Computer Science
UNIVERSITY OF TORONTO

Improving the Reliability of Next-Generation SSDs using WOM-v codes

Shehbaz Jaffer
University of Toronto, Google

Kaveh Mahdaviani
University of Toronto

Bianca Schroeder
University of Toronto

USENIX FAST ‘22

Motivation

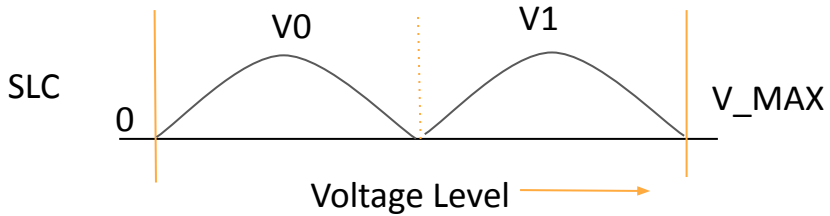

Voltage can only be increased (Program)

Voltage can be reset to zero (Erase)

Increased density – more bits per cell

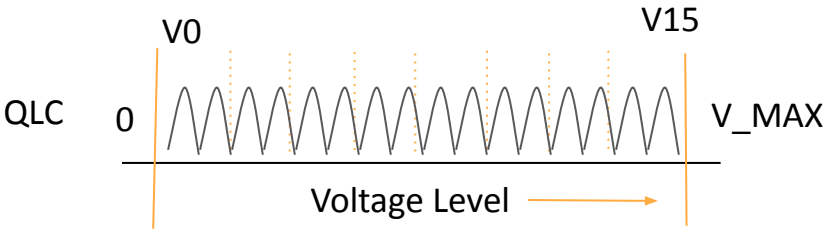

Motivation

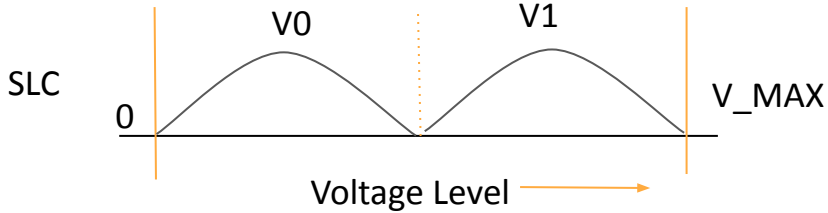

Increased density – more bits per cell

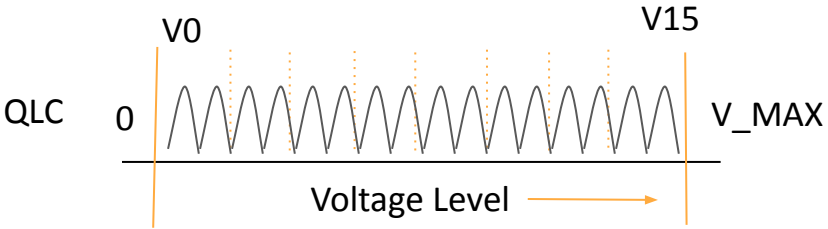

Reduced Program and Erase (P/E) cycles

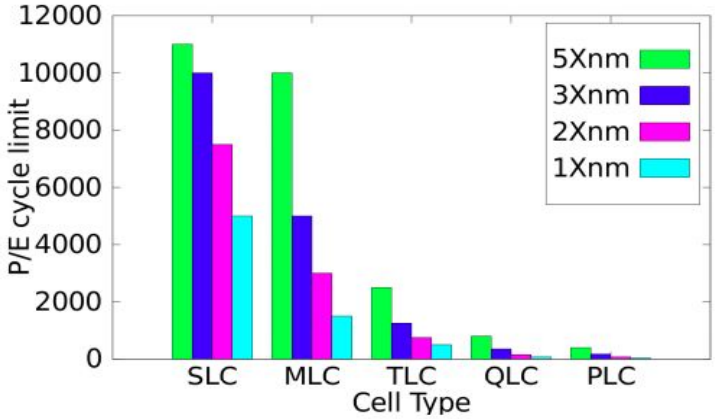

Overwrite between Erase - Write Once Memory Code

Rethinking WOM Codes to Enhance the Lifetime in New SSD Generations

Shehbaz Jaffer Kaveh Mahdaviani Bianca Schroeder
University of Toronto

WOM-v Codes

INPUT DATA

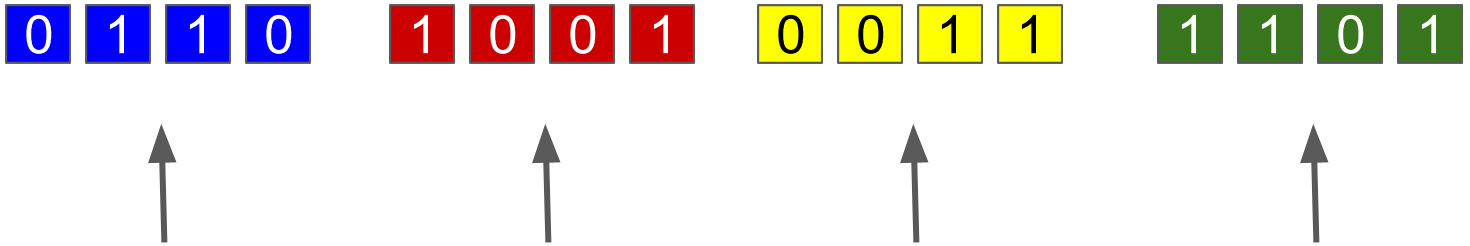

Cell Erase Count (No-WOM)

