Characterizing Flash Memory:
Anomalies, Observations, and Applications

Laura M. Grupp*, Adrian M. Caulfield*, Joel Coburn*,
Eitan Yaakobi†, Paul H. Siegel†, Jack K. Wolf†, Steven Swanson*

*Non-volatile Systems Laboratory, Department of Computer Science and Engineering
†Center for Magnetic Recording Research
Jacob's School of Engineering
University of California, San Diego

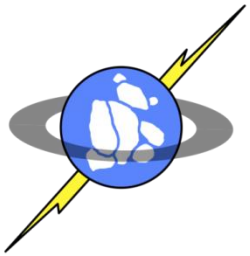

Possible Applications

- Flash Translation Layers (FTLs)
- Operating System control
- Data encodings
- Heterogeneous SSD
- Data retention time

Test Setup

- Custom-Built Daughter Board
- Xilinx XUP Board
- Full-fledge Linux
- Kernel module

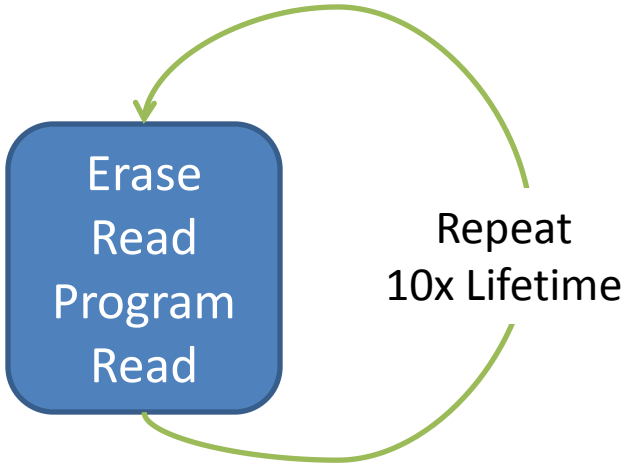

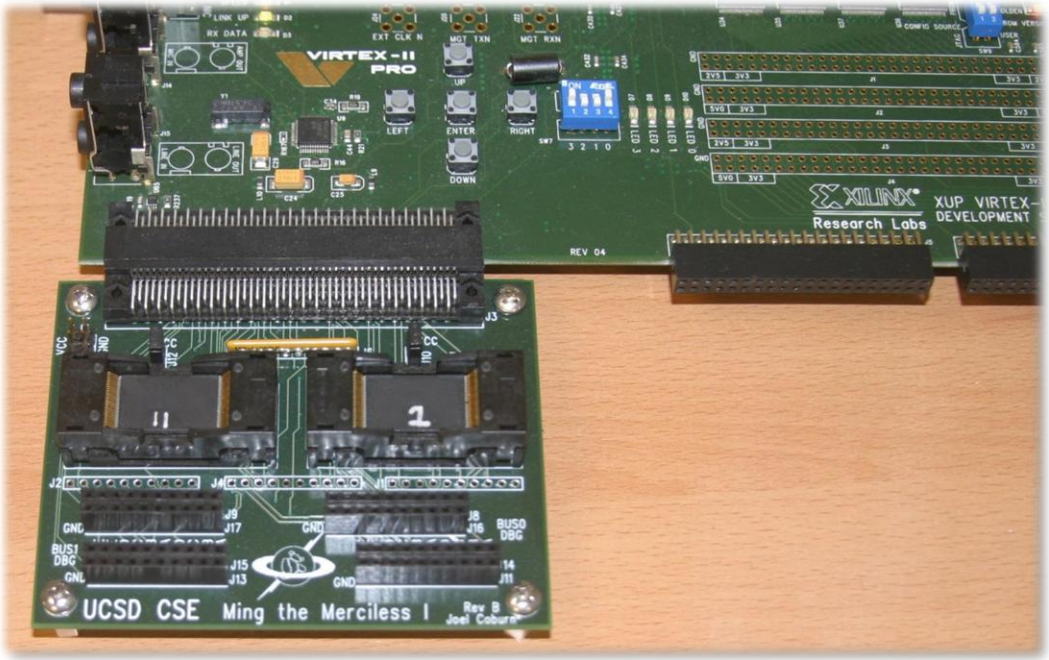

The Test Subjects

Chip Name	Manu-facturer	Tech Node	Capacity (Gb)	Page Size (B)	Pages/ Block	Block/ Plane	Planes/ Die	Dies
A-SLC2	A		2	2048	64	1024	2	1
A-SLC4	A		4	2048	64	4096	1	1
A-SLC8	A		8	2048	64	4096	2	1
B-SLC2	B	50nm	2	2048	64	2048	1	1
B-SLC4	B	72nm	4	2048	64	2048	2	1
E-SLC8	E		8	2048	64	4096	1	2
B-MLC8	B	72nm	8	2048	128	4096	1	1
B-MLC32	B	50nm	32	4096	128	2048	2	2
C-MLC64	C	43nm	64	8192	128	4096	1	2
D-MLC32	D		32	4096	128	4096	1	2
E-MLC8	E		8	4096	128	1024	1	2

The Tests

Quantify known and unknown properties

- Performance

- Energy Efficiency

- Reliability

The Tests

Quantify known complexities, look for new ones

- **Performance**

- Power & Energy

- Reliability

Erase

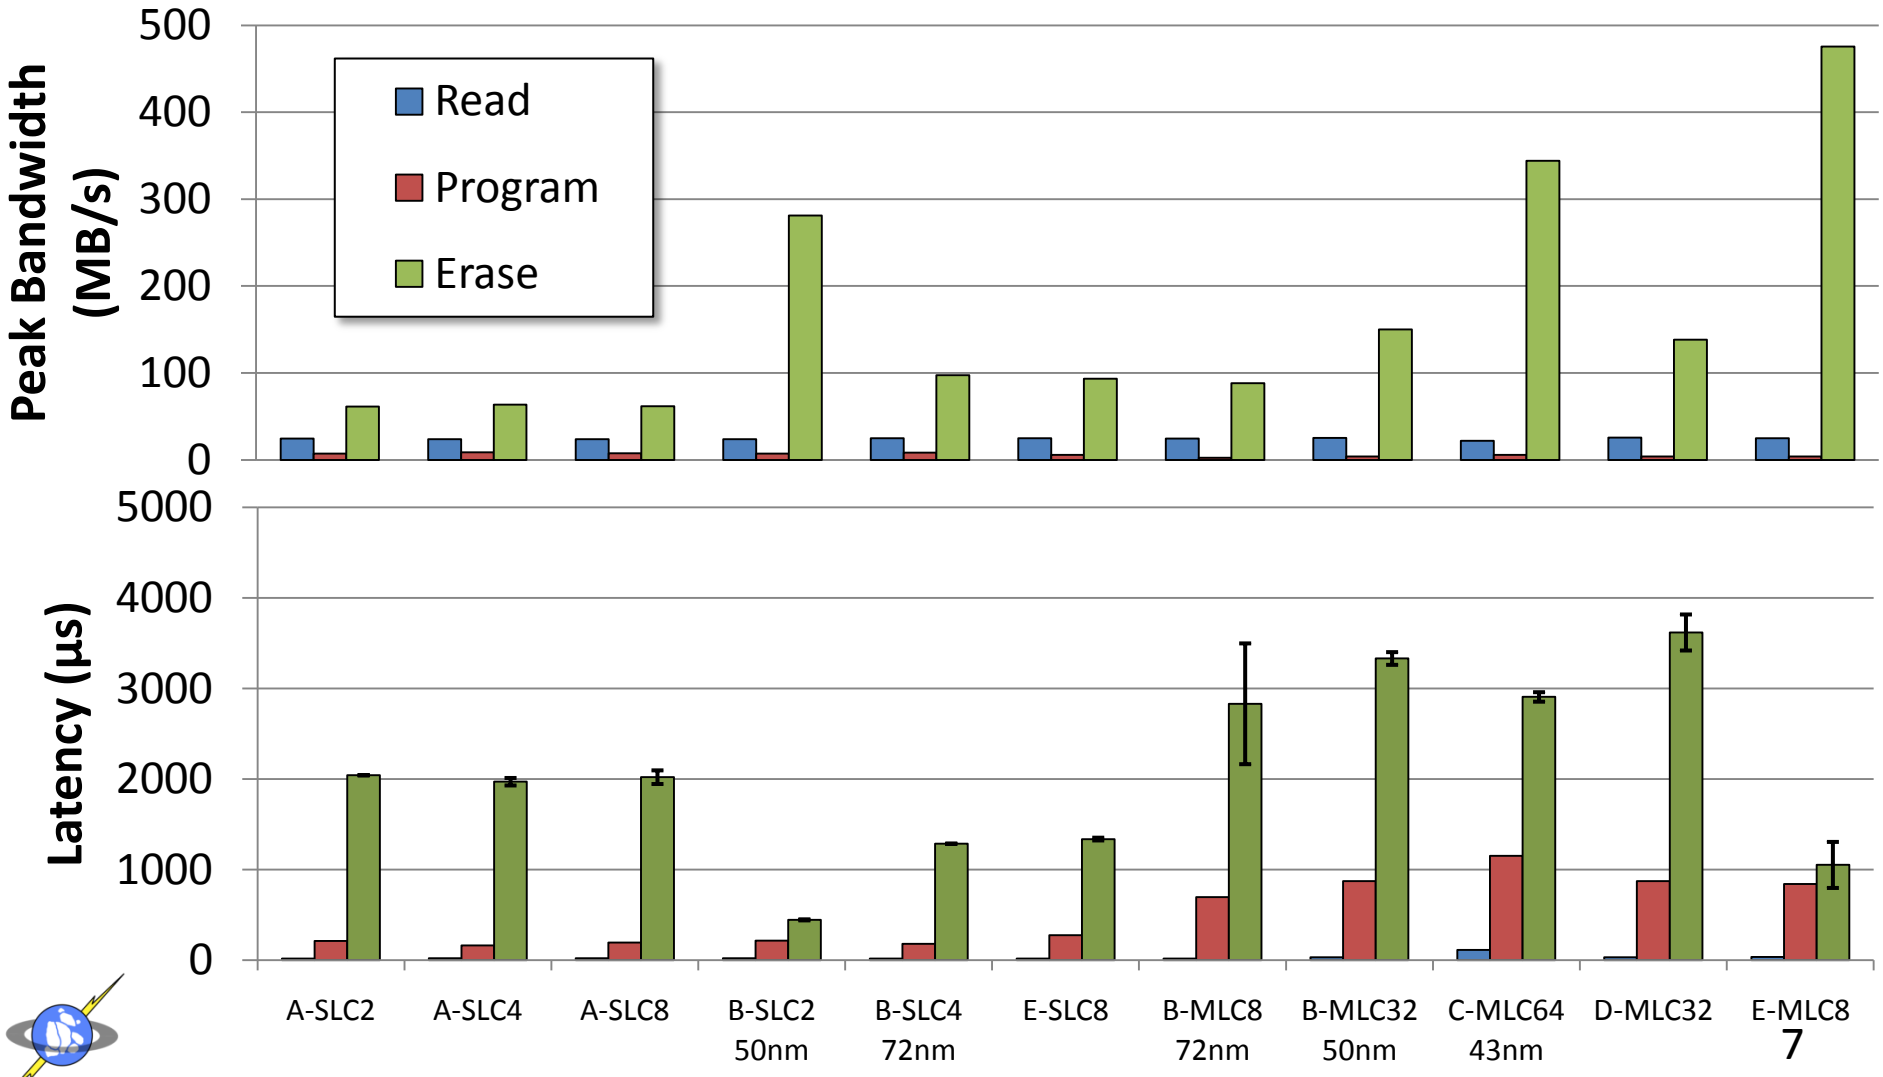
Peak Bandwidth (MB/s)
500
400
300
200
100
0
Read
Program
Erase
Latency (µs)
5000
4000
3000
2000
1000
0
A-SLC2
A-SLC4
A-SLC8
B-SLC2 50nm
B-SLC4 72nm
E-SLC8
B-MLC8 72nm
B-MLC32 50nm
C-MLC64 43nm
D-MLC32
E-MLC8

