

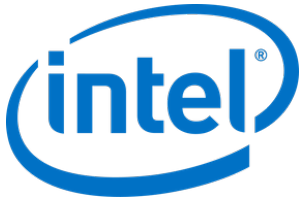

HPCA 2023
The 29th IEEE International Symposium on
High-Performance Computer Architecture (HPCA-29)

Realizing Extreme Endurance Through Fault-aware Wear Leveling and Improved Tolerance

Jiangwei Zhang∗, Chong Wang∗, Zhenhua Zhu∗, Donald Kline, Jr†, Alex K. Jones‡,
Huazhong Yang∗, and Yu Wang∗

∗ **Department of EE, Tsinghua University**; † Intel Corporation

‡ Department of ECE, University of Pittsburgh


Outline

- Background and Introduction
- RETROFIT
  - Fault-aware Wear Leveling
  - GPS to Enhance RETROFIT
  - Column-level Wear Leveling
  - Put it all together
- Results
- Conclusion

Outline

- **Background and Introduction**
- RETROFIT
  - Fault-aware Wear Leveling
  - GPS to Enhance RETROFIT
  - Column-level Wear Leveling
  - Put it all together
- Results
- Conclusion

Background and Introduction

Memory Wall
Bottlenecks in Speed and Power Consumption

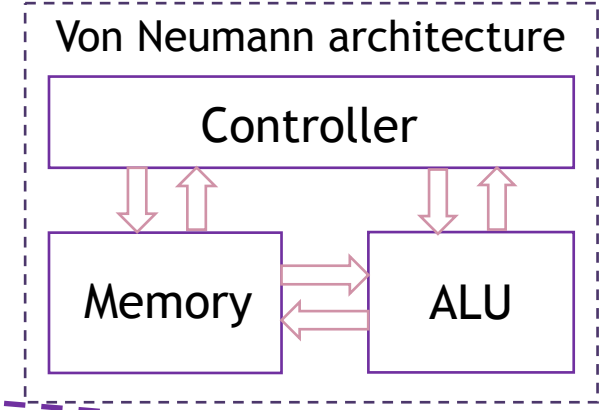

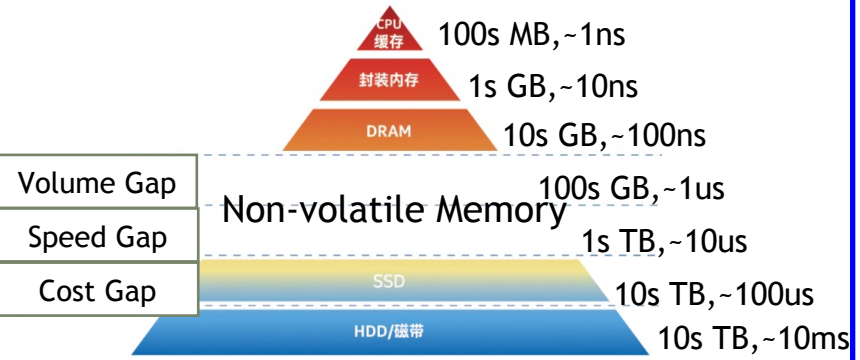

New Memory Architecture

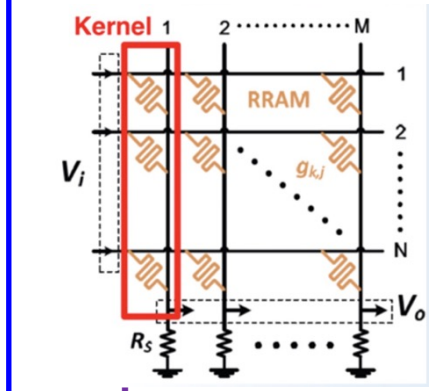

In-memory Computing

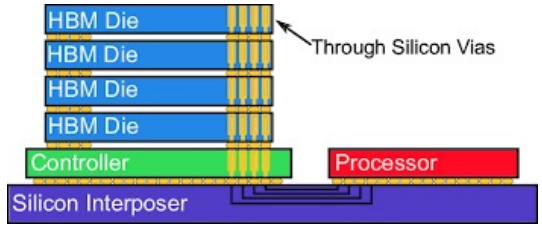

Near-memory Computing

Background and Introduction

- Non-volatile memories (such as PCM and RRAM) are highly integrated, scalable, can be accessed by bit, and could replace DRAM. Still, there are reliability problems, such as device failure and limited endurance.

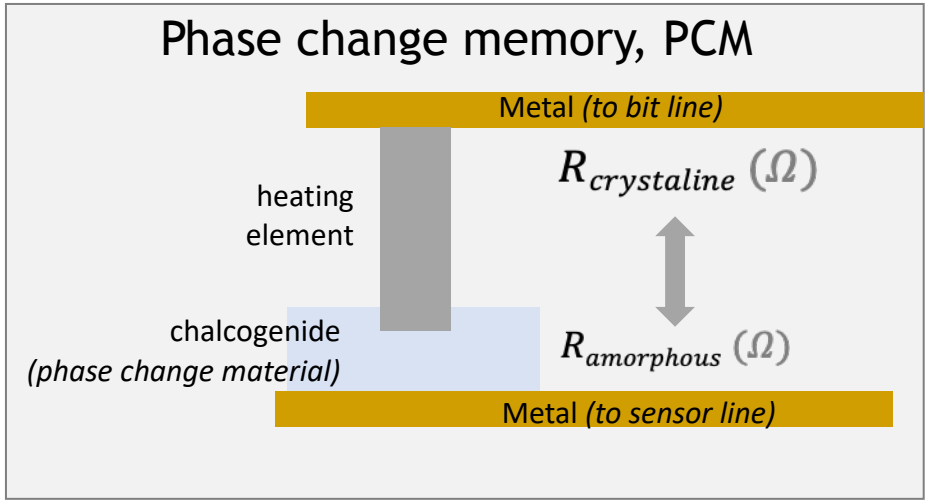

	DRAM	PCM	RRAM	NAND Flash
Nonvolatility	x	√	√	√
Area(F^2)	6	4	4	4
Write Latency	~1s ns	~10s ns	~10s ns	0.1~1s ms
Endurance	>10^{15}	10^8	10^6~10^9	10^5
3D stacking	x	√	√	√

Background and Introduction

Non-volatile memory has limited write endurance and is prone to early hard faults (Stuck-at faults):

- Stuck at '0' or '1'
- Still readable

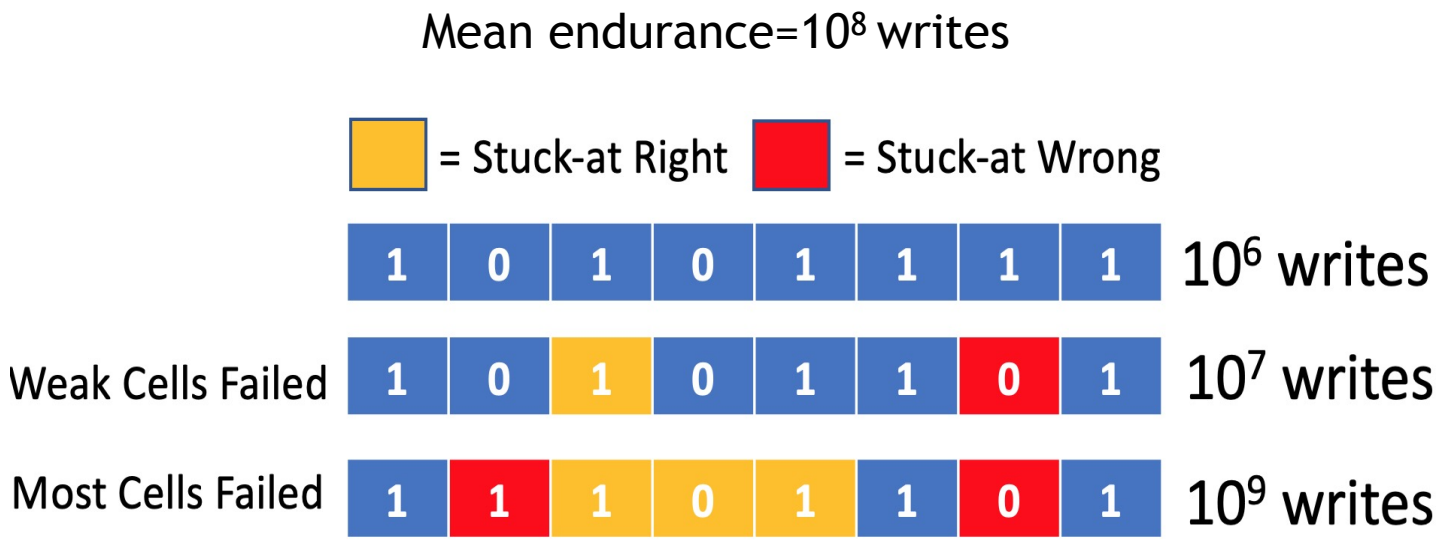

Background and Introduction

- Uneven memory access greatly reduces memory lifetime.
- Wear leveling technology improves memory lifetime by balancing the frequencies of writes between memory cells.

The function of Wear Leveling

https://www.transcend-info.com/Embedded/Essay-22

Background and Introduction

- Even with uniform write distribution, **process variation** from scaling can significantly reduce memory lifetime.

