FAST'25

Liquid-State Drive: A Case for DNA Block Device for Enormous Data

Jiahao Zhou[1], Mingkai Dong[1], Fei Wang[2], Jingyao Zeng[1], Lei Zhao[2], Chunhai Fan[2], Haibo Chen[1]

[1]Institute of Parallel and Distributed Systems (IPADS), Shanghai Jiao Tong University

[2]School of Chemistry and Chemical Engineering, Shanghai Jiao Tong University

上海交通大學 SHANGHAI JIAO TONG UNIVERSITY

IPADS INSTITUTE OF PARALLEL AND DISTRIBUTED SYSTEMS

Outline

- DNA Storage Basics
- A Naïve DNA Block Design
- LiqSD Design
- Evaluation
- Conclusion

DNA Storage Principle & Mechanism

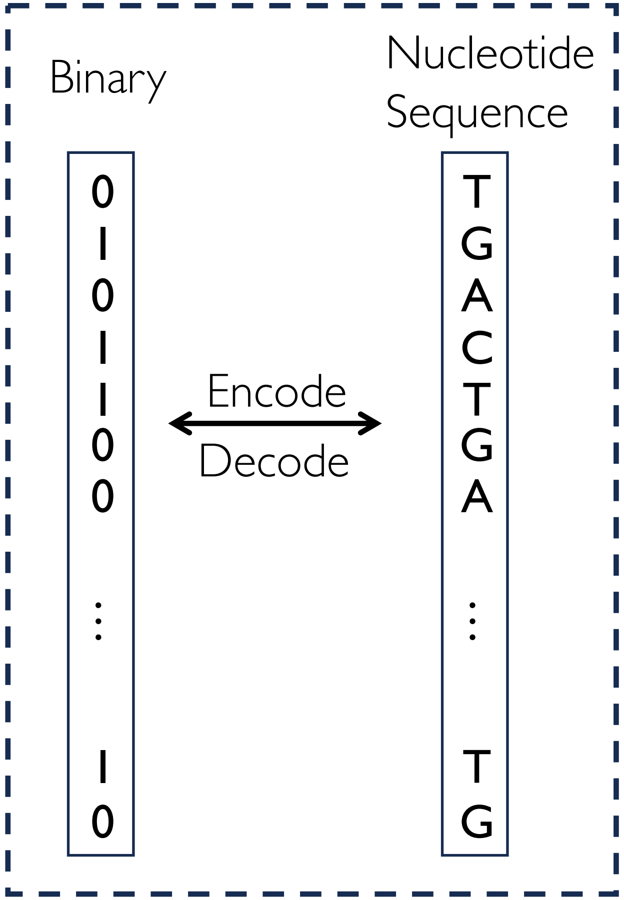

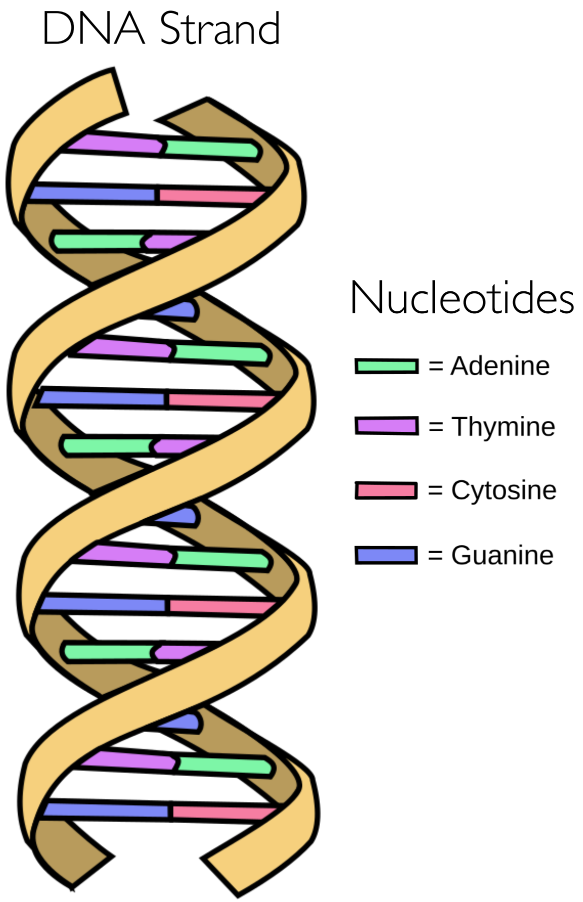

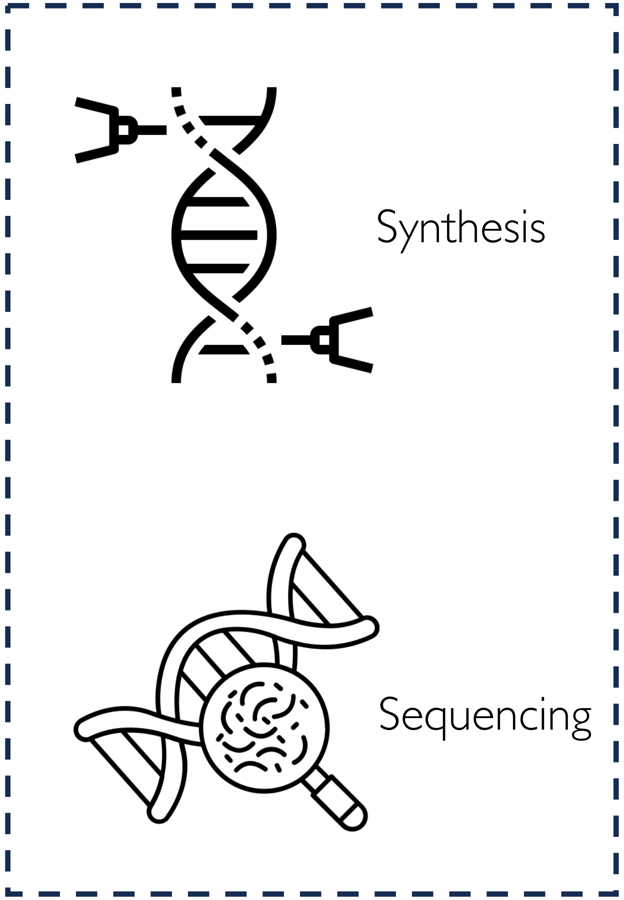

Information Presentation

Information Store/Retrieval

DNA Storage Principle & Mechanism

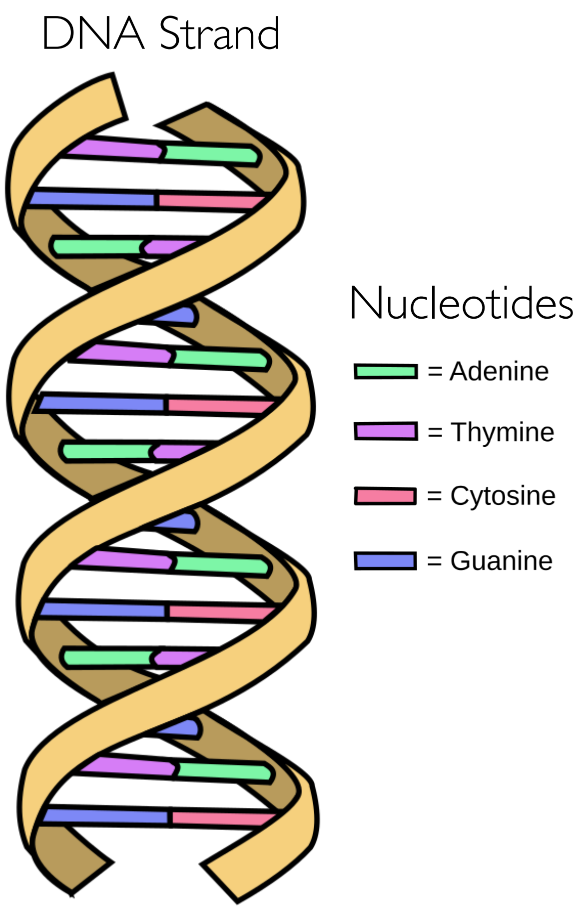

	Storage Density (bits/mm³)	*Lifespan*
HDD	10^{10}	~ 5 years
SSD	10^{13}	~ 5 years
DNA	10^{18}	Centuries

