

Data Deluge

More data
More opportunities



November 14 – 15, 2012 | San Jose, CA

Next-Generation Flash: Challenges and Solutions

Presented by: Earl T. Cohen





Storage. Networking. Accelerated.™

Agenda

- Key NAND Flash Trends
 - Higher Bit Count
 - Shrinking Geometries
- Challenges of Higher Density Flash
 - Performance
 - Usability
- Solutions to the Higher Density Flash Challenges
 - Performance
 - Usability

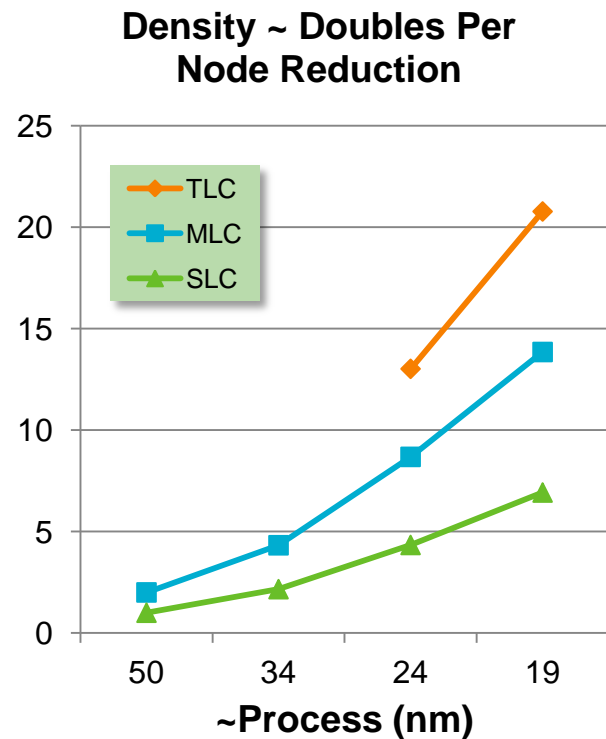
Key NAND Flash Trends



Increasing Density From Two Directions

Benefits of Higher Density Flash

More bits per cell and reduced geometry increases GB/sq mm



For the Same Qty Silicon

2GB
50nm
SLC



32GB
19nm
MLC


16x Capacity Increase for Same Cost

Enables More GB in Same Footprint and Reduced \$/GB


Challenges of Higher Density Flash

■ Performance

- Higher read latency and longer program times




- Fewer die → less parallelism at a given capacity




■ Usability

- Increasing error rates