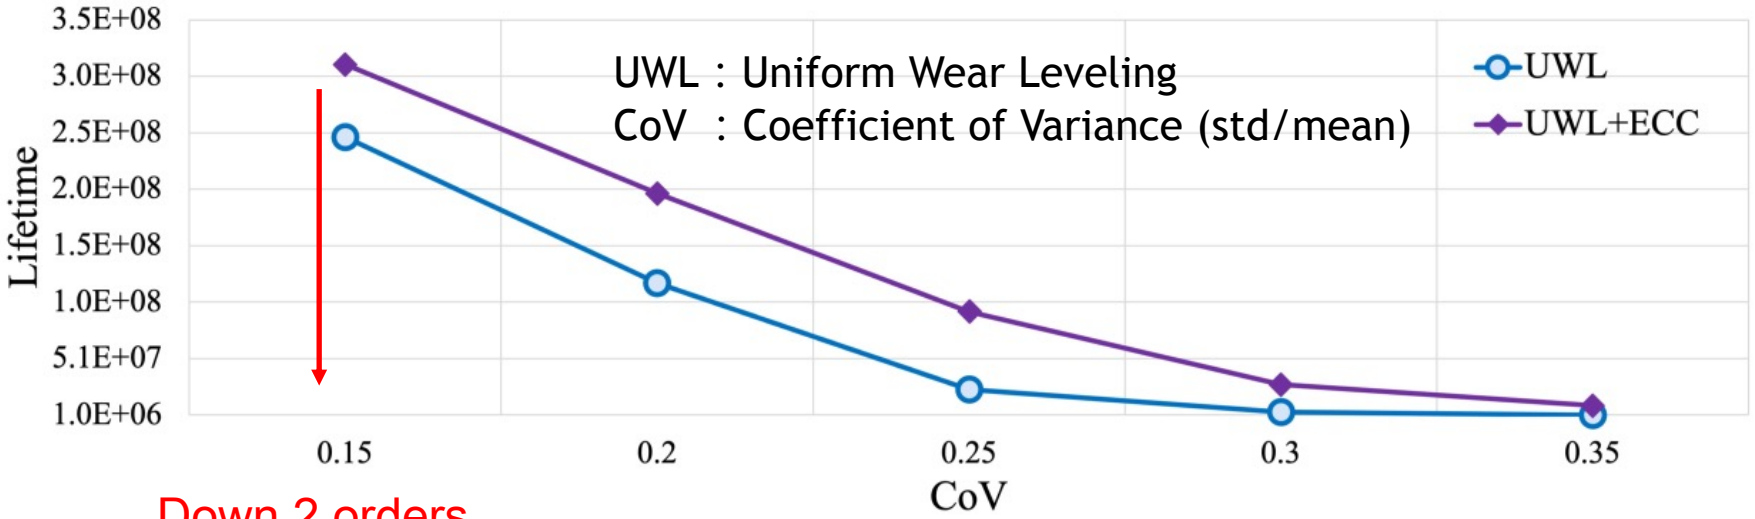

Impact of process variation on memory lifetime

Background and Introduction

- As the number of faults increases, the remaining lifetime is monotonically decreasing. Therefore, the number of faults can reflect the relative strength of the row.

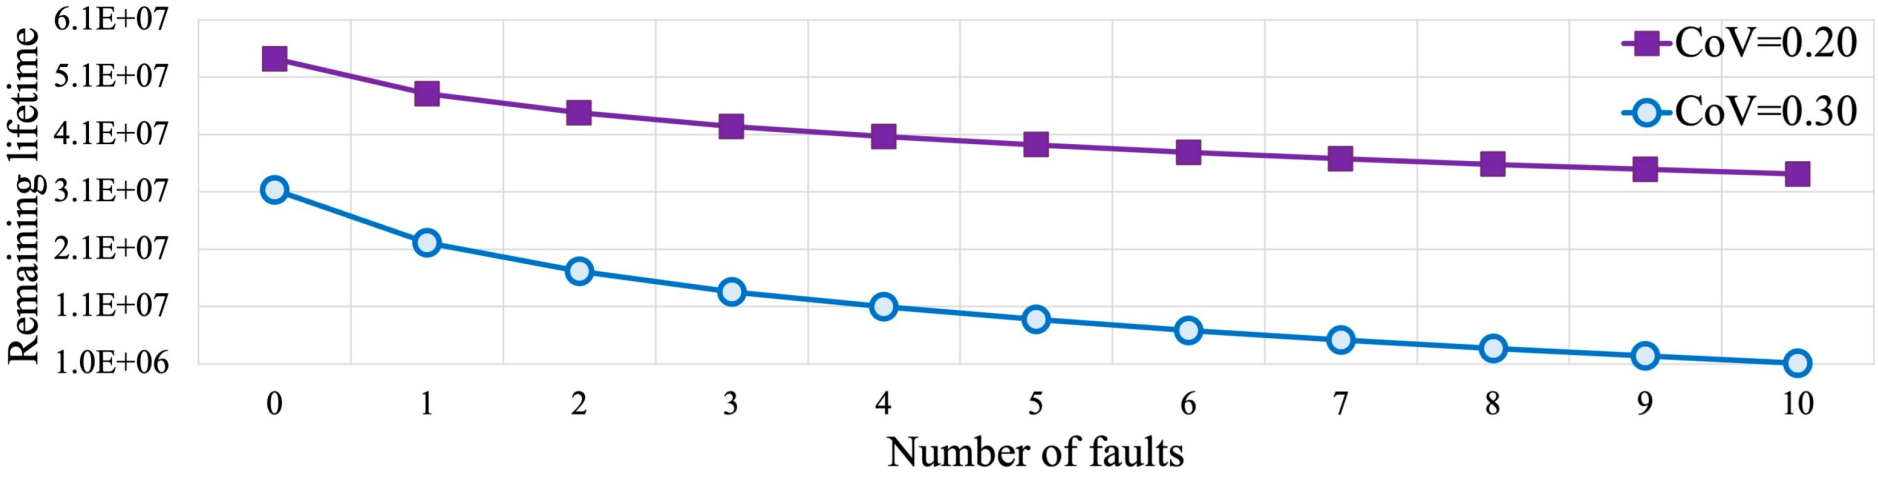

The relationship between the remaining lifetime (or write cycles) and the number of faults in a 512-bit row.

Motivation

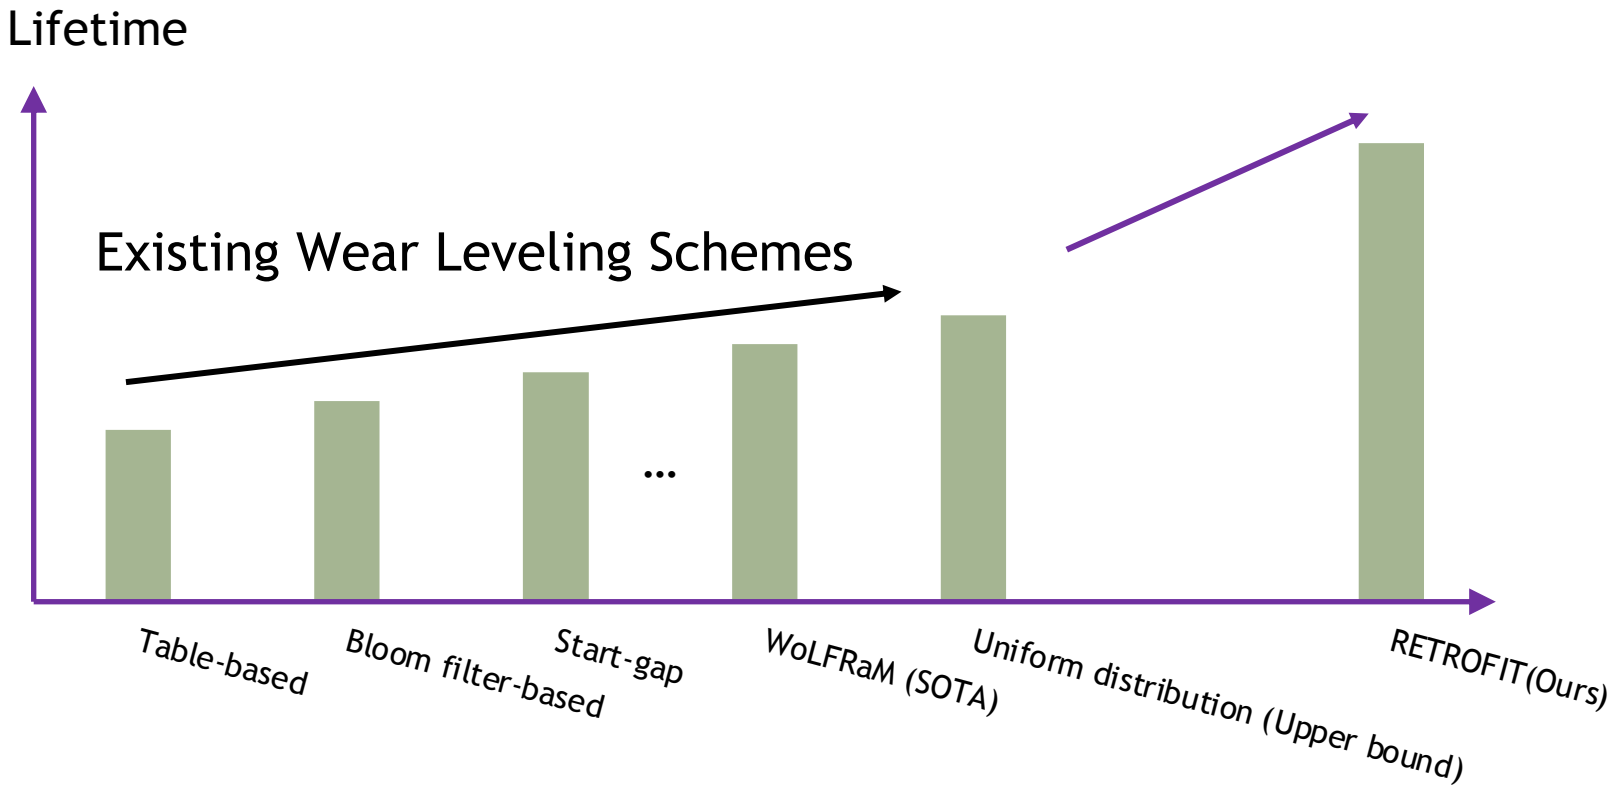
Lifetime
Existing Wear Leveling Schemes
...
Table-based
Bloom filter-based
Start-gap
WoLFRaM (SOTA)
Uniform distribution (Upper bound)
RETROFIT(Ours)