High Density & Long Lifespan

A Typical DNA Storage Hardware

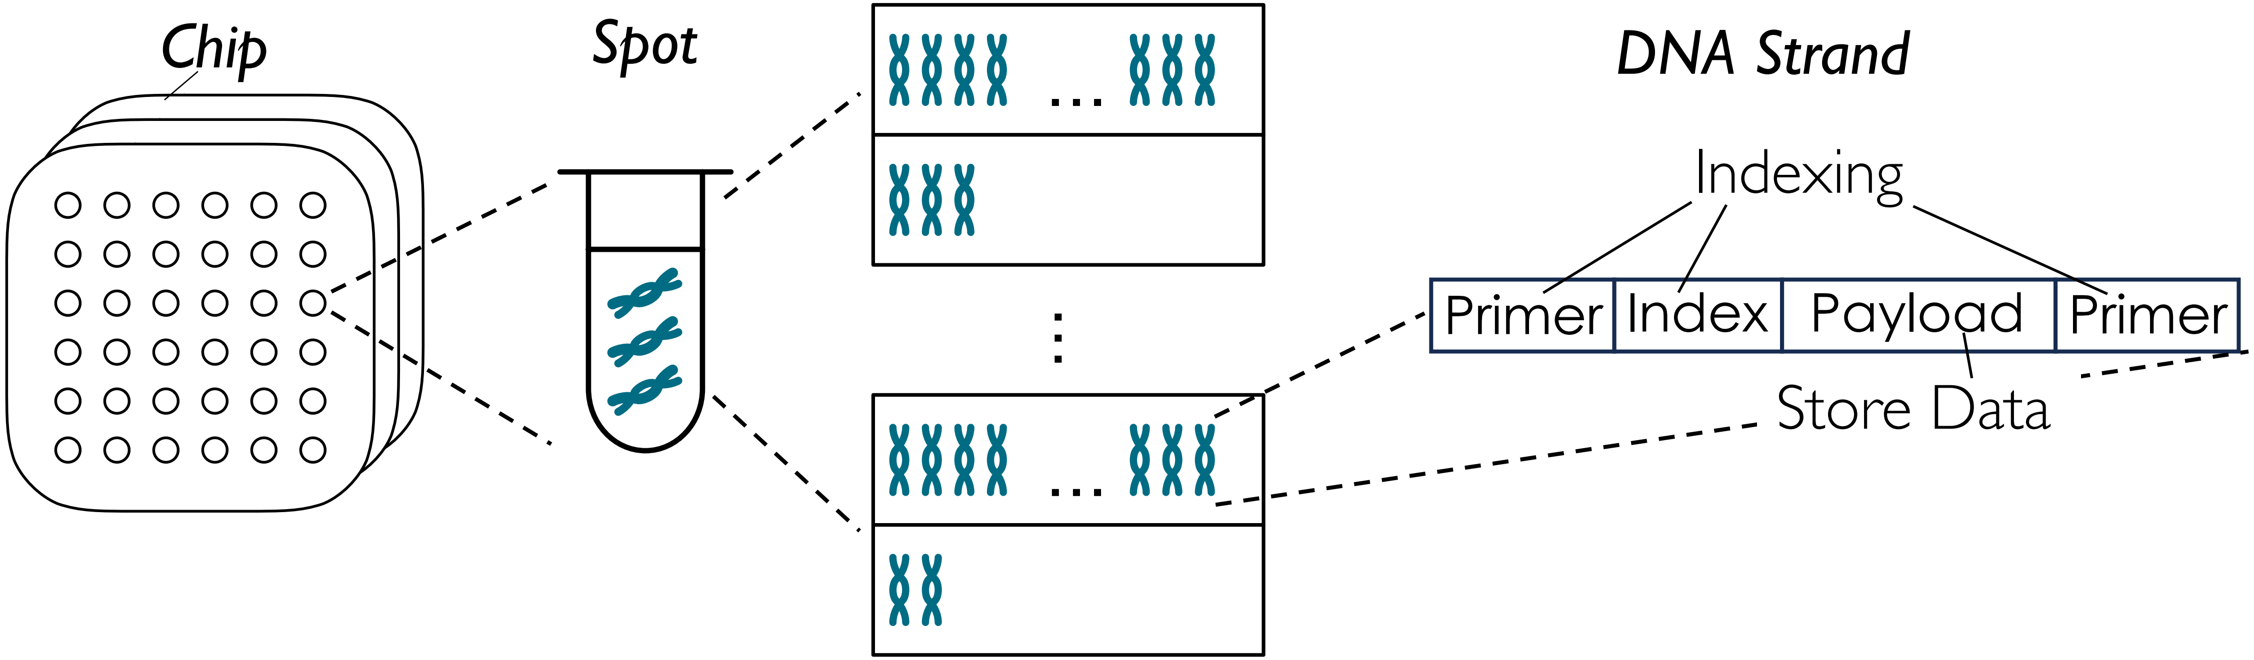

A Typical DNA Storage Hardware

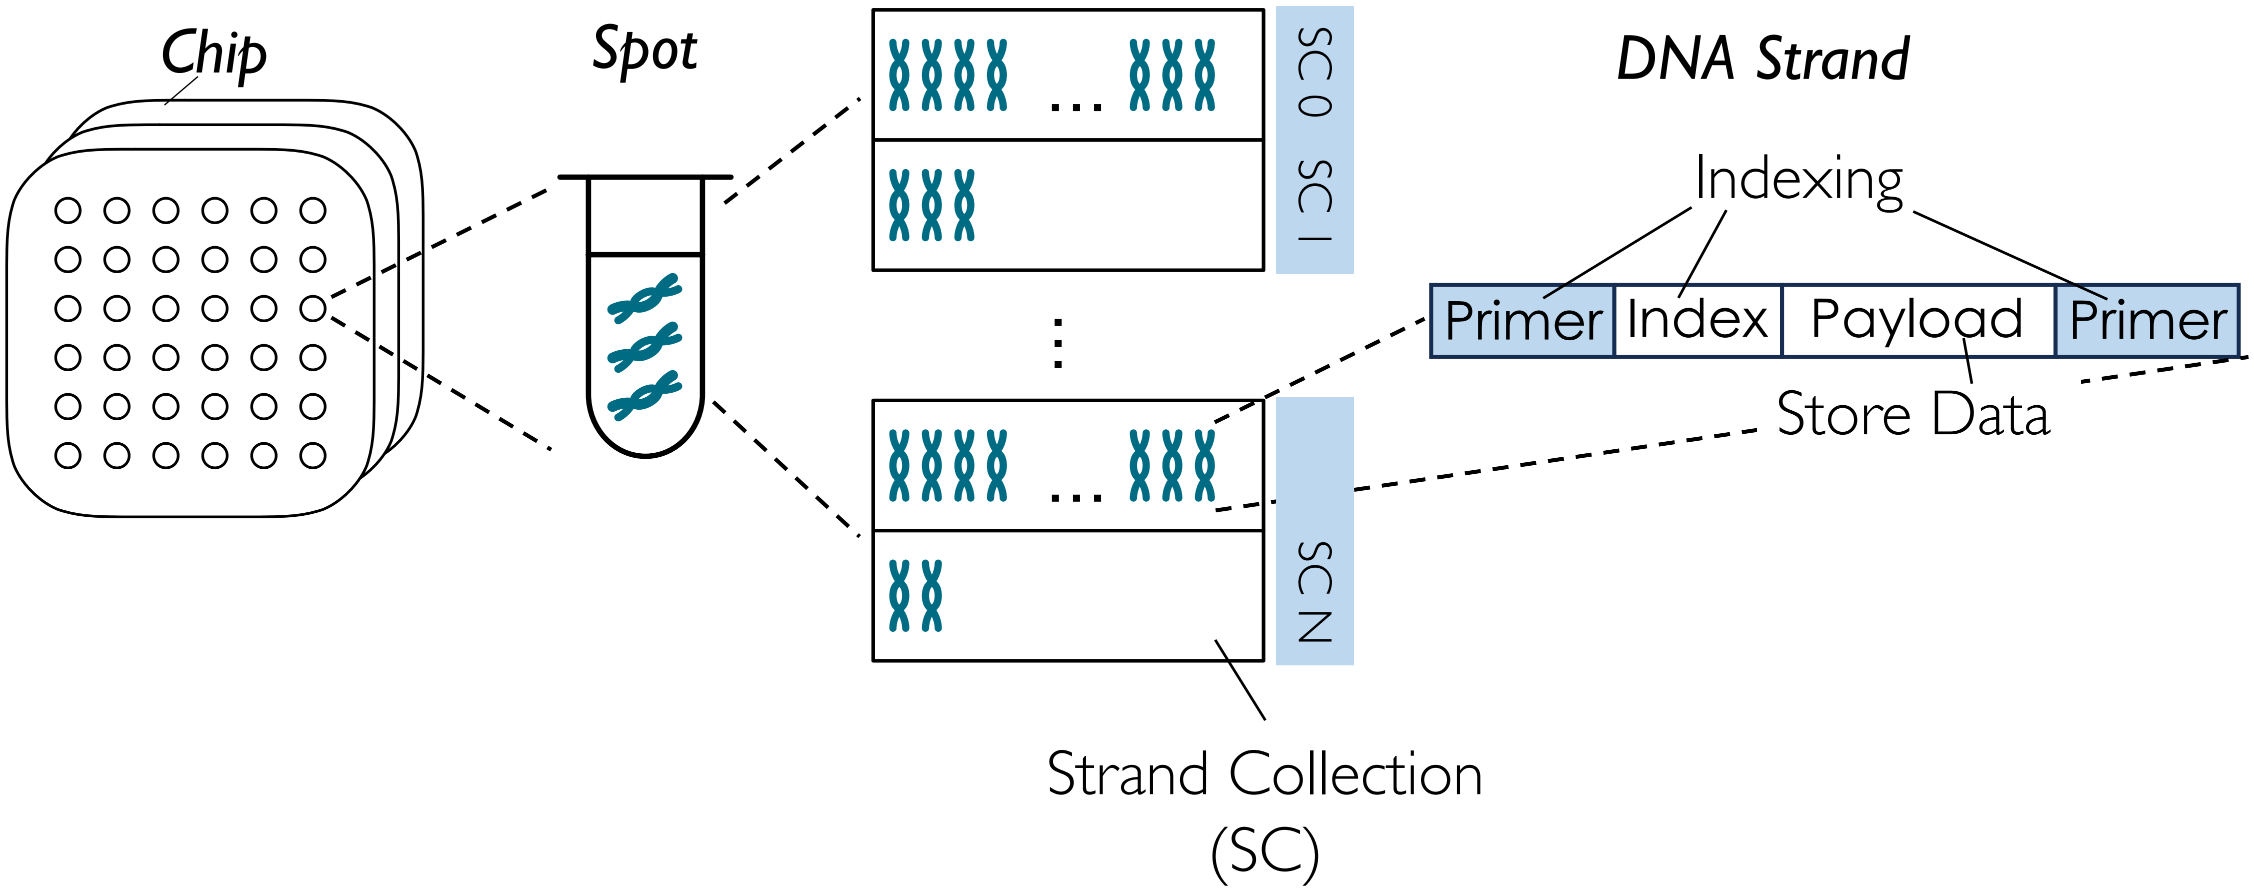

A Typical DNA Storage Hardware

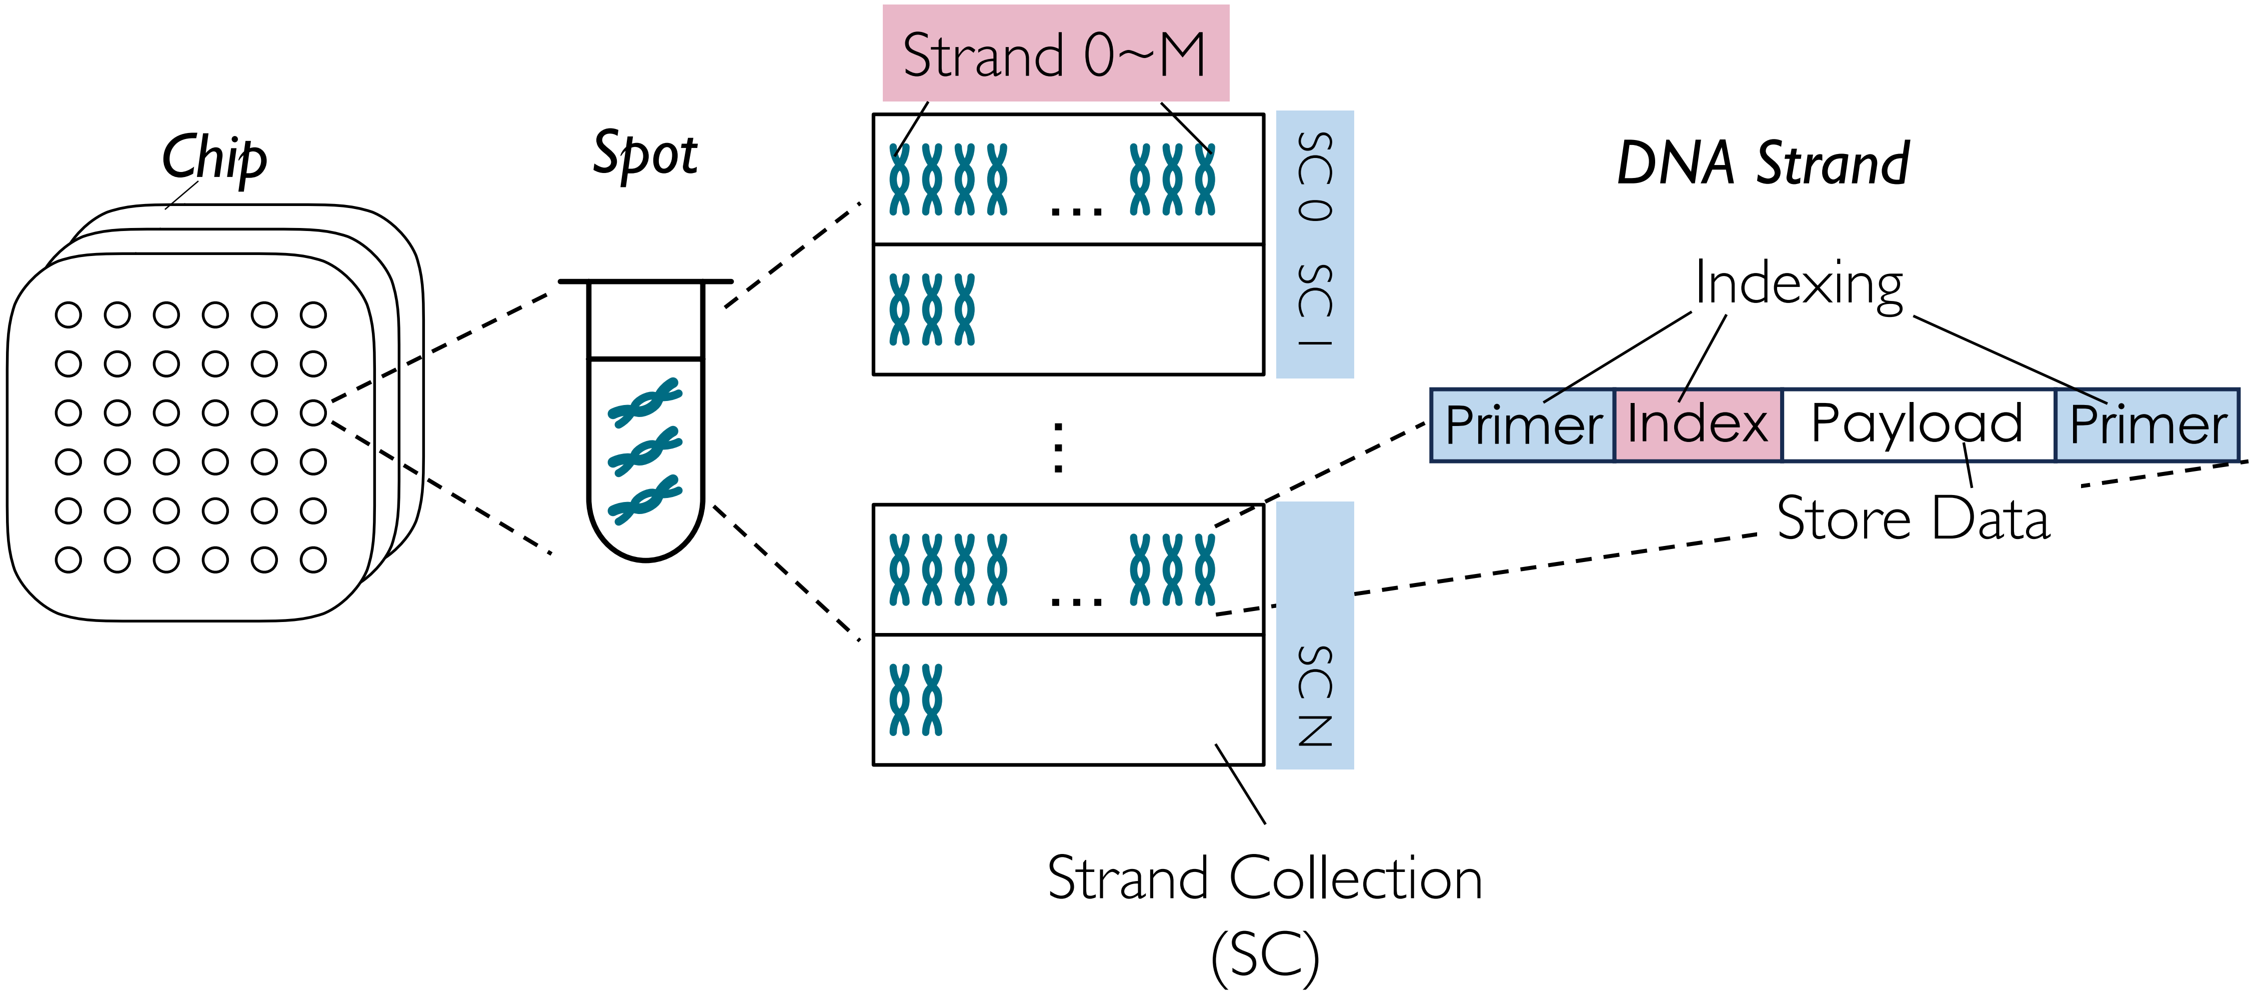

DNA Storage Operation

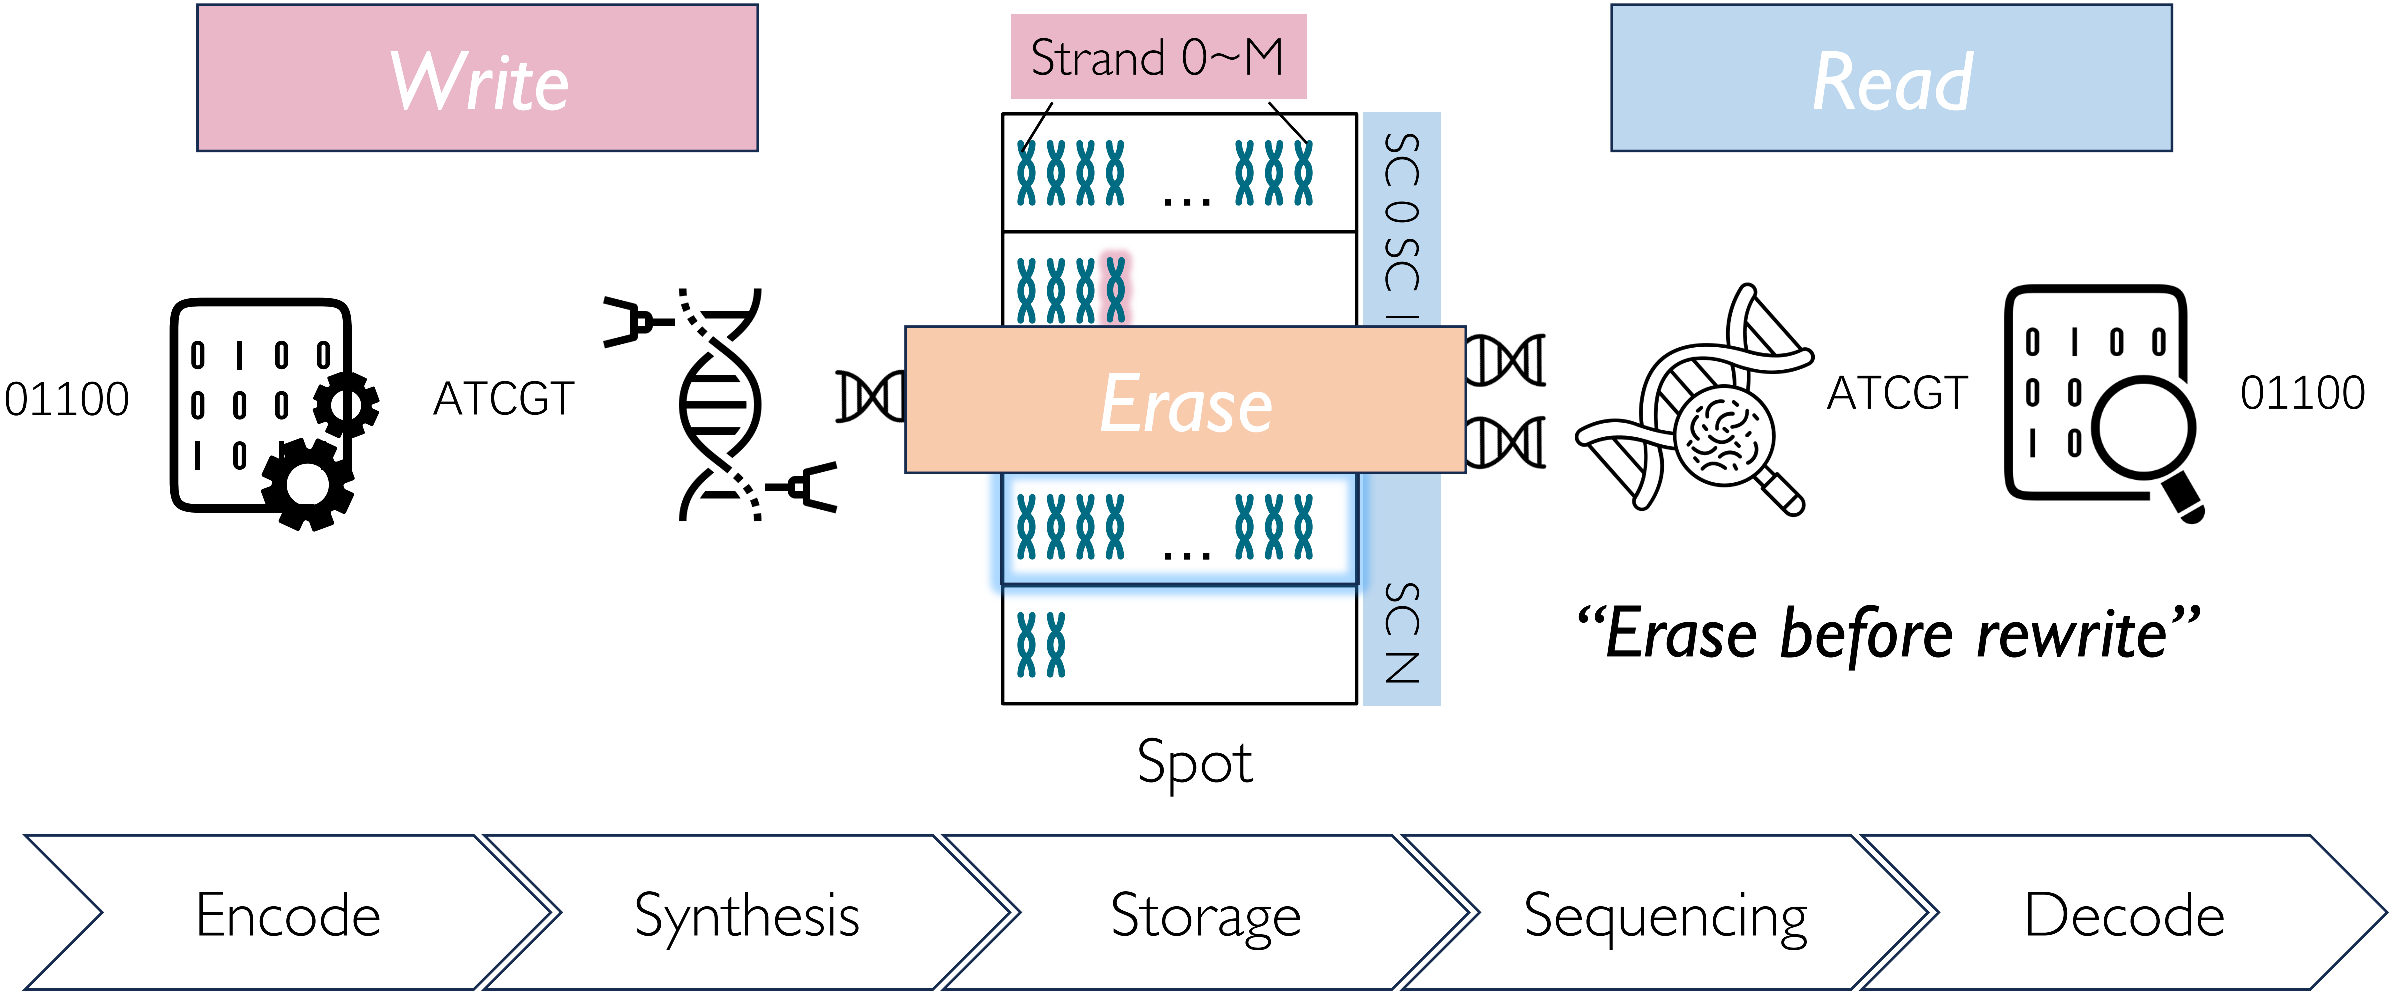

DNA Storage Operation

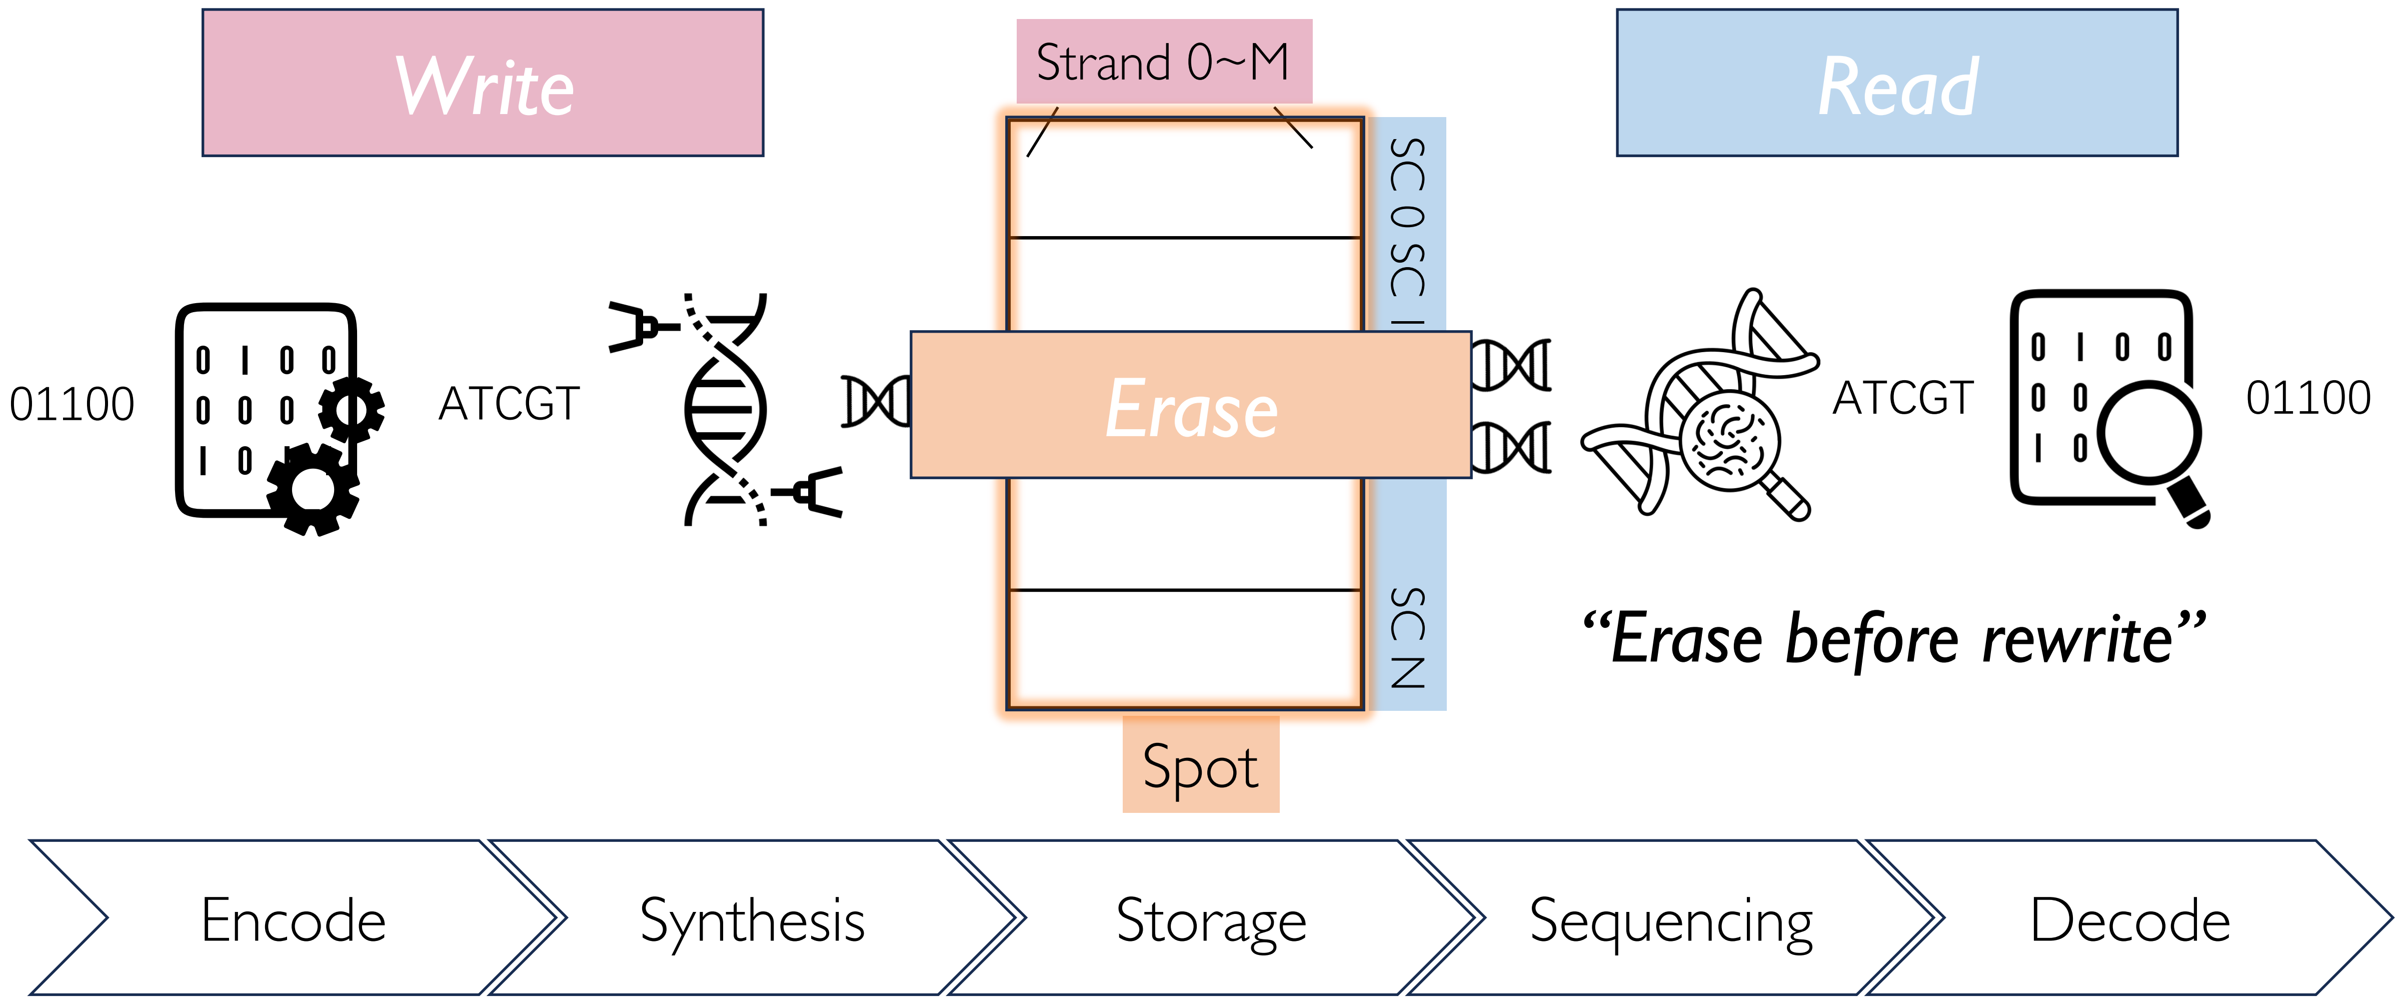

DNA Storage Operation

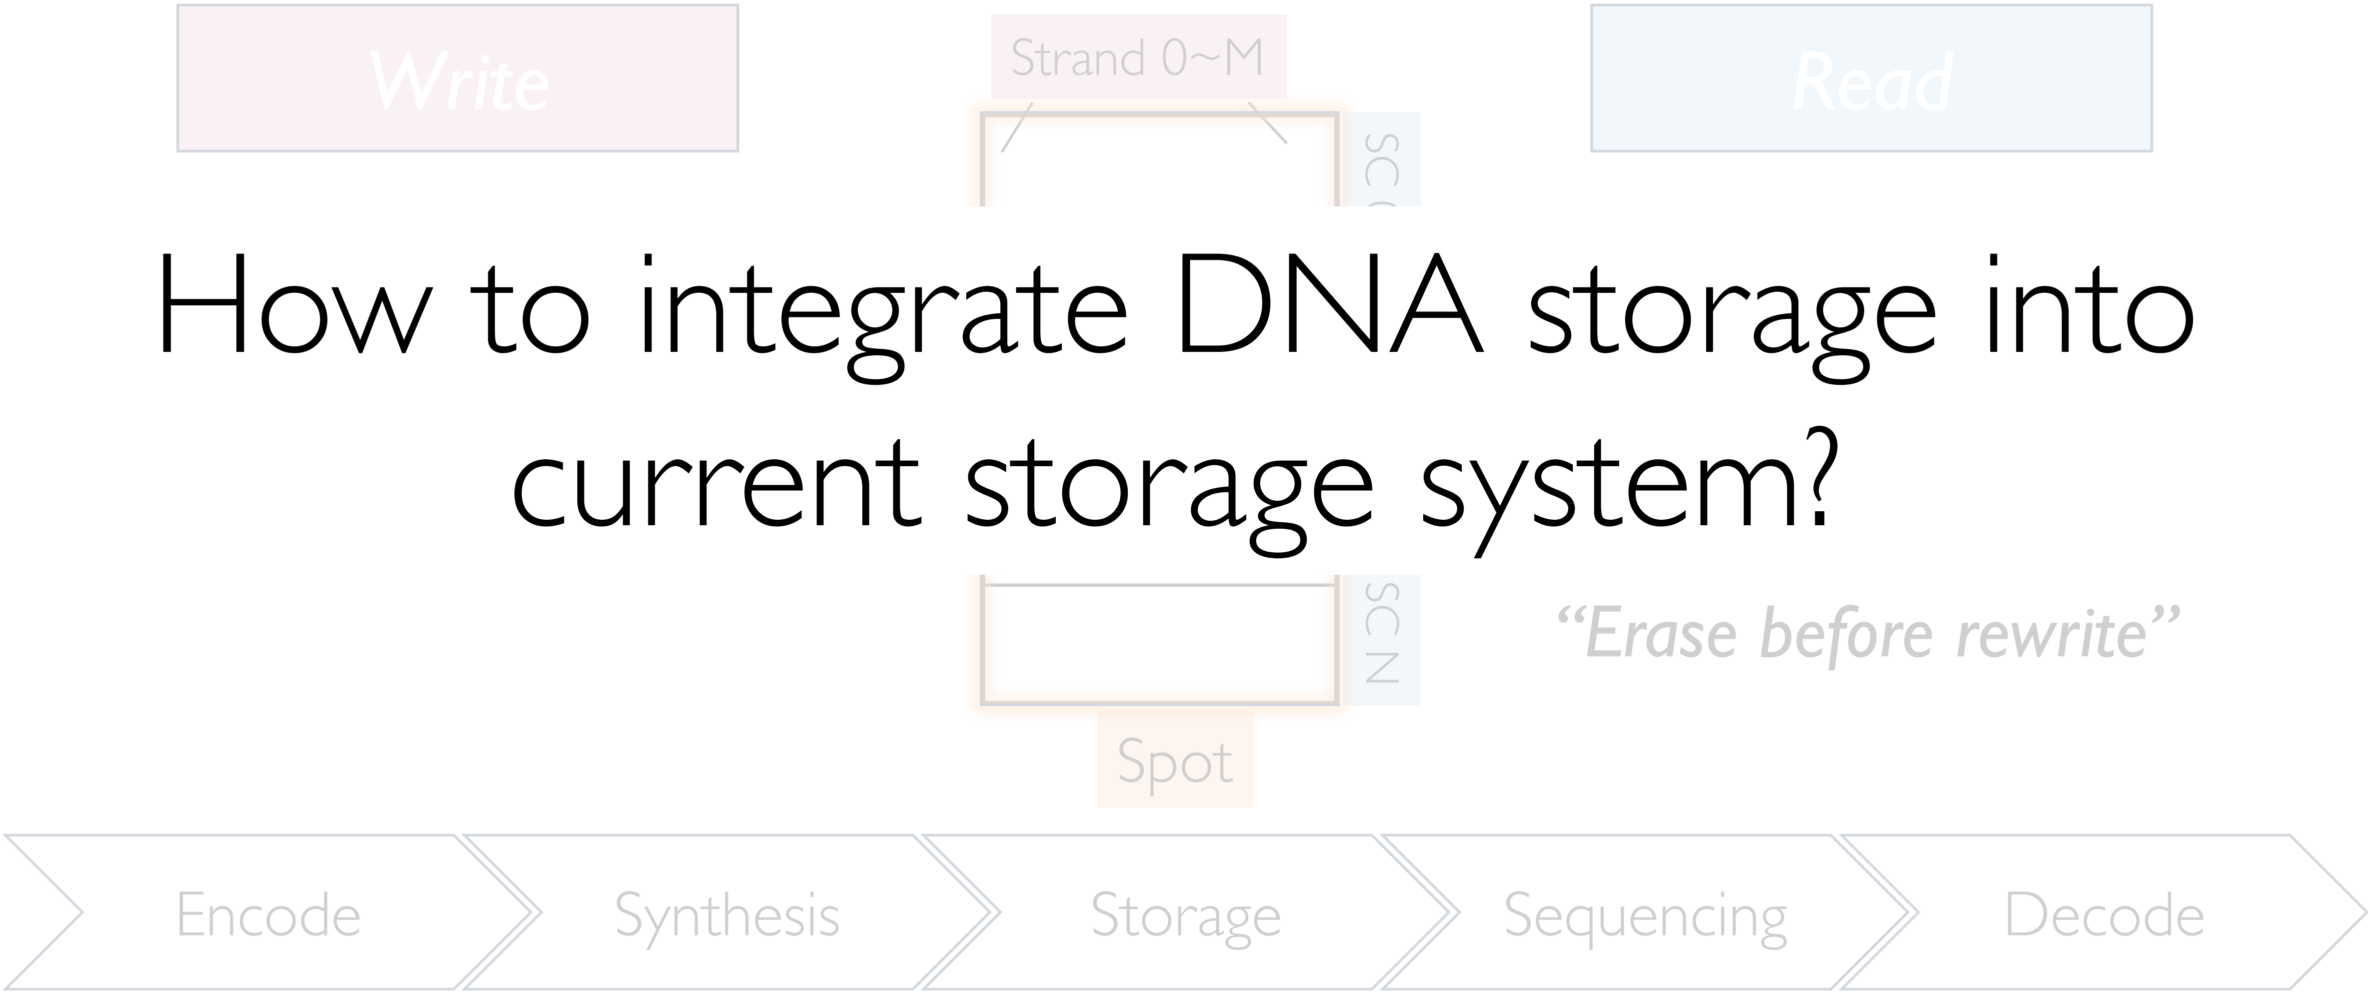

How to integrate DNA storage into current storage system?

Building a DNA Block Device

Simplicity

Random access
the fixed-size block

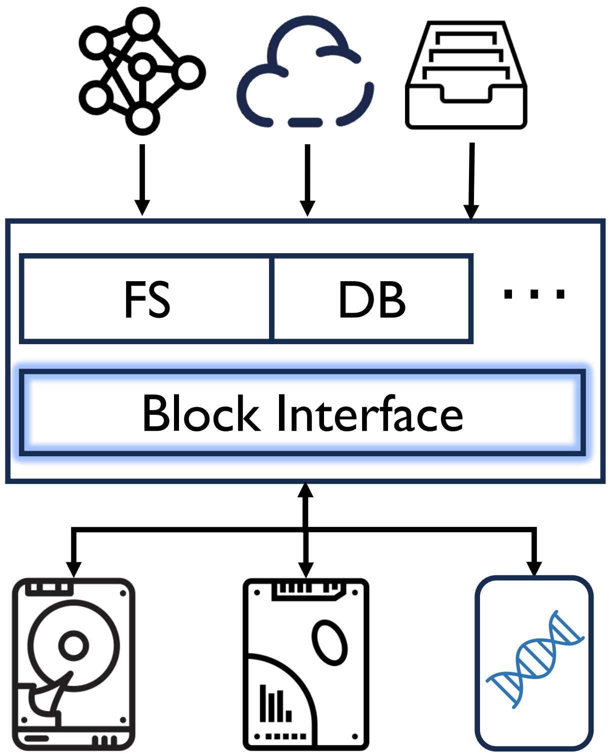

