Improving NAND Endurance by Dynamic Program and Erase Scaling

Jihong Kim

Department of Computer Science and Engineering
Seoul National University, Korea

NVRAMOS 2013

October 24, 2013

Trend 1 : NAND Capacity

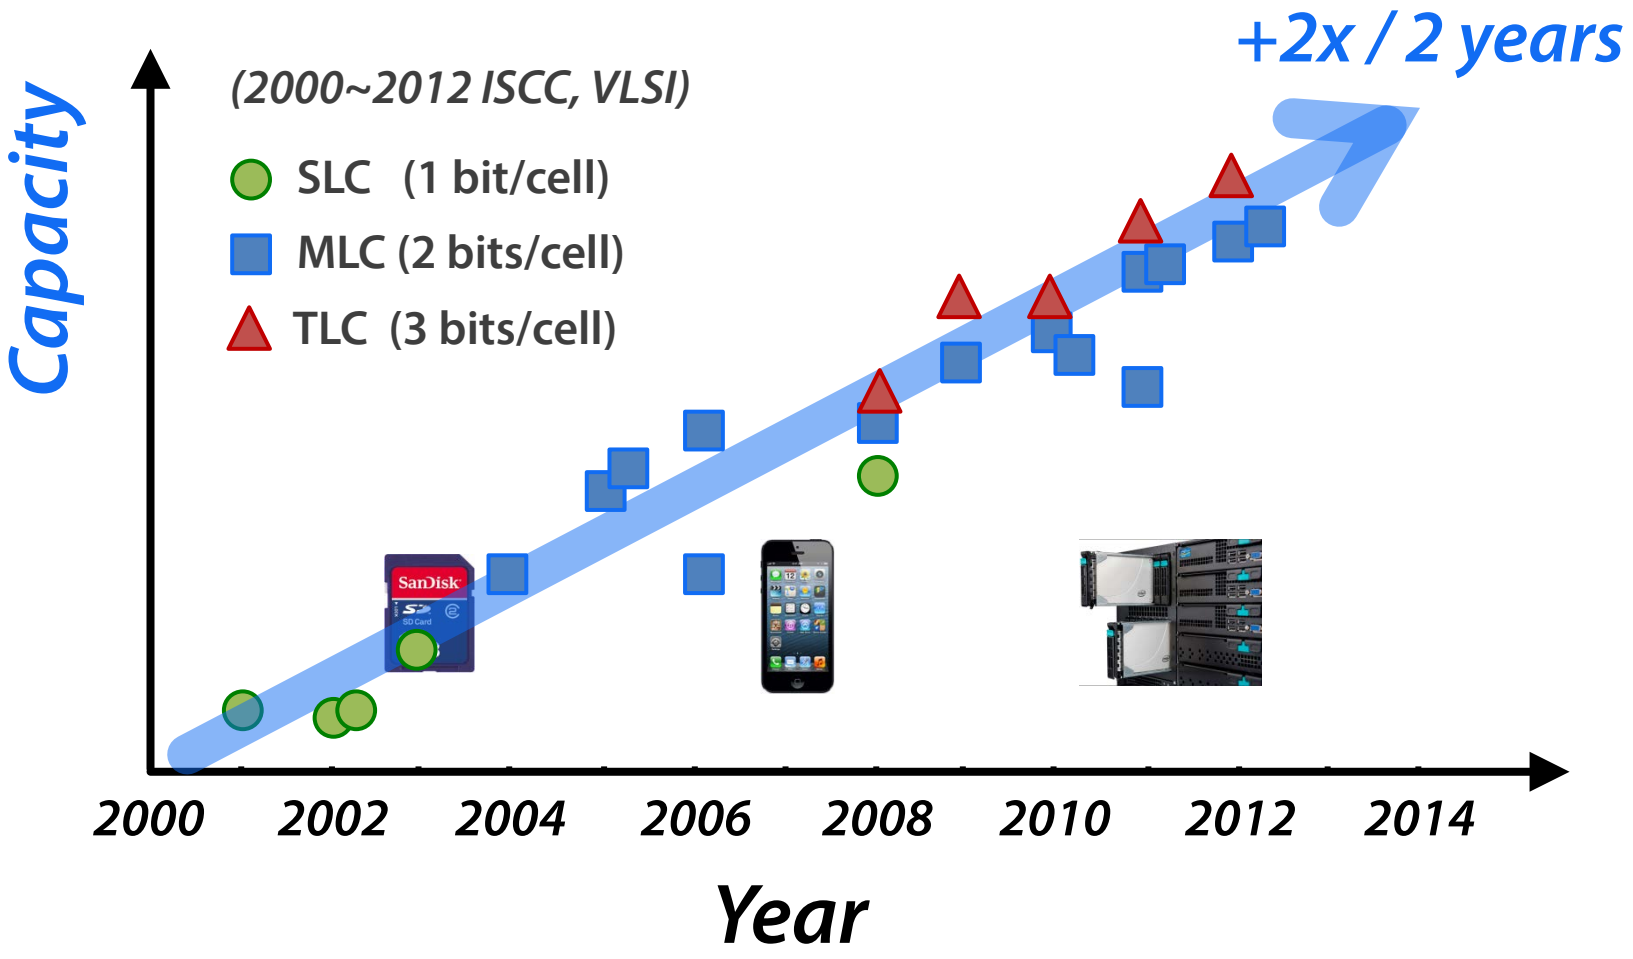

The cost-per-bit of NAND devices is continuously improving.

Trend 2 : NAND Endurance

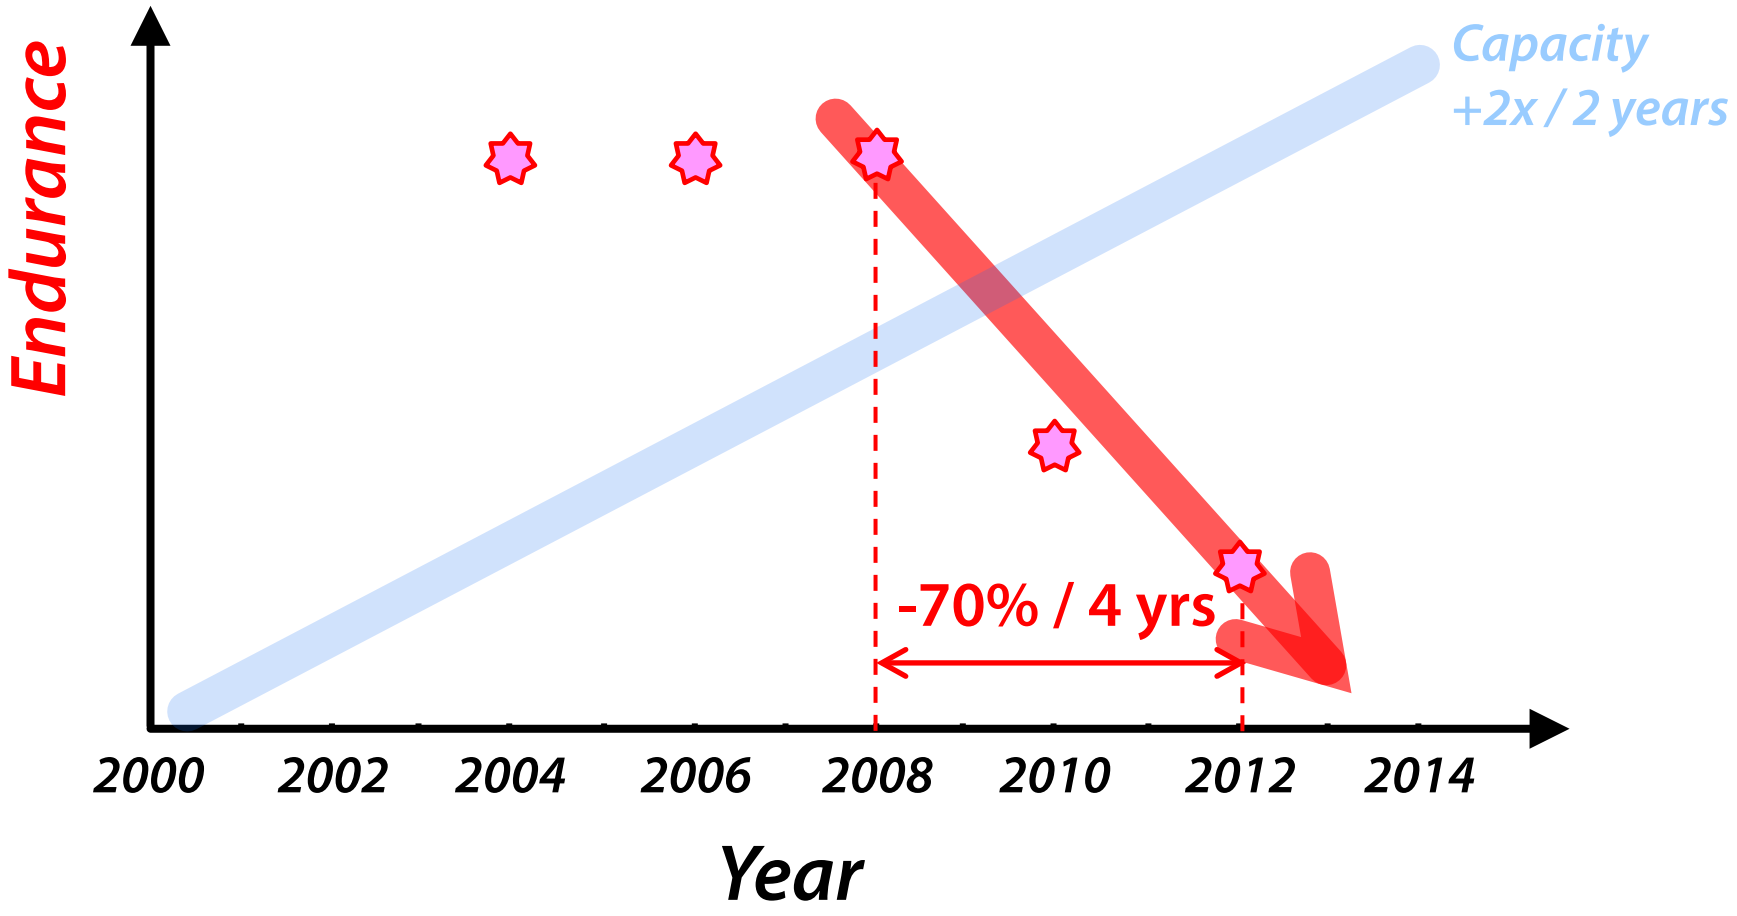

The NAND endurance is drastically decreased last 4 years as a side effect of recent advanced technologies.