Read and Program

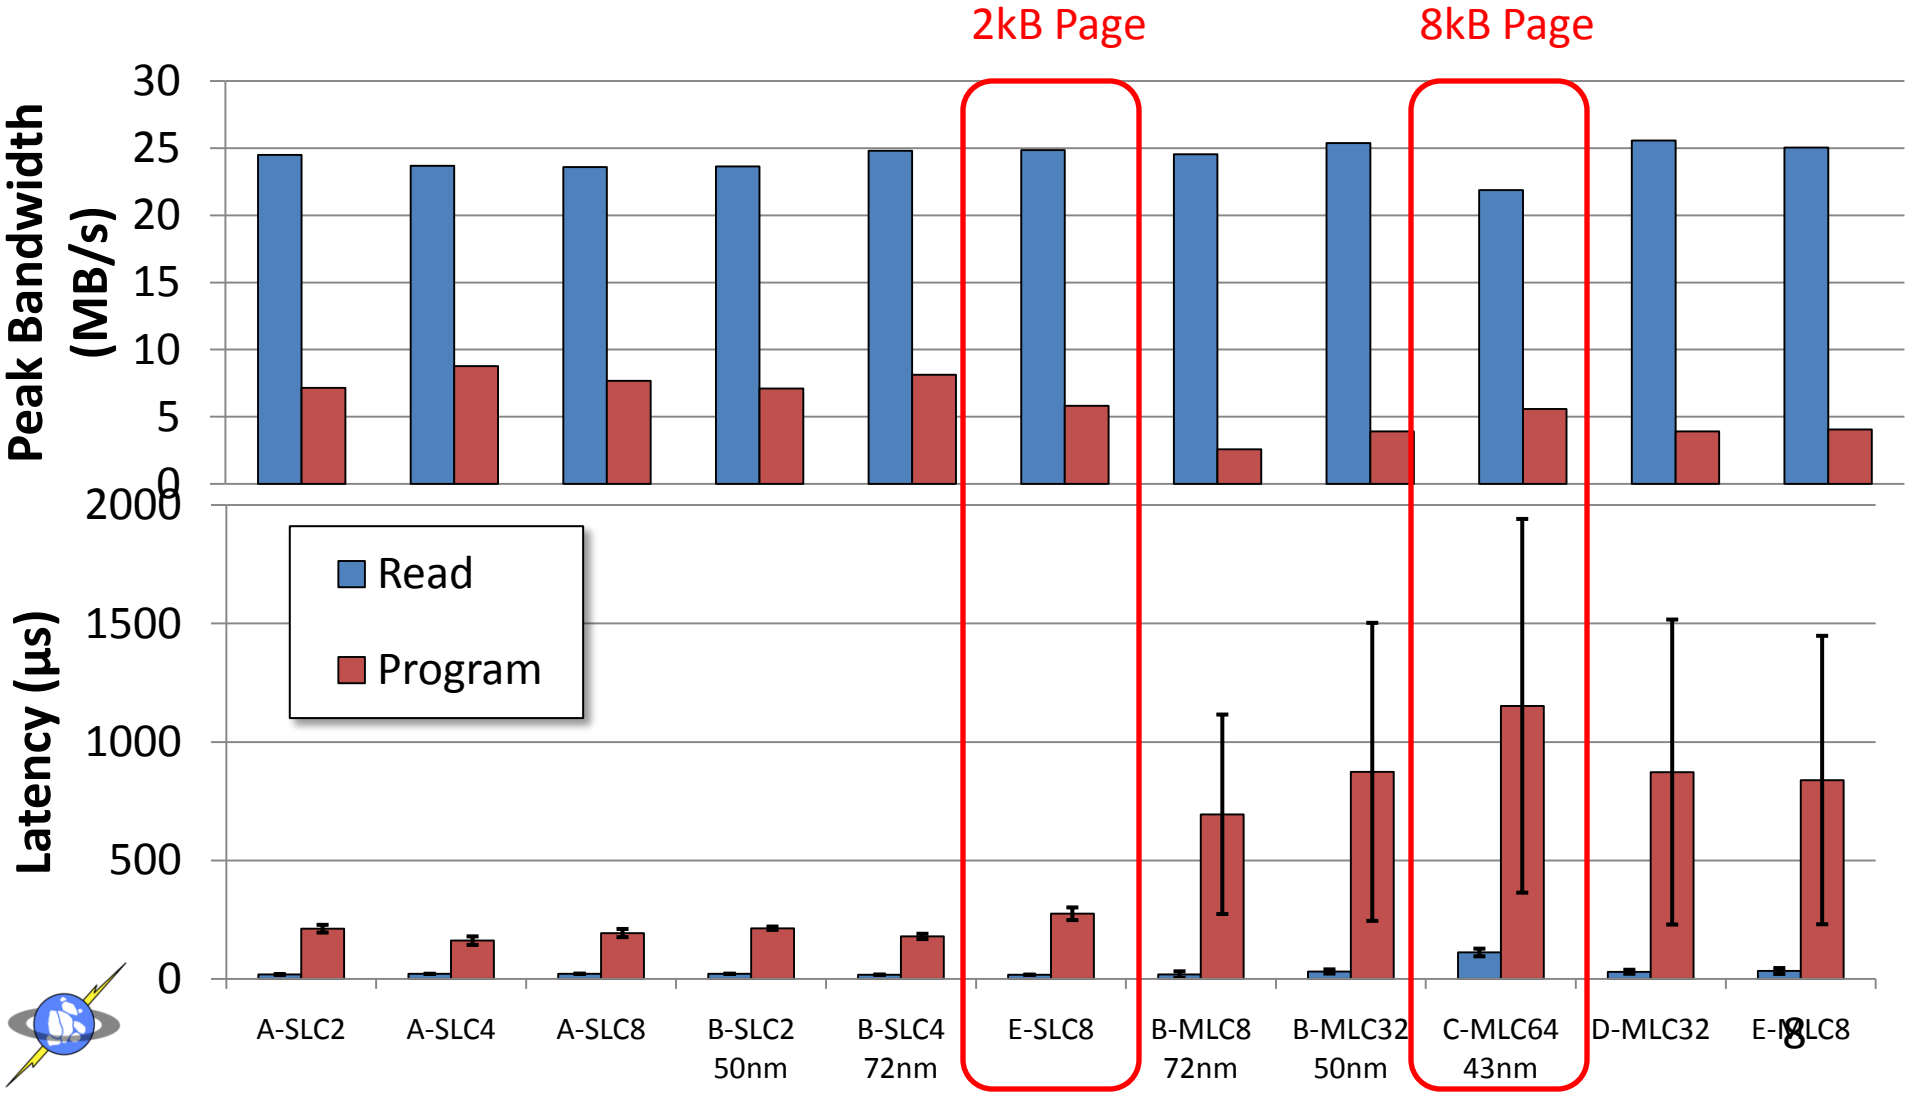

Program Latency Anomaly

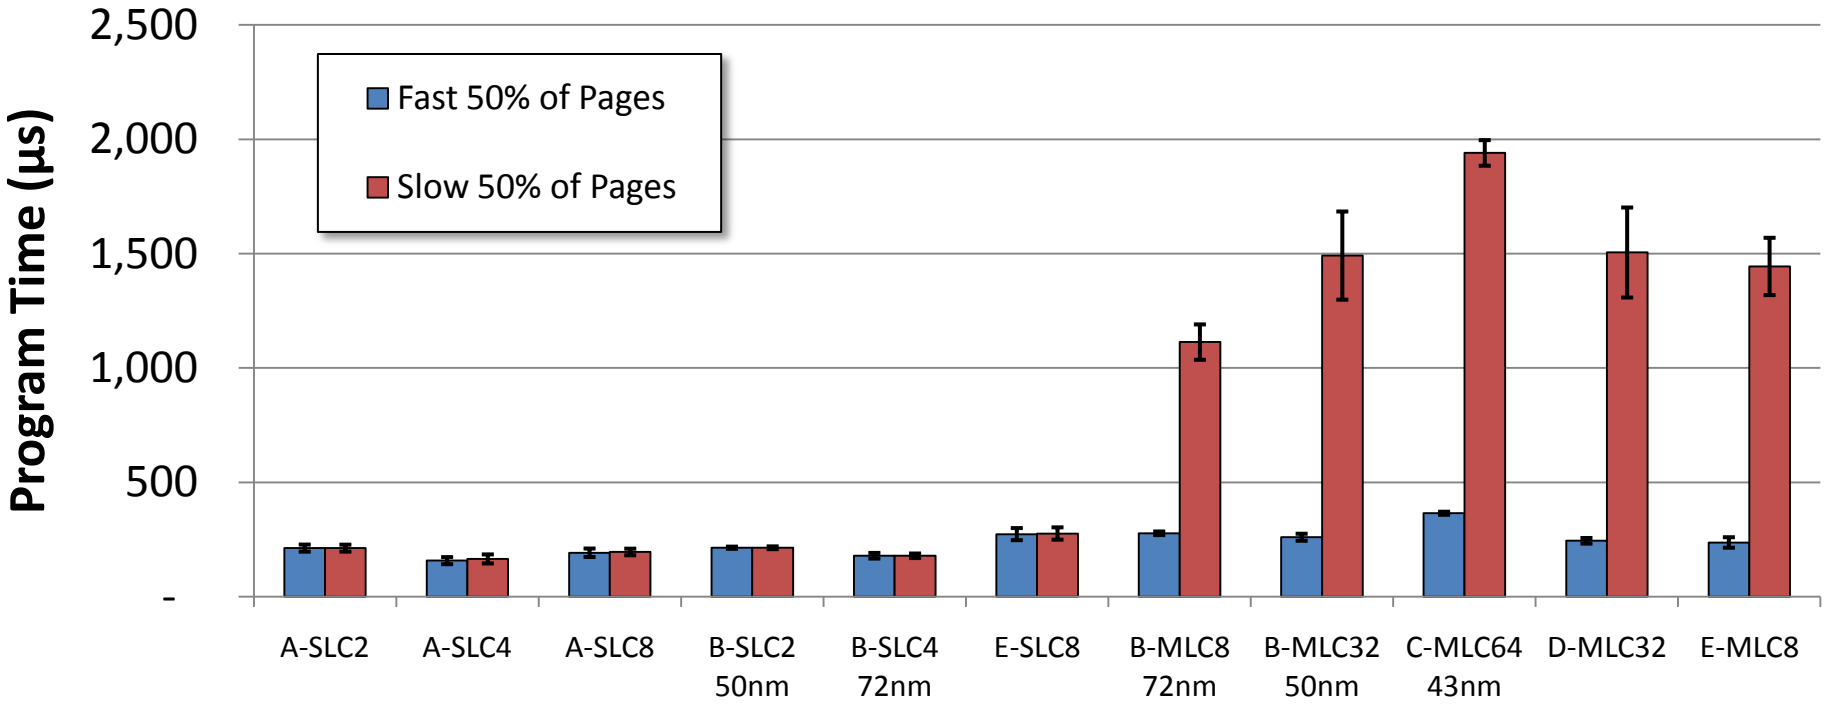

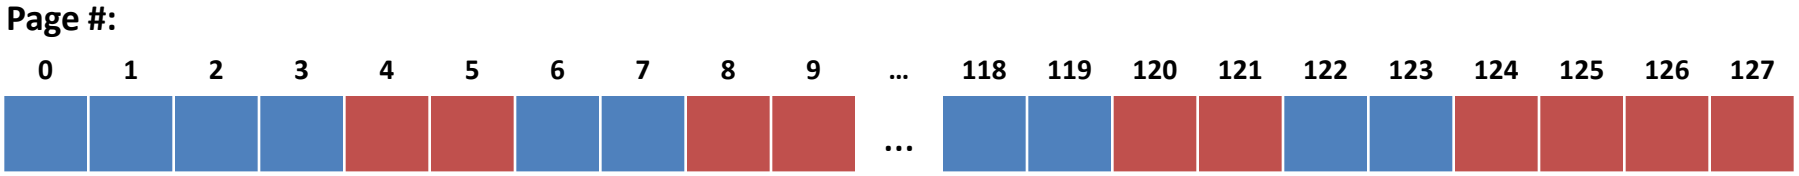

High & Low Order Bits

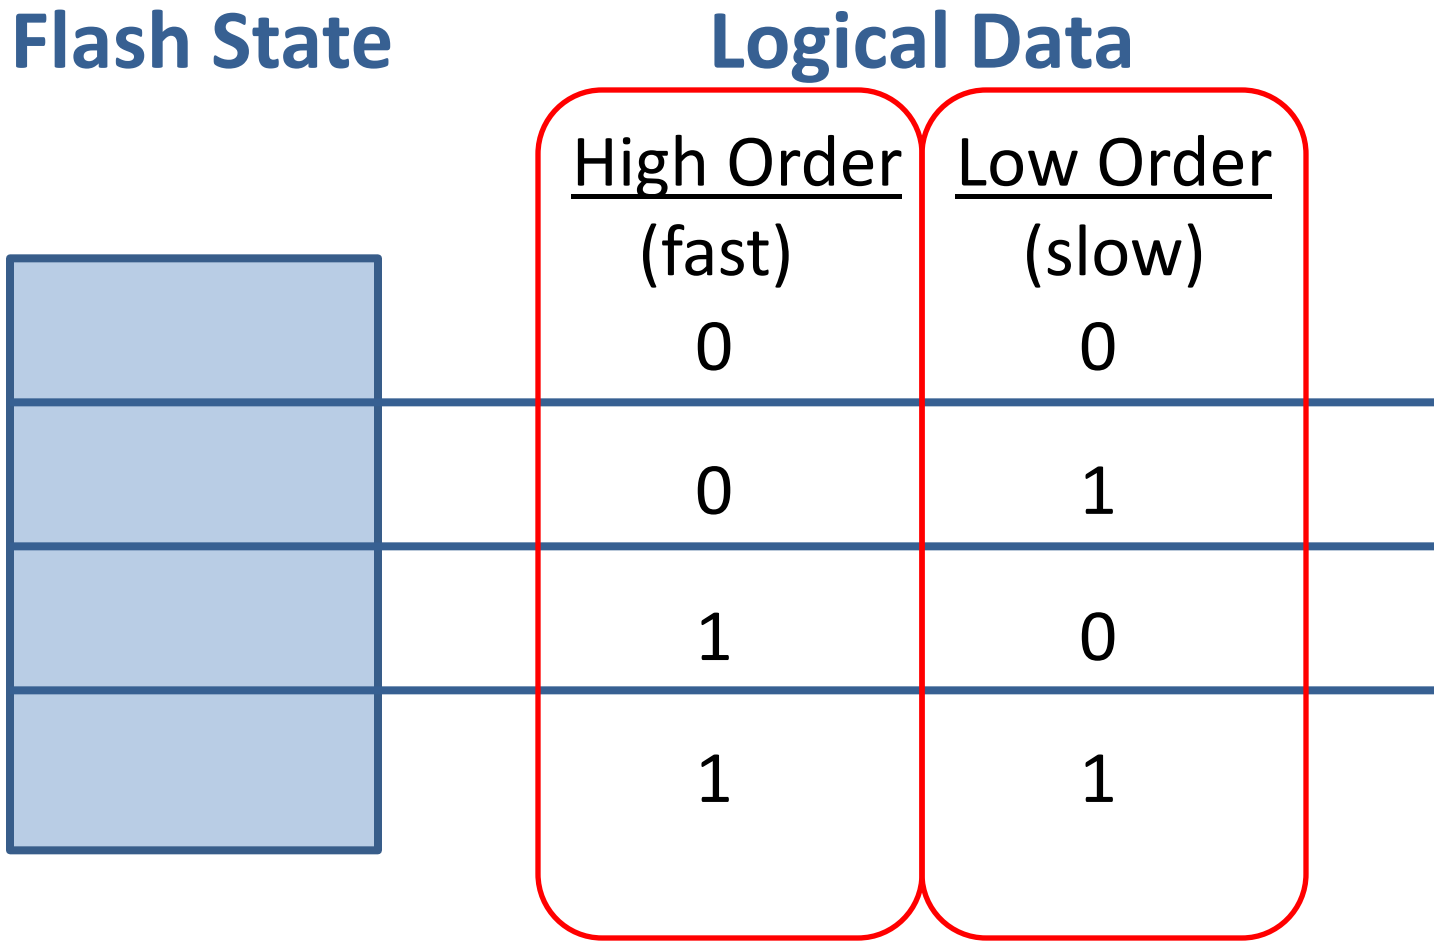

Variation-Aware Interface

Extend Interface from MSR [Birrell et. al., 2005]

"High priority" == write to fast page