Generality

Widely deployed
in various scenarios

Outline

- DNA Storage Basics
- A Naïve DNA Block Device
- LiqSD Design
- Evaluation
- Conclusion

A Naïve Design: Block Partition

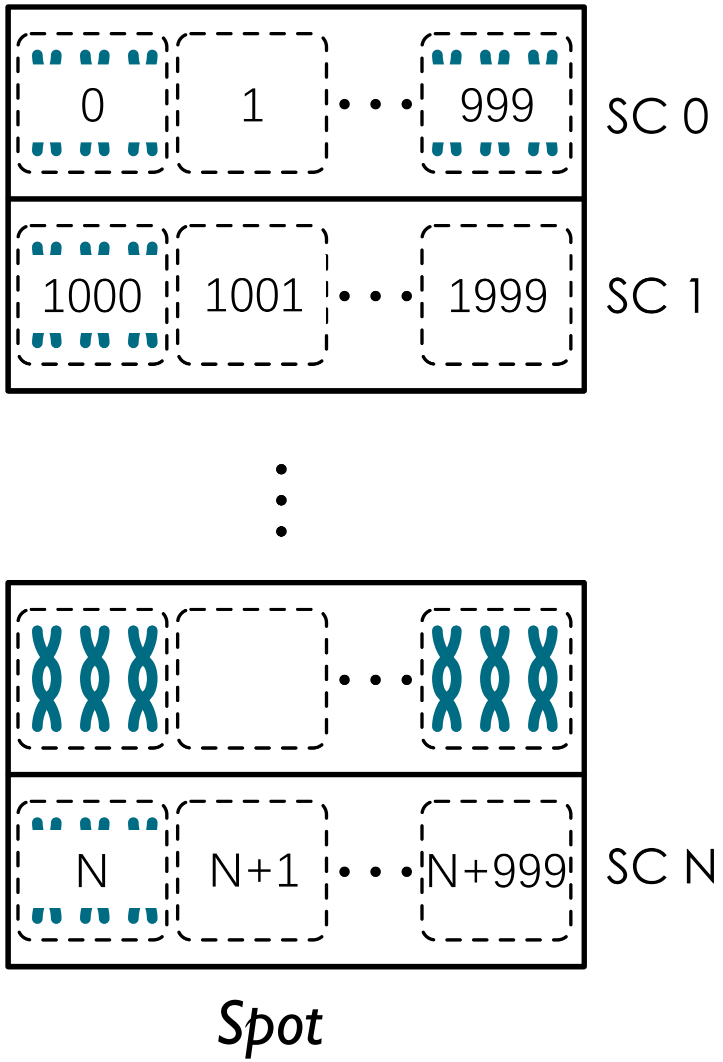

Mapping *consecutive strands* into a block

Assigning *Physical Block Address* to each block (PBA)

A Naïve Design: Block Operation

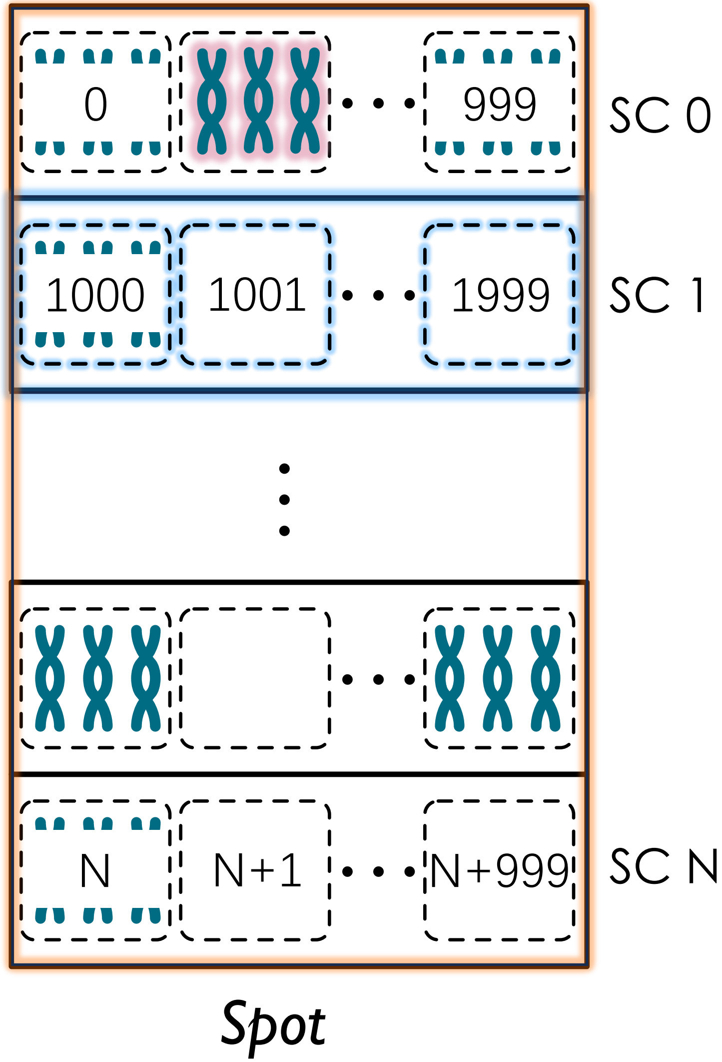

	Read	Write	Erase
DNA Hardware	SC ~1 million strands ~10s MiB/s	Strand 300 nucleotides KiB/s ~ MiB/s	Spot 2,000 SCs Erase before rewrite

Read and write 12 Million Blocks,
Extremely costly!