4

QLC Cell

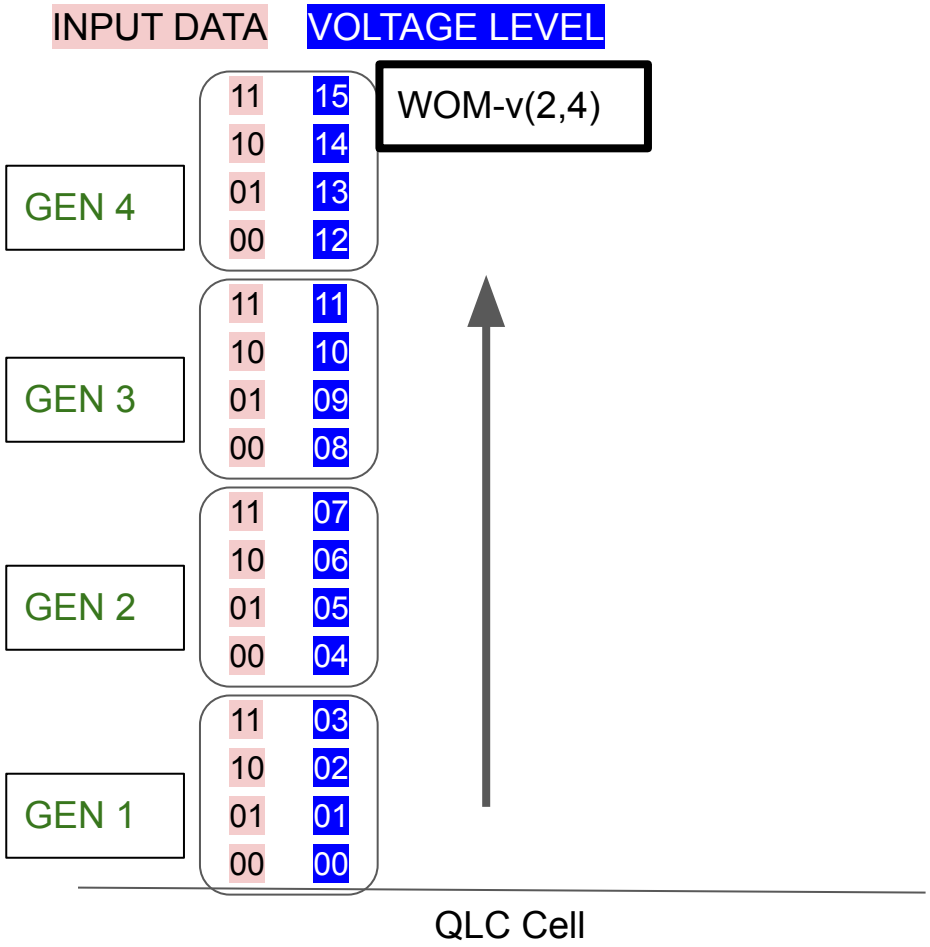

Cell Erase Count
(No-WOM)

4

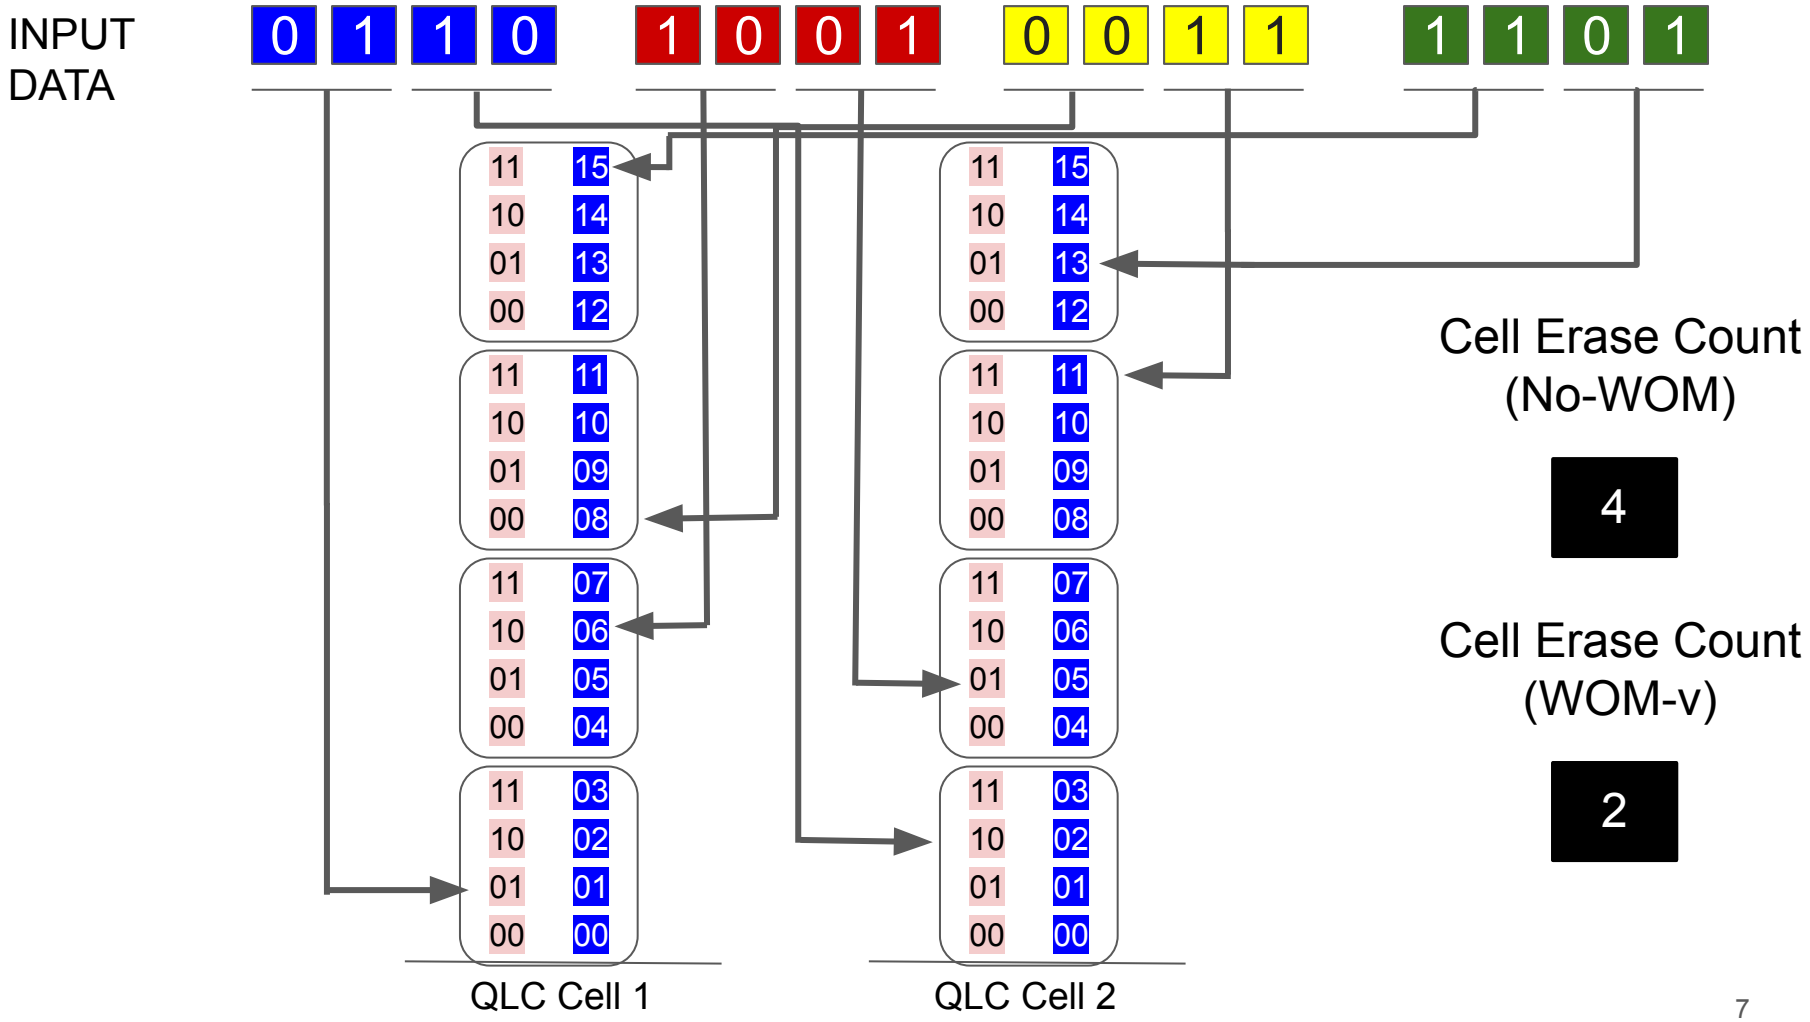
INPUT DATA
0 1 1 0 1 0 0 1 0 0 1 1 1 1 0 1
11 15
10 14
01 13
00 12
11 11
10 10
01 09
00 08
11 07
10 06
01 05
00 04
11 03
10 02
01 01
00 00
QLC Cell 1
QLC Cell 2
Cell Erase Count (No-WOM)
4
Cell Erase Count (WOM-v)
2

