LDPC-in-SSD: Making Advanced Error Correction Codes Work Effectively in Solid State Drives

Kai Zhao*, Wenzhe Zhao†, Hongbin Sun†,
Tong Zhang*, Xiaodong Zhang‡, and Nanning Zheng†

* ECSE Department, Rensselaer Polytechnic Institute, USA
† Xi'an Jiaotong University, P.R.China
‡ Department of CSE, The Ohio State University, USA

NAND Flash Memory

- Increasing Adoptions and Decreasing Cost of NAND Flash Memory

Increasingly wide adoption

consumer

client

enterprise

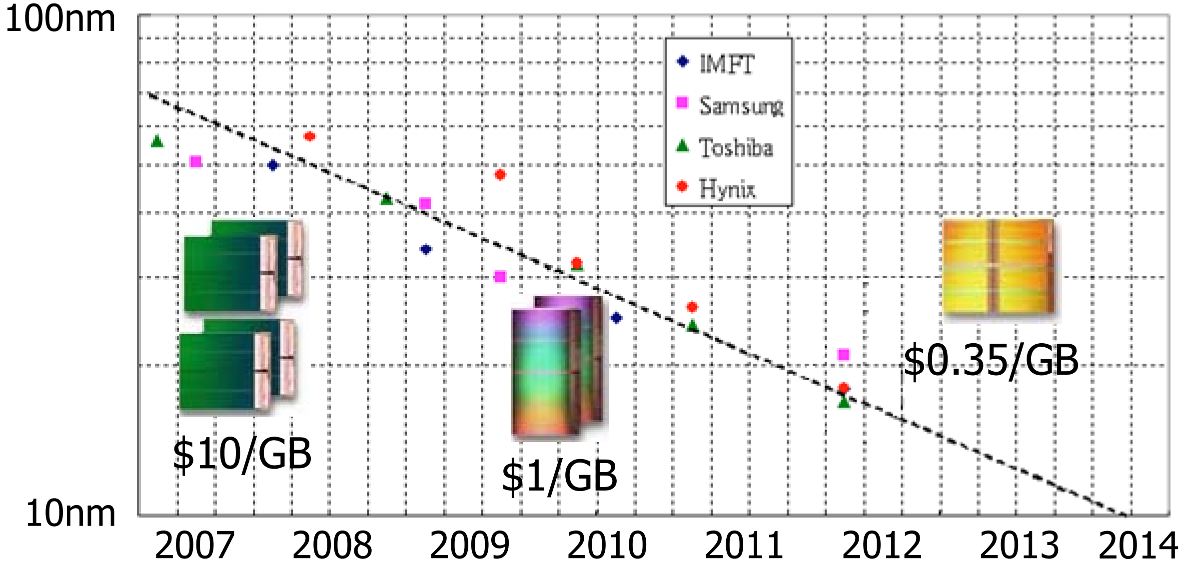

Bit cost reduction

NAND Flash Memory

- Noise in NAND flash memory increases as chip technology scaling continues

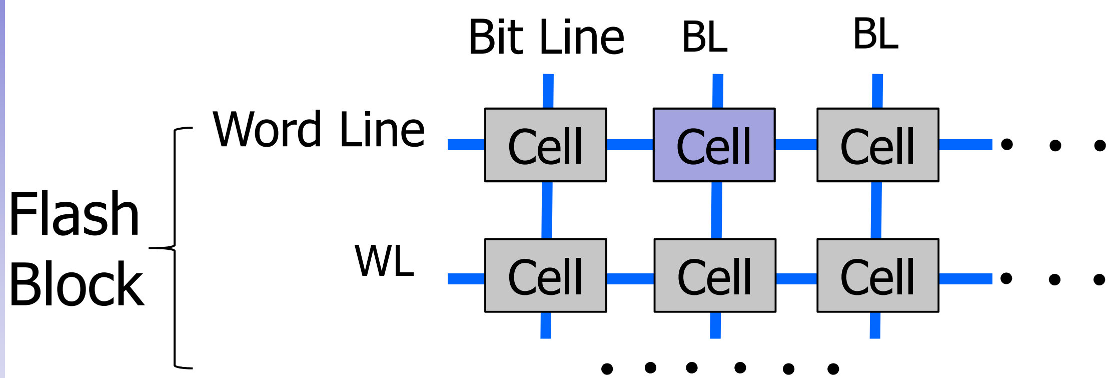

- Cell size and distance between cells shrink

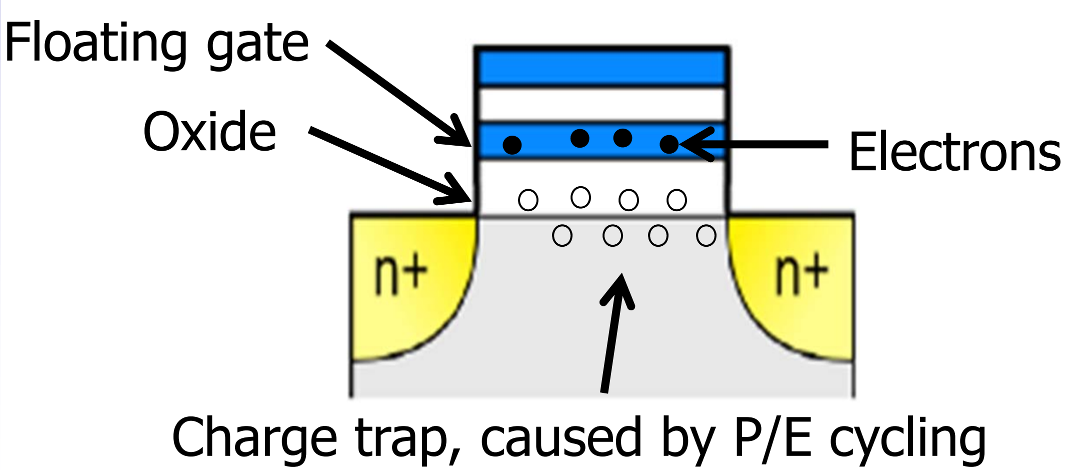

- Oxide layer becomes thinner
- Number of electrons held in floating gate reduced

NAND Flash Memory

- Smaller cell, thinner oxide layer and fewer electrons make the flash memory increasingly noisy

<u>Cell-to-cell interference</u>

Victim cell

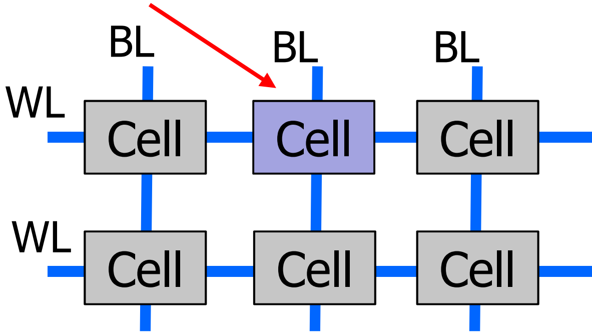

<u>Random telegraph noise (RTN)</u>

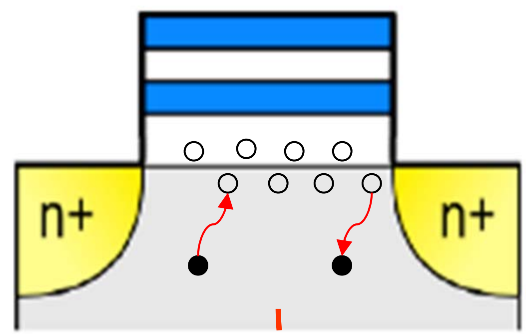

<u>Retention noise</u>

Electrons leaks from the floating gate

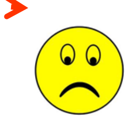

Higher raw bit error rate

Error Correction Codes in Storage Systems

- Bose-Chaudhuri-Hocquengham (BCH) code
  - A class of cyclic error-correcting codes, invented in 1959, hard-decision decoding, adopted in all SSDs
- Low-density parity-check (LDPC) code
  - A linear error-correcting code, invented by Robert.G.Gallager in 1963 and rediscovered in 1996, hard-decision/soft-decision decoding

Soft-decision LDPC code has a much stronger error correction capability

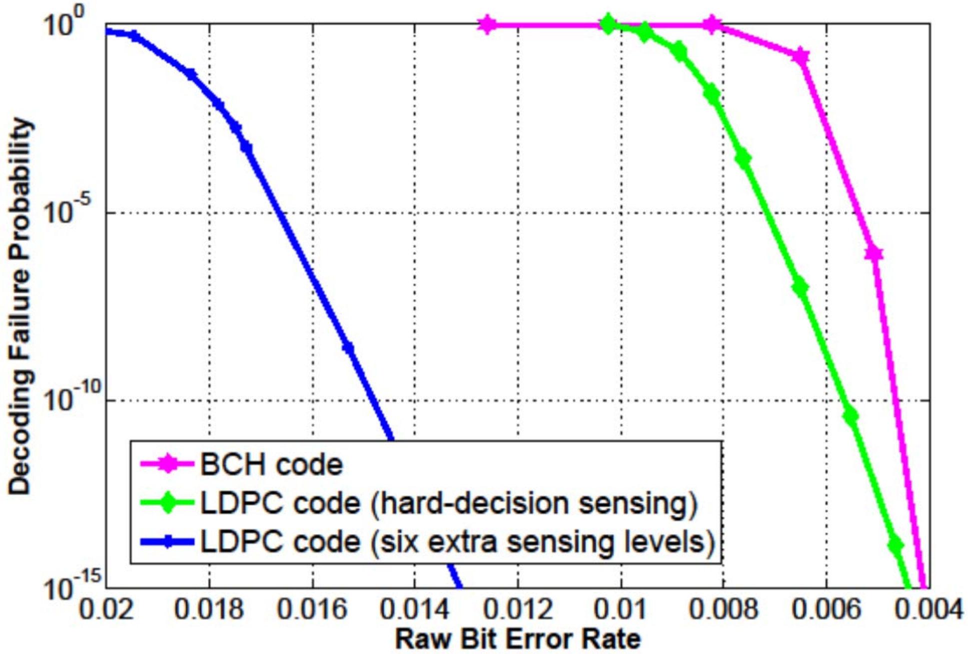

Error Correction Codes in Storage Systems

- BCH codes' relatively weak error correction capability becomes inadequate
- LDPC code can significantly improve the reliability of SSDs compared to BCH code
- Challenges for adopting LDPC codes in SSDs
  - Designing LDPC codes of good performance
  - Techniques to address the LDPC decoder input initialization problem
  - Minimizing LDPC-induced latency increase in SSDs

Outline

- LDPC code and soft-decision sensing
- Proposed techniques
- Experiment Setup and Preparation
- Experimental Results
- Conclusion

LDPC Code and Soft Sensing

- Hard-decision decoding (BCH, LDPC)
  - Inputs are in binary form
  - Simple hardware implementation
  - Relatively weak error correction capability
- Soft-decision decoding (LDPC)
  - Inputs are quantized to integers
  - Complicated hardware design
  - Strong error correction capability

For LDPC code, its error correction strength strongly depends on the *accuracy of the input information* which presents the probability information of the data stored in NAND flash memory

LDPC Code and Soft Sensing

$x_{n-1}, \cdots, x_1, x_0$

Data encoded

Noisy channel

$y_{n-1}, \cdots, y_1, y_0$

Hard-decision

1,0,1,1,0,1,0,0,0,1,1,0

Input Init.

