Reducing Read Latency Fluctuations of Flash Storage Systems Using Preemptible Programs and Erases

Jisung Park, Jaehoon Lee, Myungsuk Kim, Myungjun Chun and Jihong Kim

Department of Computer Science and Engineering, Seoul National University

Read Performance Fluctuations in SSDs 1

- Read response time can be increased **23x** over the average.
  - Despite *out-of-order I/O scheduling* and *preemptive GC*
  - Page reads can be blocked by programs and erases whose latencies are longer in 2-3 orders of magnitude.
  - In particular, the block erasure time has been significantly increased (e.g., up to 45 ms) in 3D NAND flash memory.

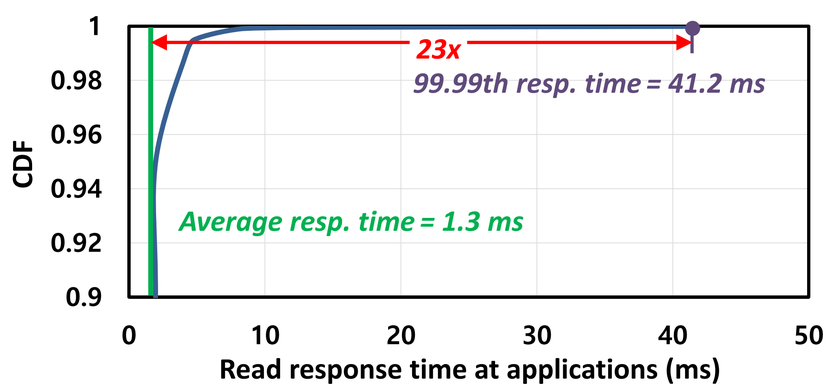

<CDF of read latency>

Immediate Suspension/Resume Technique 2

- Preemption of on-going program/erase operations
  - For avoiding reads from being delayed by program and erase operations
  - Can reduce page read latency by 99% over when a page read is blocked by an erase operation

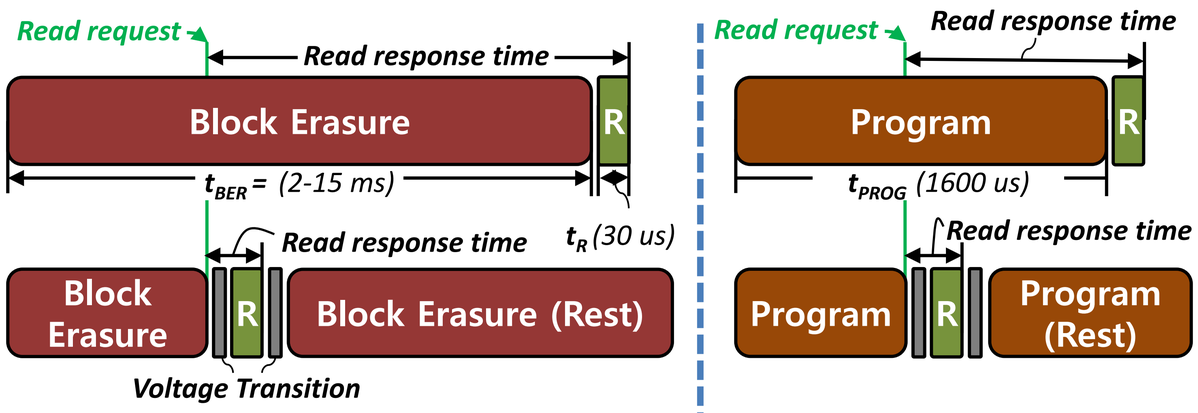

<Erase Suspension> <Program Suspension>

Problems of Immediate Suspension/Resume 3

- Degradation of NAND reliability
  - *Over-erasure* and *over-program* problems
  - Due to the *increased* number of voltage applications
  - In modern TLC NAND flash memory, it results in data loss by increasing bit errors.

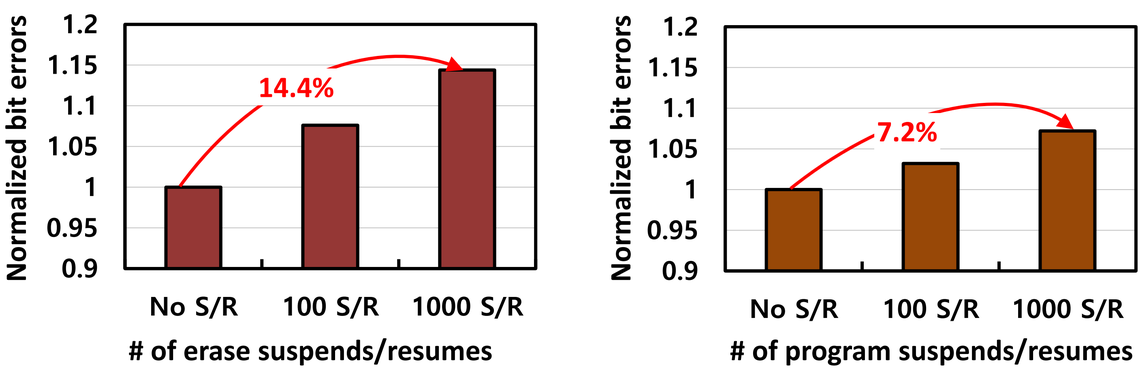

<Normalized Bit Errors After 1K P/E Cycles>
(Measured from over 1200 Blocks out of 80 NAND Chips)

