Data Integrity on 20nm SSDs

Robert Frickey
Intel Corporation

Agenda

- **Overview of data integrity metrics**
- Challenges as Nand scales to 20nm and below
- Nand-related media optimization
- Future trends

Data Integrity Metrics

- Lots of different data integrity metrics
  - Uncorrectable Bit Error Rate
    - Ex: 1 sector for every 10^{16} bits read
  - Mean Time Between Failures
    - Ex: 2 million hours
  - Annual Fail Rate
    - Ex: 0.44% of drives
  - Data Retention
    - 1 year at room temp at end-of-life

Precise definitions and measurement techniques are needed to ensure clarity and consistency.

Fear, Uncertainty, and Doubt

- Despite datasheet metrics, not easy to predict behavior of SSDs in the field
  - Validation should be considered as part of data integrity
- Need to understand all challenges relating to SSD data integrity
  - Beyond traditional NAND management

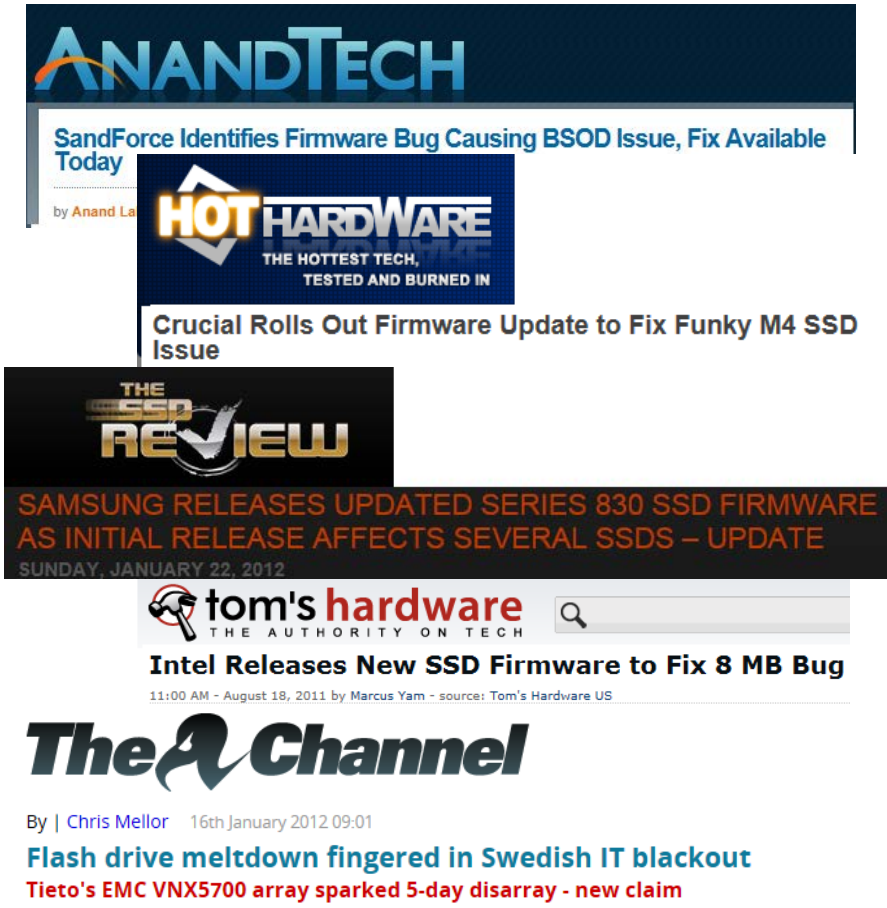

SSD Failure Points

- SSDs are made up of many components
  - Nand
  - DRAM
  - ASIC
  - Passives (caps, supercaps, etc)