Trend 3 : Total Amount of Writes

Total amount of writes = *Endurance* X *Capacity*

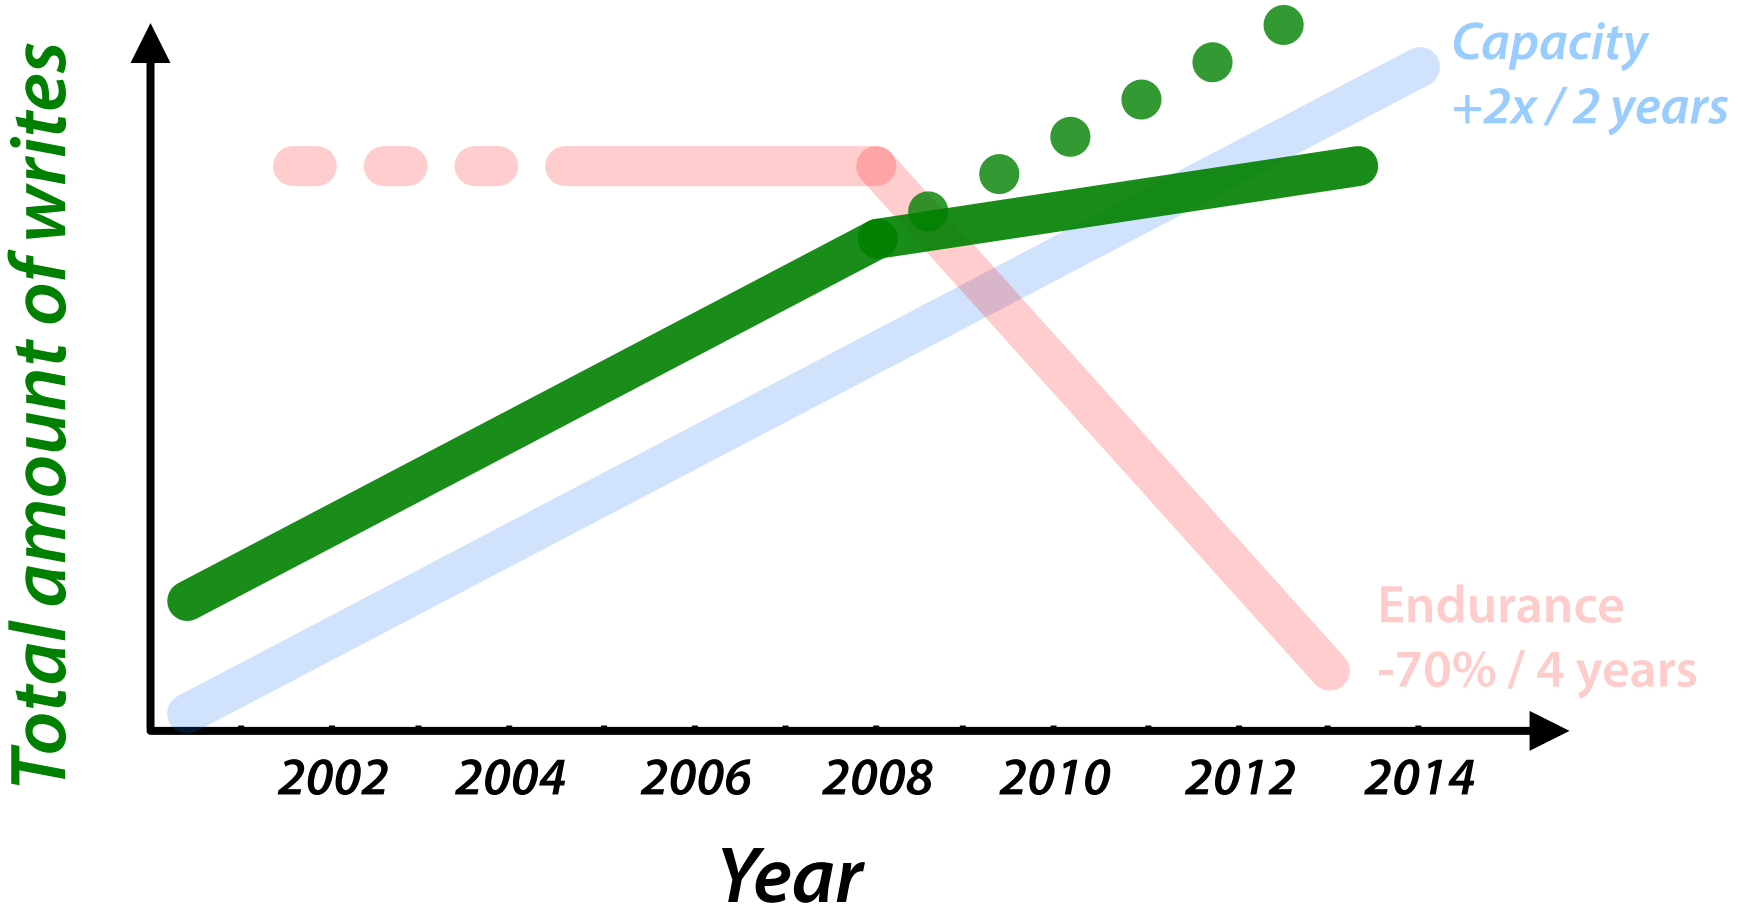

The total amount of writes of NAND-based storage does not increase as much as we expected.

Trend 4: Lifetime of NAND-Based Storages

$$Lifetime = \frac{Endurance \times Capacity}{Wday \times WAF}$$

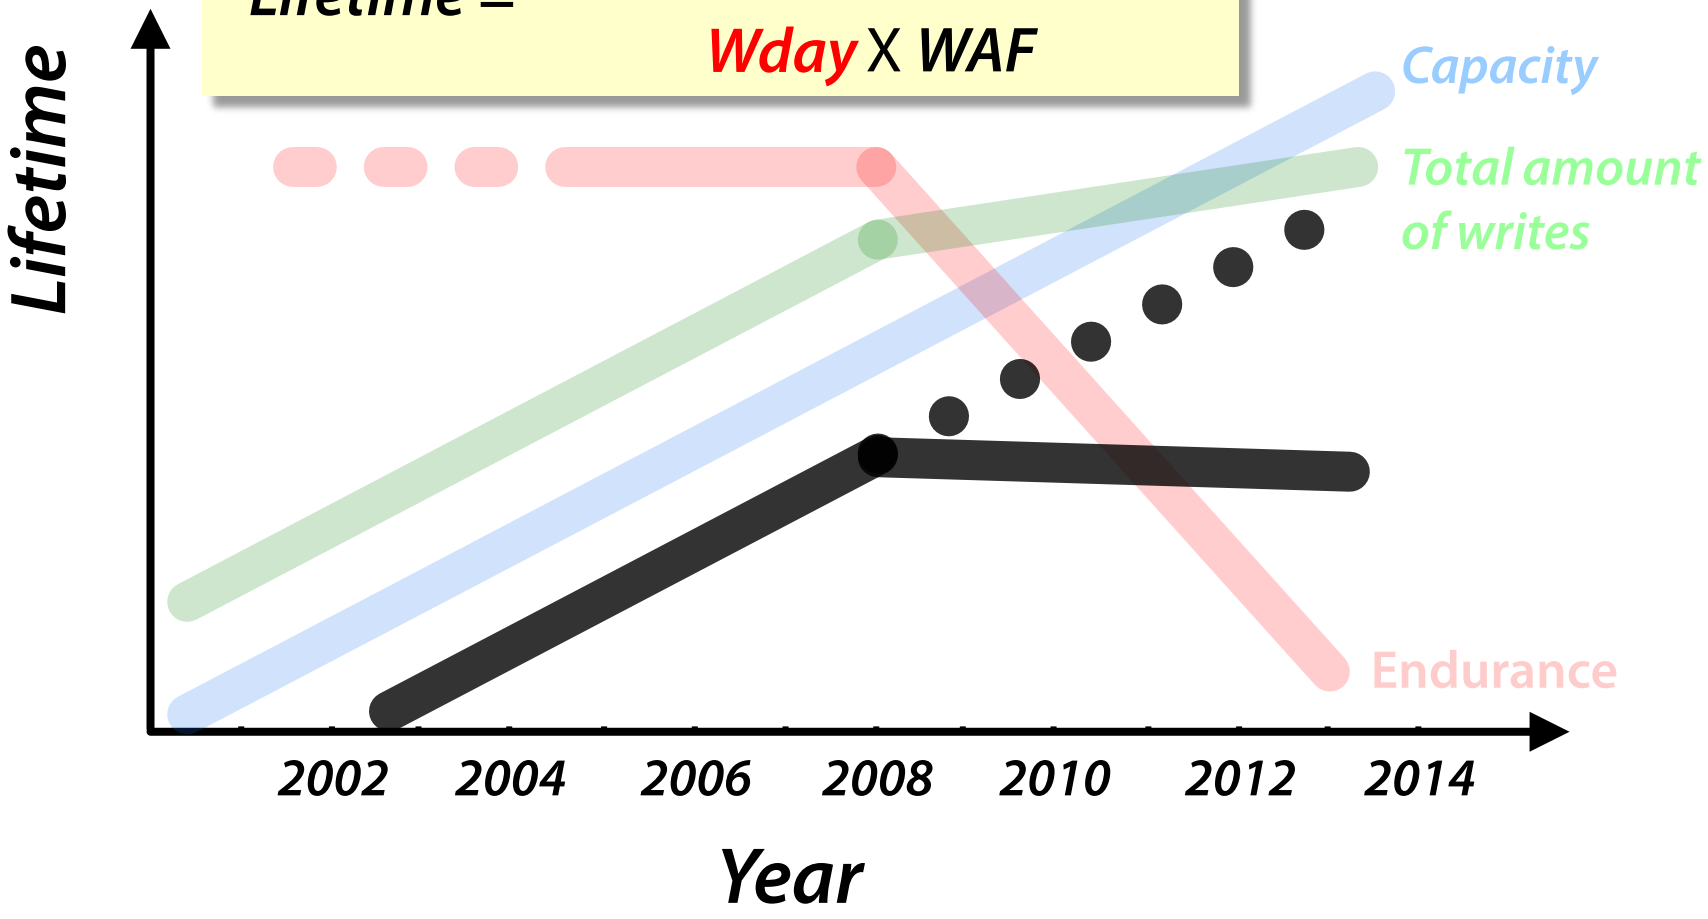

Existing Lifetime Enhancing Schemes

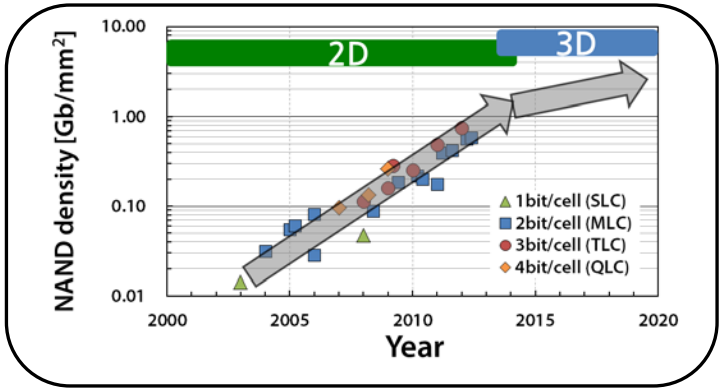

$$Lifetime = \frac{Endurance \times Capacity}{W_{day} \times WAF}$$

① *Data compression*
② *Deduplication*
③ *Dynamic throttling*

Reducing WAF by increasing the efficiency of an FTL algorithm (e.g., garbage collection, wear leveling)

Our Goal

$$Lifetime = \frac{Endurance \times Capacity}{Wday \times WAF}$$

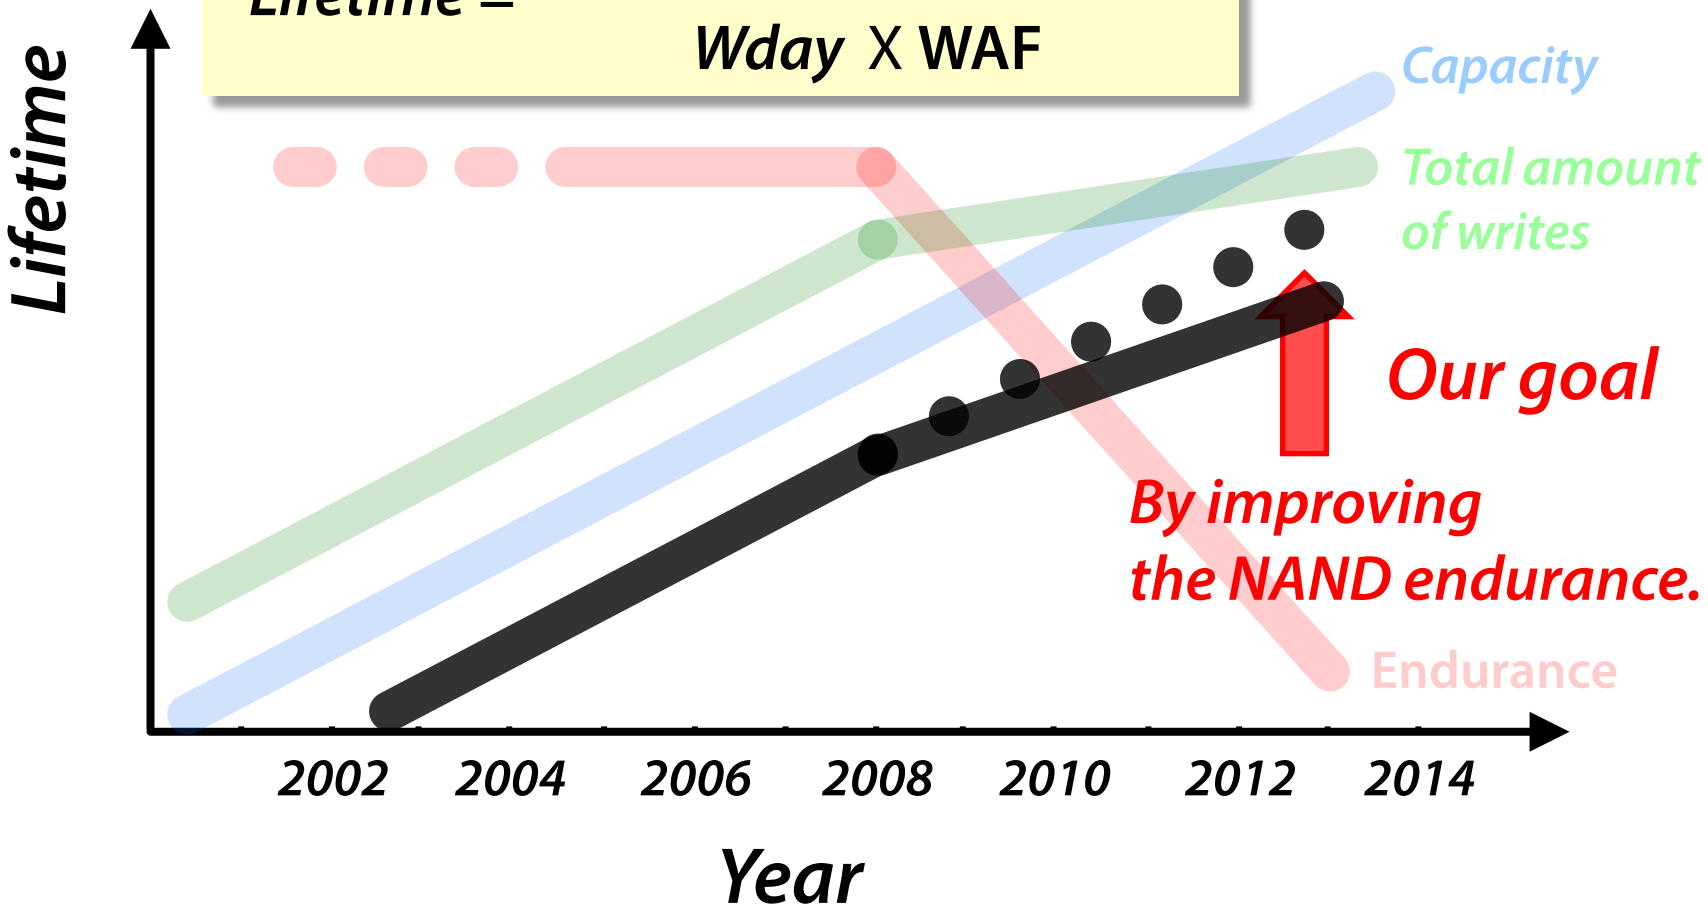

Improving the NAND endurance is required for sustainable growth in the NAND flash-based storage market.

