Characterizing Flash Memory for Power Failure

Hung-Wei Tseng*, 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

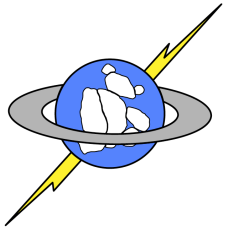

Outline

- The danger of power failure
- Experimental setup
- Results
- Conclusion and future work

Power Failure and Flash

- Solid state drives (SSDs) do not tolerate power failure well
  - FTL (Flash Translation Layer) rebuilds file system abstraction using metadata
  - Power failure may corrupt metadata

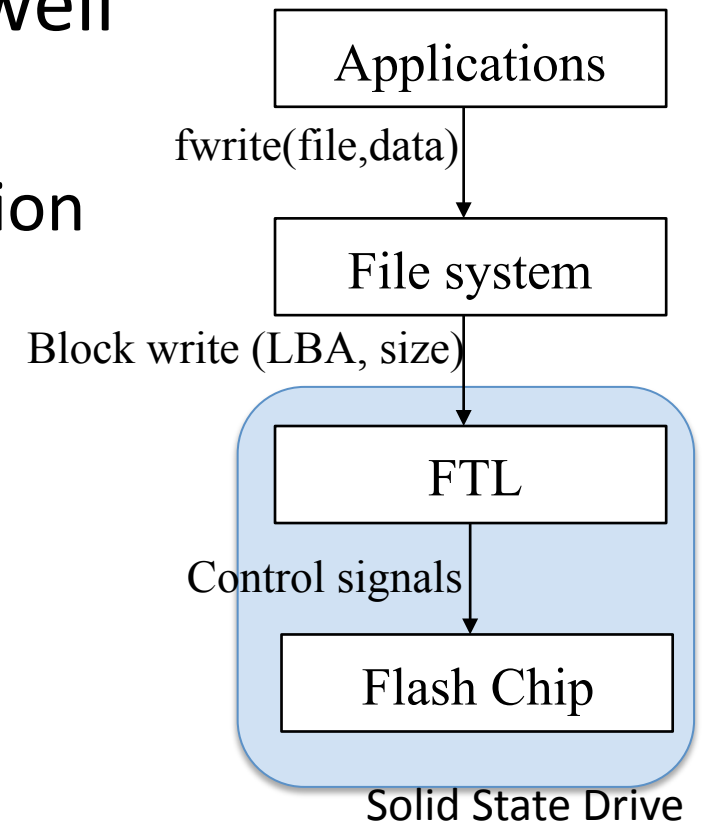

Experimental Setup

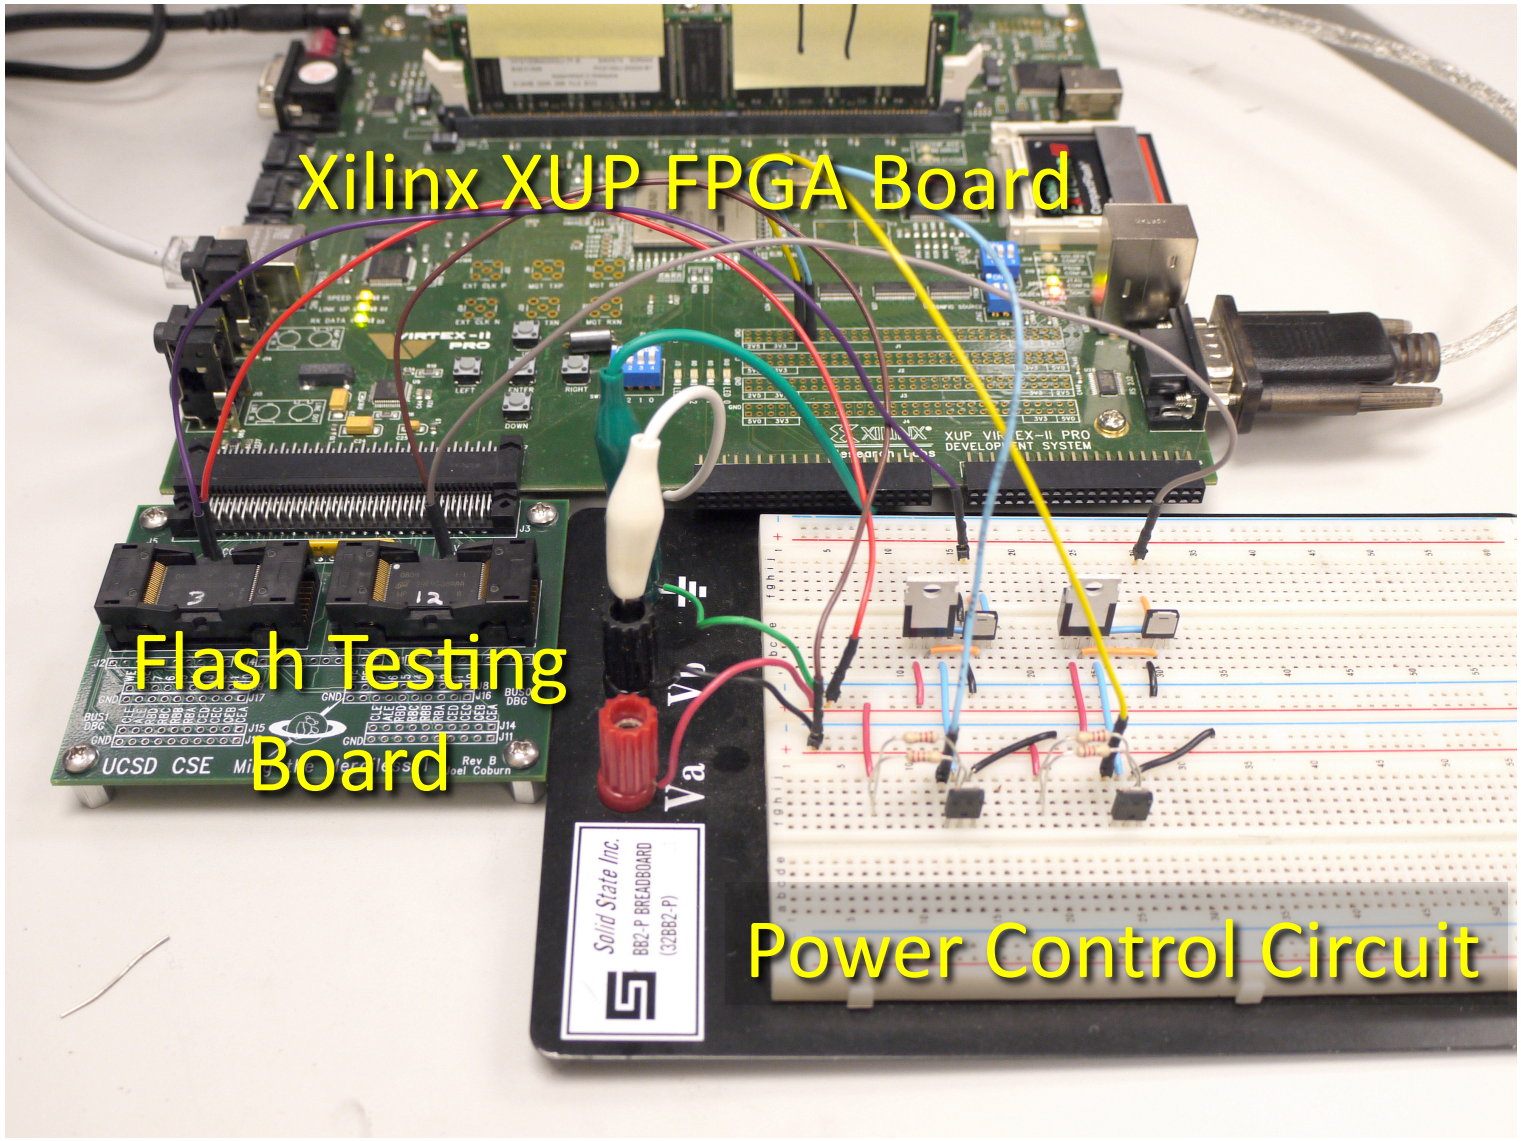

Experimental Setup

- Cutting power takes less than 5 us.