"Low priority" == write to any page

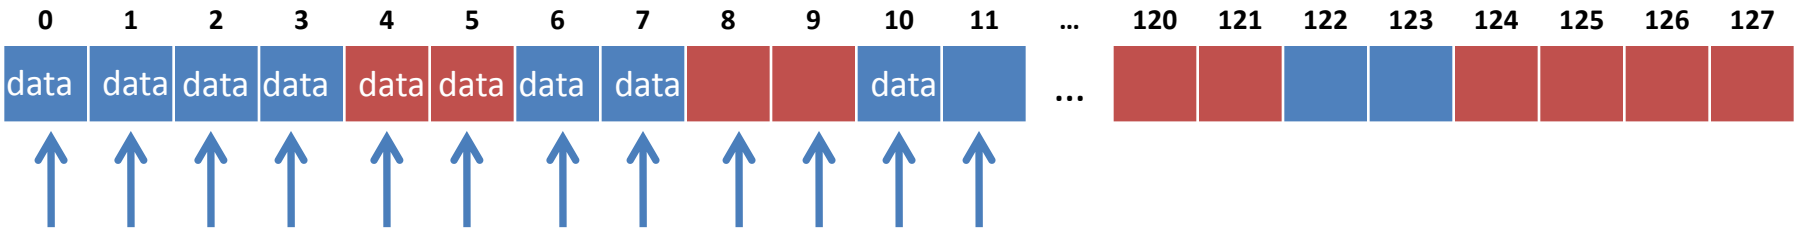

\+ Lower latency when it matters

\- Increase wear

Improving Paging Latency

- Paging in/out virtual memory == high priority
- Goal: reduce swap latency **40% faster**
- Side effect: Increased Wear **3% increase**

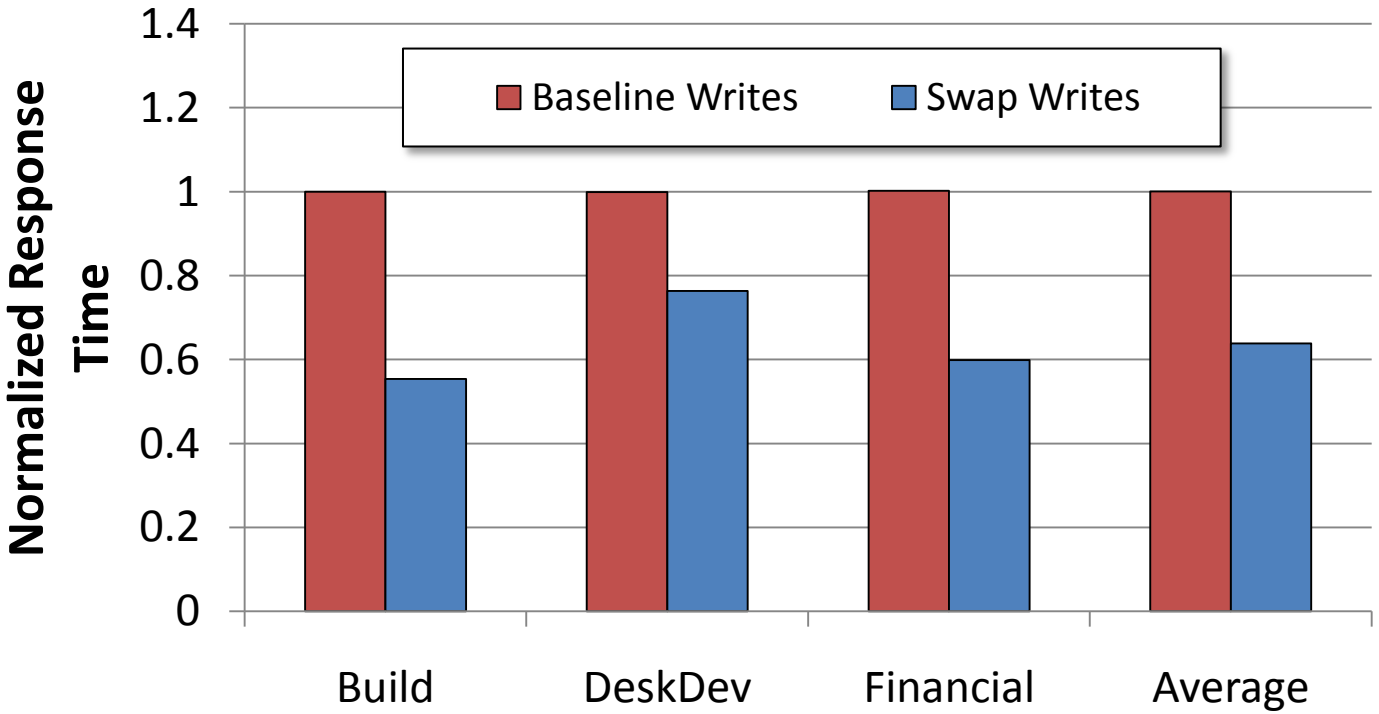

MLC as SLC: Single Mode Level Cell?

