FAST'16
FEBRUARY 22–25, 2016
SANTA CLARA, CA
Sponsored by USENIX
in cooperation with ACM SIGOPS
14th USENIX Conference on File and Storage Technologies
usenix THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

Access Characteristic Guided Read and Write Cost Regulation for Performance Improvement on Flash Memory

Qiao Li [§], Liang Shi [§], Chun Jason Xue [※]
Kaijie Wu [§], Cheng Ji [※], Qingfeng Zhuge [§], and Edwin H.-M. Sha [§]

[§] College of Computer Science, Chongqing University
[※] Department of Computer Science, City University of Hong Kong

重慶大學 CHONGQING UNIVERSITY

香港城市大學 City University of Hong Kong

Outline

- Background
- Design
- Evaluation
- Conclusion

Background

- NAND flash memory is widely used from USB to big data centers.

- Flash memory development
  - bit density: from 1 bit to 6 bits
  - technology scaling: from 65nm to 10nm
- Performance degradation

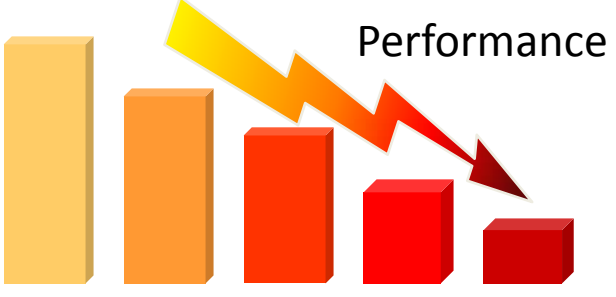

This paper's objective : improve read and write performance

Flash Write (Programming)

- Incremental Step Pulse Programing (ISPP) is used to program Flash page
- The program voltage is increased by the step size
- Finished when the voltage exceeds the threshold voltage V_{th}
- Program time → write cost

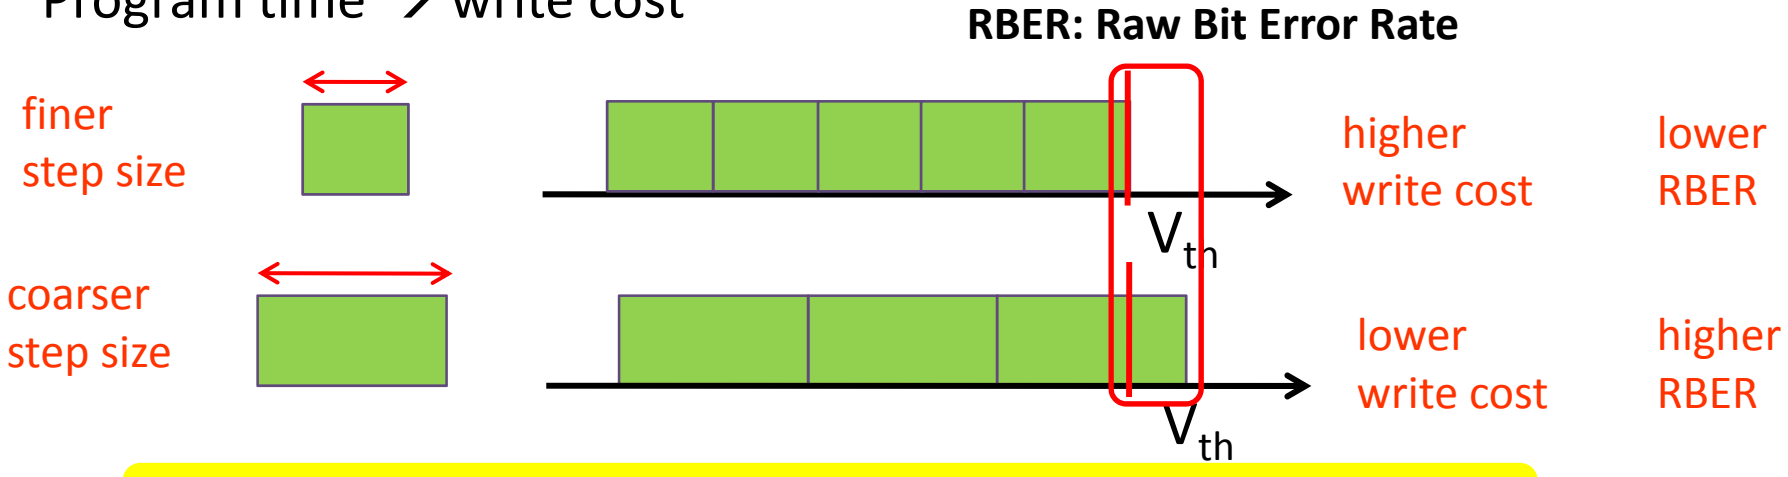

Write cost is related to RBER (Raw Bit Error Rate)