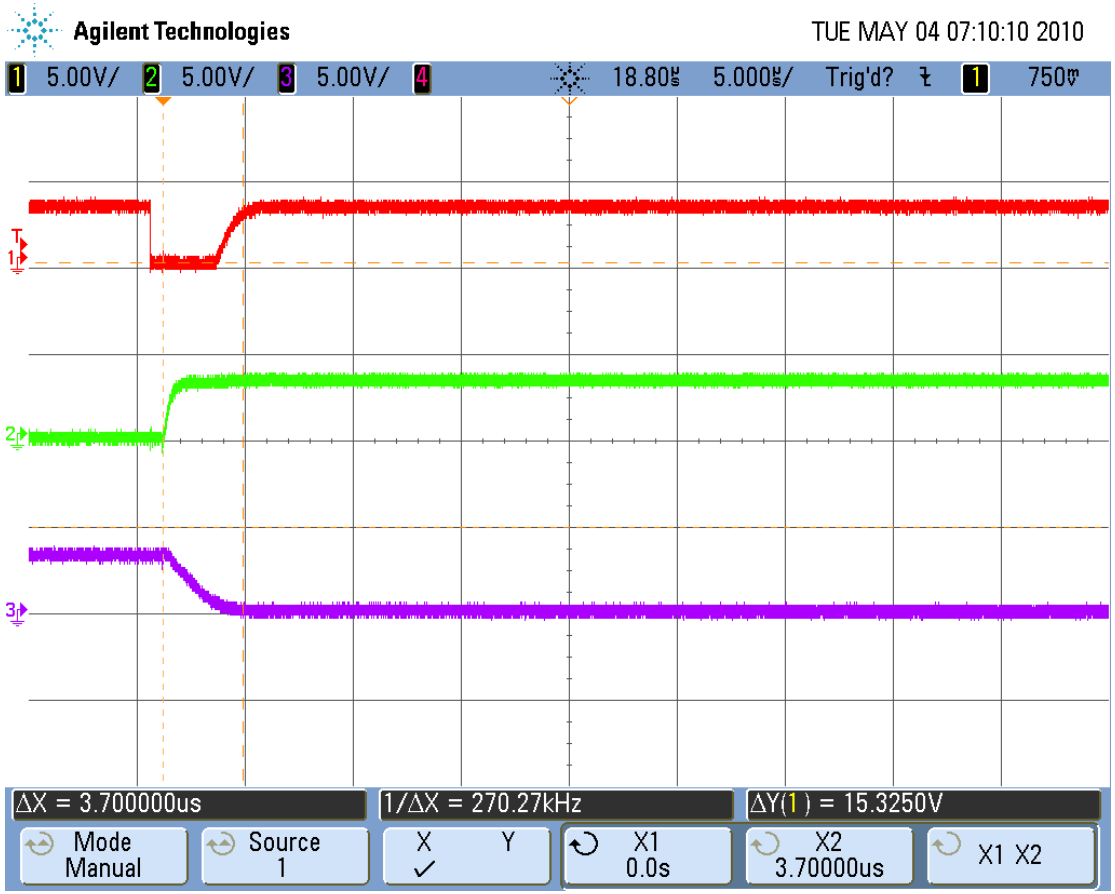

Experimental Setup

	Chip Name	Max PE Cycles	Tech Node	Capacity (Gb)	Bytes Page	Pages Block	Blocks Plane	Planes Die	Dies
MLC	A-MLC16	10,000		16	4096	128	2048	2	1
	B-MLC32	10,000	50nm	32	4096	128	2048	2	2
	B-MLC32-2	5,000	34nm	32	4096	256	2048	2	1
	B-MLC128-2	5,000	34nm	128	4096	256	2048	2	4
	D-MLC32	5,000		32	4096	128	4096	1	2
	E-MLC8	10,000		8	4096	128	1024	1	2
SLC	A-SLC2	100,000		2	2048	64	1024	2	1
	A-SLC4	100,000		4	2048	64	4096	1	1
	A-SLC8	100,000	60nm	8	2048	64	4096	2	1
	B-SLC2	100,000	50nm	2	2048	64	2048	1	1
	B-SLC4	100,000	72nm	4	2048	64	2048	2	1
	E-SLC8	100,000		8	2048	64	4096	1	2

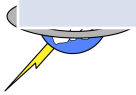

Power Failure During Program

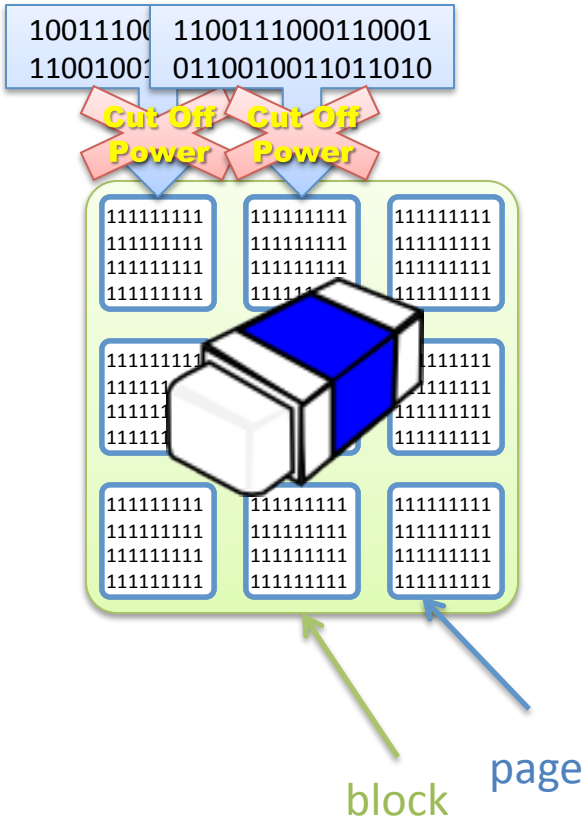

1. Erase a block
2. Program a page with data
3. Cut power after page program command begins
4. Repeat 2 and 3 until we programmed all pages in a block.
5. Measure the bit error rate

Power Failure During Program

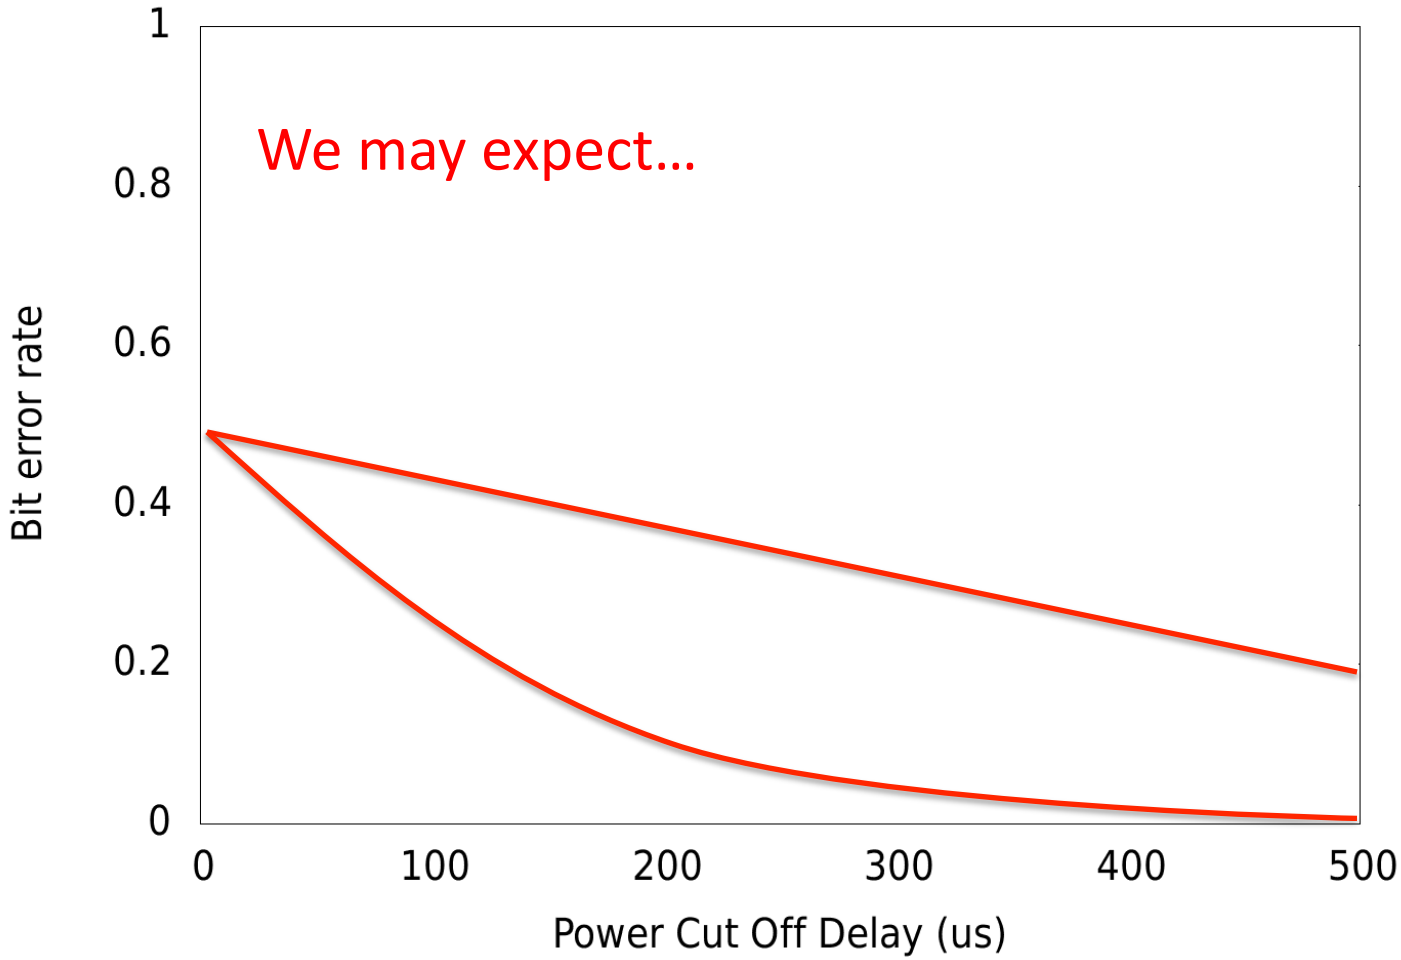

Power Failure During Program (MLC)