	Read	Write	Update
Naïve Block Device	Read all blocks within SC	Write strands of the block	1. Real all SC 2. Erase Spot 3. Write new & other block

Out-of-place Update

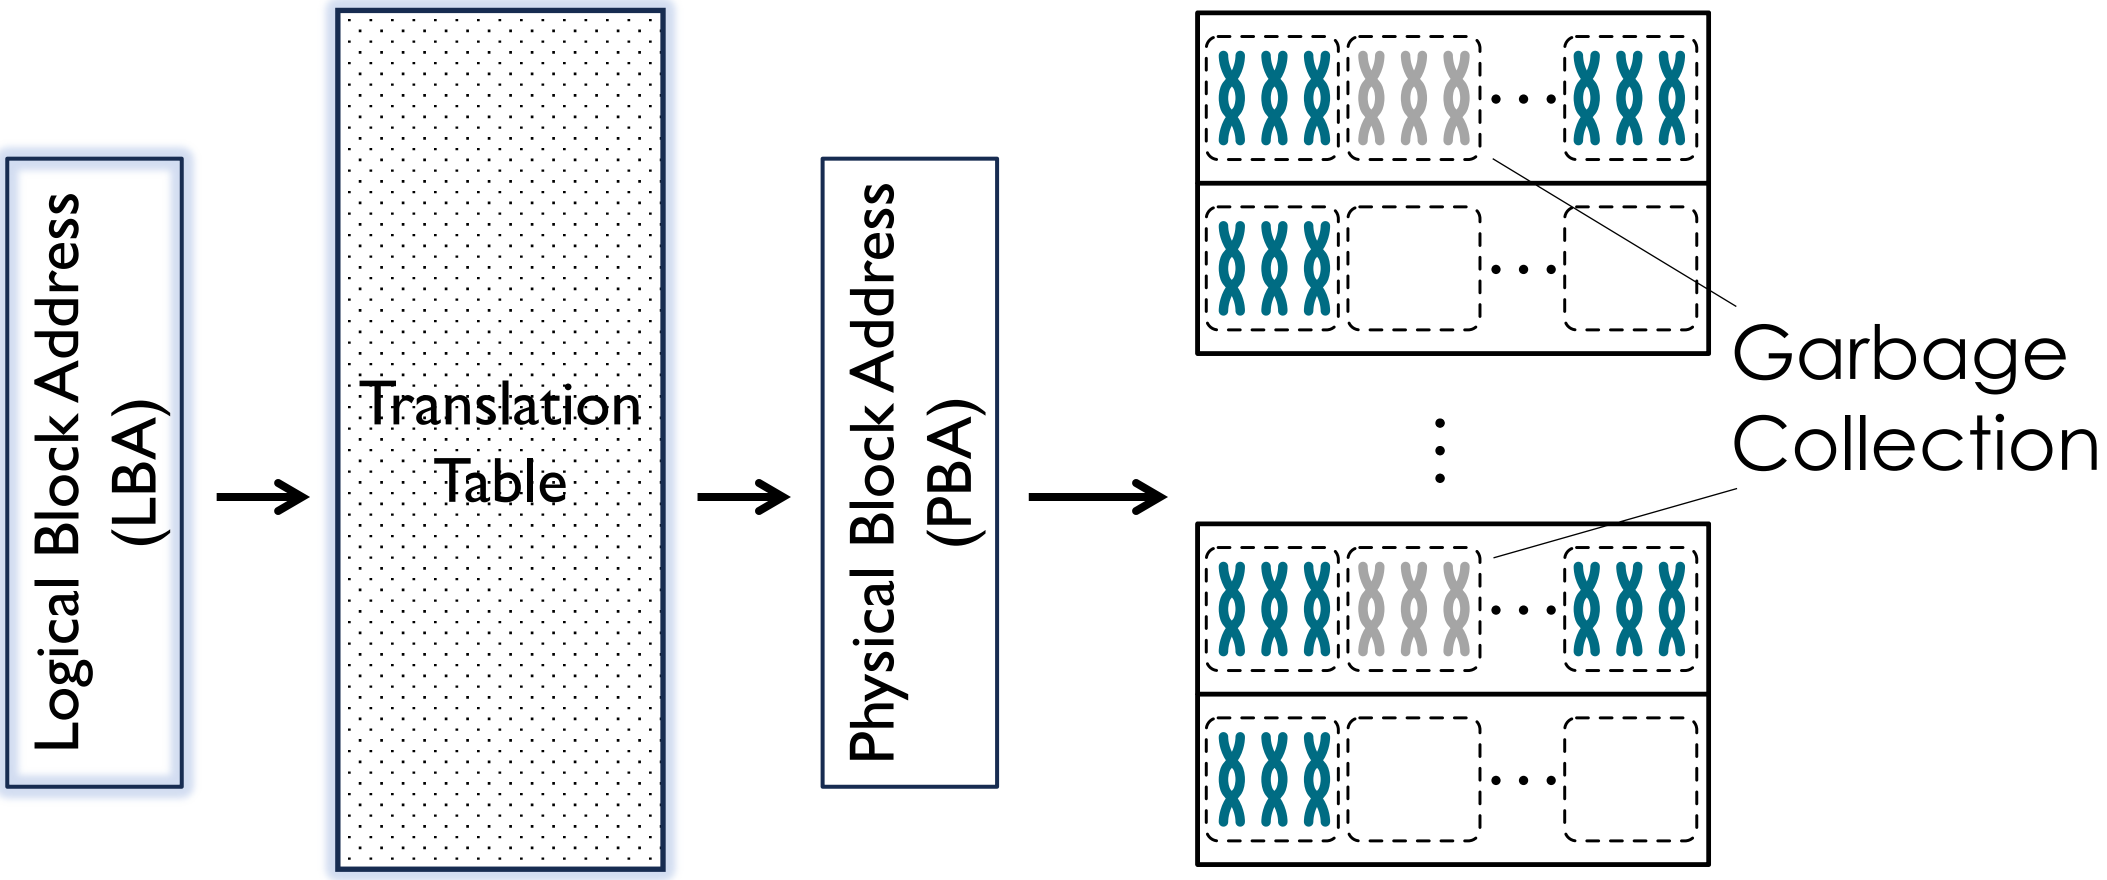

Translate LBA → PBA

Out-of-place Update

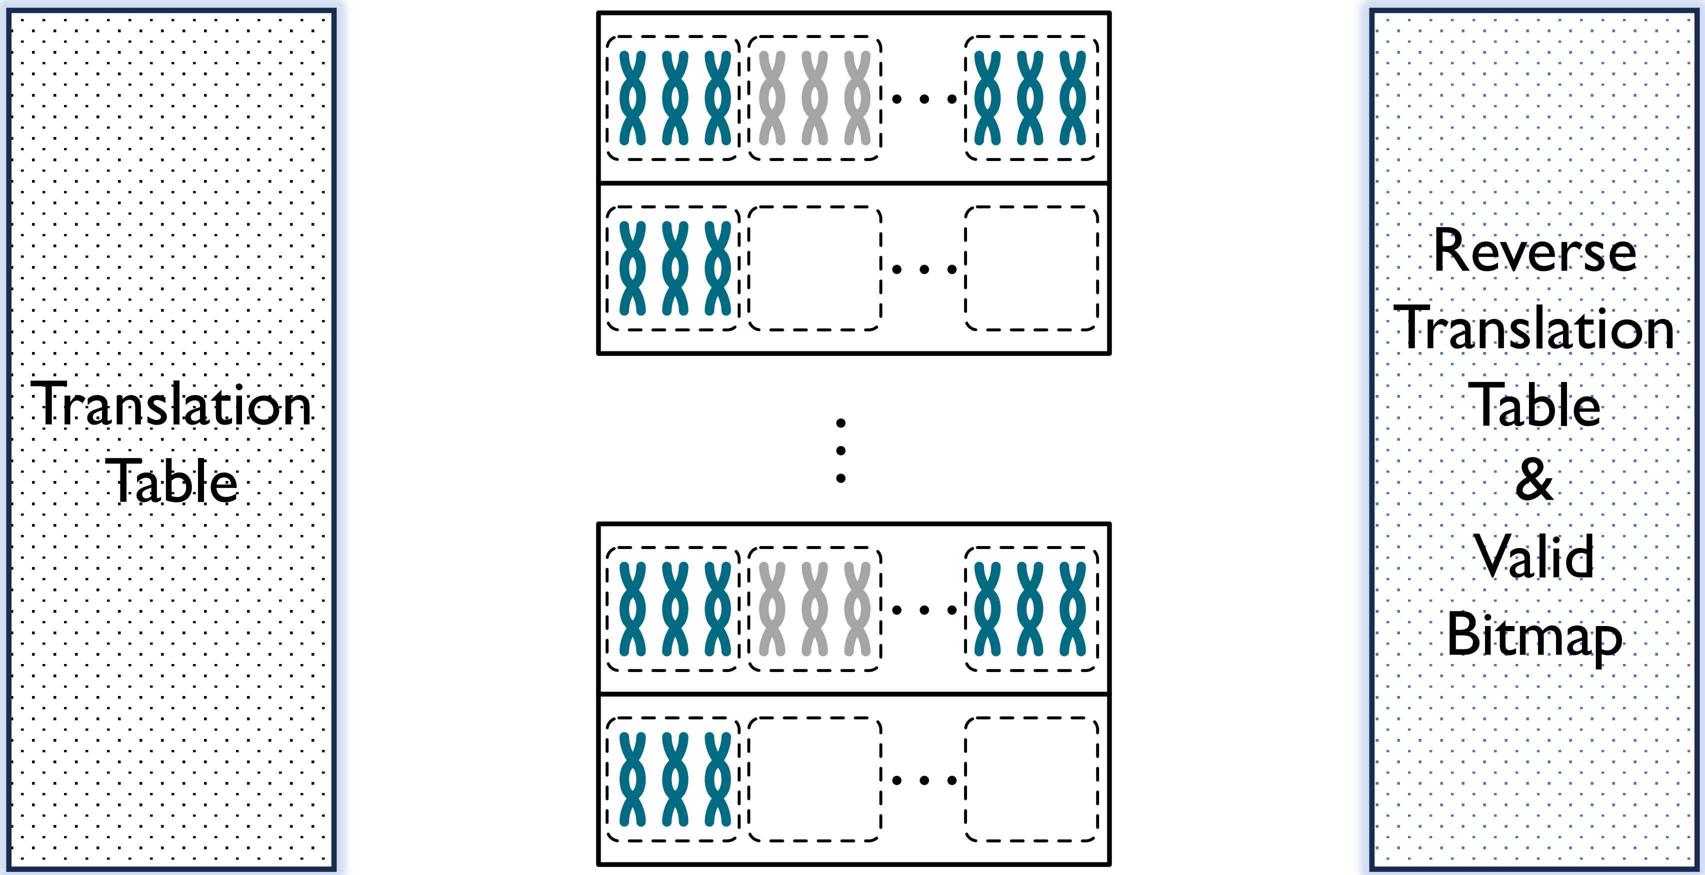

<u>*Translate LBA → PBA*</u>   <u>*Record GC Metadata*</u>

Out-of-place Update

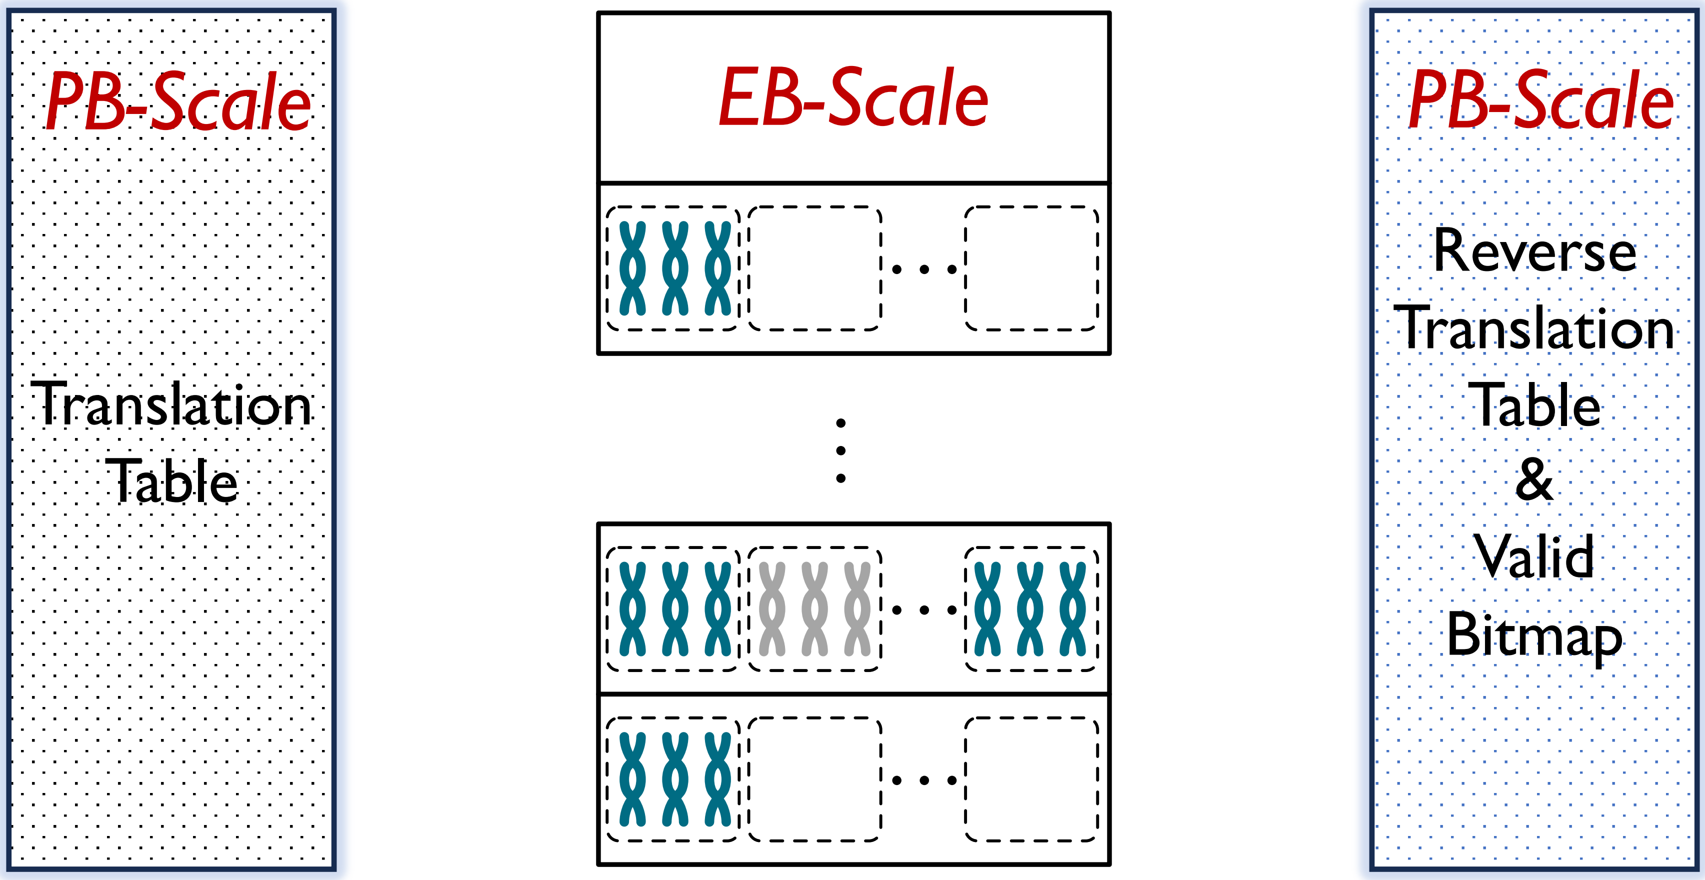

Translate LBA → PBA *Record GC Metadata*

Out-of-place Update

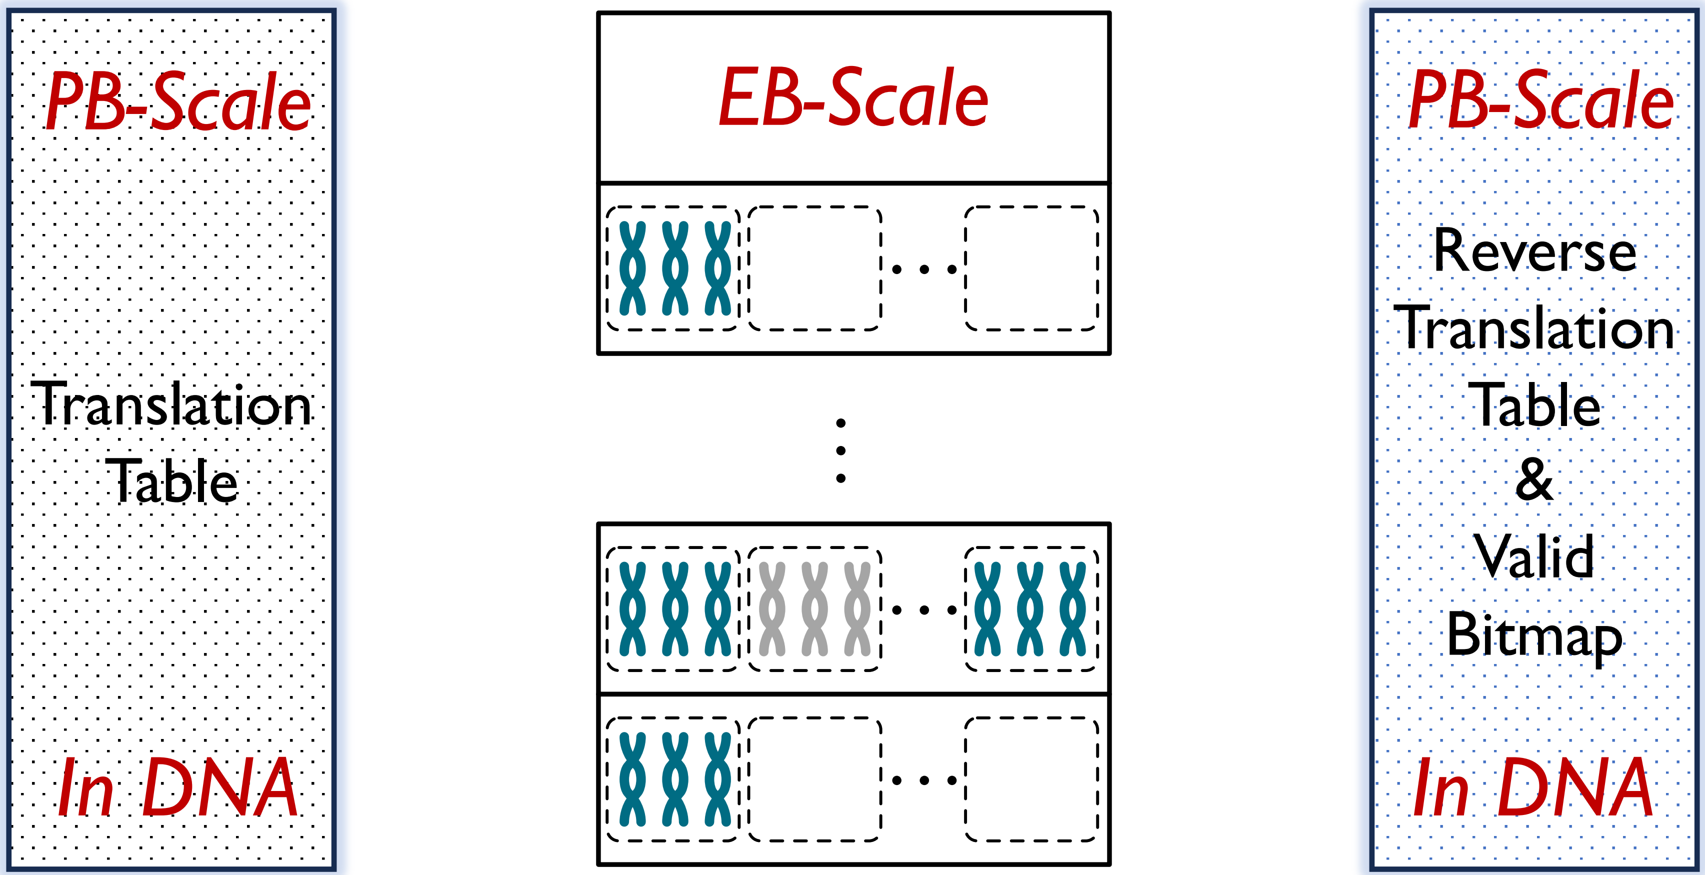

PB-scale metadata needs to be stored in DNA!