Flash Read

- Low-Density Parity Code (LDPC) is applied in Flash for strong ECC capability
- The decoding strength of LDPC depends on the accuracy of input information

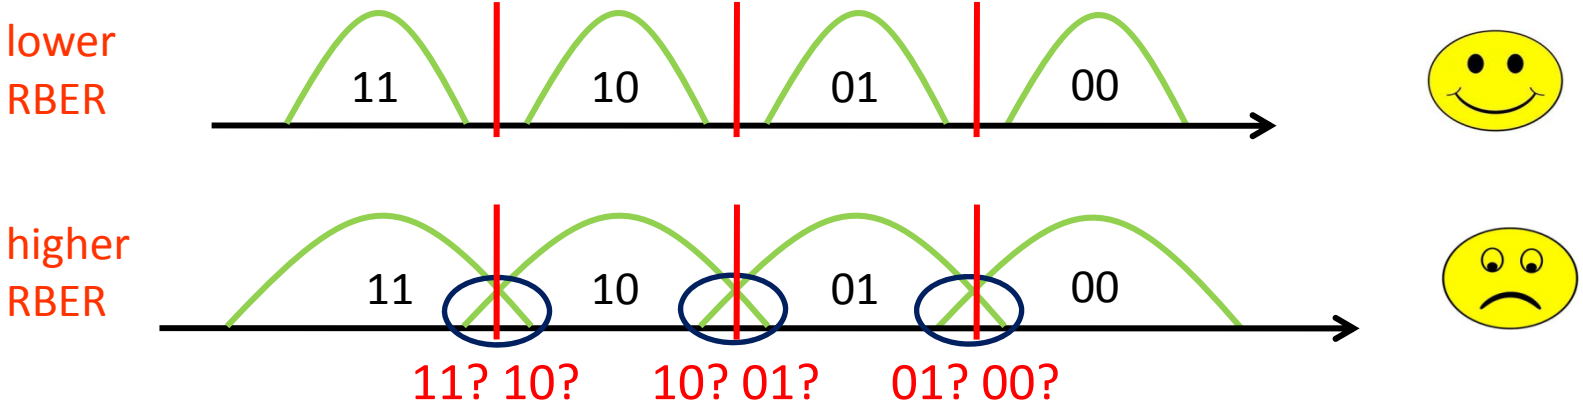

Flash Read

- Low-Density Parity Code (LDPC) is applied in Flash for strong ECC capability
- The decoding strength of LDPC depends on the accuracy of input information

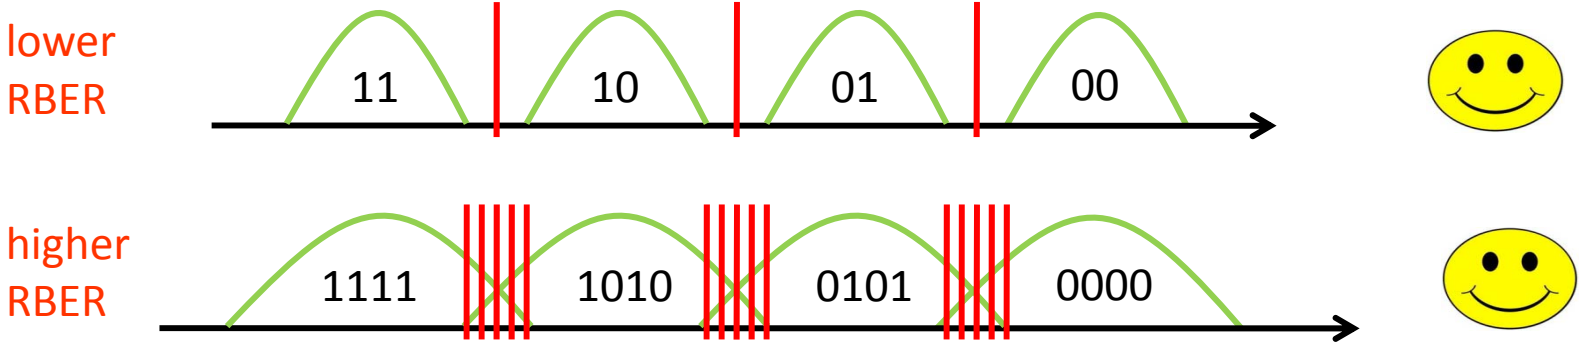

Flash read: Step1. Sensing Step2. Transfer

To decode data with higher RBER, higher read cost is needed.

