Toward Privacy-Aware Data Systems

navigating the **privacy-performance tradeoff**

Subhadeep Sarkar

Brandeis UNIVERSITY

https://subhadeep.net/

lab

Every two days we generate as much data as we did since the dawn of humanity until 2003.

— **Eric Schmidt** (CEO, Google), 2010

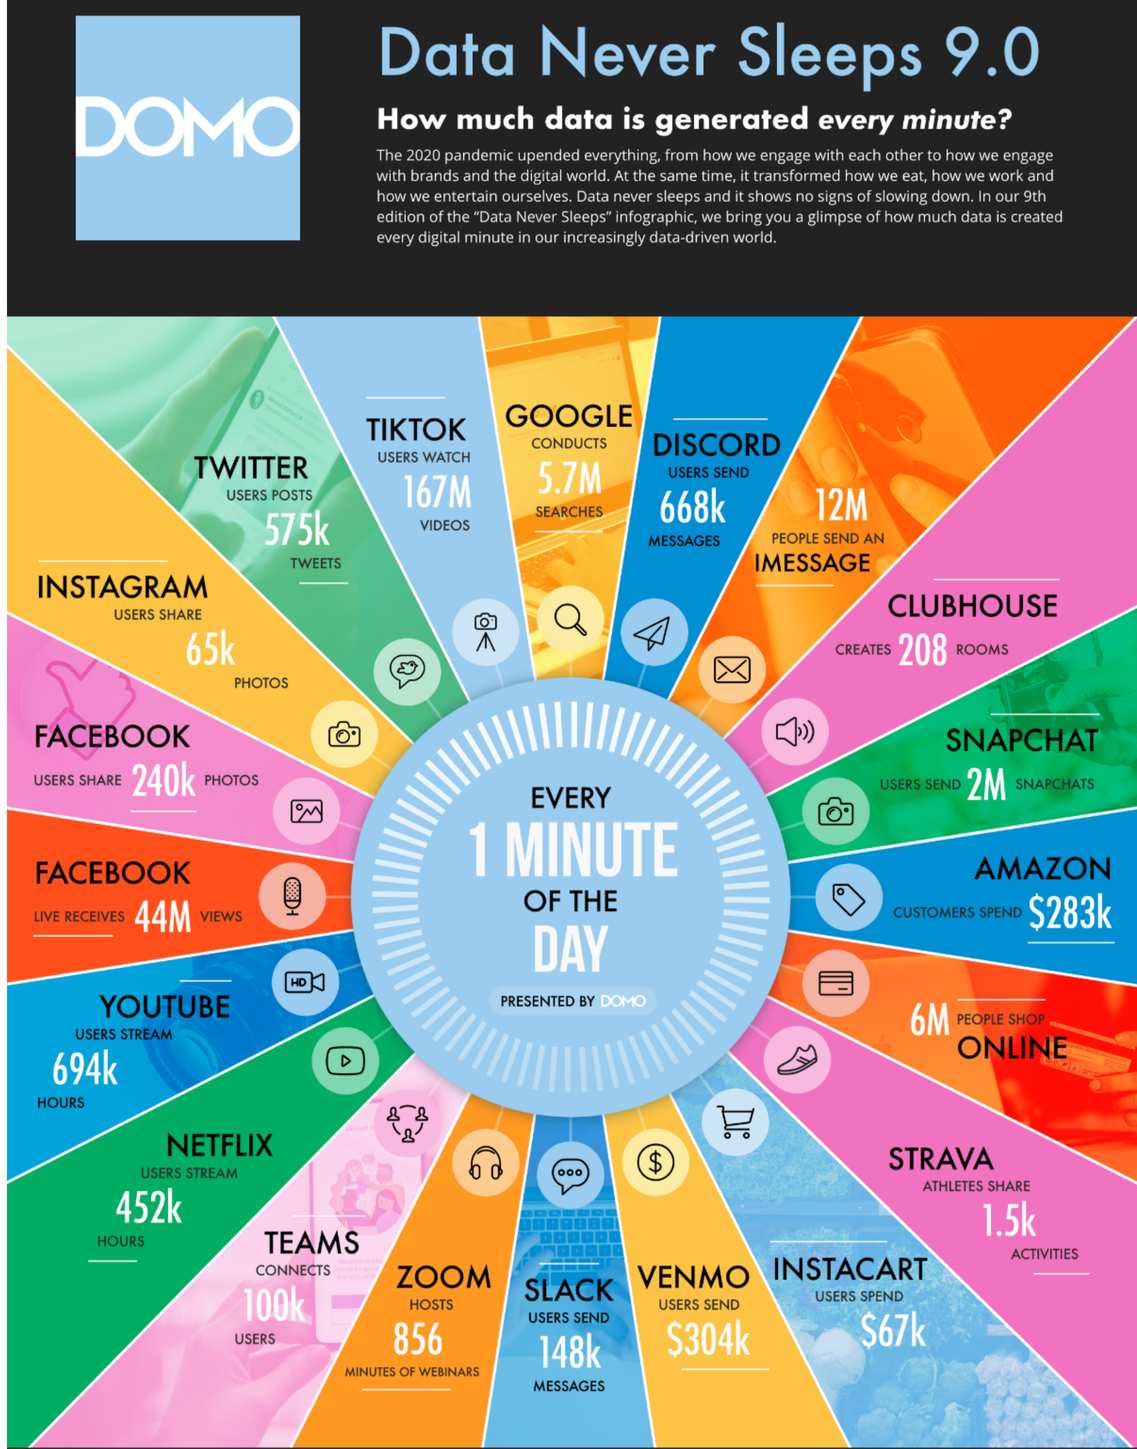

Ingestion-Optimized Systems

batched inserts

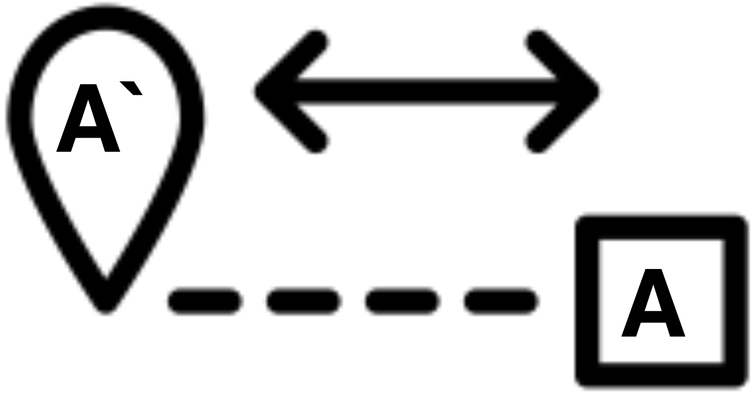

out-of-place
deletes/updates

Out-of-place Deletes/Updates

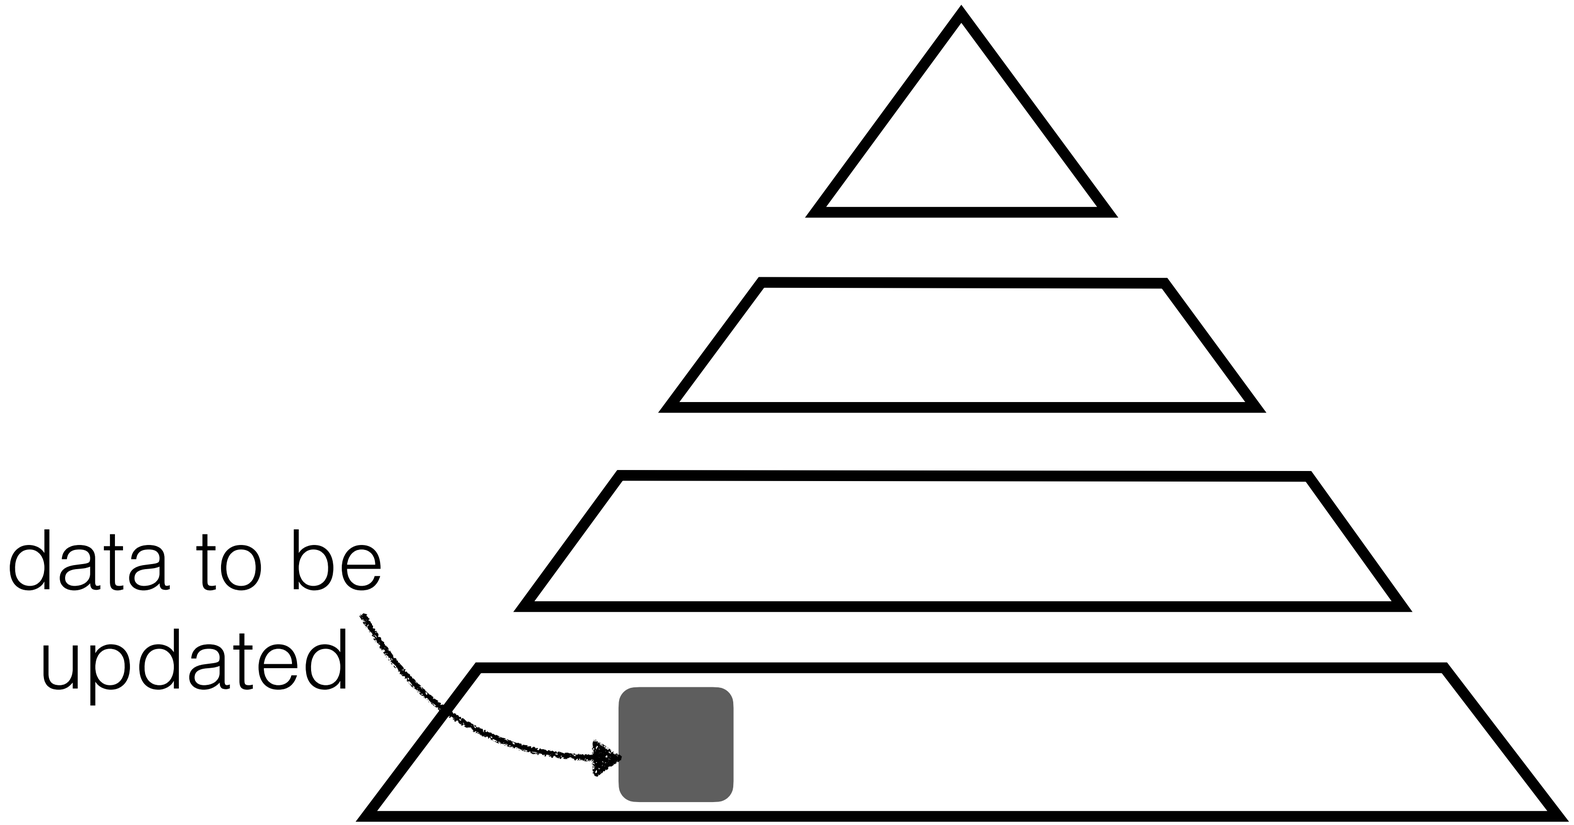

Out-of-place Deletes/Updates

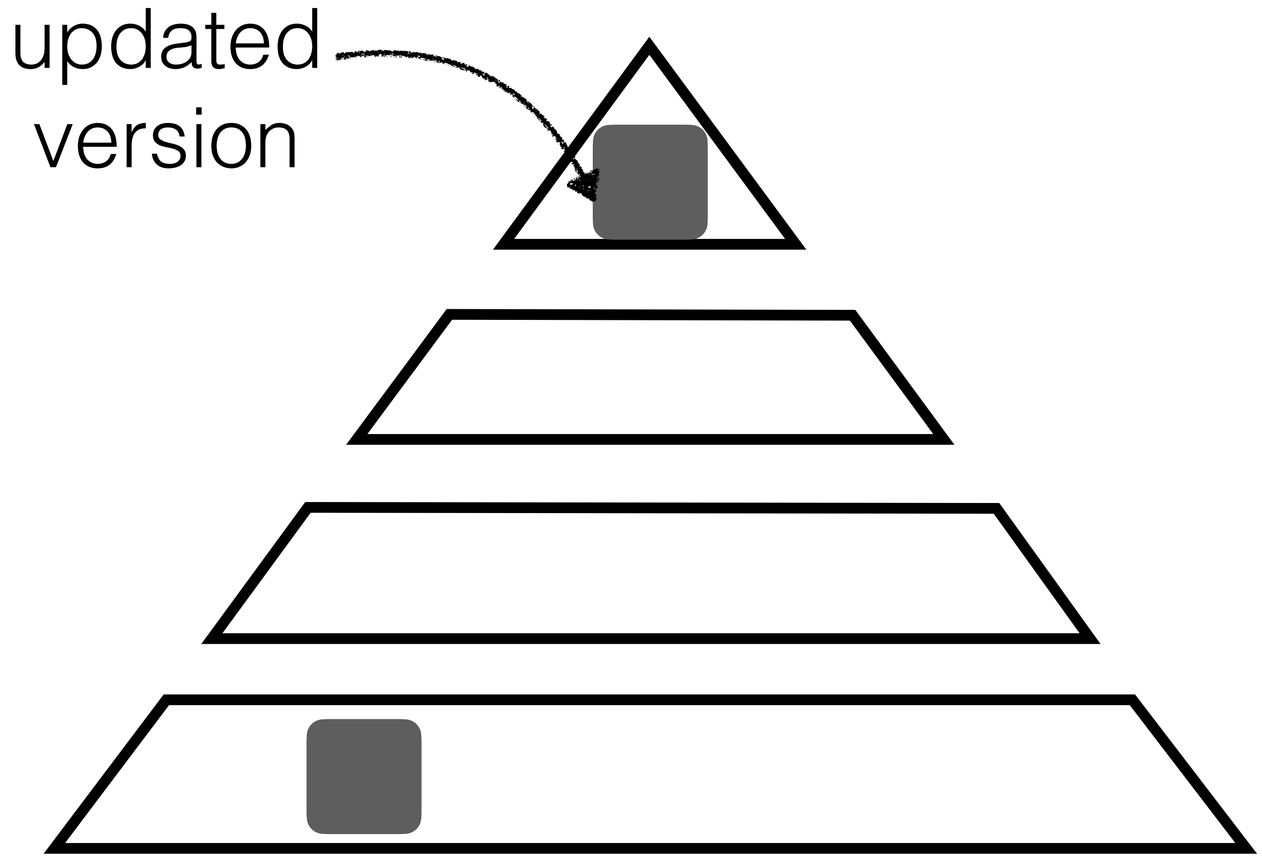

Out-of-place Deletes/Updates