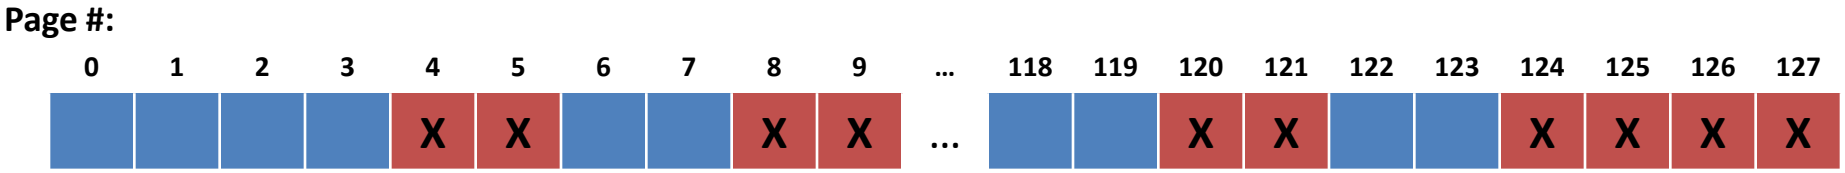

Average cost per bit	
MLC:	31 ¢/Gbit
SLC:	94 ¢/Gbit
MLC is **33%** of SLC	

DRAMeXchange
4/10/10

The Tests

Quantify known complexities, look for new ones

- Performance
- **Power & Energy**
- Reliability

From a representative datasheet →

	Typ	Max	Unit
Sequential read	45	90	mW
Program	45	90	mW
Erase	45	90	mW

Peak Power

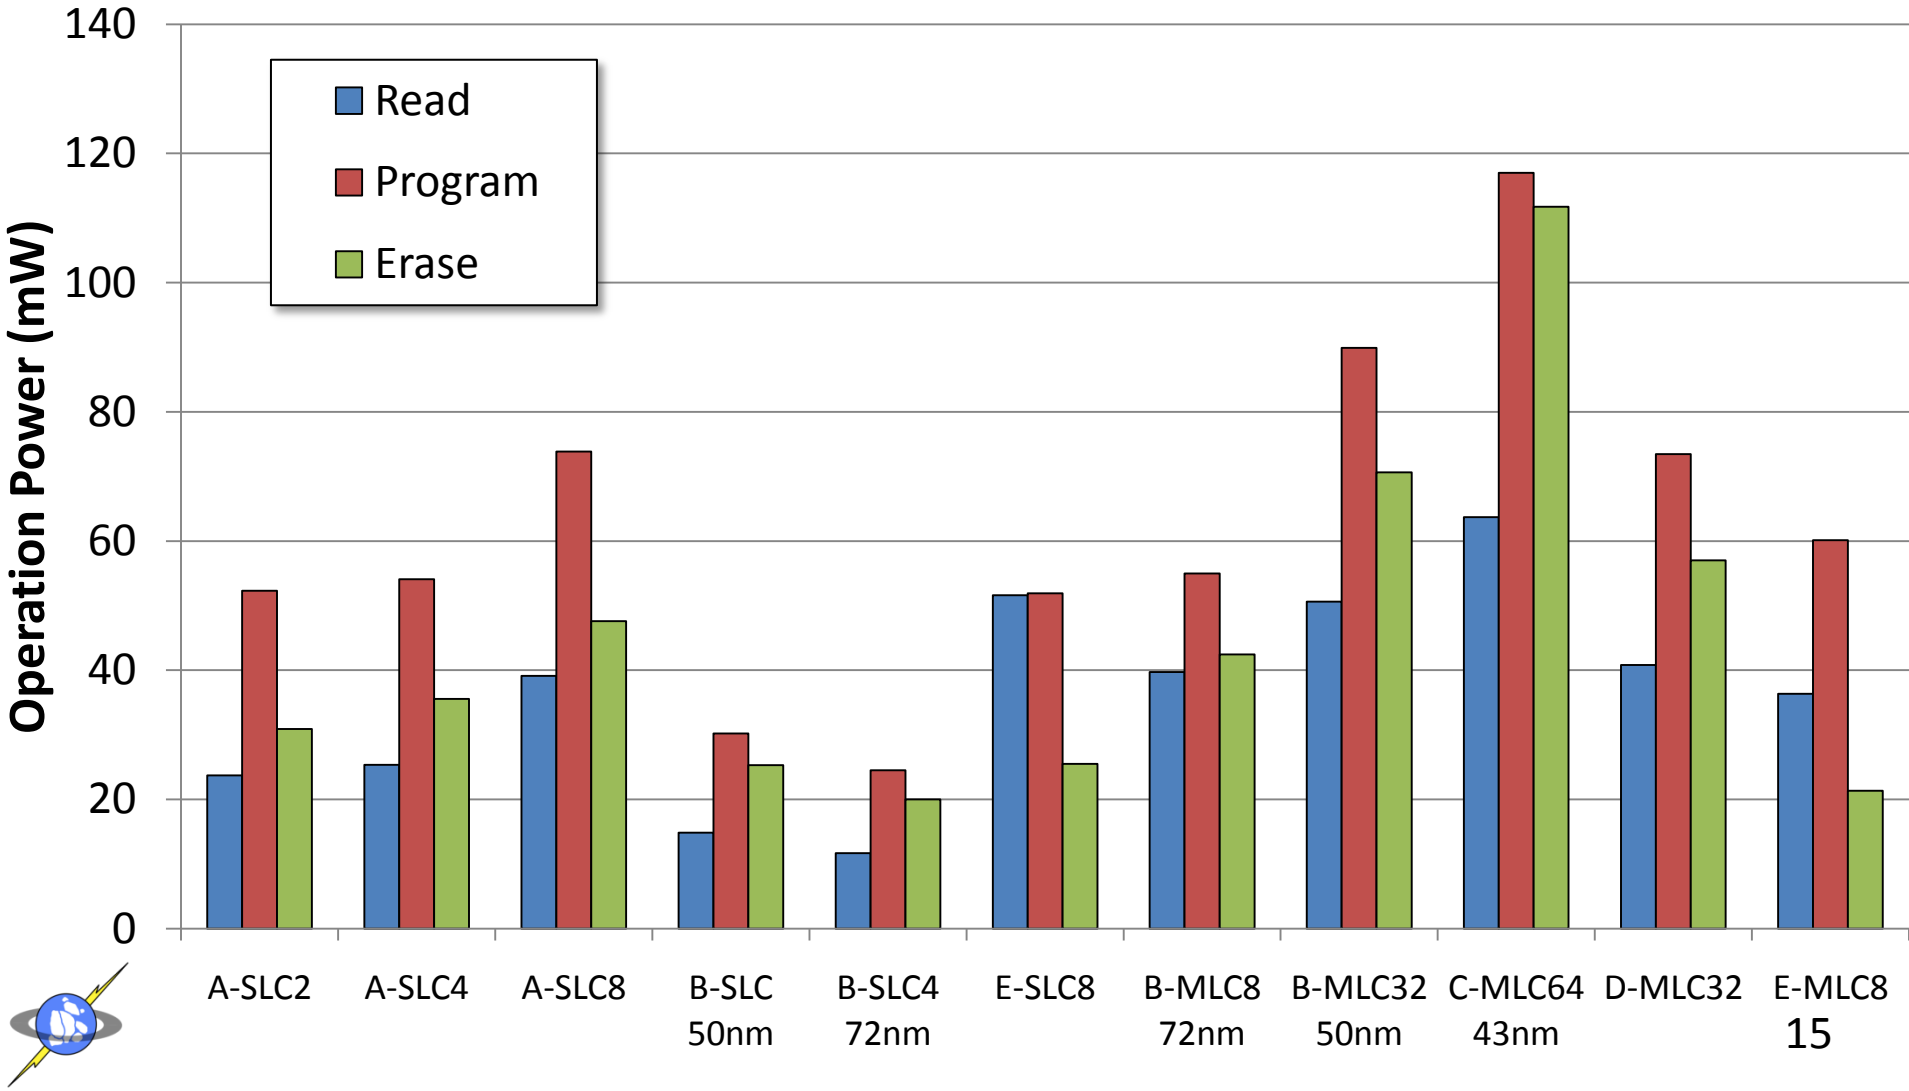

Energy: MLC

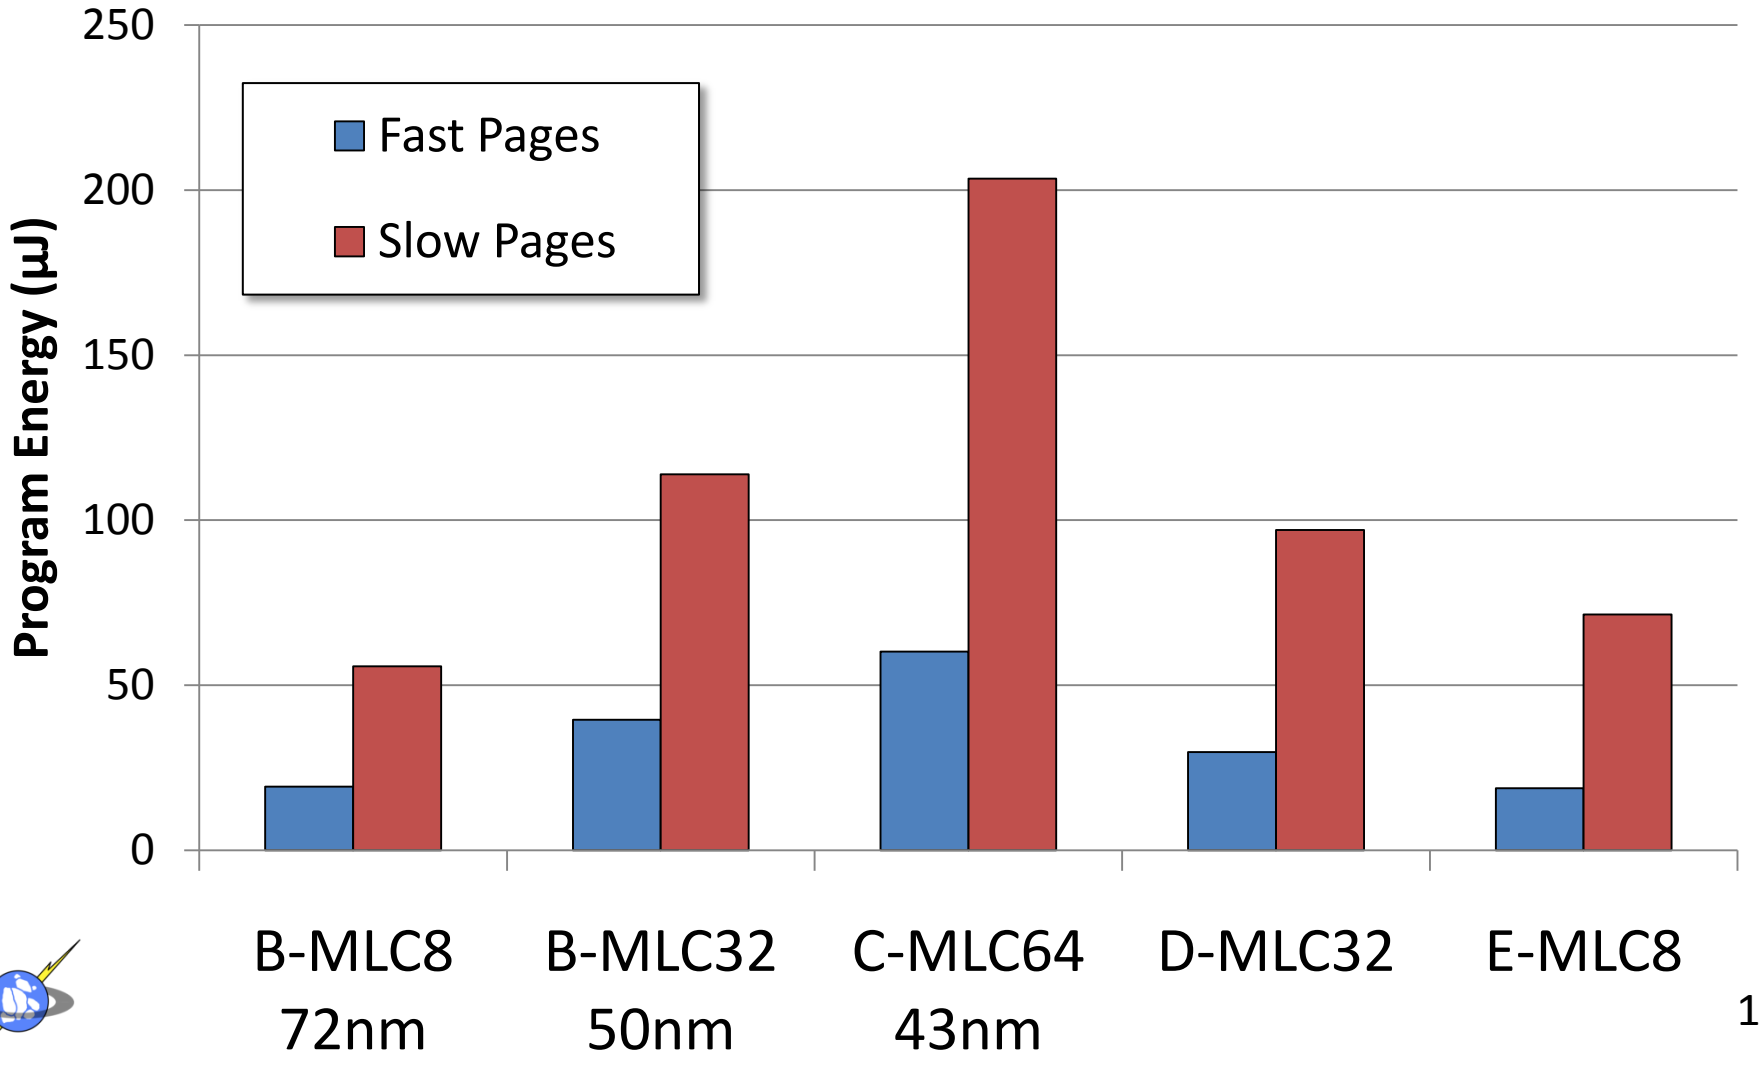
250
200
150
100
50
0
Program Energy (µJ)
Fast Pages
Slow Pages
B-MLC8
72nm
B-MLC32
50nm
C-MLC64
43nm
D-MLC32
E-MLC8