Read cost is related to RBER (Raw Bit Error Rate)

Read and Write Cost Regulation

- Read and write costs can be regulated

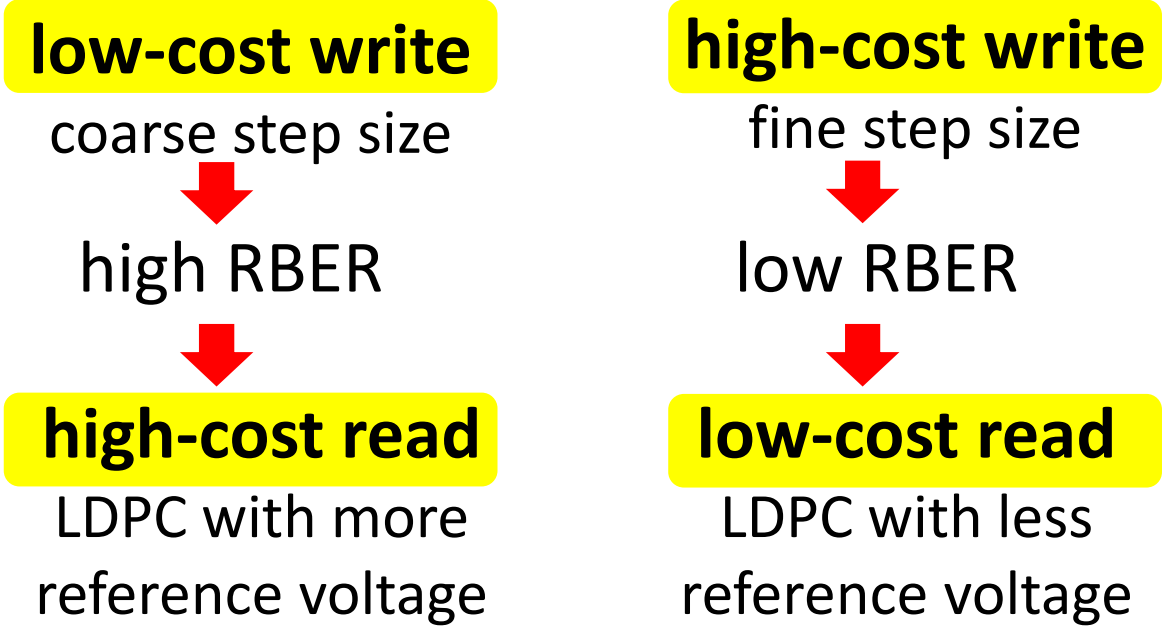

Preliminary Study

- Experiments are conducted to investigate the difference between different write and read costs in real workloads, including PROJ, USR, HM from MSR
- 3 combinations of read and write costs are evaluated

LCW/HCR Low-Cost Write/High-Cost Read
MCW/MCR Medium-Cost Write/Medium-Cost Read
HCW/LCR High-Cost Write/Low-Cost Read