- SSD reliability limited by the first component to fail

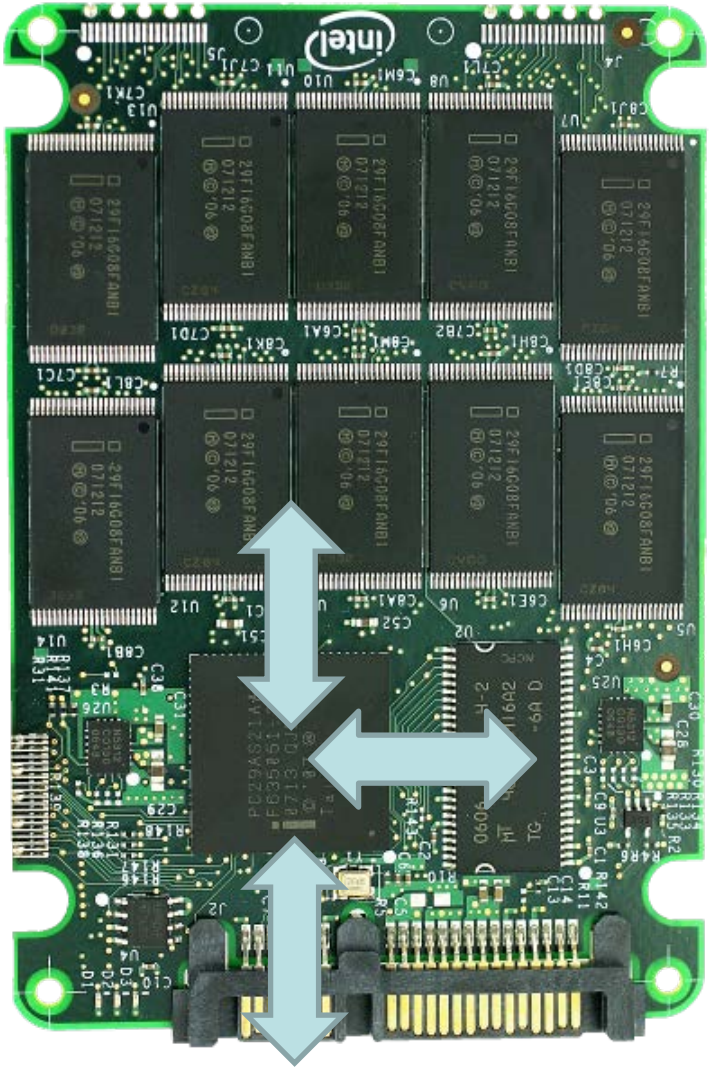

Part II

- Overview of data integrity metrics
- **Challenges as Nand scales to 20nm and below**
- Nand-related media optimization
- Future trends

Review of Intrinsic Nand Limitations

Reliability Issue	Mechanism
Basic endurance: bad blocks	Trapping
Write bit errors	Program Disturb and overprogram
Data retention and read disturb	Detrapping
	SILC (retention and read disturb)
Gross errors	Defects and shorts

What's inside the NAND?