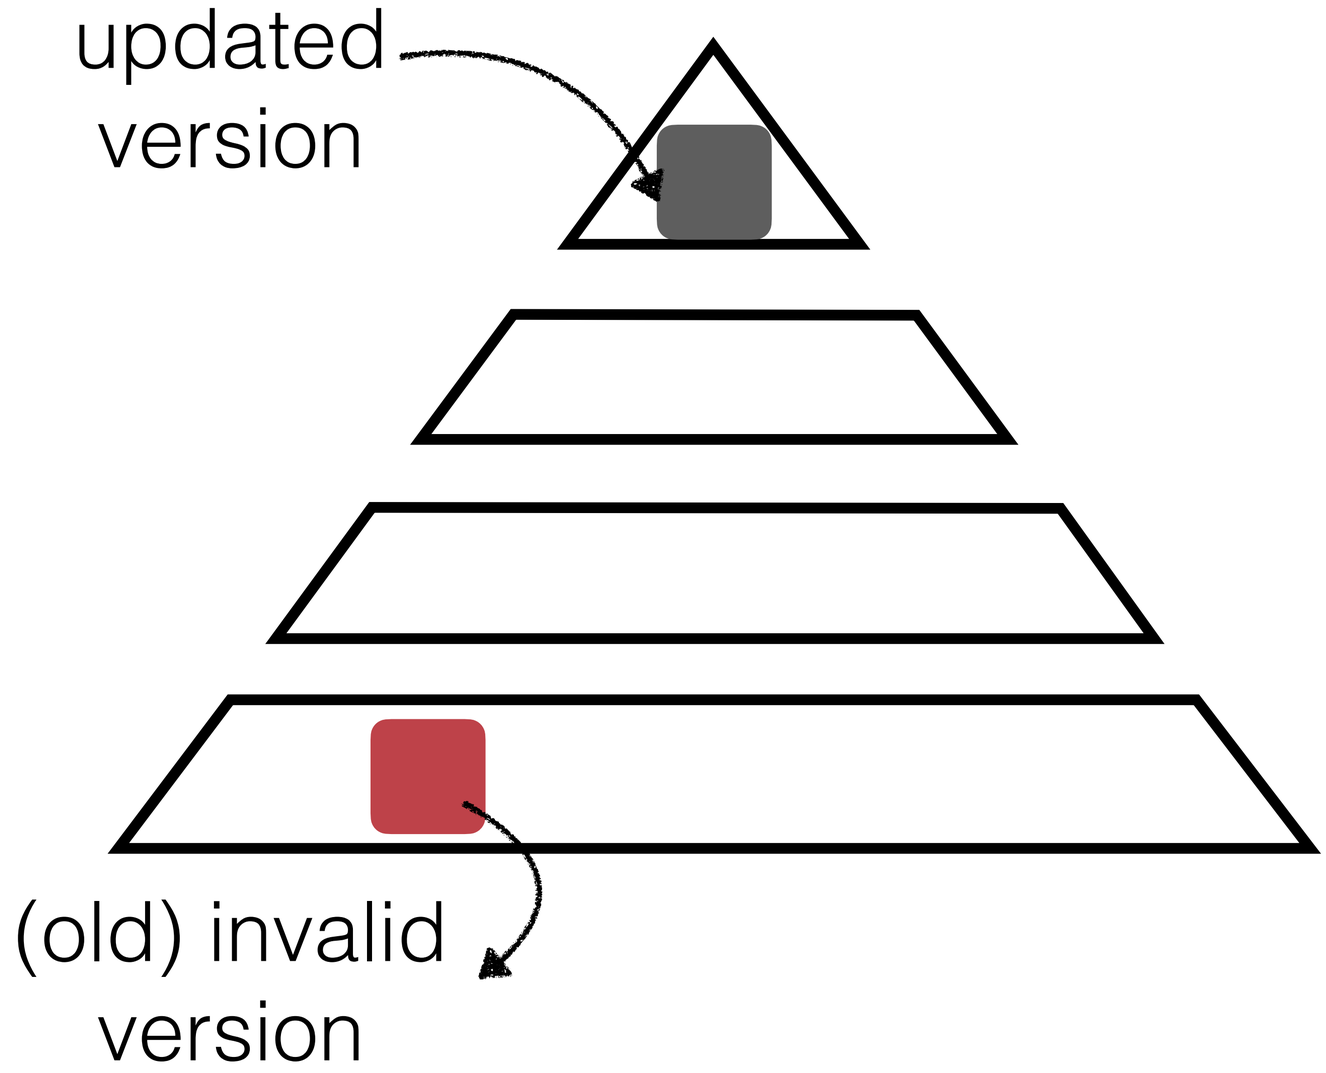

Out-of-place Deletes/Updates

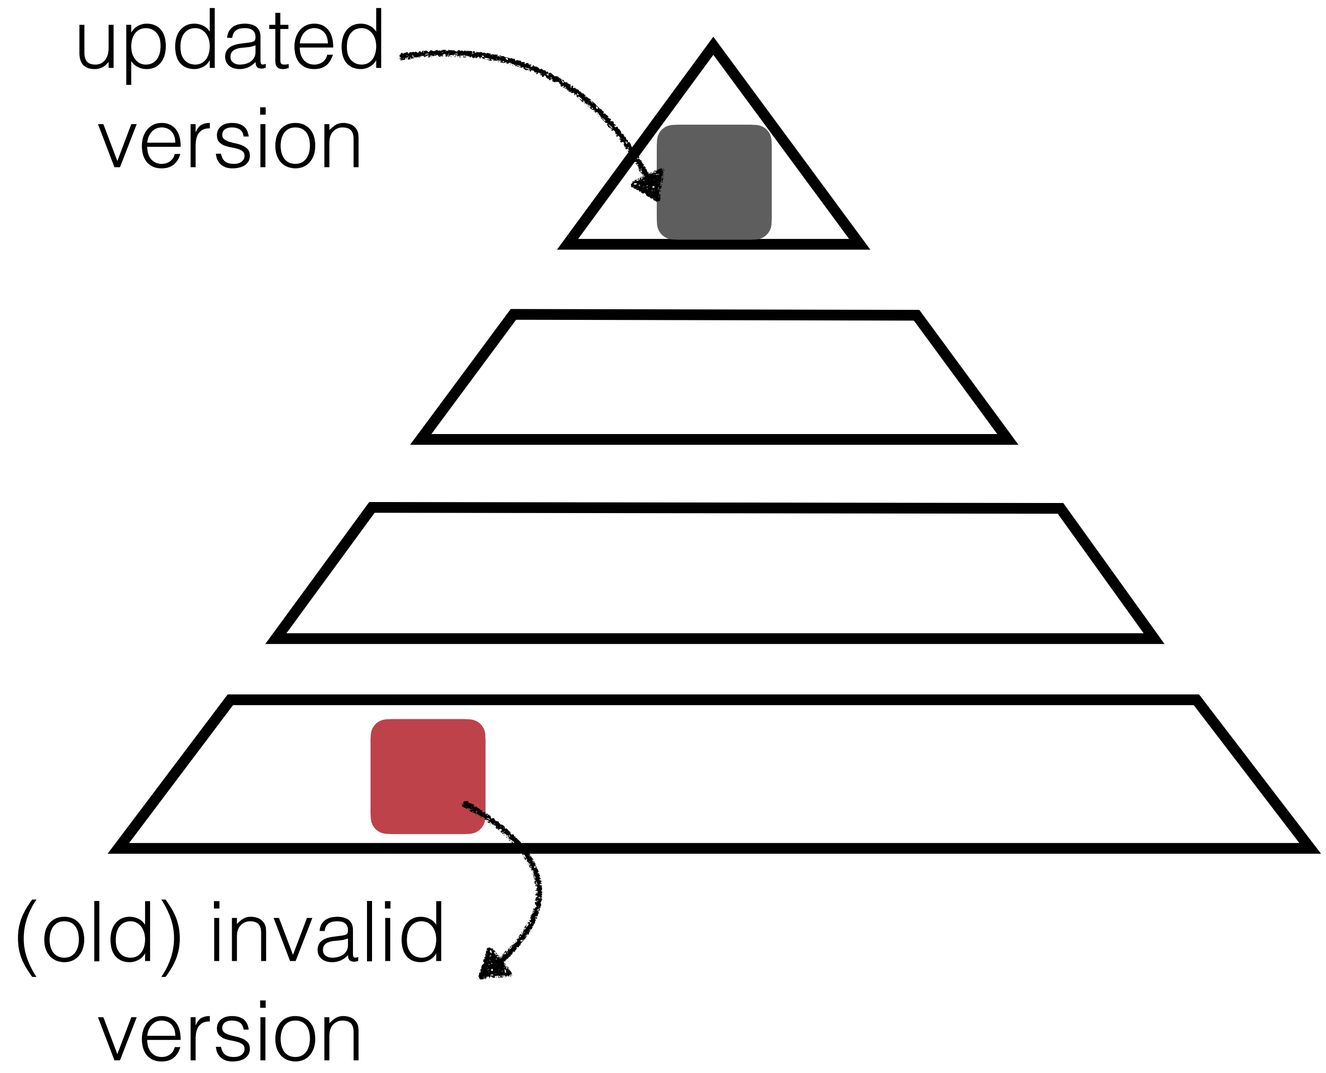

Out-of-place Data Systems

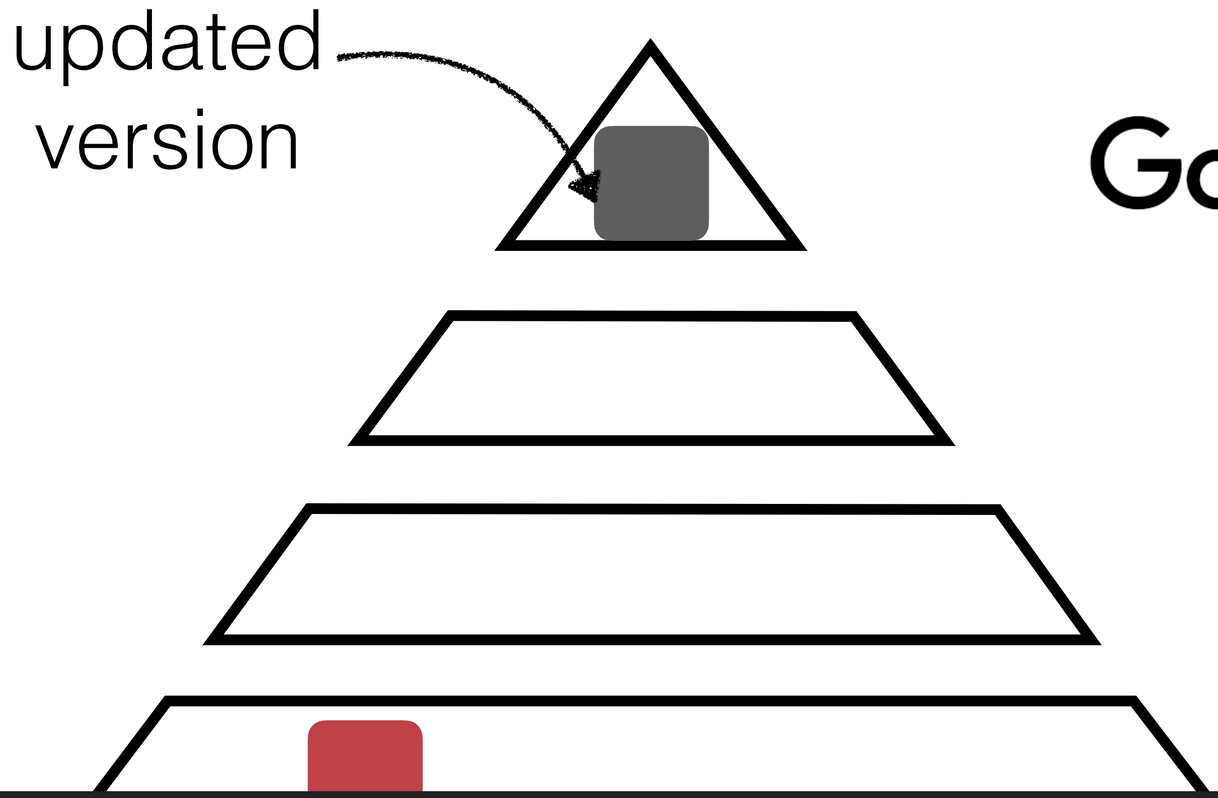

Google Meta amazon

THE APACHE® SOFTWARE FOUNDATION SAP®

mongoDB

Hidden Cost: does not scale with deletes!