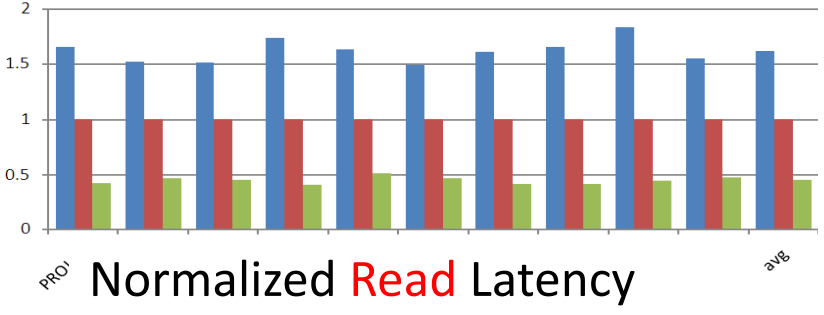

Normalized Read Latency

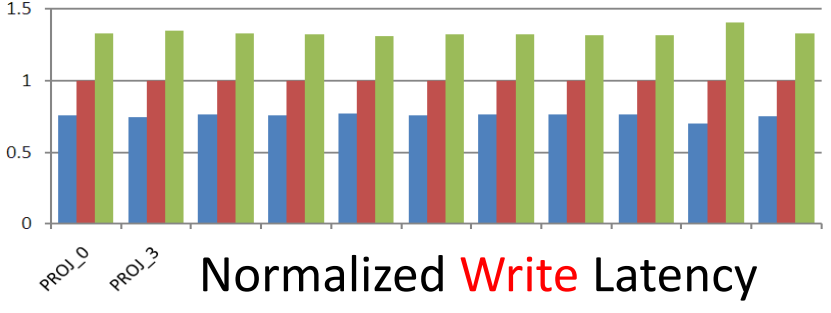

Normalized Write Latency

- Co… …CW/LCR:
  - …
  - …

The performance gap indicates that the read and write cost regulation should be applied carefully

Ideal Case

- All reads are performed with low-cost read
- All Writes are performed with low-cost write
- Ideal but **impossible**!
- How to regulate cost for performance improvement?

Our approach is based on the access characteristic of workloads

Observation

- Key Observation of Access Characteristics:
  - **Most read requests access read-only pages, more than 85% on average!**
  - **Most write requests access write-only pages, more than 91% on average!**
  - Only a small part of reads and writes access interleaved-access pages

Distribution of Reads

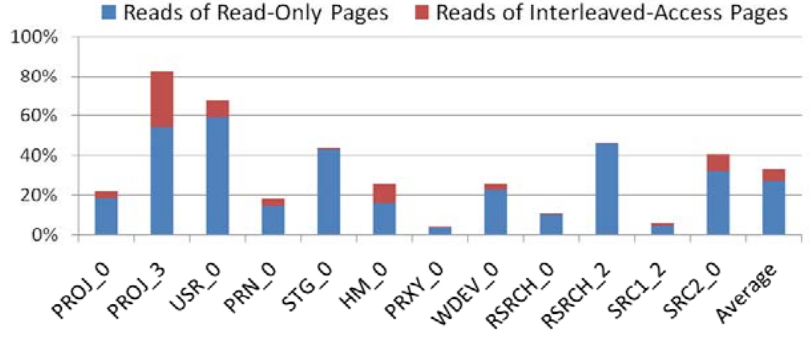

Distribution of Writes

Approach

Access Cost Regulation

- Read-only pages --- low-cost read
- Write-only pages --- low-cost write
- Interleaved-access pages --- medium-cost access (default)

Identification

- Identify read-only pages and write-only pages

Re-Write

- The cost of read is determined by the write on the data.
- Re-write read-only pages that are accessed by high-cost read.

Access Characteristic Identification

- Access history per page

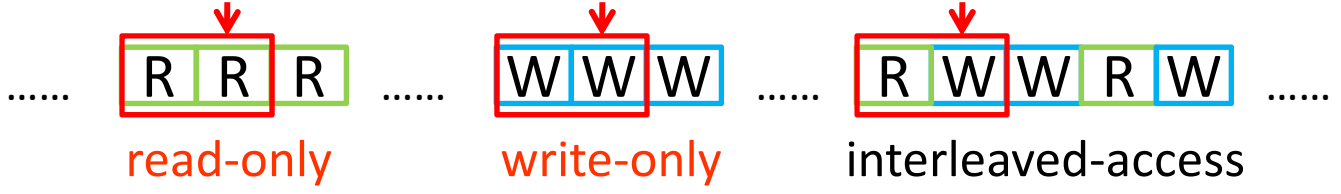

- History window
  - Upcoming access + most recent access
- Re-write Read-only pages with high-cost during idle time
- Each mapping entry in the FTL is extended with two fields
  - 1-bit low-cost write tag, and the access history