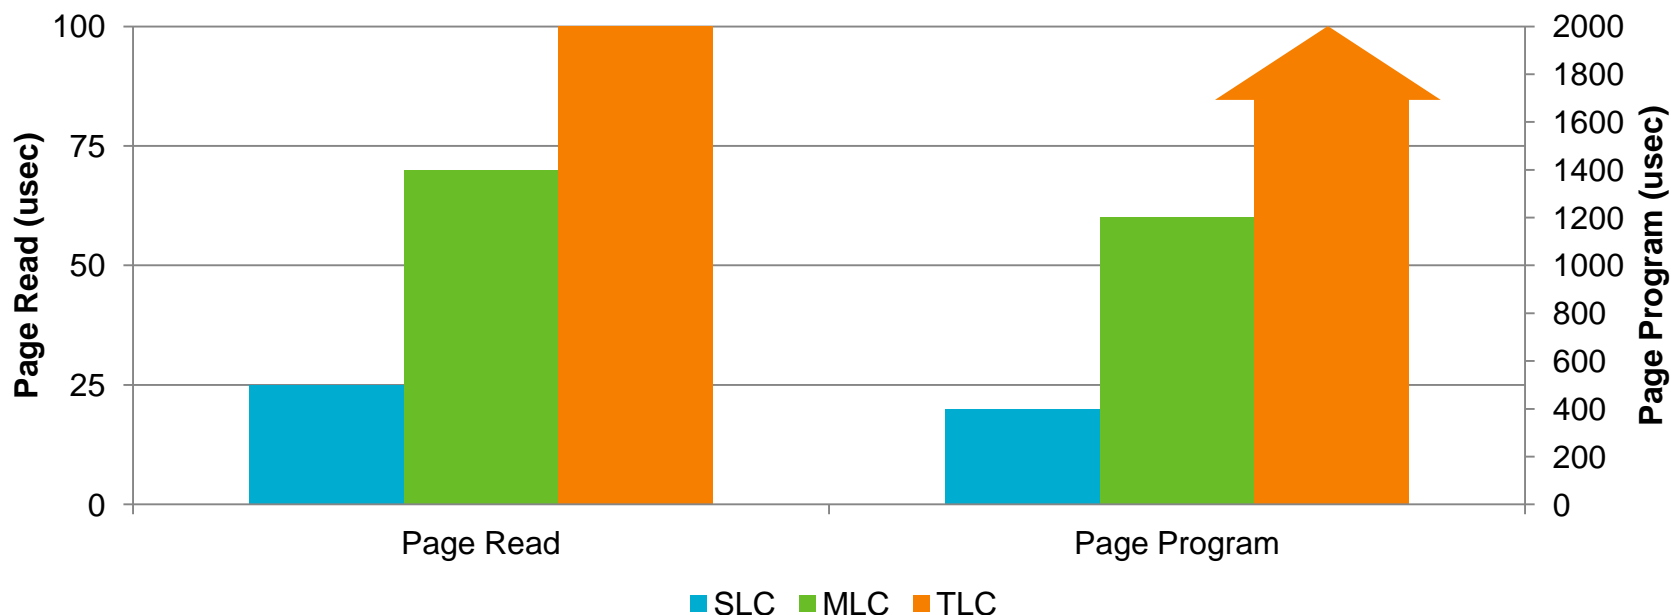


- Shorter endurance




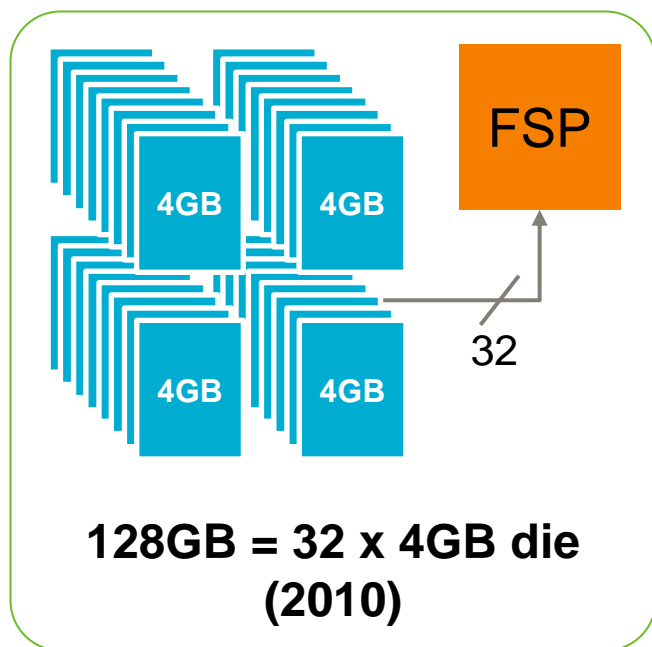
Challenges of Higher Density Flash (cont'd)

- Performance: Higher Read Latency and Longer Program Times
 - As geometries shrink, the flash read access and program times are increasing
 - MLC and TLC are cost-effective, but more bits per cell slow operations down
 - It takes more time to deal with the finer resolutions
 - SLC Tr (page read) is ~25 usec, vs. ~70 usec for MLC and ~100 usec for TLC
 - Page program time is ~400 usec for SLC vs. ~1200 usec for MLC
 - And even more for TLC



Challenges of Higher Density Flash (cont'd)



- Performance: Fewer Die → Less Parallelism at a Given Capacity
 - 128GB SSD in 2010 had 32, 4GB die
 - 128GB SSD in 2013 has 8, 16GB die
 - Instead of 32 parallel accesses, the 2013 SSD can only support 8
 - If Queue Depth of I/O operations is never more than 1 it won't matter
 - But that's not an interesting, real-world condition for most SSDs