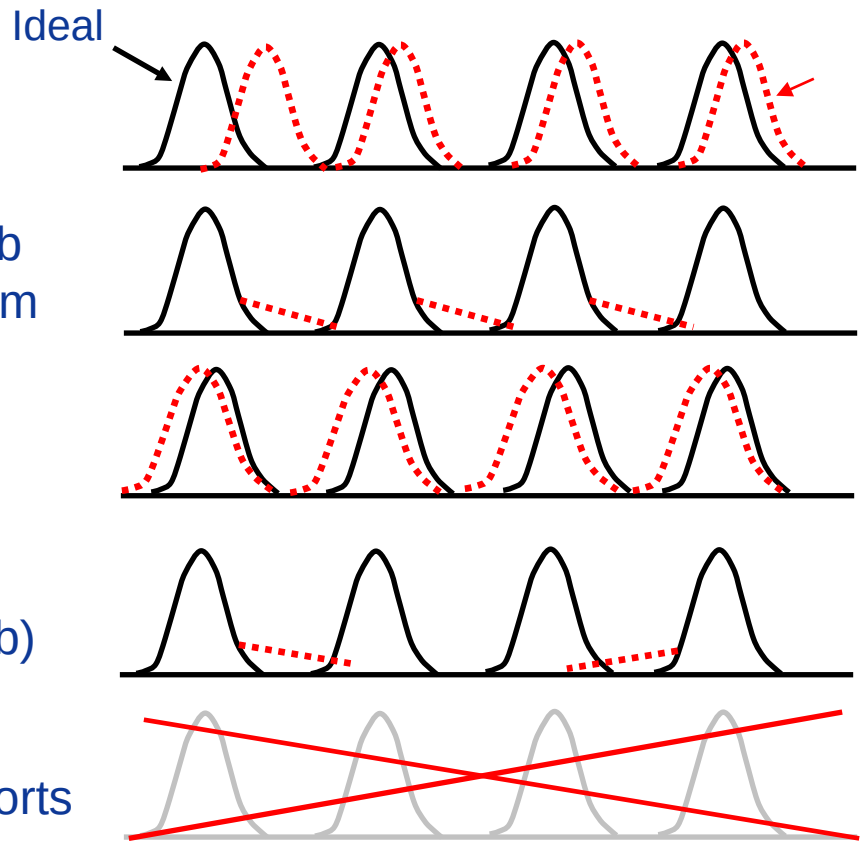

Scaling Challenges
The Fine Print

- Bad news: Raw Bit Error Rate generally increasing with smaller feature size

Litho (nm)	Page Size	Spare Bytes per Sector
72	2112	16
50	4314	27.25
34	4320	28
25	8640	28
20	8936	46.5

- Good news: Challenges with 20nm and smaller NAND are similar to high endurance challenges

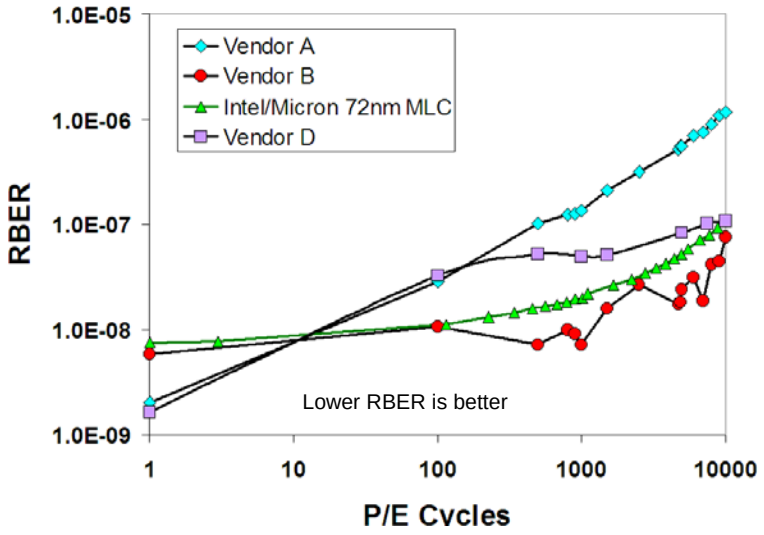

Chart Source: Intel Nand Flash Memory datasheets
Chart Source: Mielke, N. et al, "Bit Error Rate in NAND Flash Memories" IEEE International Reliability Physics Symposium 2008

Neighboring Effects

- In real estate, when a house is sold for a high price, the value of similar homes nearby also goes up.

- NAND flash has a similar property called cell-to-cell interference, where programming a cell slightly increases the V_T of neighboring cells.

