

Improving I/O Performance of Large-Page Flash Storage Systems Using Subpage-Parallel Reads

Jisung Park, Myungsook Kim, Youngsun Song, Sungjin Lee[†] and Jihong Kim

Department of Computer Science and Engineering, Seoul National University

[†]Department of Information and Communication Engineering, DGIST



Page Size Trend in NAND Flash Memory 1

- Page size has increased with the NAND capacity.
 - 256 B in 300-nm SLC → 16 KB in recent NAND devices
- Advantages of a larger page size
 - Less peripheral circuits for addressing pages: **high density**



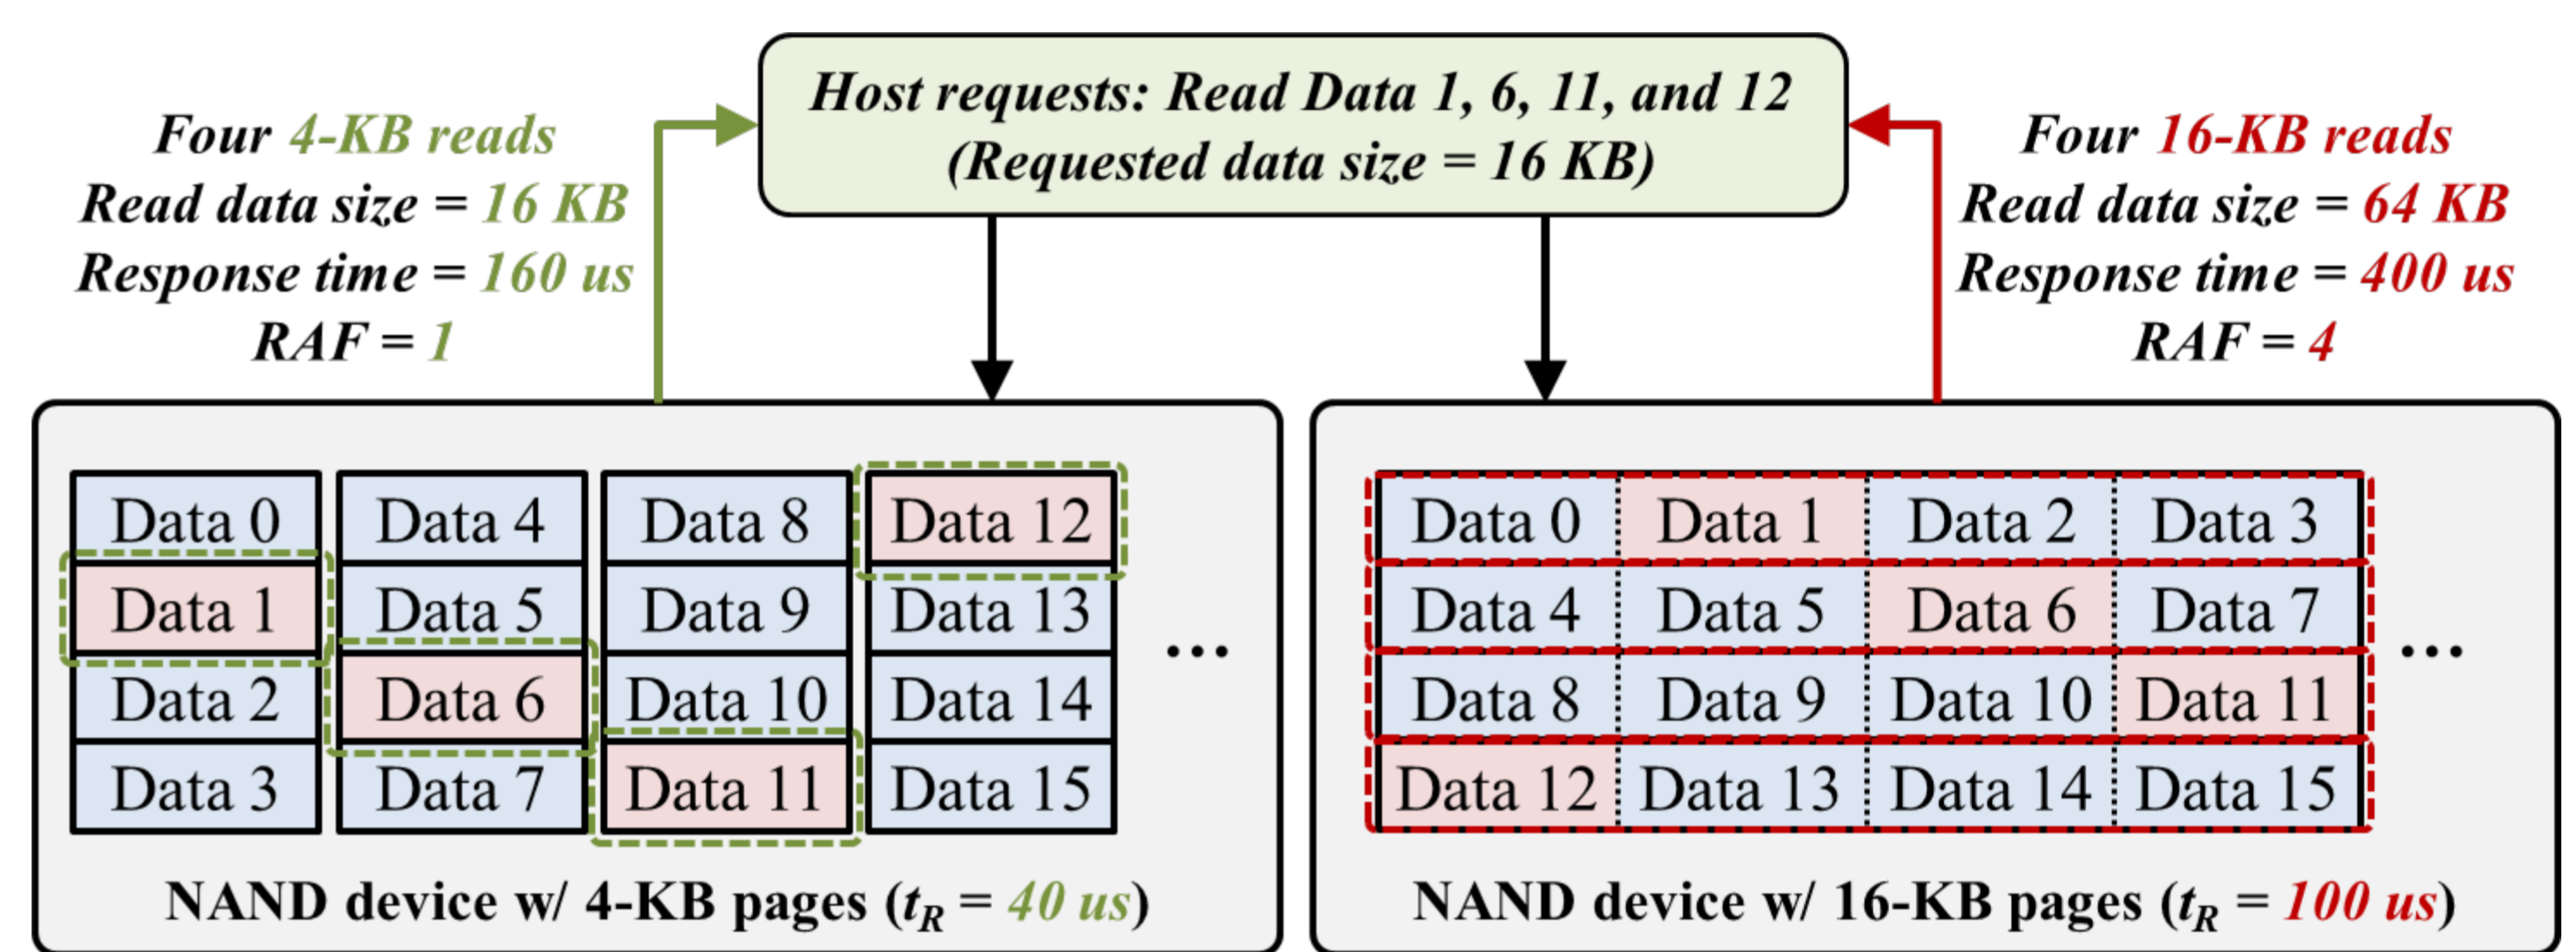
- More cells read/written at the same time: **high bandwidth**