Outline

- Background and Introduction
- **RETROFIT**
  - **Fault-aware Wear Leveling**
  - **GPS to Enhance RETROFIT**
  - **Column-level Wear Leveling**
  - **Put it all together**
- Results
- Conclusion

RETROFIT: Fault-aware Wear Leveling

Fault-free case

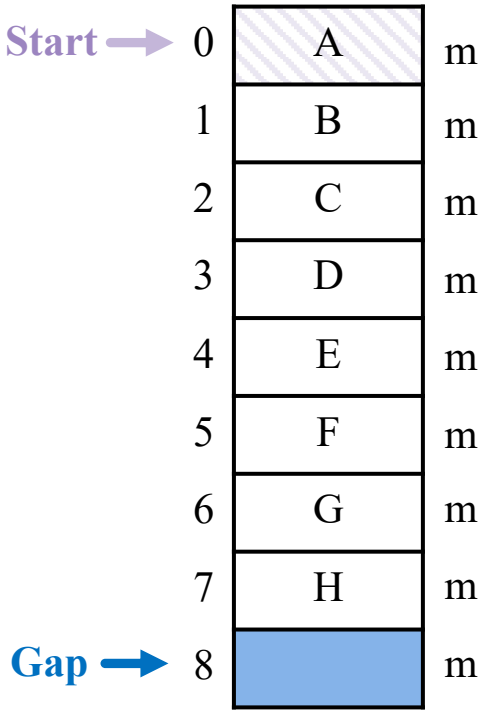

Start Gap (single gap)

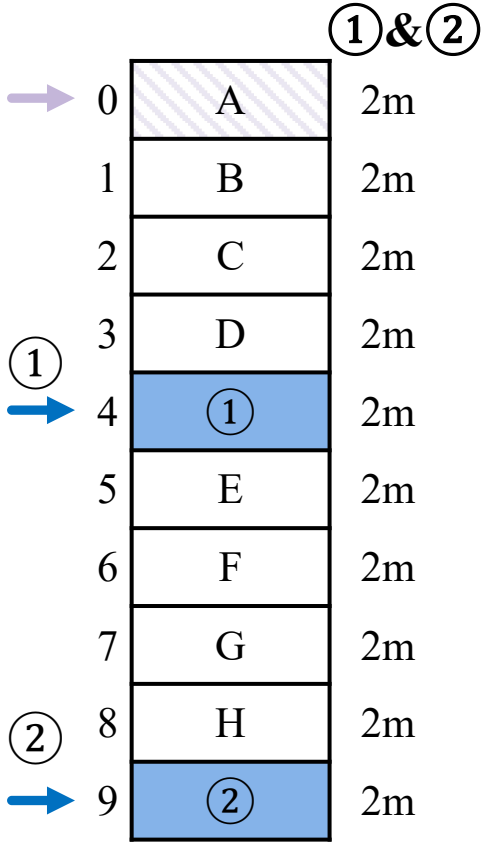

RETROFIT(two gaps)

M. K. Qureshi et al, MICRO 2009

RETROFIT: Fault-aware Wear Leveling

Fault-free case

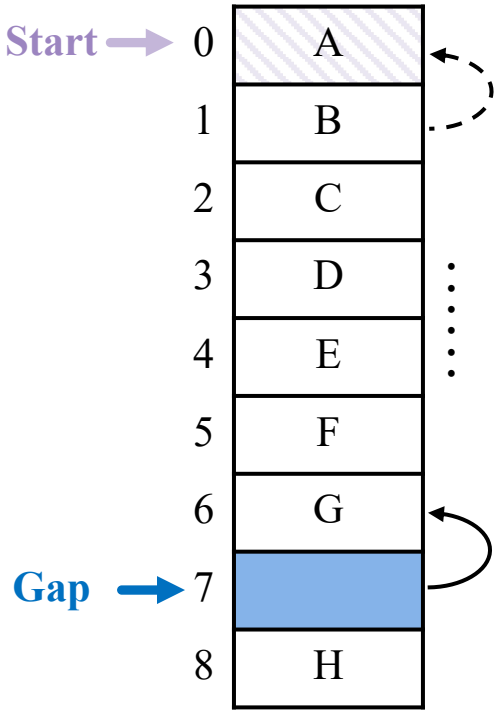

Start Gap (single gap)

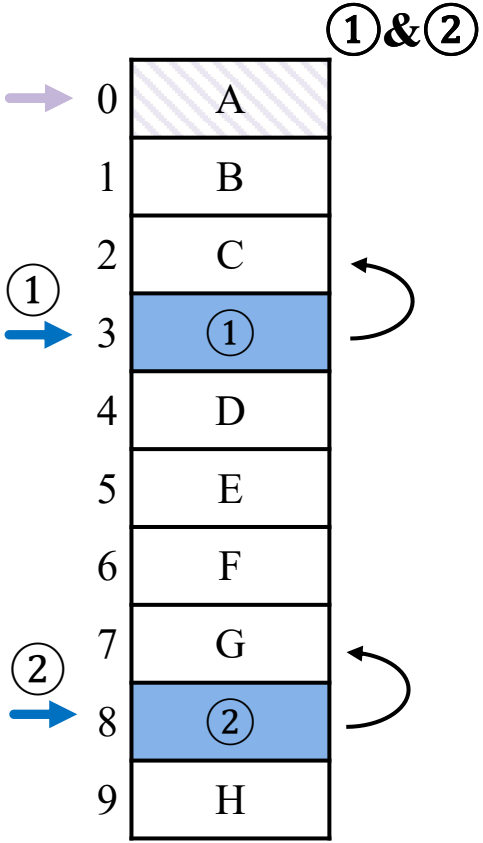

RETROFIT(multiple gaps)

M. K. Qureshi et al, MICRO 2009

RETROFIT: Fault-aware Wear Leveling

Fault-free case

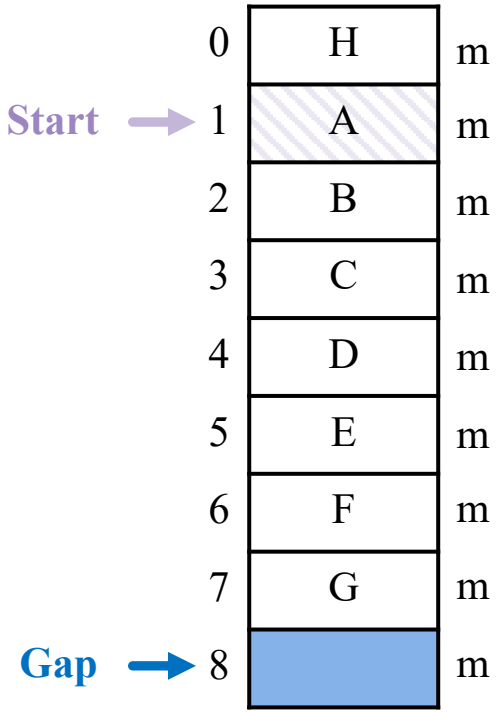

Start Gap (single gap)

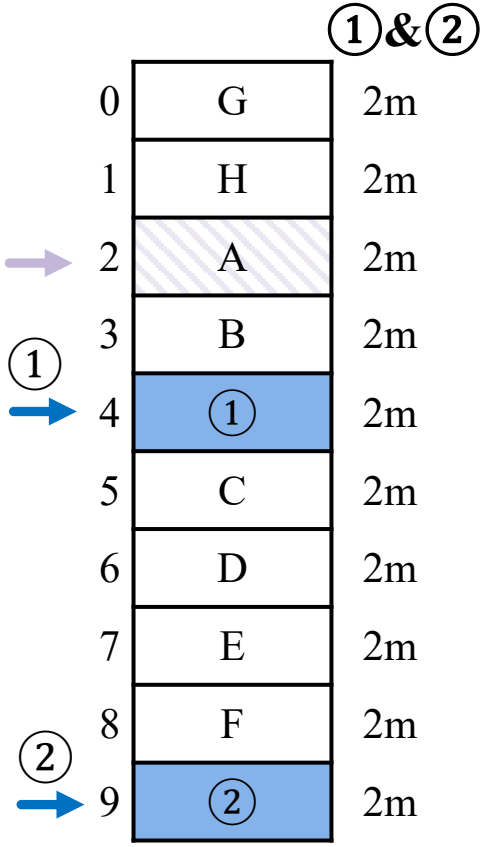

RETROFIT(multiple gaps)

M. K. Qureshi et al, MICRO 2009

RETROFIT: Fault-aware Wear Leveling

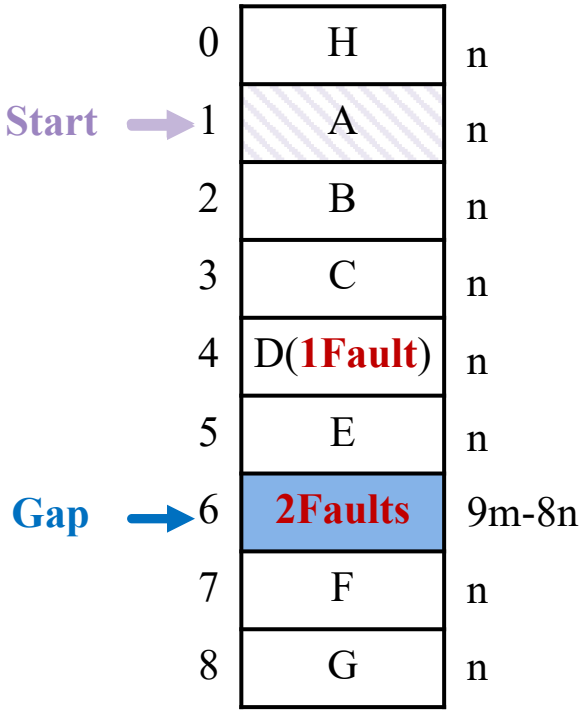

RETROFIT(single gap)