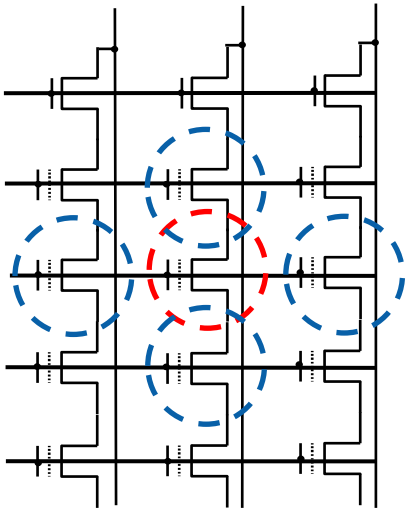

Cell-to-Cell Interference

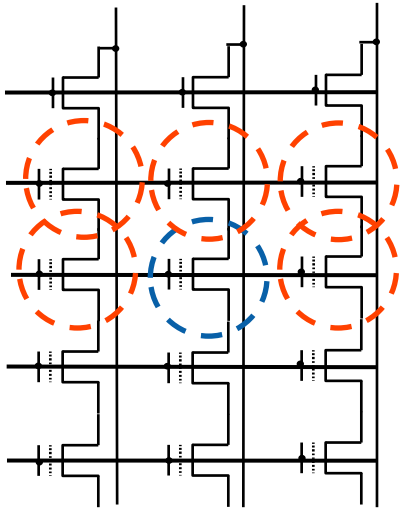

- Due to the nature of the page mapping within a block, there are different types of interference:
  - BL-BL
  - WL-WL
  - Diagonal
- Continued scaling increases the effect of cell-to-cell interference, which widens the program state width.

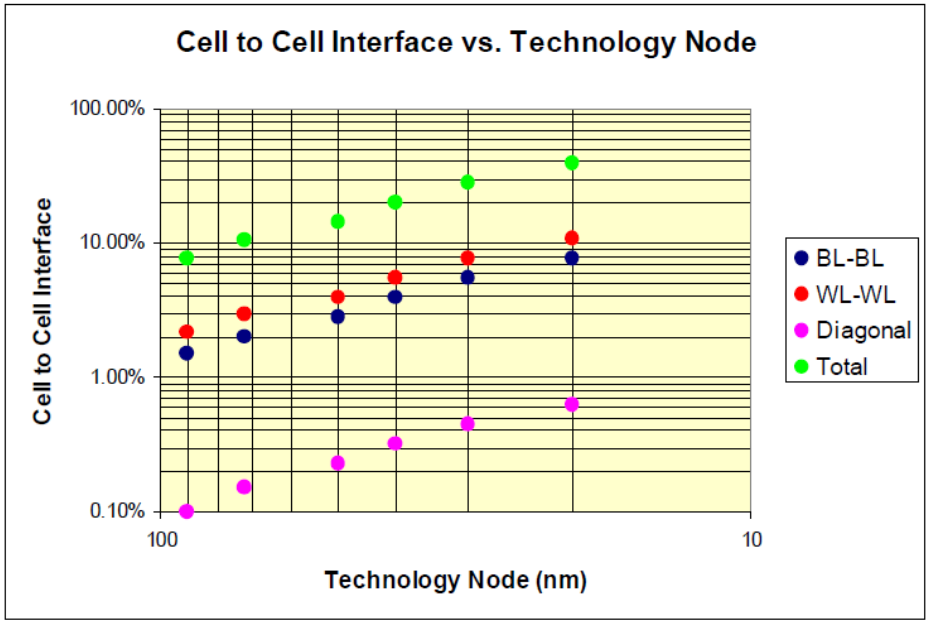

Chart Source: Prall, K., "Scaling Non-Volatile Memory Below 30 nm" IEEE Non-Volatile Semiconductor Workshop 2007

Part III

- Overview of data integrity metrics
- Challenges as Nand scales to 20nm and below
- **Nand-related media optimization**
- Future trends

Redundancy Protects Against NAND Defects

NAND A	NAND B	NAND C	NAND D	NAND E	NAND F
0110	0111	0100	1101	1001	

Parity = 0110 Parity = 0001 Parity = 0101 Parity = 1000 Parity = 0001

- As user data is written to the NAND, parity is calculated and eventually written to the NAND
- If a catastrophic defect occurs, ECC cannot recover data from the failing die
- Data can be recalculated from the parity of the other dice

A RAID-4 style redundancy system can be used to tolerate extrinsic NAND defects.