Outline

- Introduction
- Motivation
- Key Components of the DPES Approach
  - Erase Voltage Scaling
  - Program Time Scaling
  - Dynamic Program and Erase Scaling
- Implementation of DPES-Aware FTL
- Experimental Results
- Conclusion

Motivation : Device Physics Model

Cross-section view of
NAND flash memory cells

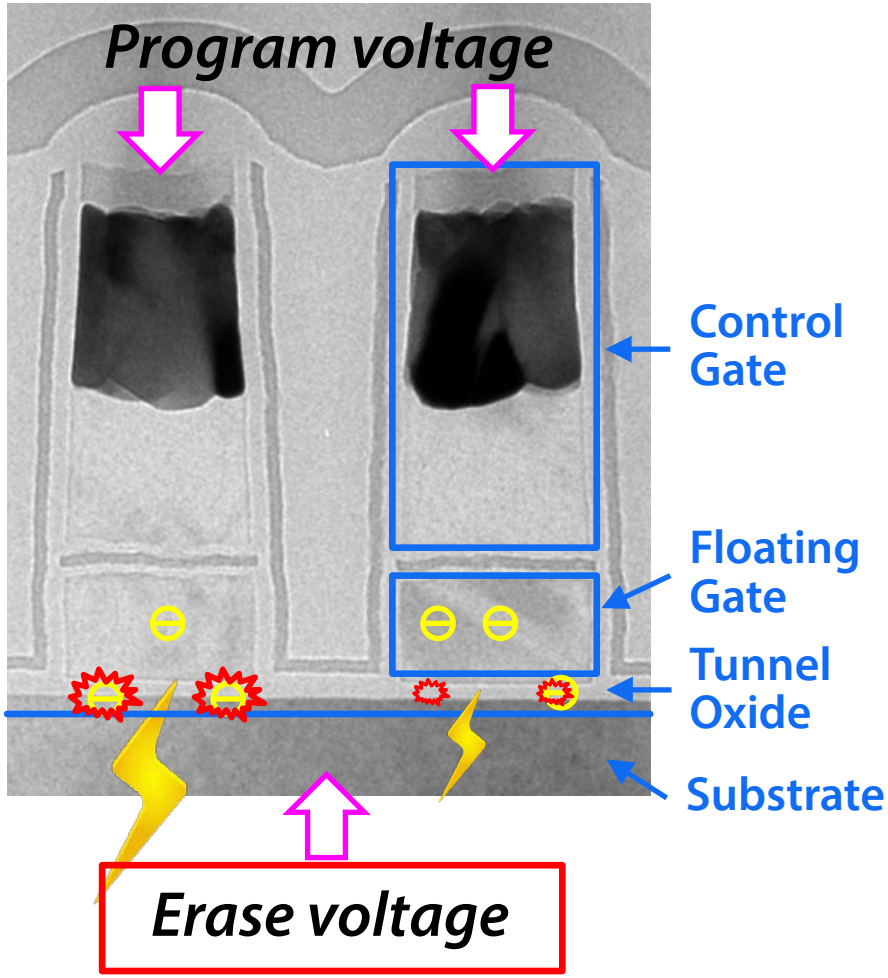

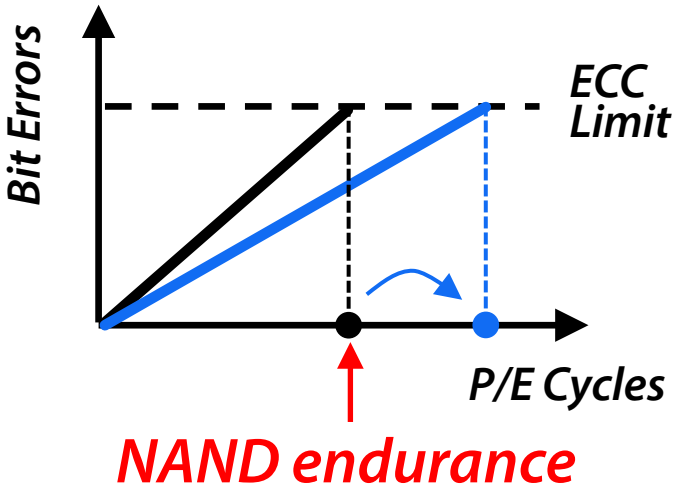

Erase voltage ↓

↓

Endurance ↑

Overview of Our Proposed Approach

DPES (Dynamic Program and Erase Scaling) approach

Erase voltage/time scaling

(Tradeoff : endurance and erase voltage/time)

Program time scaling

(Tradeoff : erase voltage and program time)

Dynamically changes program and erase voltage/time

Improves the NAND endurance

without degradation in the overall write throughput

Outline

- Introduction
- Motivation
- **Key Components of the DPES Approach**
  - **Erase Voltage Scaling**
  - **Program Time Scaling**
  - **Dynamic Program and Erase Scaling**
- **Implementation of DPES-Aware FTL**
- **Experimental Results**
- **Conclusion**

Erase Voltage Scaling

"Effective wearing represents the effective degree of NAND wearing after one P/E cycle."

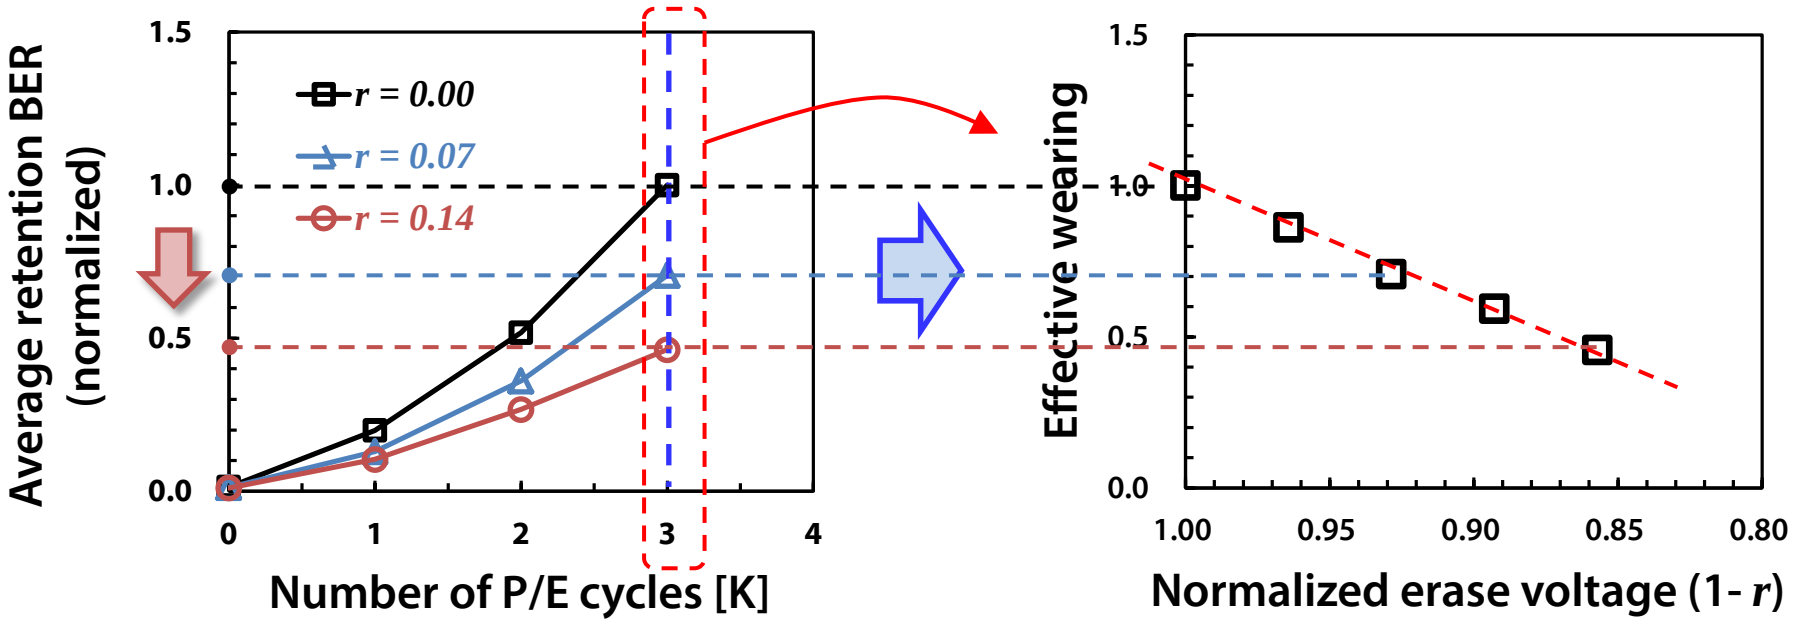

Lowering the erase voltage can reduce the effective wearing.