-4,+2,-1,-1,+3,-2,+2,+3,+2,-4,-3,+1

Soft-decision

$$\log\left(\frac{P(y_i \mid x_i = 0)}{P(y_i \mid x_i = 1)}\right)$$

Log-likelihood ratio (LLR)

Positive: 0
Negative: 1
Magnitude presents the probability

LDPC Decoder

LDPC Code and Soft-sensing

Hard sensing vs. Soft sensing

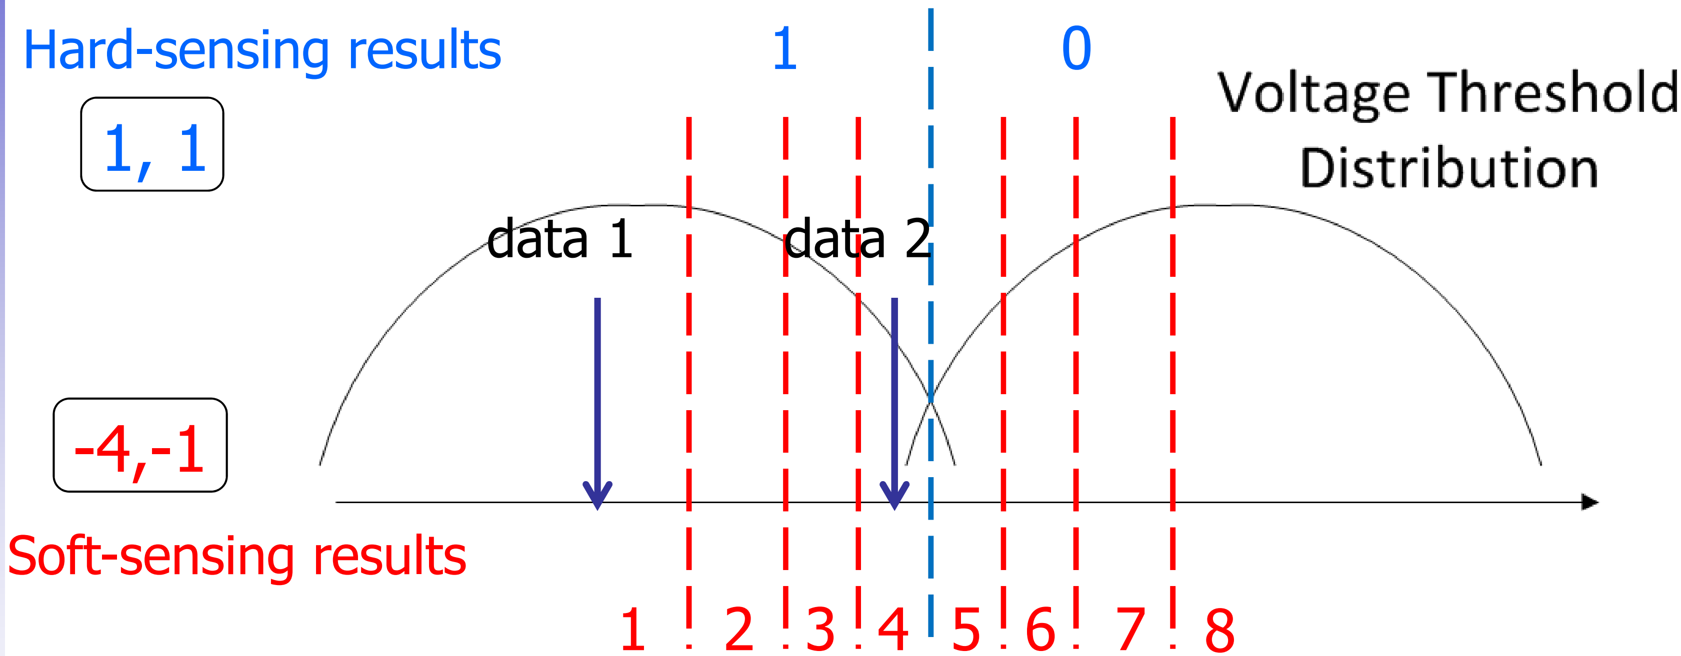

Quantization	1 bit	2 bits	3 bits
Times of Sensing	1	3	7

Soft-decision decoding suffers from a huge latency

LDPC Code and Soft Sensing

- Large latency of soft-decision sensing and corresponding data transfer to SSD controller is destructive to read performance
- A basic two-step read strategy
  - Hard-decision decoding still works in most of the times