Hidden Cost of Logical Deletes

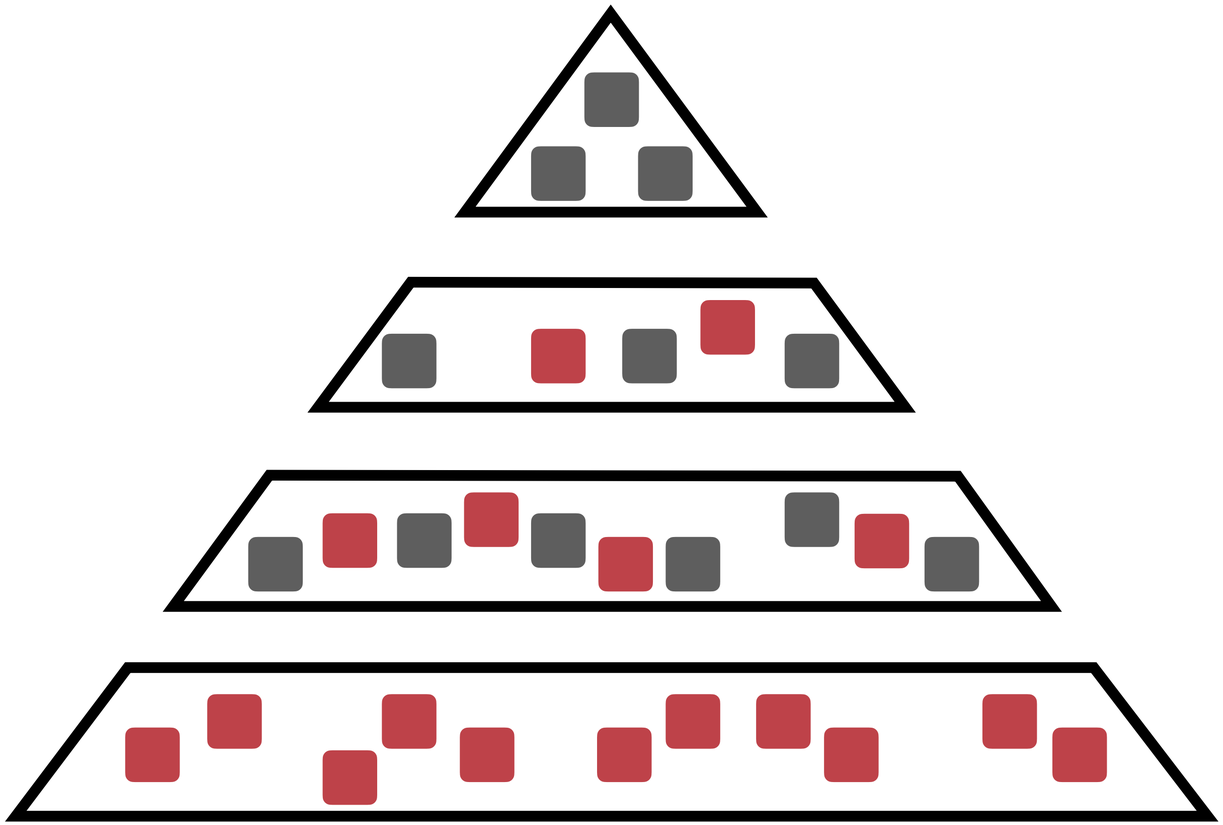

delete-markers / tombstones

logically invalidated data

up to 2x

high space amplification

1.5x - 5x

poor read performance

increased operational cost