Read Retry

- MLC placements can shift beyond the reference voltage, causing read errors
- As long as the distributions are not overlapping, the data should be recoverable
- Read Retry can shift the reference until a passing read point is found.

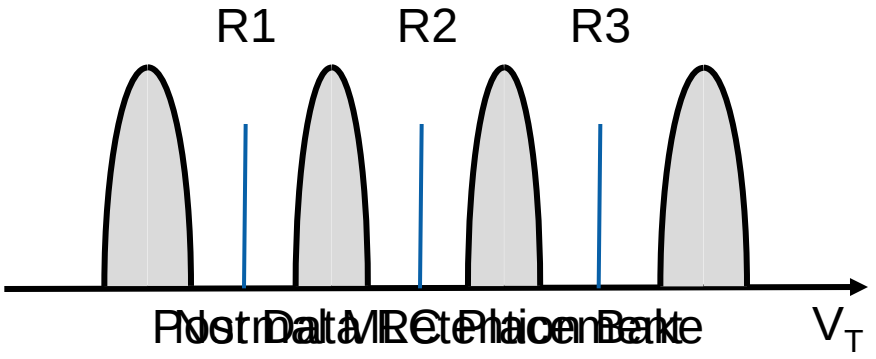

Read Retry can extend endurance capability, with a read latency penalty.

Background Data Refresh

- Data retention is guaranteed even for a power-off condition
- Under normal SSD usage, some static data may sit idle
- BDR ensures that static data is refreshed regularly

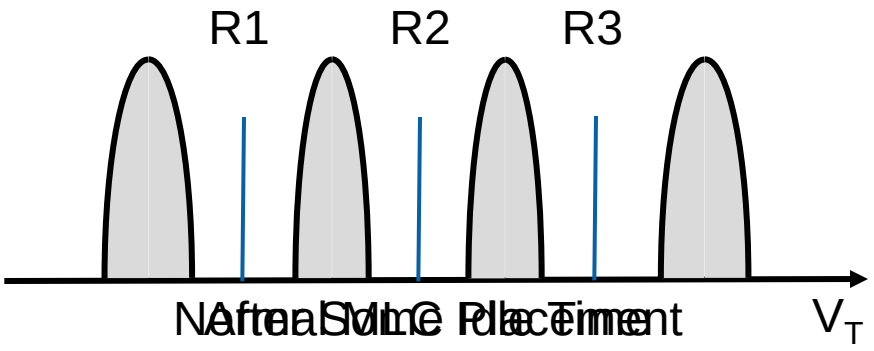

Background Refresh can alleviate concerns about power-on data retention.

Advanced Read Retry

- Advanced Read Retry is an attempt to mitigate the effect of cell-to-cell interference
- The neighboring cells are read first to determine their state
- Next the desired cells to read out are read twice, with a high and lower reference
- The NAND decides which reference data to return based on the neighbor data