Effect of Erase Voltage Scaling

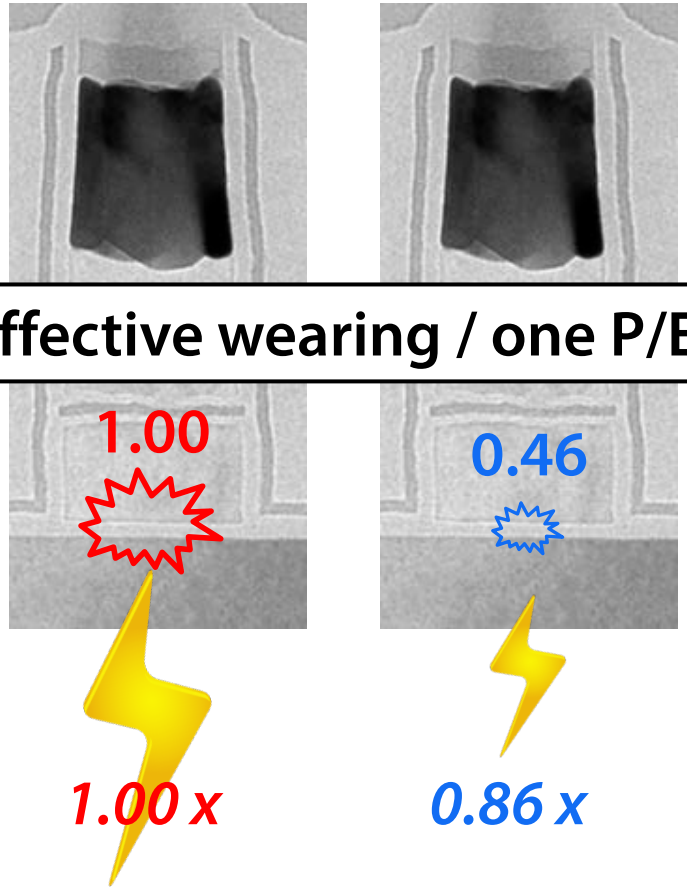

	Total sum of effective wearing	Endurance
3K P/E	3.00 K	3.00 K
	1.38 K	6.52 K

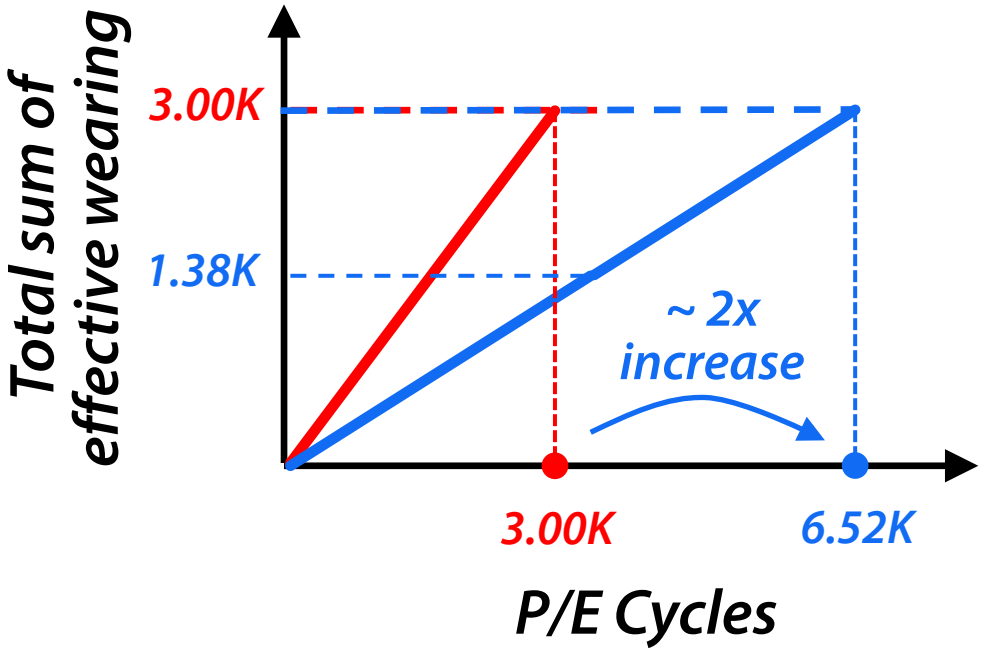

Writing Data to a Shallowly Erased Block

Erase voltage ∝ **Width of Vth distributions**

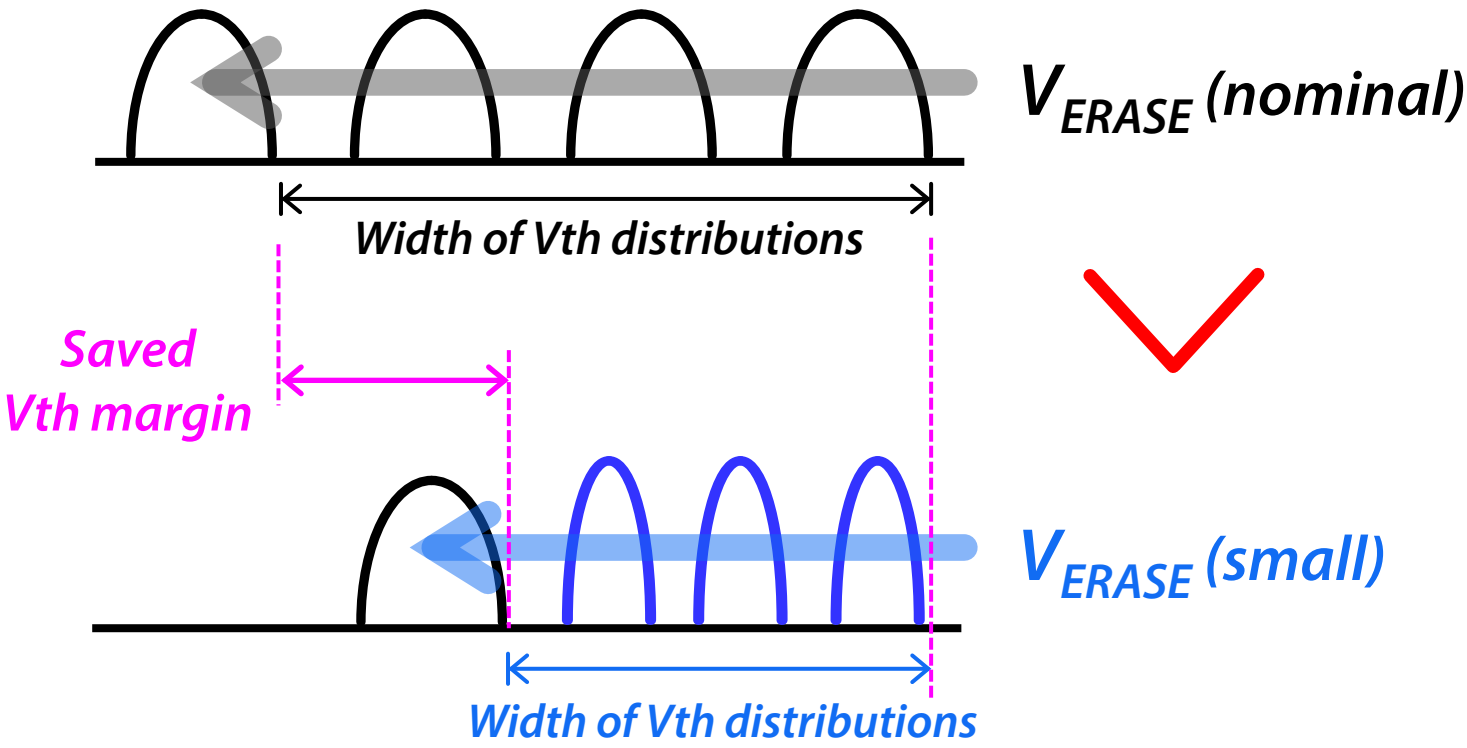

To write data to a shallowly erased NAND block,
it is necessary to shorten the width of Vth distributions.

Program Time Scaling

Tradeoff : {Program time} vs. {Width of Vth distributions}

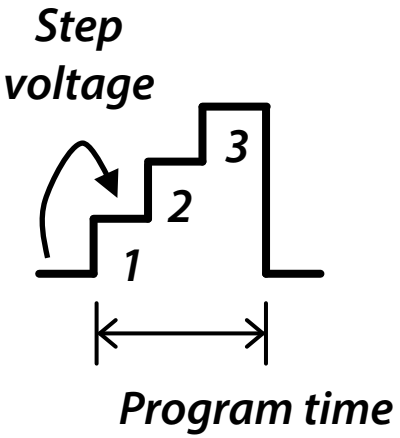

Incremental Step Pulse Program Scheme

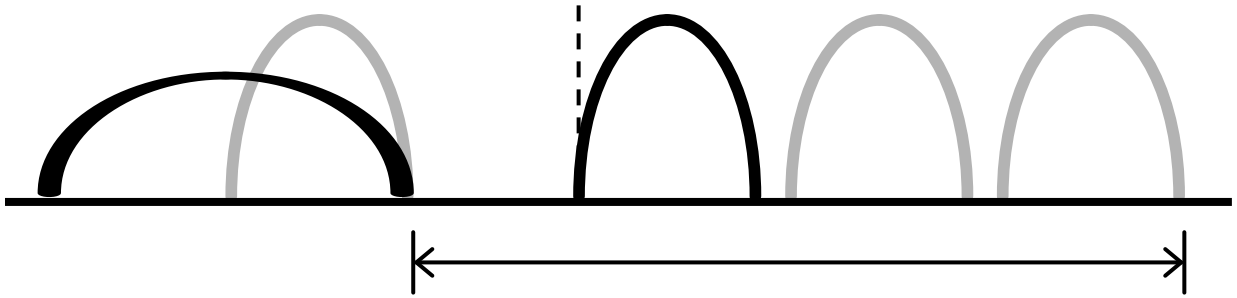

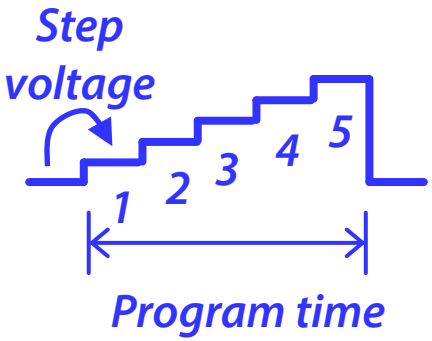

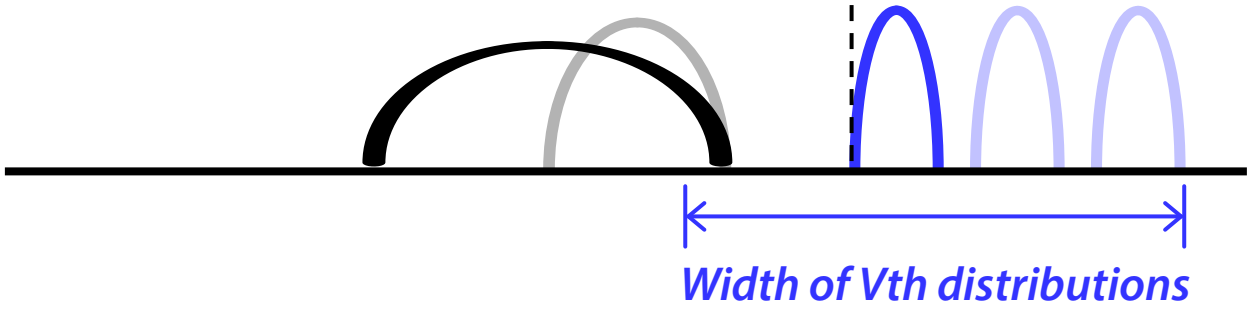

To shorten the width of Vth distributions, the program time is increased.

