

# Solid State Drives

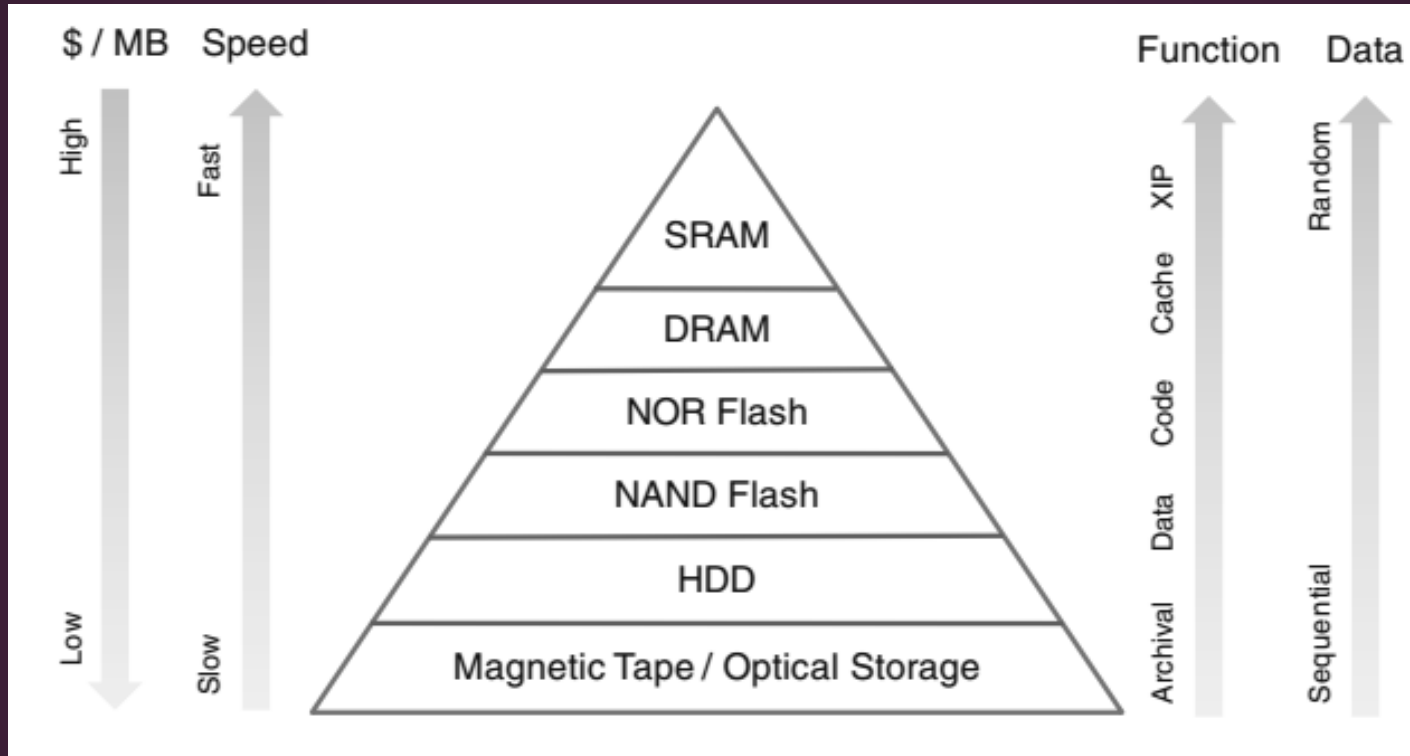
*By: Amar Bhatt and Mateo Diz*

# Agenda

- Overview
- History
- Technology
- Performance Limitations
- Solutions to Limitations
- Current Applications
- Future Research

# Overview

- Solid state drives are the newest solution to fast, high-capacity storage
- Reliable
- One of the fastest growing research fields in the last decade
- Popular choice among laptop and mobile device manufacturers



Source: Micheloni, R., Marelli, A., & Eshghi, K. (2012). *Inside Solid State Drives*. New York City, New York: Springer.