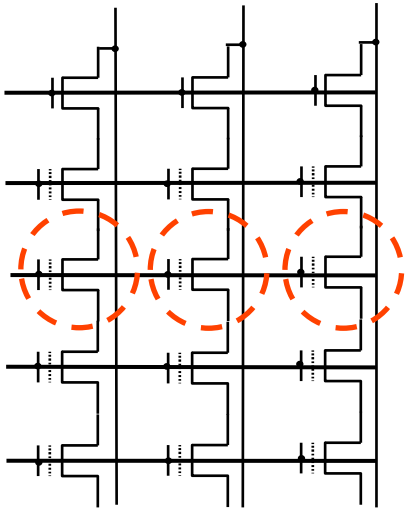

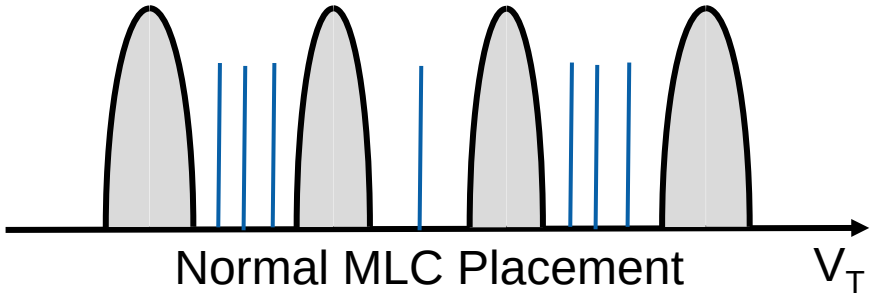

Reads with Soft Information

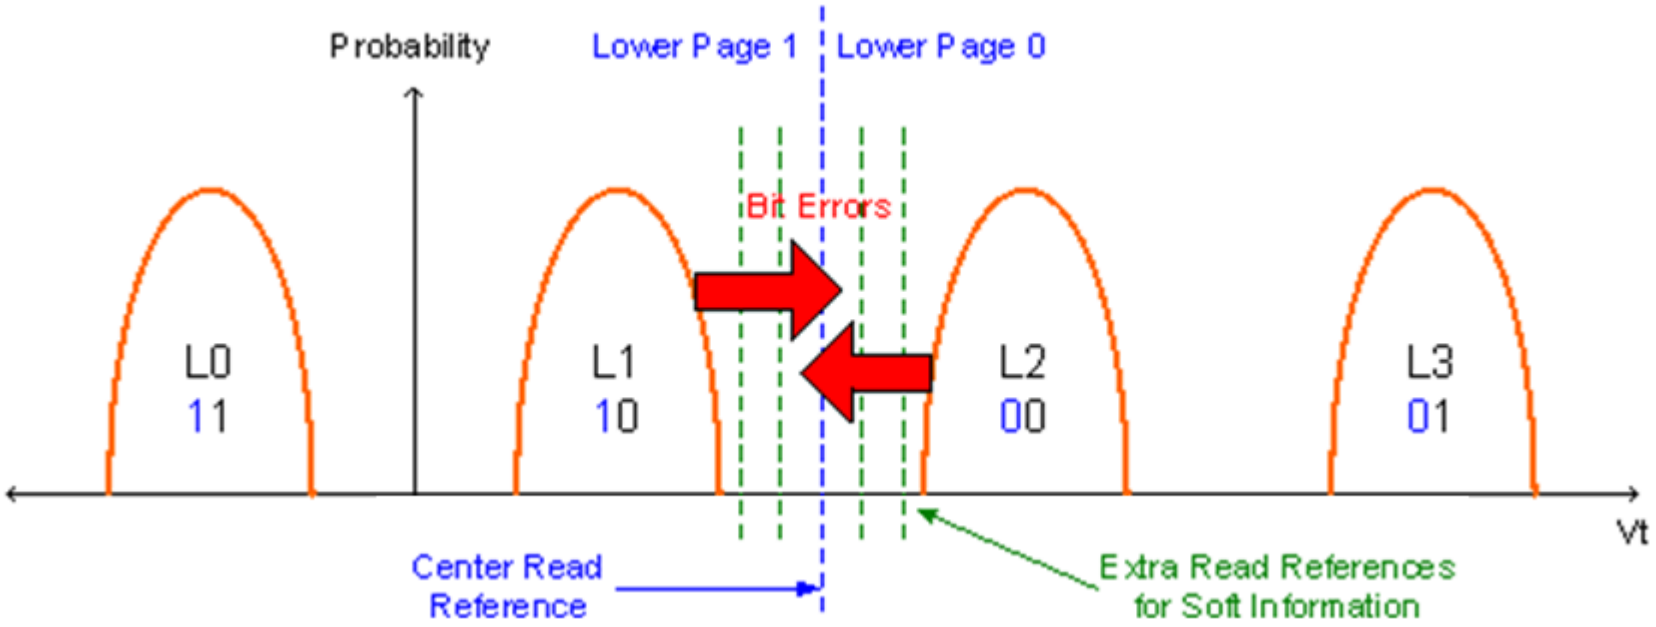

Soft information allows for advanced ECC engines, such as Low Density Parity Check (LDPC) codes.