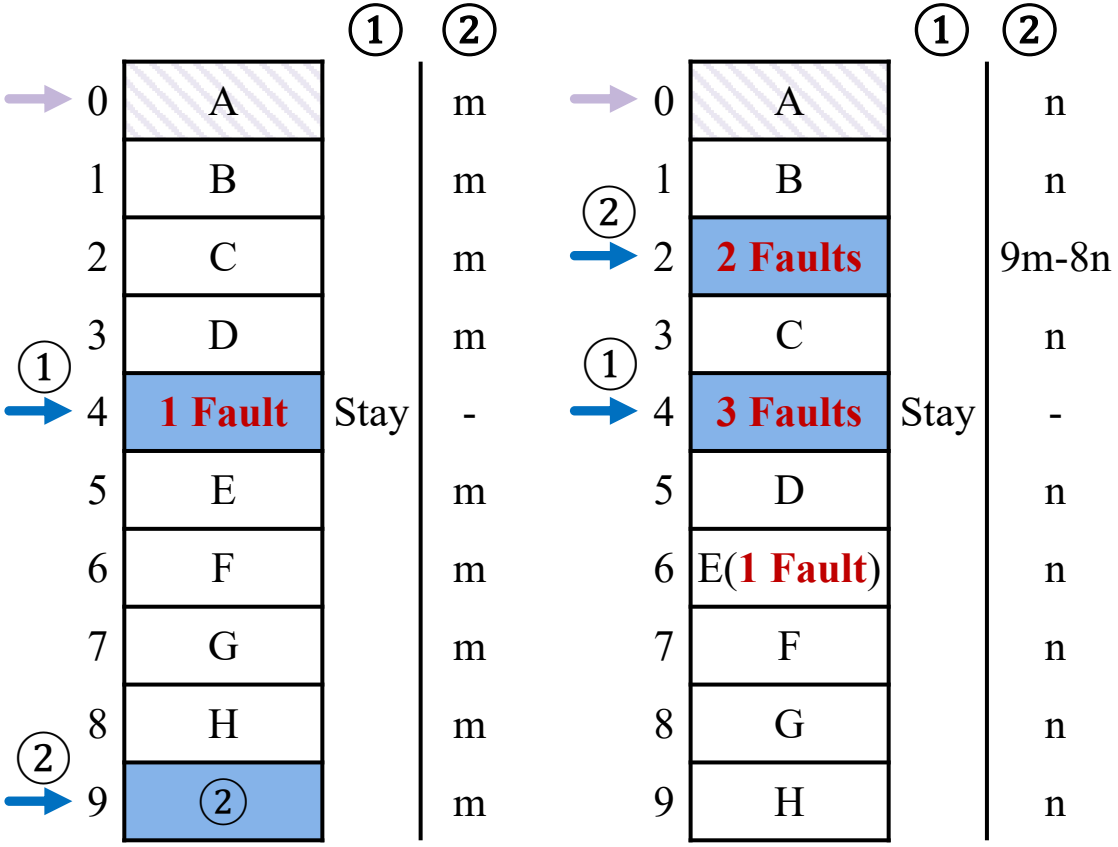

RETROFIT(multiple gaps)

M. K. Qureshi et al, MICRO 2009

RETROFIT: Fault-aware Wear Leveling

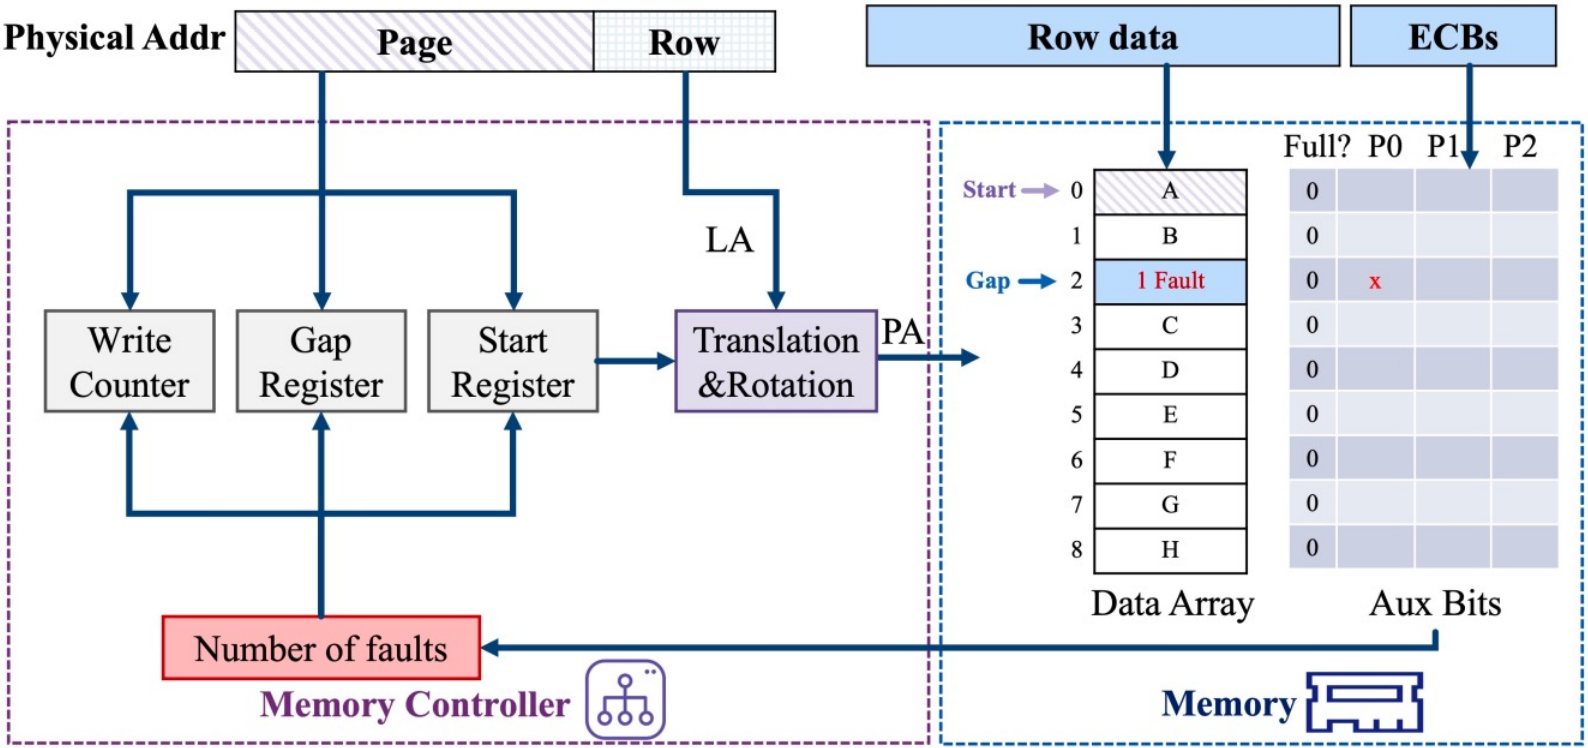

LA: Logical address
PA: Physical address

Memory controller design to implement RETROFIT

GPS to Enhance RETROFIT

- Retired rows often have **only one uncorrected fault. A big waste of resources!**
- Redesign retired rows as **Protecting-page pointers (PPPs).** A 512-bit row can be designed as 28 PPPs, which can additionally correct 28 in-page faults.