# History


- Used initially to improve boot-time on a machine
- Created as an alternative to solve memory problems caused by moving parts on a Hard Disk Drive
- Increased popularity among mobile device manufacturers

# Leaders in manufacturing

- 2003
  - Samsung controlled 60% of the SSD market in their sole production of SLC-SSD
  - Toshiba controlled 30% of the SSD market in the production of SLC and 2X MLC
- 2008
  - Samsung controlled 40% of market
  - FlashVision, Hynix, IM Flash
    - Manufacturers for SanDisk, Toshiba, Intel, Micron

# Cost Overtime

- As time goes on SSDs become exponentially cheaper
- 2004 8GB SLC-SSD, \$4,000
- 2009 16GB MLC-SSD, \$200



Source: Grupp, L., Davis, J., & Swanson, S. (n. d.). *The Bleak Future of NAND Flash Memory*. Retrieved April 30, 2015, from <http://cseweb.ucsd.edu/~swanson/papers/FAST2012BleakFlash.pdf>

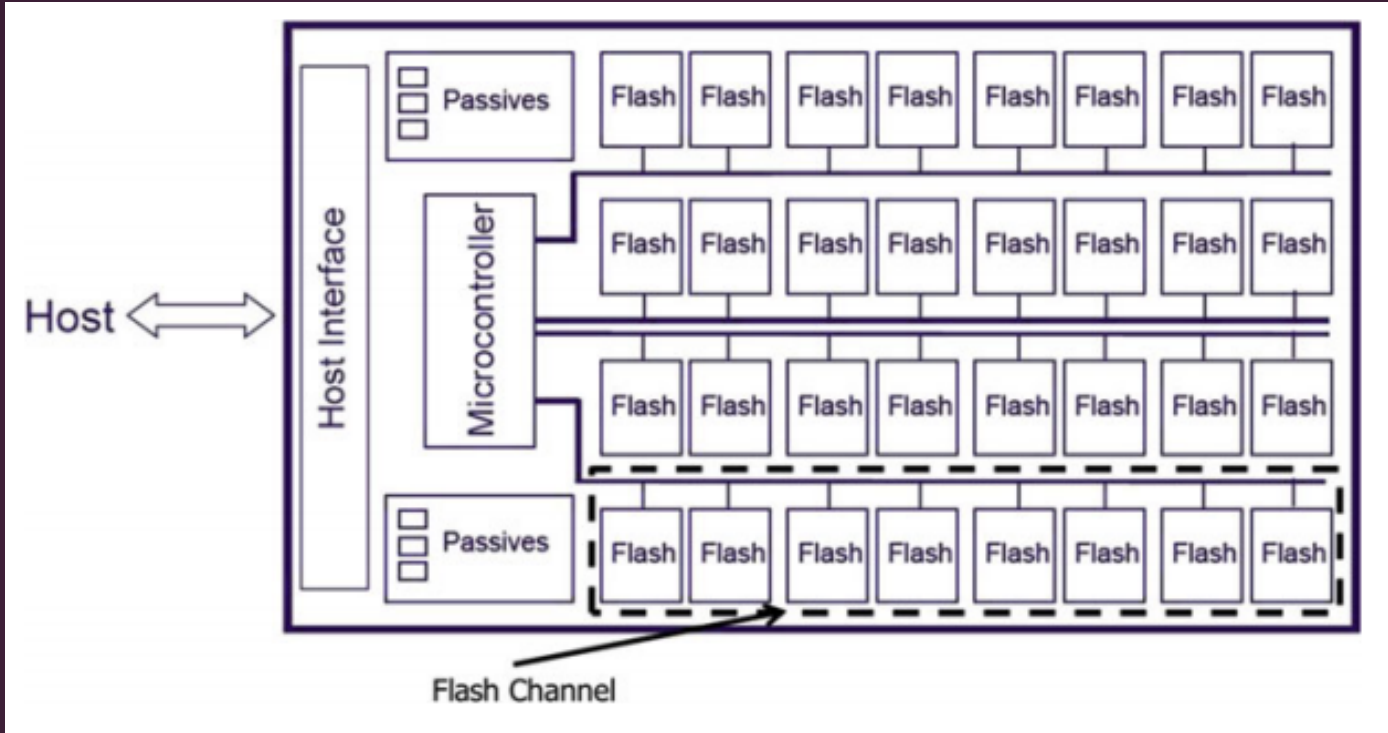
# How it works

- NAND-Based Architecture
- Bits per cell
  - Single-Level Cell
  - Multi-Level Cell
  - Triple-Level Cell



# NAND-Based Architecture

- Cheaper to manufacturer
- Non-volatile memory
- Parallel architecture suitable for storing and executing code (unlike NOR-based architectures)
- Durable, shock-resistant
- Low Power Consumption




Source: Micheloni, R., Marelli, A., & Eshghi, K. (2012). *Inside Solid State Drives*. New York City, New York: Springer.