Logical Deletes & **Data Privacy**

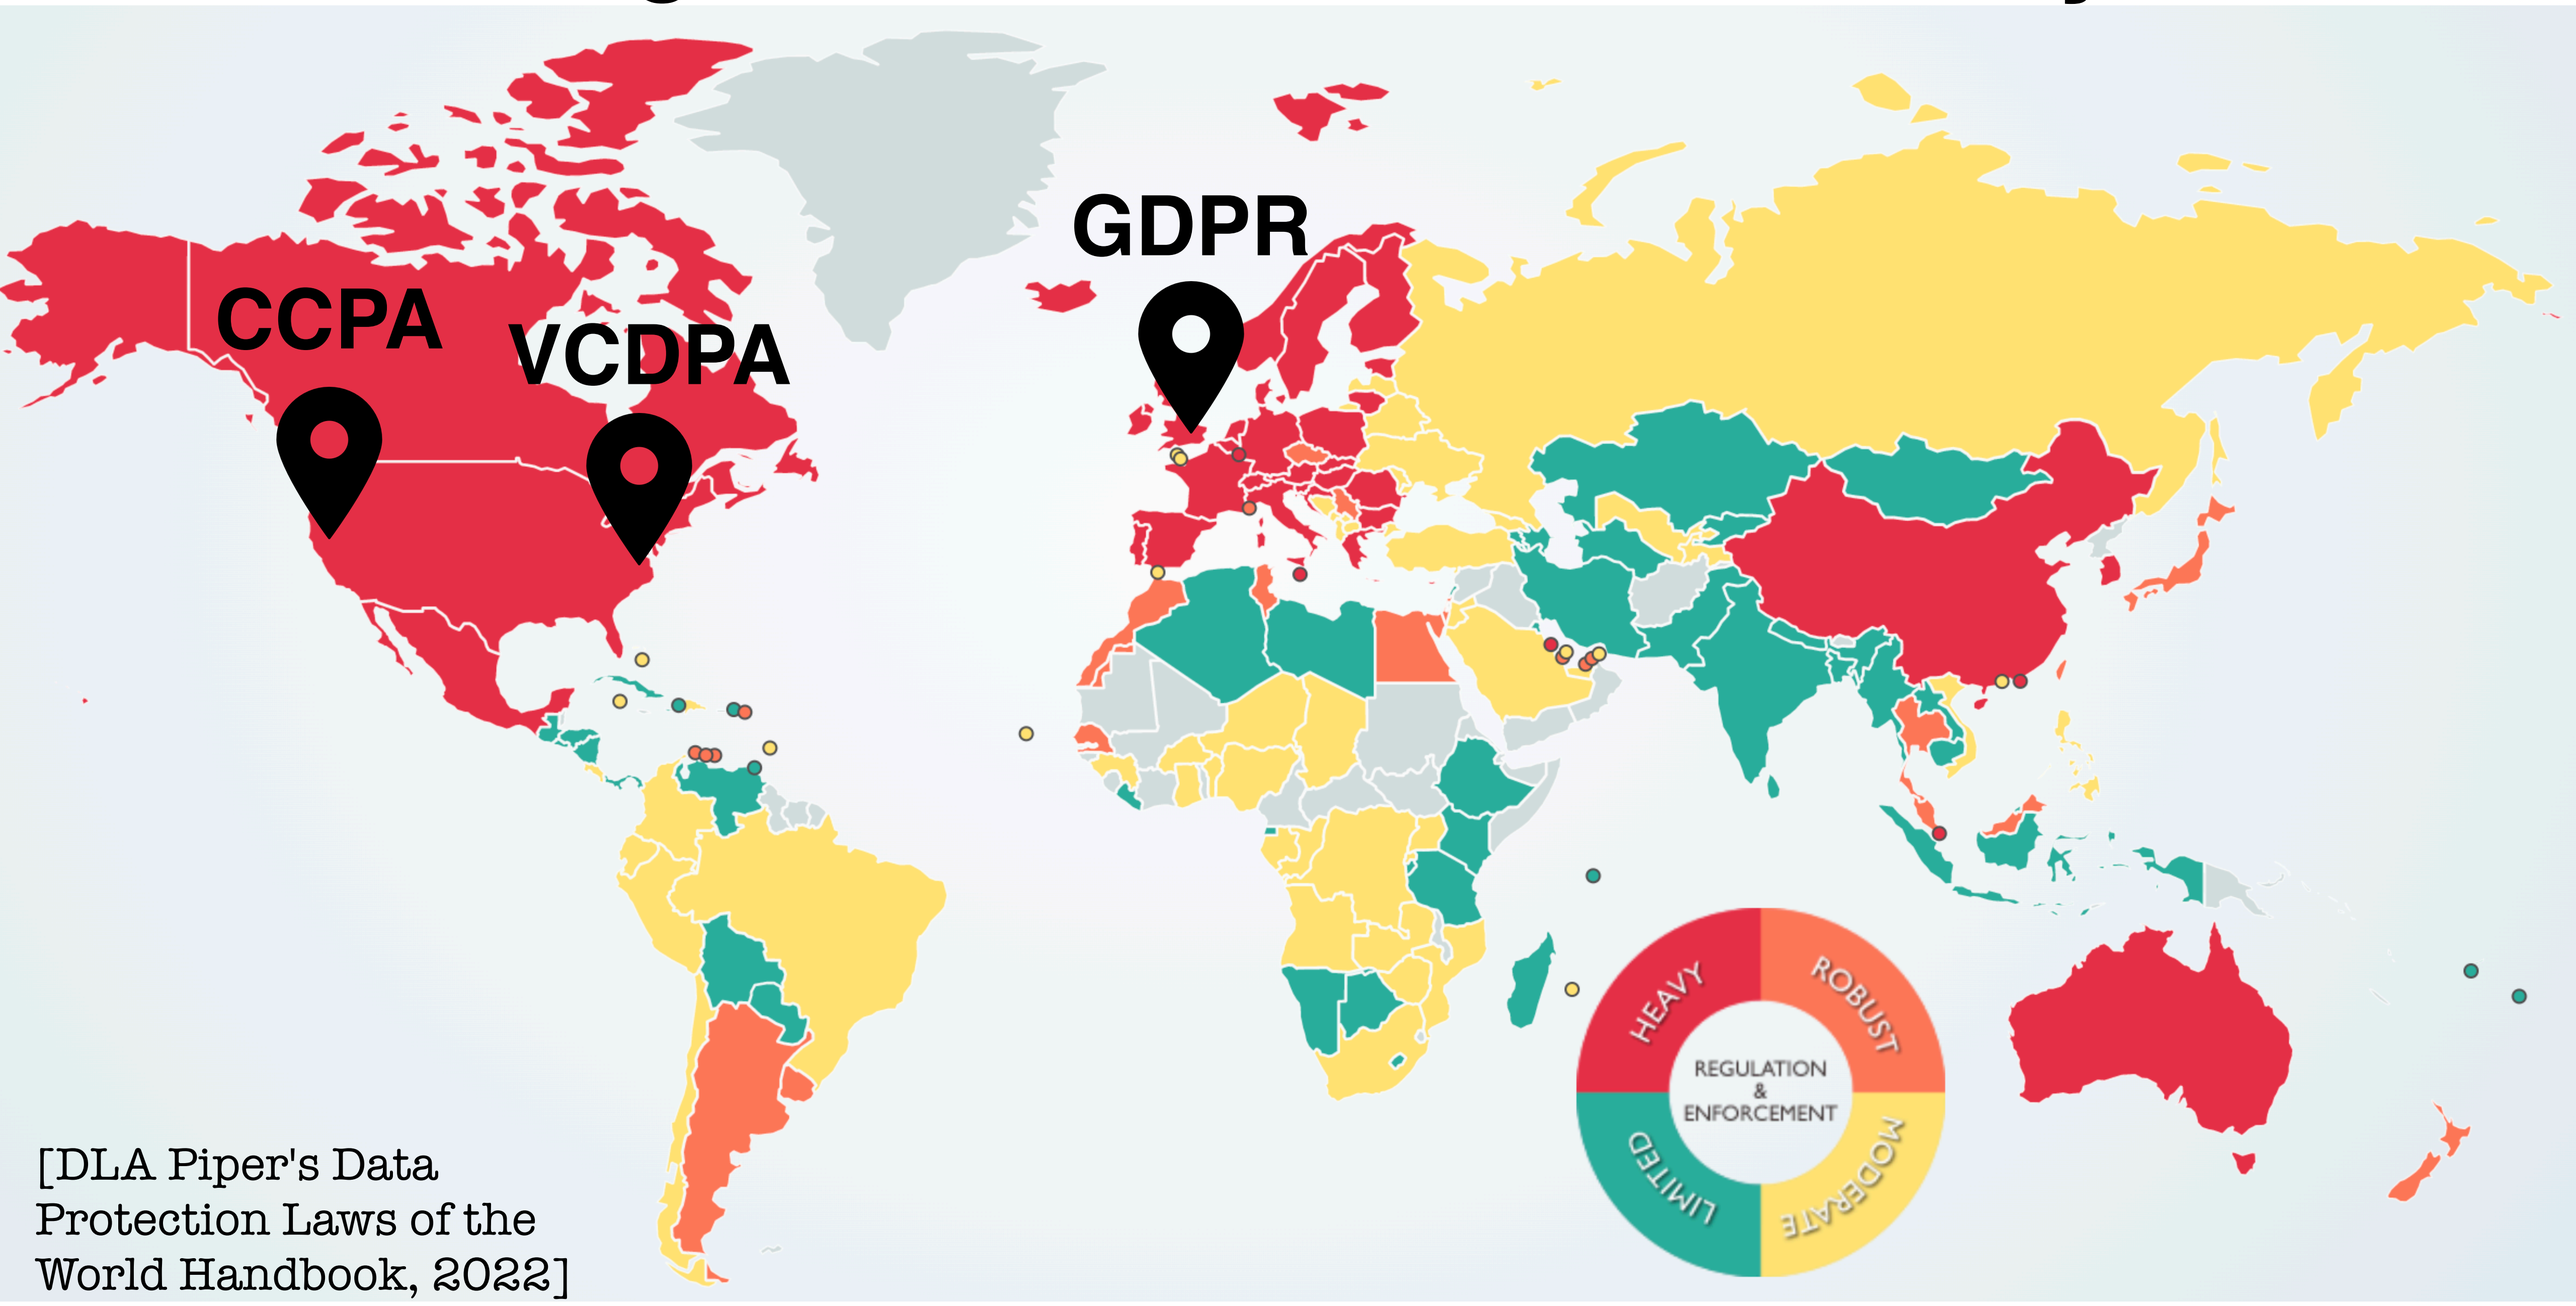

[DLA Piper's Data Protection Laws of the World Handbook, 2022]

GDPR
(EU, UK)

Right to be forgotten

CCPA
(California)

Right to delete

VCDPA
(Virginia)