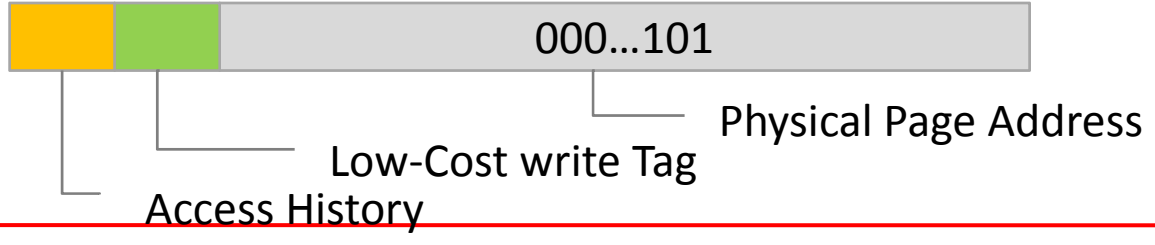

Experiments

- Simulator: SSDSim [15][16]
- 12 workloads from MSR [17]
- 8 channels, 8 chips per channel and 4 planes per chip
- Default FTL, page mapping, garbage collection and wear leveling
- Comparing these 3 techniques
  - Traditional: Normal Flash without Cost Regulation
  - Li et. al: Cost Regulation for Access Conflict Minimization
  - AGCR : The Proposed Technique

Experiment Results

- Read and write performance are evaluated.

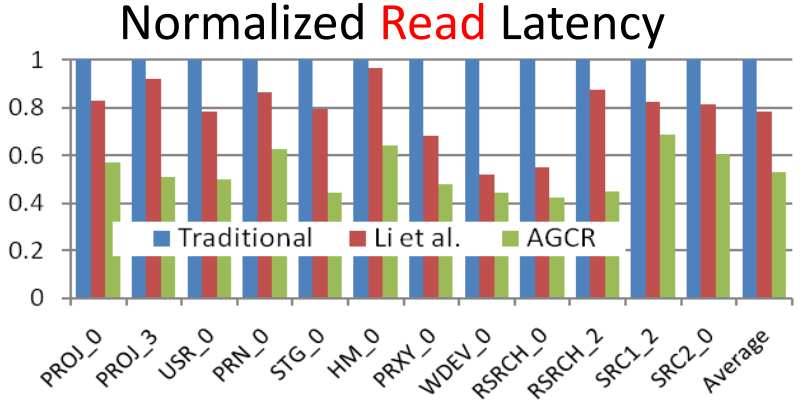

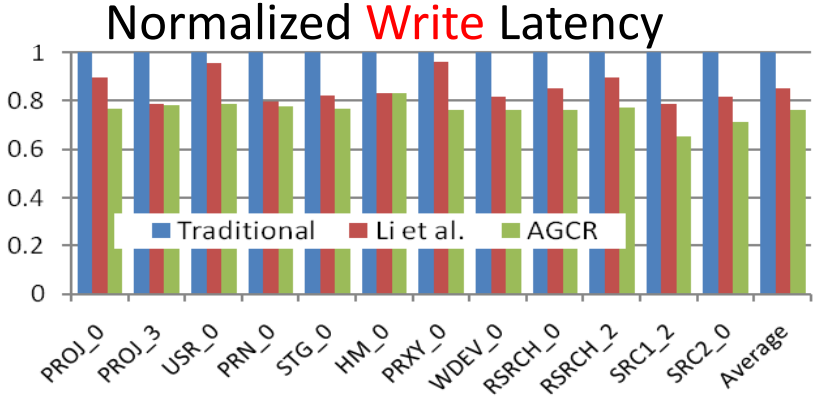

- Compared to the state-of-the-art approach:
  - AGCR improves read performance by 32% on average

  And at the same time
  - AGCR improves write performance by 22% on average

Experiment Results

- The distributions for operations of different costs.