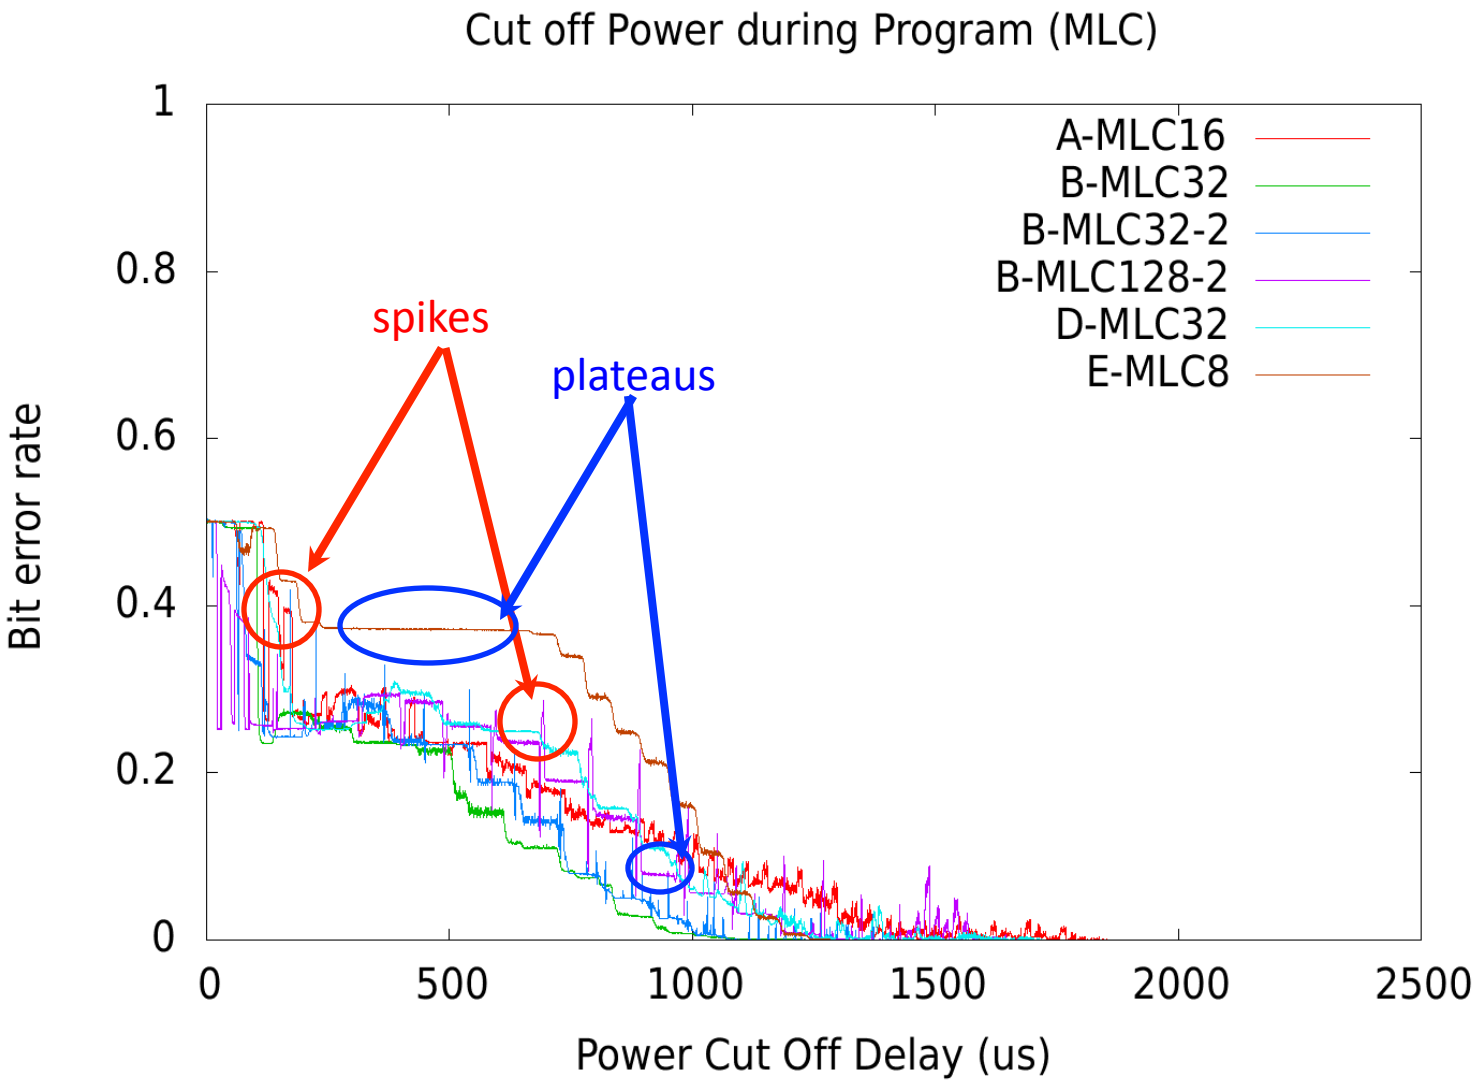

Power Failure During Program (SLC)

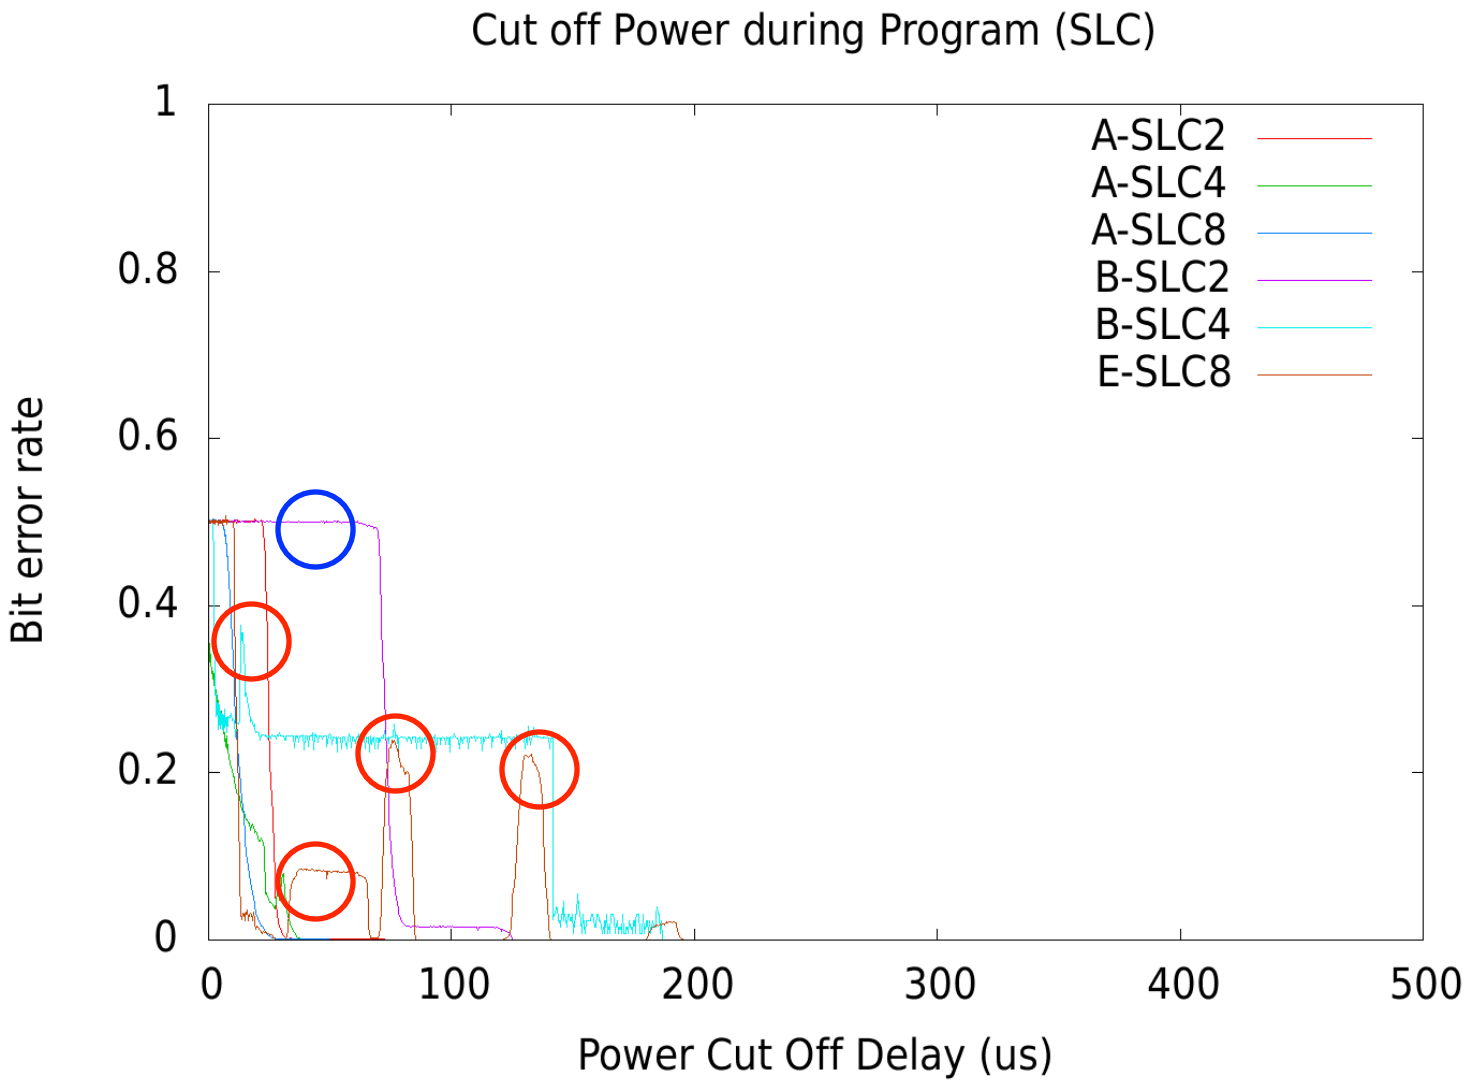

Power Failure During Program

- The bit error rate does not decrease monotonically
  - Plateaus and Spikes
  - What happened?

MLC Basics

MLC: Multi-Level Cell

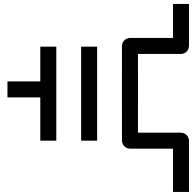

== 2+ bits

Fast Page (Fast Bit)

Slow Page (Slow Bit)