Deletion right

timely
deletes

+

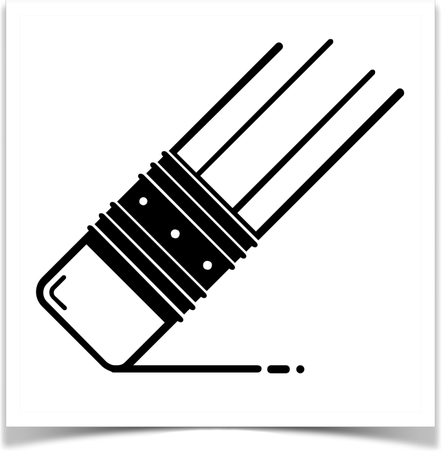

persistent
deletes

Right to be forgotten

Right to delete

Deletion right

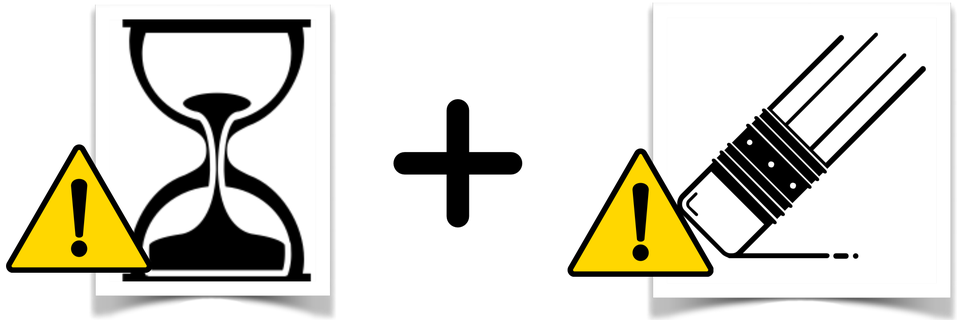

timely deletes + persistent deletes

Even years later, Twitter doesn't delete your direct messages

Zack Whittaker, Natasha Lomas / 1:57 PM EST • February 15, 2019

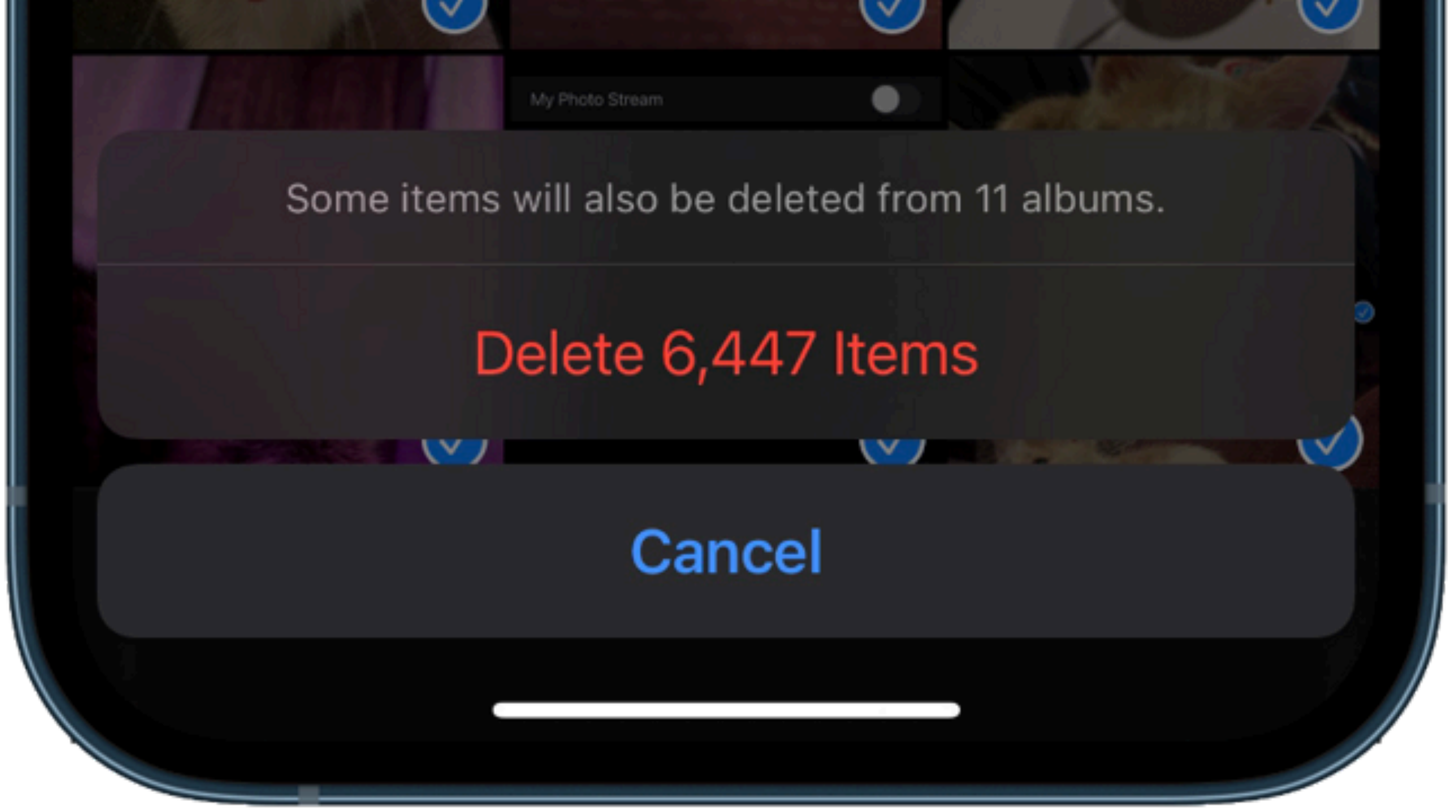
My Photo Stream
Some items will also be deleted from 11 albums.
Delete 6,447 Items
Cancel