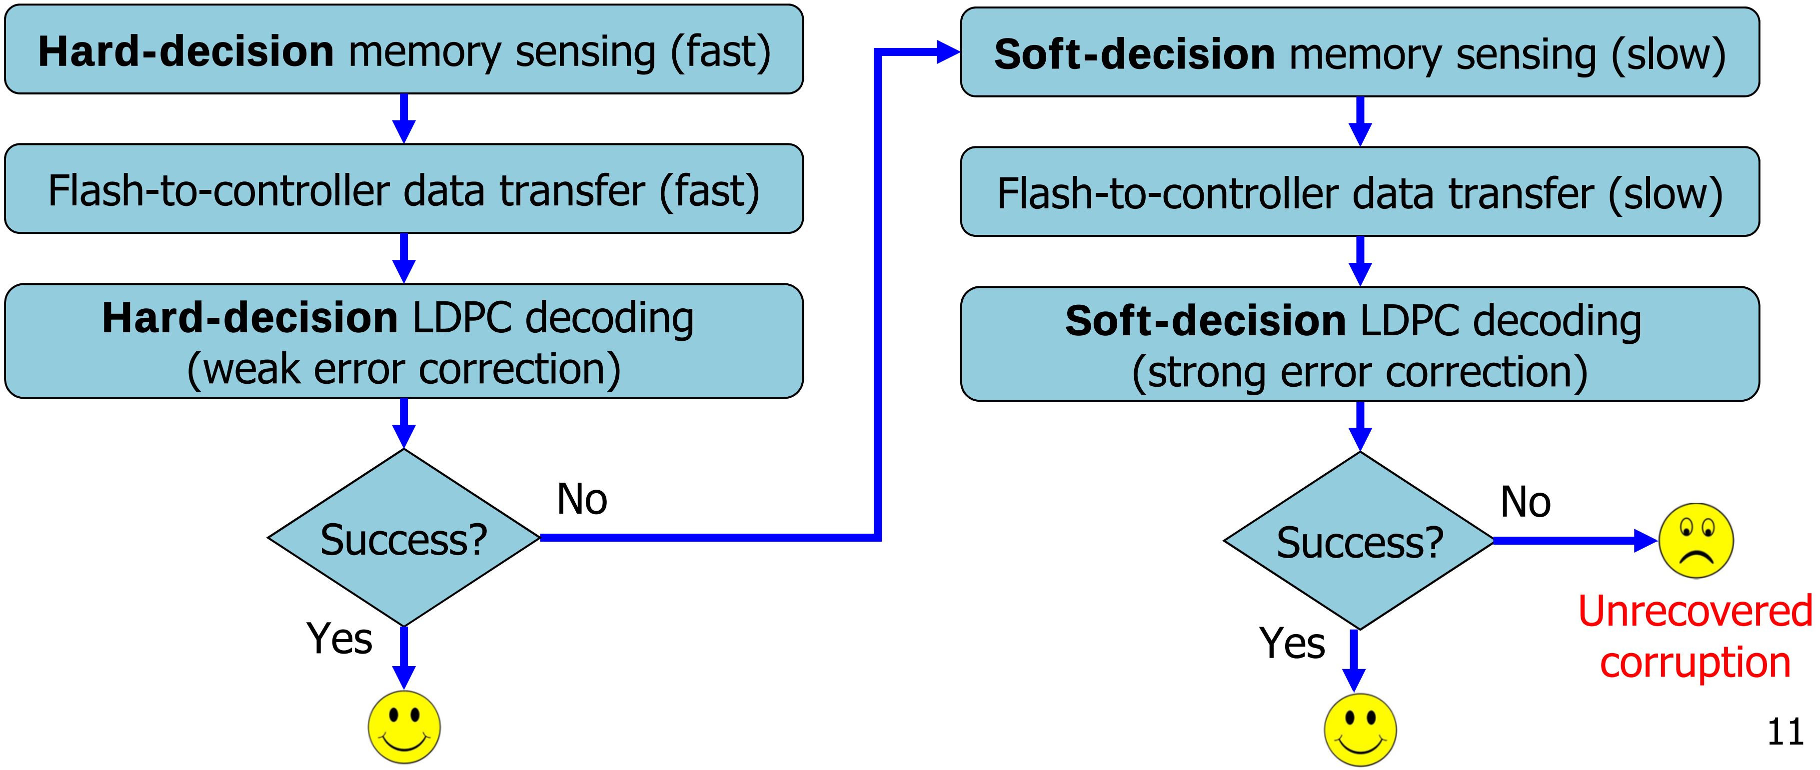

LDPC and Soft Sensing

- Trace-based simulation shows that further improvement is necessary

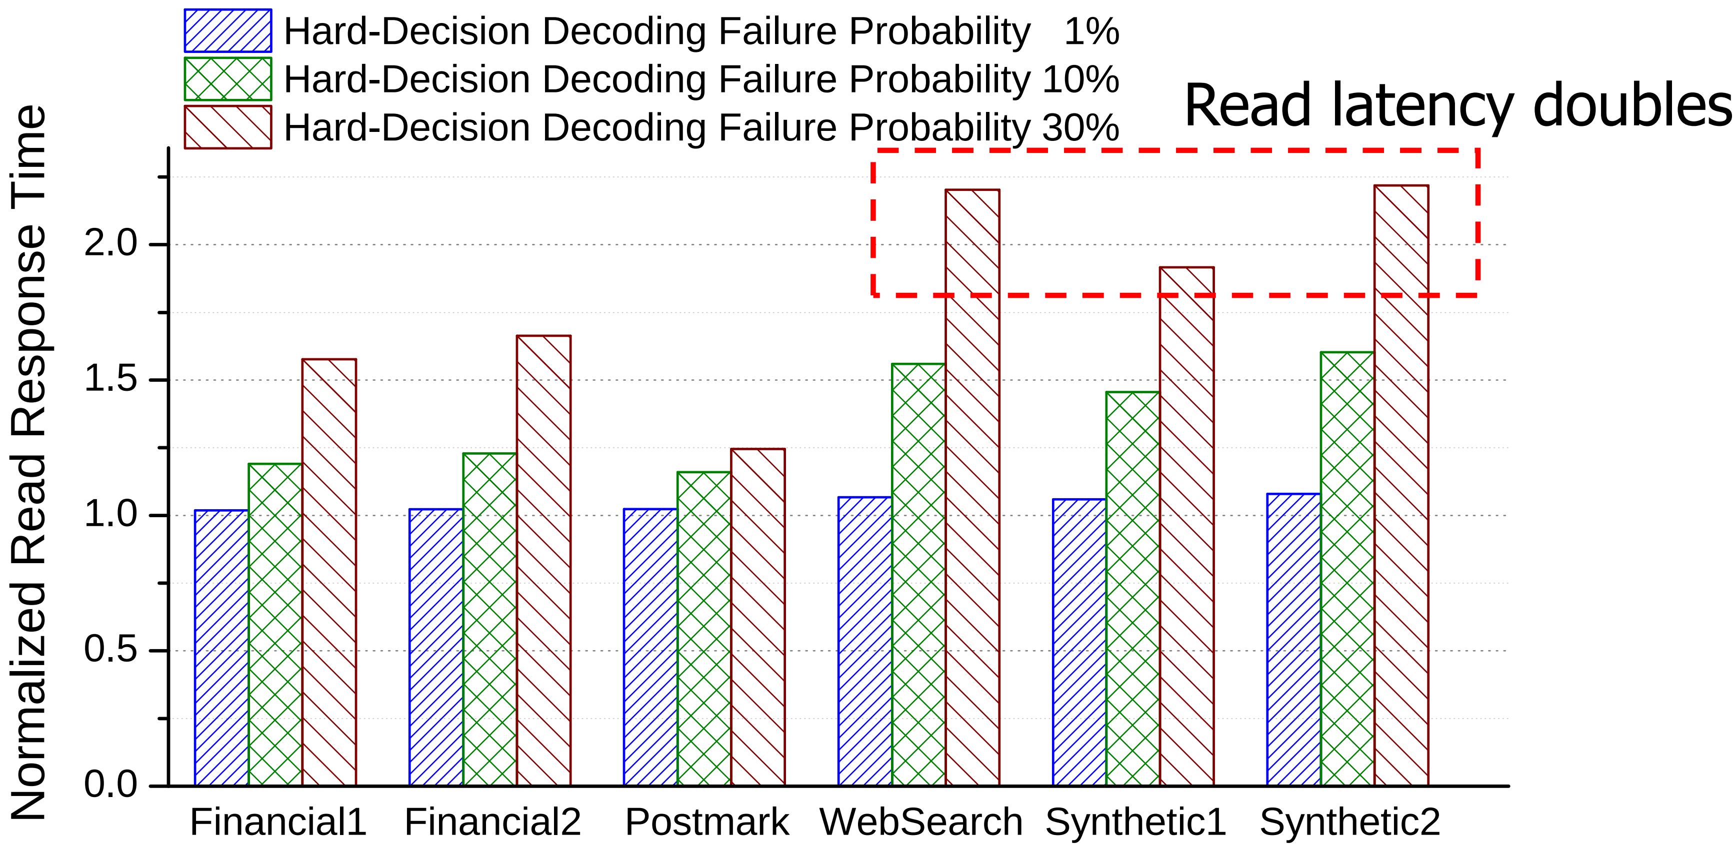

Proposed Techniques

- Two orthogonal aspects to improve read performance
  - Minimize the latency of soft-decision sensing and data transfer
  - Minimize the unnecessary number of high-precision soft-decision sensing

- Our approaches
  - Look-Ahead Memory Sensing
  - Fine-Grained Progressive Sensing and Decoding
  - Data Placement Interleaving

Look-ahead Memory Sensing

- Start soft-decision sensing in advance

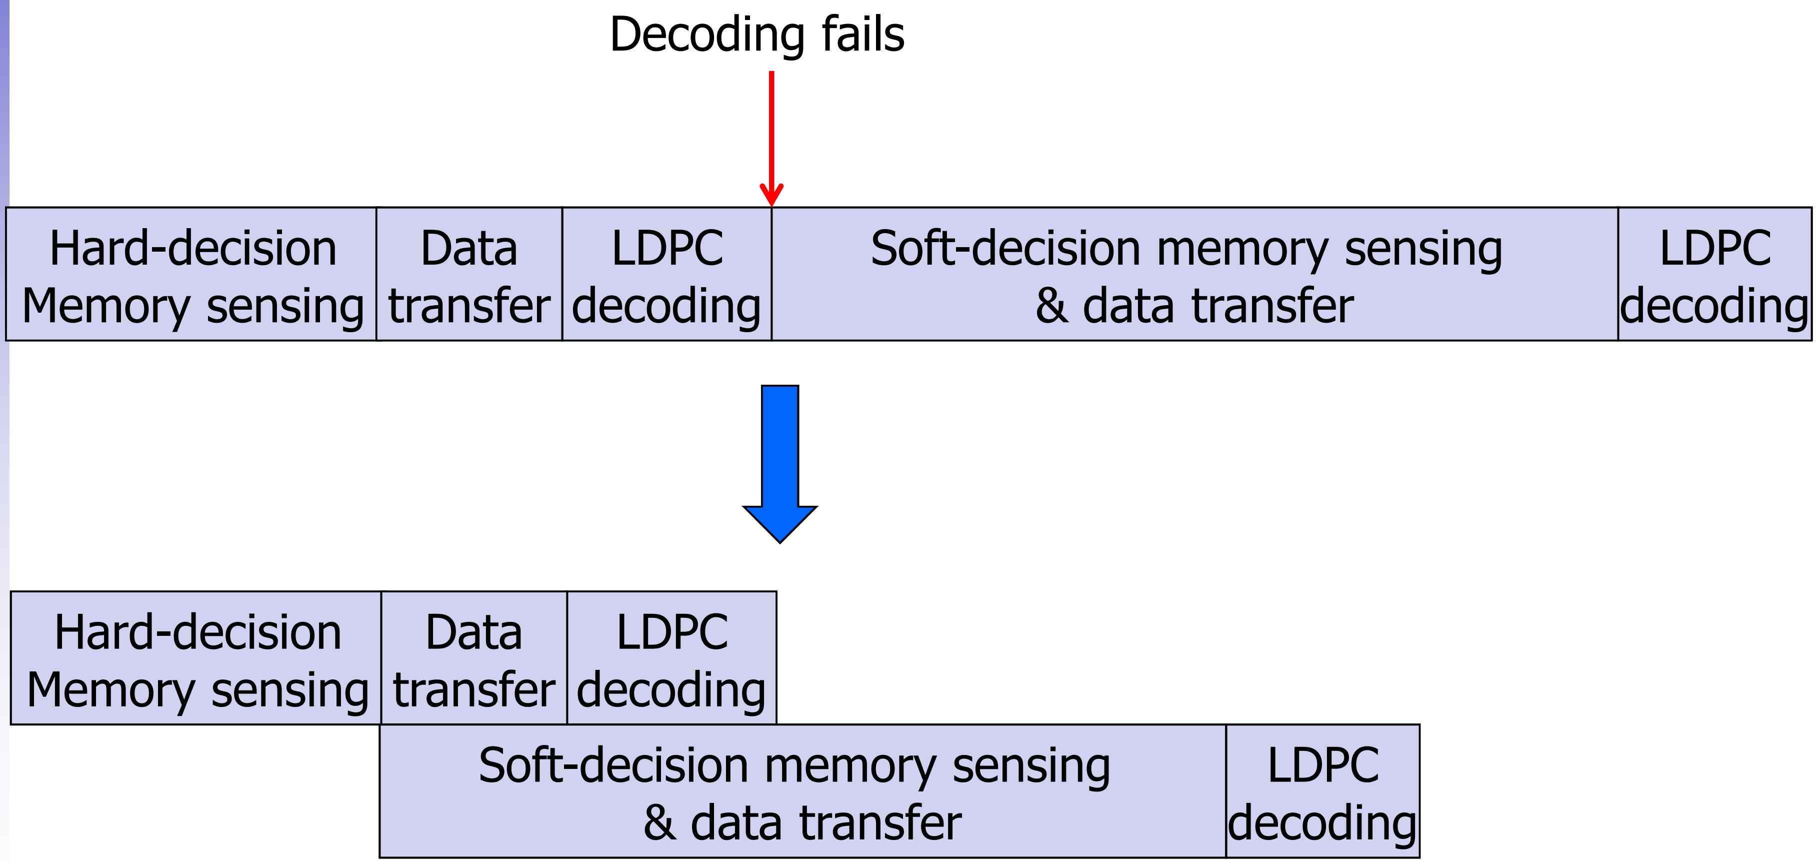

Fine-Grained Progressive Sensing and Decoding

- Why not exploit the trade-off space between latency and error correction capability