Controller-Managed Nand Algorithms

- NAND programming algorithms are generally static through cycling
- Intrinsic NAND limitations will cause program distributions to widen during cycling.
- Controller-managed programming has the potential to optimize the NAND better.

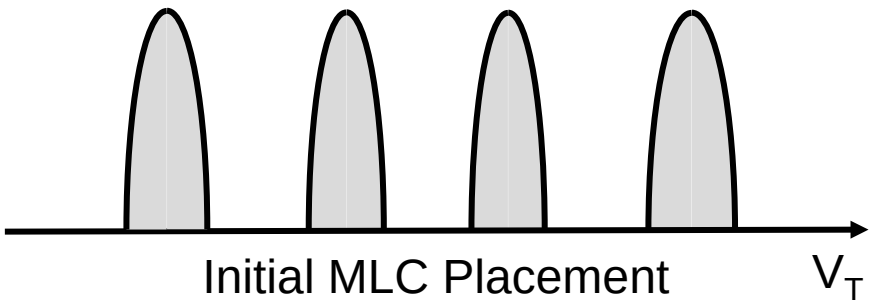

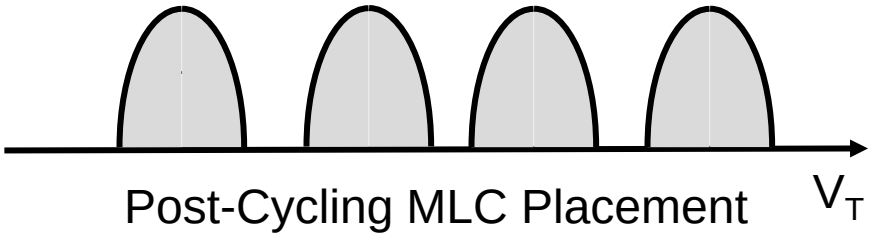

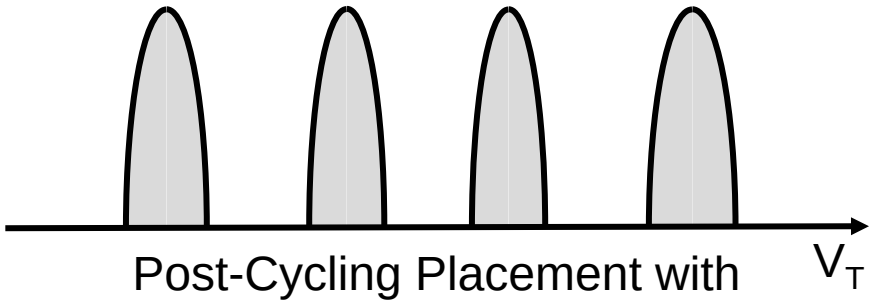

Part IV

- Overview of data integrity metrics
- Nand reliability basics
- Challenges as Nand scales to 20nm and below
- Nand-related media optimization
- **Future trends**

Scaling Trends

- Use of a planar cell (vs traditional wrap-around control gate) can mitigate some scaling challenges
  - Thinner/shorter floating gate reduces cell-to-cell interference
  - Planar cell can reduce the high fields present to enable greater endurance
  - Shorter cell is more fab friendly and easier to scale
- The search will continue for a next-generation non-volatile memory with DRAM-like speed and further ability to scale with lithography advances