Delayed Suspension & Immediate Reset 4

- Delayed suspension/resume
  - Delays suspensions until finishing an atomic inner loop to limit the maximum number of suspensions
  - No lifetime loss, but increased suspension time

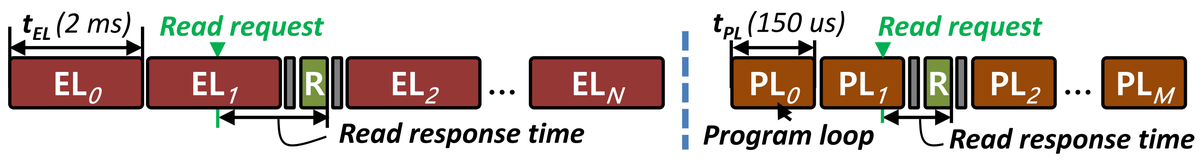

- Immediate reset
  - Immediately stops operations and starts again from the beginning (at another page for program operation)
  - Fast preemption, but additional P/E increases

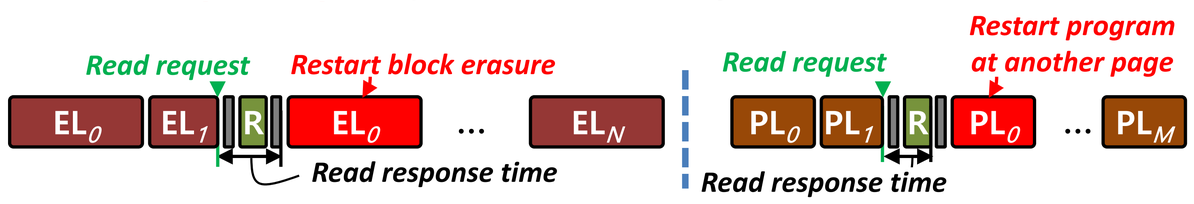

ppFTL: Priority-aware Preemption 5

- Hybrid approach based on reset and suspension
  - With priority information of read requests

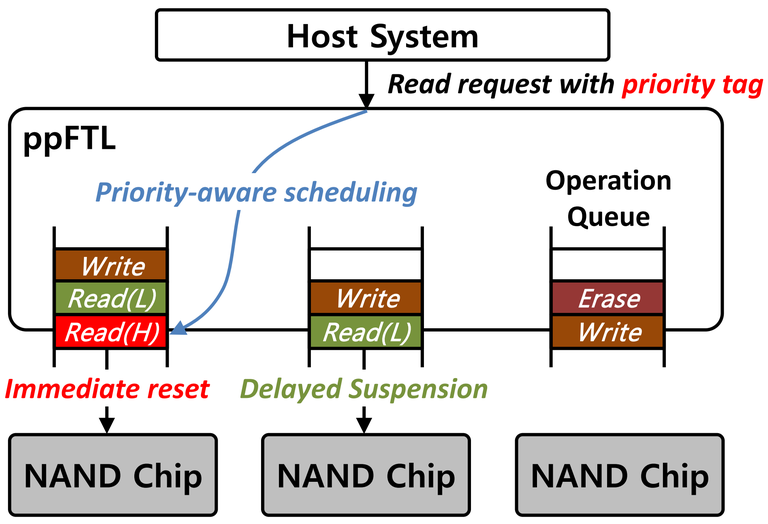

<An Overview of ppFTL)>

Experimental Results 6

- Performance improvement: 50% over baseline
- Lifetime improvement: 20% over unconditional reset

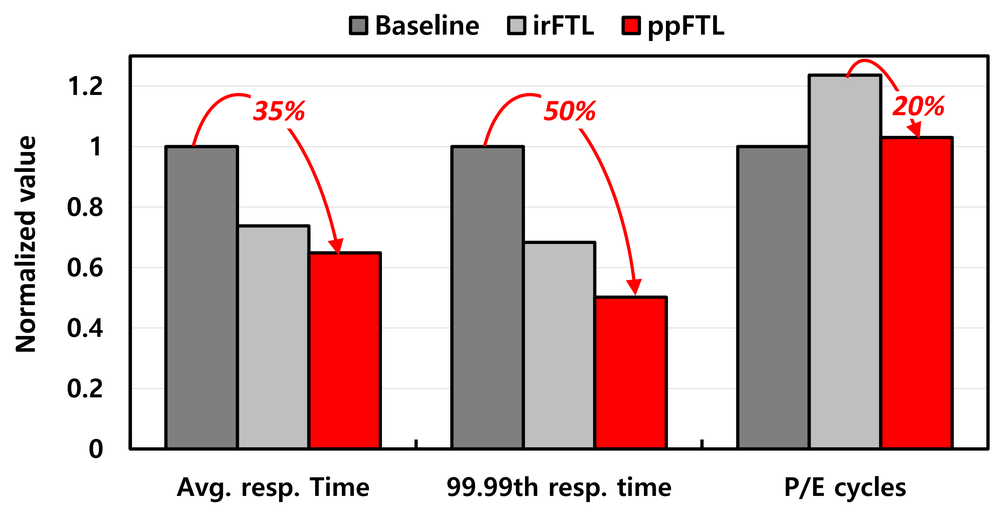

<Comparisons of three different FTLs>
(Baseline: No Preemption, irFTL: unconditional reset)