# Improving Endurance in 3D-NAND Flash

Roman Pletka, Ioannis Koltsidas, Nikolas Ioannou, Sasa Tomic, Nikolaos Papandreou, Thomas Parnell, Haralampos Pozidis
IBM Research – Zurich Research Laboratory

Aaron Fry, Tim Fisher
IBM Systems

# Agenda

- Background on Flash Management
  - How can different Flash management components benefit from each other ?
- 3D NAND characteristics
  - What are the similarities and differences between 2D and 3D NAND Flash ?
- Consequences on Flash Management
  - Data placement and garbage collection strategies
  - Threshold voltage shifting
  - Wear leveling techniques



**Disclaimer: Results in this presentation are not specific to a particular product or a Flash memory vendor**

# Enterprise Flash Commitment

Enterprise all-Flash arrays

- Consistent sub-ms latency
- 5 – 7 years of endurance under heavy writes
- Enterprise reliability

## How?

- Purpose-built hardware (FPGAs)
- Hardware-only data path
- Purpose-built ECC schemes
- Advanced Signal Processing
- State-of-the-art Garbage Collection and Data Placement
- Fine-grained Health Management

# Flash Management Components

## Dependencies and Synergies

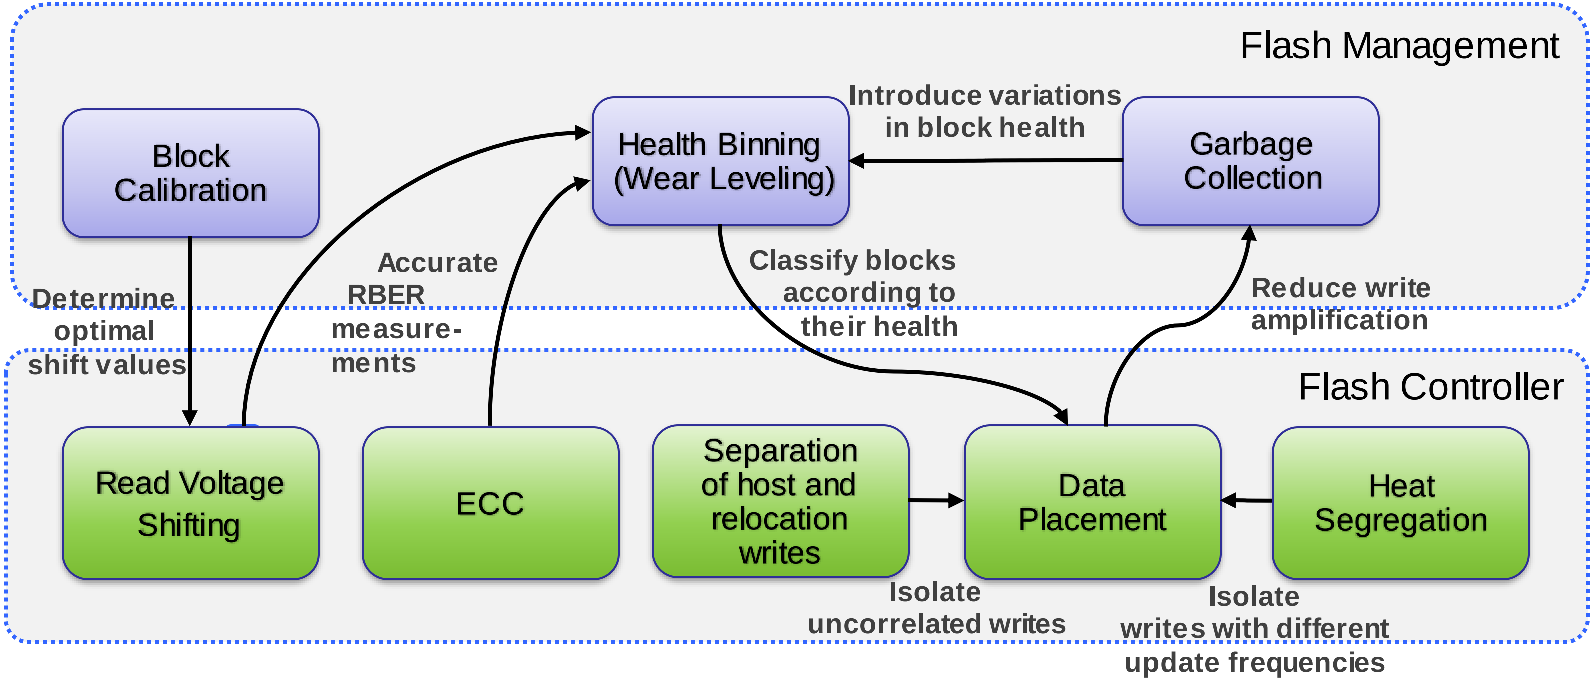

# Data Placement Strategies

Significant write amplification reductions can be achieved for skewed workloads through adequate data placement strategies:

- Separation of host and relocation writes:
  - Typical workloads show no correlation between the update frequencies of host and relocation writes
- Heat Segregation:
  - Heat: Tracking update frequency at LBA granularity

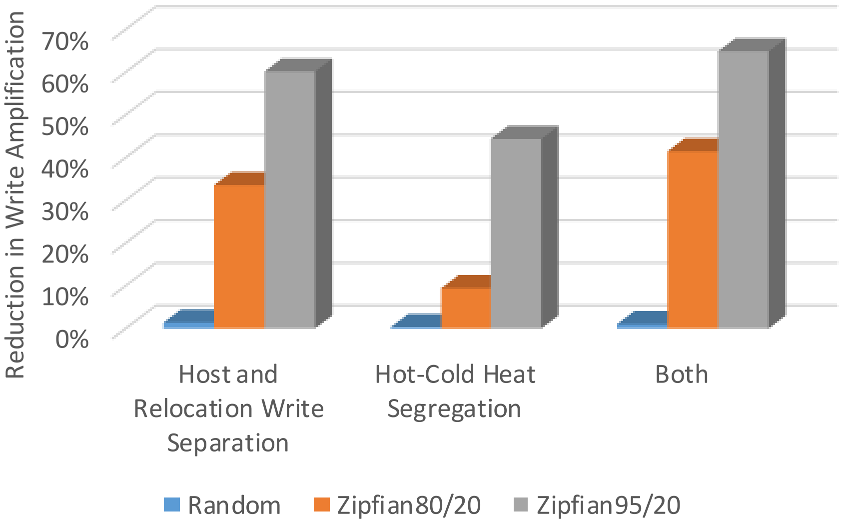

# Endurance Characteristics

- In 2D NAND Flash, significant differences in the RBER over time have been observed among blocks [1]
  - Some blocks have almost twice the endurance of others
  - Low RBER at early life does not indicate a good block, and an early high RBER not a weak one!

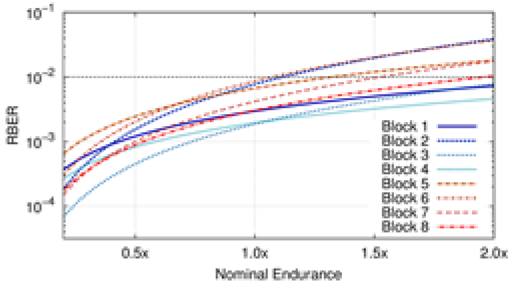


RBER of different consumer-level flash blocks in the same device as a function of P/E cycles


[1] Health Binning: Maximizing the Performance and Endurance of Consumer-level NAND Flash R. Pletka, S. Tomic, SYSTOR 2016


- The same effects exist in 3D NAND:
  - Comparison of the average block endurance using the same ECC capable of correcting in the order of $10^{-2}$ errors

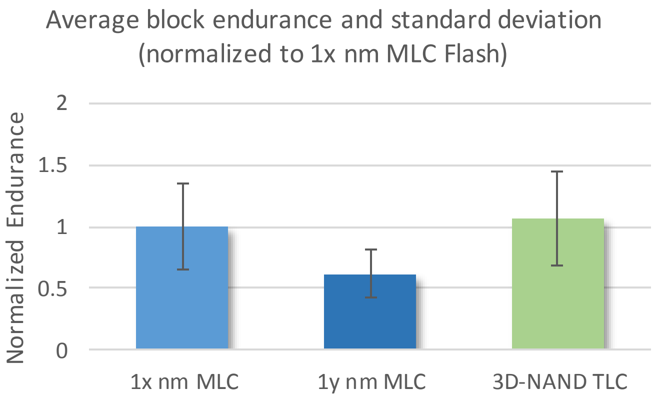


Flash Memory Summit

# From Traditional Wear Leveling to Health Binning

IBM

**Dynamic WL**

Balance P/E cycles across blocks upon overwrites and relocations. Typically uses the least worn available block to place new data.

- Introduce data placement with stream segregation
- Use better blocks for hotter data

**Static WL**

Identifies the least worn blocks holding static data in the background. Still valid data is relocated to another block causing an increase in write amplification.

- Reduce Static WL to address retention and read disturb limitations of Flash
- Perform relocations instead of block swapping

**P/E Cycle-based WL**

Balances wear of blocks based on their program-erase cycle count only.