The Tests

Quantify known complexities, look for new ones

- Performance
- Energy Efficiency
- **Reliability**

SLC Wear Out

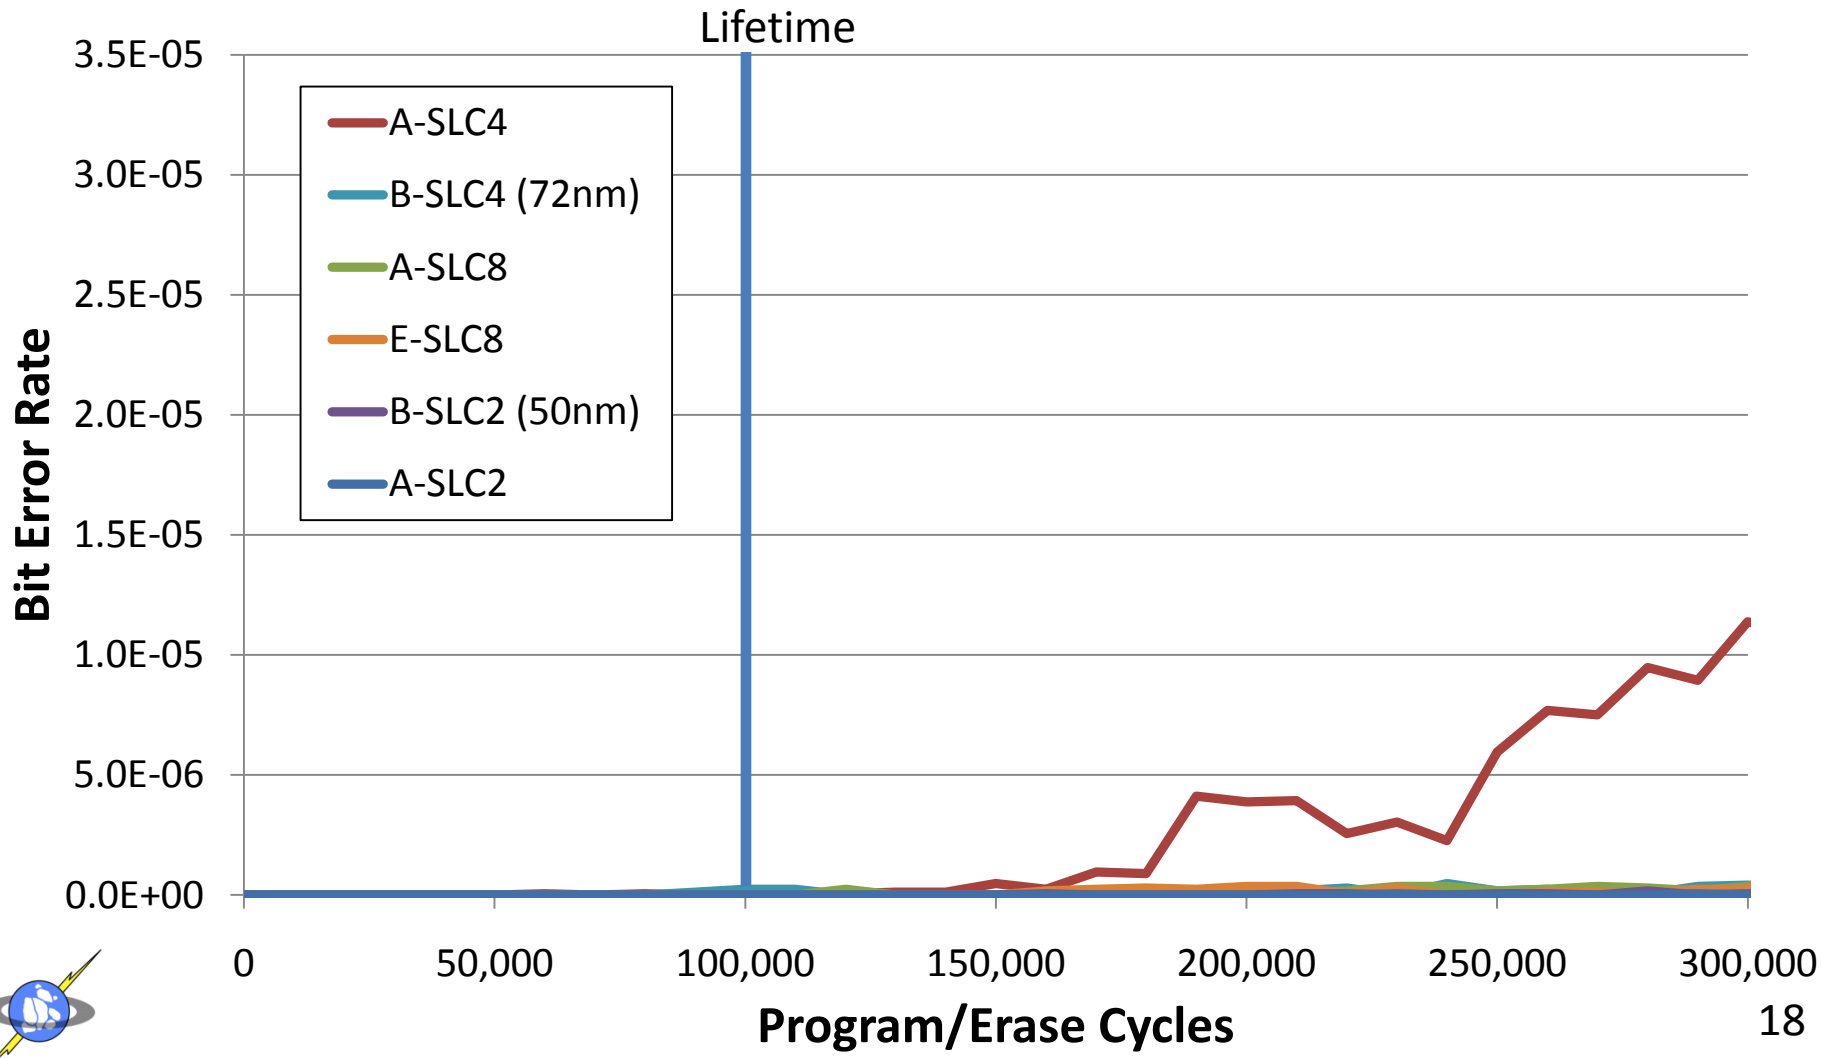
Lifetime
3.5E-05
3.0E-05
2.5E-05
2.0E-05
1.5E-05
1.0E-05
5.0E-06
0.0E+00
Bit Error Rate
A-SLC4
B-SLC4 (72nm)
A-SLC8
E-SLC8
B-SLC2 (50nm)
A-SLC2
0
50,000
100,000
150,000
200,000
250,000
300,000
Program/Erase Cycles

MLC Wear Out

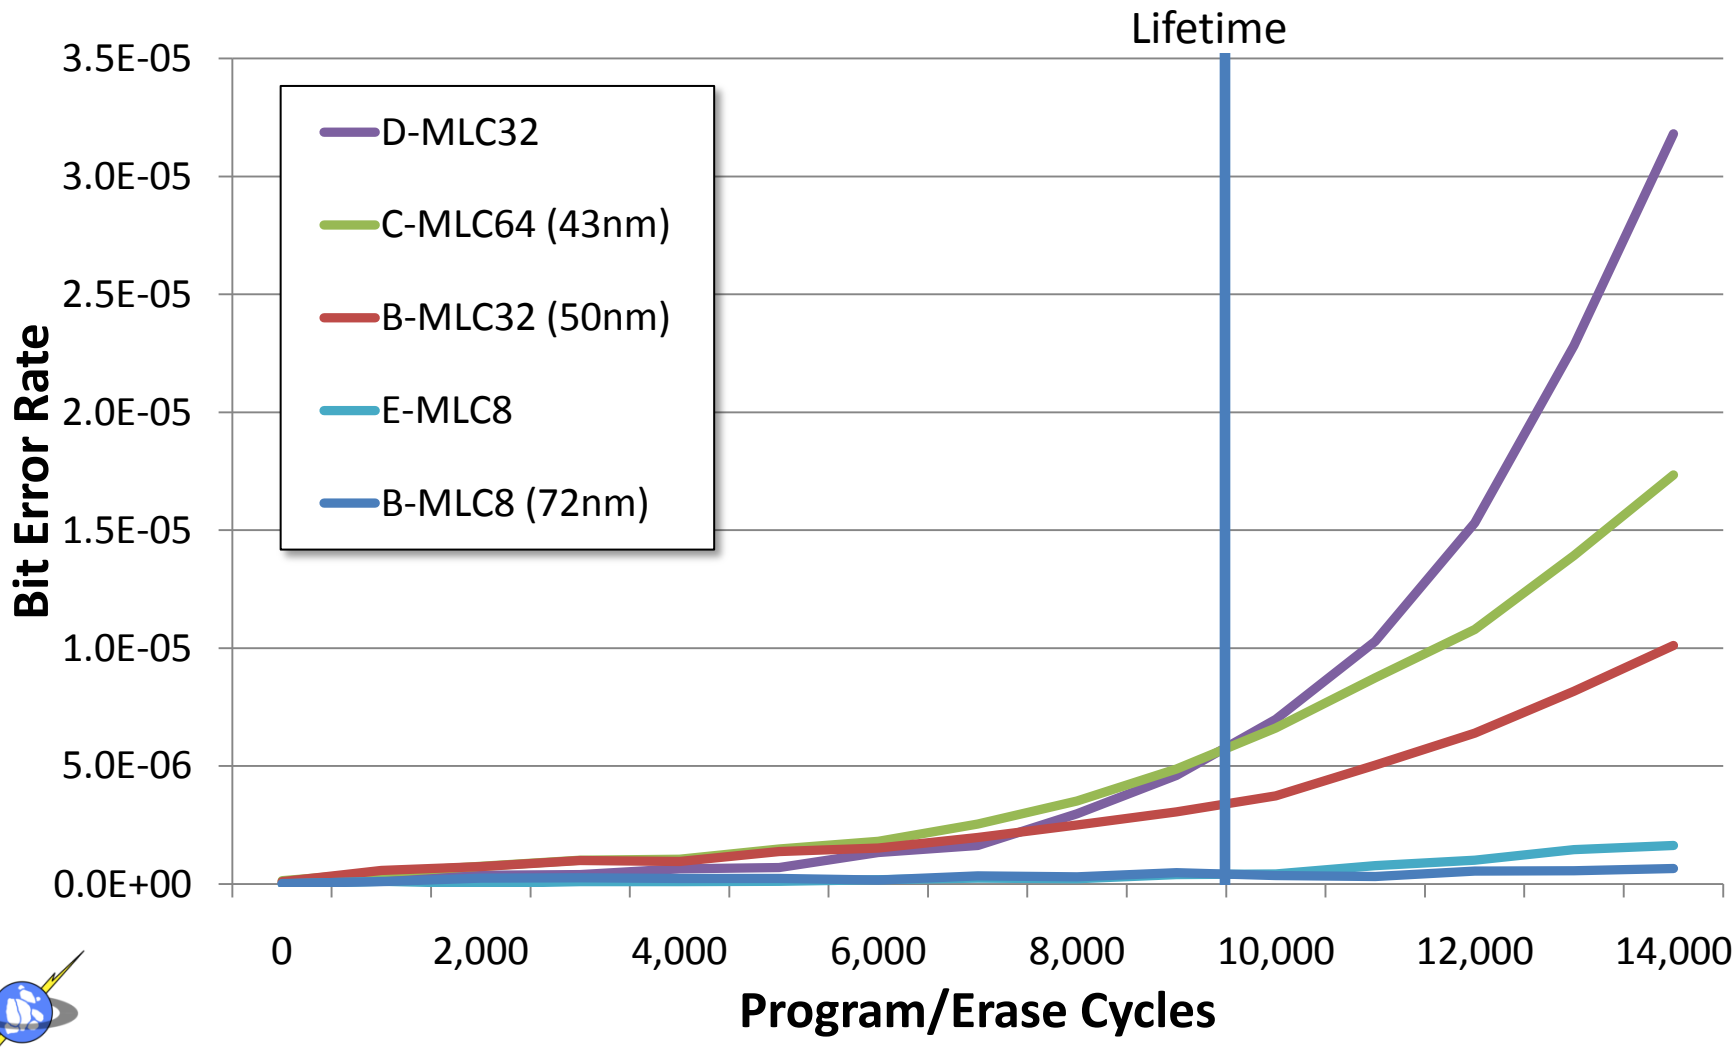
Lifetime
3.5E-05
3.0E-05
2.5E-05
2.0E-05
1.5E-05
1.0E-05
5.0E-06
0.0E+00
Bit Error Rate
D-MLC32
C-MLC64 (43nm)
B-MLC32 (50nm)
E-MLC8
B-MLC8 (72nm)
0
2,000
4,000
6,000
8,000
10,000
12,000
14,000
Program/Erase Cycles