100M+ deletes/day

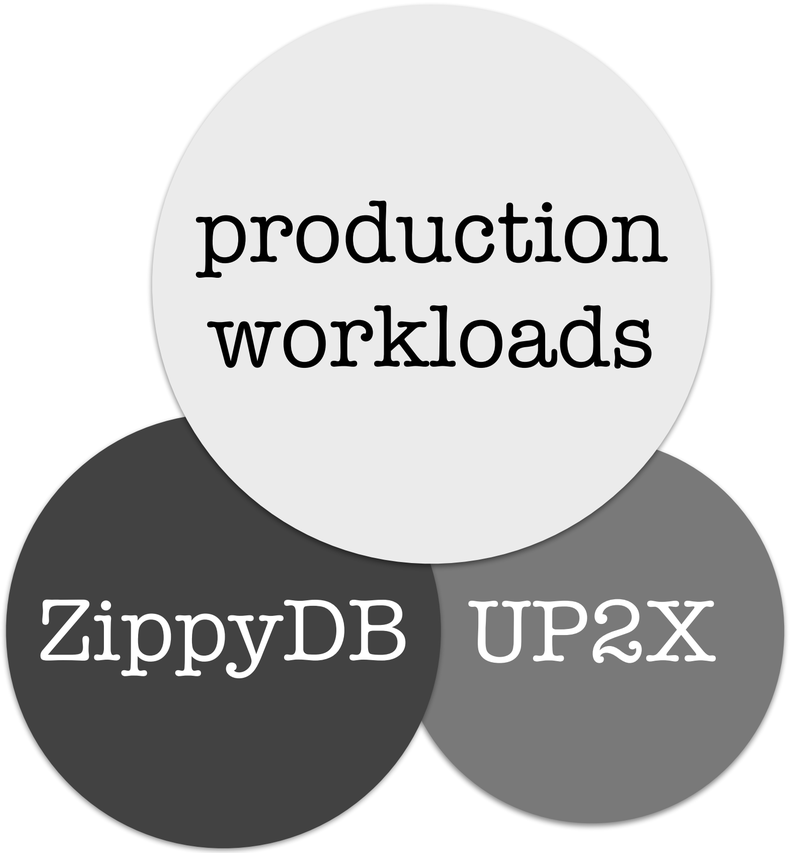

deletes in **batches**

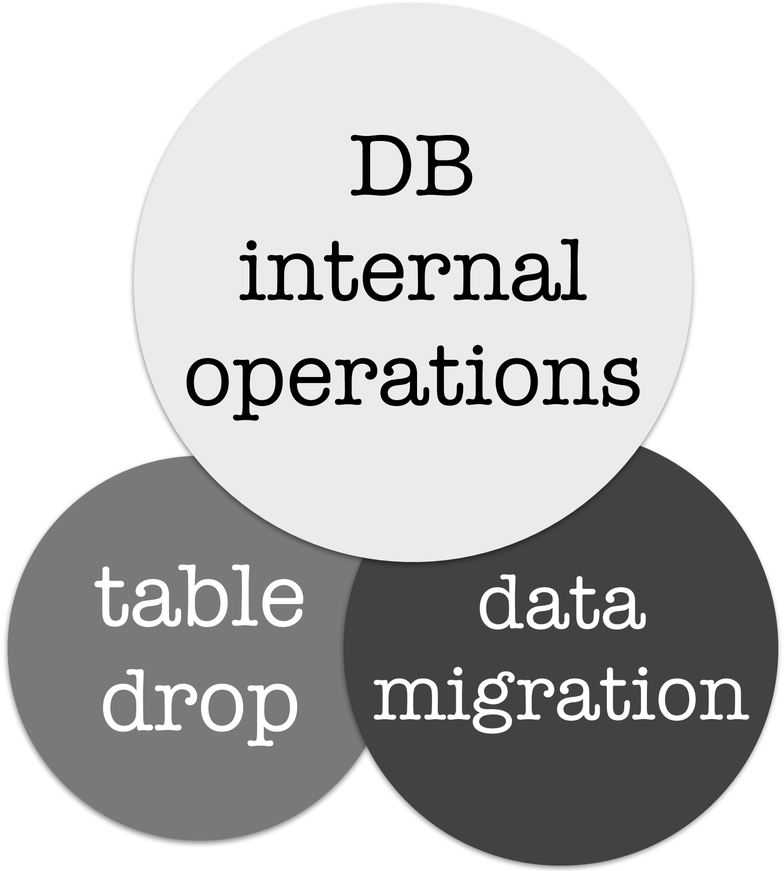

user-generated deletes

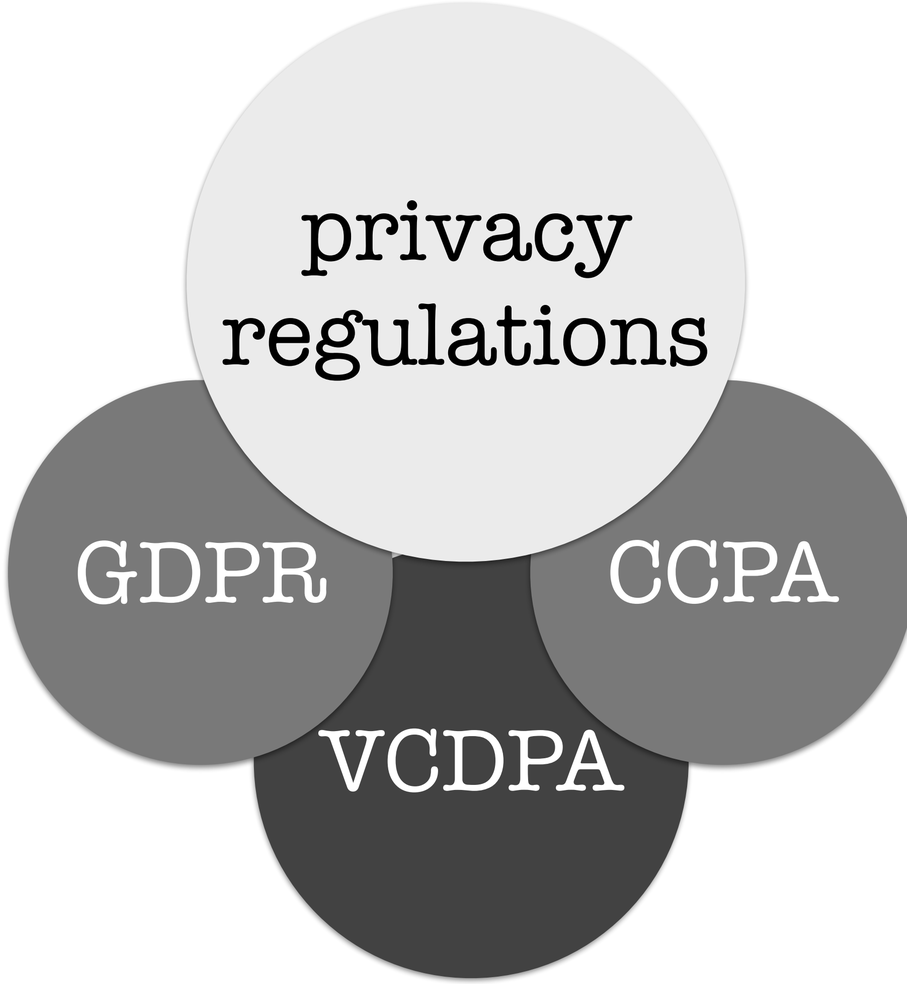

Goal: Enabling **Privacy through Deletion**

dissecting privacy policies and regulations

IEEE DEBull 2022

EDBT 2022

IEEE iThings 2018

fundamental design changes in data layouts and access methods

ICDE 2023-a

TPCTC 2022

SIGMOD 2022-b

VLDB 2021

navigating the privacy-performance tradeoff

ICDE 2023-c

ACM TODS 2023-b

SIGMOD 2022-a

SIGMOD 2020

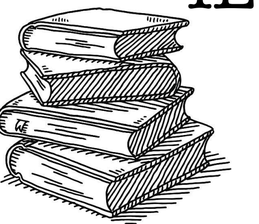

Goal: Enabling **Privacy through Deletion**

dissecting privacy policies and regulations

fundamental design changes in data layouts and access methods

navigating the privacy-performance tradeoff

IEEE DEBull 2022

EDBT 2022

IEEE iThings 2018

ICDE 2023-a

TPCTC 2022

SIGMOD 2022-b

VLDB 2021

ICDE 2023-b

SIGMOD 2022-a

SIGMOD 2020

ACM TODS (under review)

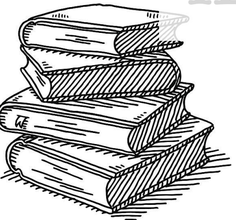

Log-**S**tructured **M**erge-tree

LSM-tree

LSM-tree

The Log-Structured Merge-Tree (LSM-Tree)

1996

Patrick O'Neil[1], Edward Cheng[2]
Dieter Gawlick[3], Elizabeth O'Neil[1]
To be published: Acta Informatica

LSM-tree

NoSQL

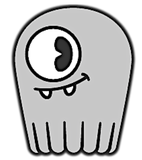

relational

time-series

2023

LSM-tree

NoSQL

relational

time-series

2023

LSM **Basics**

key-value pairs

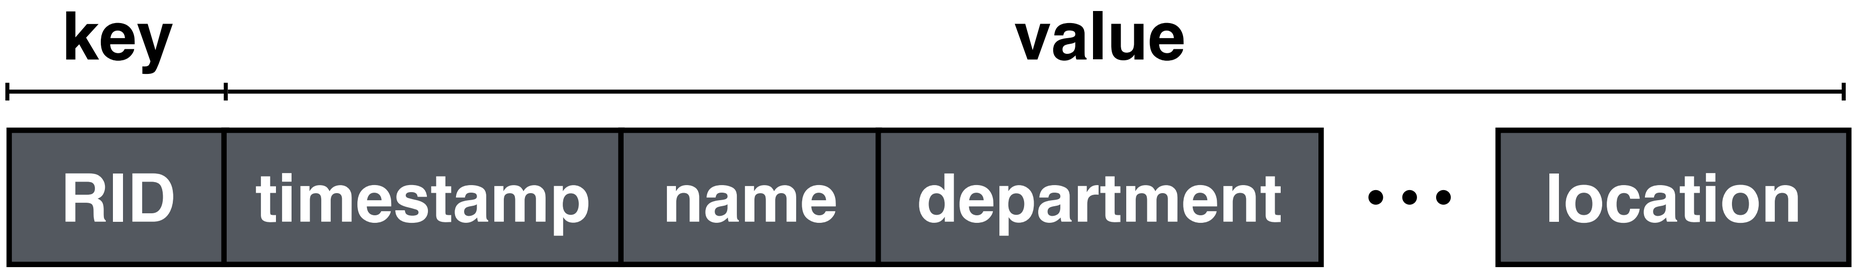

LSM **Basics**

key-value pairs

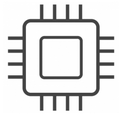

buffer

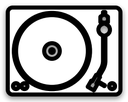

put(6)

put(2)

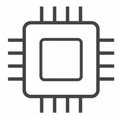

buffer

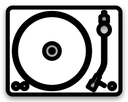

put(1)

put(6)

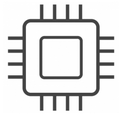

buffer 2

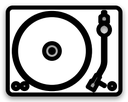

put(4)

put(1)

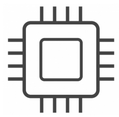

buffer

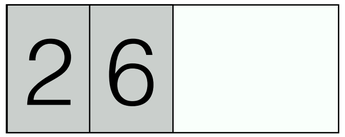

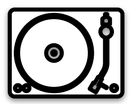

put(4)