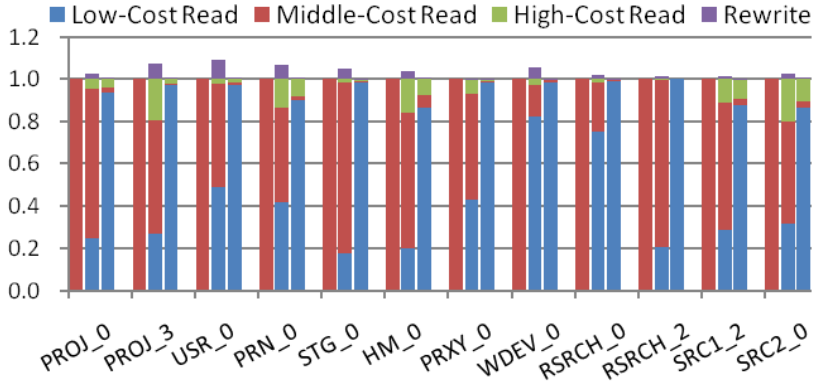

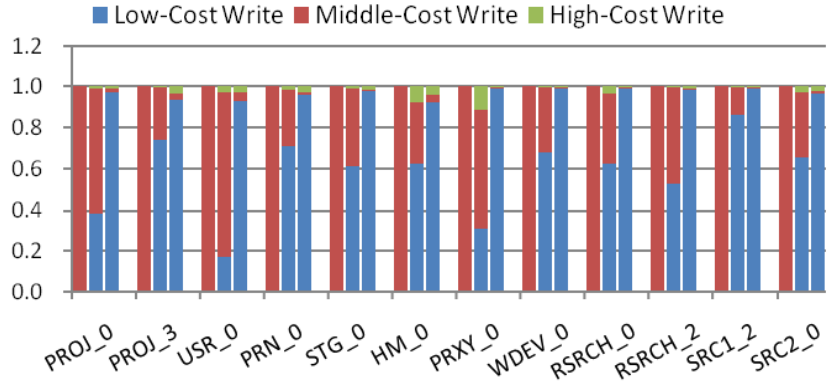

- Comparing to Li et al.'s work, AGCR issues considerably more low-cost reads and writes
- The percentage of re-write operations is no more than 1% of all accesses issued by the host

Identification Accuracy

- Impact of window size on identification accuracy.

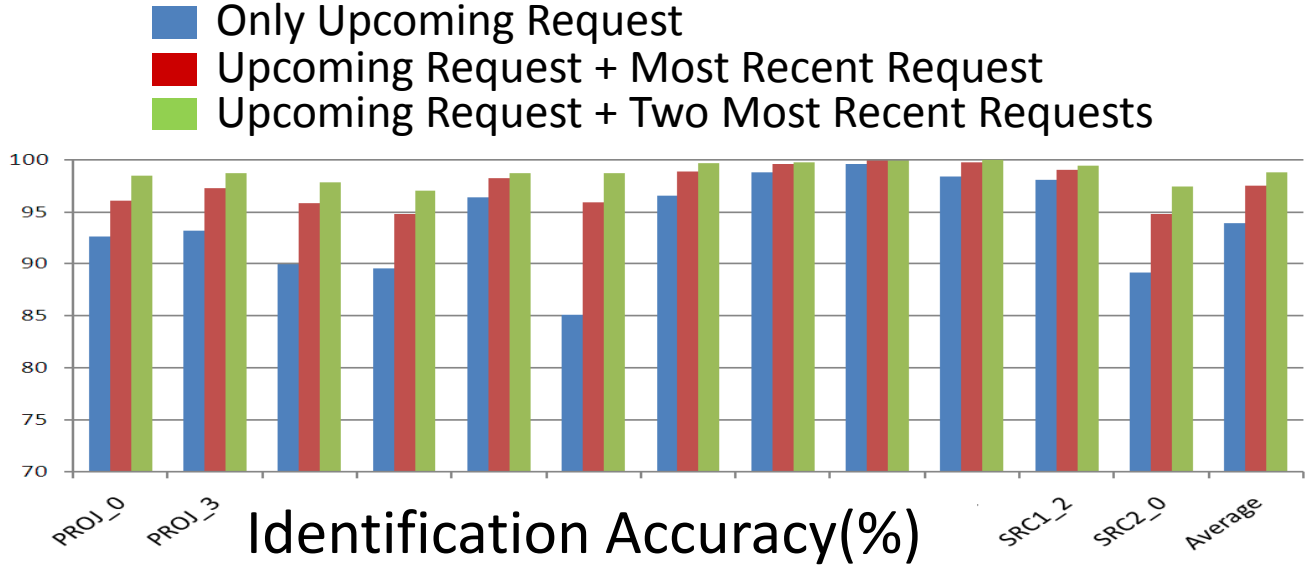

- The Identification method achieves high accuracy.
- A larger window results in higher accuracy.

Conclusion

Preliminary Study

- We presented a preliminary study to show the potential performance improvement of our approach.

Observation

- We made the observation that most reads (writes) access read-only (write-only) pages.

Approach

- We proposed a comprehensive approach to regulate the cost of reads and writes.

Evaluation

- Results show that the proposed approach achieves significant performance improvement.

Thank you!

Questions?