Read Disturb

The degradation of a bit
by **reading** that cell or another.

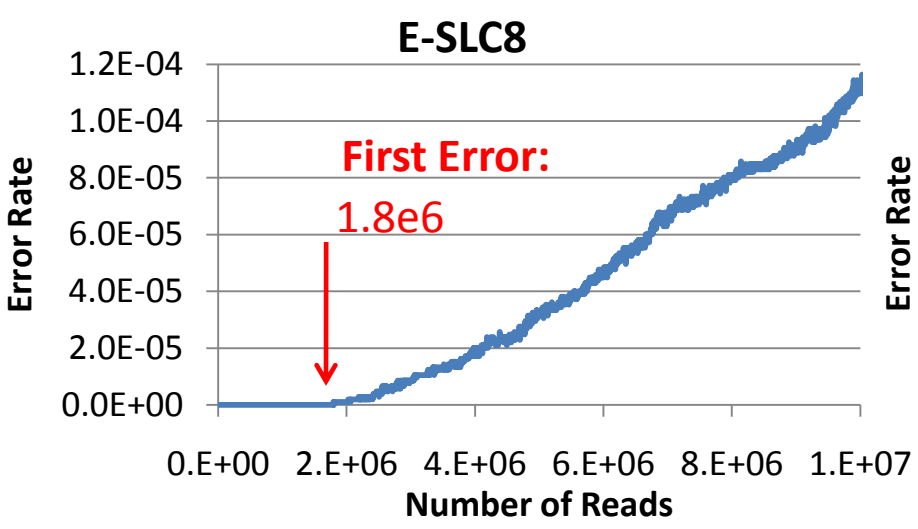

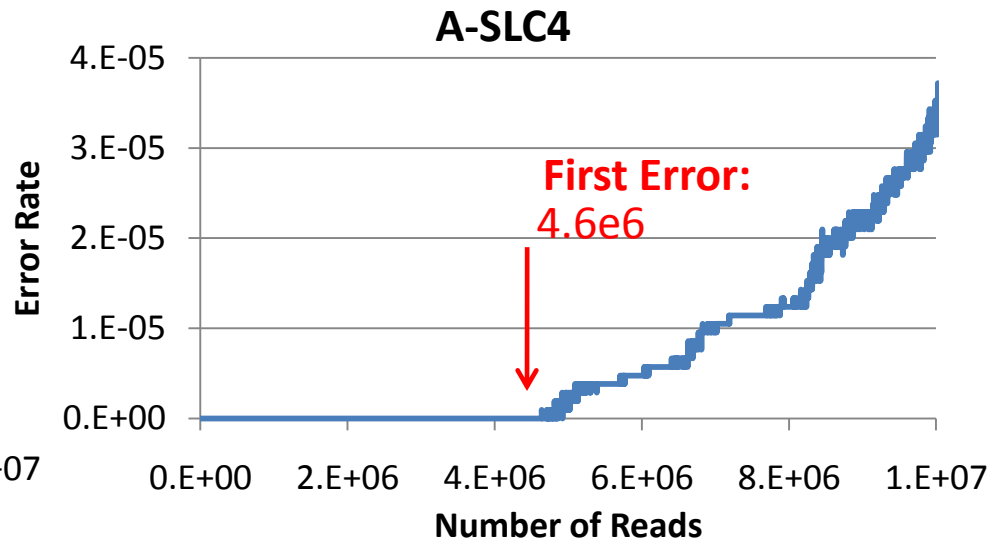

Program Disturb

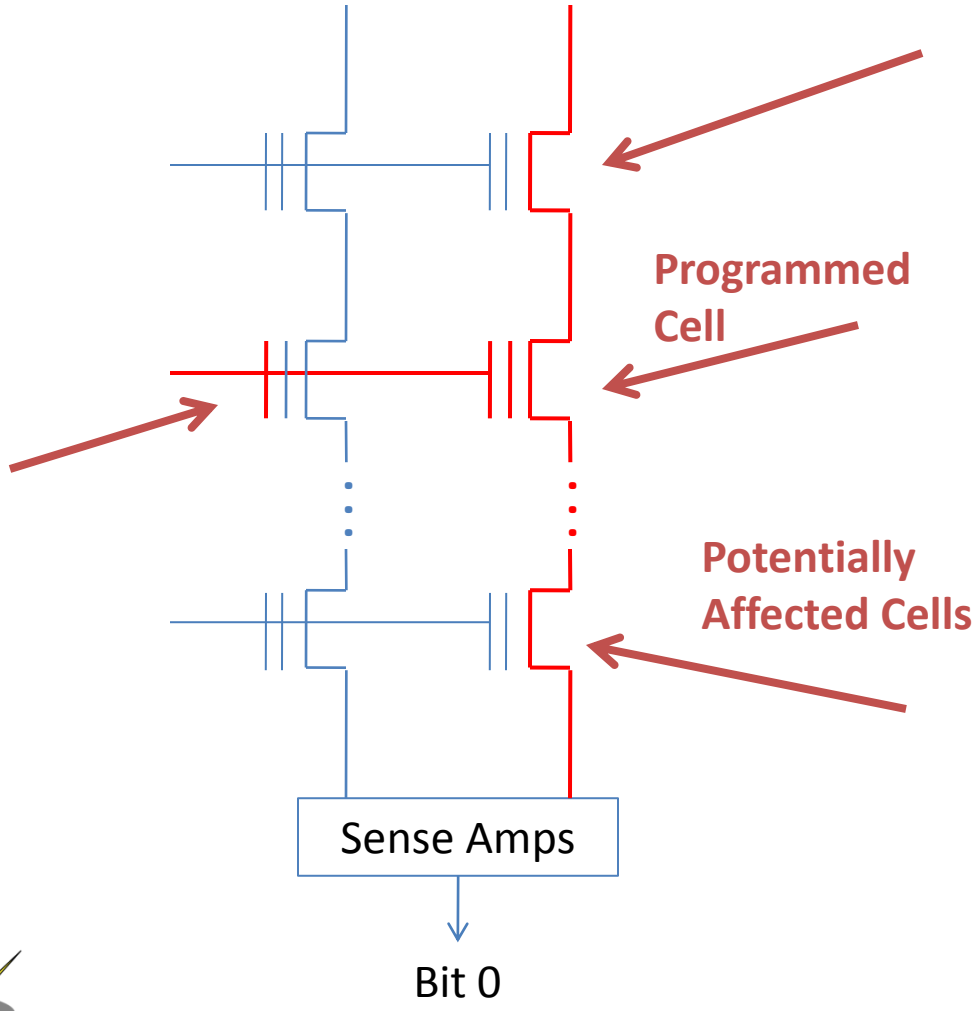

The degradation of a bit by **programming** that cell or another

Procedure:

Erase Block

Repeat:

Program new data to one page

Program Disturb

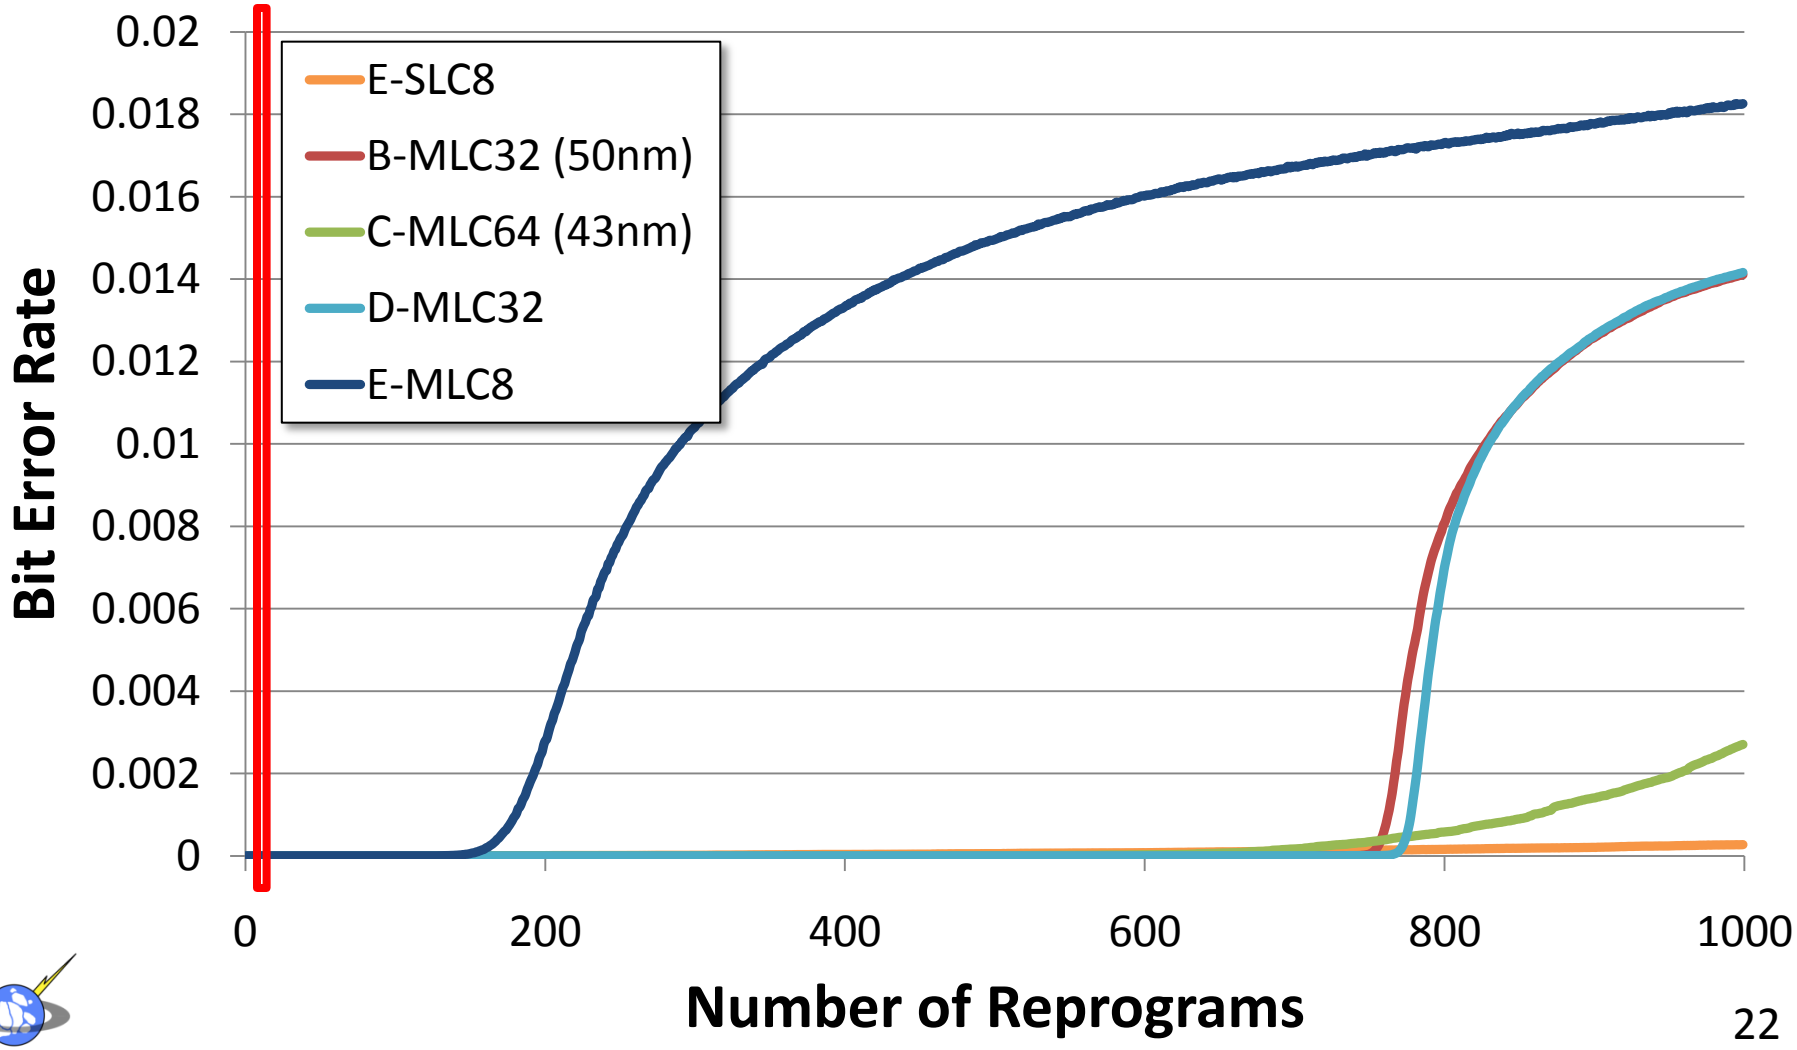

Program Disturb: 1 reprogram

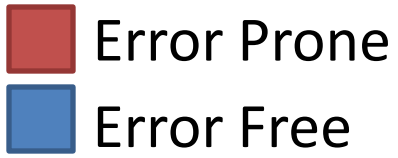

- SLC: no errors for at least one reprogram

- MLC: errors for reprograms of certain pages

Write-Once Memory (WOM) Codes

Old Byte: 01110110

New Byte: 10000111

Logical Bits	First Generation	Second Generation
00	111	000
01	110	001
10	101	010
11	011	100

Program:	01	11	01	10
Reprogram:	10	00	01	11

Physical	**110**	**011**	**110**	**101**	**1st Gen.**
Physical	**010**	**000**	**110**	**100**	**2nd Gen.**

WOM Codes and Flash

Procedure: Erase, Program, Reprogram, Repeat

WOM-safe: can be reprogrammed – use WOM encoding

WOM-unsafe: can't be reprogrammed – don't encode

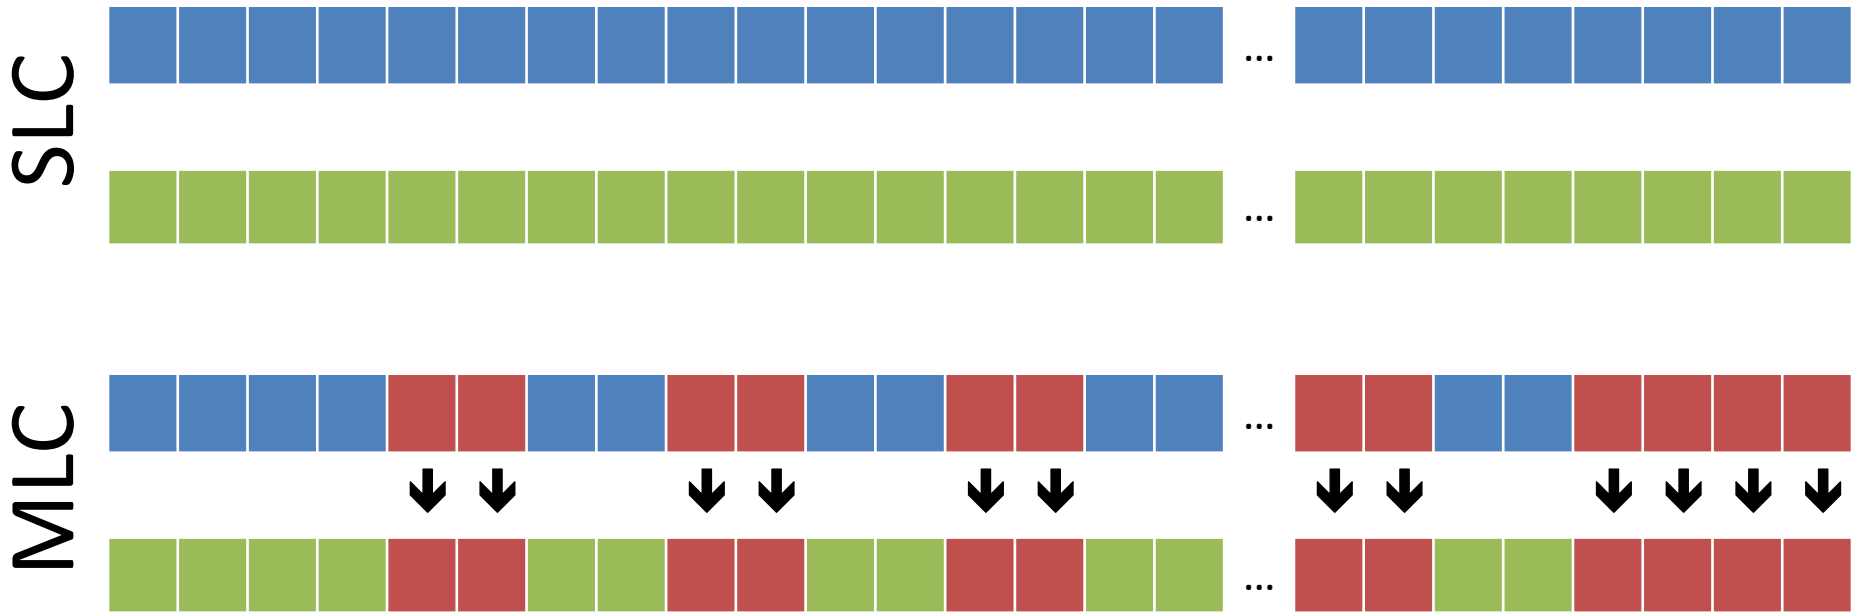

Effective Lifetime

How many logical bits can we program per erase?

SLC

2 * ⅔

= 33% increase

MLC

½(2 * ⅔) + ½(1)

