconnect Process Integration](#)

San Jose, CA
June 25, 2010

[Photovoltaics Overview](#)

San Francisco, CA
July 12, 2010

[Photovoltaics Technology and Manufacturing](#)

San Francisco, CA
July 13, 2010

[Reliability Challenges](#)

San Francisco, CA
July 15, 2010

[RF & Mixed Signal Testing](#)

Phoenix, AZ
July 26-28, 2010

[IC Packaging Design and Modeling](#)

San Jose, CA
August 9-11, 2010

[Wafer Fab Processing](#)

San Jose, CA
August 9-12, 2010

[IC Packaging Technology](#)

San Jose, CA
August 12-13, 2010

[Failure and Yield Analysis](#)

Munich, Germany
September 7-10, 2010

[Semiconductor Reliability](#)

Munich, Germany
September 13-15, 2010

[Wafer Fab Processing](#)

San Jose, CA
October 11-14, 2010

[Dynamic Random Access Memory](#)

Dallas, TX
November 19, 2010