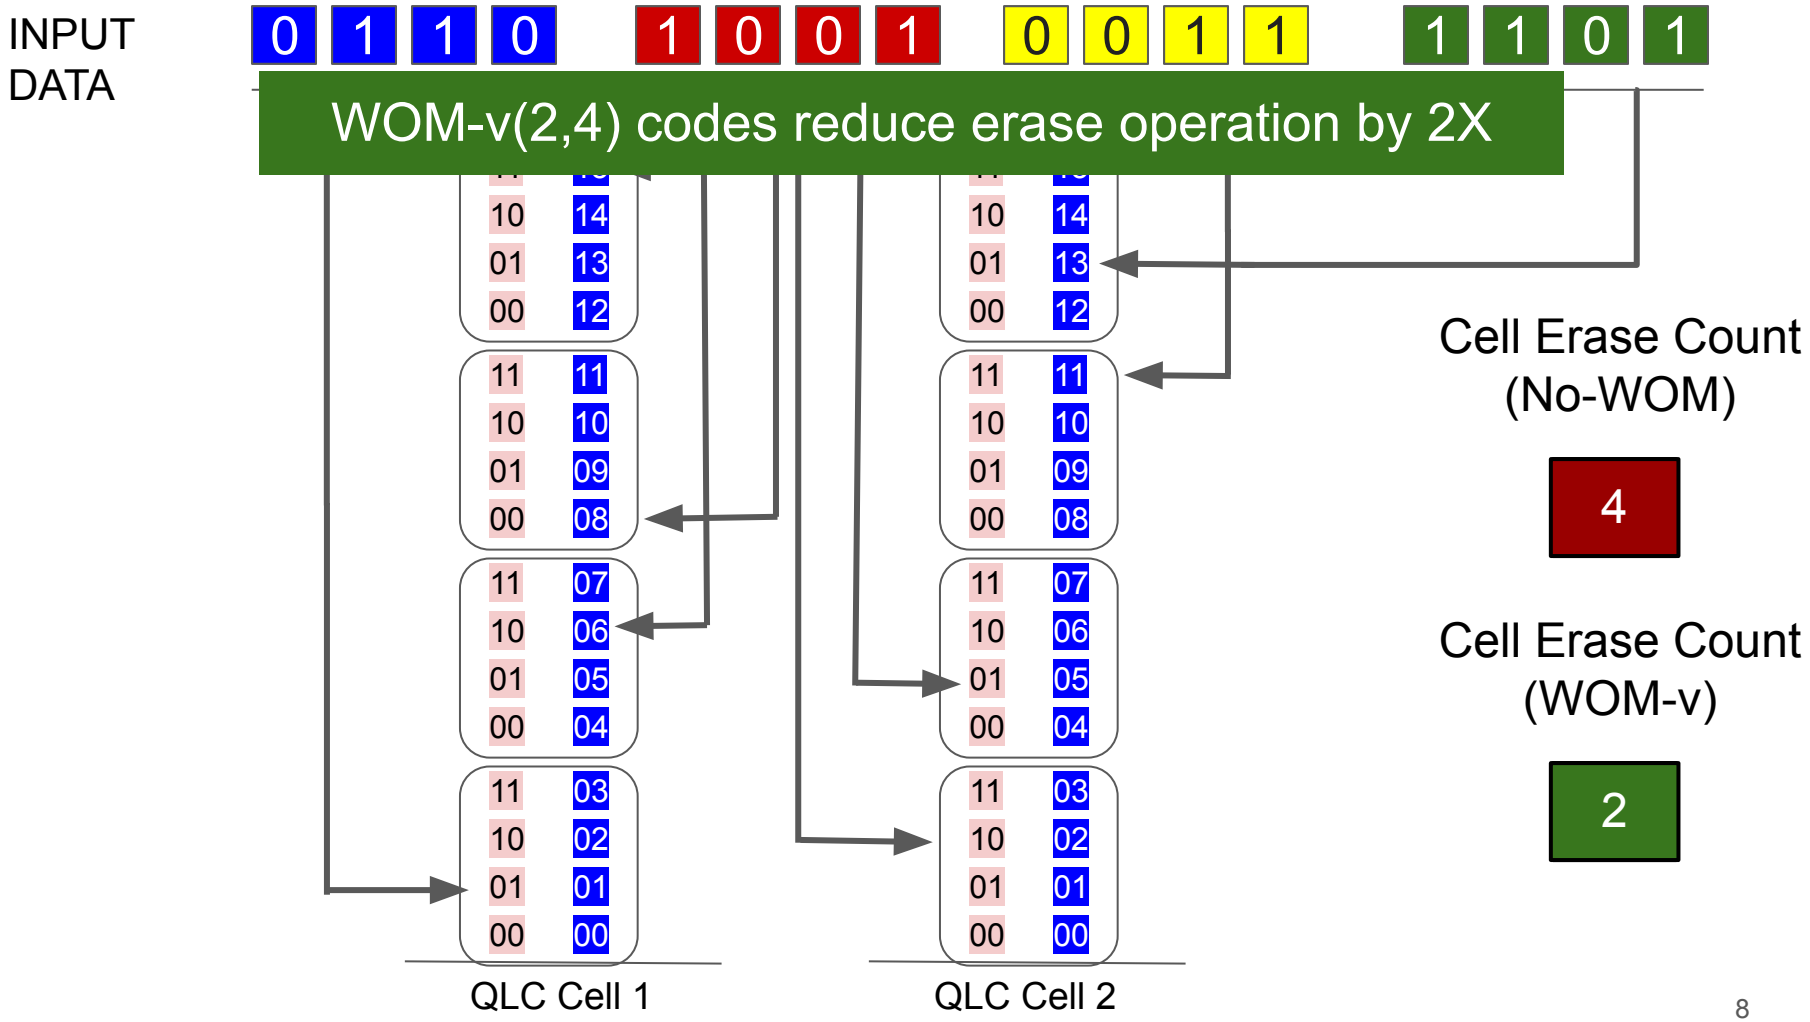
INPUT DATA
0 1 1 0 1 0 0 1 0 0 1 1 1 1 0 1
WOM-v(2,4) codes reduce erase operation by 2X
10 14
01 13
00 12
11 11
10 10
01 09
00 08
11 07
10 06
01 05
00 04
11 03
10 02
01 01
00 00
QLC Cell 1
QLC Cell 2
Cell Erase Count (No-WOM)
4
Cell Erase Count (WOM-v)
2