Challenge:

Maintaining in-DNA metadata reintroduces costly DNA update

For each write,

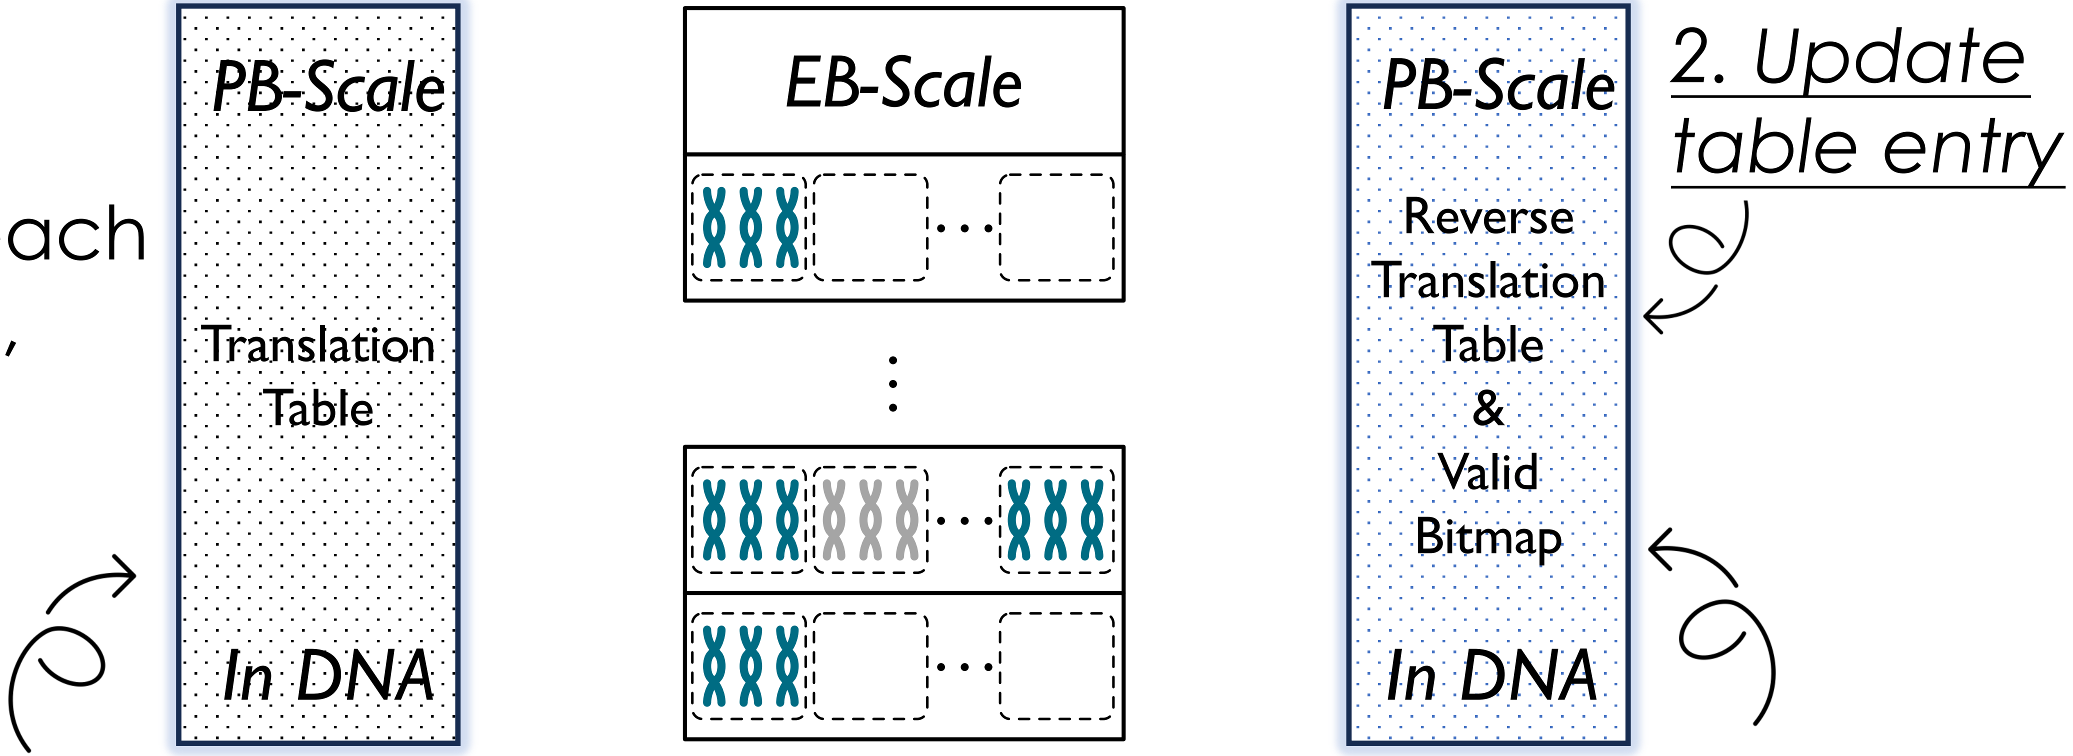

1. Update table entry

3. Invalidate old block

Challenge:

Maintaining in-DNA metadata reintroduces costly DNA update

For each write,

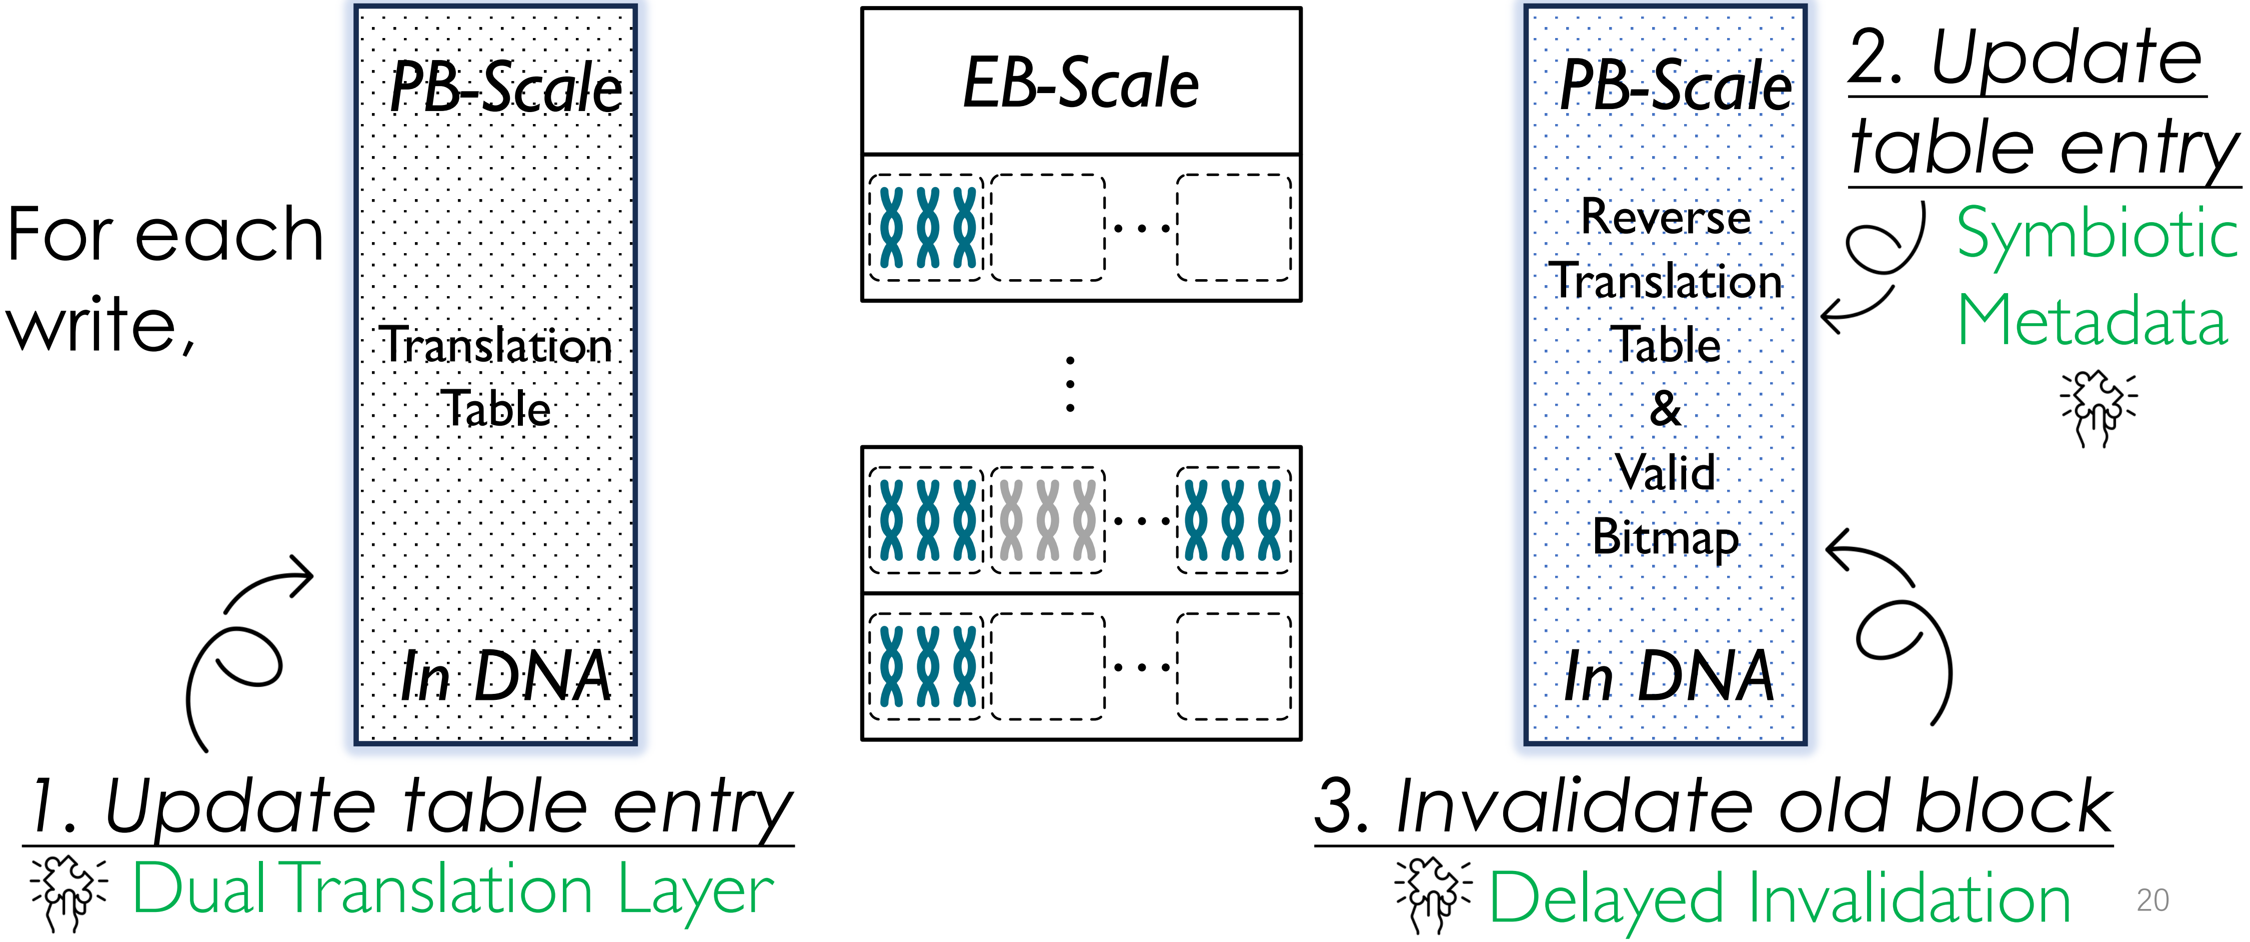

Outline

- DNA Storage Basics
- A Naïve DNA Block Design
- LiqSD Design
- Evaluation
- Conclusion

Dual Translation Layer

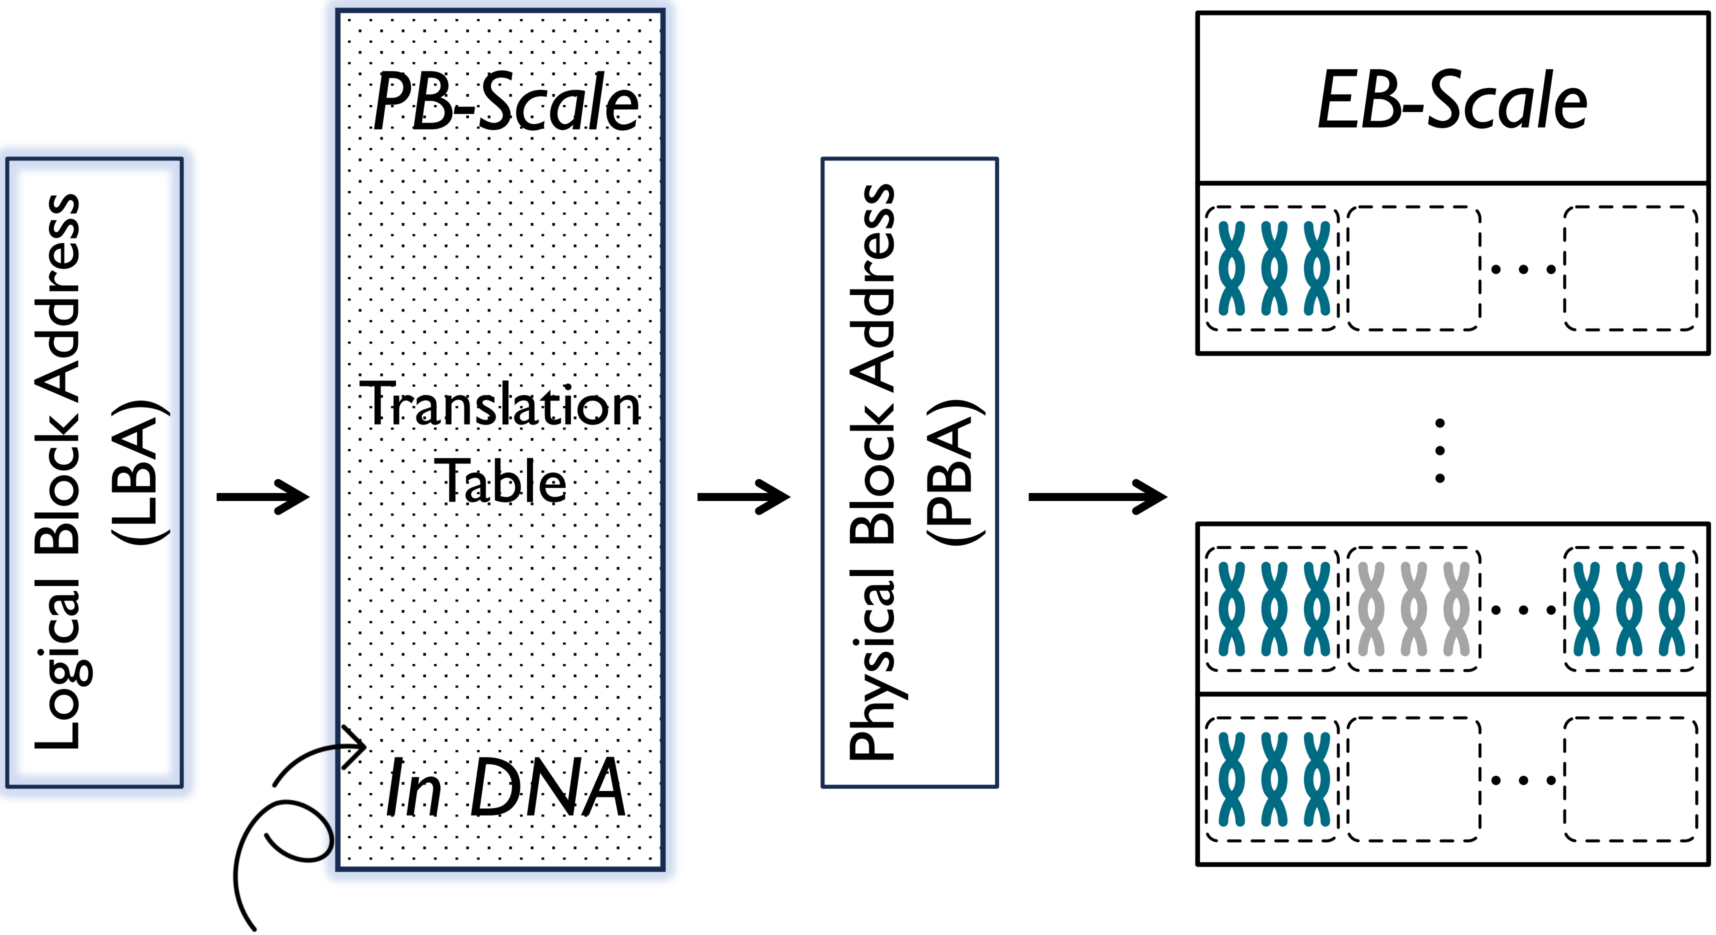

Also out-of-place update the translation table!