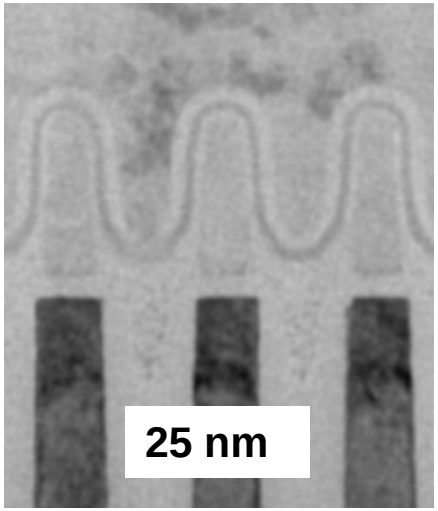

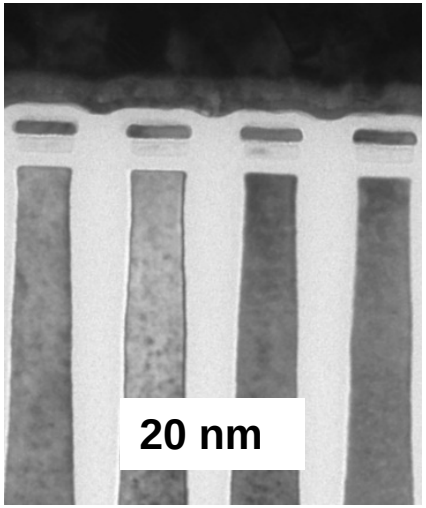

Media Management Predictions

- SSDs continue to differentiate, with no change in the availability of high-endurance MLC drives
- ECC engines will take advantage of soft information
- Media management complexity continues to increase, but possibly becoming more vendor specific
- Quality of Service will become more relevant and limiting

Scaling NAND to 20nm and below will blur the line between component and system.

Call to Action

- Avoid skepticism and seek understanding.
- 20nm NAND and SSDs are coming. New price points will enable even further adoption of SSDs into the storage stack.
- Understand your usage model and endurance requirements. Innovate around application needs.
- With knowledge and understanding, embrace 20nm SSDs and proceed with confidence.

Legal Disclaimer

- INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.
- Intel may make changes to specifications and product descriptions at any time, without notice.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
- Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to: http://www.intel.com/products/processor_number
- Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.
- Intel, Sponsors of Tomorrow and the Intel logo are trademarks of Intel Corporation in the United States and other countries.
- *Other names and brands may be claimed as the property of others.
-

Risk Factors

The above statements and any others in this document that refer to plans and expectations for the second quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Words such as “anticipates,” “expects,” “intends,” “plans,” “believes,” “seeks,” “estimates,” “may,” “will,” “should,” and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel’s actual results, and variances from Intel’s current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be the important factors that could cause actual results to differ materially from the company’s expectations. Demand could be different from Intel's expectations due to factors including changes in business and economic conditions, including supply constraints and other disruptions affecting customers; customer acceptance of Intel’s and competitors’ products; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Potential disruptions in the high technology supply chain resulting from the recent disaster in Japan could cause customer demand to be different from Intel’s expectations. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel’s response to such actions; and Intel’s ability to respond quickly to technological developments and to incorporate new features into its products. The gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; product mix and pricing; the timing and execution of the manufacturing ramp and associated costs; start-up costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; product manufacturing quality/yields; and impairments of long-lived assets, including manufacturing, assembly/test and intangible assets. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel's products and the level of revenue and profits. The majority of Intel’s non-marketable equity investment portfolio balance is concentrated in companies in the flash memory market segment, and declines in this market segment or changes in management’s plans with respect to Intel’s investments in this market segment could result in significant impairment charges, impacting restructuring charges as well as gains/losses on equity investments and interest and other. Intel's results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Intel’s results could be affected by the timing of closing of acquisitions and divestitures. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. An unfavorable ruling could include monetary damages or an injunction prohibiting us from manufacturing or selling one or more products, precluding particular business practices, impacting Intel’s ability to design its products, or requiring other remedies such as compulsory licensing of intellectual property. A detailed discussion of these and other factors that could affect Intel’s results is included in Intel’s SEC filings, including the report on Form 10-Q for the quarter ended April 2, 2011.

Thank You!