Optimizations

- Code Word Sharing.
- Same Generation Transition.
- Programming beyond maximum Generations.
- Optimal GC (GC_OPT) Mode.
- No-Read (NR) Mode.

Code Word Sharing

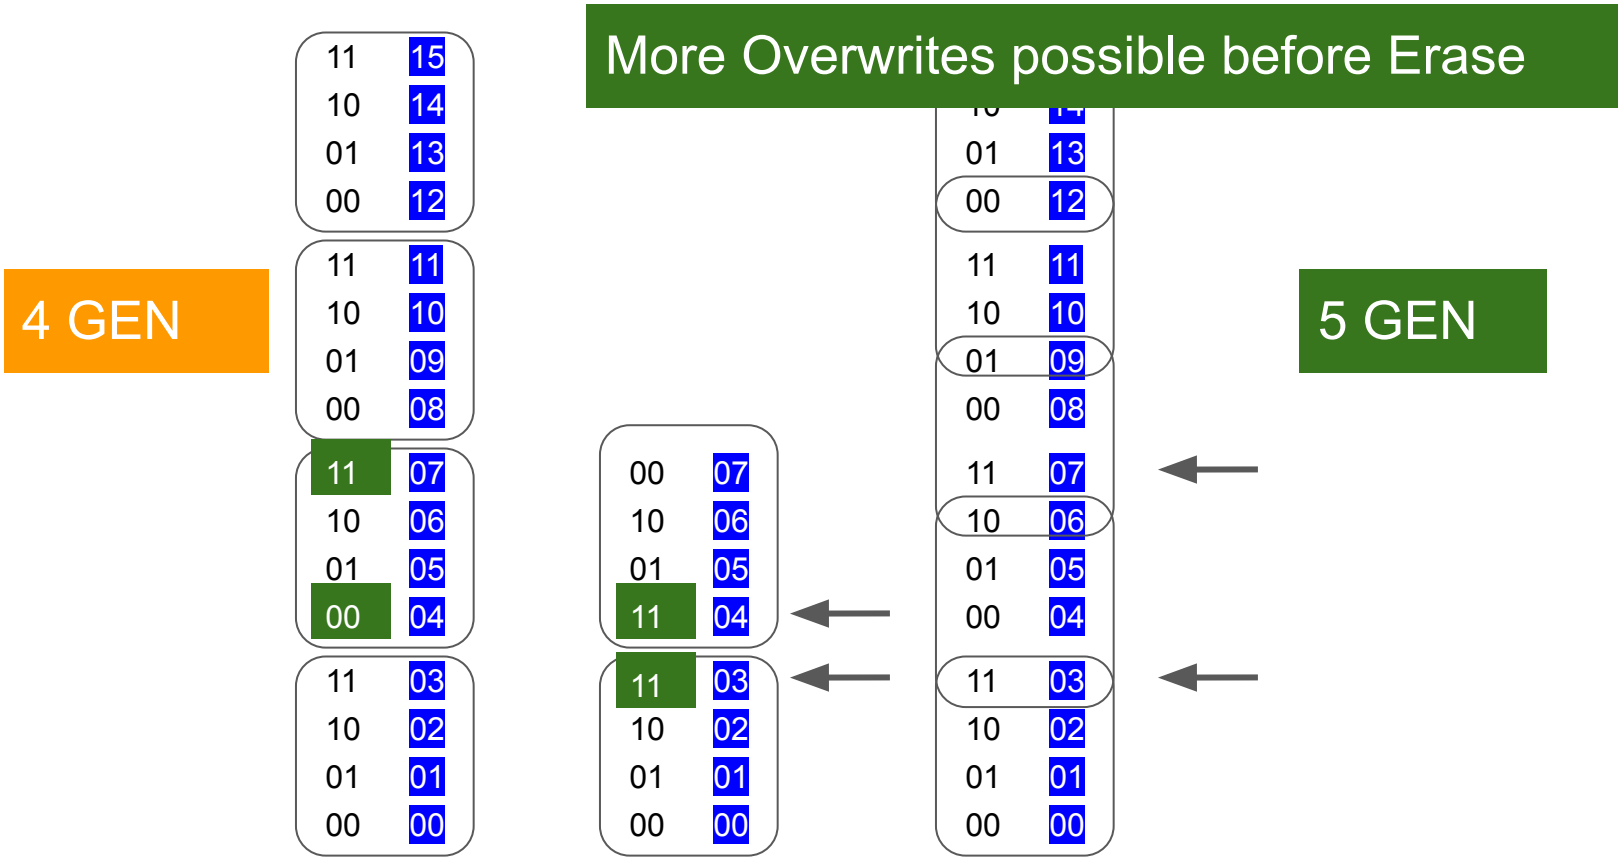

Challenge in introducing WOM-v codes to SSDs

SSD Architecture