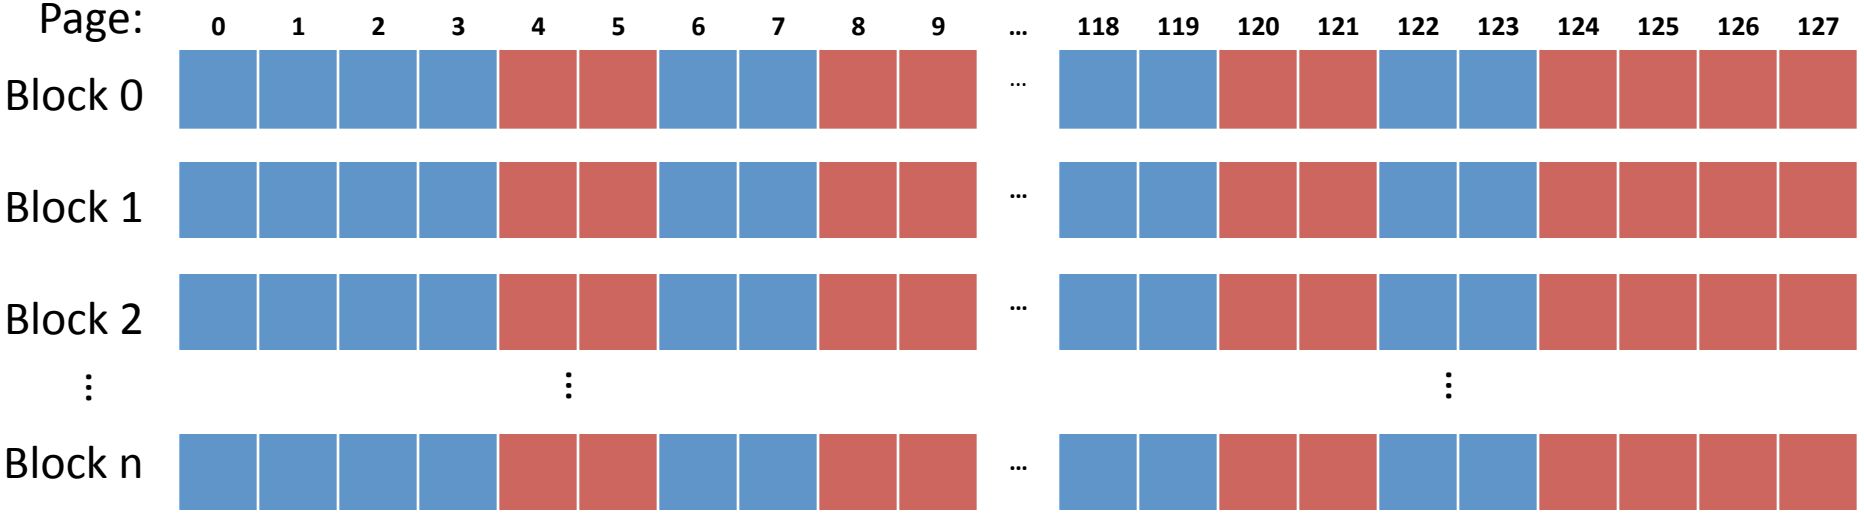

Detail of Power Failure During Program

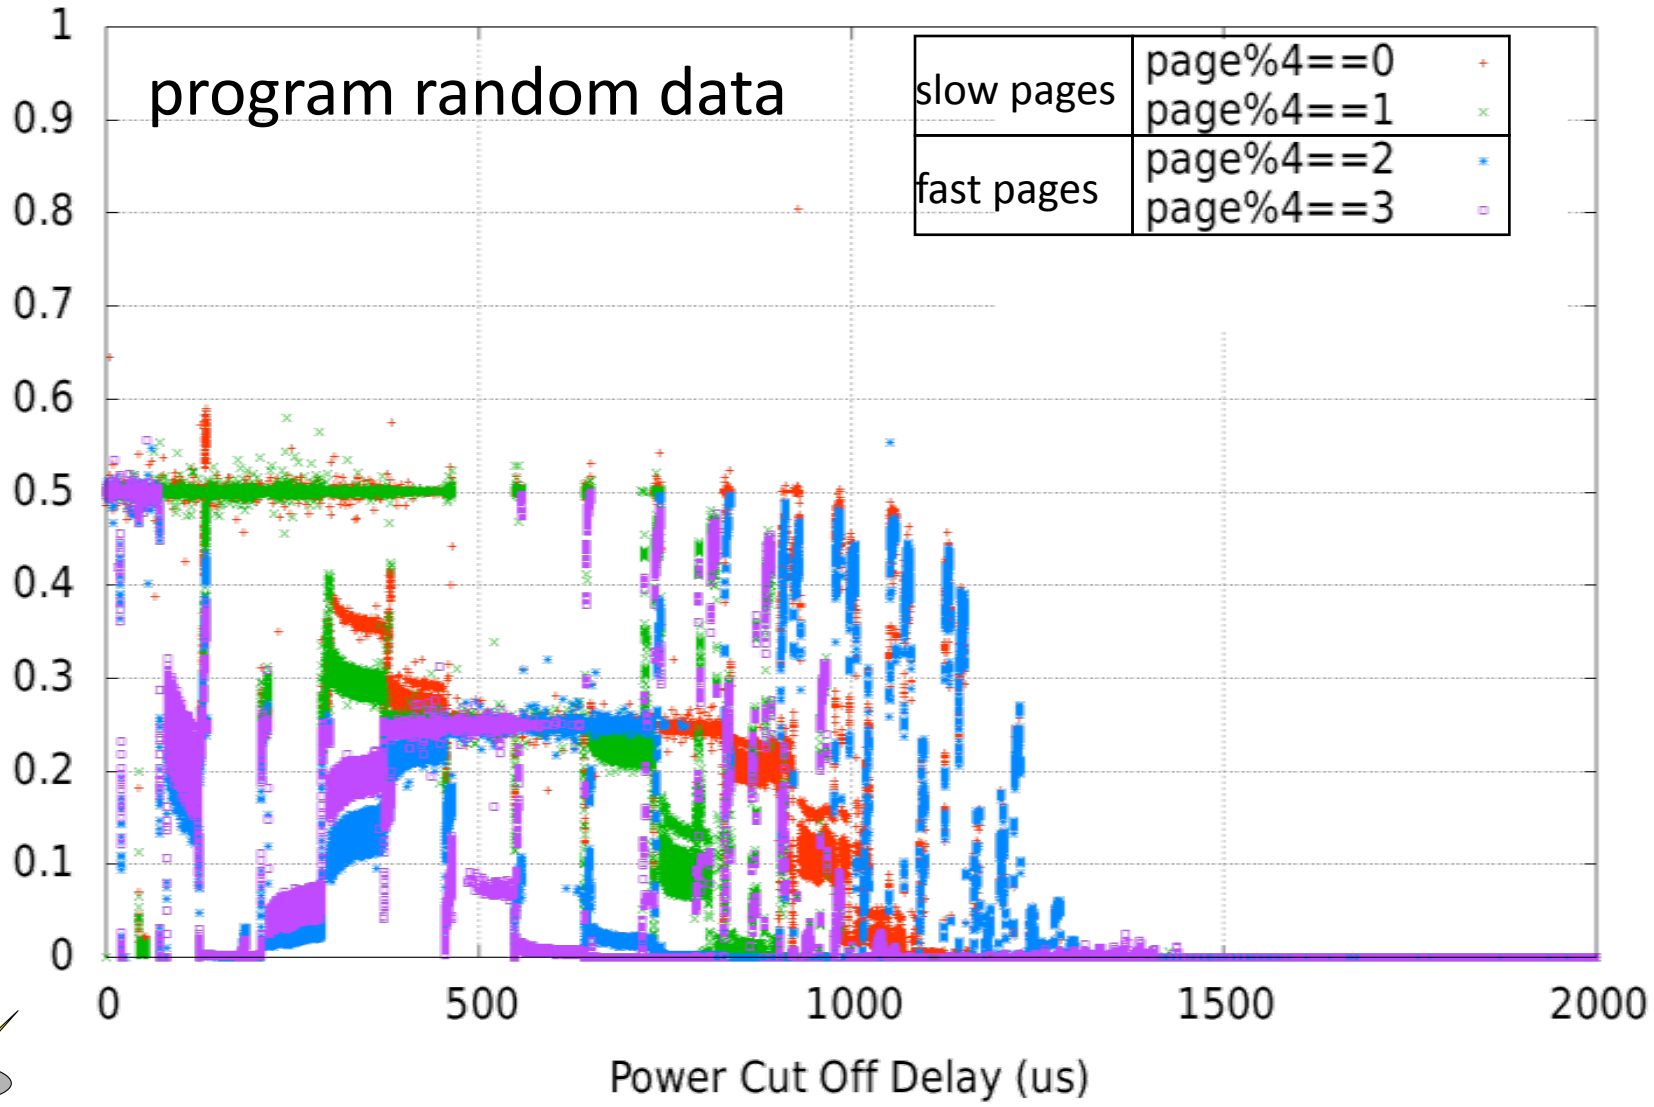

Decompose the Pattern

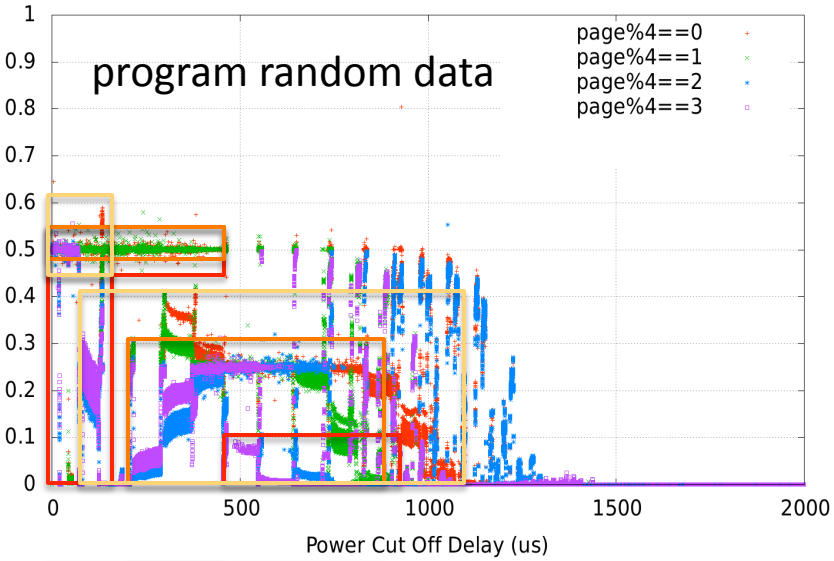

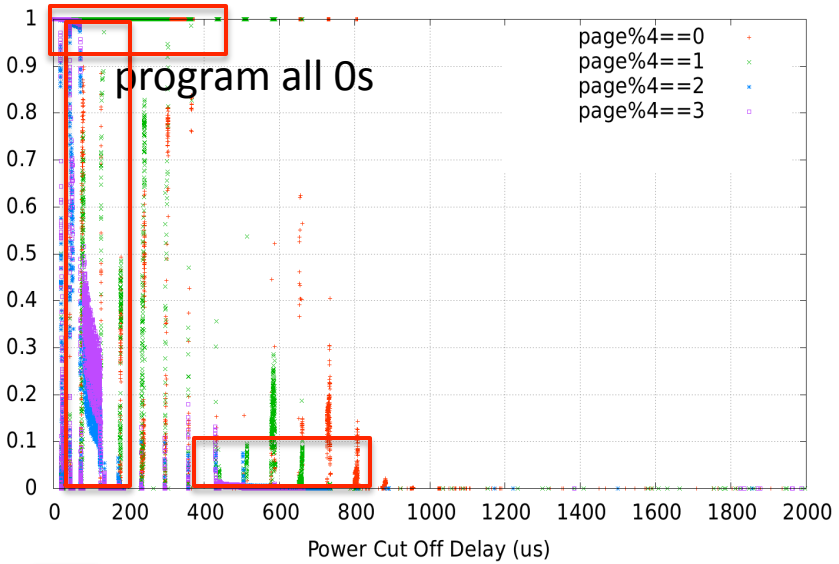

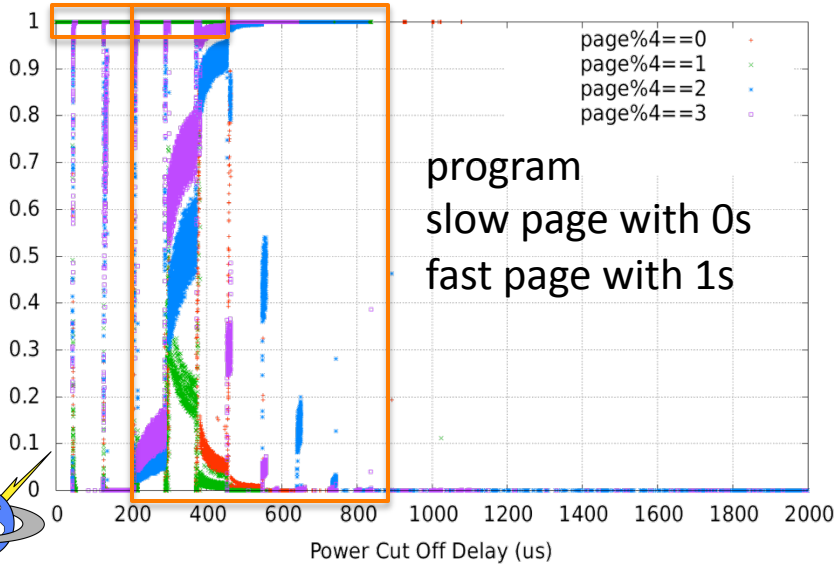

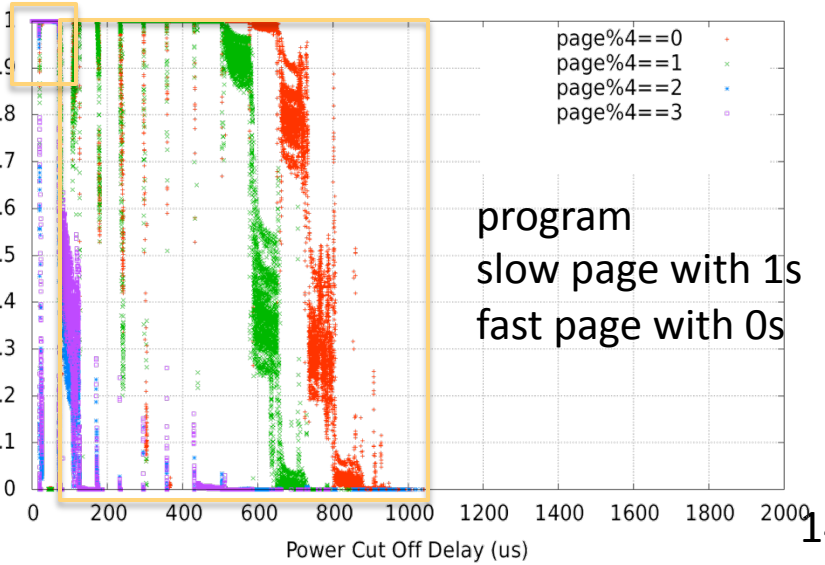

Decompose the Pattern (Cell State)