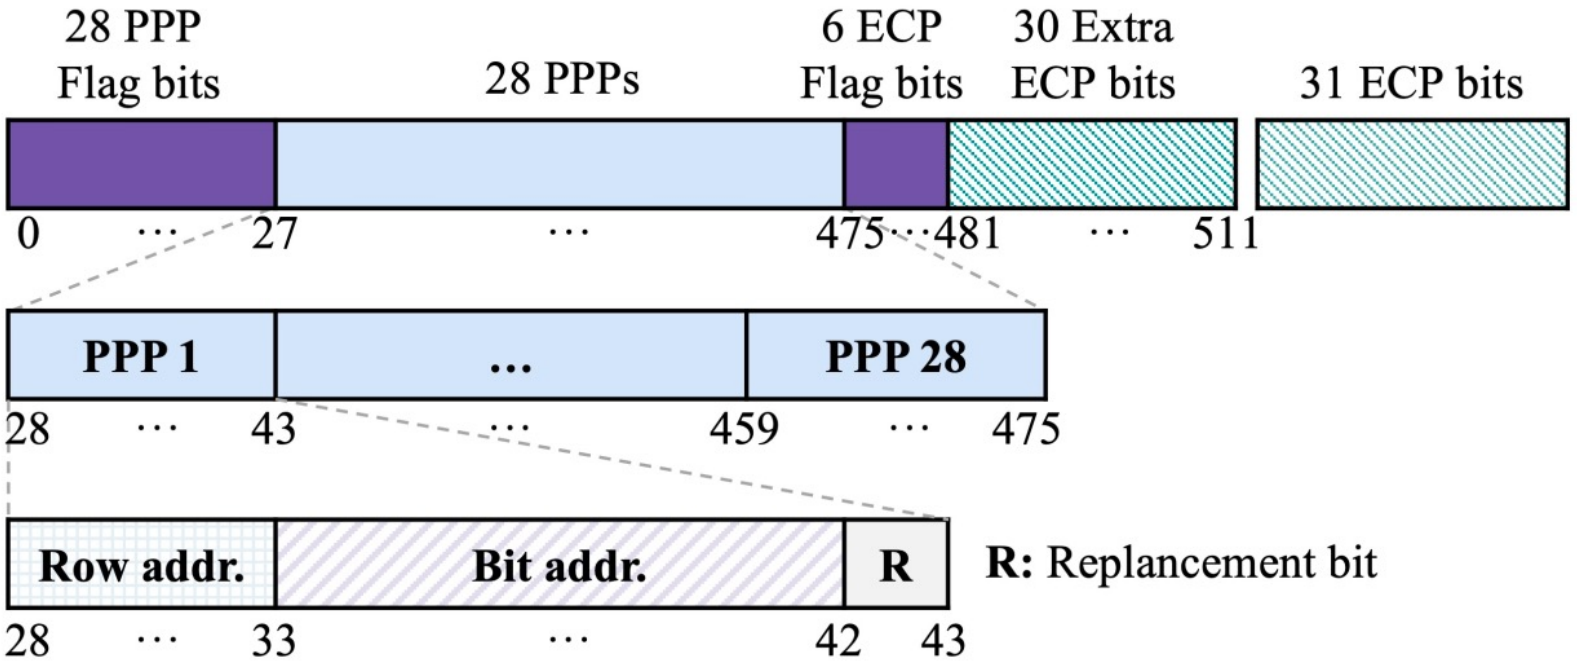

GPS to Enhance RETROFIT

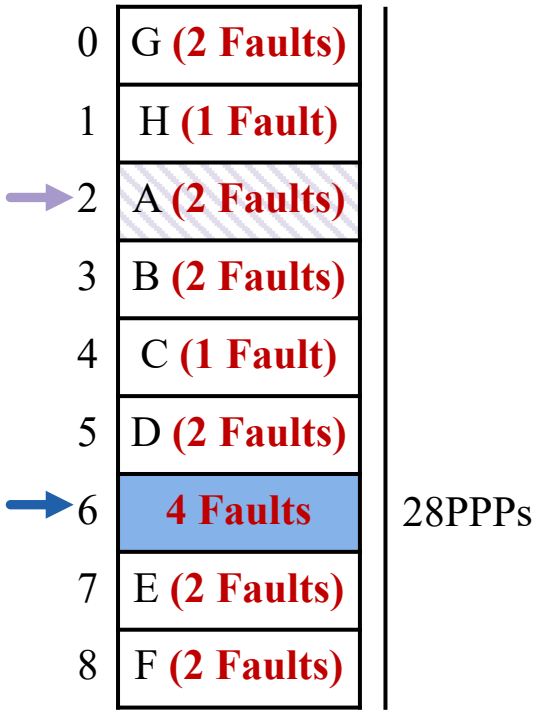

RETROFIT+GPS(1 gap)

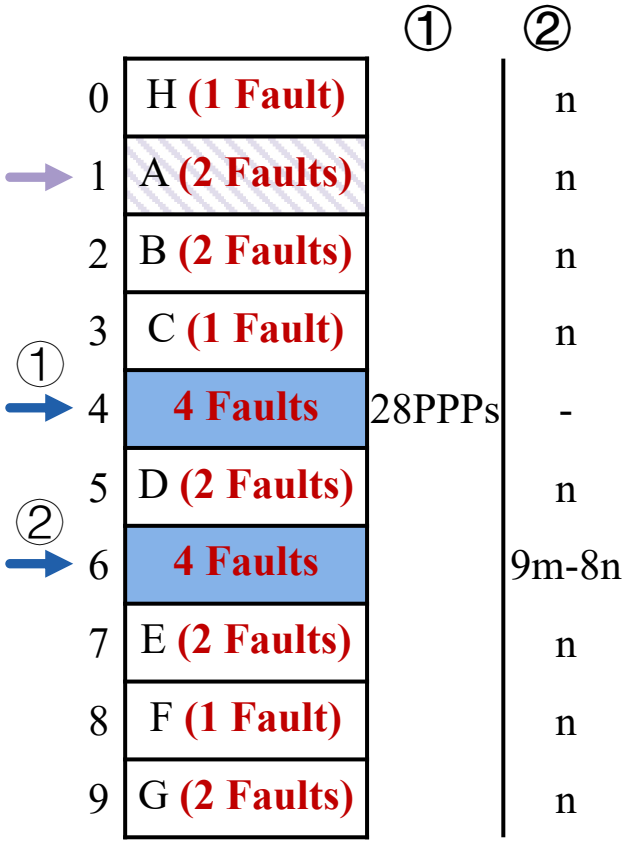

RETROFIT+GPS(2 gaps)

GPS to Enhance RETROFIT

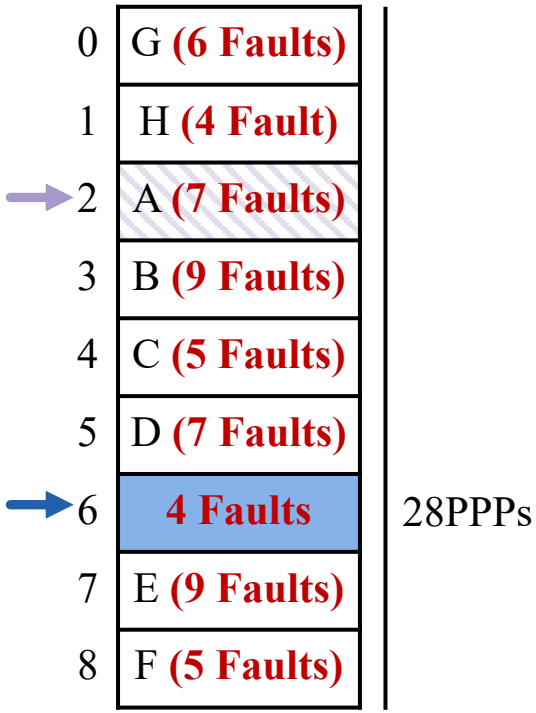

RETROFIT+GPS(1 gap)

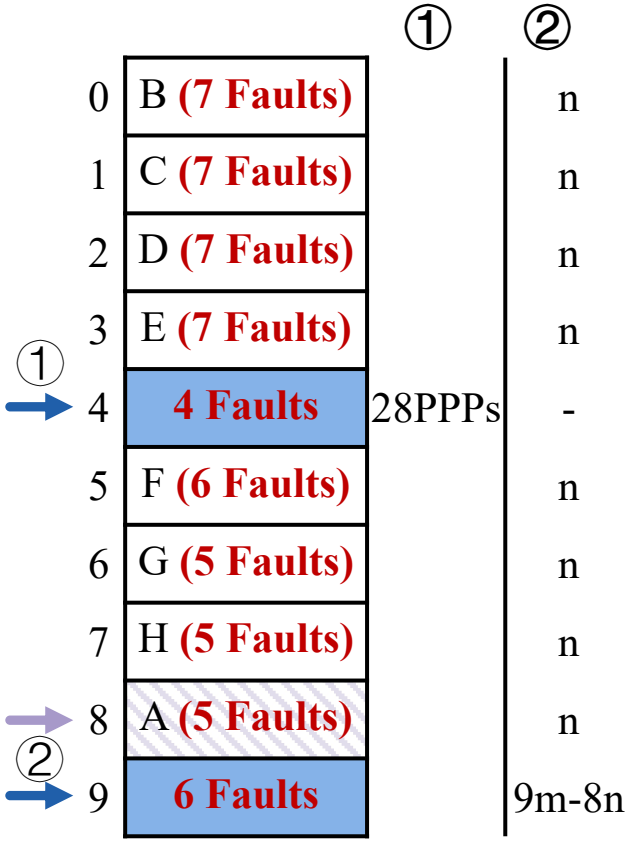

RETROFIT+GPS(2 gaps)

GPS to Enhance RETROFIT

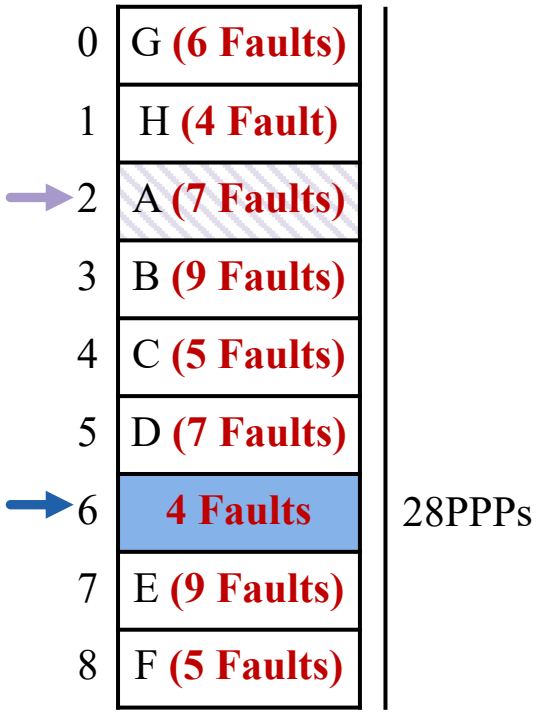

RETROFIT+GPS(1 gap)