# Workings of a Flash Cell

- Cell value is obtained by using a test voltage
- Setting a cell
  - Programming Voltage - sets cell to logic 0
  - Erasure Voltage - sets cell to logic 1
  - Higher than test voltage because they need to force electrons over an oxide substrate between two gates


# Inside of a Flash Cell

- Floating Gate Transistor
- Bitline connected to Source
- Word line connected to gate
- Voltage on the gate allows for reading a value
- At threshold voltage, gate is used to program/erase a cell



*Source: NAND Evolution and its Effects on Solid State Drive (SSD) Useable Life. (2009). Western Digital White Paper. Retrieved May 4, 2015, from [http://www.wdc.com/WDPProducts/SSD/whitepapers/en/NAND\\_Evolution\\_n\\_0812.pdf](http://www.wdc.com/WDPProducts/SSD/whitepapers/en/NAND_Evolution_n_0812.pdf)*

# Programming vs. Erasing



*Source: NAND Evolution and its Effects on Solid State Drive (SSD) Useable Life. (2009). Western Digital White Paper. Retrieved May 4, 2015, from [http://www.wdc.com/WDProducts/SSD/whitepapers/en/NAND\\_Evolution\\_0812.pdf](http://www.wdc.com/WDProducts/SSD/whitepapers/en/NAND_Evolution_0812.pdf)*

# Types of SSDs

- Single-Level Cell (SLC)
  - Stores one bit {0,1}
  - Expensive to make
  - Fast read times
- Multi-Level Cell (MLC)
  - Stores two bits {00,01,10,11}
  - Cheaper to make
  - Most used in retail SSDs
  - Slower read times

# Types of SSDs (cont...)

- MLC-1
  - Can act as an SLC to trade off capacity for performance
- Triple-Level Cell (TLC)
  - Stores three bits {000,001,010, ... 110, 111}
  - Slowest read times
  - Newest to the market

# Technology Growth of SSDs

- NAND technology density is exceeding at a faster rate than Moore's law
- SLC-SSD components are at 64 Gb using stacked monolithic die
- New packaging techniques allow MLC capacities to be between 512 GB - 1 TB
- Closing in on HDD capacity



# Performance

	<b>Solid State Drive</b>	<b>Hard Disk Drive</b>
<b>Start-up time</b>	few milliseconds	seconds
<b>Random Access Time</b>	< 0.1 ms	2.9 - 12ms
<b>Data transfer rates</b>	100 MB/s - 600 MB/s	140 MB/s
<b>Defragmentation</b>	Not needed	Needed
<b>Noise (sound)</b>	No moving parts	Moving parts
<b>Temperature limit</b>	High Temperatures	< 131°F
<b>Size</b>	1.8-3.5"	3.5"
<b>Life Expectancy</b>	Limited based on usage	9 - 11 years
<b>Capacity</b>	2 TB	8 TB



# Performance Goals

- Larger Capacity
- Smaller Area
- Better reliability



Is it practical?

# Capacity

- As capacity increases, latency and instability increase, read bandwidth decreases
- Unstable around 16TB, stops scaling at ~4.6TB
- By 2024 (when 16TB is achieved)
  - Latency will increase 2.5x
  - SSDs will no longer be a speed advantage



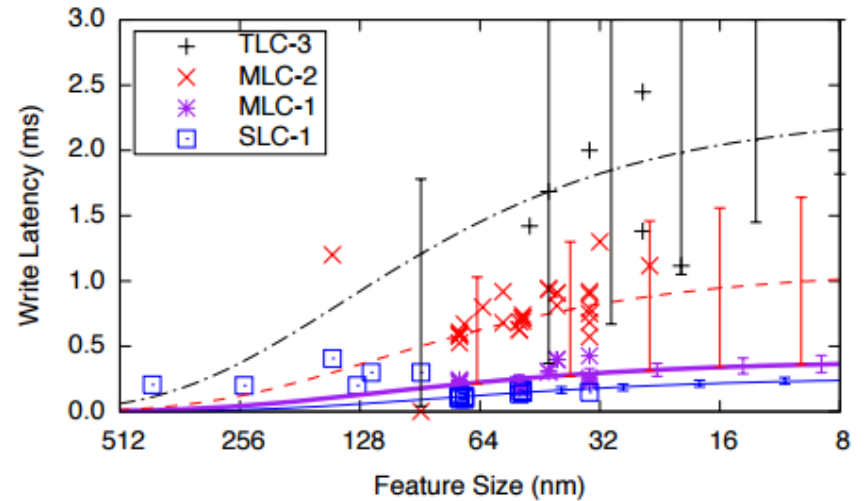
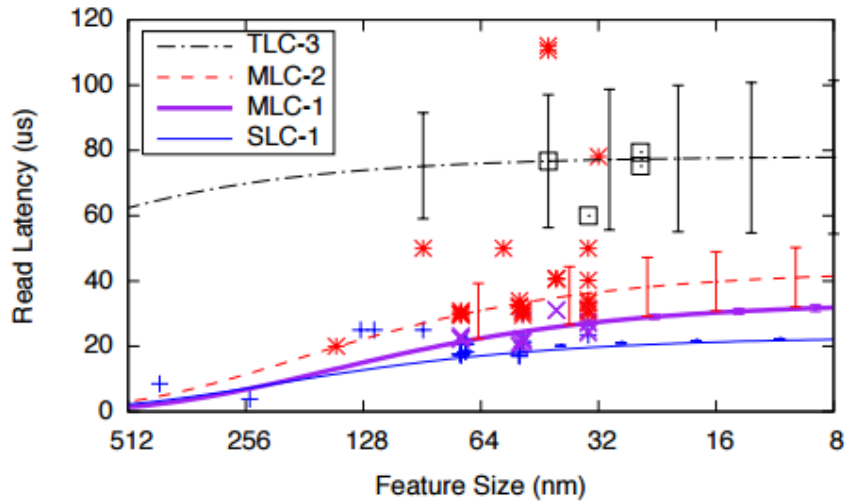
Source: Grupp, L., Davis, J., & Swanson, S. (n.d.). *The Bleak Future of NAND Flash Memory*. Retrieved April 30, 2015, from <http://cseweb.ucsd.edu/~swanson/papers/FAST2012BleakFlash.pdf>



