

Replicate, Reduce, Recycle: Extending the Lifetime of Flash Memory

Kuei Sun, Jamie Liu, Ryan Yang, Angela Demke Brown, Ashvin Goel

University of Toronto

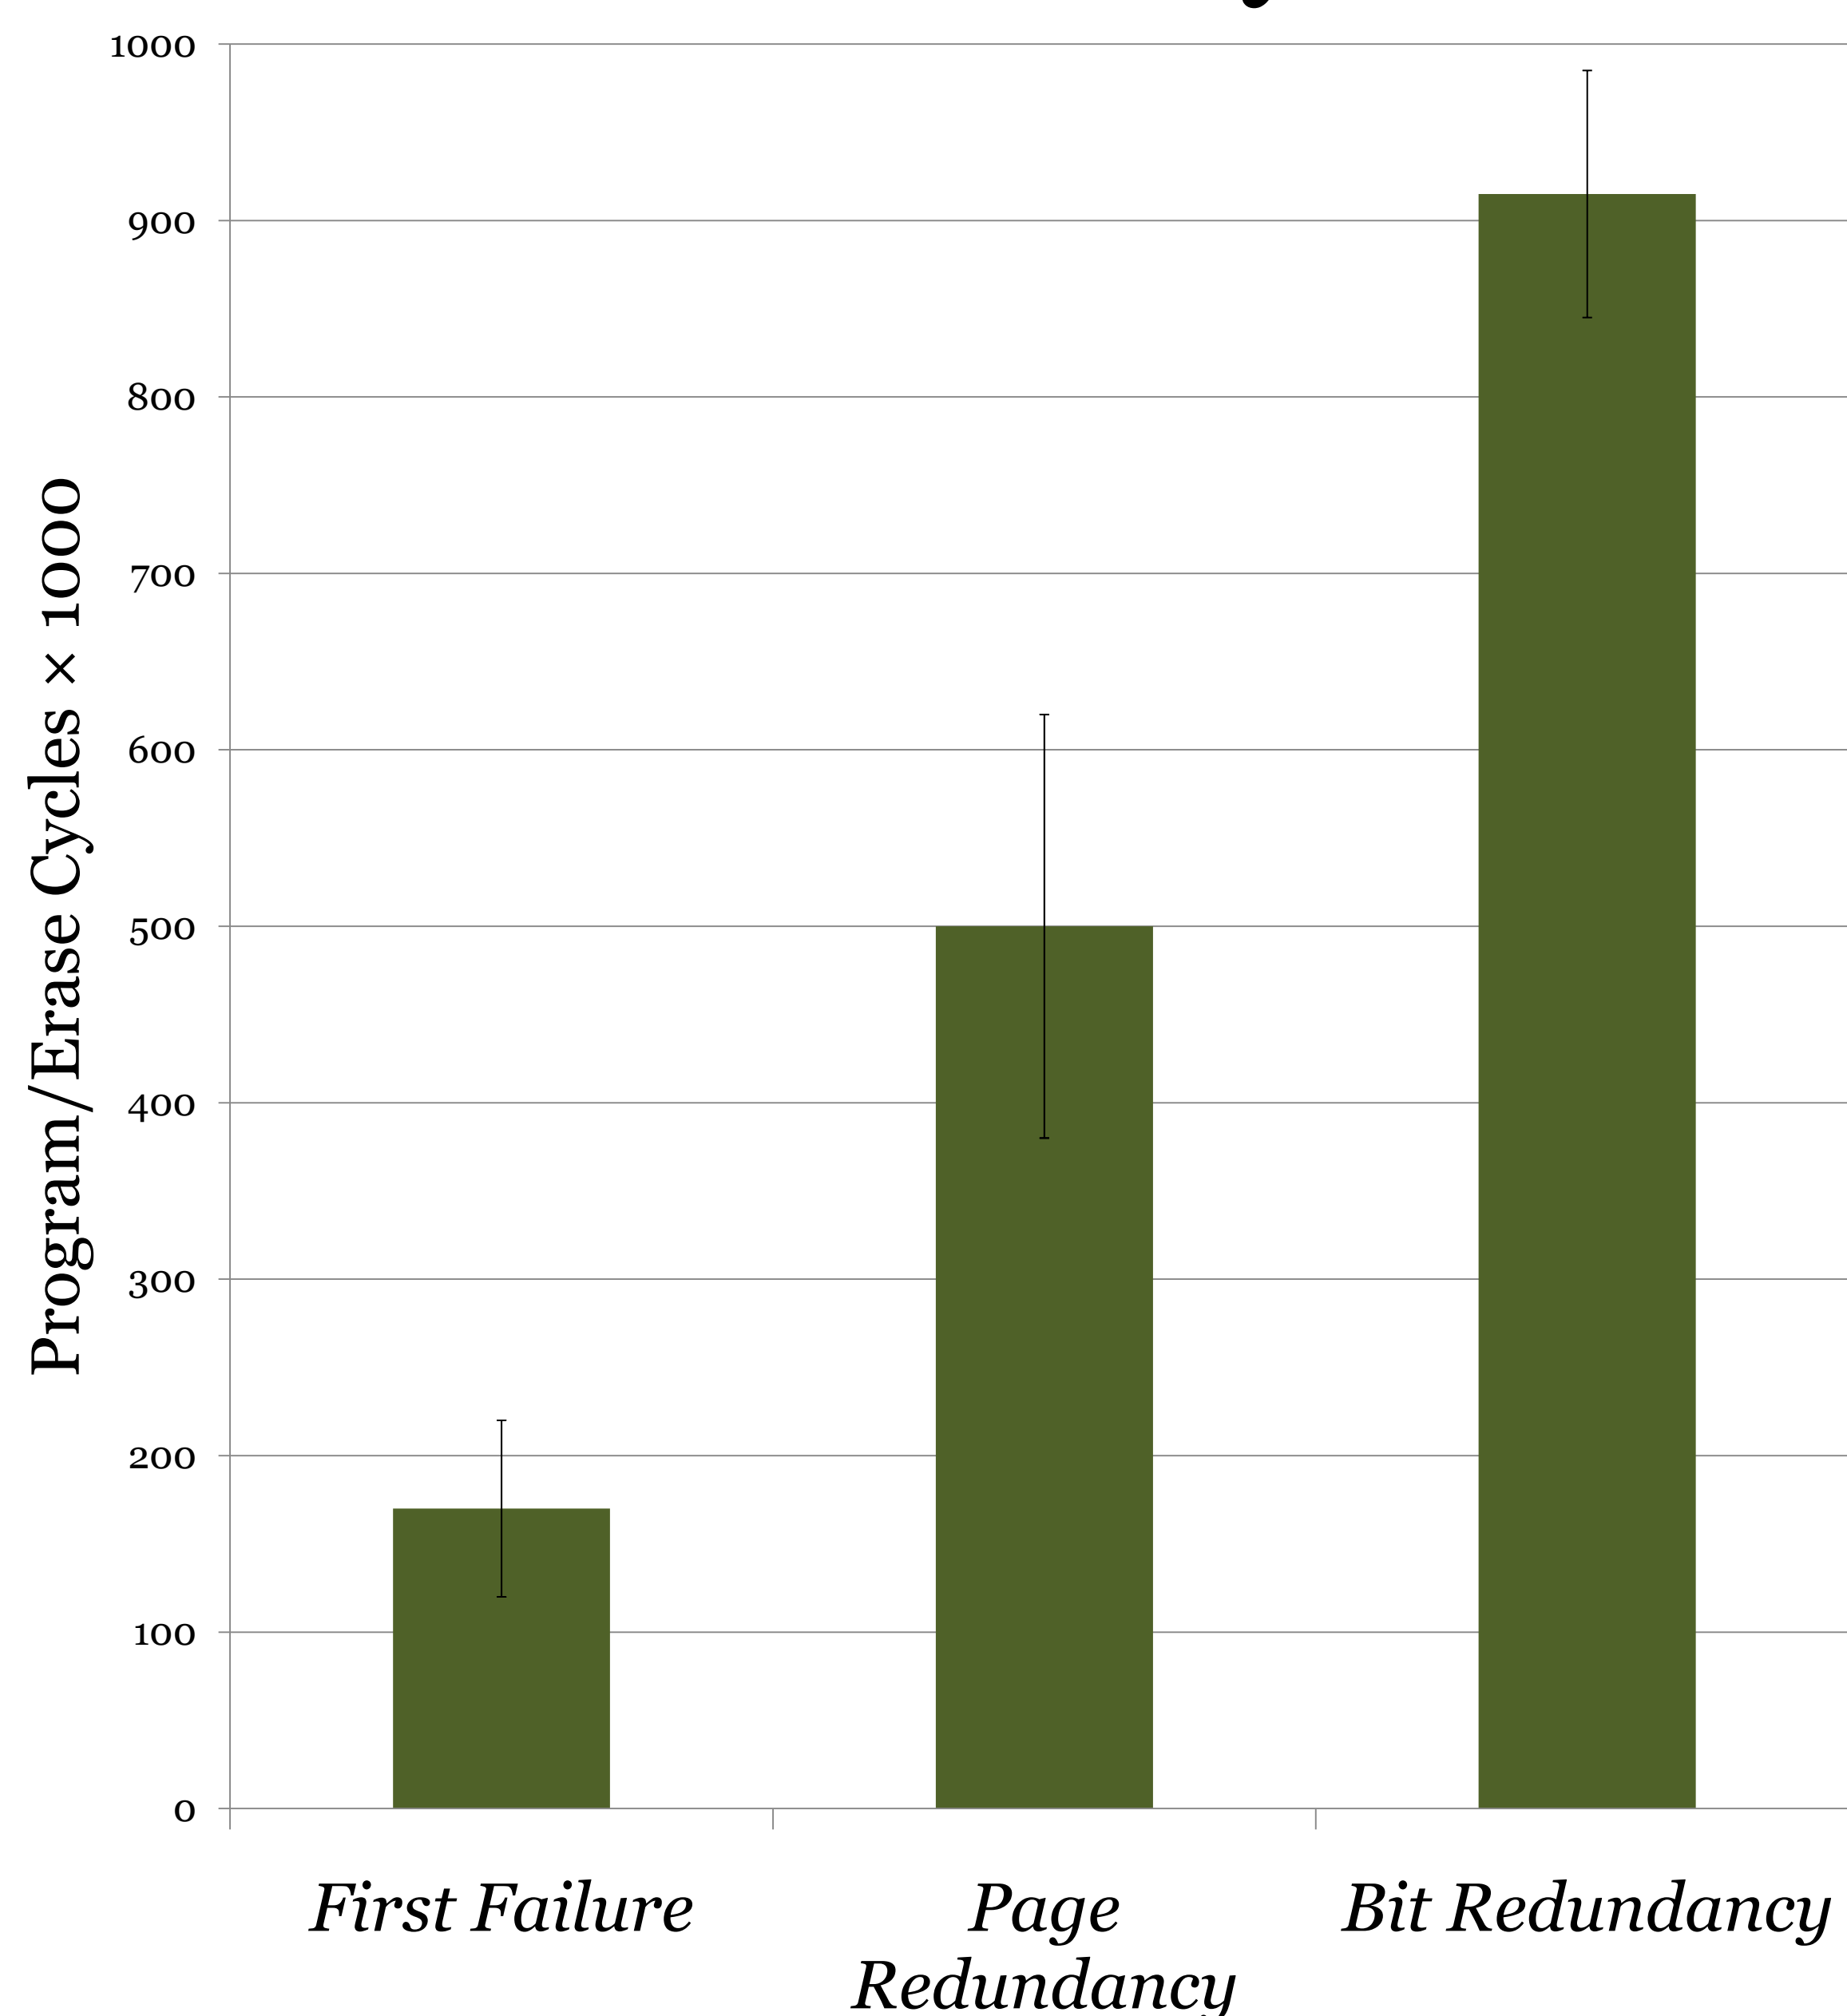
Problem

- Flash memory offers many advantage over disk drive
 - ✓ Faster speed, low energy consumption, declining cost, and higher density
- Multi-level cell (MLC) NAND flash
 - ✓ High density, can store multiple bits within a single memory unit
 - ✗ Wears out quickly (~10K p/e cycles) within weeks in write-intensive environments
- Single-level cell (SLC) NAND flash
 - ✓ Has higher durability (~100K p/e cycles)
 - ✗ Can only store one bit per cell, and is more expensive

GOAL: High Density, Low Cost, and Reliable Storage