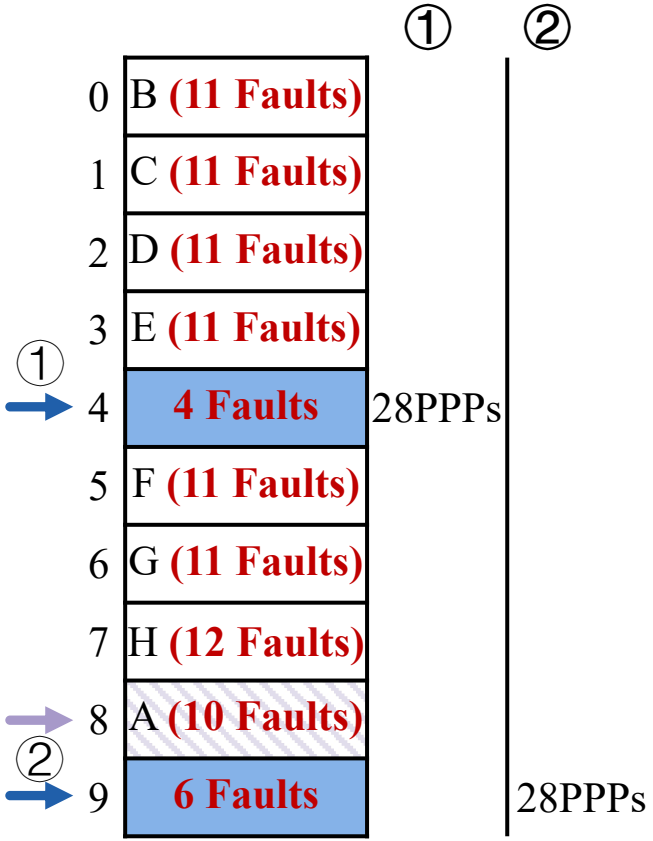

RETROFIT+GPS(2 gaps)

Column-level Wear Leveling

- Rotate the row being copied whenever the Gap moves.
- The error correction bits are reused in the rotation to reduce the bitflips. They are idle for a significant amount of time before a fault occurs.
- In order to support multiple Gap rows, the Start and Gap registers are extended to 10 bits and 5 bits.

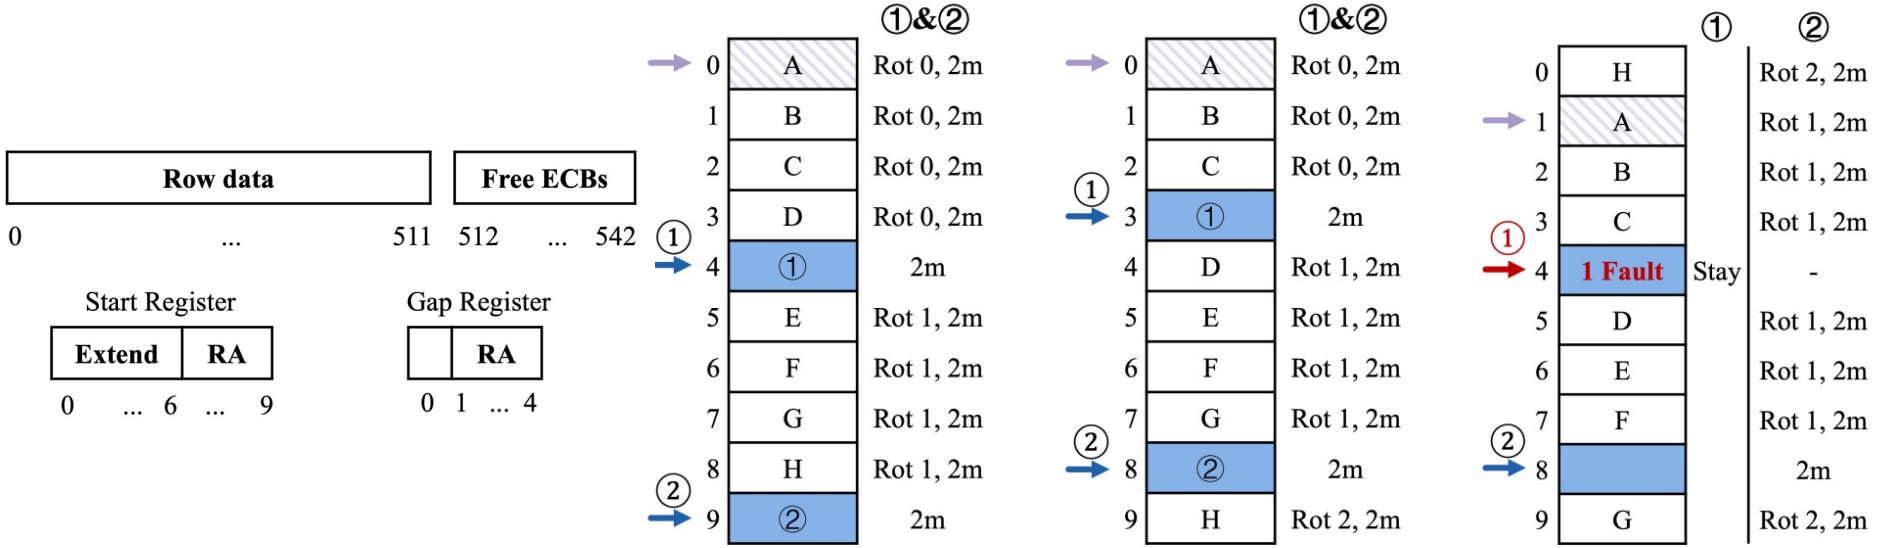

Put it all together

Fault-free

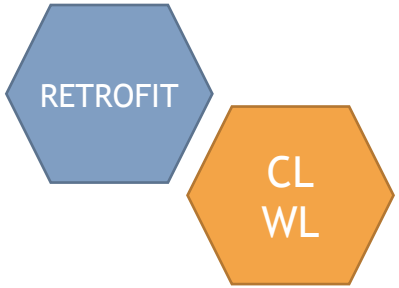

Move every S×m writes
Rotate the row simultaneously

With faults

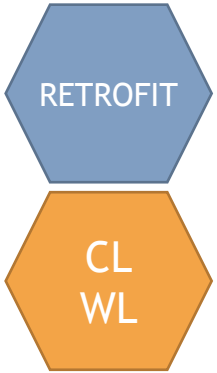

Faulty rows protected by spare rows or biased wear leveling

RETROFIT fails
GPS activates

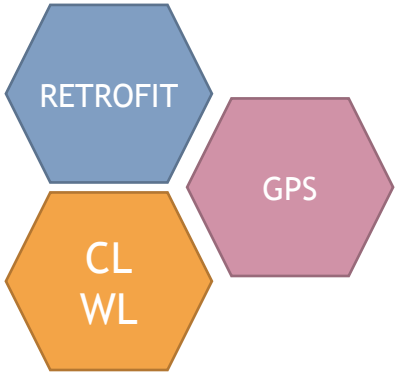

Continue wear leveling

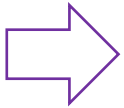

Life ends

Outline

- Background and Introduction
- RETROFIT
  - Fault-aware Wear Leveling
  - GPS to Enhance RETROFIT
  - Column-level Wear Leveling
  - Put it all together
- **Results**
- Conclusion

Results

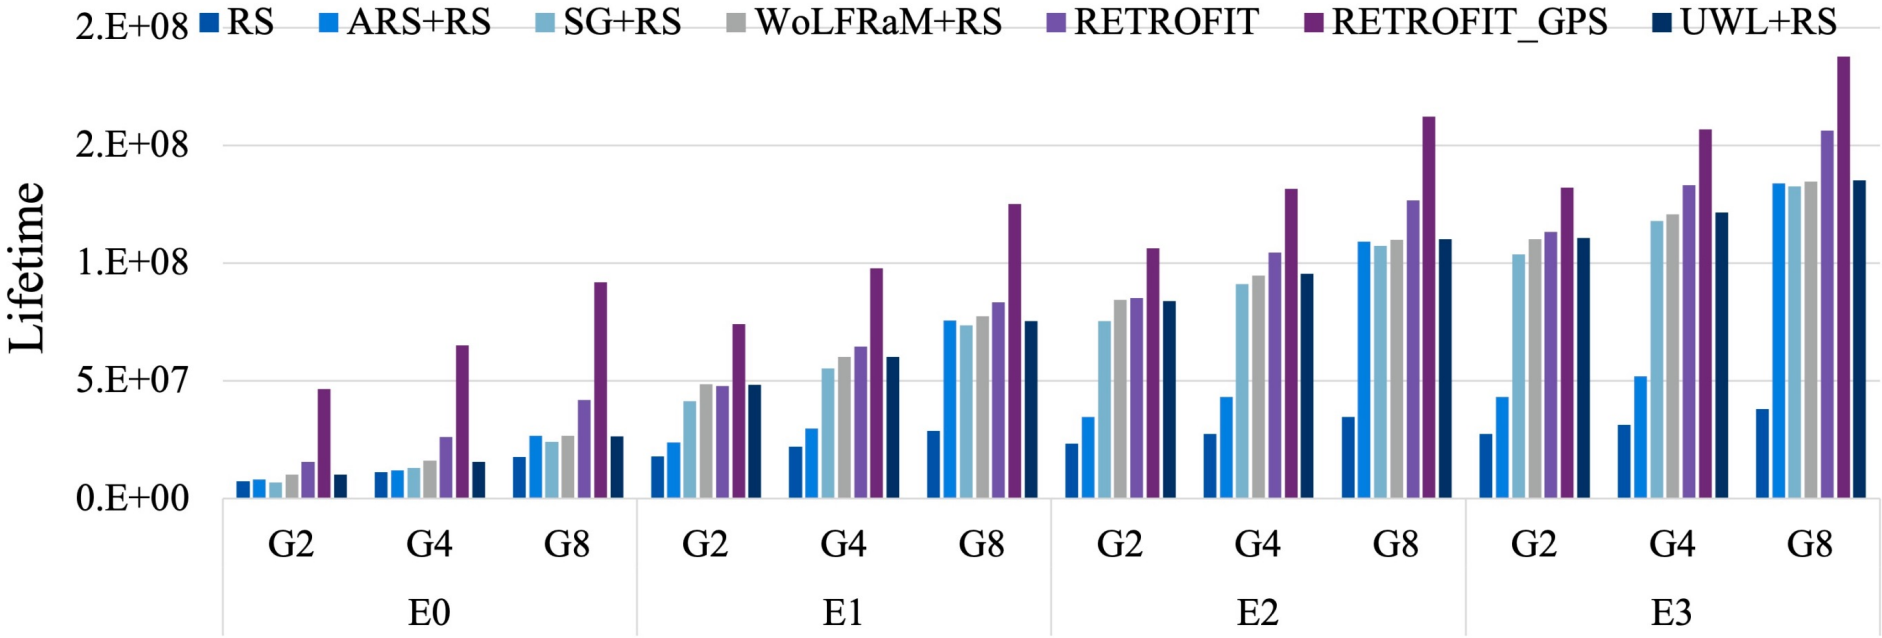

Ex stands for ECPx, Gx stands for x Gap lines per page

Based on Monte Carlo simulation, for the average results of 1000 endurance maps, RETROFIT and RETROFIT_GPS are 0.6x and 3.5x higher than SOTA respectively, and RETROFIT is 0.6x higher than UWL+RS (previous upper bound).