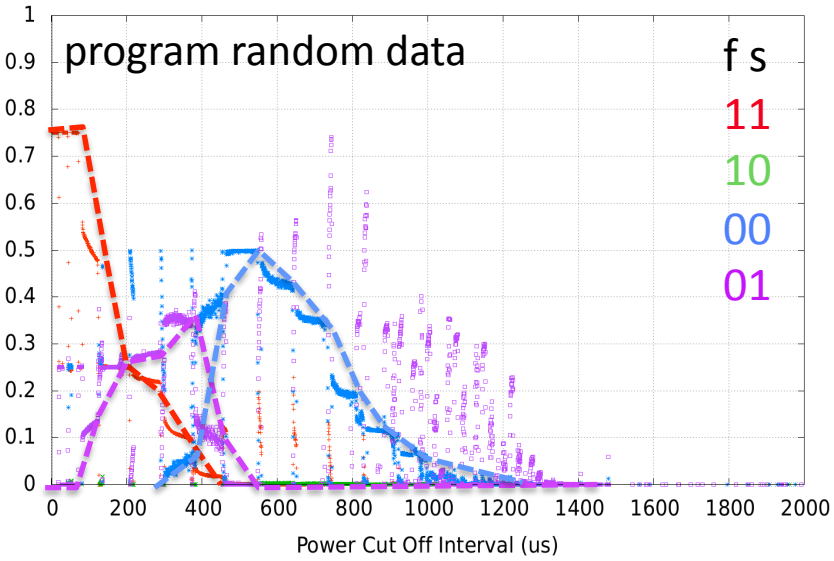

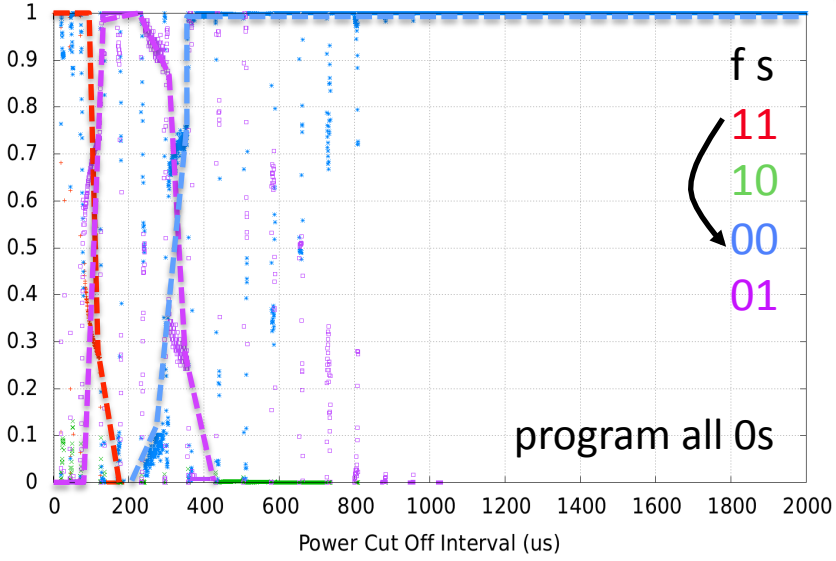

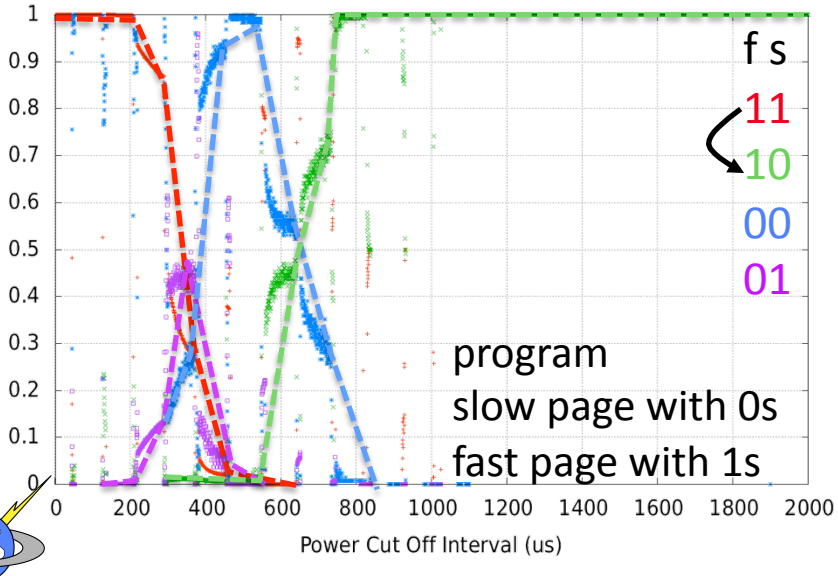

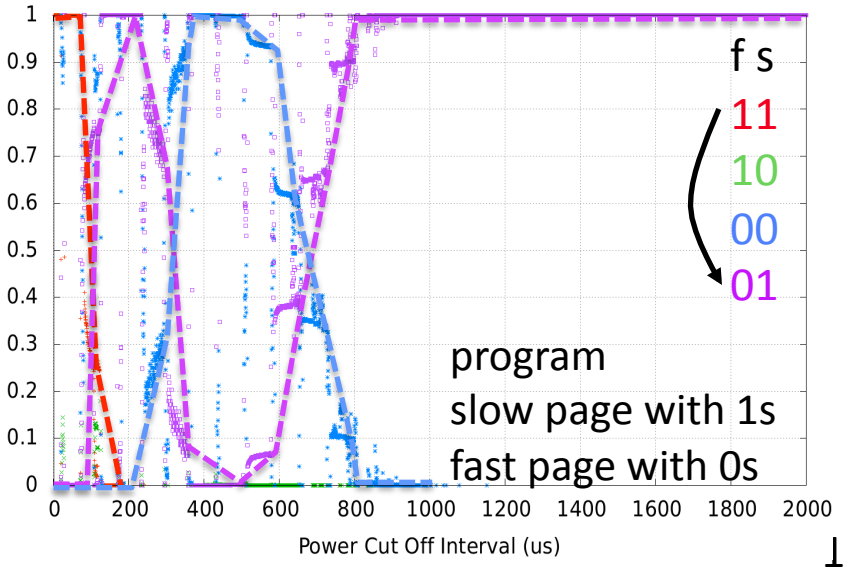

Power Failure During Erase

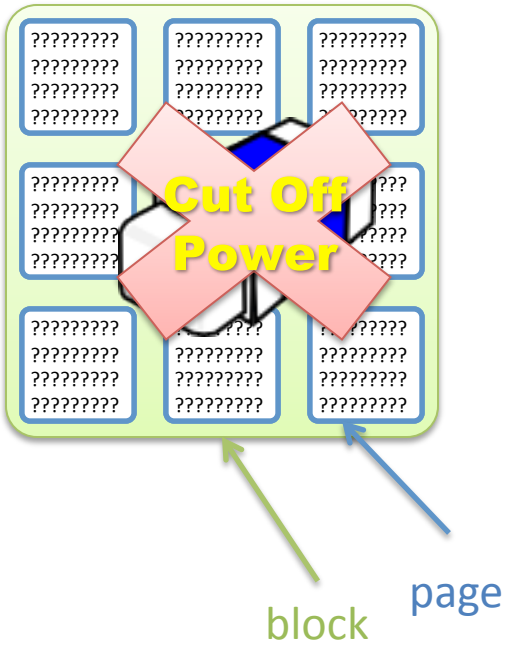

1. Erase a block
2. Program each page with data
3. Erase a block
4. Cut power after erase command begins
5. Measure the bit error rate