Write Amplification

Performance and Parallelism

Workload Pattern

Implementation

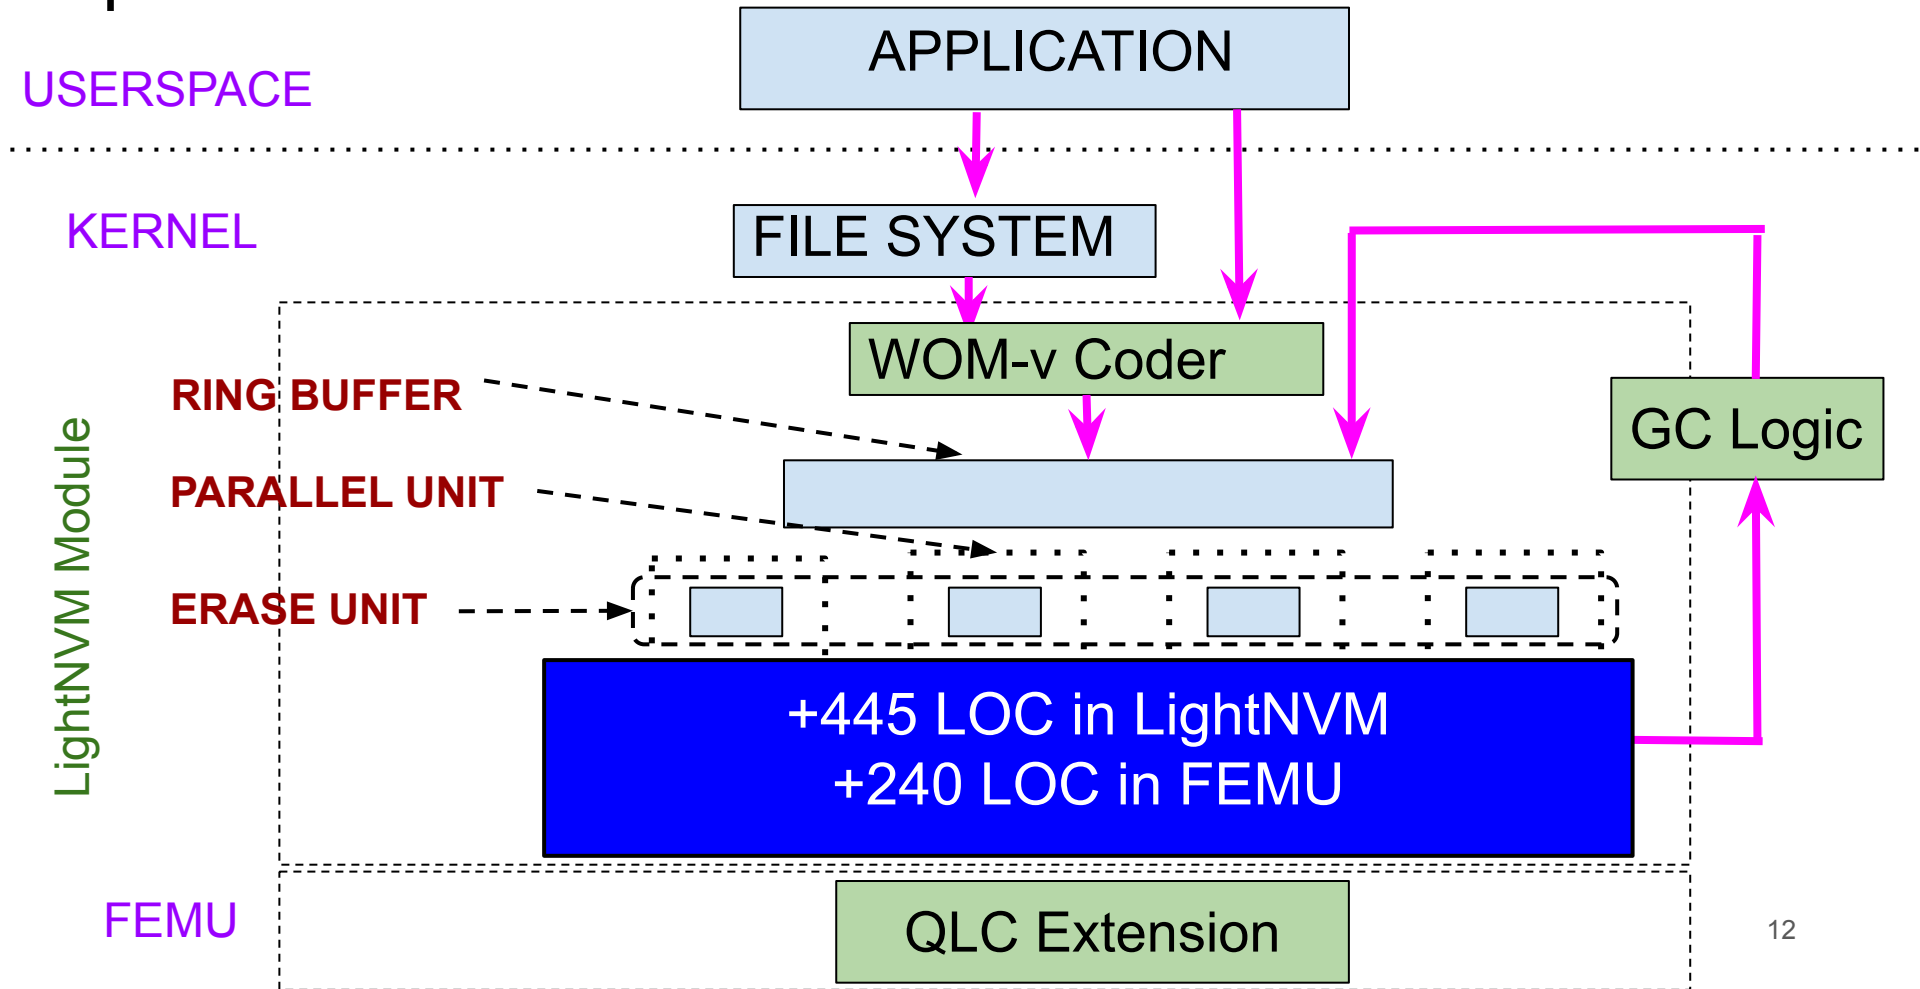

Implementation

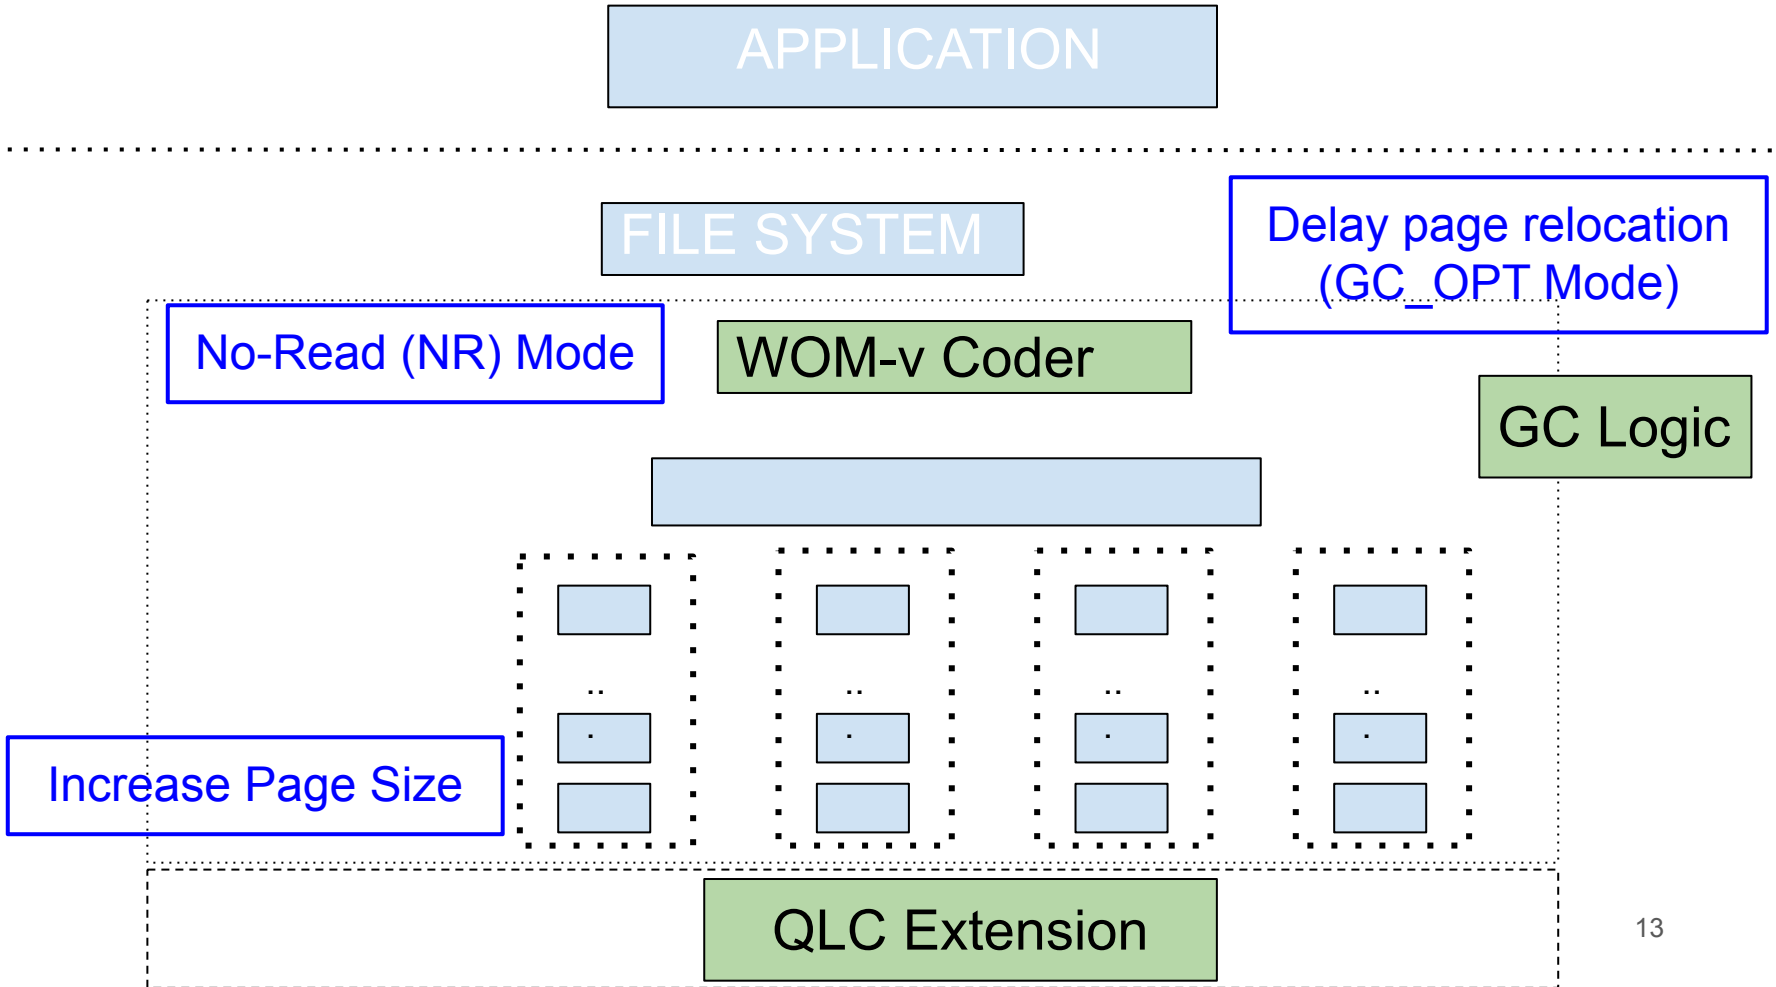

Evaluation: Real World Traces

Erase Operation Reduction

4.4 - 11.1x reduction in erase cycles