Challenges of Higher Density Flash (cont'd)

■ Usability: Increasing Error Rates


- Decreasing flash geometries have made cells hold less charge
 - From >1K electrons per cell (in 2004) down to ~100 electrons in 2013
- Less charge means more susceptibility to errors
 - Inter-Cell Interference (ICI) when programming
 - Read Disturb
 - Reading one page affects its neighbors
 - Leakage (retention)
 - Loss of electrons over time
 - Charge trapping / oxide breakdown
 - And more
- More states (e.g., MLC vs. SLC) means greater accuracy needed
 - States are closer together



Challenges of Higher Density Flash (cont'd)


■ Usability: Shorter Endurance

- Decreasing flash geometries and increasing bits per cell increase error rates
- The higher error rates reduce endurance
 - As the cells wear they cannot hold the charge as well and eventually the errors cannot be recovered leading to EOL
- Endurance of 100K or more cycles with older SLC down to 3K cycles with newer MLC
 - And perhaps 500 cycles with newer TLC



Solutions to Higher Density Flash Issues


Performance: Higher Read Latency & Longer Program Times



- Improve Flash Storage Processors (FSP) architecture and F/W to reduce average and max latency
 - Except for the QD=1 latency freaks, real-world latencies are a complex function that calls for balancing all foreground and background latencies of the SSD, *and* for efficient management of flash usage
 - Checkpoint Latency is a large area to be pursued to reduce outliers
- Maximize usage of write bandwidth
 - Reduce write amplification
 - Design to be flash-limited on write

Solutions to Higher Density Flash Issues


Performance: Fewer Die → Less Parallelism at a Given Capacity



- Add higher levels of parallelism in the FSP
 - Ensure that all die *and* all flash channels are efficiently used
- Optimize for newer flash architectures
 - Take advantage of increasing multi-plane page size
 - More write parallelism
 - Take advantage of new flash features
 - Partial page read, erase suspend, etc.

Solutions to Higher Density Flash Issues


Usability: Increasing Error Rates



- Enhance FSP to use LDPC (Low-Density Parity-Check) error correction (*beyond RS and BCH*)
 - Flash is becoming more of an analog vs. digital media
- Leverage both hard- and soft-decision LDPC decoding
 - Soft-decision uses analog “voltage level” information obtainable from flash
- Employ signal processing techniques, such as via DSPs
 - Understand the flash channel model – not just simple AWGN type errors
- ECC space in the flash needs to be changed by the FSP in some cases
 - Many more endurance/reliability trade-offs possible if you can do this
- Extend technologies such as RAISE™ data protection

Solutions to Higher Density Flash Issues

Usability: Shorter Endurance



- Take technologies like DuraWrite™ endurance enhancement to the next level
- Reduce writes in the first place
 - The lower the write amplification, the higher the endurance
- Optimize the required garbage collection and wear-leveling processes
 - A complex, multi-dimensional problem to wear evenly while trying to reduce write-amplification due to garbage collection
- New features to make better use of all available space in flash
 - Variable-Size Flash Translation Layers that enable re-use of spare data not needed for error correction
- Work with flash vendors on performance/endurance trade-offs
 - Tune flash for desired operating points

Key Takeaways...

- Decreasing geometries make usage of flash memory more complicated in compute environments, such as SSDs
 - There are more than just density issues to manage
 - The Flash Storage Processor is key to solving these issues
 - LSI is well positioned to solve these problems
-
- Come see the live demonstrations of current Flash Storage Processors in the exhibition area



Storage. Networking. **Accelerated.**[™]

Some of the views expressed herein are opinion and suggestions only. Individual performance or testing results may vary. LSI, the LSI & Design logo, DuraWrite, RAISE, and Storage. Networking. Accelerated. are the trademarks or registered trademarks of LSI Corporation. All other brand or product names may be the trademarks or registered trademarks of their respective companies.