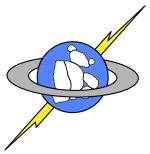

Power Failure During Erase (MLC)

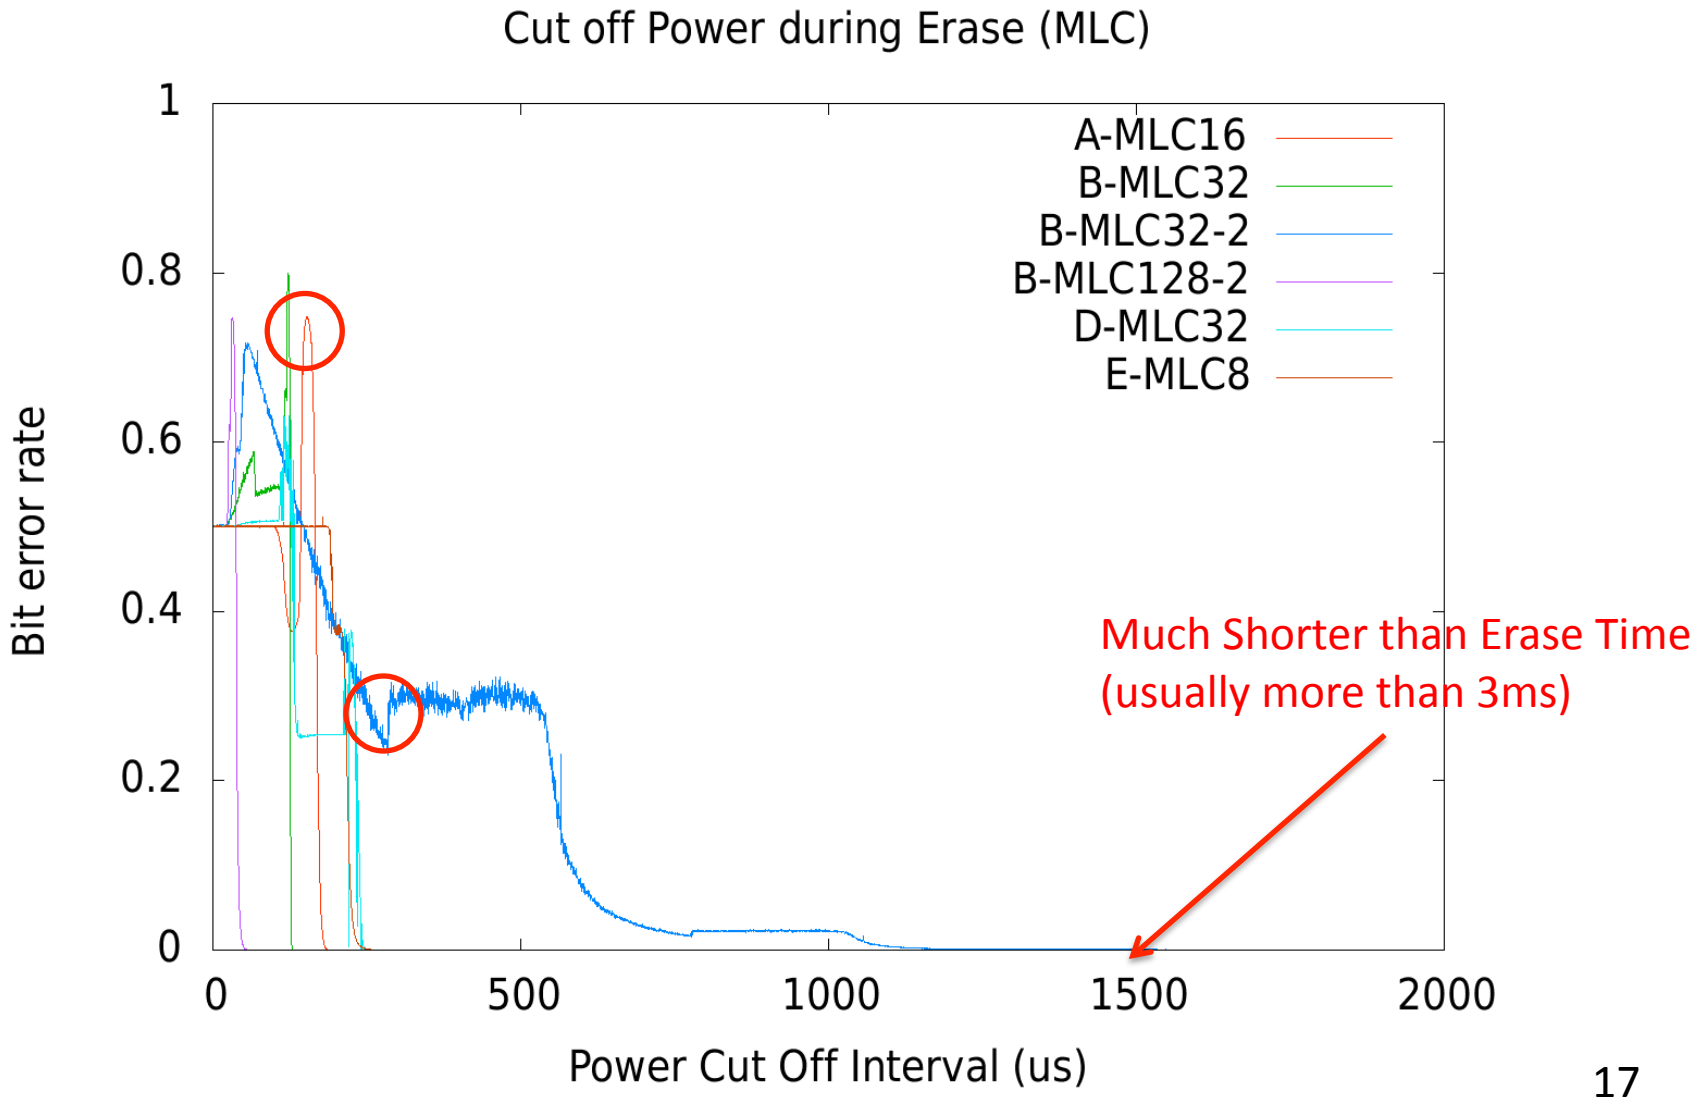

Power Failure During Erase (SLC)