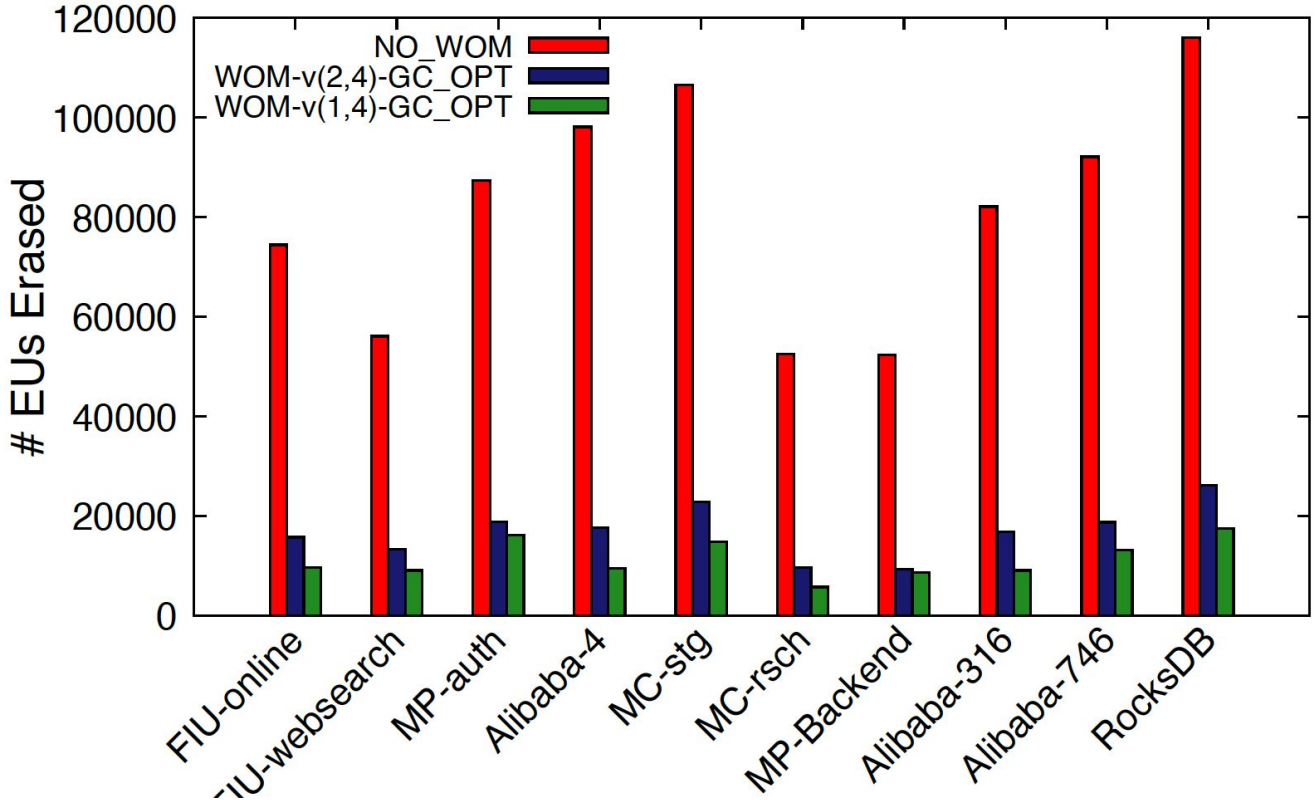

Write Performance

< 8% write performance overhead

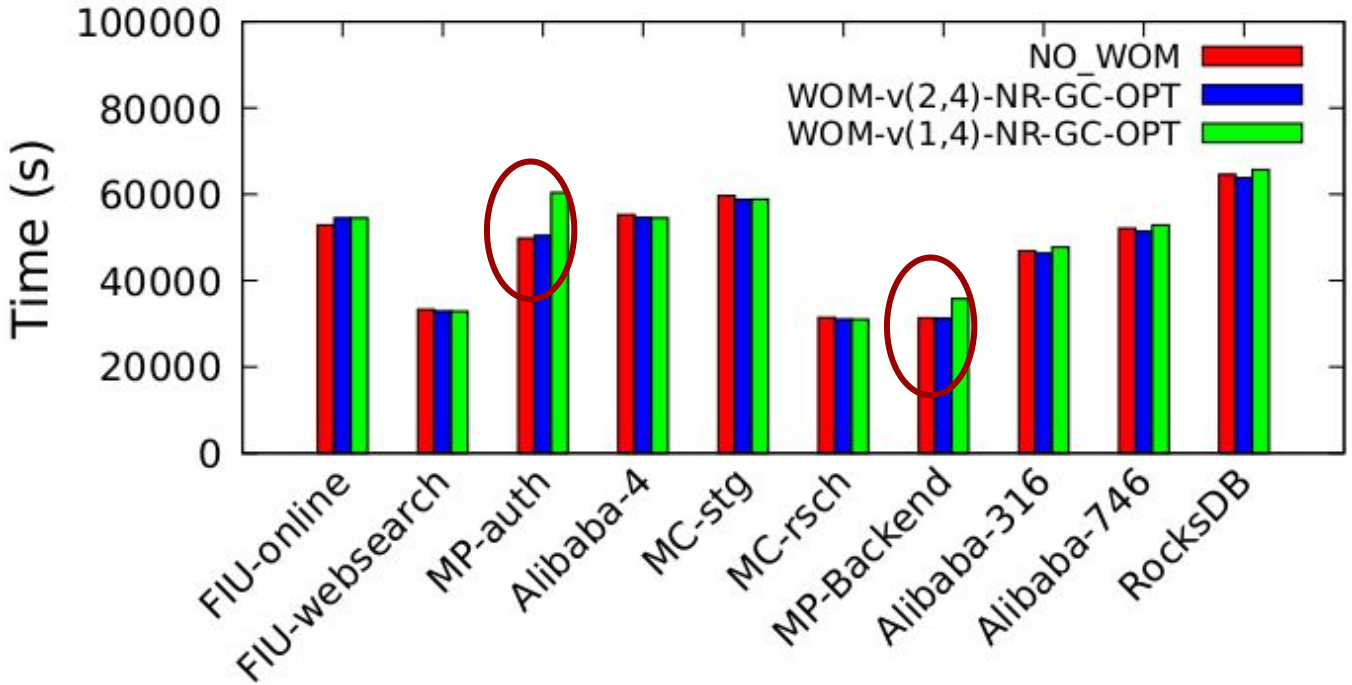

Read Performance

Negligible average read latency overhead

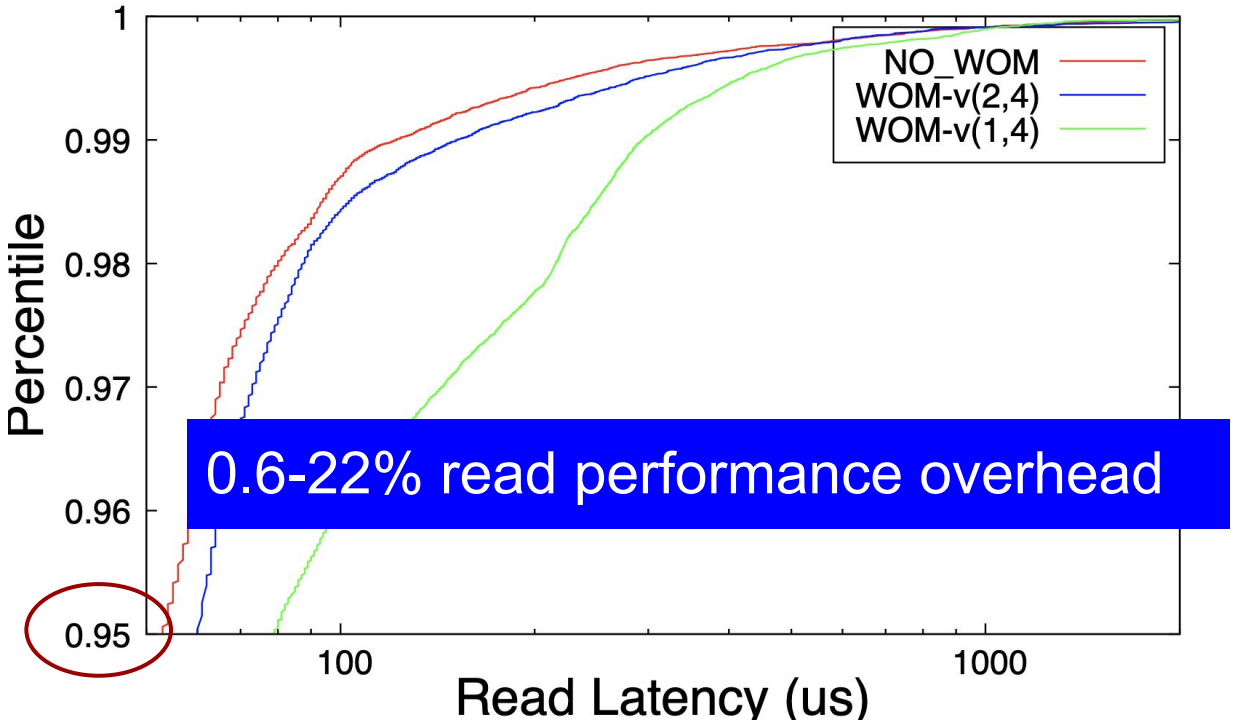

0.6-22% read performance overhead

WOM-v(2,4) QLC v/s MLC Drive