Page Size	4 KB	16 KB
t_R	40 us	100 us
t_{PROG}	1300 us	
Read Bandwidth	100 MB/s	160 MB/s
Write Bandwidth	3 MB/s	12 MB/s

Large Page Problem: Amplified Reads 2

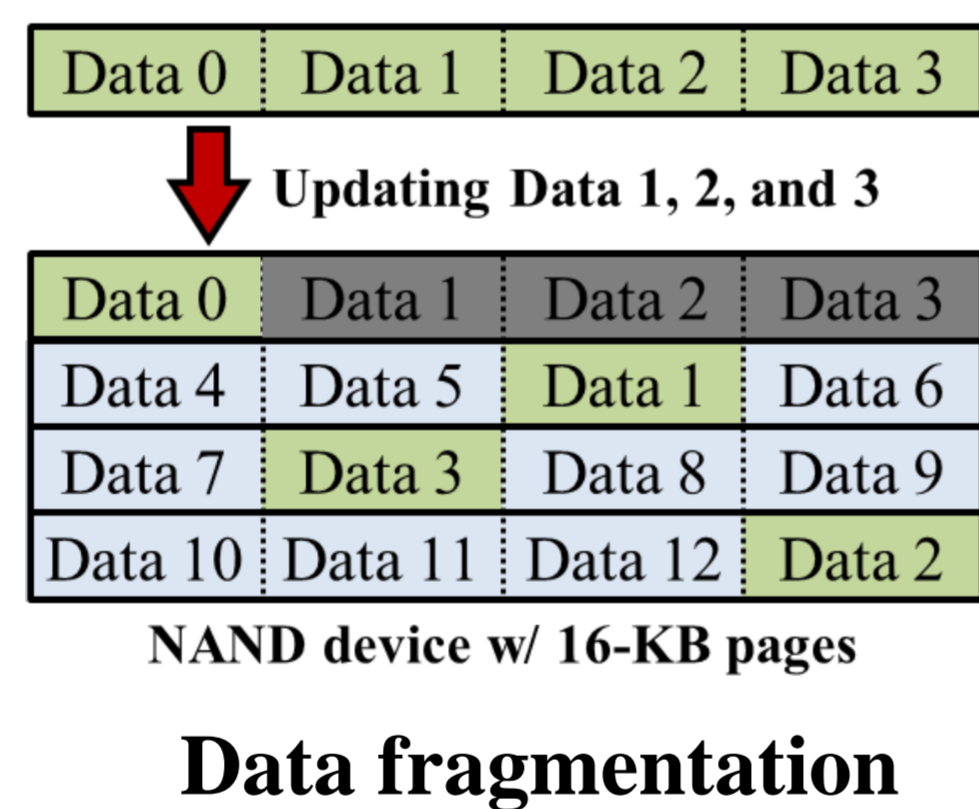