Minimum Program Time Requirement

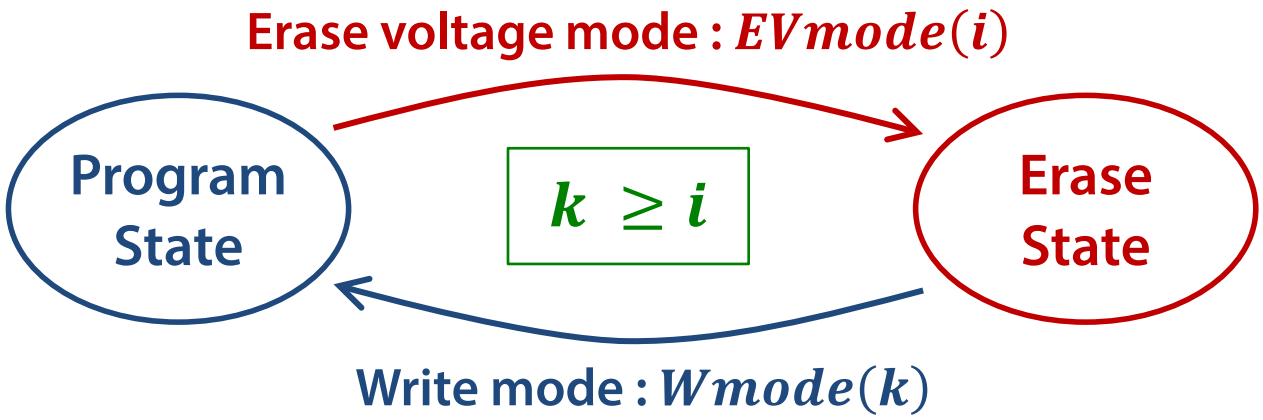

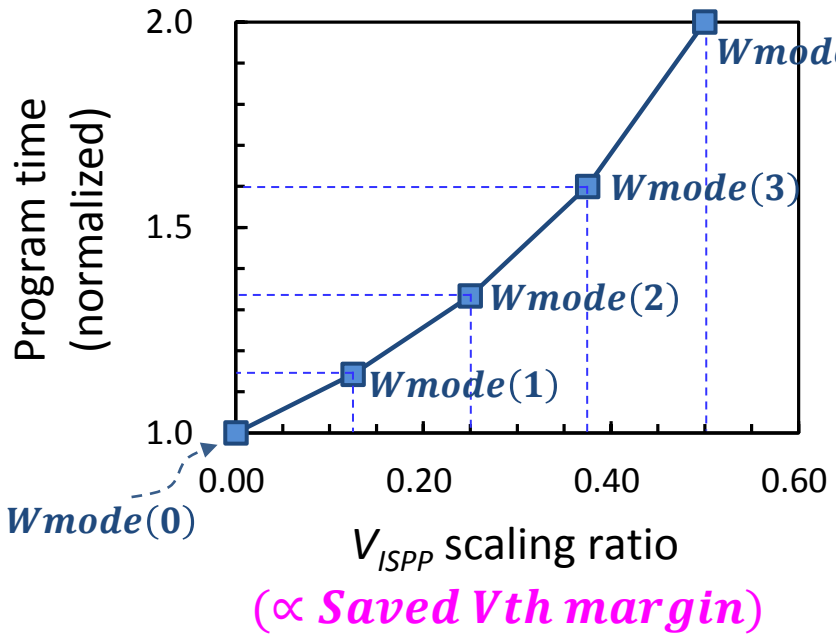

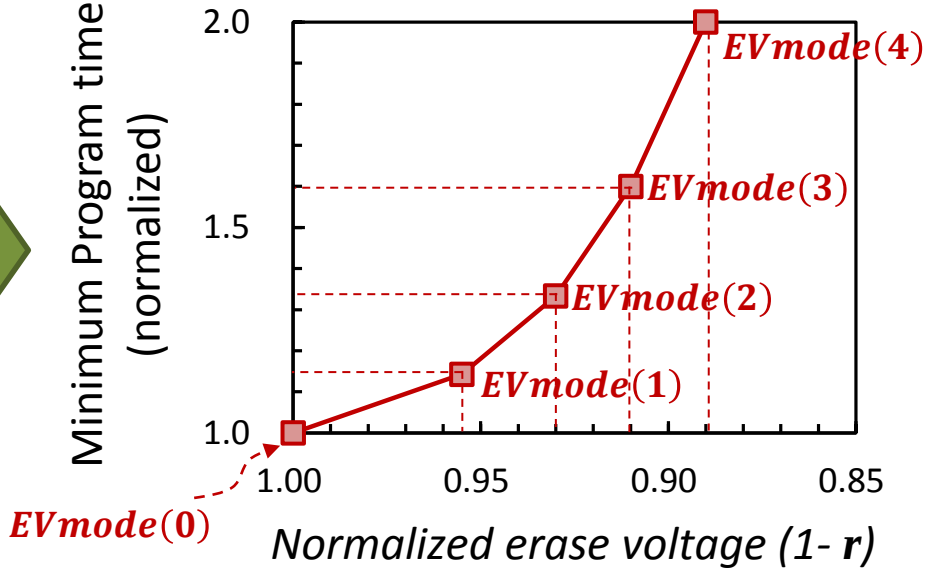

Example of EVmode and Wmode Selection

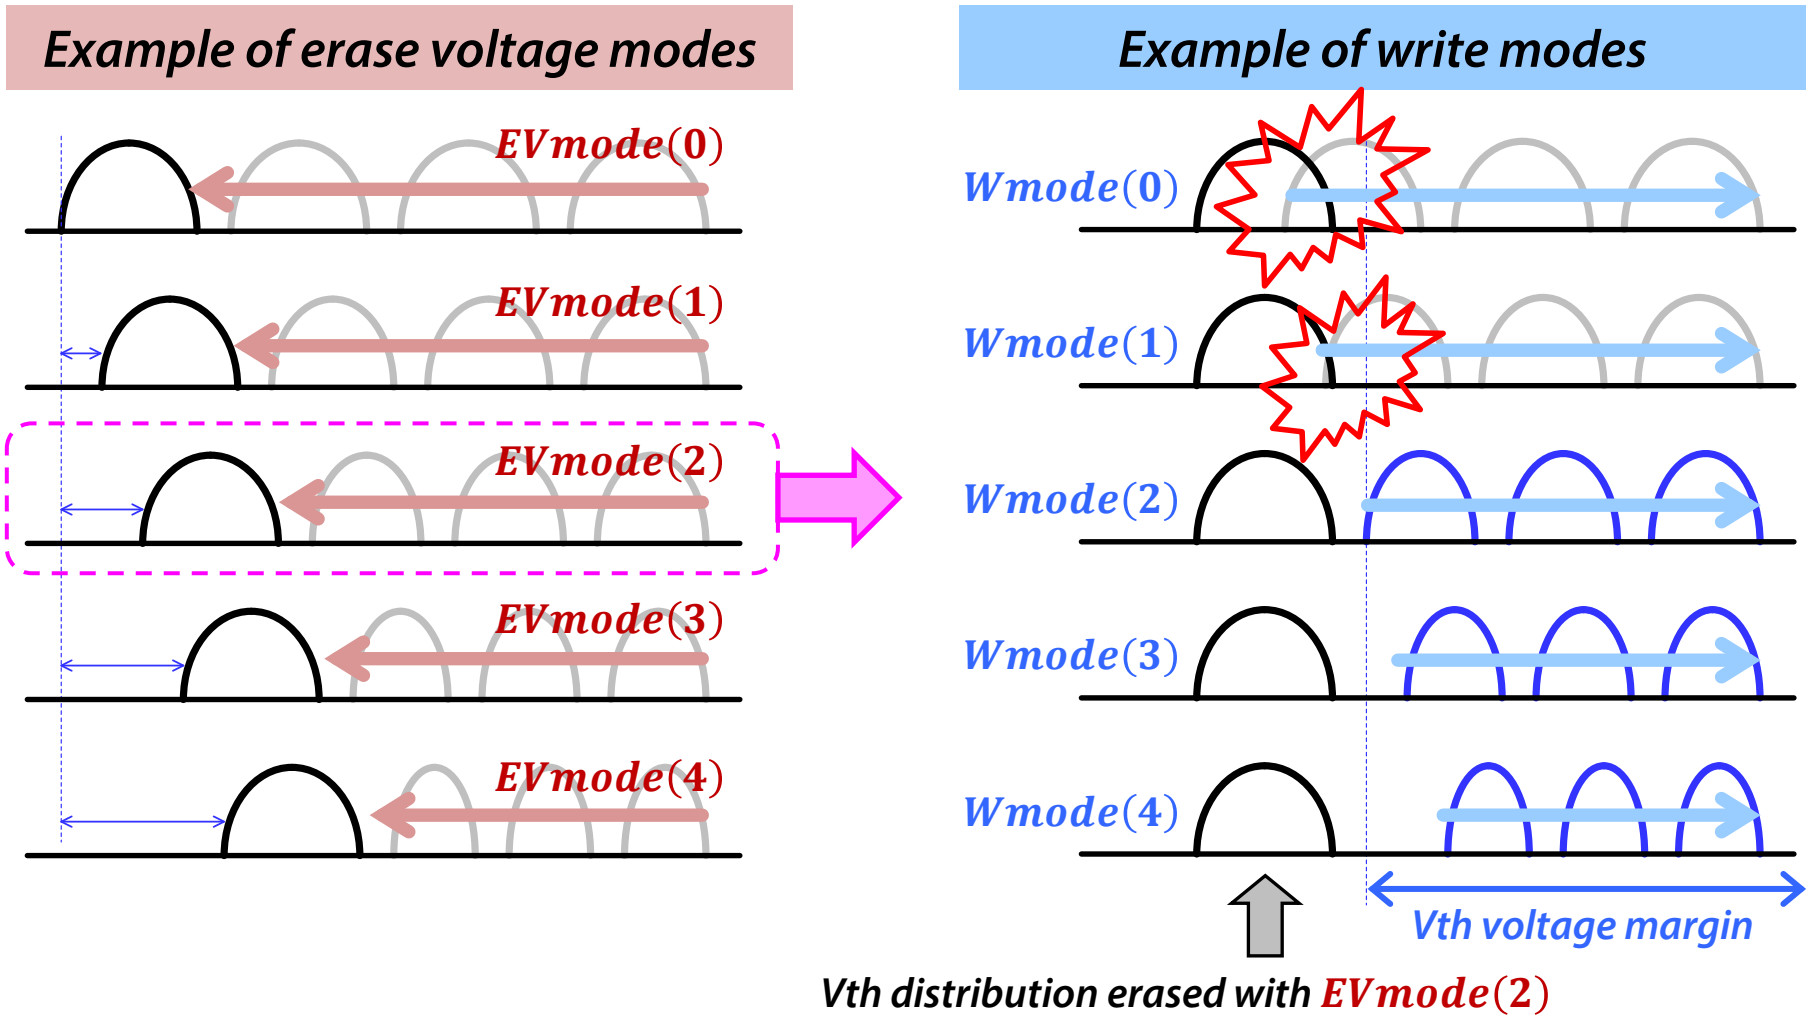

For writing a block erased with $EVmode(2)$*,*

$Wmode(2)$, $Wmode(3)$, $Wmode(4)$ *should be used.*

Lazy Erase Scheme

Program state

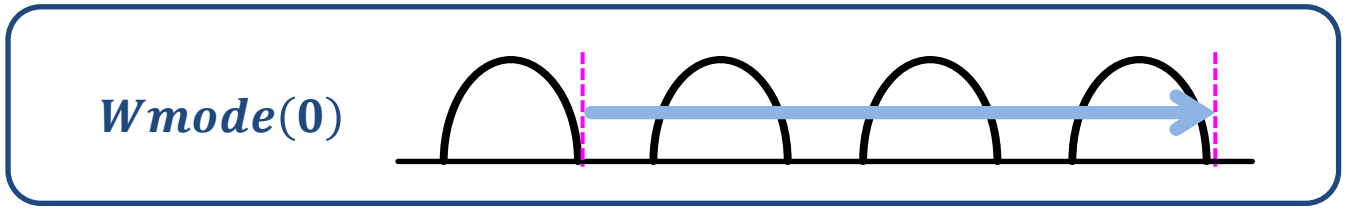

Erasing with $EVmode(2)$.

Erase state

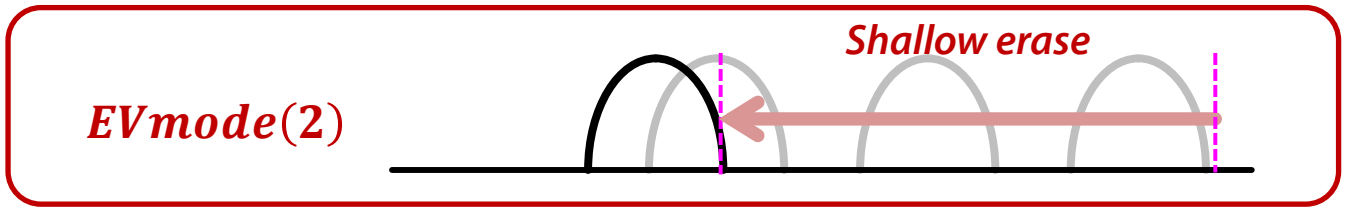

To write data with a faster write mode (e.g., $Wmode(0)$ *) to the shallowly erased block,*

Erase state

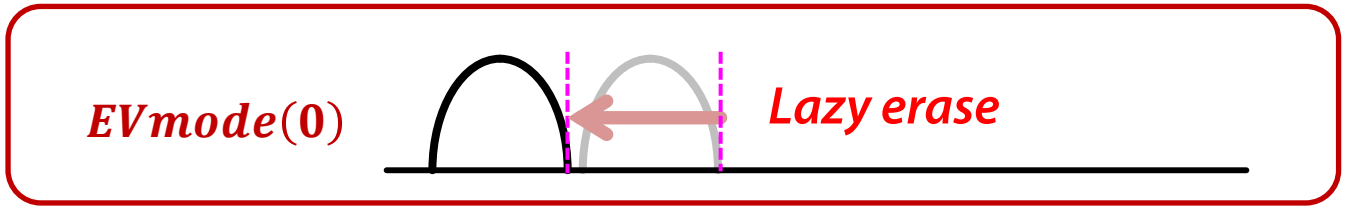

Writing with $Wmode(0)$

Program state

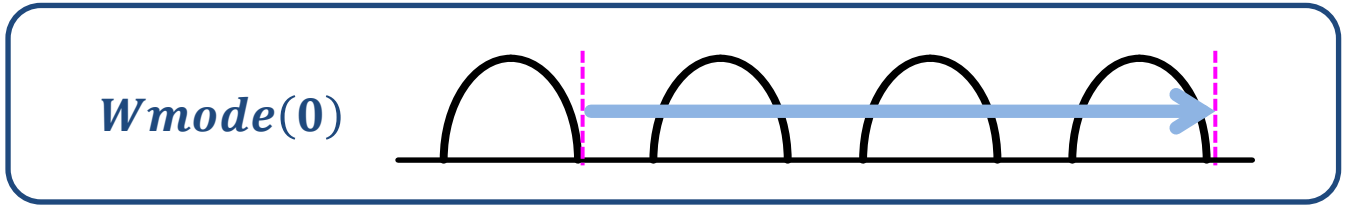

Dynamic Program and Erase Scaling

Program times and erase voltages are dynamically changed for improving the NAND endurance .