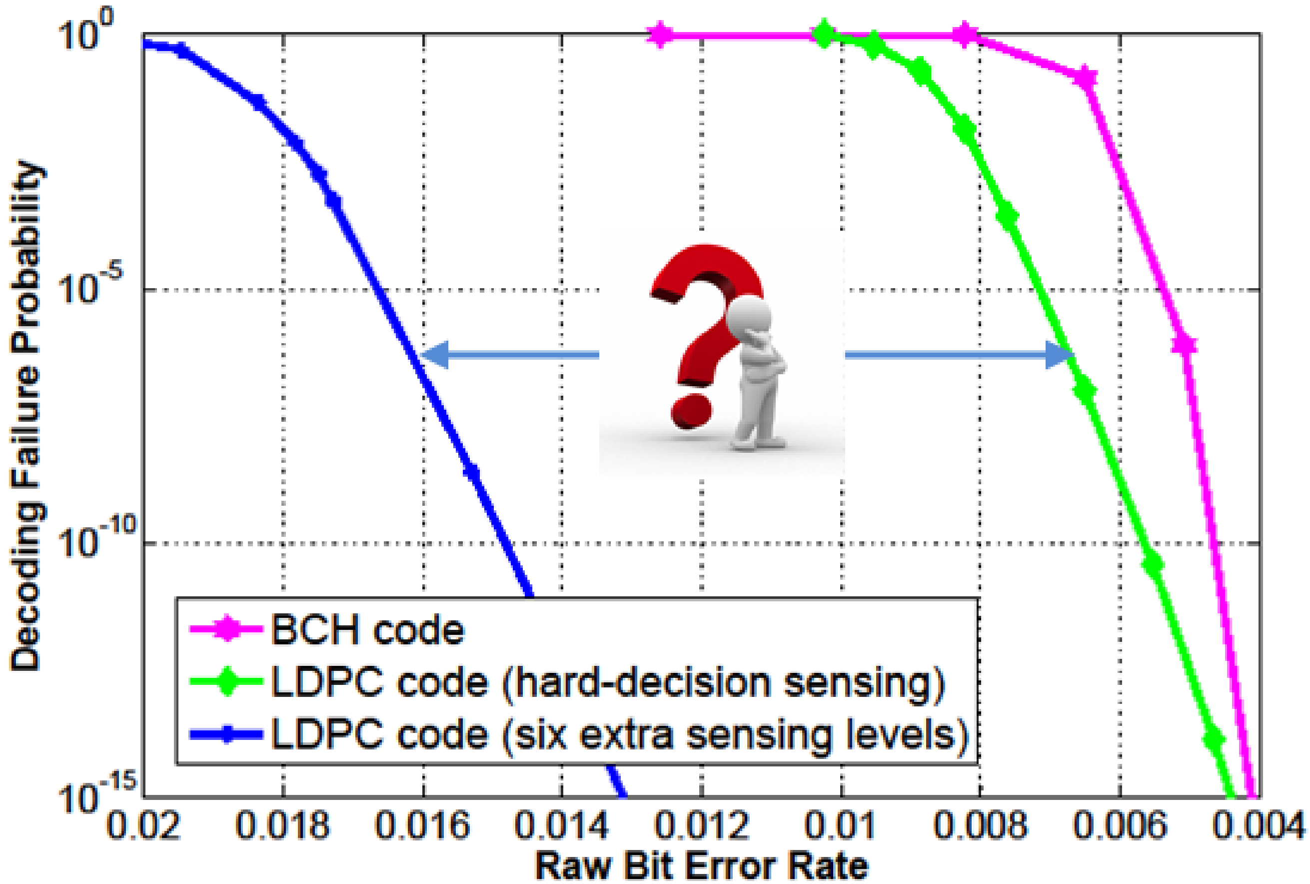

Fine-Grained Progressive Sensing and Decoding

- Fully exploit LDPC's error correction capability

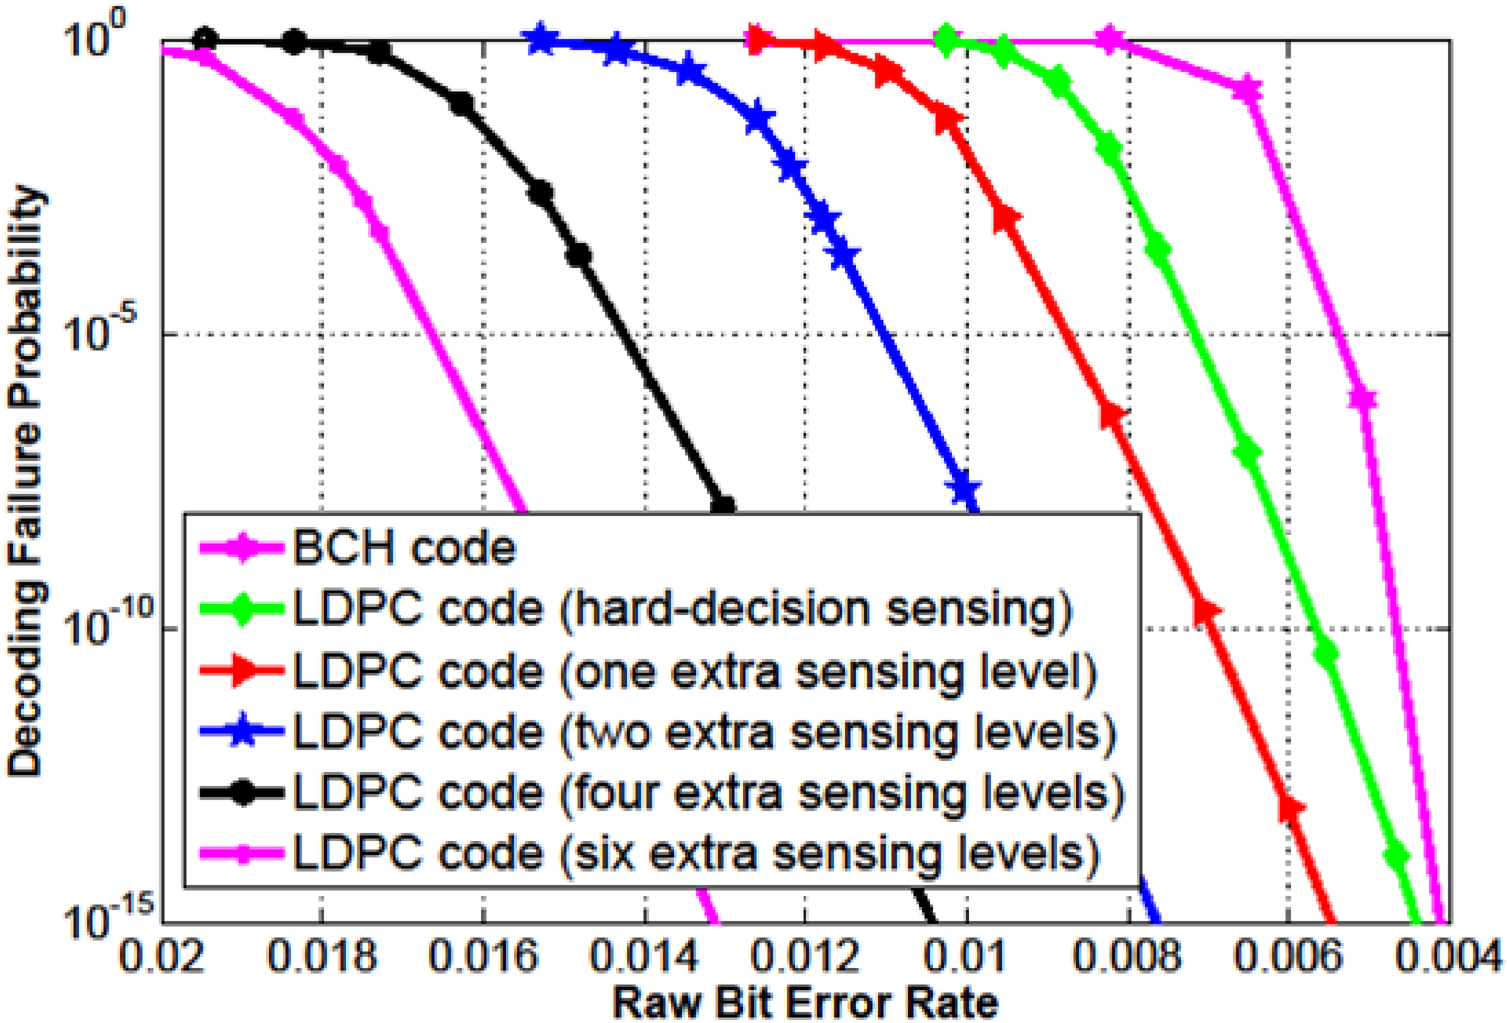

Fine-Grained Progressive Sensing and Decoding

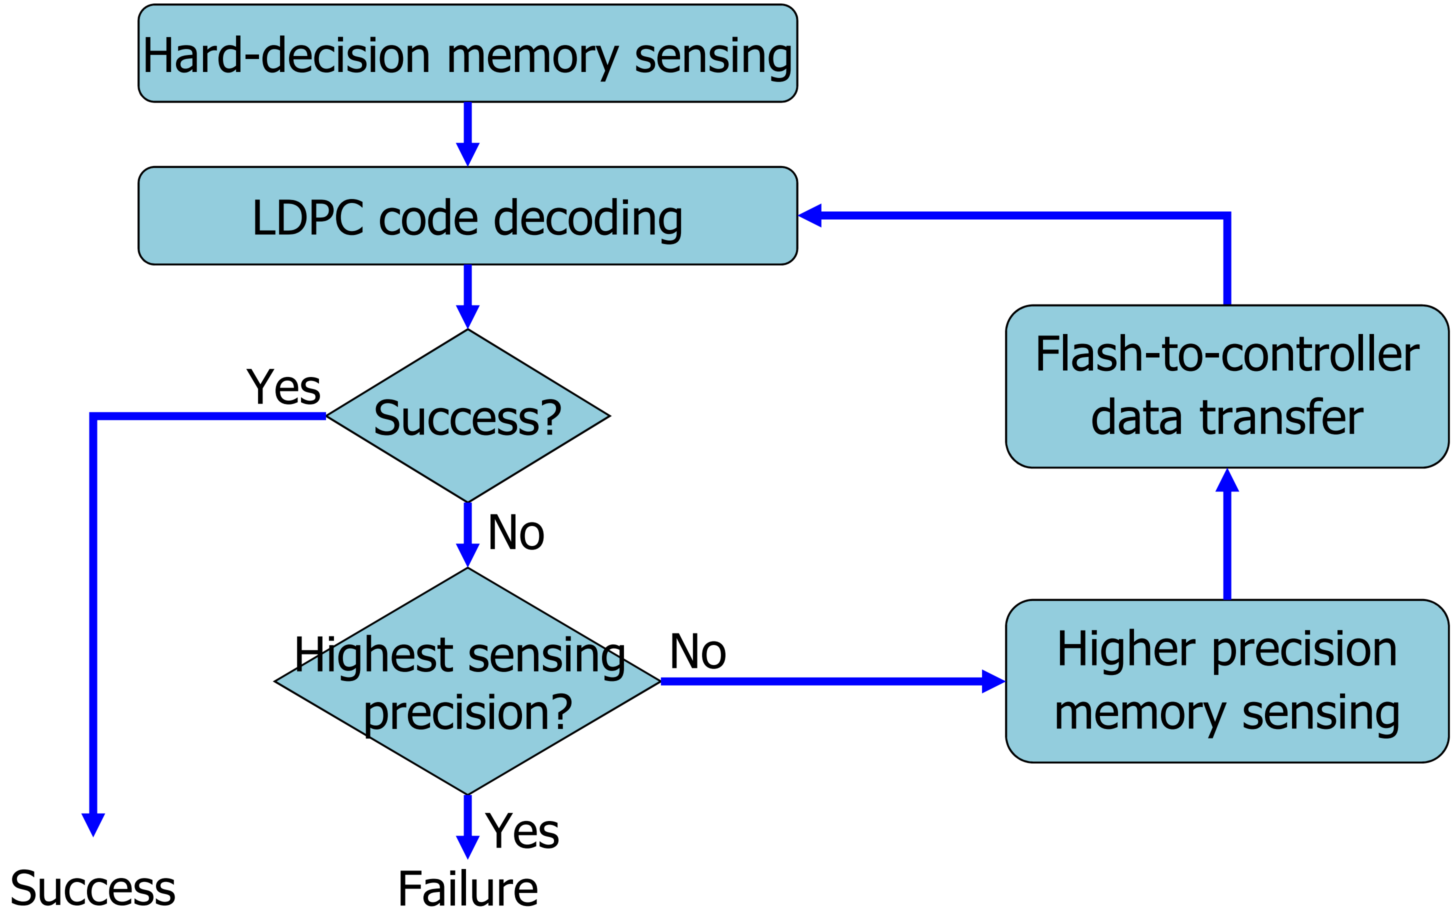

Can be combined with look-ahead memory sensing design strategy

Data Placement Interleaving

- Noticeable reliability variance among different NAND flash memory chips

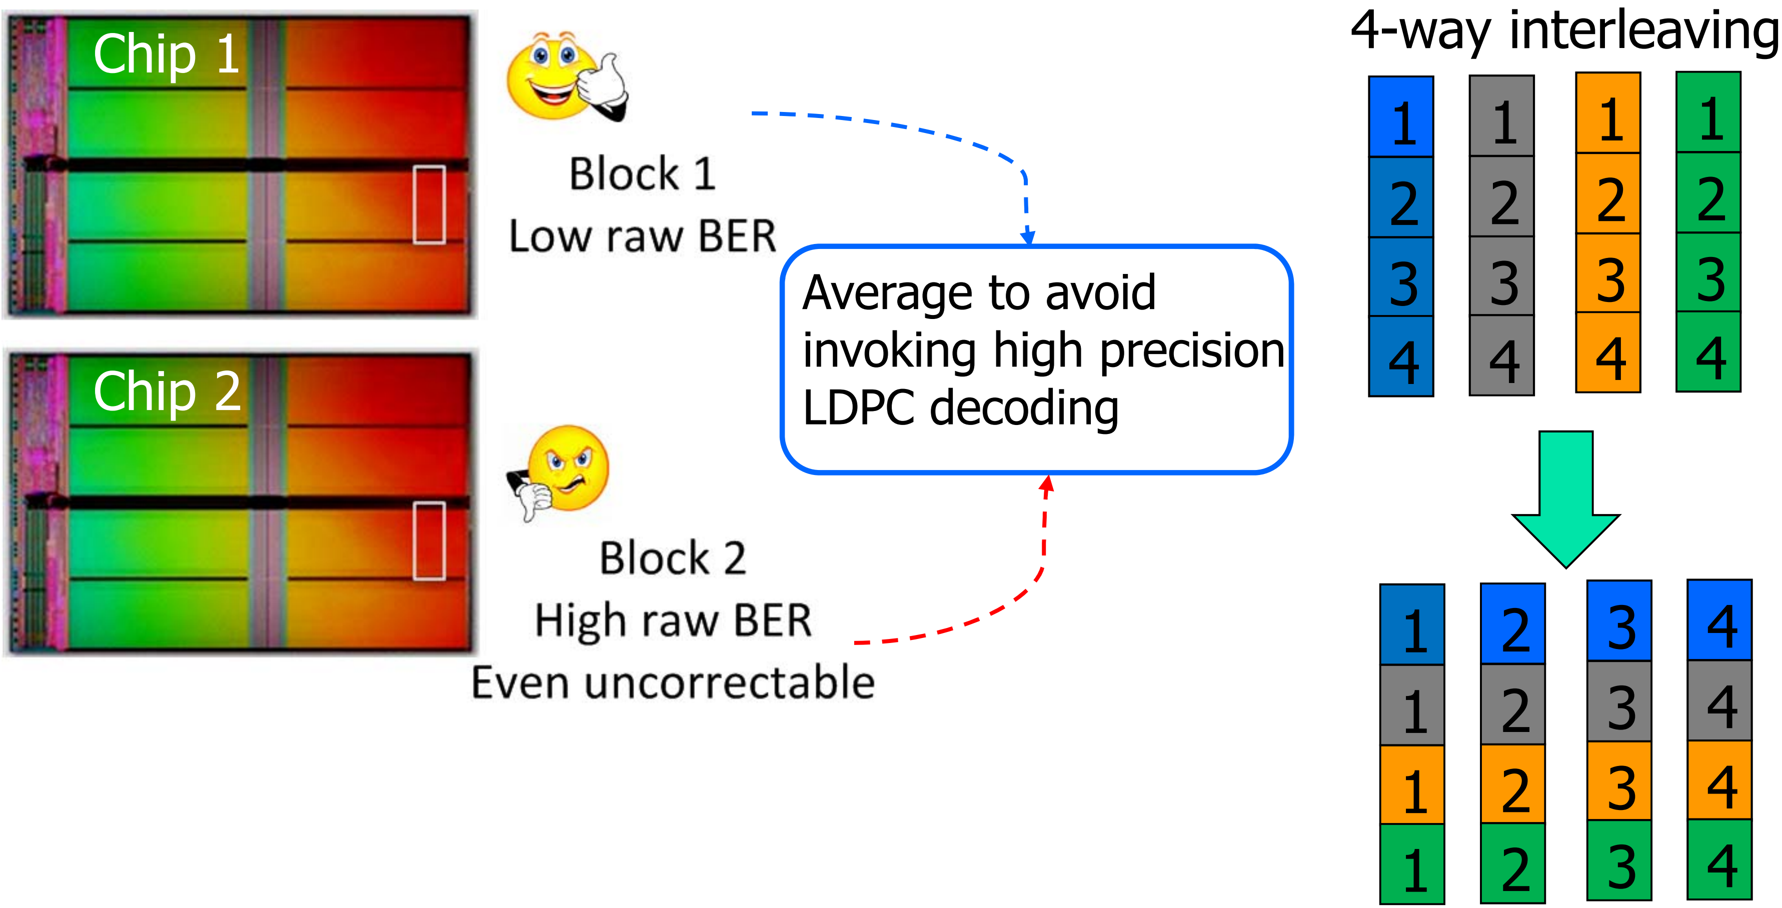

Experiment Setup and Preparation

- SSD Simulator: SSD module in DiskSim
- Workload traces: Fiancial1&2, Postmark, WebSearch, Synthetic workload1&2
- MLC NAND flash memory chip configuration
  - 4KB per page, 64 pages per block, 2048 blocks per plane, 4 planes per die, and 2 dies per chip.
  - 8 chips X 8 channels
  - 200 MBps chip I/O bandwidth
- Separate parameters for upper/lower pages
  - Read/write/erasure latency
  - Page raw bit error rate

Experiment Setup and Preparation

- Our measurement results are based on 25 nm MLC NAND flash chips

	Read	**Write**
Upper page	55 µs	1.45 ms
Lower page	41 µs	121 µs

- Latency of sensing one extra level is set to 14 µs

- 10k PE cycles and 1 month retention time

Experiment Setup and Preparation

- Experiment flow to get the BER statistics

Continuous P/E cycling over a short time

↓

Baking to emulate wear-out recovery

↓

Program once

↓

Baking to emulate data retention

↓

Bit error rate statistics

Use Arrhenius law to determine the baking time

$$t_{rc} = A \cdot t_{cyc} \cdot e^{E_A\left(\frac{1}{k \cdot T_{rc}} - \frac{1}{k \cdot T_{cyc}}\right)}$$

$$t_{rt} = t_{cyc} \cdot e^{E_A\left(\frac{1}{k \cdot T_{rt}} - \frac{1}{k \cdot T_{cyc}}\right)}$$