IDEA: Use MLC flash until it fails, then recycle it as a low-density SLC flash

Results for Page and Bit Redundancy



Recycling Failed Drives

- Page Redundancy**
 - ✓ Each MLC cell is used by a pair of data pages: most significant bit (MSB) page and least significant bit (LSB) page
 - ✓ By writing a single page of data to both pages in the pair, we are writing to the same physical cells
 - ✓ We decode '00' as '0' and '01', '10', and '11' as '1'
 - ✓ The usable lifetime is increased by over **3 times**
- Bit Redundancy**
 - ✓ When each flash cell is programmed, neighboring cells's data may be "accidentally" flipped due to charge leakage, this is *program disturb*
 - ✓ We replicate each bit that is written. E.g. "010" becomes "001100"
 - ✓ The usable lifetime is also increased by over **5 times**

ALTHOUGH THE CAPACITY IS HALVED, WE HAVE MANAGED TO RECOVER AN OTHERWISE FAILED DRIVE!

Challenges

- Need to bypass Flash Translation Layer
 - FTL can mask certain block failures
 - Built memory controller based on Boboila & Desnoyers
- Need FPGA implementation
 - For precise timing and fast testing
 - Block failure testing takes time with our redundancy scheme!
- Need efficient redundancy schemes
 - Why do the redundancy schemes work so well!
 - Network coding could be used
 - Compare our approach with ECC
- Need to handle variability
 - Different SSD memory blocks and chips have significantly varying reliability characteristics
- Need to handle reduced capacity
 - Perhaps used when SSD is a caching device