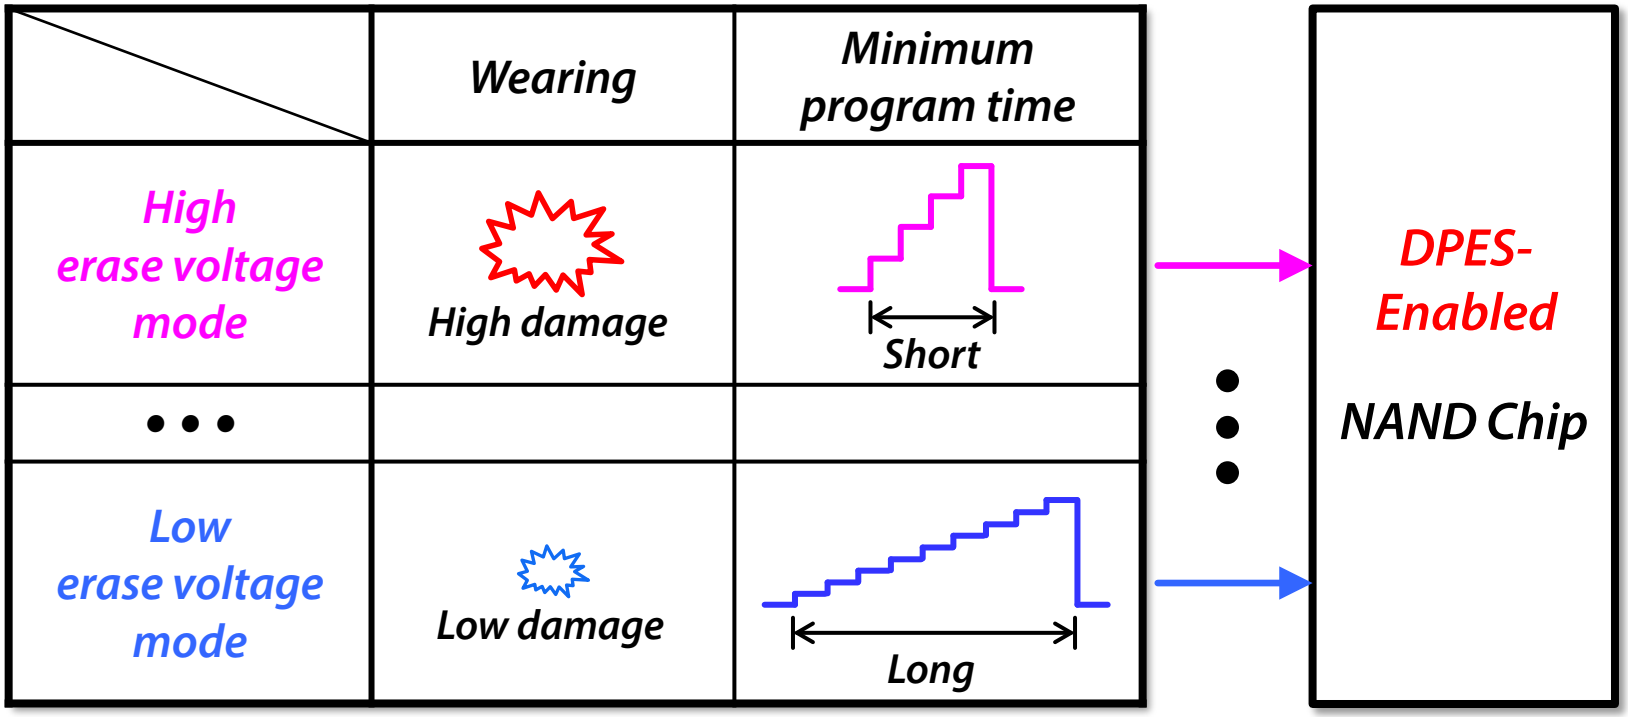

DPES enables S/W to exploit the tradeoff relationship between the NAND endurance and the erase voltage/time.

Outline

- Introduction
- Motivation
- Key Components of the DPES Approach
  - Erase Voltage Scaling
  - Program Time Scaling
  - Dynamic Program and Erase Scaling
- **Implementation of DPES-Aware FTL**
- **Experimental Results**
- **Conclusion**

Overview of DPES-Aware FTL (AutoFTL)

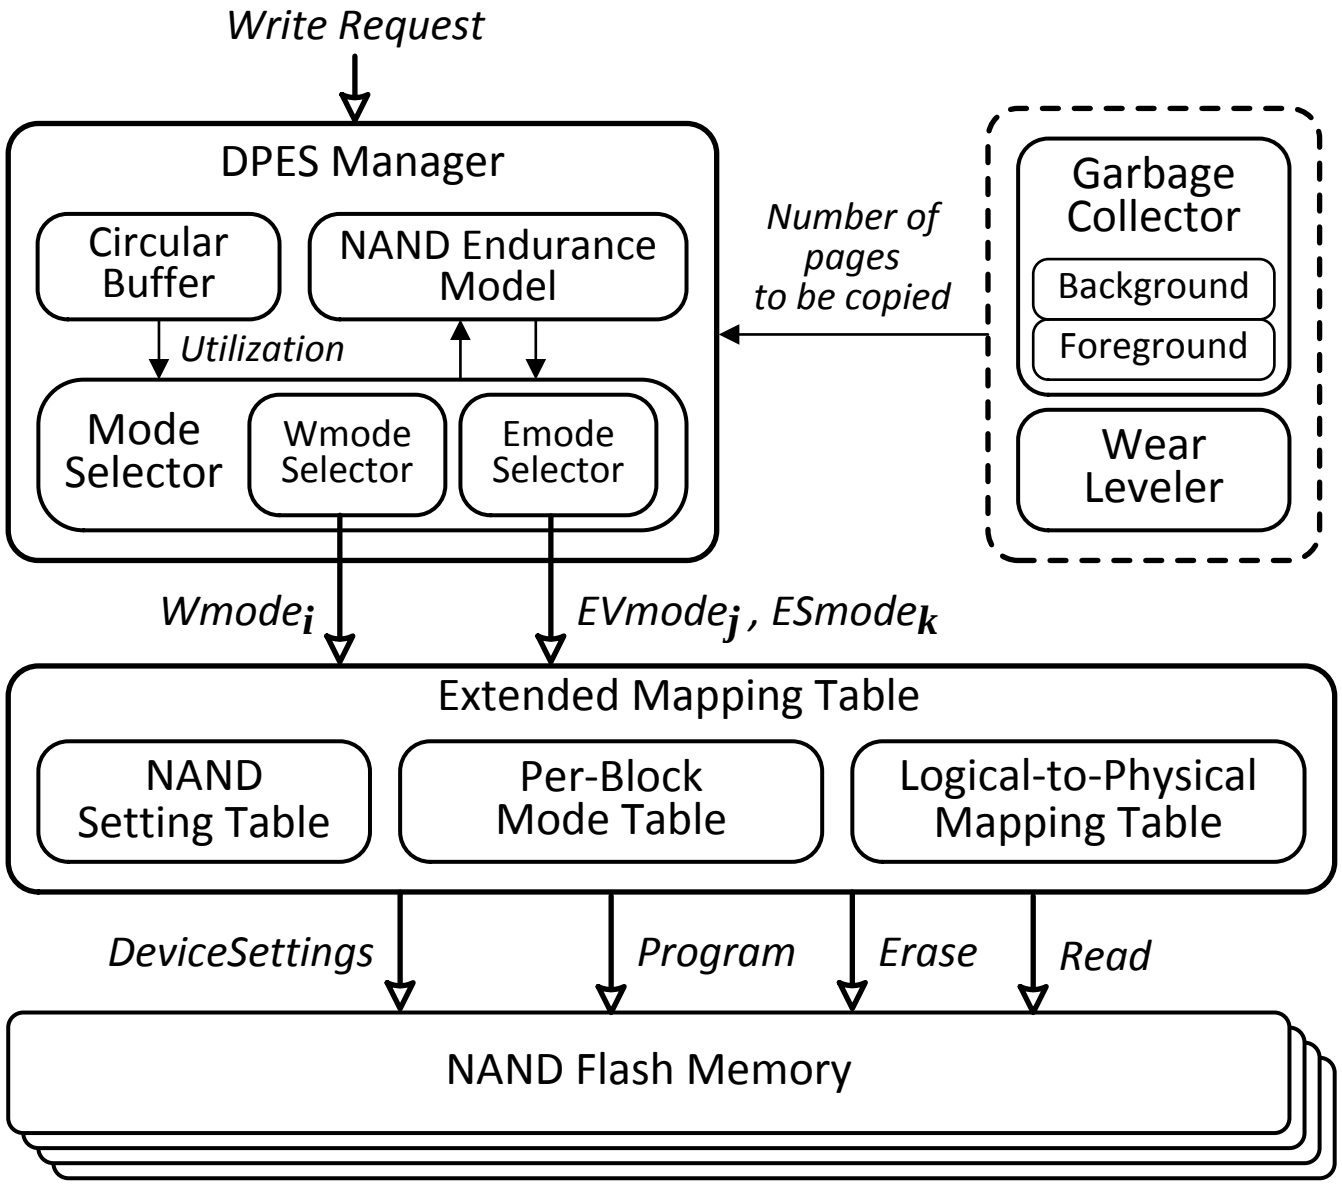

AutoFTL : Write Mode Selection

The DPES manager chooses the most appropriate write mode depending on the buffer utilization ratio.

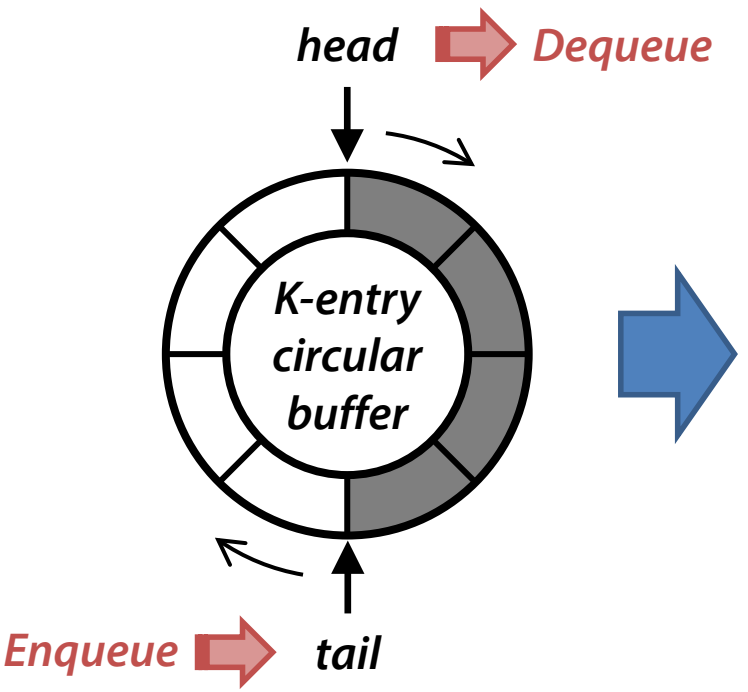

[Write-mode selection rules]

Buffer utilization (*u*)	Write mode
$u \le 20\%$	4
$20\% < u \le 40\%$	3
$40\% < u \le 60\%$	2
$60\% < u \le 80\%$	1
$u > 80\%$	0