Results

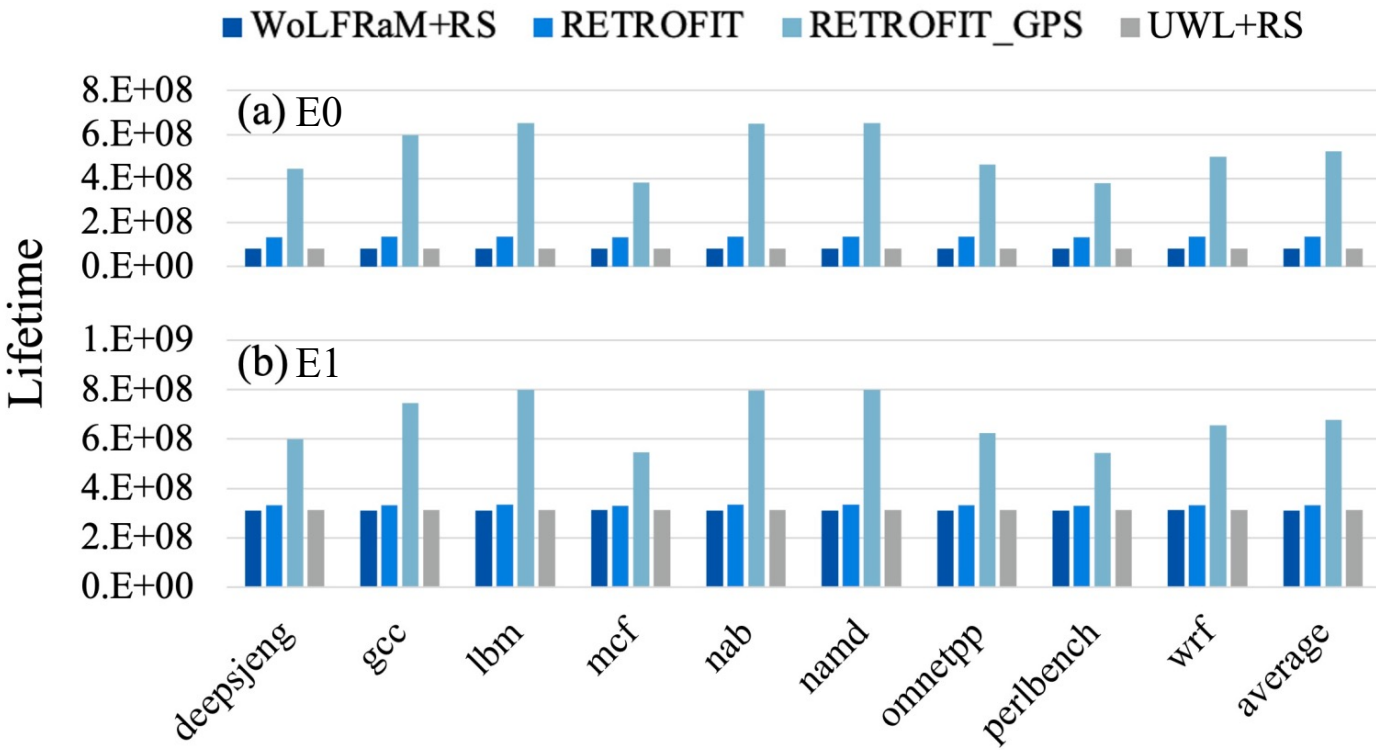

Based on the results of the SPEC2017 benchmarks, for the average results of 1000 endurance maps, RETROFIT and RETROFIT_GPS are 0.6x and 5.4x higher than SOTA, respectively.

Results

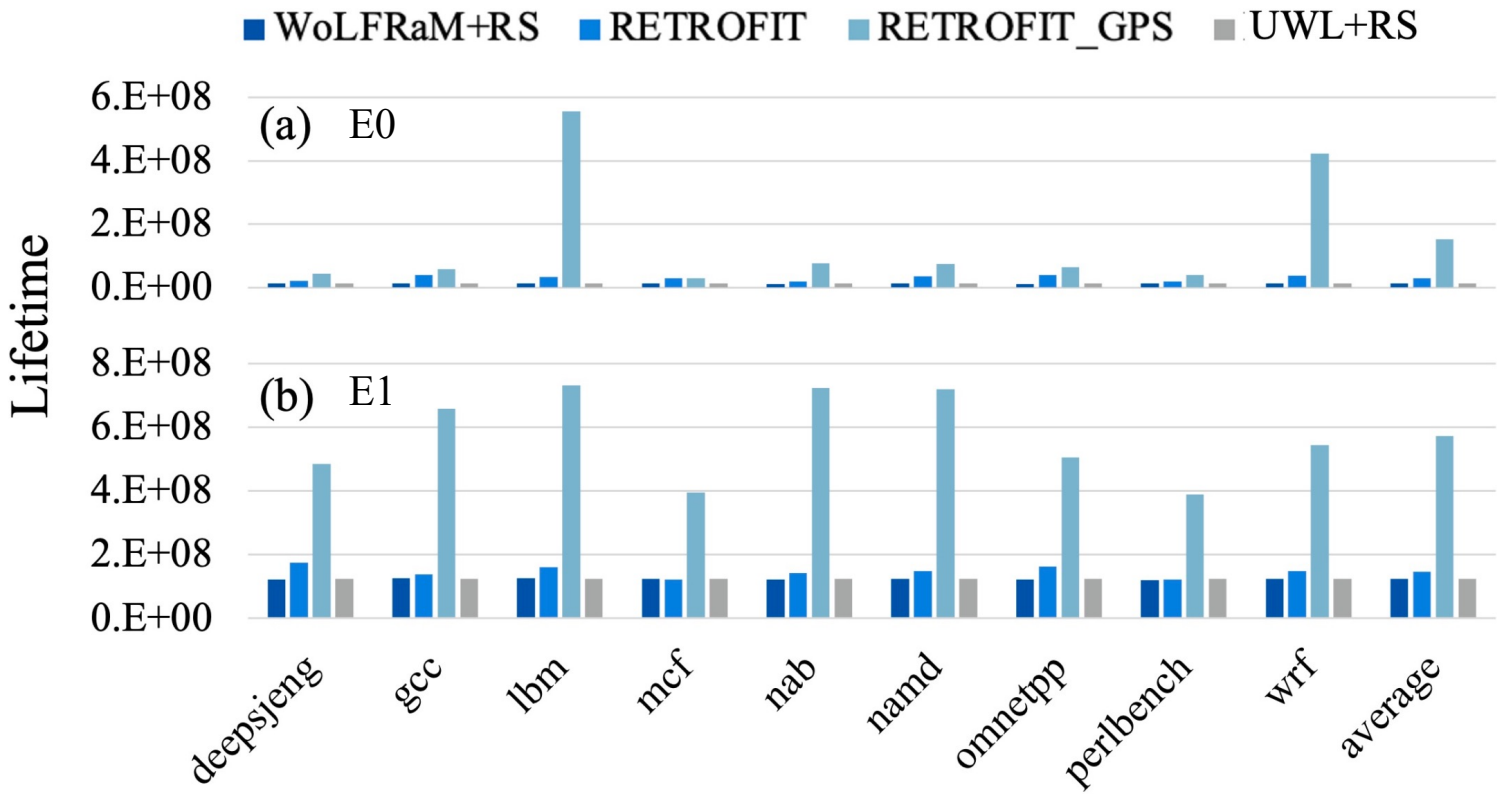

For the worst-case results of 1000 endurance maps, RETROFIT and RETROFIT_GPS are 2.6x and 16.0x higher than SOTA, respectively.