Experiment Setup and Preparation

- Page raw BER distribution

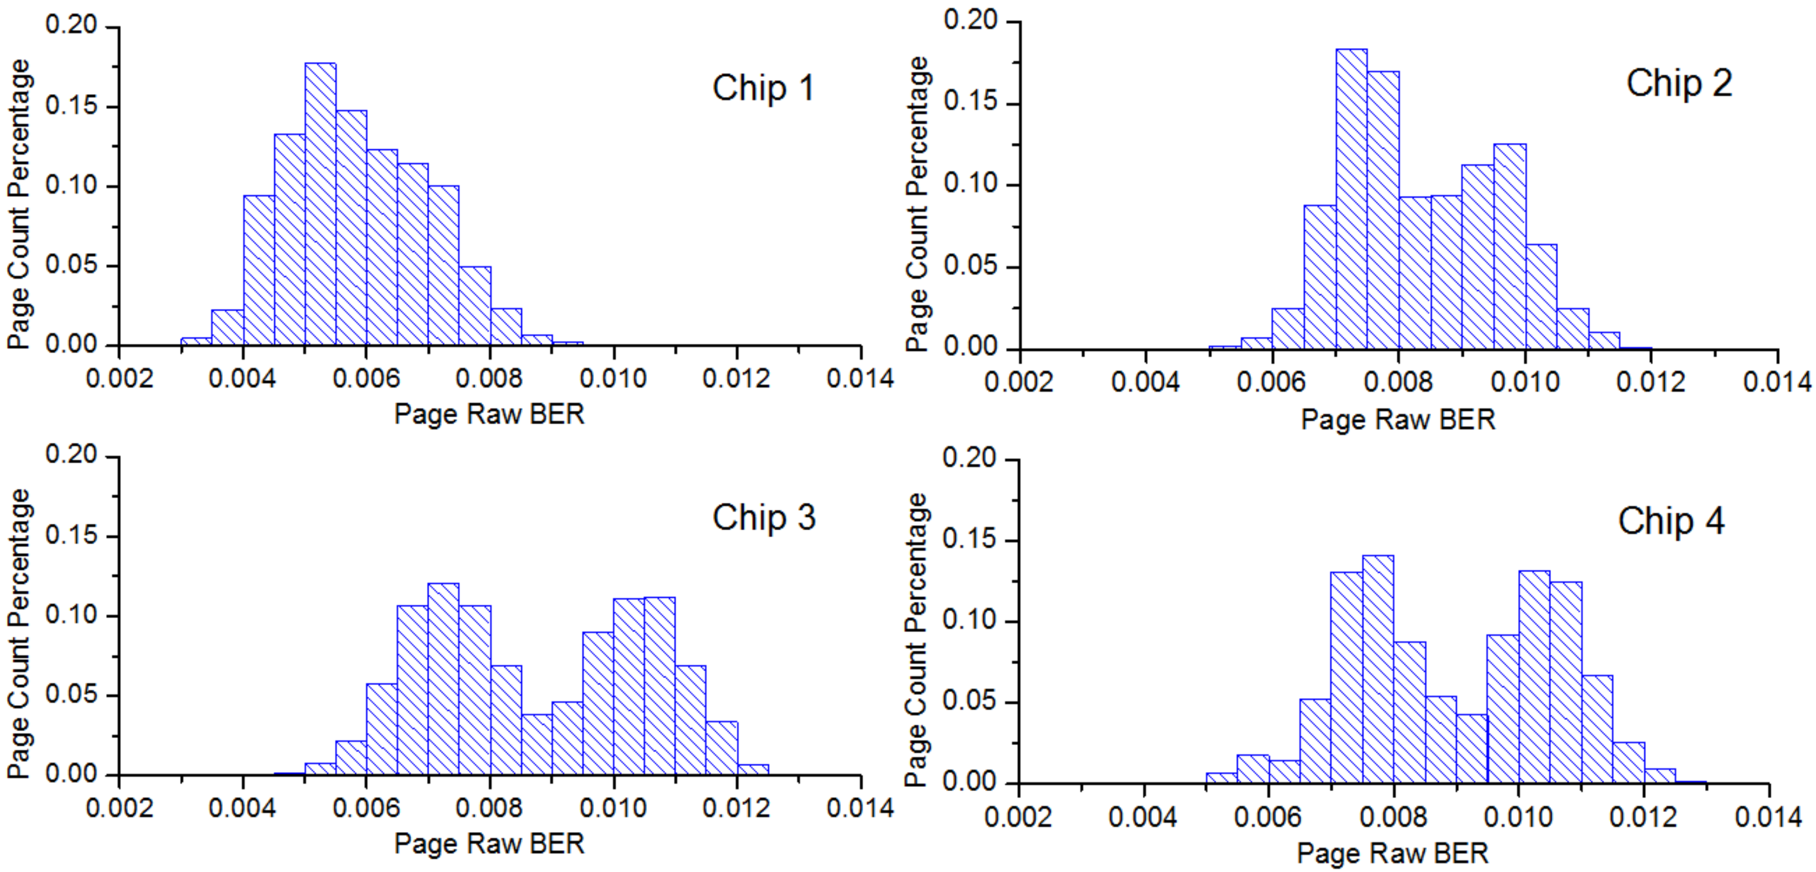

Experimental Results

- Baseline: two-step sensing and decoding strategy

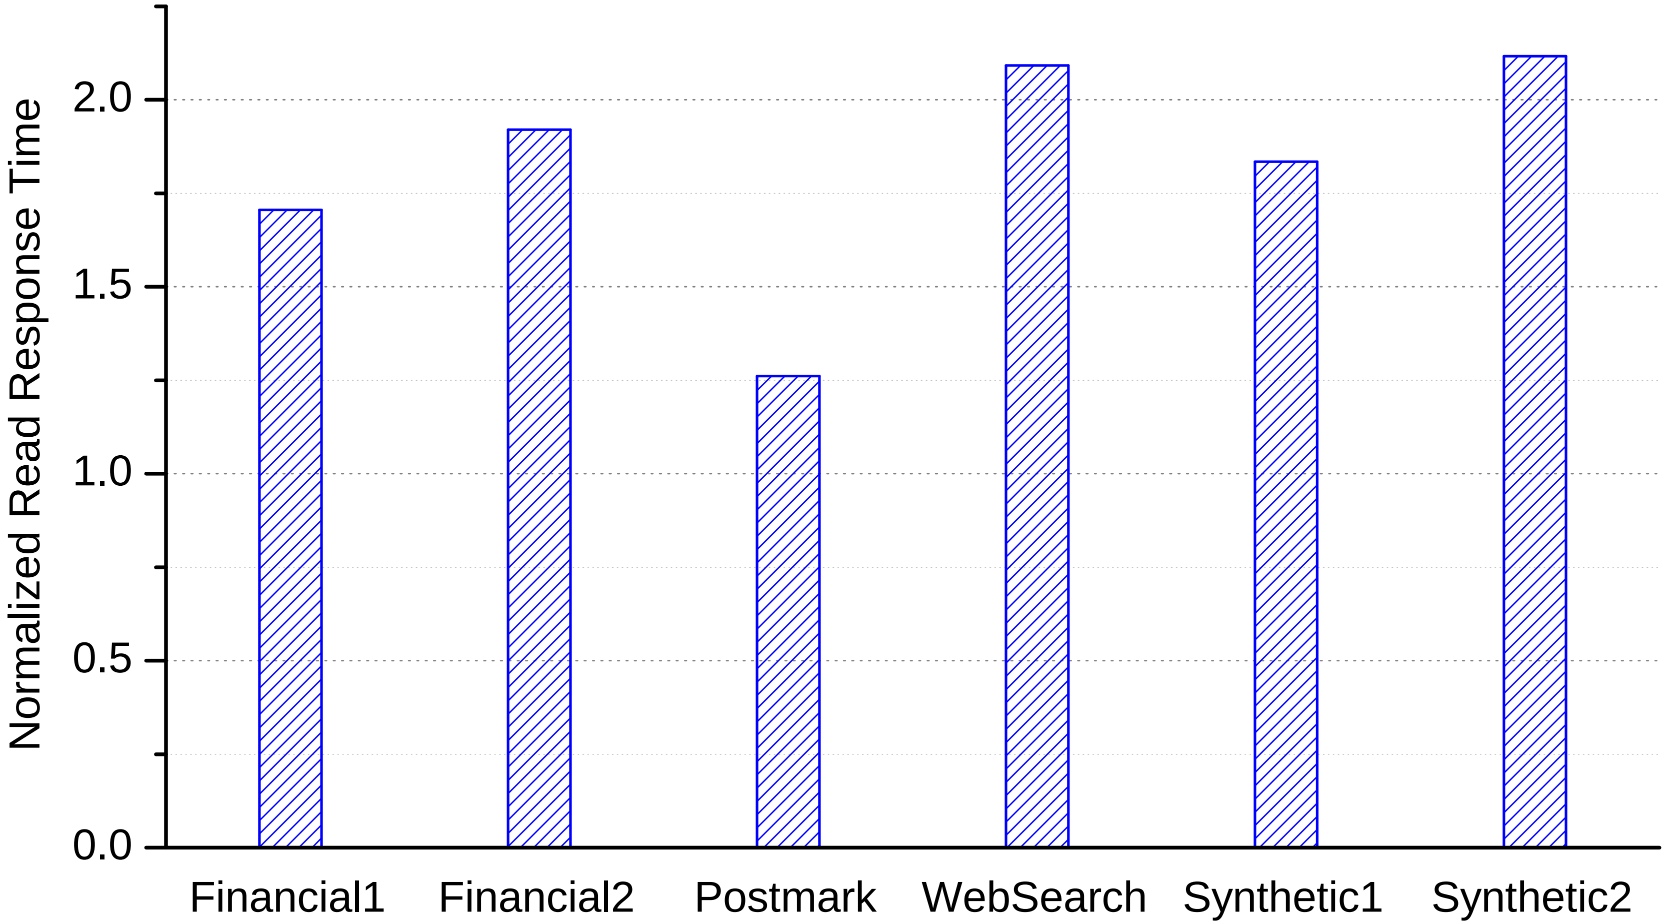

Experimental Results

- Look-ahead memory sensing

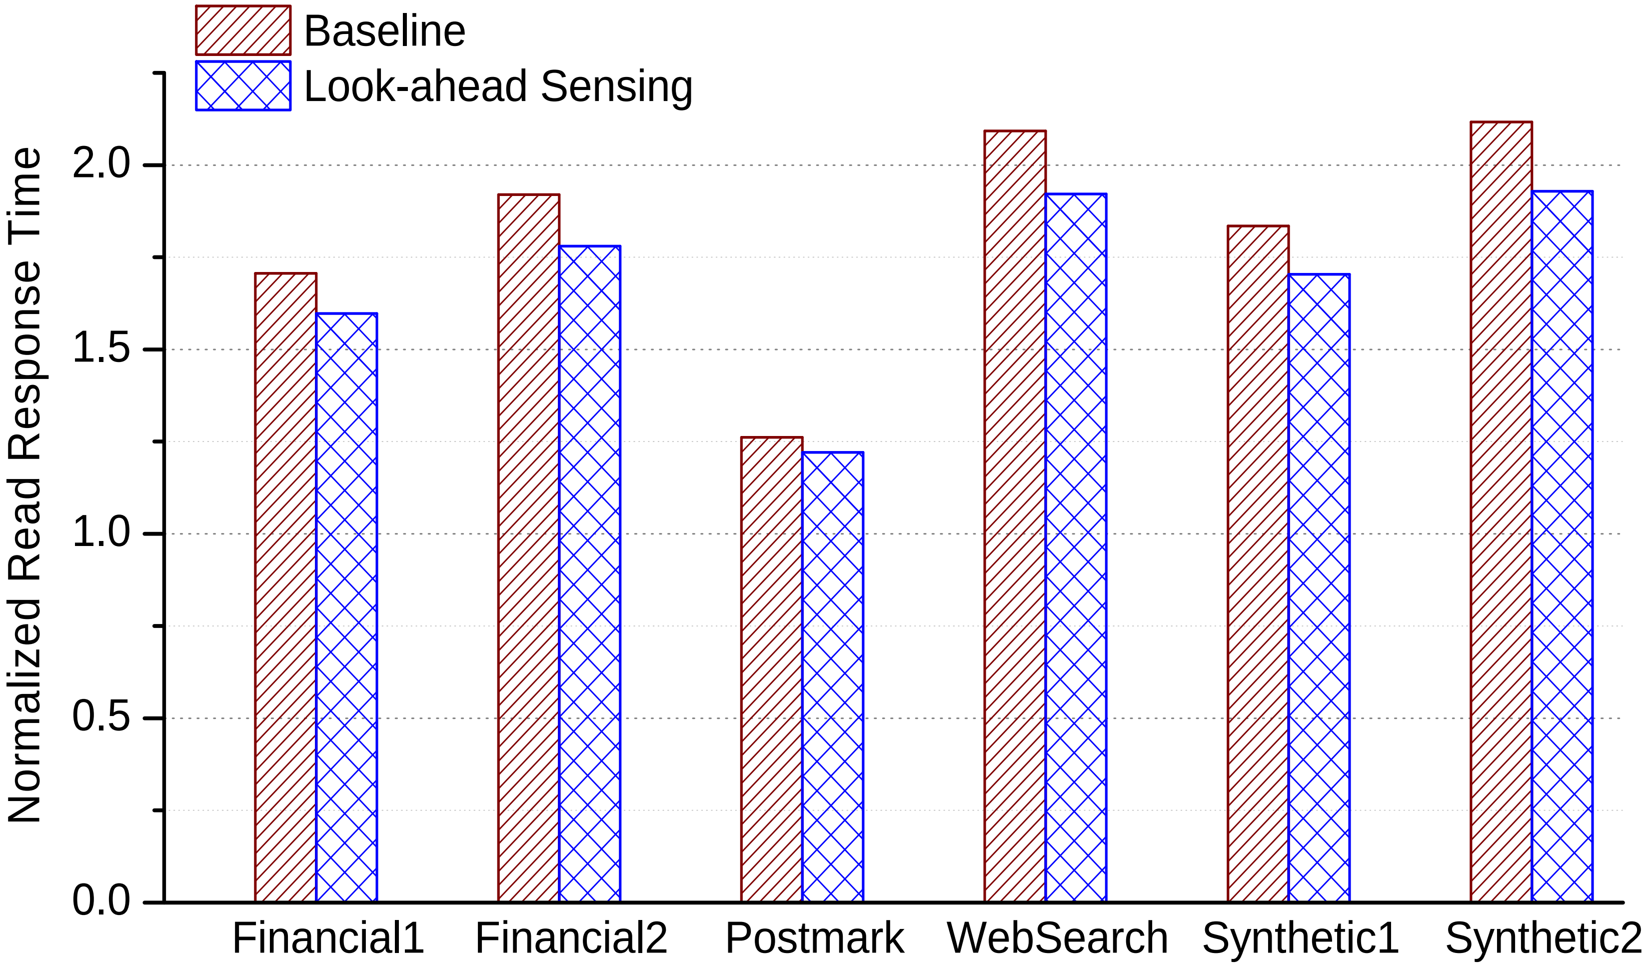

Experimental Results

- Fine-grained progressive sensing

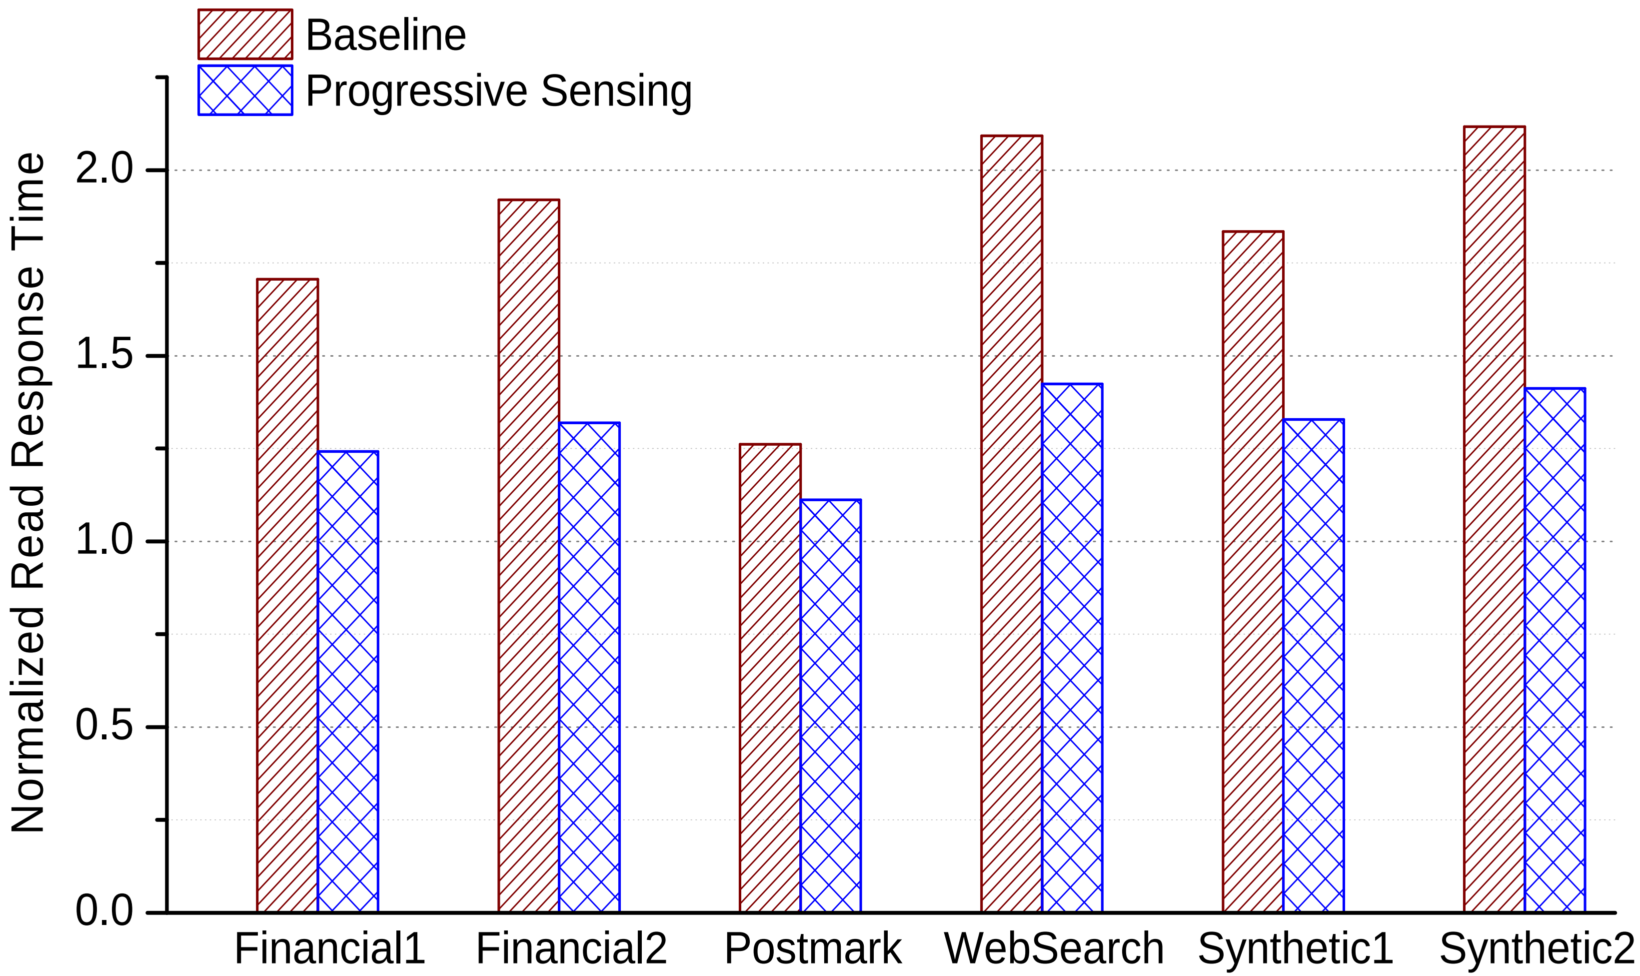

Experimental Results

- Combined look-ahead and progressive sensing