Program time

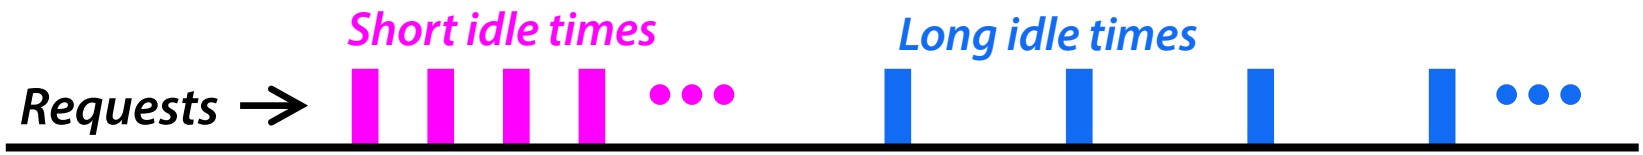

DPES-Aware Write and Read Operations

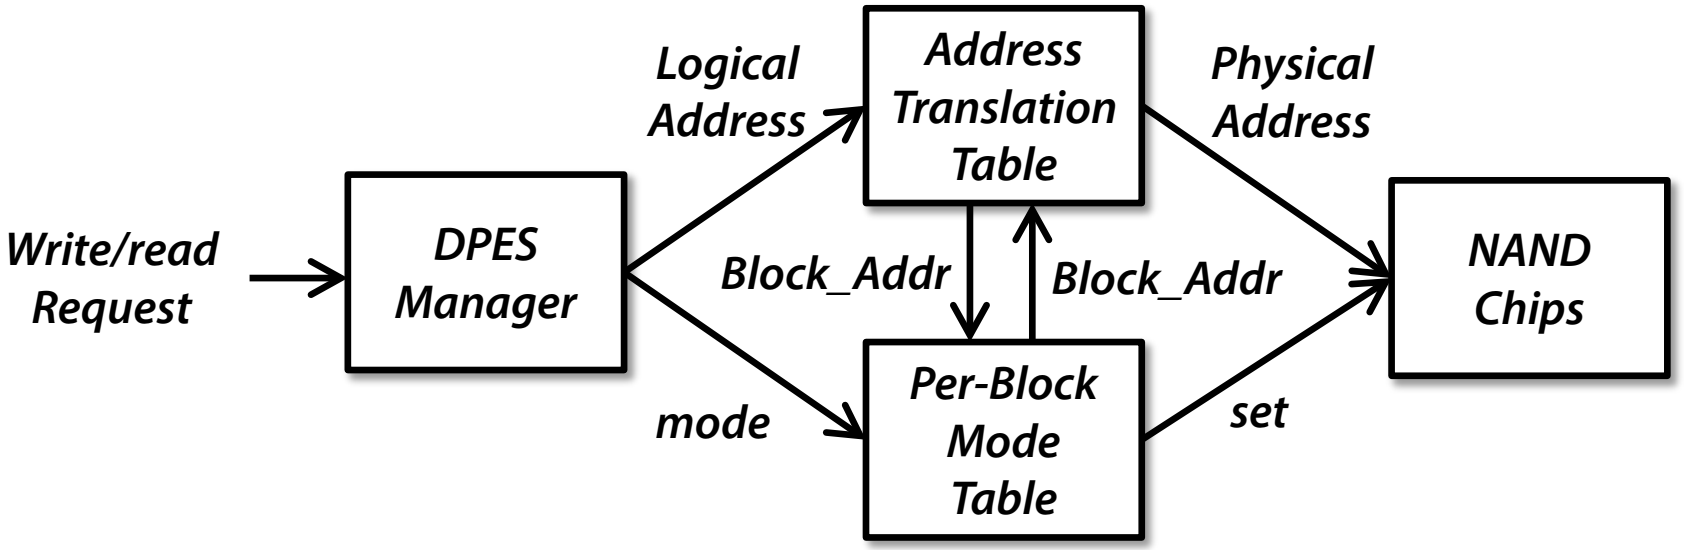

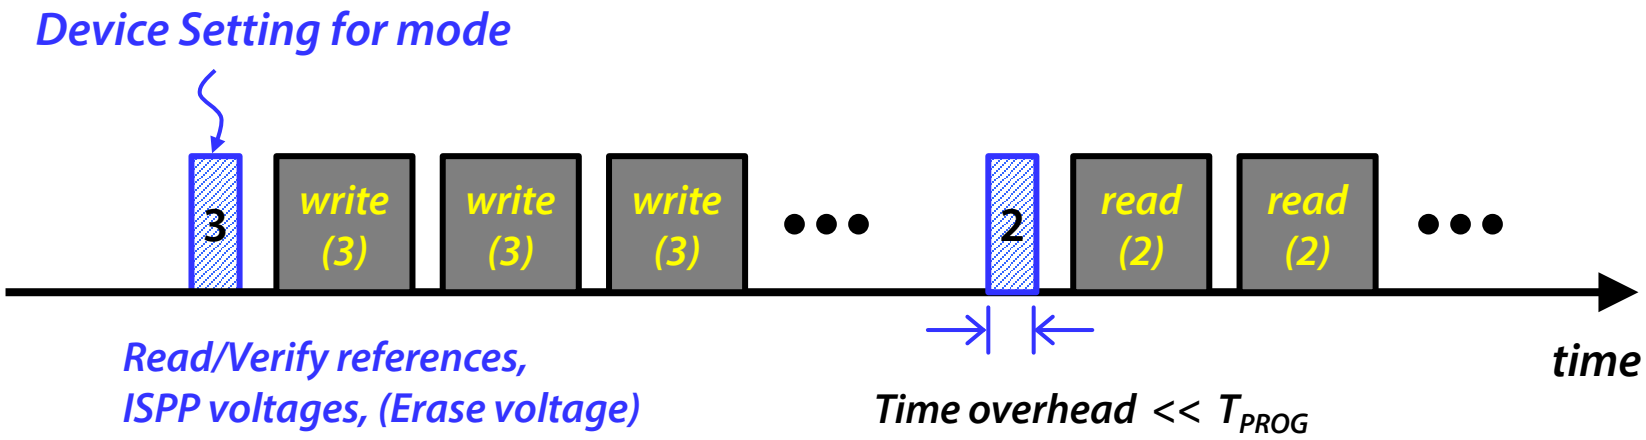

AutoFTL : Erase Voltage Mode Selection

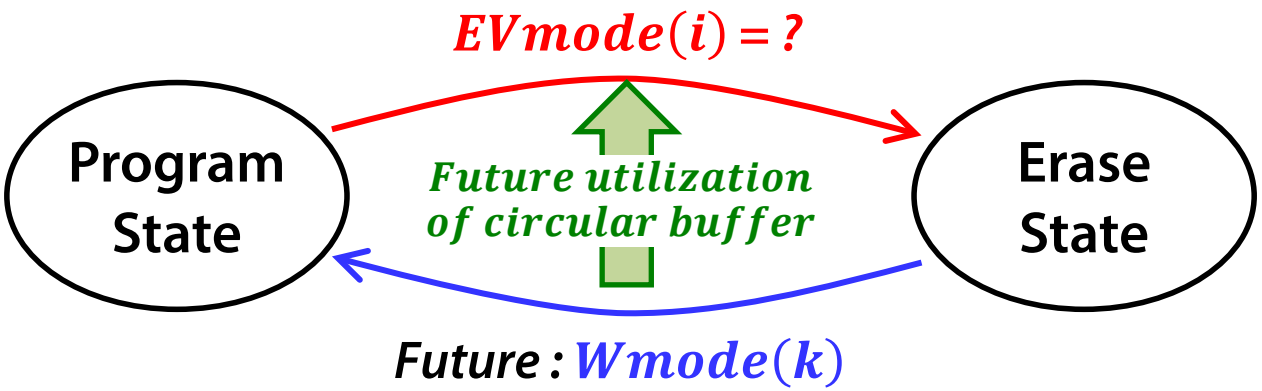

	If we know $Wmode(k)$ before a block is erased	If we don't know $Wmode(k)$ before a block is erased
Cases	Foreground garbage collection	Background garbage collection, wear leveling
$EVmode(i)$	$i = k$	Prediction based on the past utilization history
		Incorrect prediction → Lazy erase

AutoFTL : DPES-Aware Garbage Collection

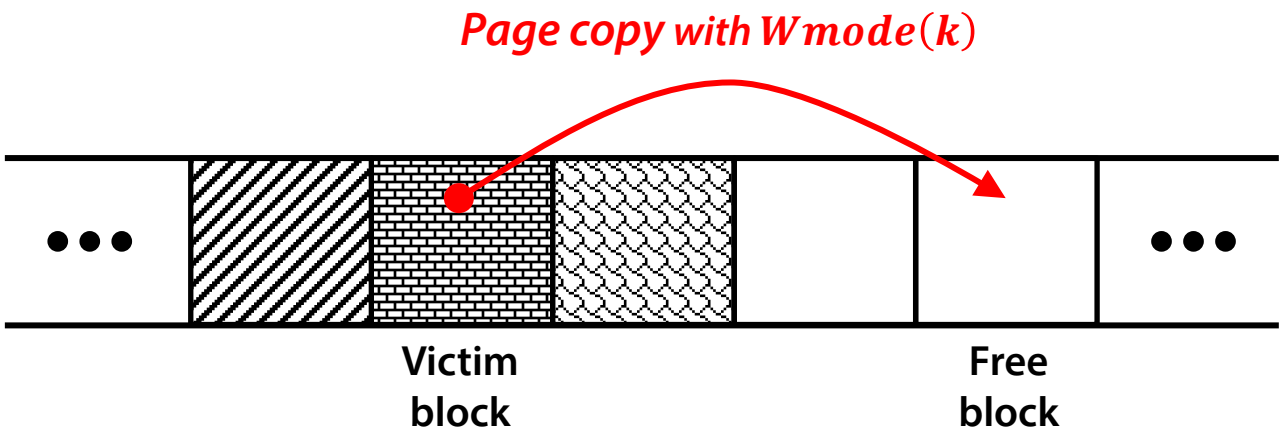

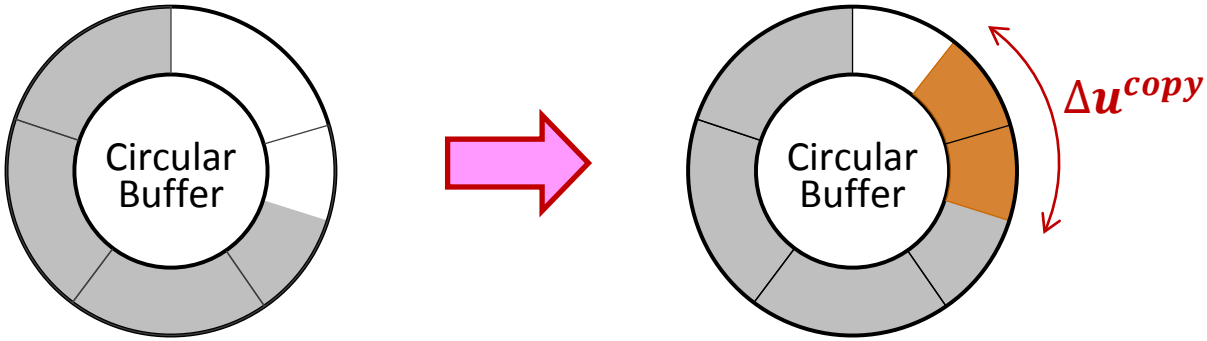

Current utilization, u

Effective utilization,

$$u^e = u + \Delta u^{copy}$$

Outline

- Introduction
- Motivation
- Key Components of the DPES Approach
  - Erase Voltage Scaling
  - Program Time Scaling
  - Dynamic Program and Erase Scaling
- Implementation of DPES-Aware FTL
- **Experimental Results**
- **Conclusion**

Experimental Settings

		Configuration 1	Configuration 2
NAND flash chip		128 blocks/chip, 8 KB/page	
*Extended FlashBench** configuration	Chips/channel	2	8
	# of channels	1	4
	Size of circular buffer	80 KB	32 MB
	NAND timing model	*Timing accurate emulation model using hrtimers (variation < 1%)*	
I/O traces		*Mobile (2ea)*	*Server (6ea)*

* *S. Lee et al., "FlashBench: A Workbench for a Rapid Development of Flash-Based Storage Devices," IEEE Int. Symp. Rapid System Prototyping, 2012.*