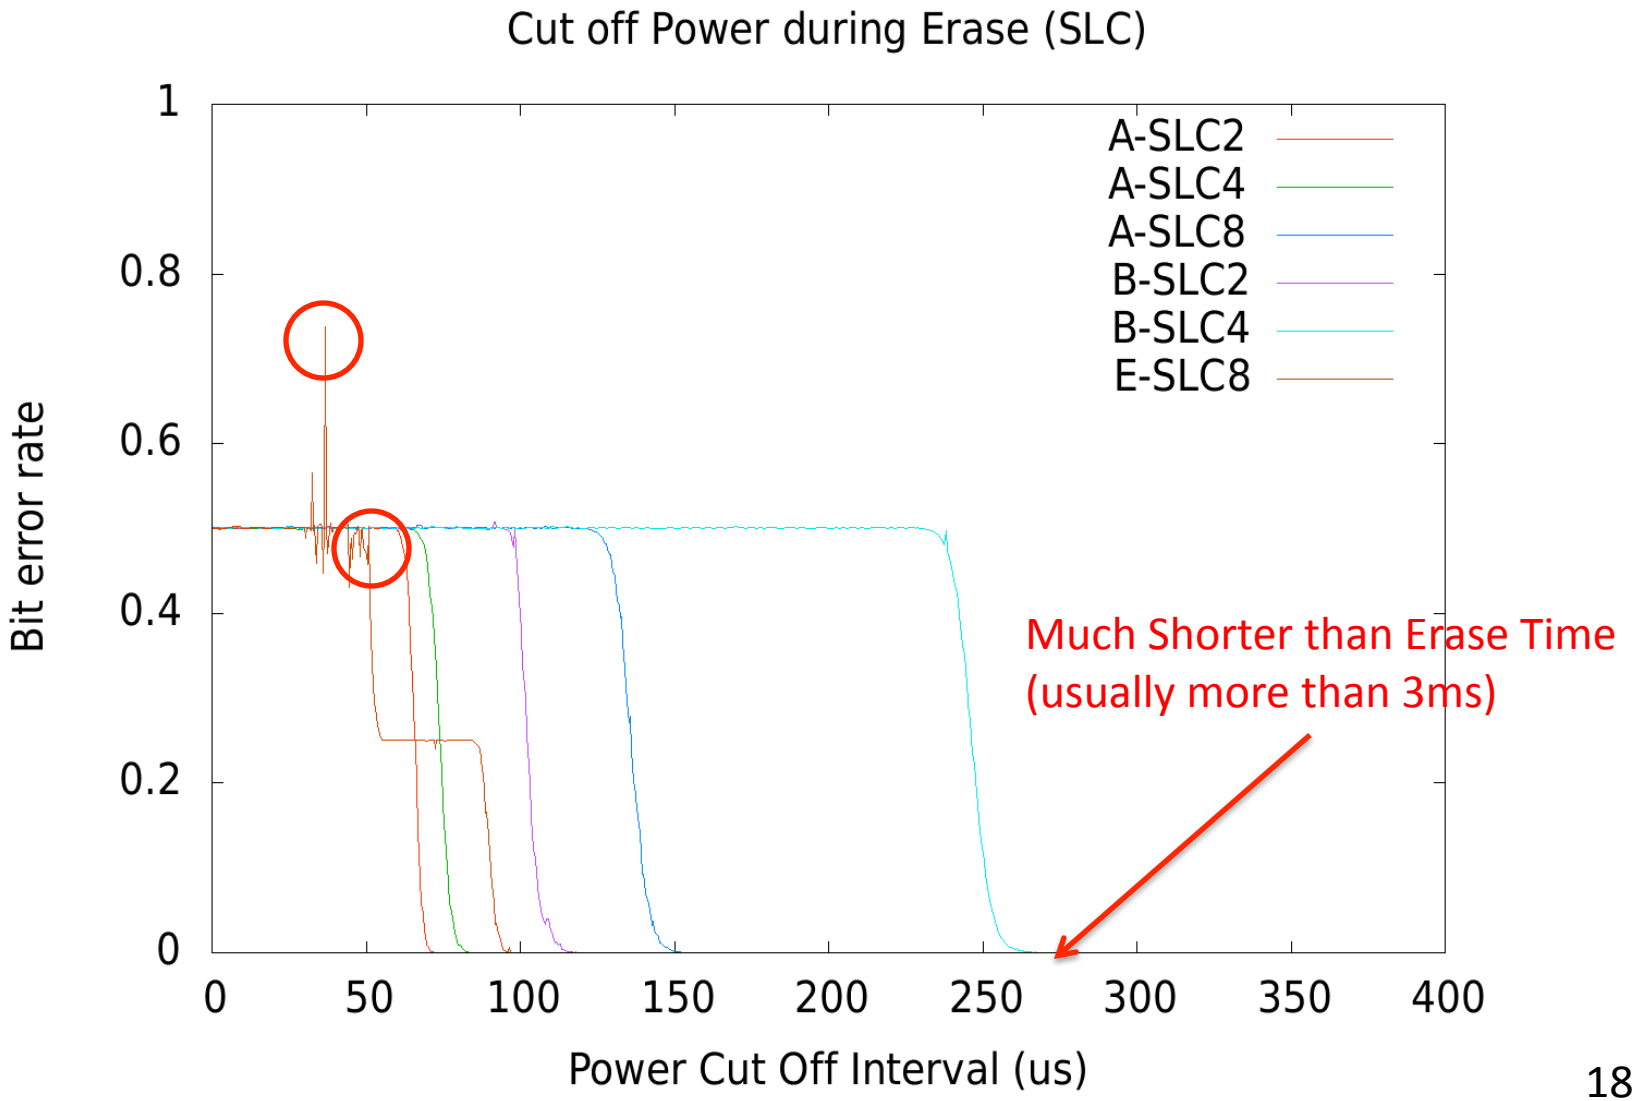

Power Failure During Erase

- The block can return to all 1s with much shorter time than regular erase time
  - Can we successfully program the block?

Program a Block Erased with Power Failure

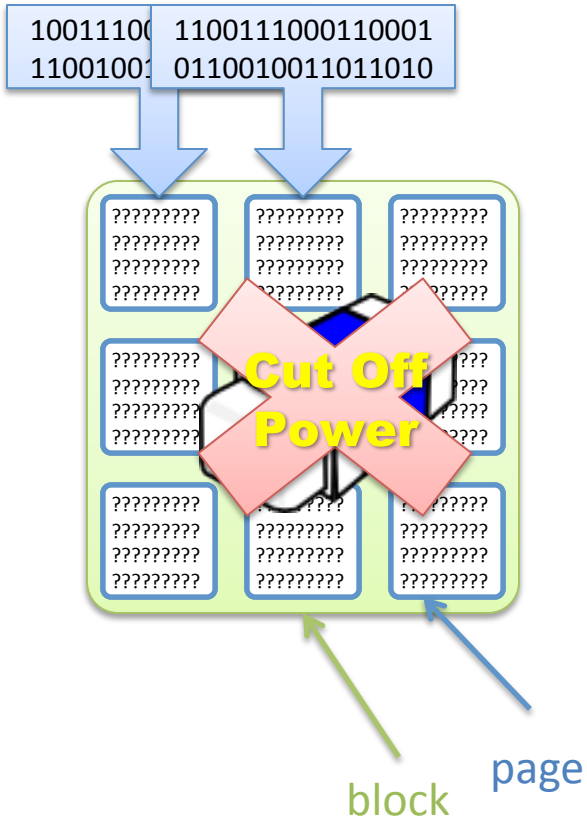

1. Erase a block
2. Program each page with data
3. Erase a block
4. Cut power after erase command begins
5. Program the each page in the block with random data
6. Measure the bit error rate