- Background grading of blocks based on RBER
- RBER estimation based on ECC feedback


**Endurance gains of up to 60% with 3D TLC NAND !**

# Read Voltage Shifting

Read Voltage Shifting has been proposed in the past [2]:

- Dynamic Read Level Shifting requires special access modes to Flash.
- Extensive characterization is required to determine behavior of read levels under different conditions.
- Read level shift values depend on:
  - Number of P/E cycles of the block
  - Number of reads a page has seen since programmed
  - Retention time
  - Individual block/page characteristics
- Block Calibration: Optimal read levels must be continuously updated in the background which takes a non-negligible amount of time.

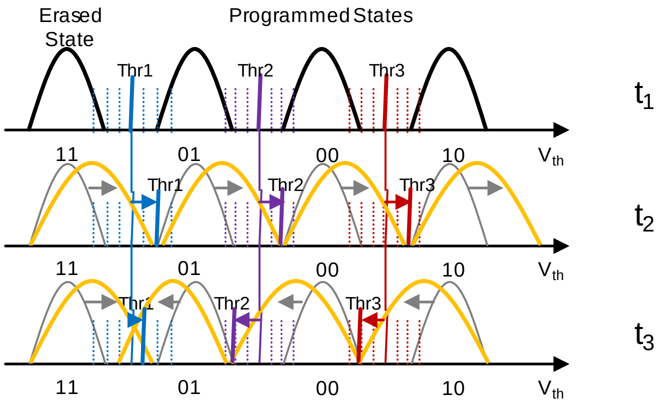


Benefits:

- Dynamic read level shifting significantly contributes to maximize flash endurance: Gains of 3x in endurance achievable!
- Calibration in the background does not impact host reads and writes.
- Use special techniques to reduce meta-data overhead.

Typical 2D NAND behavior:

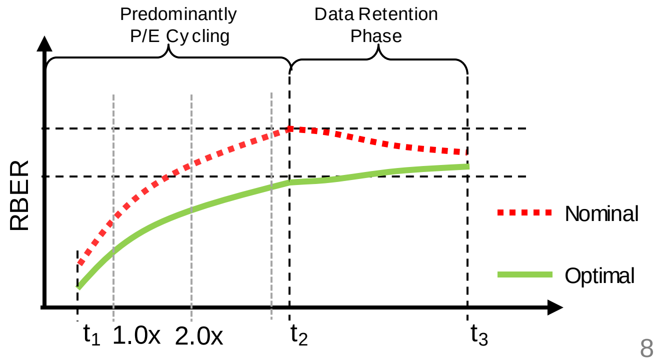


[2] Using Adaptive Read Voltage Thresholds to Enhance the Reliability of MLC NAND Flash Memory Systems, N. Papandreou, T. Parnell, H. Pozidis, T. Mittelholzer, E. Eleftheriou, C. Camp, T. Griffin, G. Tressler, A. Walls, GLSVLSI 2014

# Block Calibration Challenges

- Results from a characterization experiment showing the evolution of the RBER at specific points in time using optimal shift values.
  - RBER increases during cycling and further deteriorates during the retention phase
  - If read voltage levels are not adapted periodically as well as after erasing and reprogramming the block, the RBER continues to increase!

Typical 3D NAND behavior:

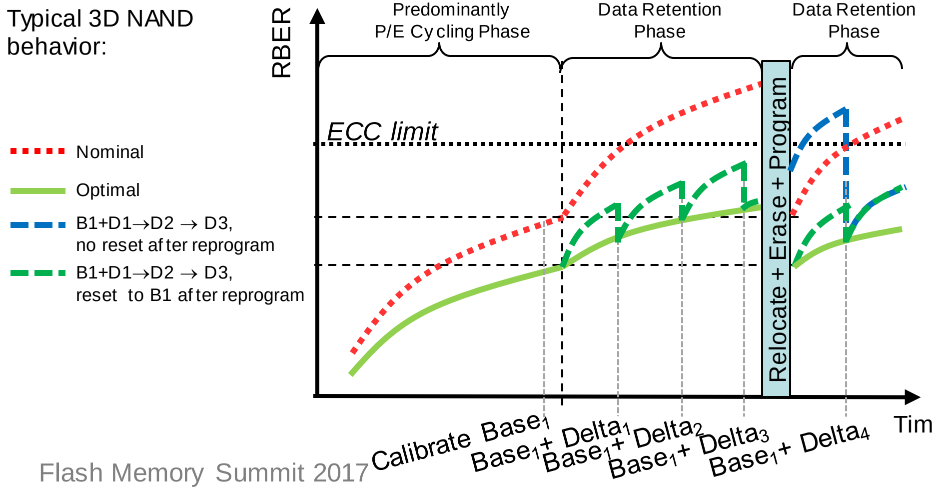


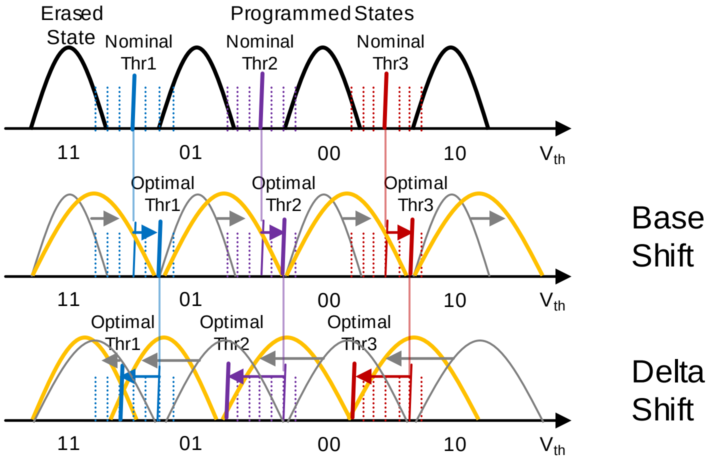


How can this be addressed ?

- Maintain separate shift values to track contributions to the threshold voltage distribution from:
  - permanent changes (e.g., P/E cycling) ⇨ Base shift
  - temporary changes w.r.t. the block erase count (e.g., retention, read disturb effects, …) ⇨ Delta Shift
- Reset Delta shift values after a block erase

# Transient effects on the RBER

- Relative contributions to the RBER from retention and read disturbs
- Observations:
  - 2D MLC Flash @ EOL behaves similarly as 3D TLC @ beginning of life
  - 3D TLC has much higher relative contributions to the RBER from transient effects
- Consequences on Flash 3D Flash Management:
  - Equalize block health with Health Binning to keep RBER differences between blocks due to permanent effects as small as possible
  - Additional data relocations are needed which may affect overall endurance

2D MLC NAND: 3D TLC NAND:

Retention Characteristics:

Read Disturb characteristics:

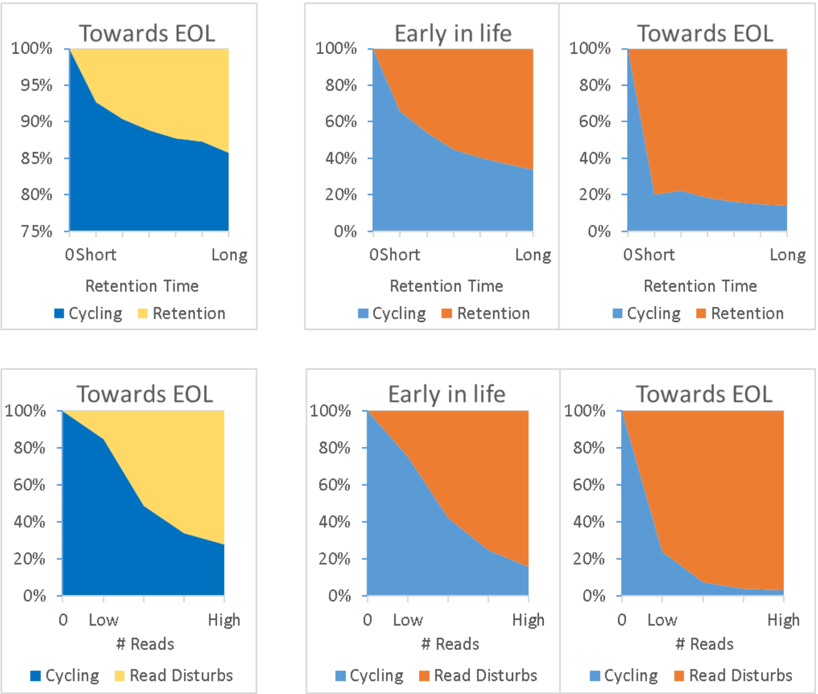

# Conclusion

- 2D vs. 3D NAND Flash characteristics
  - Similar block variability observed in 3D NAND compared to 2D NAND 1x and 1x nm generations
  - Transient effects on the RBER dominate even early in life and require careful Flash management in 3D NAND
- Efficient Flash management techniques to address these challenges have been outlined:
  - Data placement with separation of host and relocation writes combined with heat segregation
  - Health Binning using ECC feedback to grade blocks for data placement
  - Read voltage shifting with background block calibration

# Thanks !

## Questions ?

www.research.ibm.com/labs/zurich/cci/