Dual Translation Layer

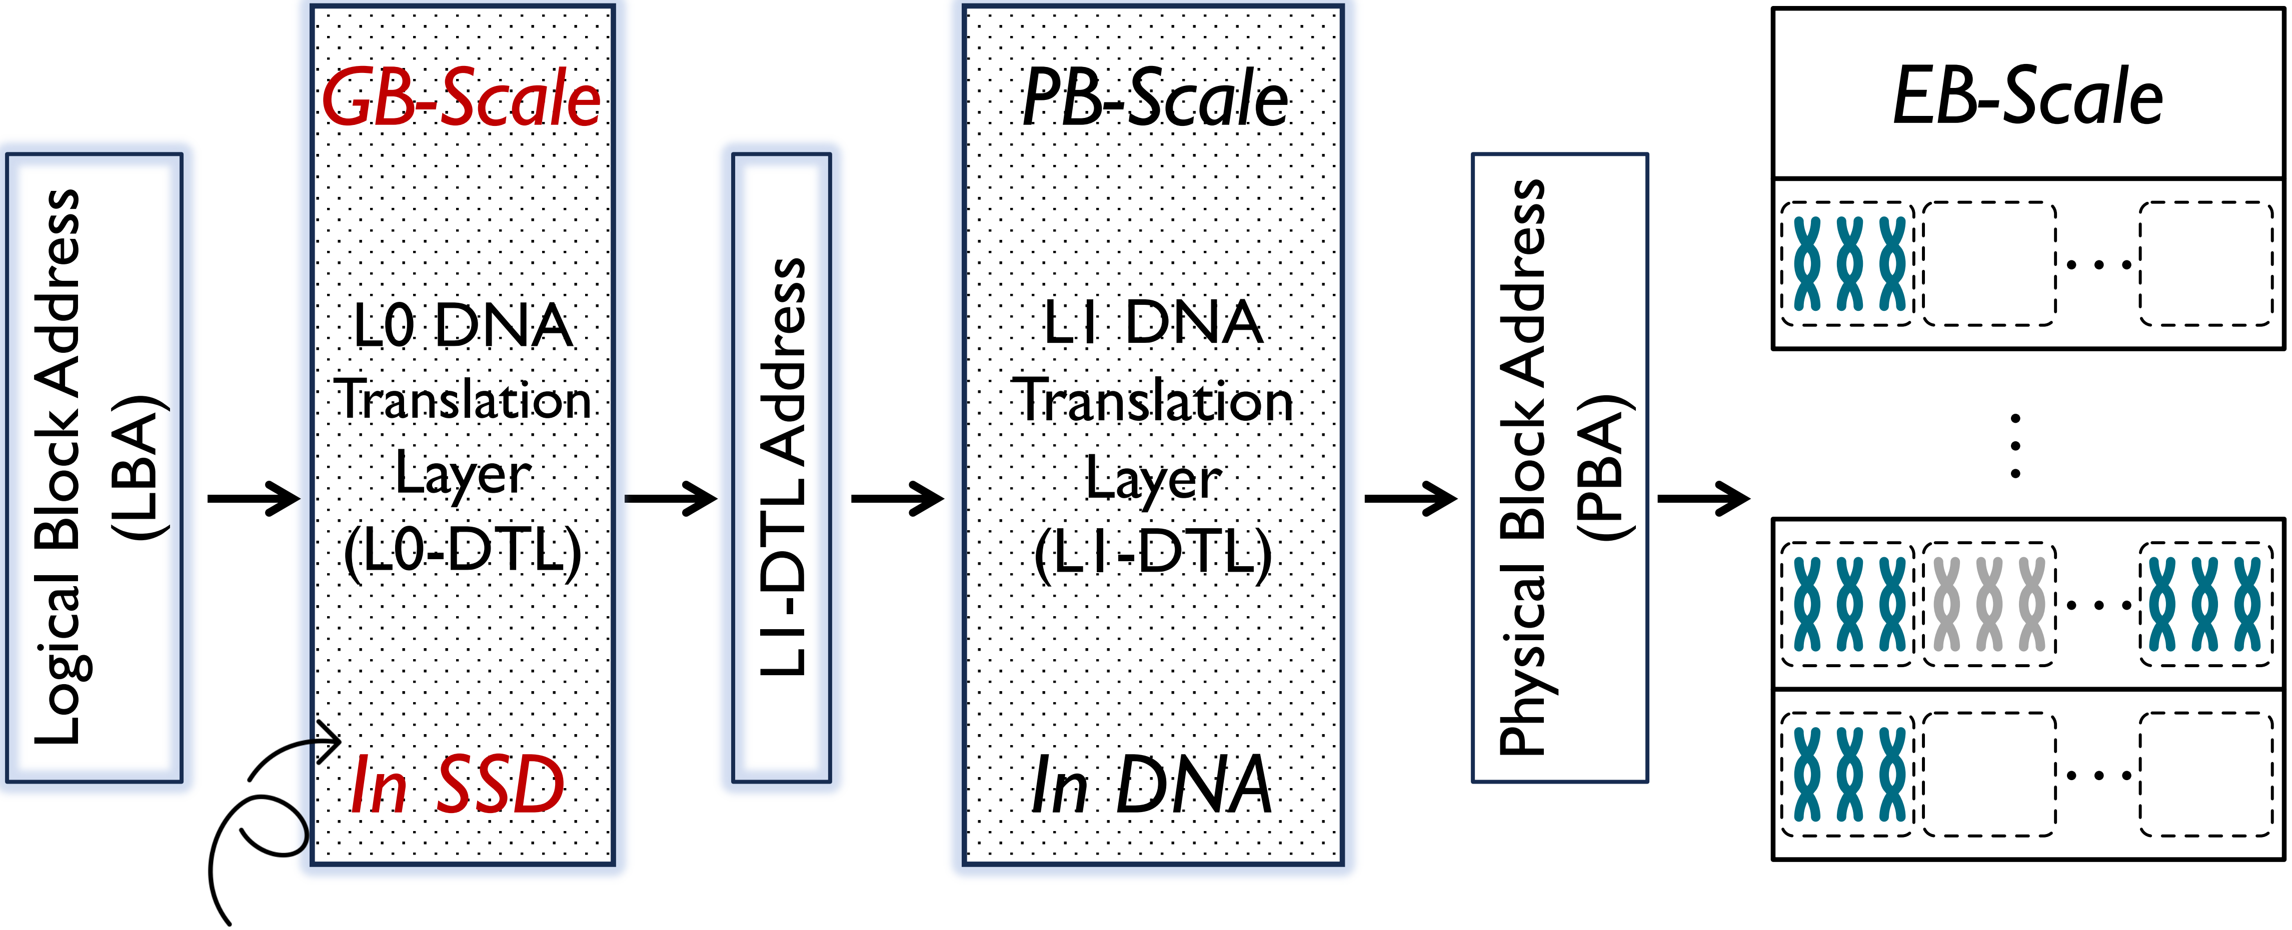

Small enough to store in fast SSDs

Symbiotic Metadata

Access along with physical block data

Store GC metadata with physical block symbiotically

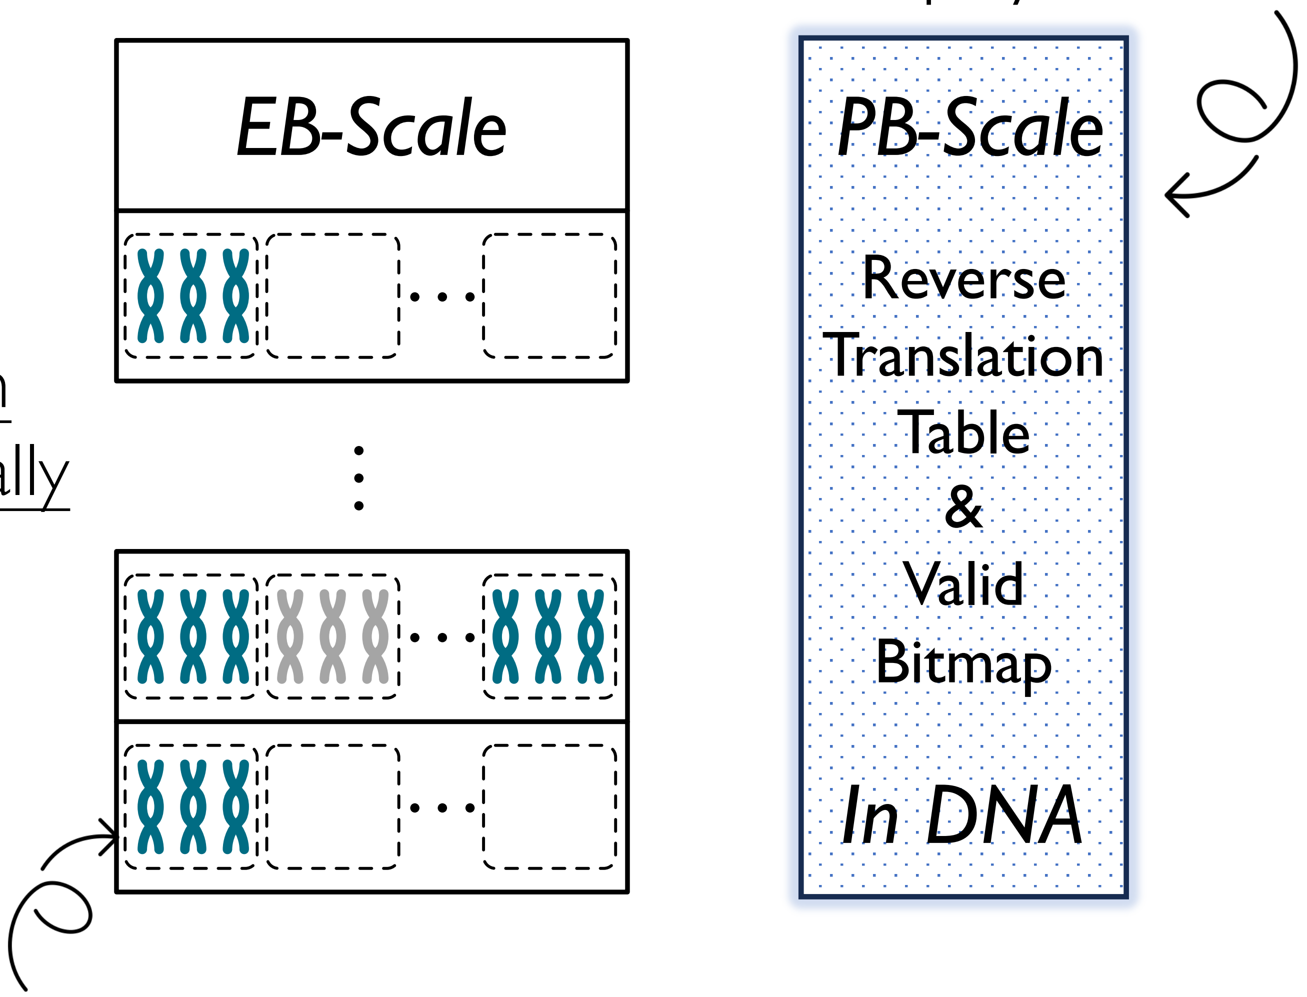

No DNA update for physical block

Symbiotic Metadata

Access along with physical block data

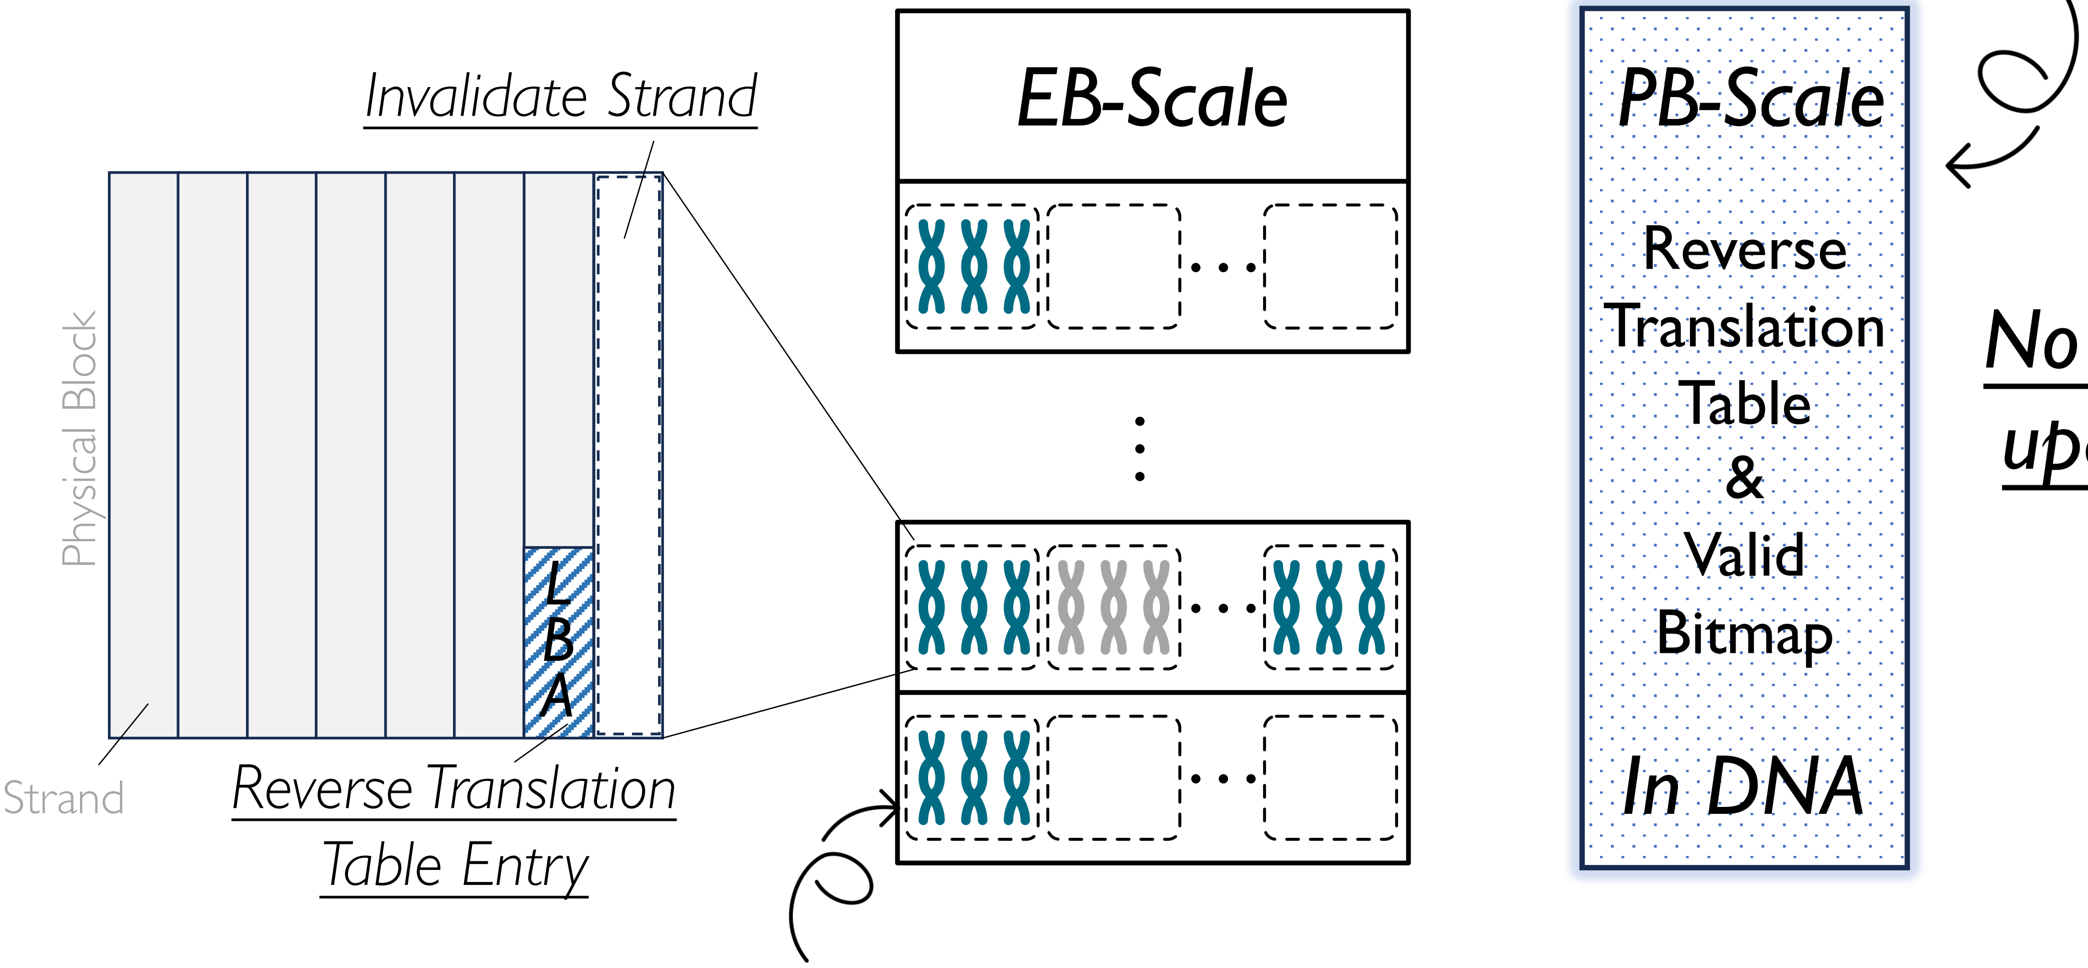

No DNA update!