Program a Block Erased with Power Failure

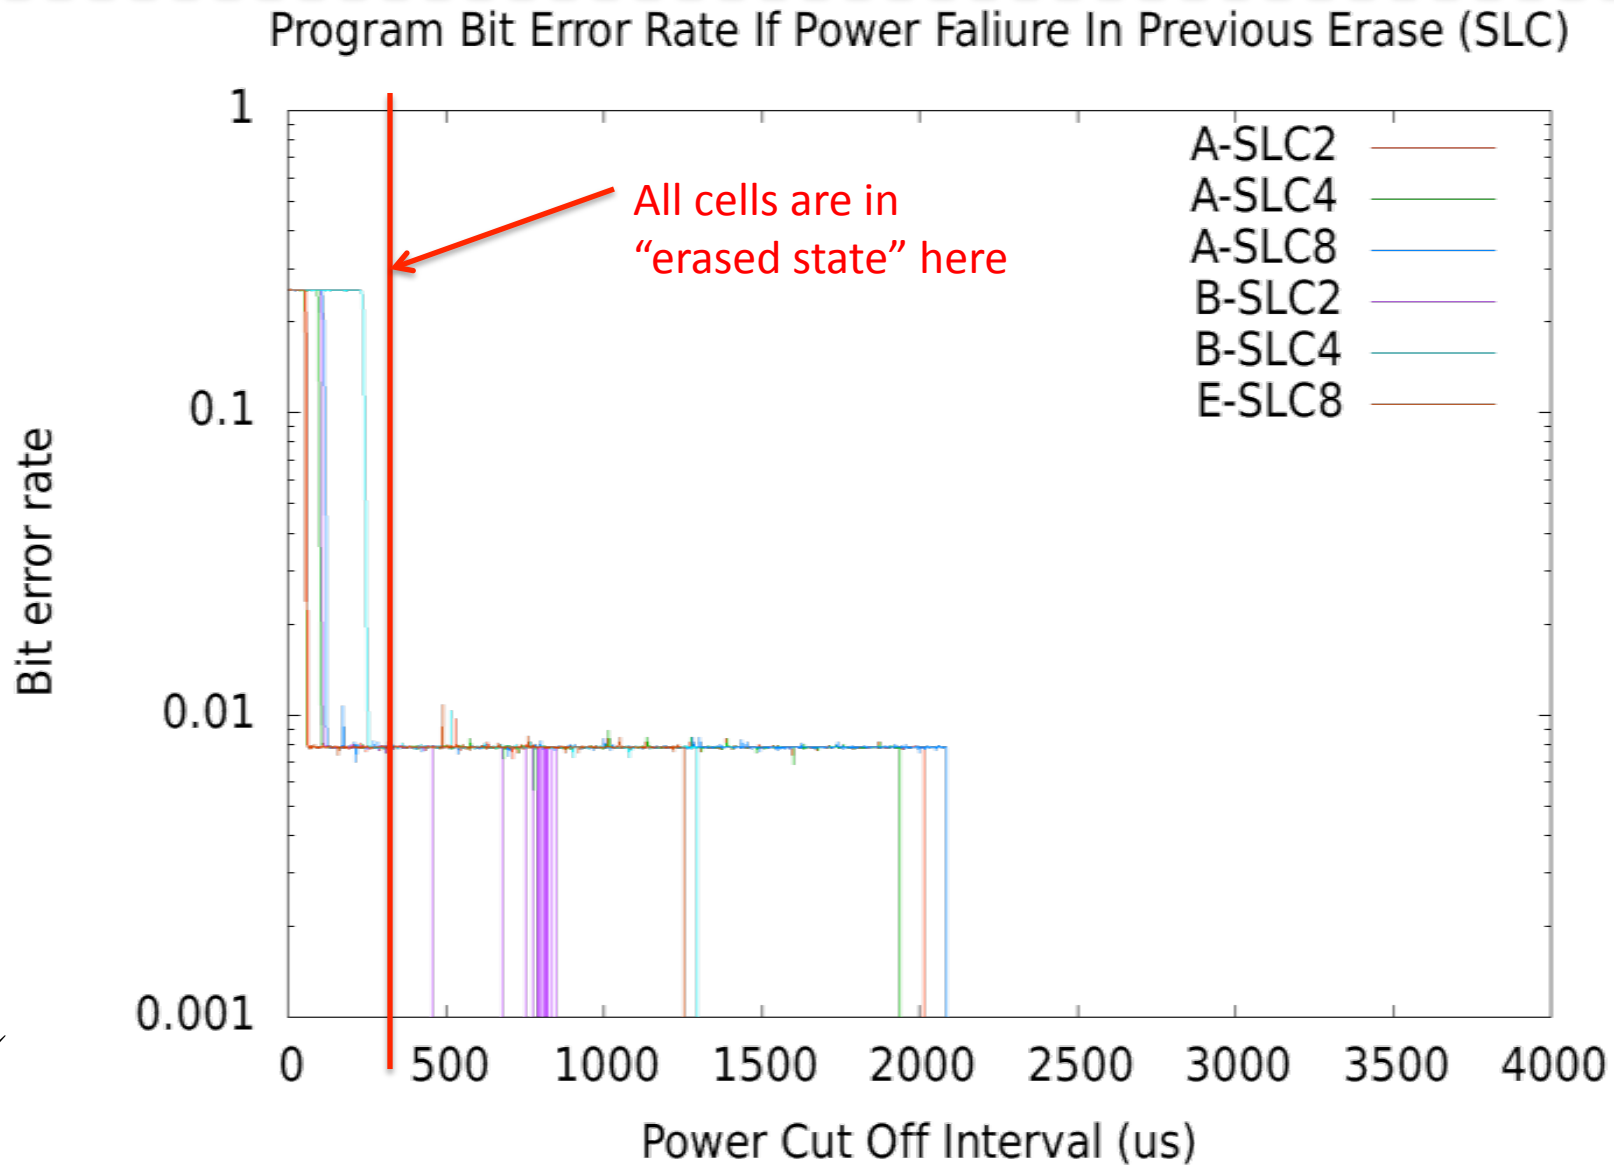

Program a Block Erased with Power Failure

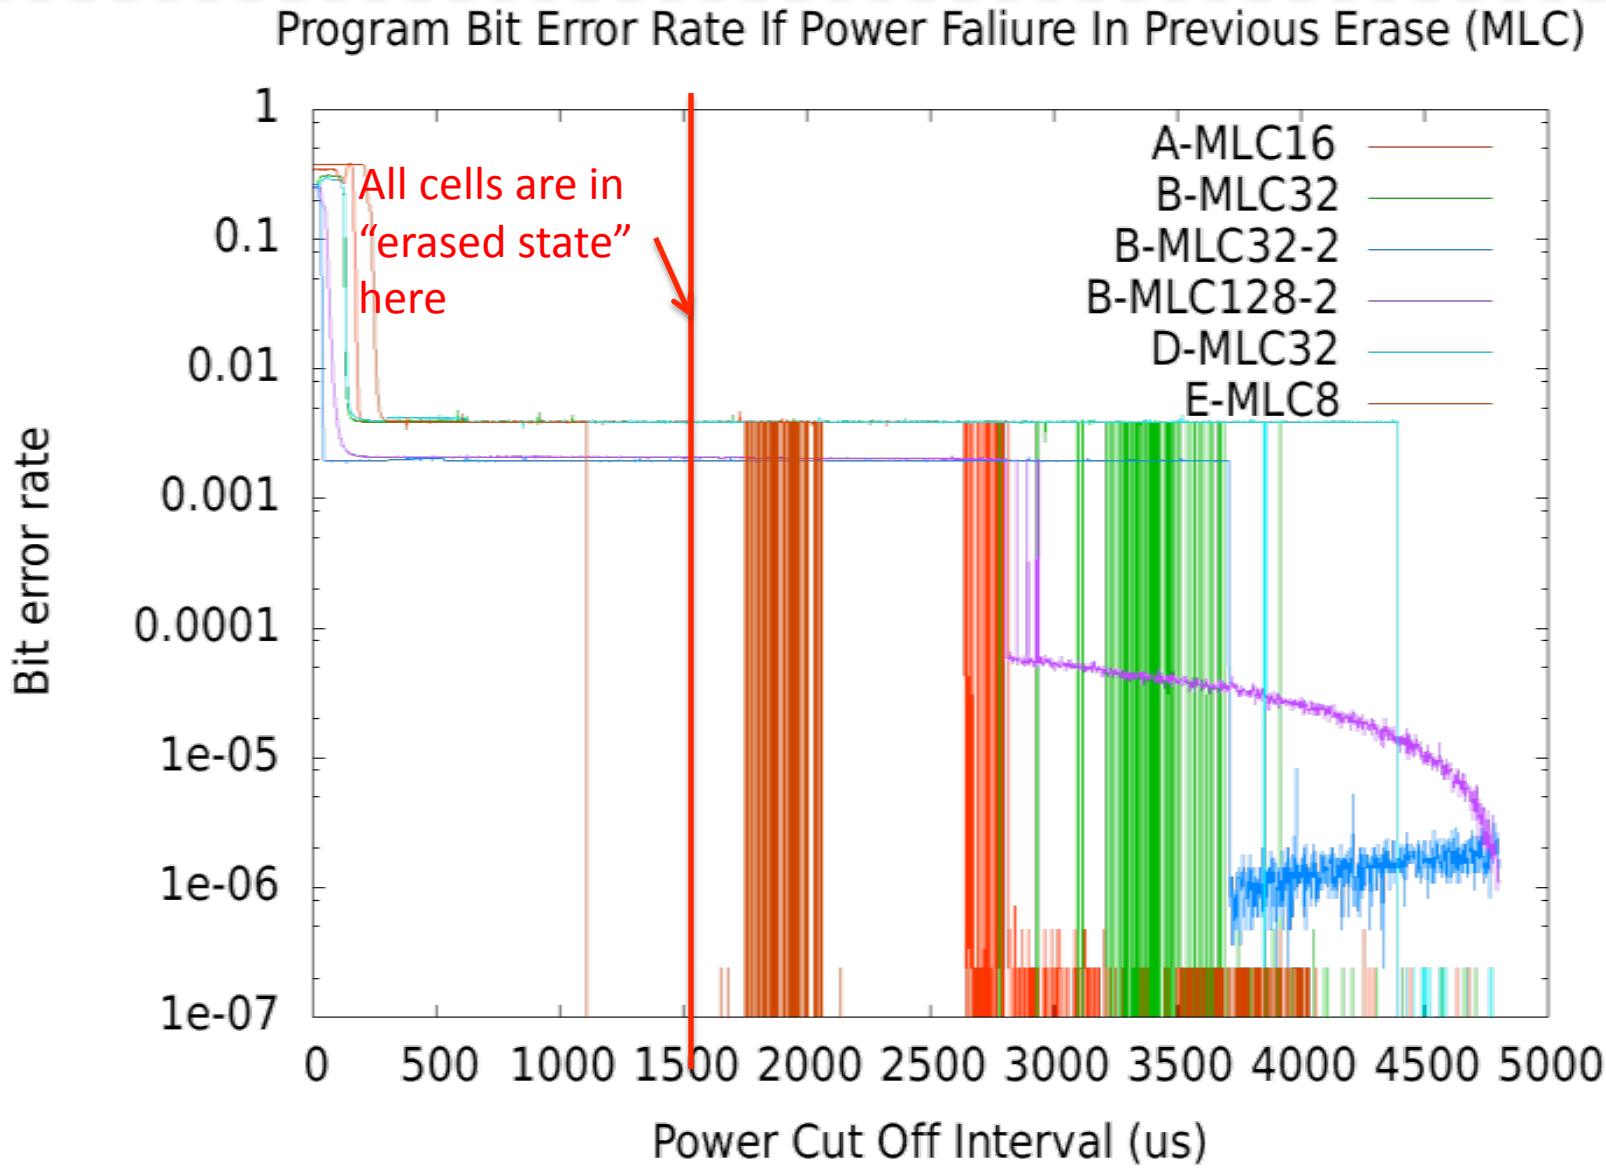

Conclusion & Future Work

- Power loss is a serious problem in SSD
- Decrease of bit error rate is not a monotone function to power cut off interval.
- The bit error rate of program/erase without error is data-related.
  - On-chip coding
  - Power-failure tolerate data encoding
- An erase may appeared to succeed, but the block may not be to program.

Thank You

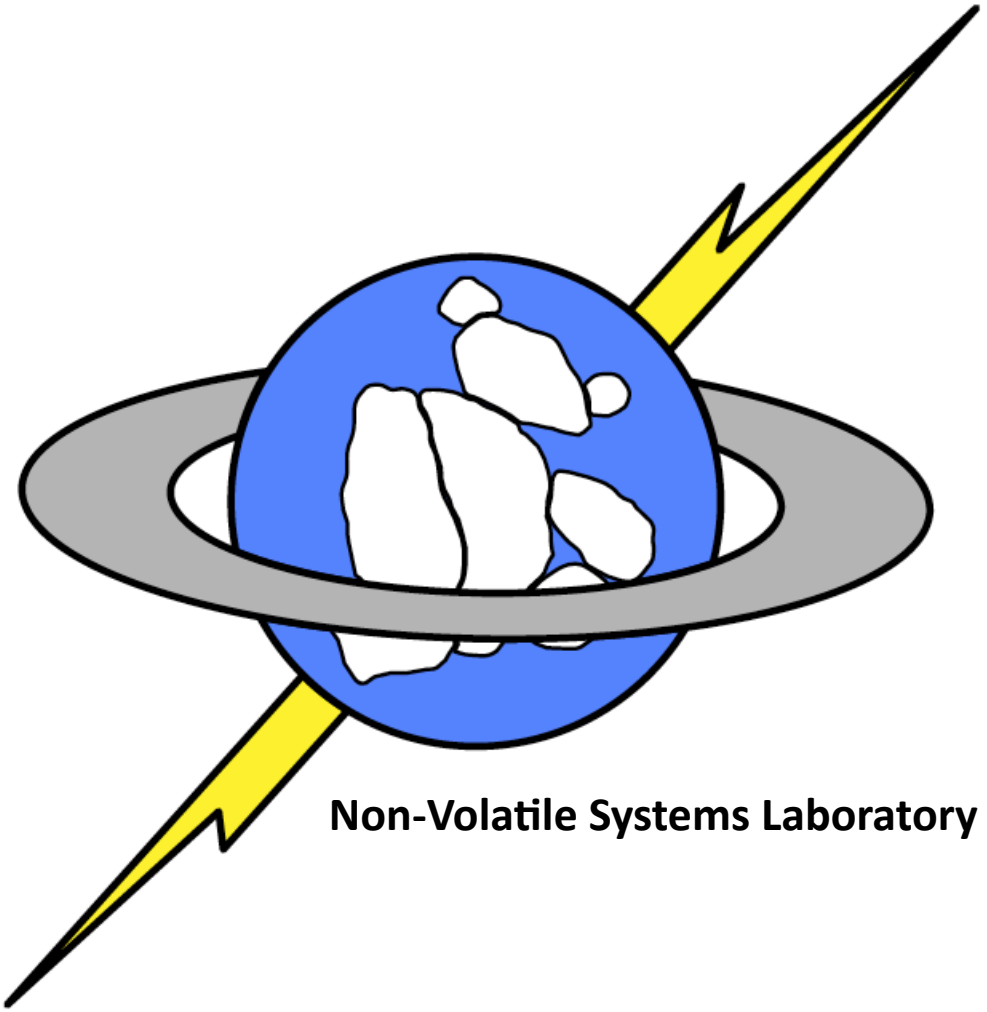

Non-Volatile Systems Laboratory