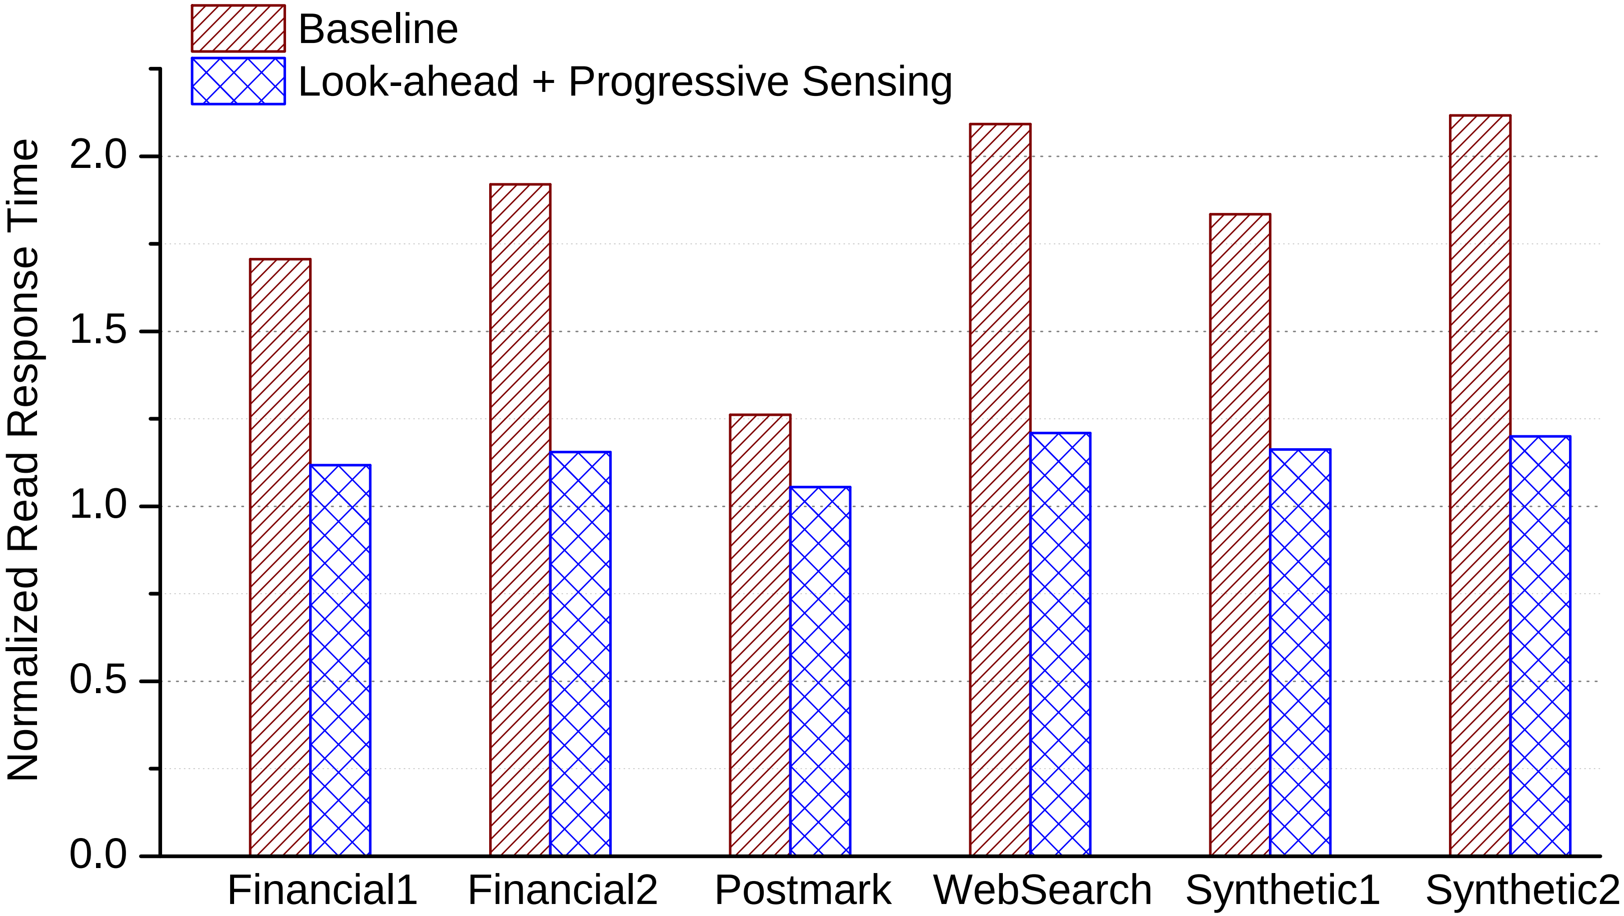

Experimental Results

- Data placement interleaving

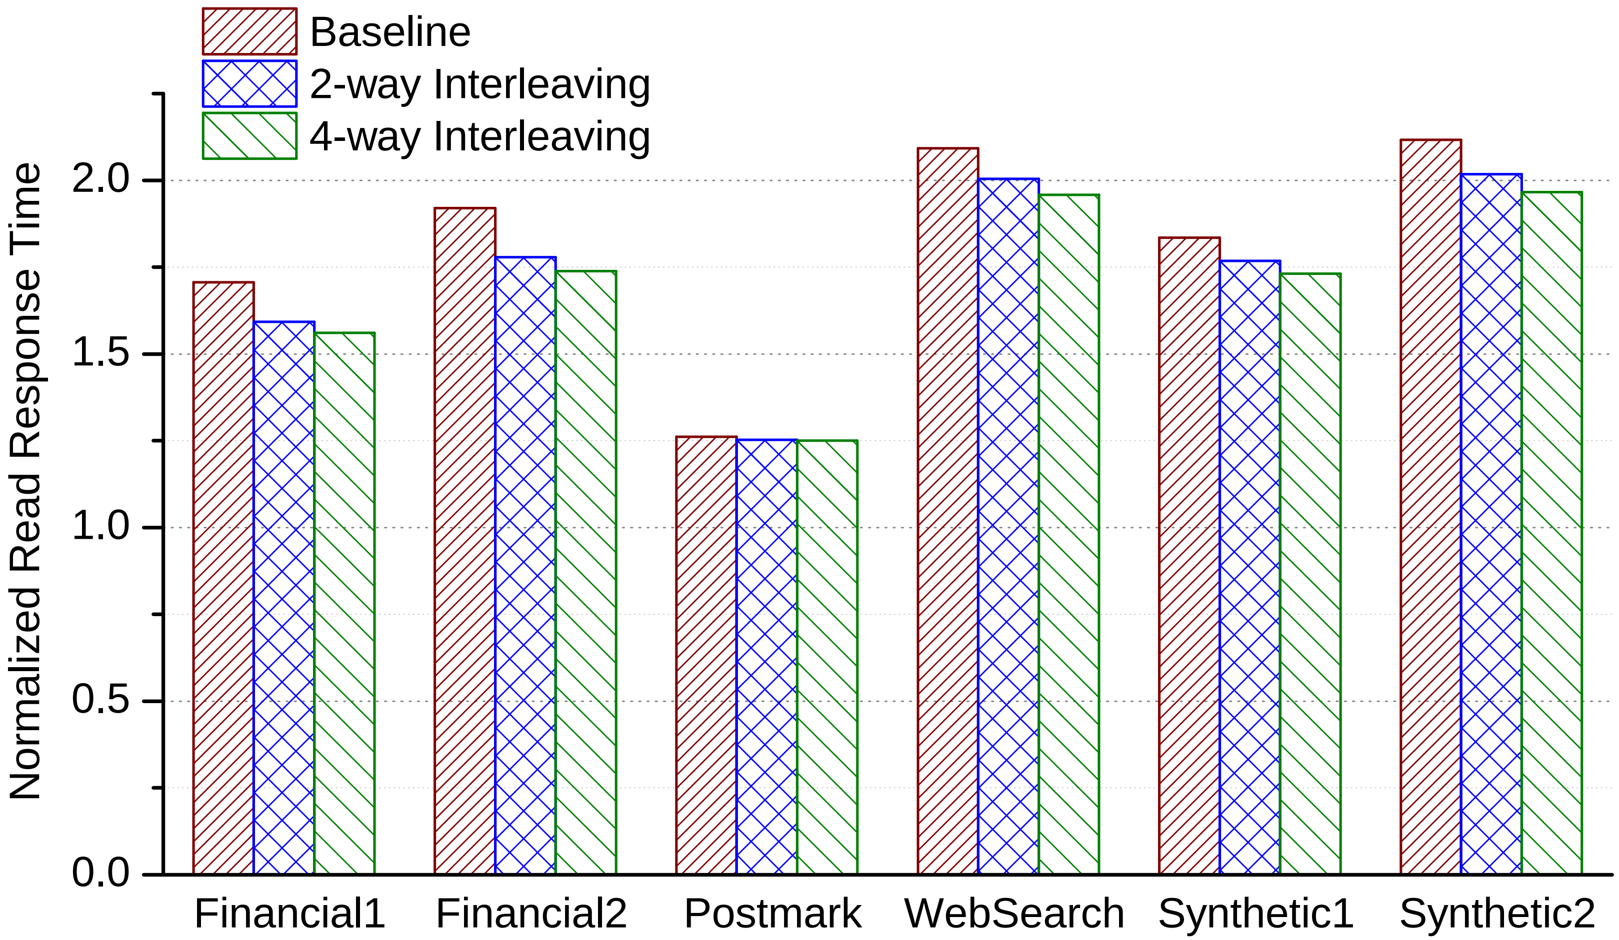

Experimental Results

- Aggregated read response time latency reduction

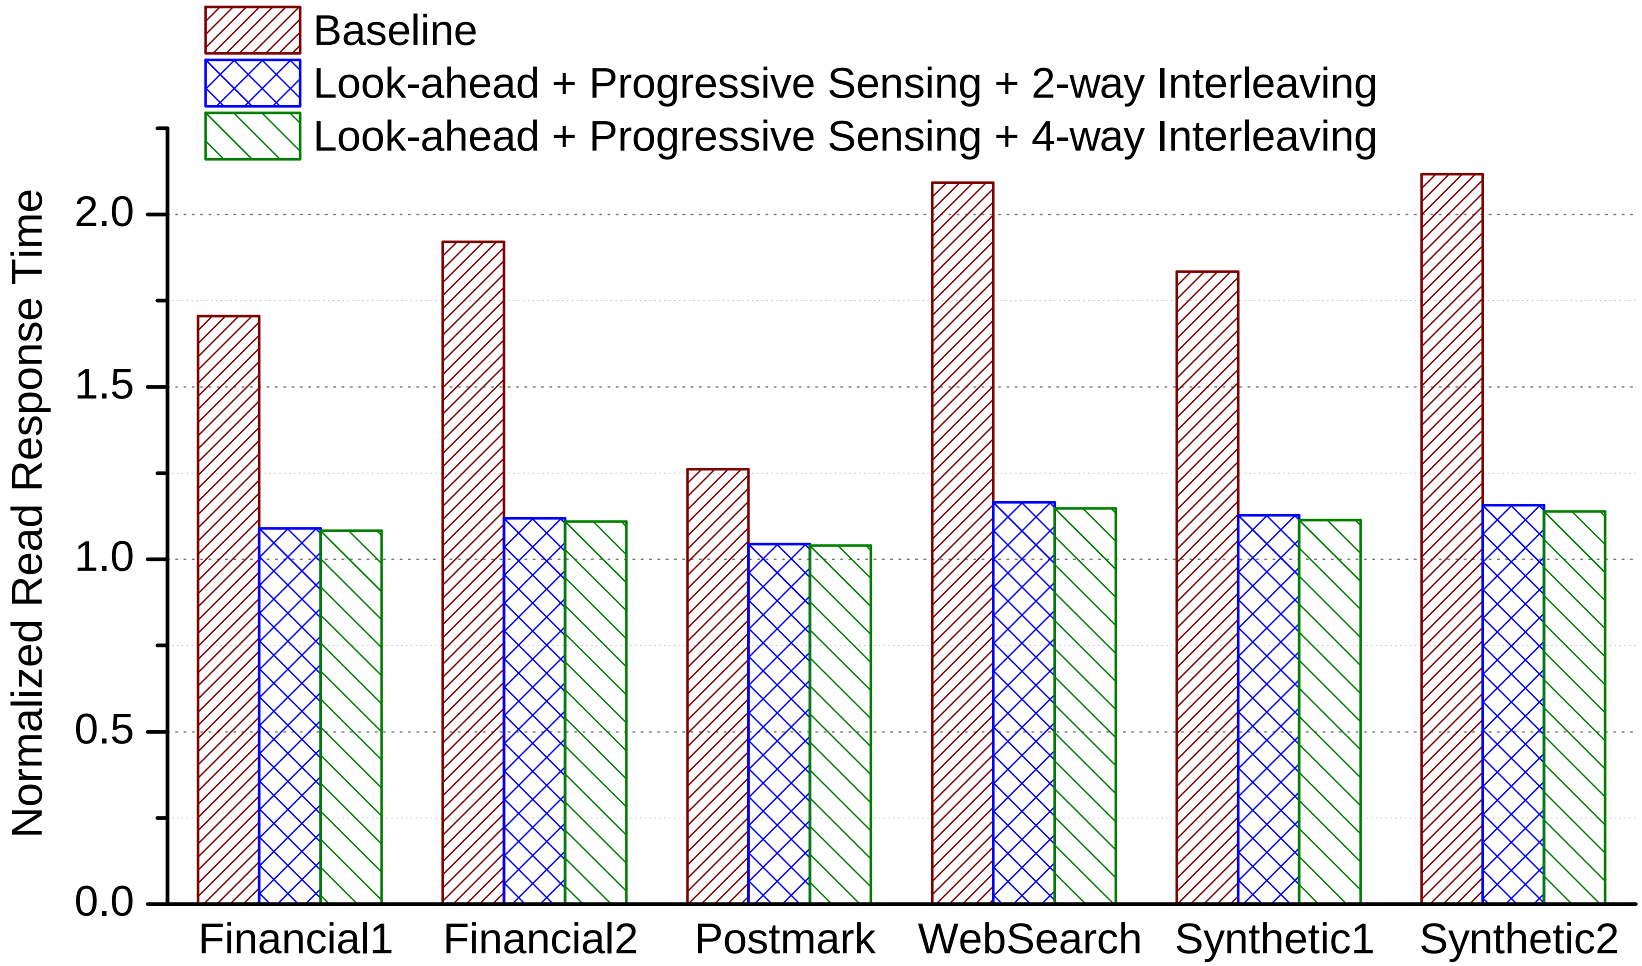

Experimental Results

- Overhead
  - Higher read energy consumption
  - Complicating the controller design
  - Write amplification caused by interleaving

Conclusion

- Increased noise in NAND flash memory caused by technology scaling is a main reason for increasingly high error rates in SSD
- Conventional ECC, such as BCH, does not have sufficient ability to make error code corrections
- LDPC is a strong ECC candidate for future SSDs to address its reliability issues under high noises
- We proposed three techniques to make LDPC work effectively in SSDs
- With LDPC-in-SSD, SSD can continue to increase its capacity, retain a high reliability, and reduce its prices.