No DNA update for physical block

Delayed Invalidation

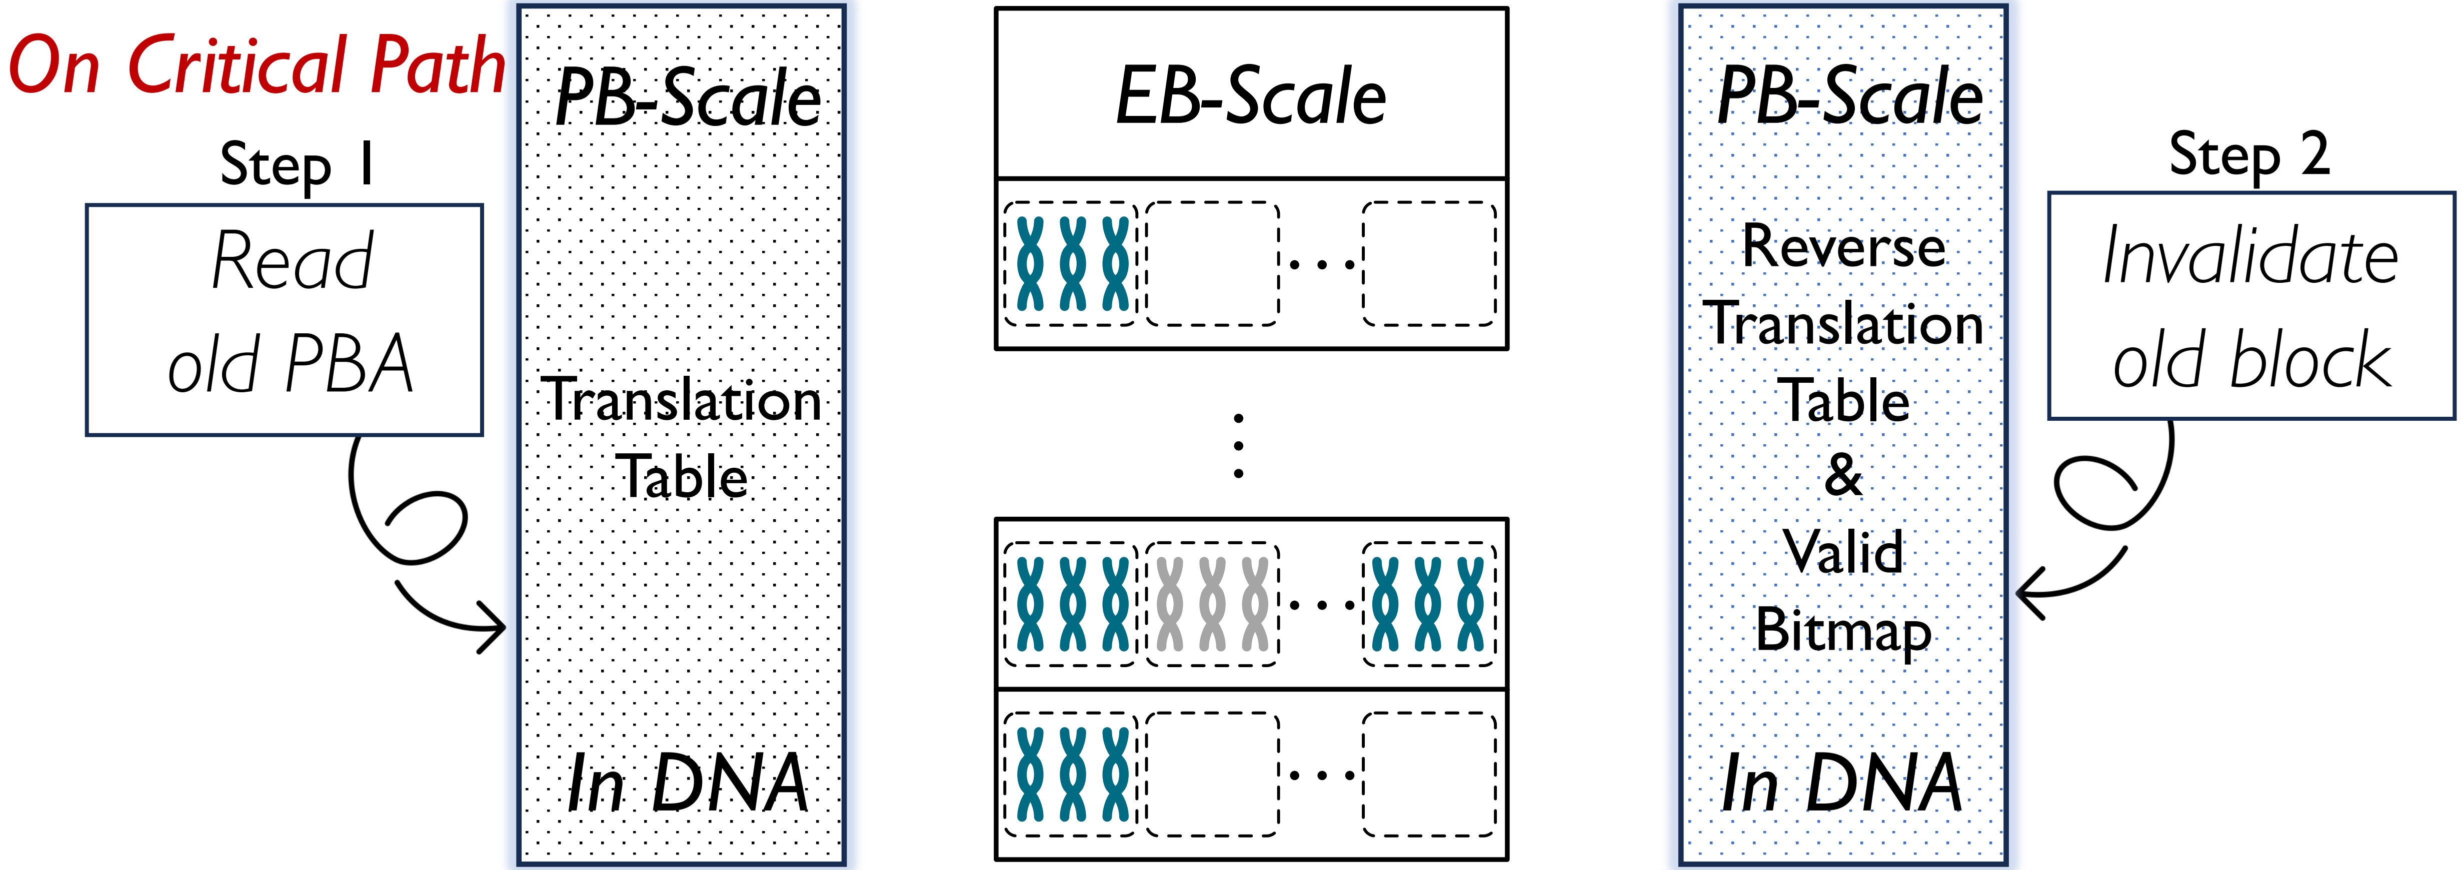

Delayed Invalidation

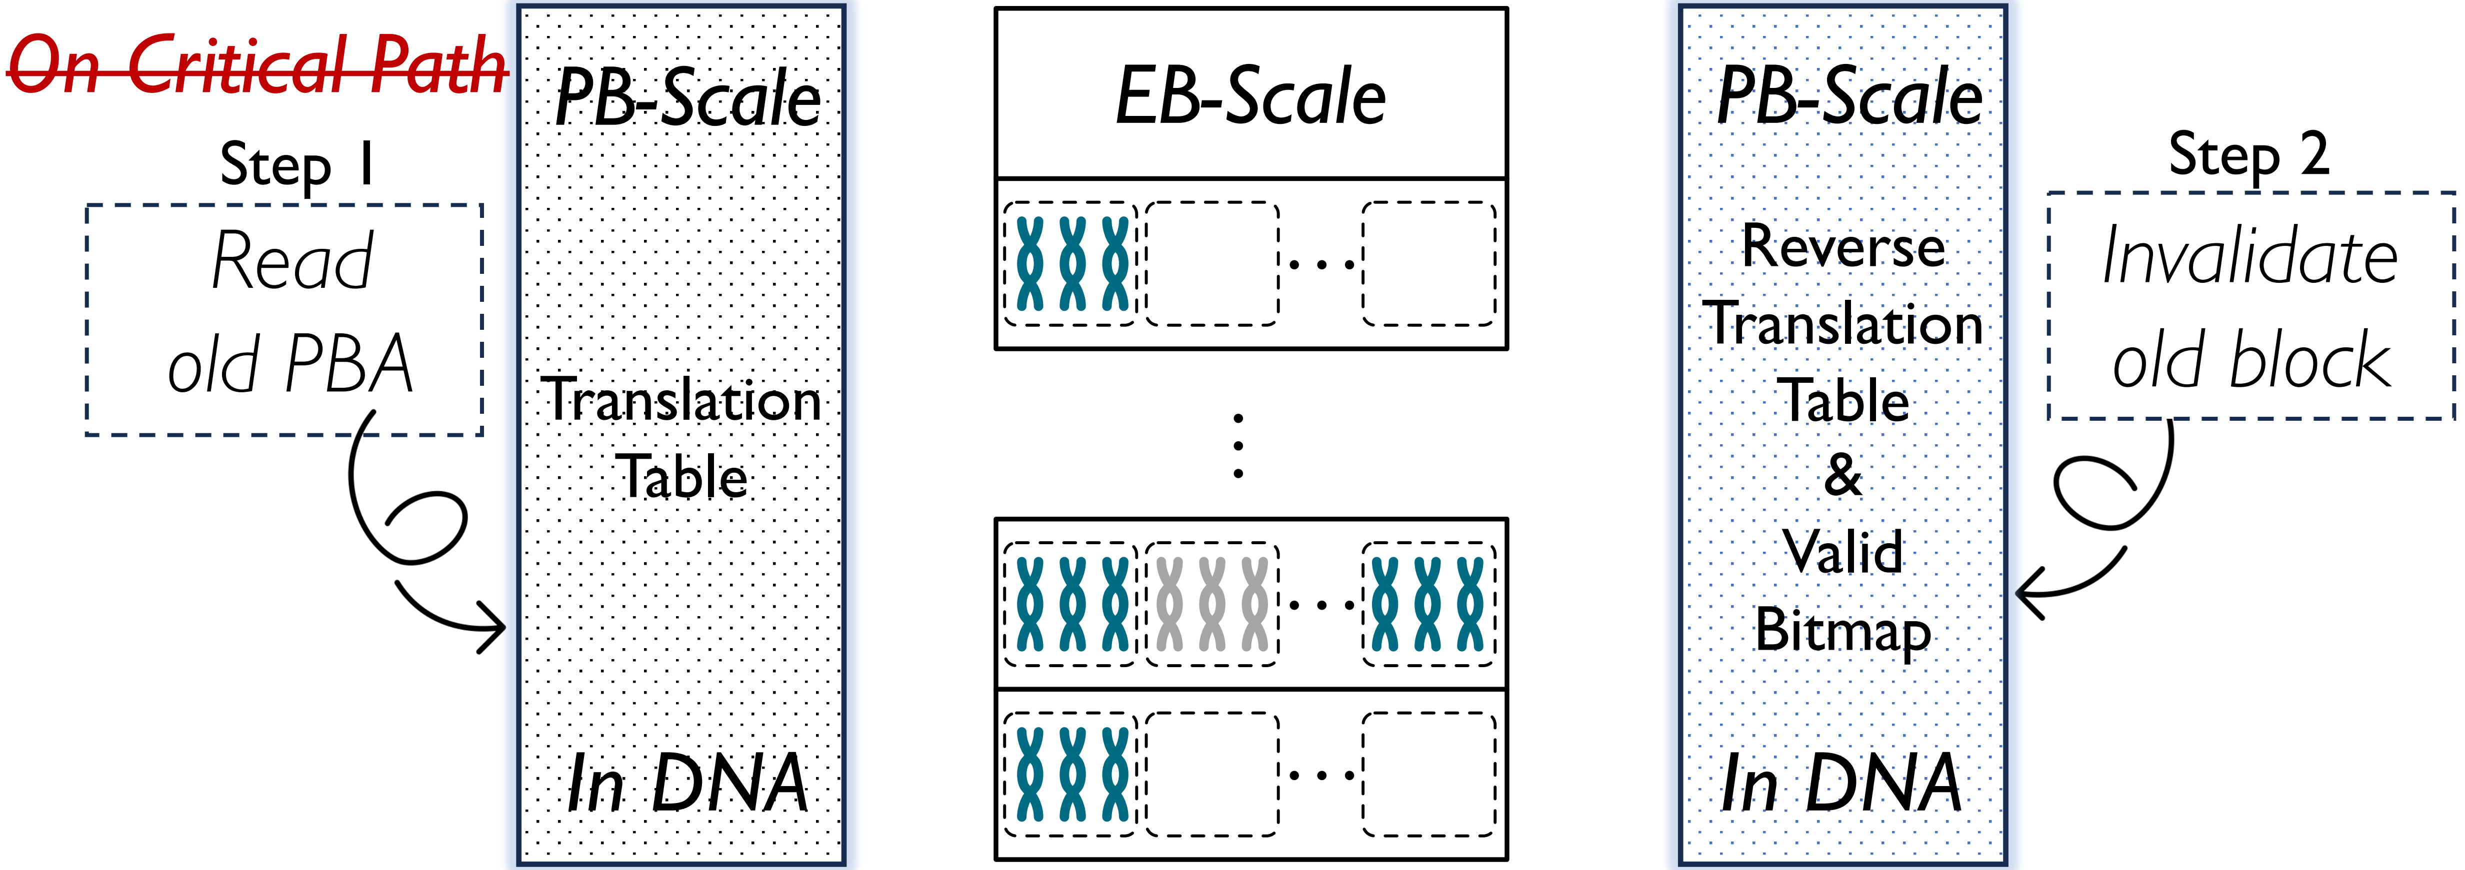

Delay invalidation to block read operation

Delayed Invalidation

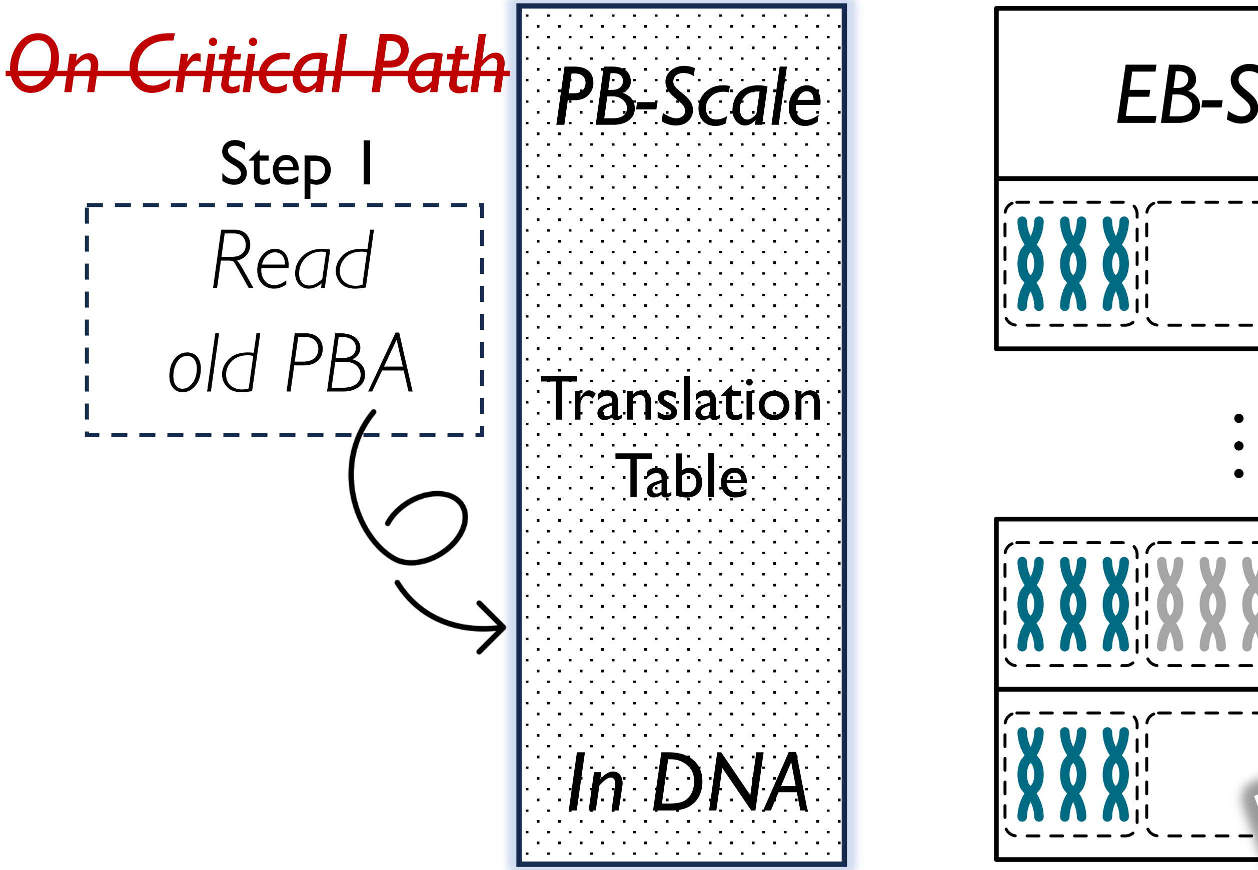

Delay invalidation to b

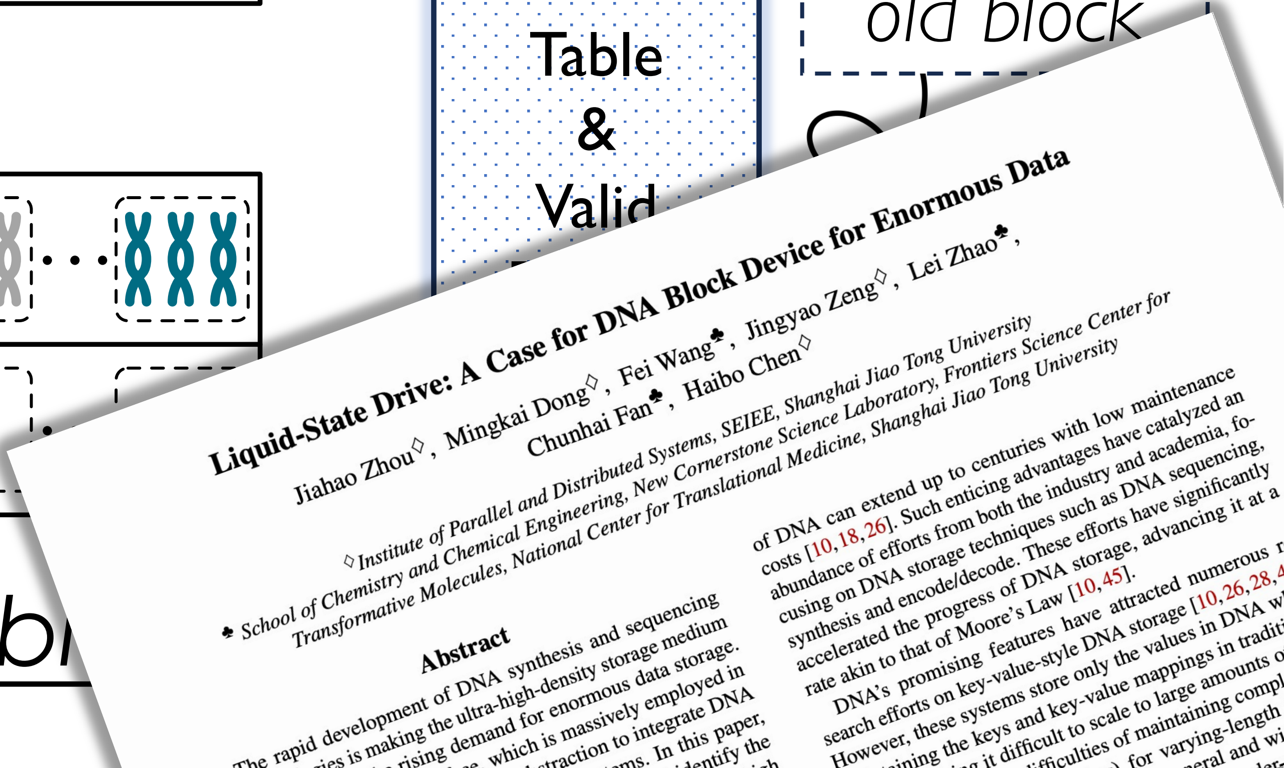

Liquid-State Drive: A Case for DNA Block Device for Enormous Data

Jiahao Zhou◇, Mingkai Dong◇, Fei Wang♣, Jingyao Zeng◇, Lei Zhao♣,
Chunhai Fan♣, Haibo Chen◇

◇Institute of Parallel and Distributed Systems, SEIEE, Shanghai Jiao Tong University
♣ School of Chemistry and Chemical Engineering, New Cornerstone Science Laboratory, Frontiers Science Center for
Transformative Molecules, National Center for Translational Medicine, Shanghai Jiao Tong University

Abstract

The rapid development of DNA synthesis and sequencing ... ies is making the ultra-high-density storage medium ... rising demand for enormous data storage ... which is massively employed in ... traction to integrate DNA ... ems. In this paper, ... identify the

of DNA can extend up to centuries with low maintenance costs [10,18,26]. Such enticing advantages have catalyzed an abundance of efforts from both the industry and academia, focusing on DNA storage techniques such as DNA sequencing, synthesis and encode/decode. These efforts have significantly accelerated the progress of DNA storage, advancing it at a rate akin to that of Moore's Law [10,45].