= 16.5% increase

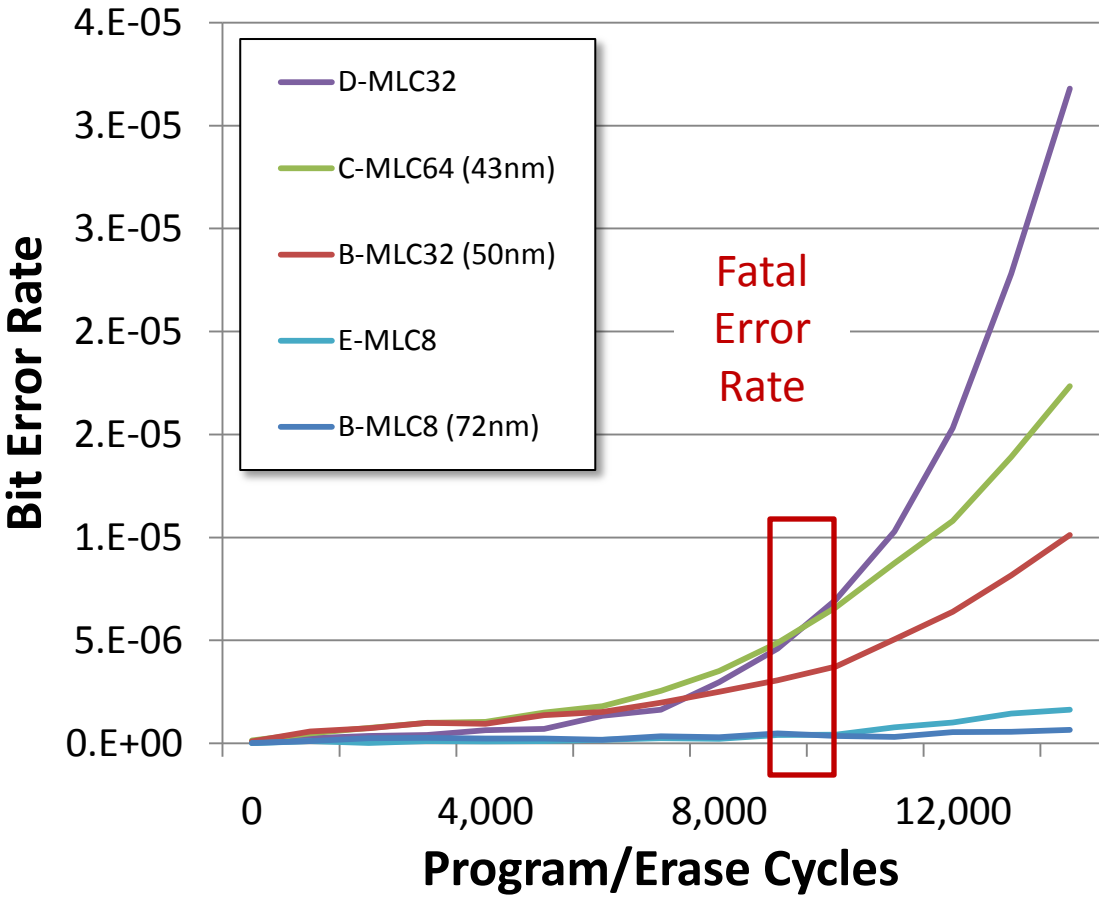

WOM Codes – Lifetime Extension

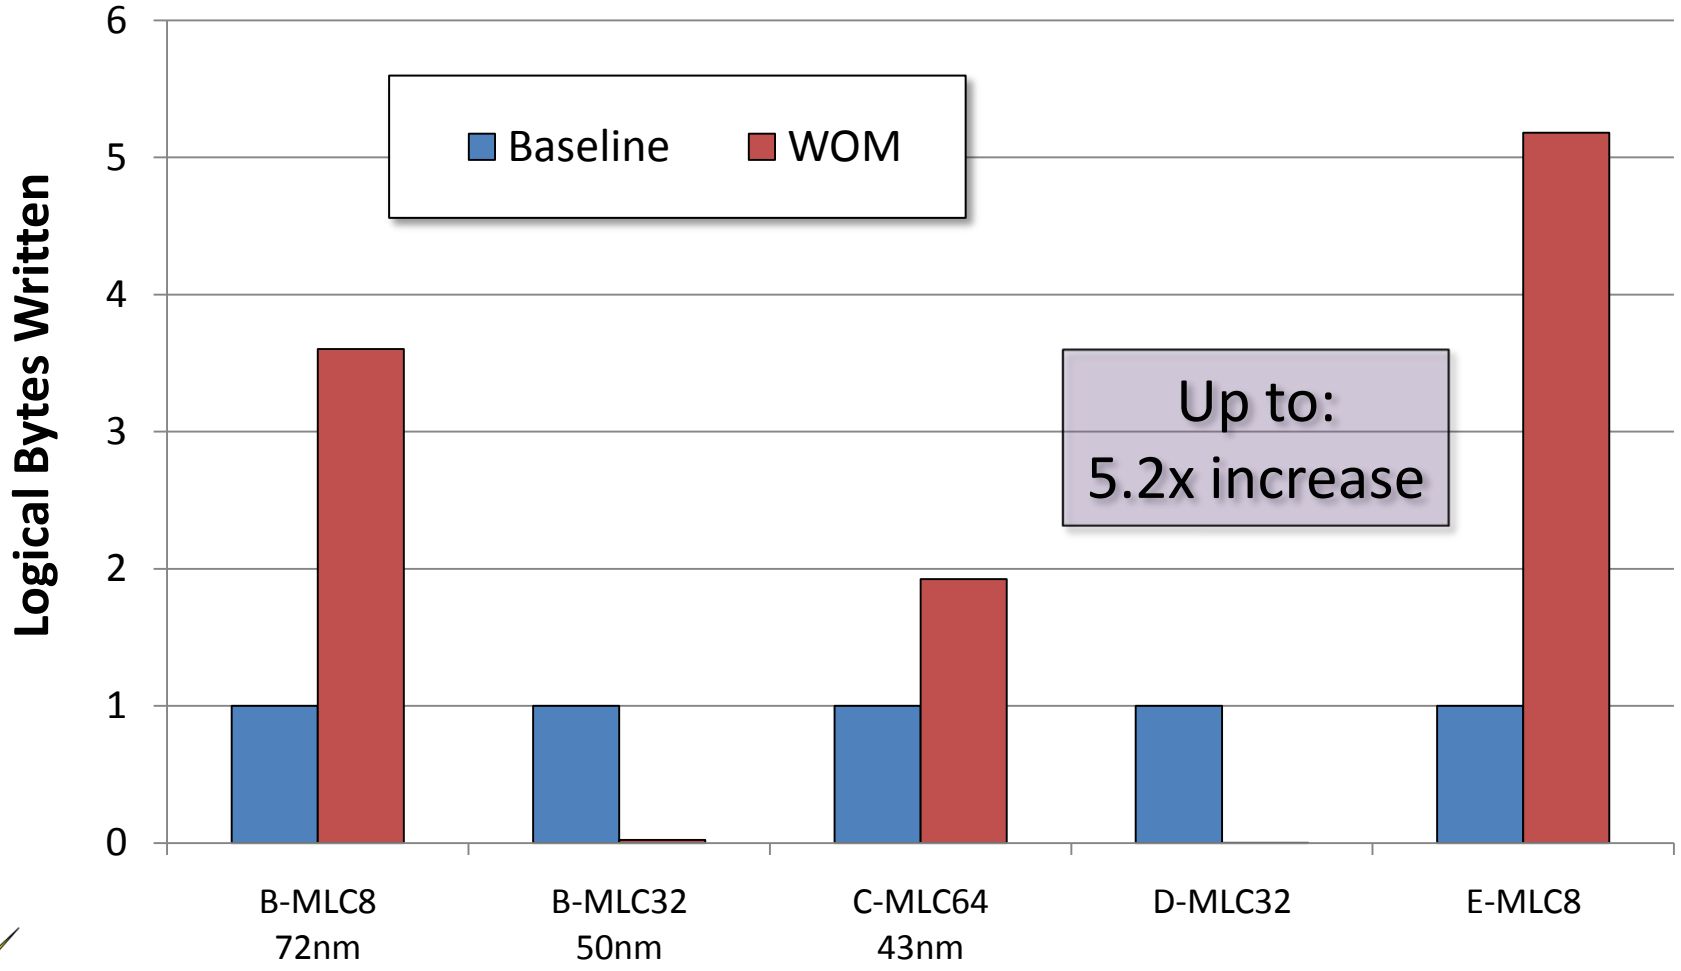

WOM Codes – Energy Reduction

Fewer erases per written bit

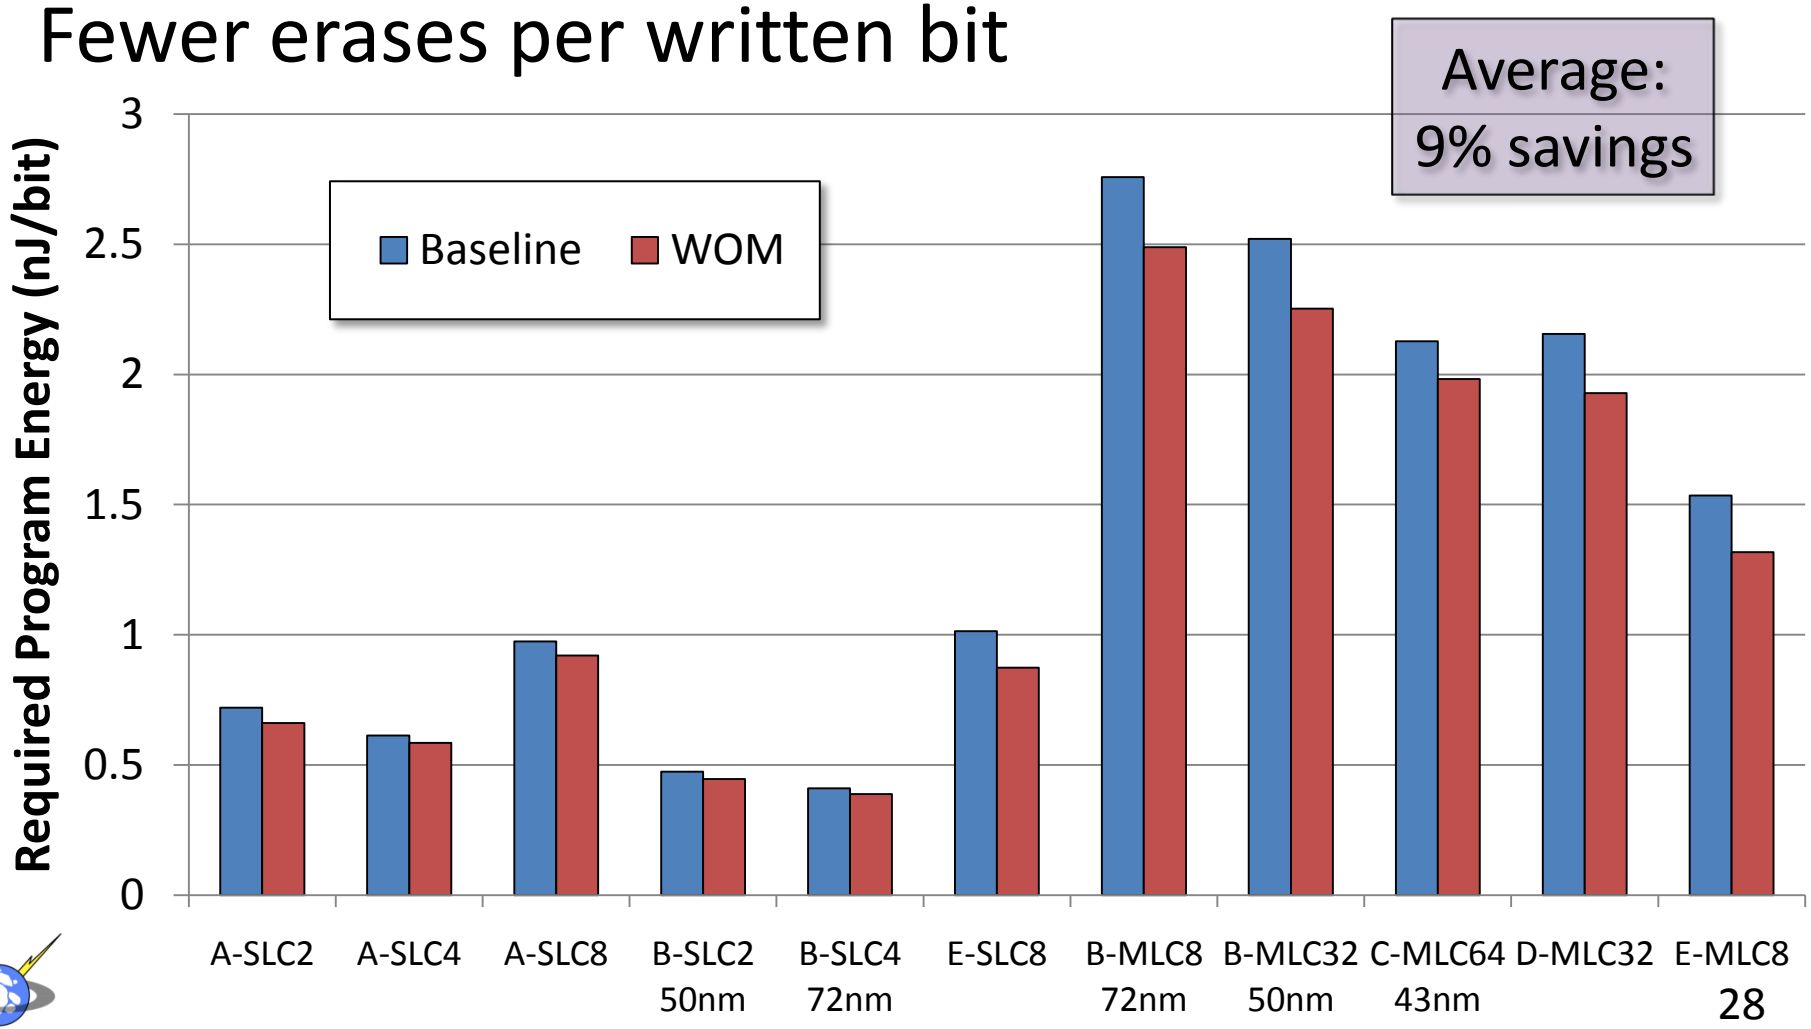

Future Work

- Chips & Commands
  - more chips, higher densities
  - Other operations: cached, multi-plane
- Effects
  - Data retention
  - Data Dependence
  - Radiation Tolerance
  - Delay between Operations
- Applications
  - Error correction codes
  - Data encodings
  - New Applications
  - Interface Modifications

Conclusion

- Flash Memory beyond the datasheet
- Variation-Aware Interface:
  - **Latency** Reduction: average of **40%**
  - **Energy** Savings: average of **13%**
- Write Once Memory (WOM) Encoding
  - **Lifetime** Extension: up to **5.2x**
  - **Energy** Savings: average of **9%**

Thank You