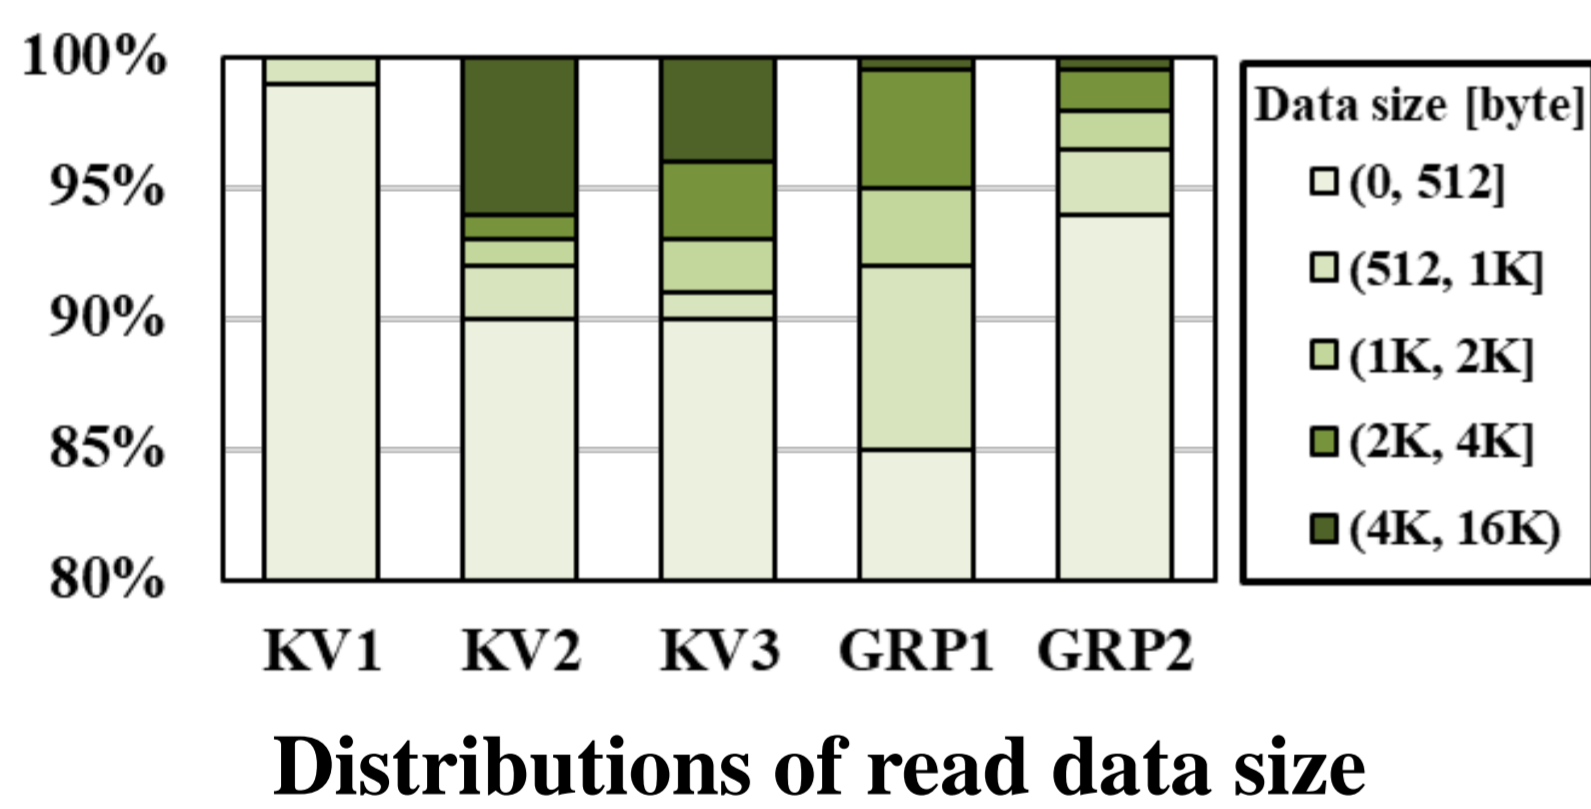
- A large portion of data is unnecessarily read.
 - When requested data are smaller than the page size
 - Increasing the read amplification factor (RAF)