Results

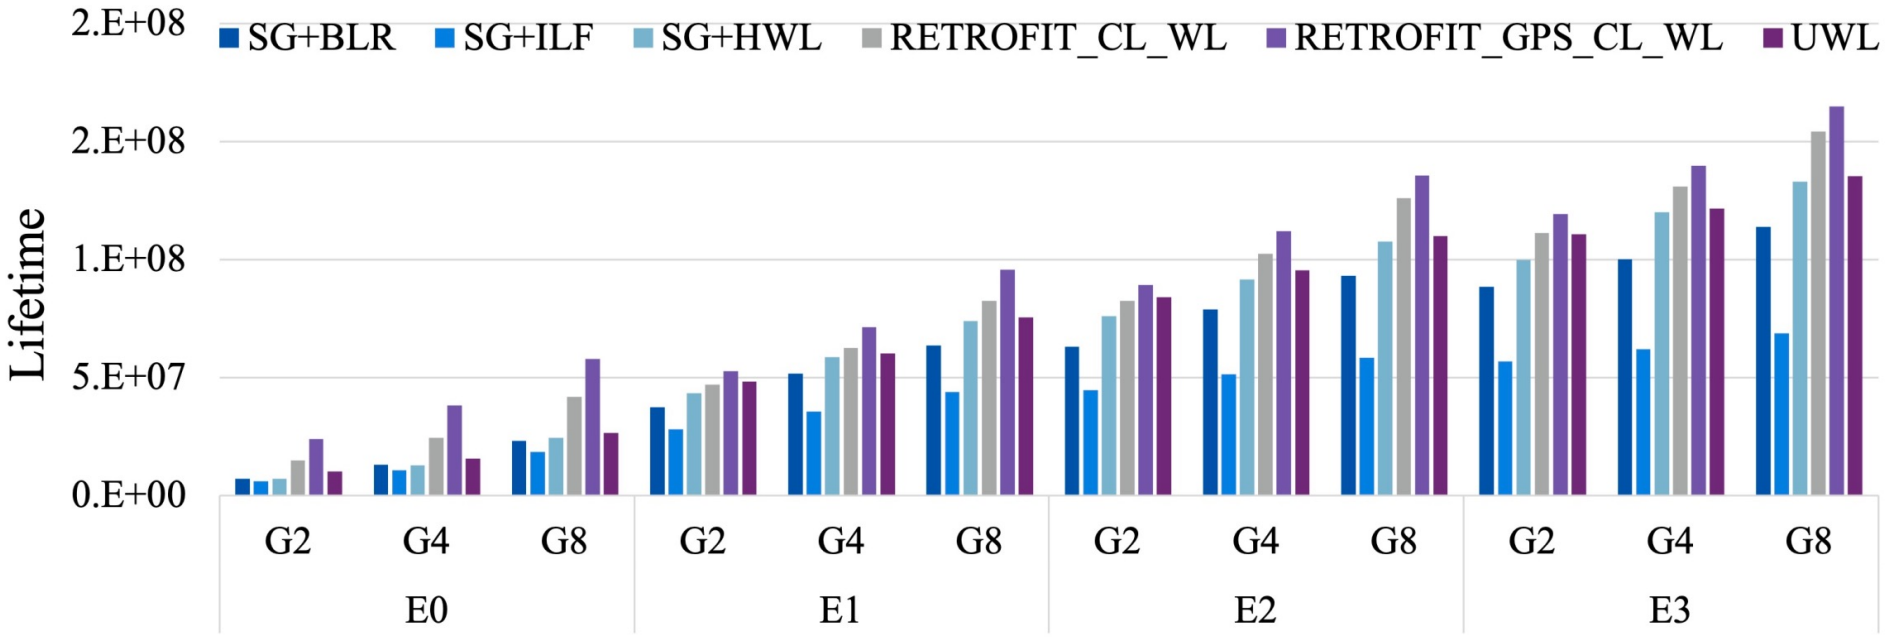

Ex stands for ECPx, Gx stands for x Gap lines per page

Based on the Monte Carlo simulation, for the average results, RETROFIT_CL_WL and RETROFIT_GPS_CL_WL are 1.1x and 2.4x higher than SOTA, respectively, and 0.6x and 1.5x higher than UWL.

Results

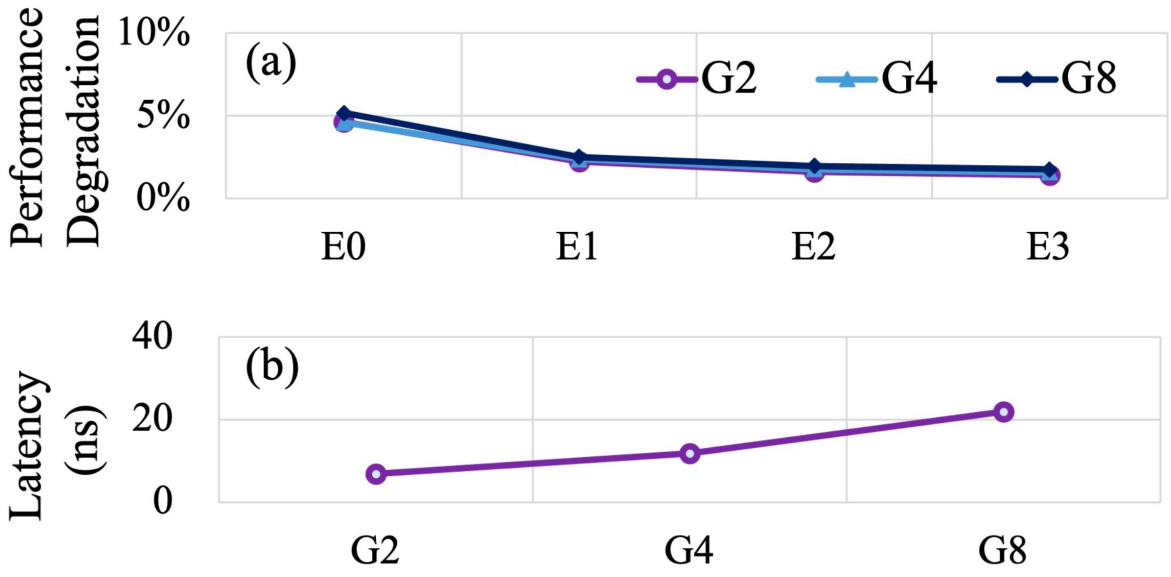

Ex stands for ECPx, Gx stands for x Gap lines per page

The proposed schemes reduce performance by 1.41%~5.17% and produce a delay of up to 22ns. The overall overhead is affordable.

Outline

- Background and Introduction
- RETROFIT
  - Fault-aware Wear Leveling
  - GPS to Enhance RETROFIT
  - Column-level Wear Leveling
  - Put it all together
- Results
- **Conclusion**

Conclusion

- **Process variation aggravates the uneven strength** between rows and bits in a row. Even if the memory is accessed uniformly, the memory lifetime will decrease by two orders of magnitude;
- RETROFIT uses **the number of faults in rows to reflect the strength of the rows** and implements biased wear leveling to break through the theoretical upper limit of uniform distribution and increase lifetime by 0.68x;
- GPS further improves lifetime by **reusing retired rows**, prolonging wear leveling, and improving the utilization rate of error correction capabilities;
- Column-level wear leveling reduces the average bitflips by **reusing idle error correction bits** and supports RETROFIT compatibility with multiple sparing rows;
- The results show that RETROFIT_GPS can increase the memory lifetime over SOTA by **5 times**, exceeding the uniform wear leveling by **1.5 times**;
- RETROFIT can also address the challenges of weak cells caused by DRAM bitline/wordline crosstalk, noise, single event upsets (SEUs), etc.

References

1. D. Kline, J. Zhang, R. Melhem, and A. K. Jones, “Flower and fame: A low overhead bit-level fault-map and fault-tolerance approach for deeply scaled memories,” in HPCA. IEEE, 2020, pp. 356–368.
2. C.-K. Luk, R. Cohn, R. Muth, H. Patil, A. Klauser, G. Lowney, S. Wallace, V. J. Reddi, and K. Hazelwood, “Pin: building customized program analysis tools with dynamic instrumentation,” in Acm sigplan notices, vol. 40, no. 6, 2005, pp. 190–200.
3. M. K. Qureshi, J. Karidis, M. Franceschini, V. Srinivasan, L. Lastras, and B. Abali, “Enhancing lifetime and security of pcm-based main memory with start-gap wear leveling,” in MICRO. IEEE, 2009, pp. 14–23.
4. S. Rashidi, M. Jalili, and H. Sarbazi-Azad, “A survey on pcm lifetime enhancement schemes,” CSUR, vol. 52, no. 4, pp. 1–38, 2019.
5. S. Schechter, G. H. Loh, K. Strauss, and D. Burger, “Use ecp, not ecc, for hard failures in resistive memories,” in ISCA, 2010, pp. 141–152.
6. N. H. Seong, D. H. Woo, V. Srinivasan, J. A. Rivers, and H.-H. S. Lee, “Safer: Stuck-at-fault error recovery for memories,” in MICRO, 2010, pp. 115–124.
7. L. Yavits, L. Orosa, S. Mahar, J. D. Ferreira, M. Erez, R. Ginosar, and O. Mutlu, “Wolfram: Enhancing wear-leveling and fault tolerance in resistive memories using programmable address decoders,” in ICCD. IEEE, 2020, pp. 187–196.
8. V. Young, P. J. Nair, and M. K. Qureshi, “Deuce: Write-efficient encryption for non-volatile memories,” ACM SIGARCH Computer Architecture News, vol. 43, no. 1, pp. 33–44, 2015.
9. J. Yun, S. Lee, and S. Yoo, “Bloom filter-based dynamic wear leveling for phase-change ram,” in DATE, 2012, pp. 1513–1518.
10. J. Zhang, D. Kline, L. Fang, R. Melhem, and A. K. Jones, “Yoda: Judge me by my size, do you?” in ICCD, 2017, pp. 395–398.
11. W. Zhang and T. Li, “Characterizing and mitigating the impact of process variations on phase change based memory systems,” iMICRO. IEEE, 2009, pp. 2–13.

References

12. X. Zhang and G. Sun, “Toss-up wear leveling: Protecting phase-change memories from inconsistent write patterns,” in DAC, 2017.
13. M. Zhao, B. Gao, J. Tang, H. Qian, and H. Wu, “Reliability of analog resistive switching memory for neuromorphic computing,” Applied Physics Reviews, vol. 7, no. 1, p. 011301, 2020.
14. M. Zhao, L. Shi, C. Yang, and C. J. Xue, “Leveling to the last mile: Near-zero-cost bit level wear leveling for pcm-based main memory,” in ICCD. IEEE, 2014, pp. 16–21.
15. V. Young, P. J. Nair, and M. K. Qureshi, “Deuce: Write-efficient encryption for non-volatile memories,” ACM SIGARCH Computer Architecture News, vol. 43, no. 1, pp. 33–44, 2015.

Thank you for your attention!
Q&A

Email: jwzhang.pitt@gmail.com