DNA's promising features have attracted numerous r... search efforts on key-value-style DNA storage [10,26,28,4... However, these systems store only the values in DNA w... taining the keys and key-value mappings in tradit... ing it difficult to scale to large amounts o... difficulties of maintaining compl...) for varying-length ... general and wi...

Outline

- DNA Storage Basics
- A Naïve DNA Block Design
- LiqSD Design
- Evaluation
- Conclusion

Evaluation

- Simulator-based
- Compared Systems
  - LiqSD – Our system
  - Coarse-DTL – 24 MB block size, DTL in SSD
  - Naïve Design
- Metrics
  - Read/Write Amplification

$$\frac{Request\ Read/Write\ Volume}{Actual\ Read/Write\ Volume}$$

  - Extra Read Ratio

$$\frac{Request\ Read\ Volume\ Caused\ by\ Write}{Actual\ Write\ Volume}$$

Evaluation: Microbenchmark

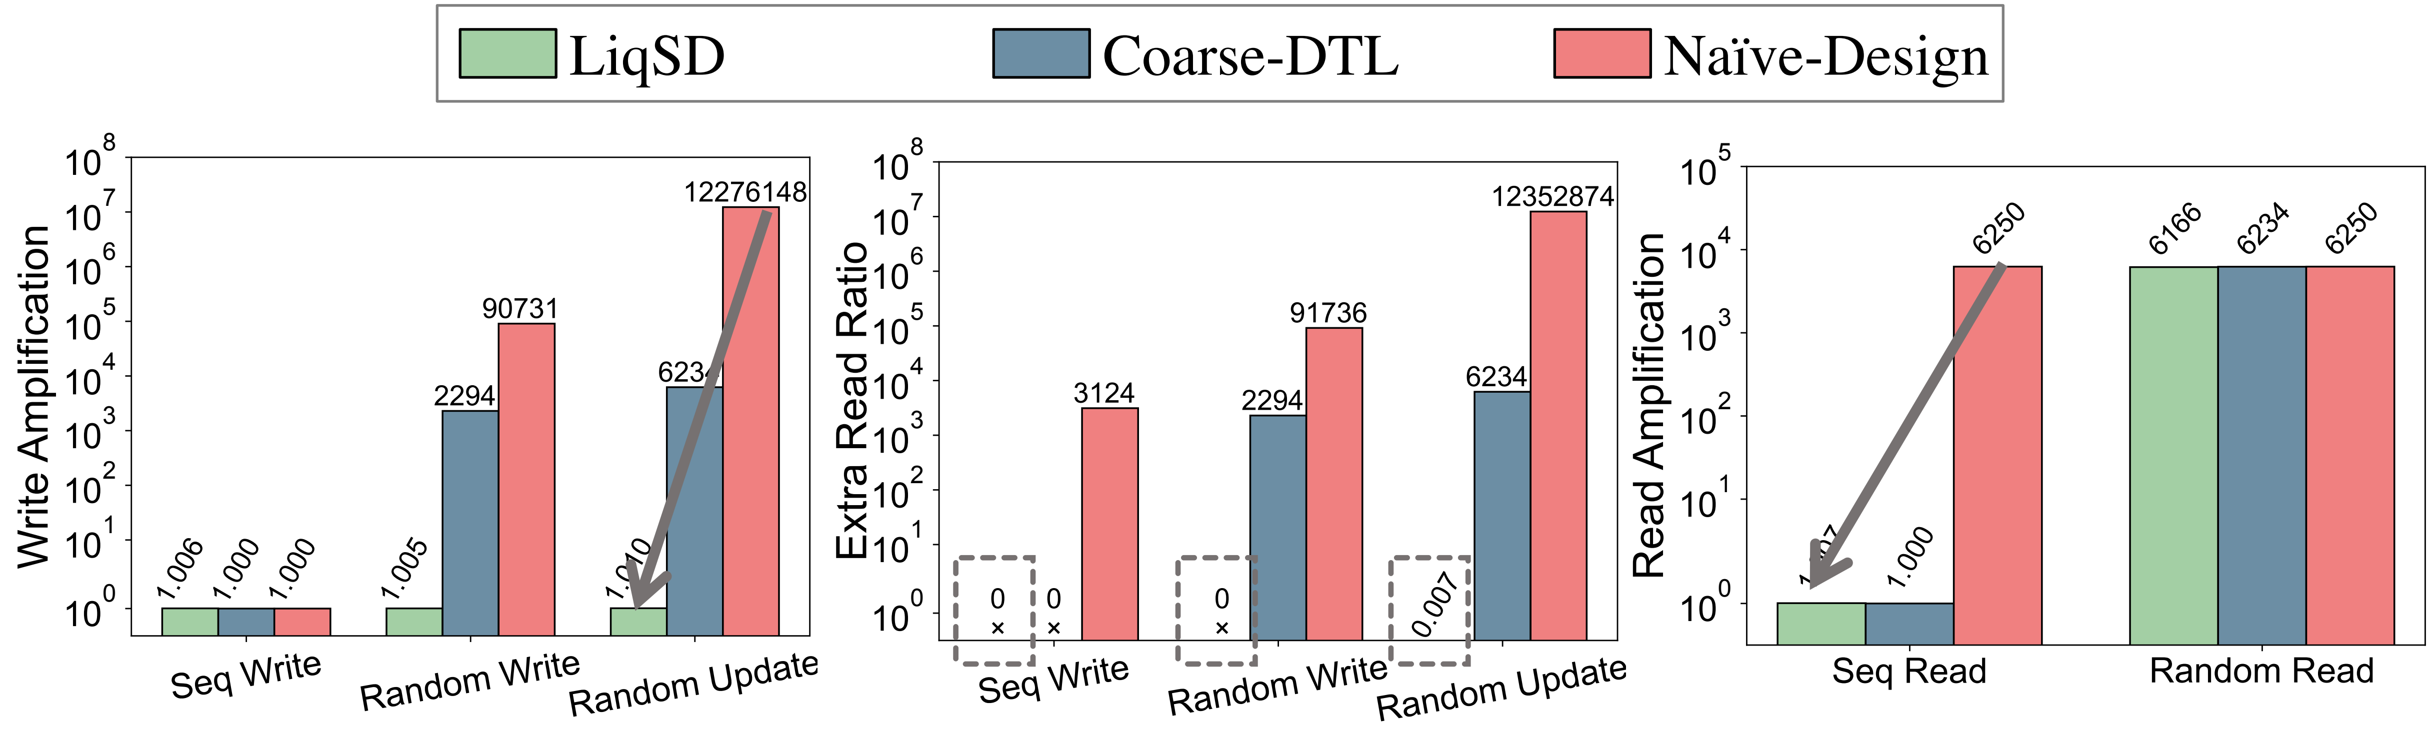

Write amplification reduce up to $10^7\times$

Extra read ratio Almost zero

Read amplification Reduce up to 6,206 ×

Evaluation: Real-world Traces

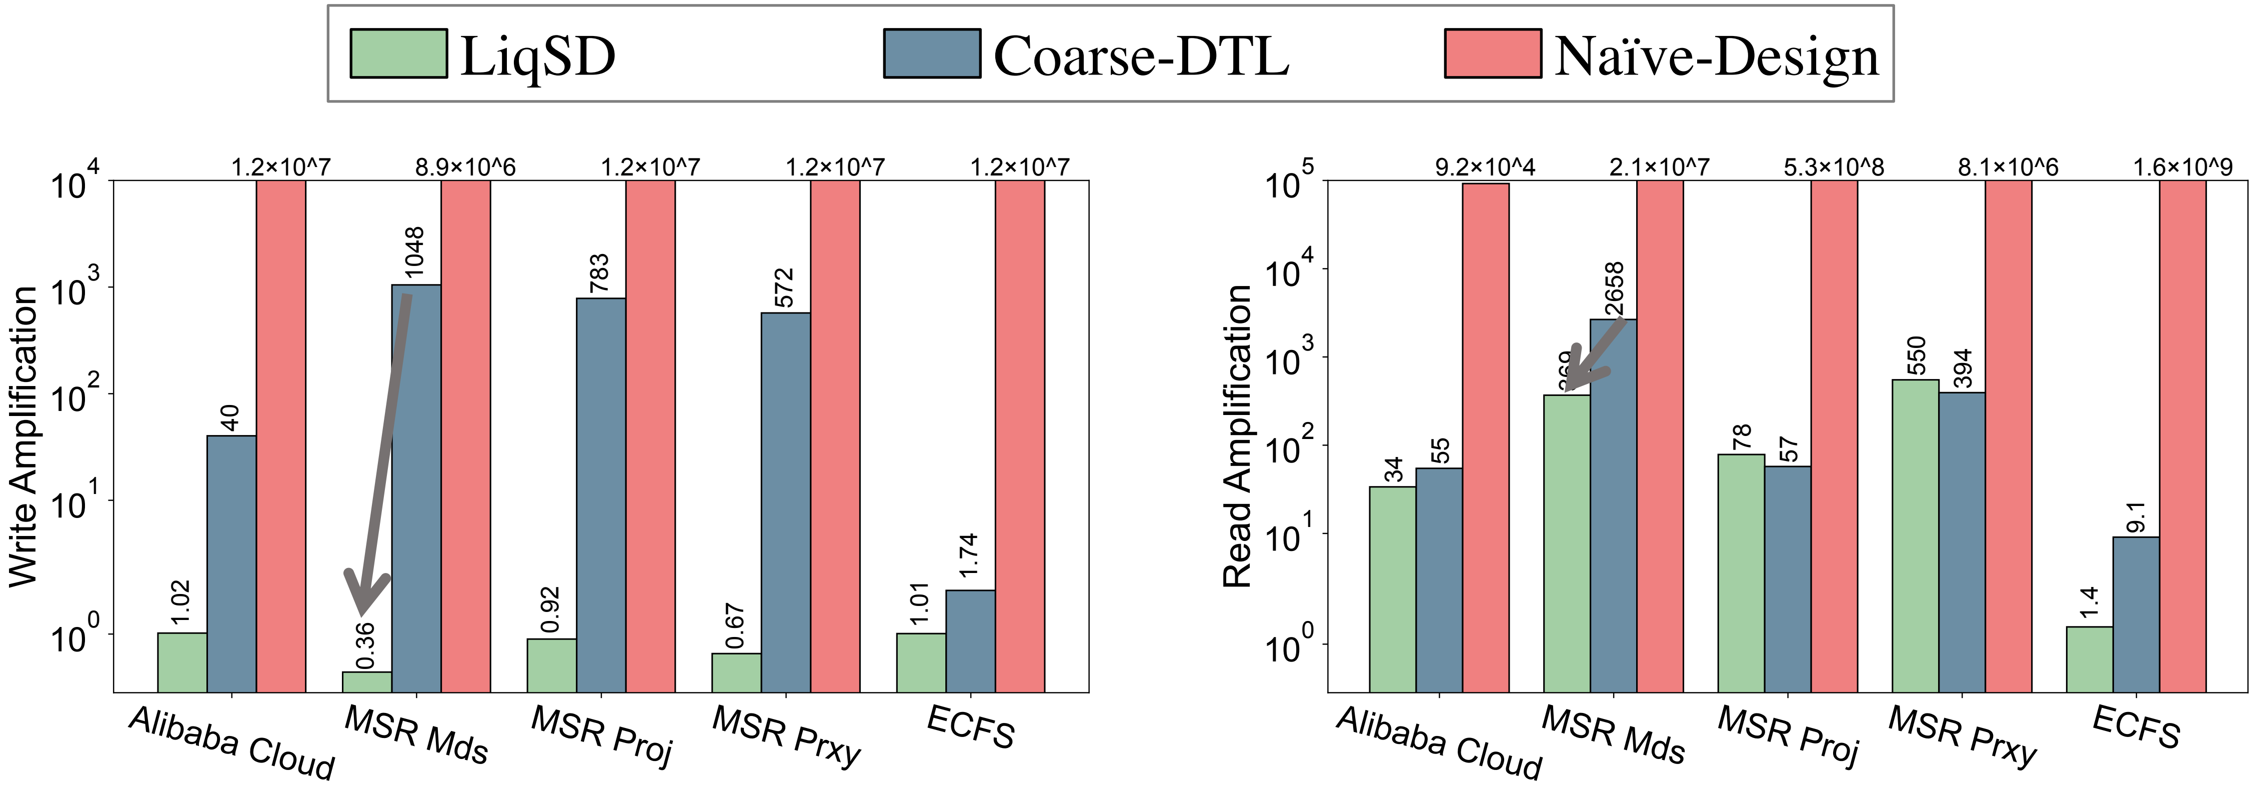

Compared to Coarse-DTL,
Reduce up to 2,927×

Compared to Coarse-DTL,
Reduce up to 7×

Evaluation: Different Block Size

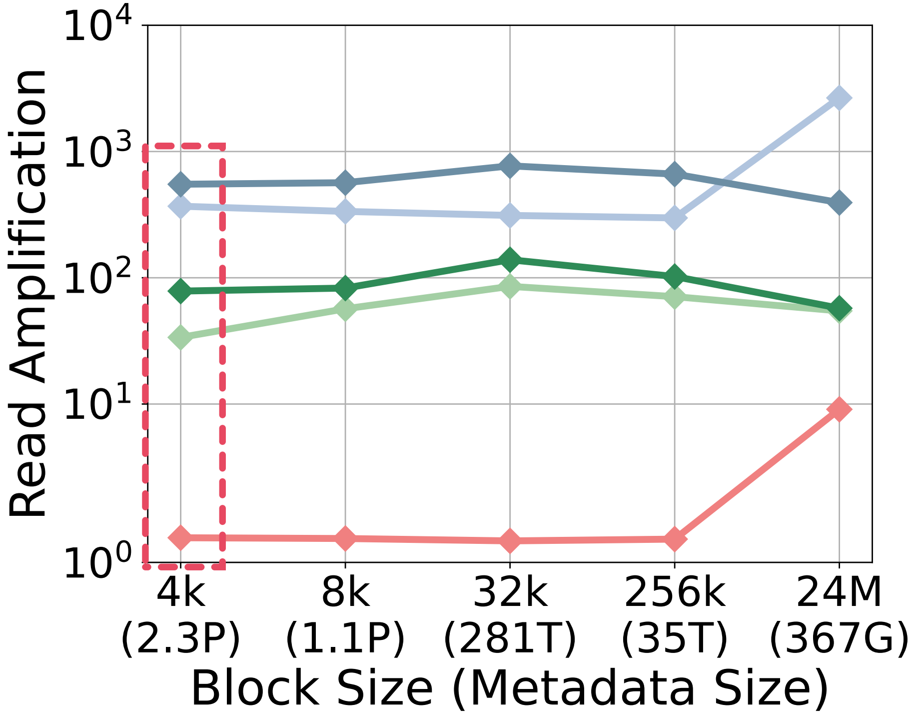

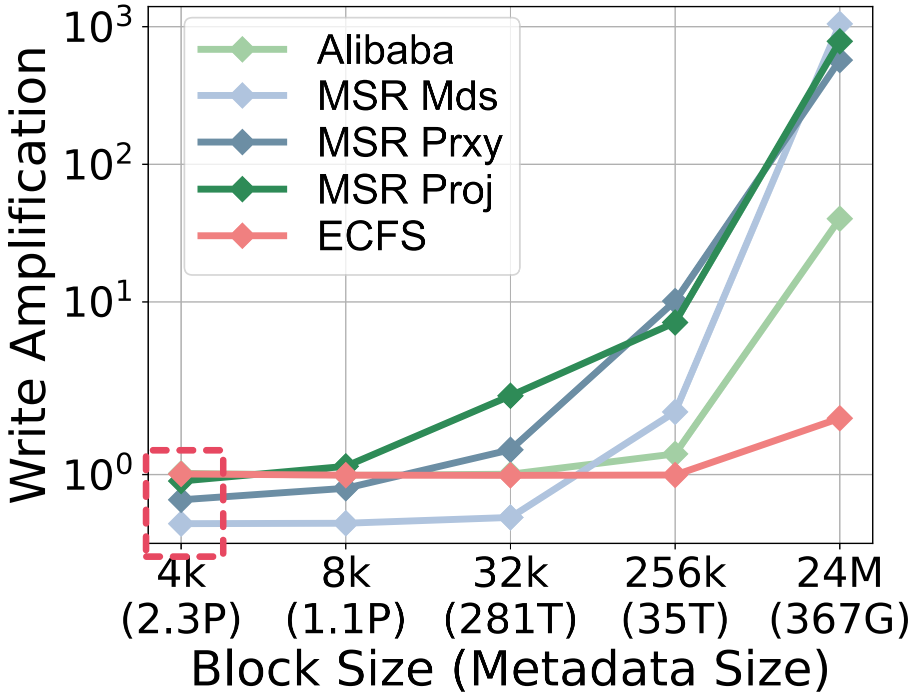

4 KiB block size achieves the best performance

Conclusion

- A holistic DNA block device design
- Key idea: Out-of-place update
- Challenge: Maintain the in-DNA metadata is costly
- Key techniques:
  - Dual Translation Layer
  - Symbiotic Metadata
  - Delayed Invalidation