Characteristics of I/O Traces

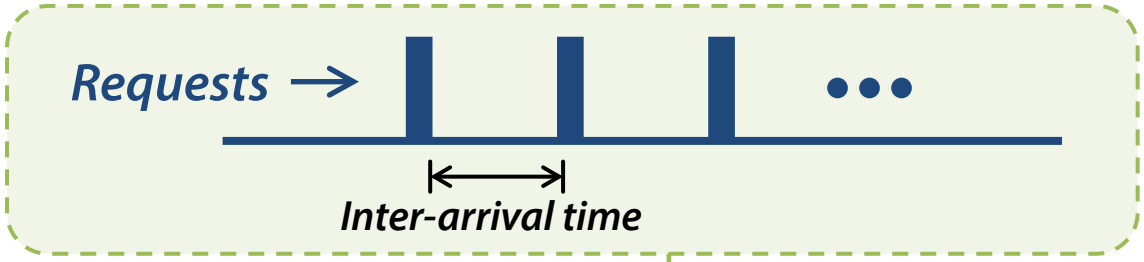

$$= \frac{T_{PROG}}{Total\ \#\ of\ chips}$$

$T_{PROG}^{effective}$

Distributions of normalized inter-arrival times (t) over

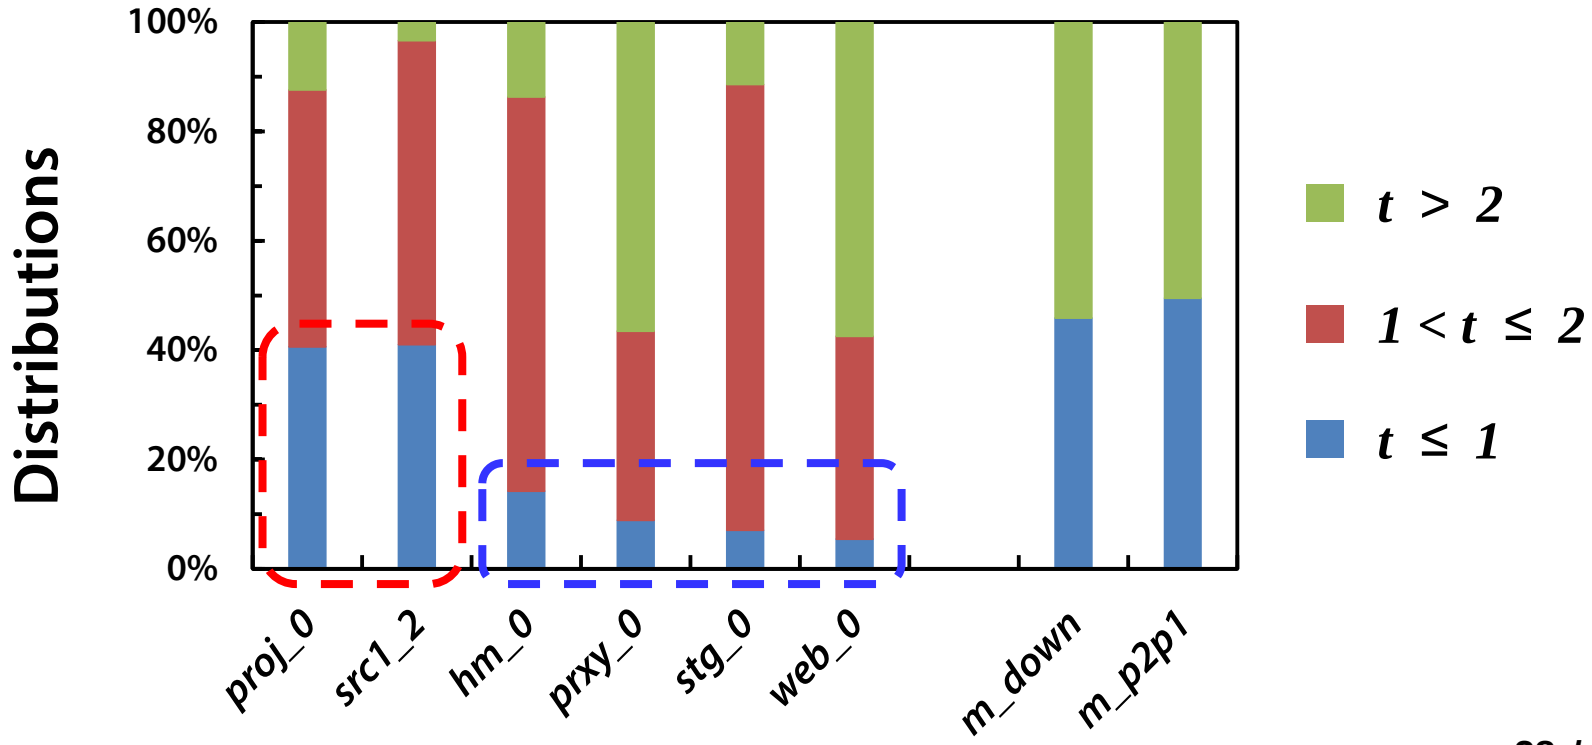

Result 1 : Normalized Endurance Gain

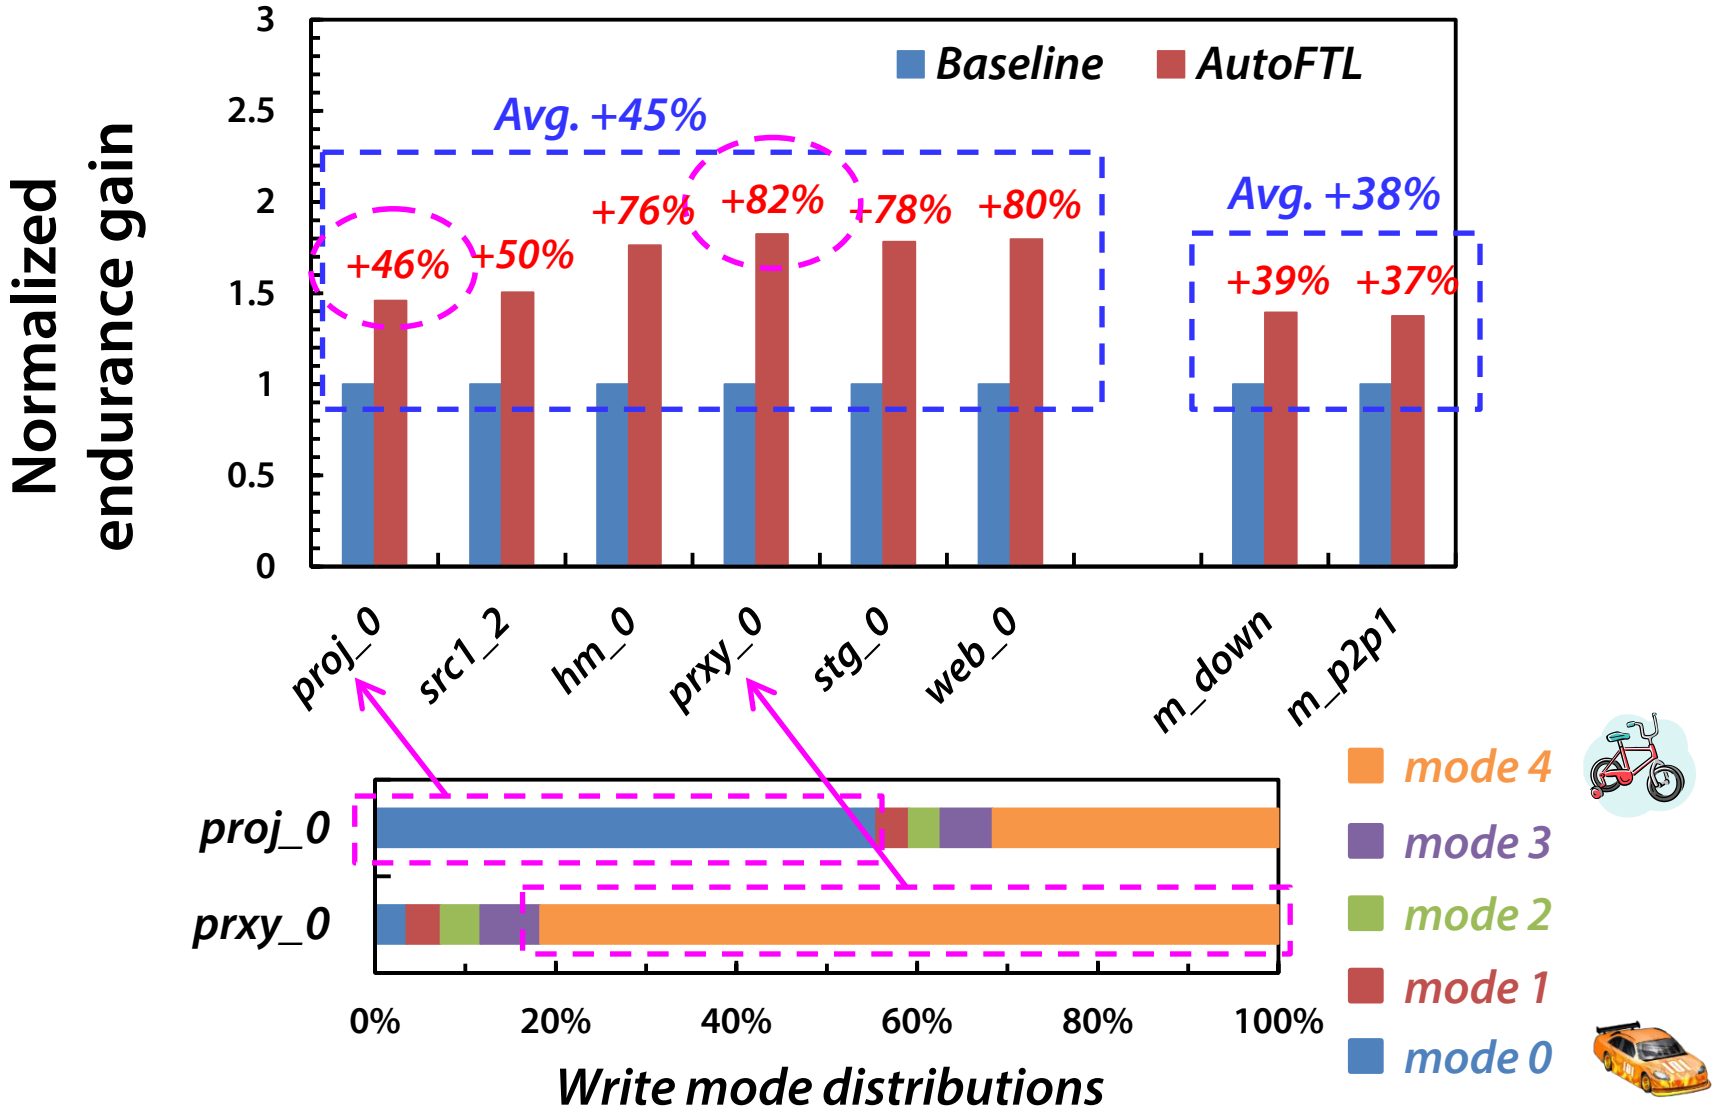

Result 2 : Overall Write Throughput

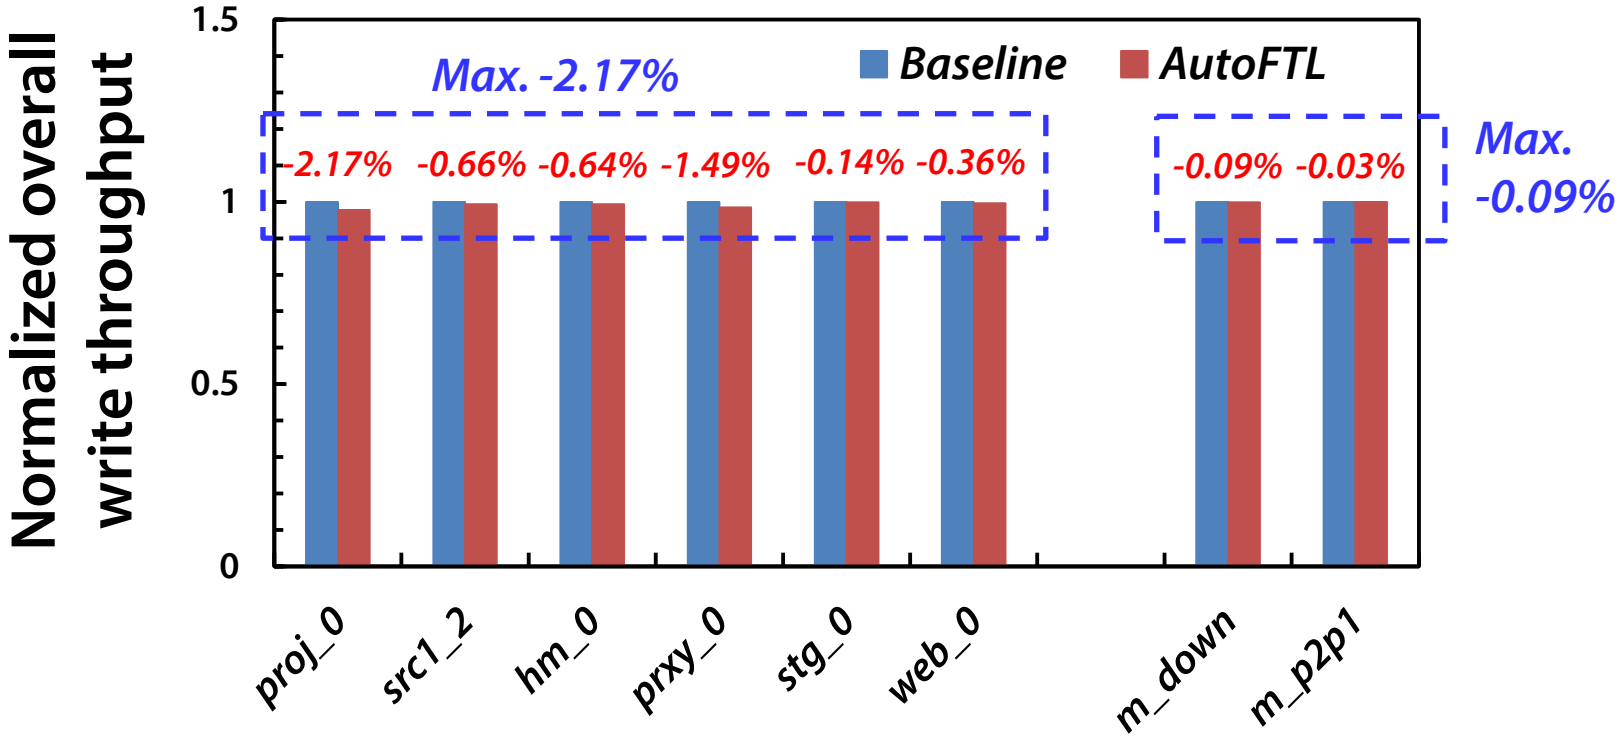

The decrease in the overall write throughput over baseline was less than 2.2%.

Conclusion

- ✓ *We have presented a system-level approach for improving the lifetime of flash-based storage systems using DPES.*
  - *Actively exploits the tradeoff relationship between the NAND endurance and the erase voltage*
  - *Automatically changes the erase voltage and the program time*
  - *Makes the key FTL modules DPES-aware*
  - *Improves the NAND endurance by 61.2% on average (with less than 2.2% decrease in the overall write throughput)*

- ✓ *Future Work*
  - *Develop adaptive mode selection rules for adequately reflecting the varying characteristics of I/O workload*