- QLC WOM-v(2,4) can store 2 logical bits per cell.

- MLC drives also store 2 logical bits per cell.

- Without WOM-v coding, MLC drive has better endurance than QLC drive.

WOM-v(2,4) QLC v/s MLC endurance?

WOM-v(2,4) QLC v/s MLC Drive

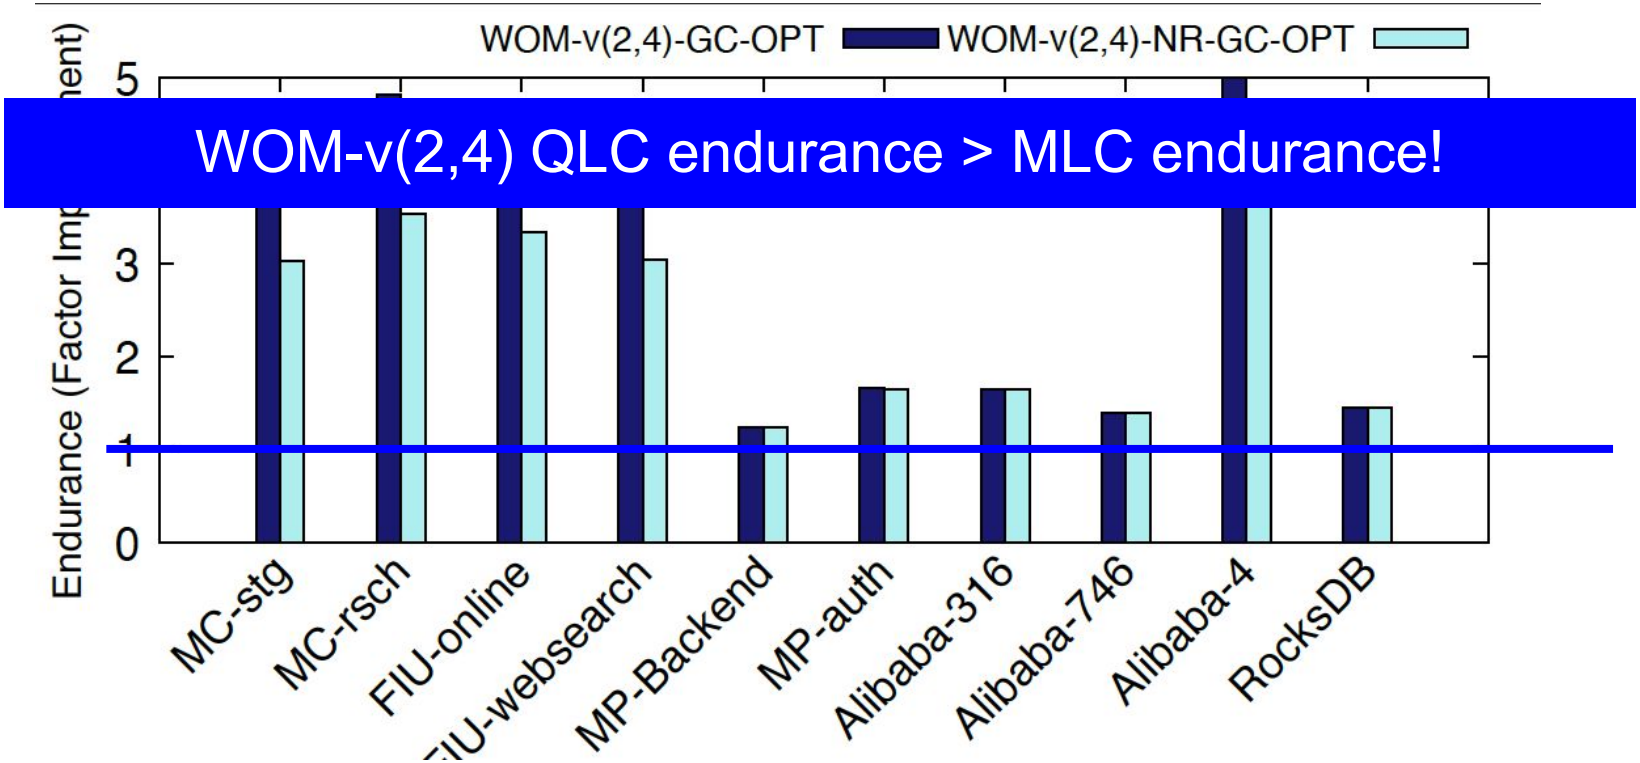

Summary

- First design and implementation of Non-Binary WOM code for QLC Drives.

- 4.4-11.1x reduced erase operations with minimal performance overheads.

- TLC+QLC FEMU Extension merged upstream with master.

- WOM-v simulator code https://github.com/uoftsystems/womv

Q & A