$$RAF = \frac{\text{Amount of data requested to read}}{\text{Amount of data actually read}}$$



Performance Impact of Amplified Reads 3

- Frequent amplified reads in large-page NAND storages
 - Workloads dominated by small random reads
 - Data fragmentation from small updates



- Read bandwidth can be significantly wasted.

workload	KV1	KV2	GRP1	SU	Ideal
Norm. bandwidth	0.031	0.034	0.035	0.64	1
RAF	31.8	29.2	28.7	1.56	1

Supporting Size-Proportional Reads 4


- Key observation: Shorter latencies can be provided for smaller read sizes.

$$t_{READ} = t_{PRE} + \frac{t_{DISCH}}{\text{# of related bit lines}} + t_{DMA}$$

Proportional to # of related bit lines

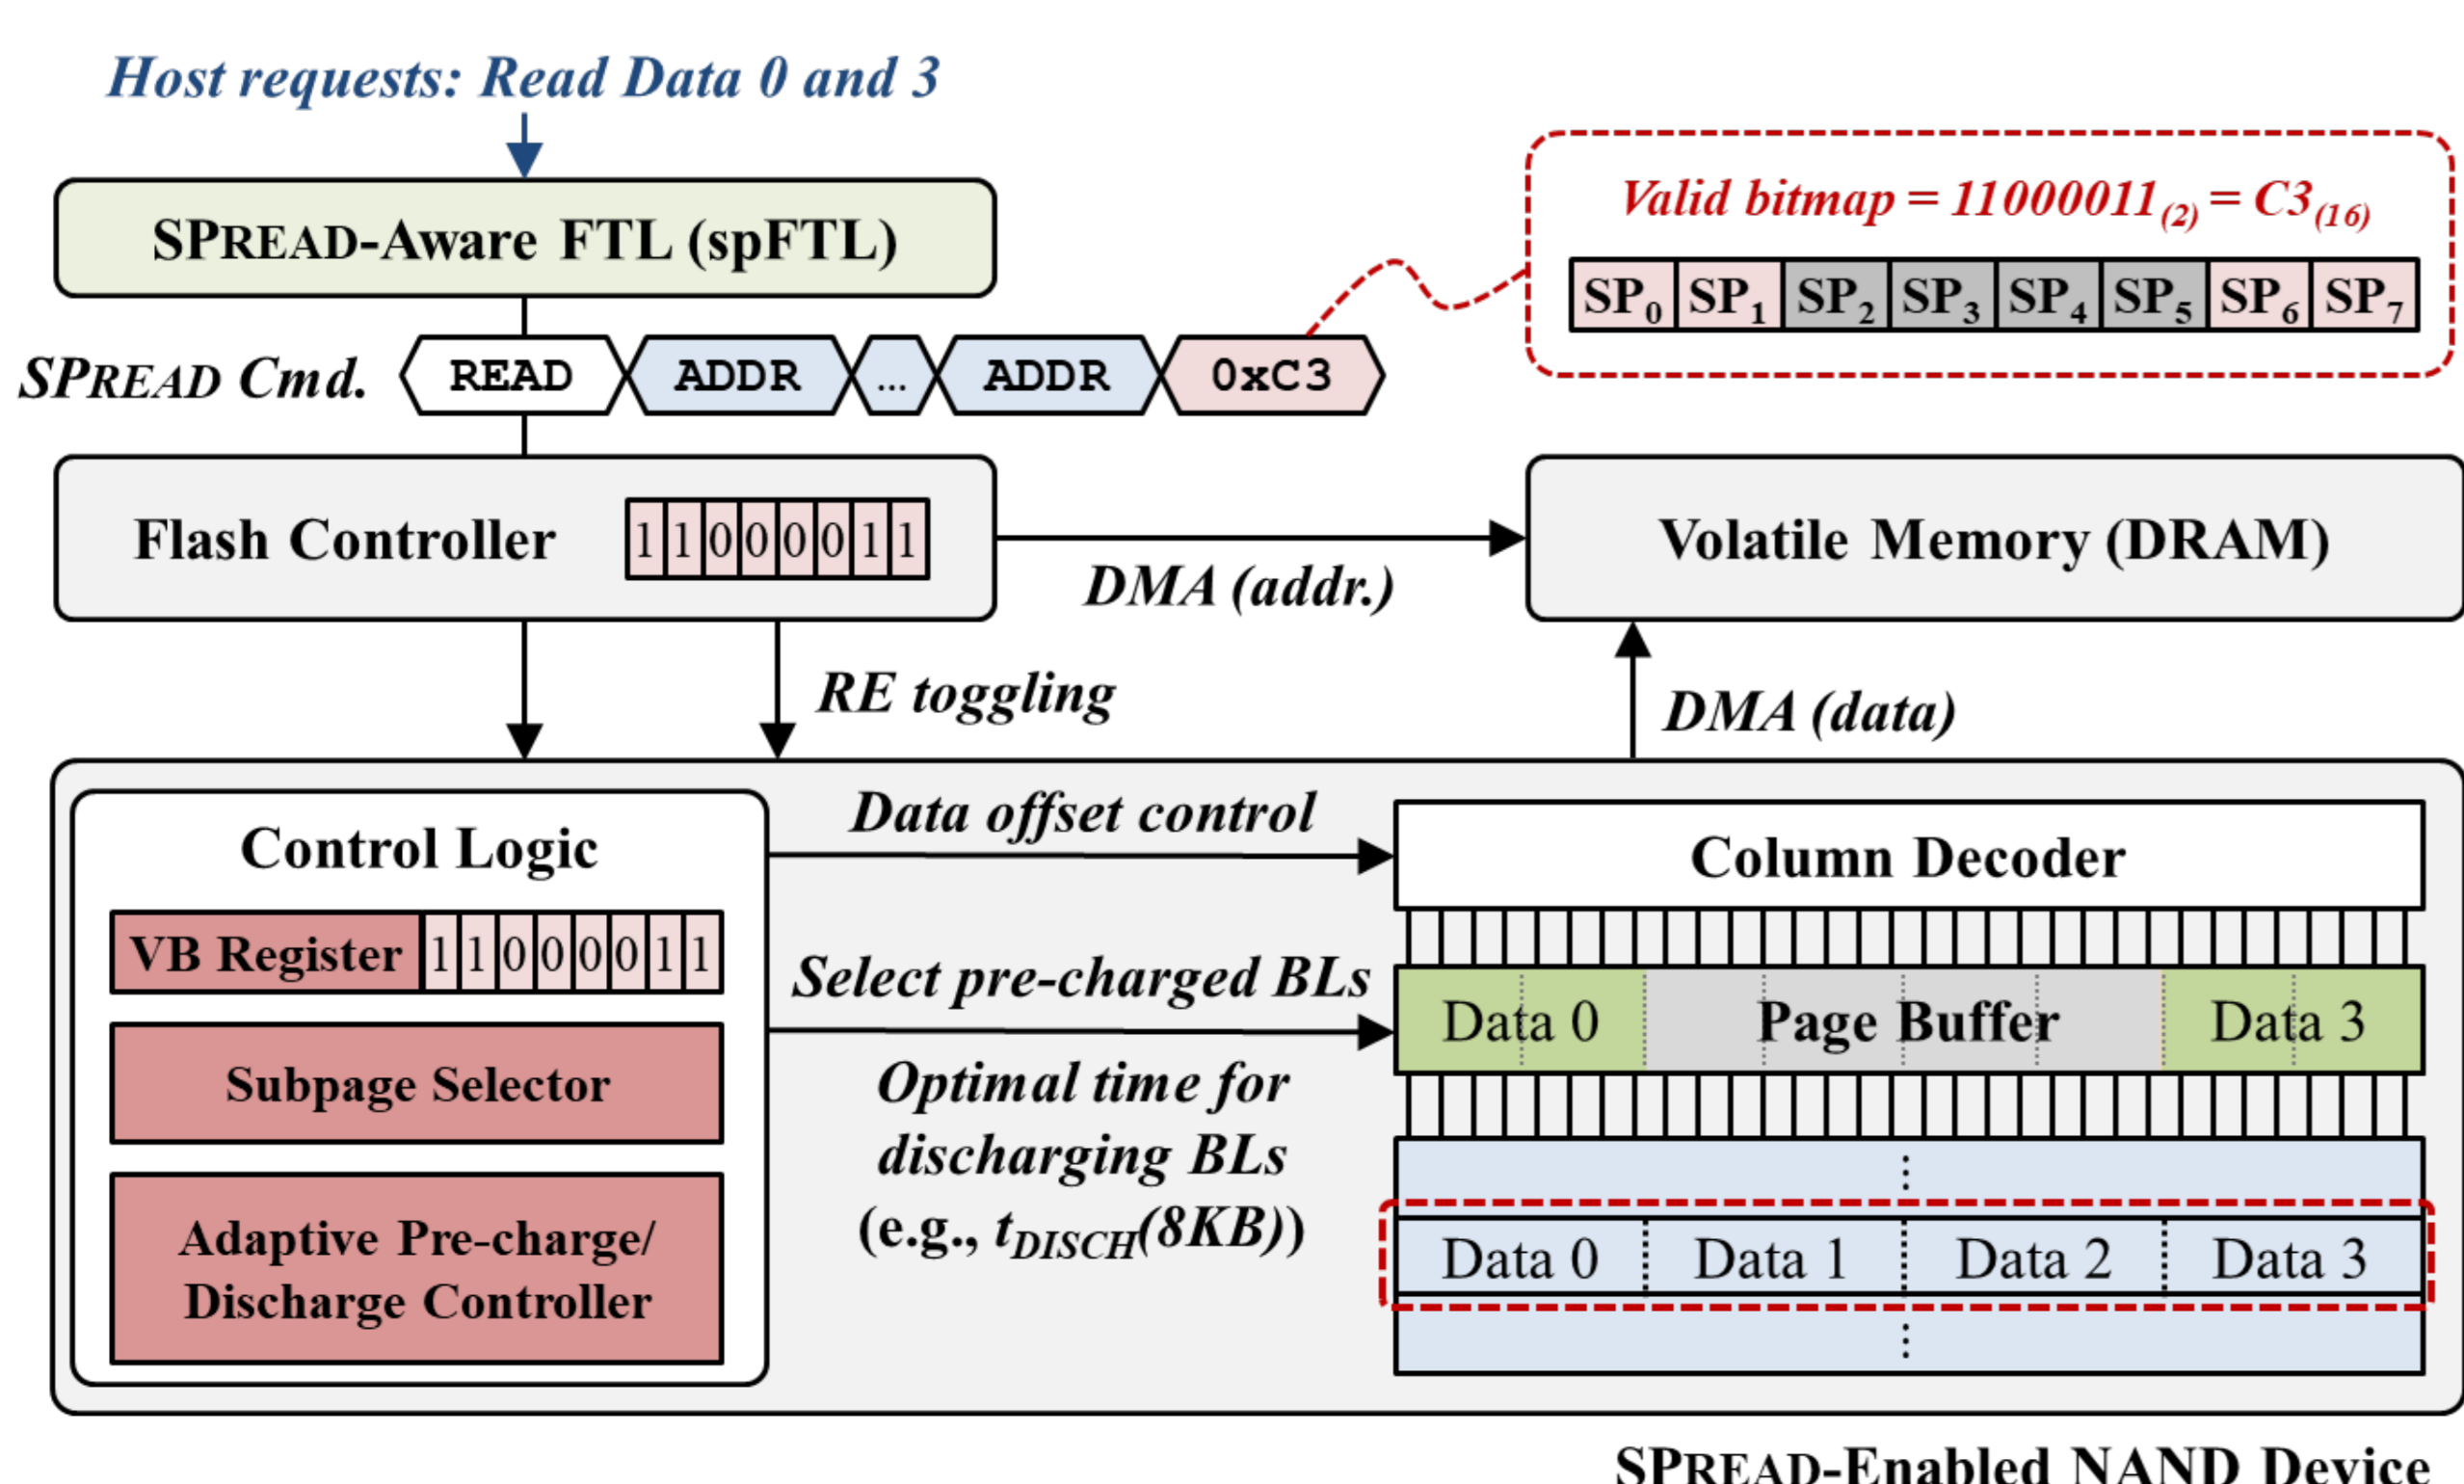
- Design requirements of subpage reads

- Subpage size = ECC encoding unit (e.g., 2 KB for LDPC)
- High flexibility for various subpage patterns



Subpage-Parallel Reads (SPREAD) 5

- Supports optimal latencies for small reads
 - By **selectively sensing and transferring** only demanded data
 - With a high flexibility to choose subpages to read



Experimental Results 6

- IOPS increased by up to **122%** (52.3% on average)
 - With low RAF values closed to 1
- High benefits even under many writes (USR and STG)
 - Reducing GC overheads by 13% and 7% on average