Source: Grupp, L., Davis, J., & Swanson, S. (n.d.). *The Bleak Future of NAND Flash Memory*. Retrieved April 30, 2015, from <http://cseweb.ucsd.edu/~swanson/papers/FAST2012BleakFlash.pdf>

# Area

- SSDs degrade in performance at smaller sizes
- The threshold voltage needed for programming/erasing becomes too close to the reading voltage and degrade the oxide layer
- Currently the limit is at 6.5nm silicon



Source: Grupp, L., Davis, J., & Swanson, S. (n.d.). *The Bleak Future of NAND Flash Memory*. Retrieved April 30, 2015, from <http://cseweb.ucsd.edu/~swanson/papers/FAST2012BleakFlash.pdf>

# Reliability

- Multiple writes and erases break down the oxide barrier, making it harder to keep a charge
- MLC endurance is 10,000 erase cycles
- SLC endurance is 100,000 erase cycles
- As usage increases, the endurance decreases



# Limitation Solutions

- **Wear-leveling**
  - Spread data evenly across SSD
  - Static wear leveling
    - Operating System, look-up tables, executables
  - Block Erasing Table (BET)
  - Evenness-Aware Algorithm
  - Dual-Pool Algorithm

# Limitation Solutions (cont...)

- Error Correction Codes (ECC)
  - Trade off between number of bits corrected, controller cost, complexity, and performance
  - Limited by how many bytes from NAND architecture are dedicated to ECC
    - Usually 16 bits of ECC data for every 512-byte sector
  - Properties

# Current Applications

- Ultrathin Laptops (MacBook Air)
- Database Transactions
- Fast-Changing Data Platforms
  - Takes advantage of short-term reliability
- High Quality Video Streaming

# Future Research

- 3D Memory
- Planar Memory Cells
- SAS Interface

# References

1. Grupp, L., Davis, J., & Swanson, S. (n.d.). The Bleak Future of NAND Flash Memory. Retrieved April 30, 2015, from <http://cseweb.ucsd.edu/~swanson/papers/FAST2012BleakFlash.pdf>
2. How SSDs work. (n.d.). Retrieved May 1, 2015, from <http://www.techradar.com/us/news/computing-components/storage/how-ssds-work-969473>
3. Hudlet, V., & Schall, D. (n.d.). SSD ≠ SSD – An Empirical Study to Identify Common Properties and Type-specific Behavior. Databases and Information Systems Group. Retrieved April 29, 2015, from <http://www.lgis.informatik.uni-kl.de/cms/fileadmin/publications/2011/HS2011.BTW.pdf>
4. Micheloni, R., Marelli, A., & Eshghi, K. (2012). Inside Solid State Drives. New York City, New York: Springer.
5. NAND Evolution and its Effects on Solid State Drive (SSD) Useable Life. (2009). Western Digital White Paper. Retrieved May 4, 2015, from [http://www.wdc.com/WDPProducts/SSD/whitepapers/en/NAND\\_Evolution\\_0812.pdf](http://www.wdc.com/WDPProducts/SSD/whitepapers/en/NAND_Evolution_0812.pdf)
6. Oswald, E. (2012, February 20). Current solid-state drive technology is doomed, says Microsoft Research | ExtremeTech. Retrieved May 1, 2015, from <http://www.extremetech.com/computing/118909-current-solid-state-drive-technology-is-doomed-researchers-say>
7. Solid-state drive. (n.d.). Retrieved May 5, 2015, from [http://en.wikipedia.org/wiki/Solid-state\\_drive](http://en.wikipedia.org/wiki/Solid-state_drive)