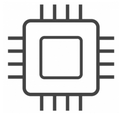

buffer

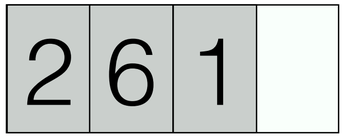

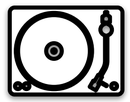

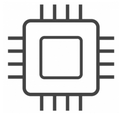

buffer

2	6	1	4

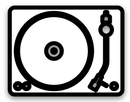

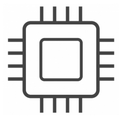

buffer

1	2	4	6

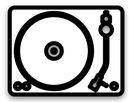

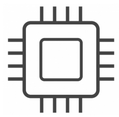

buffer

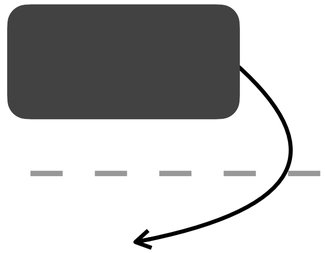

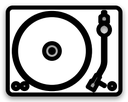

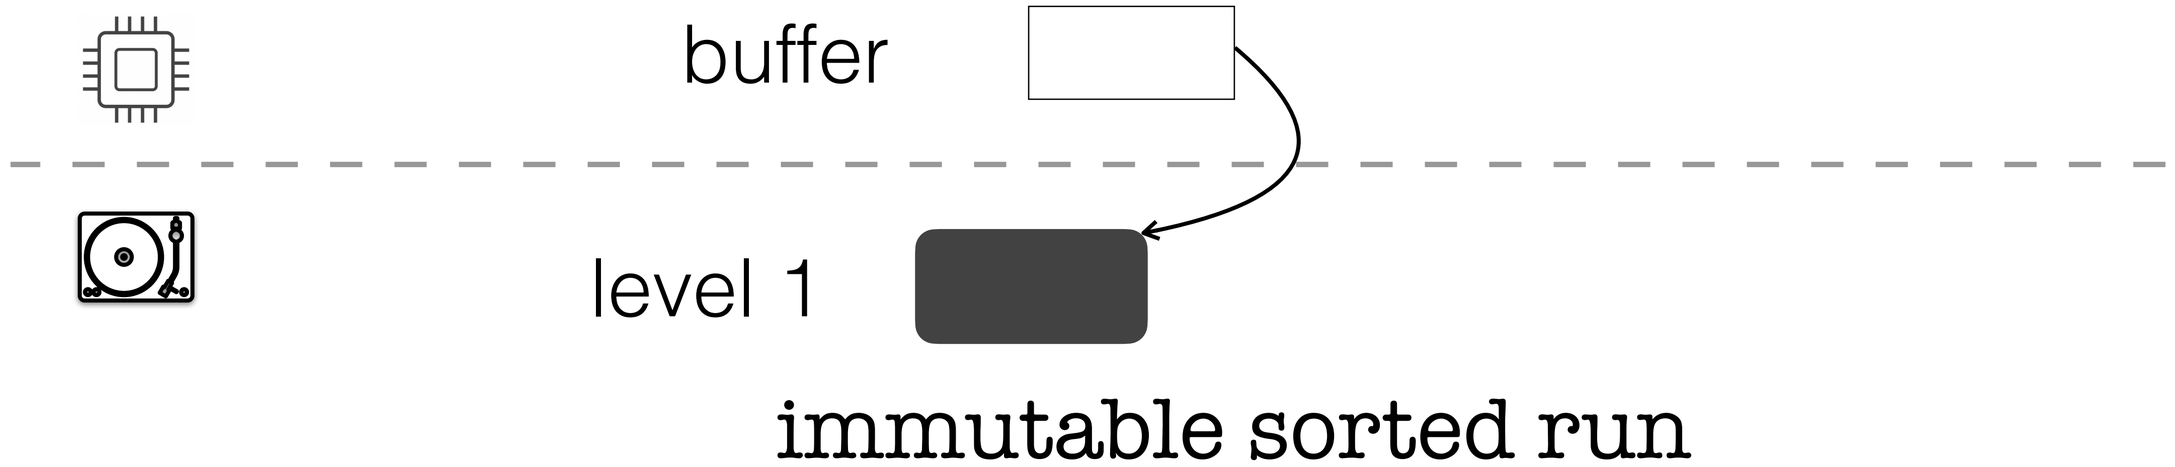
buffer
level 1
immutable sorted run

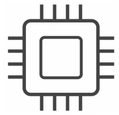

buffer

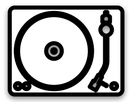

level 1

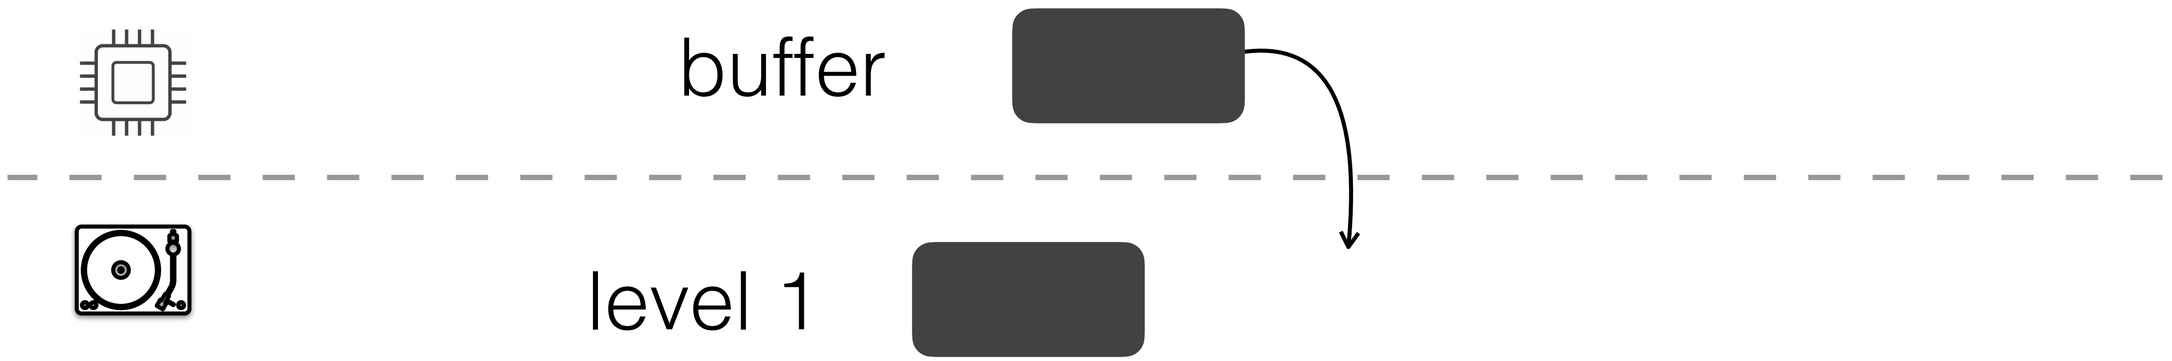
buffer
level 1

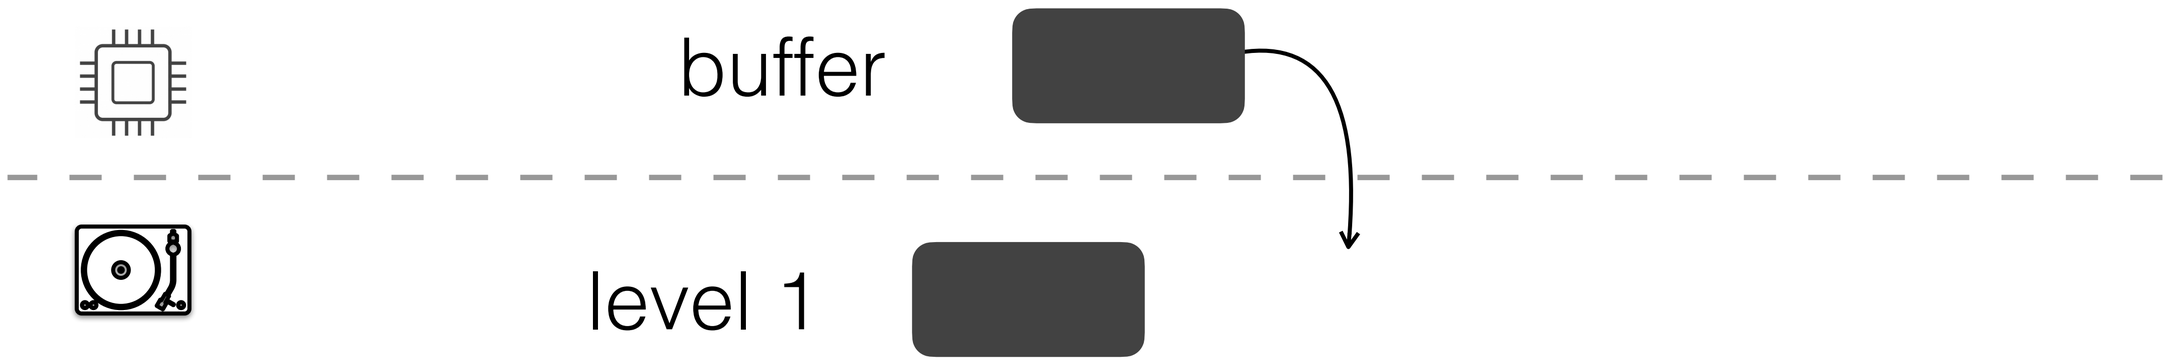
buffer
level 1

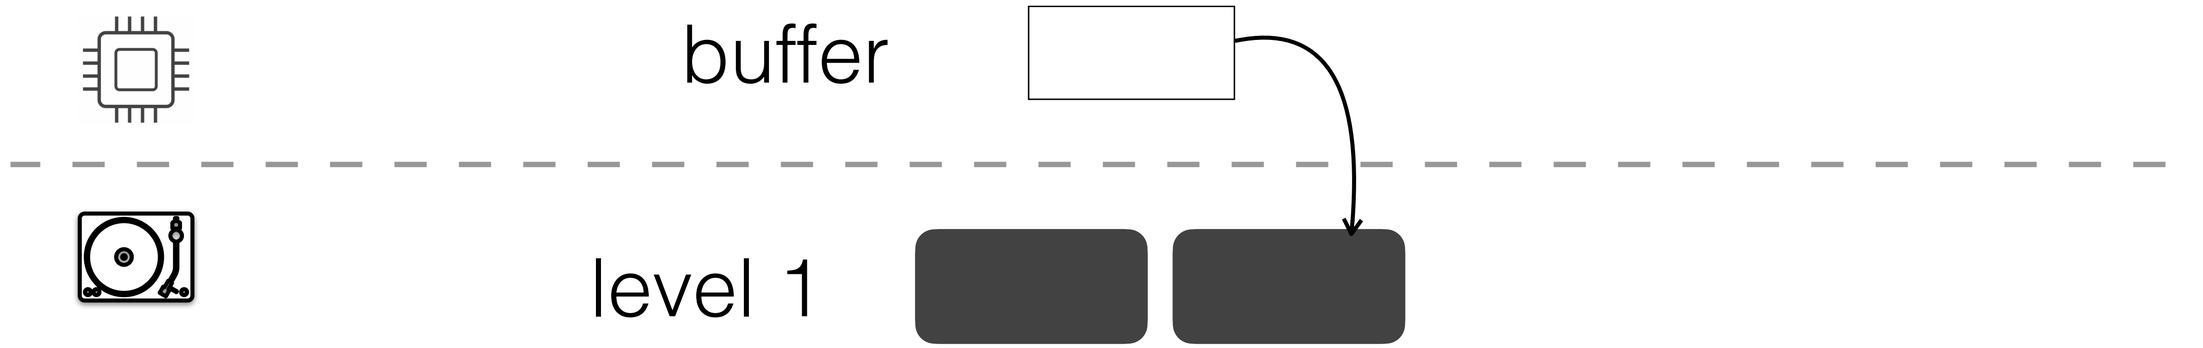
buffer
level 1

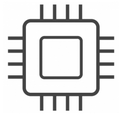

buffer

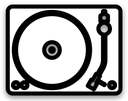

level 1

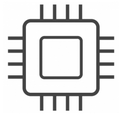

buffer

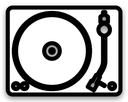

level 1

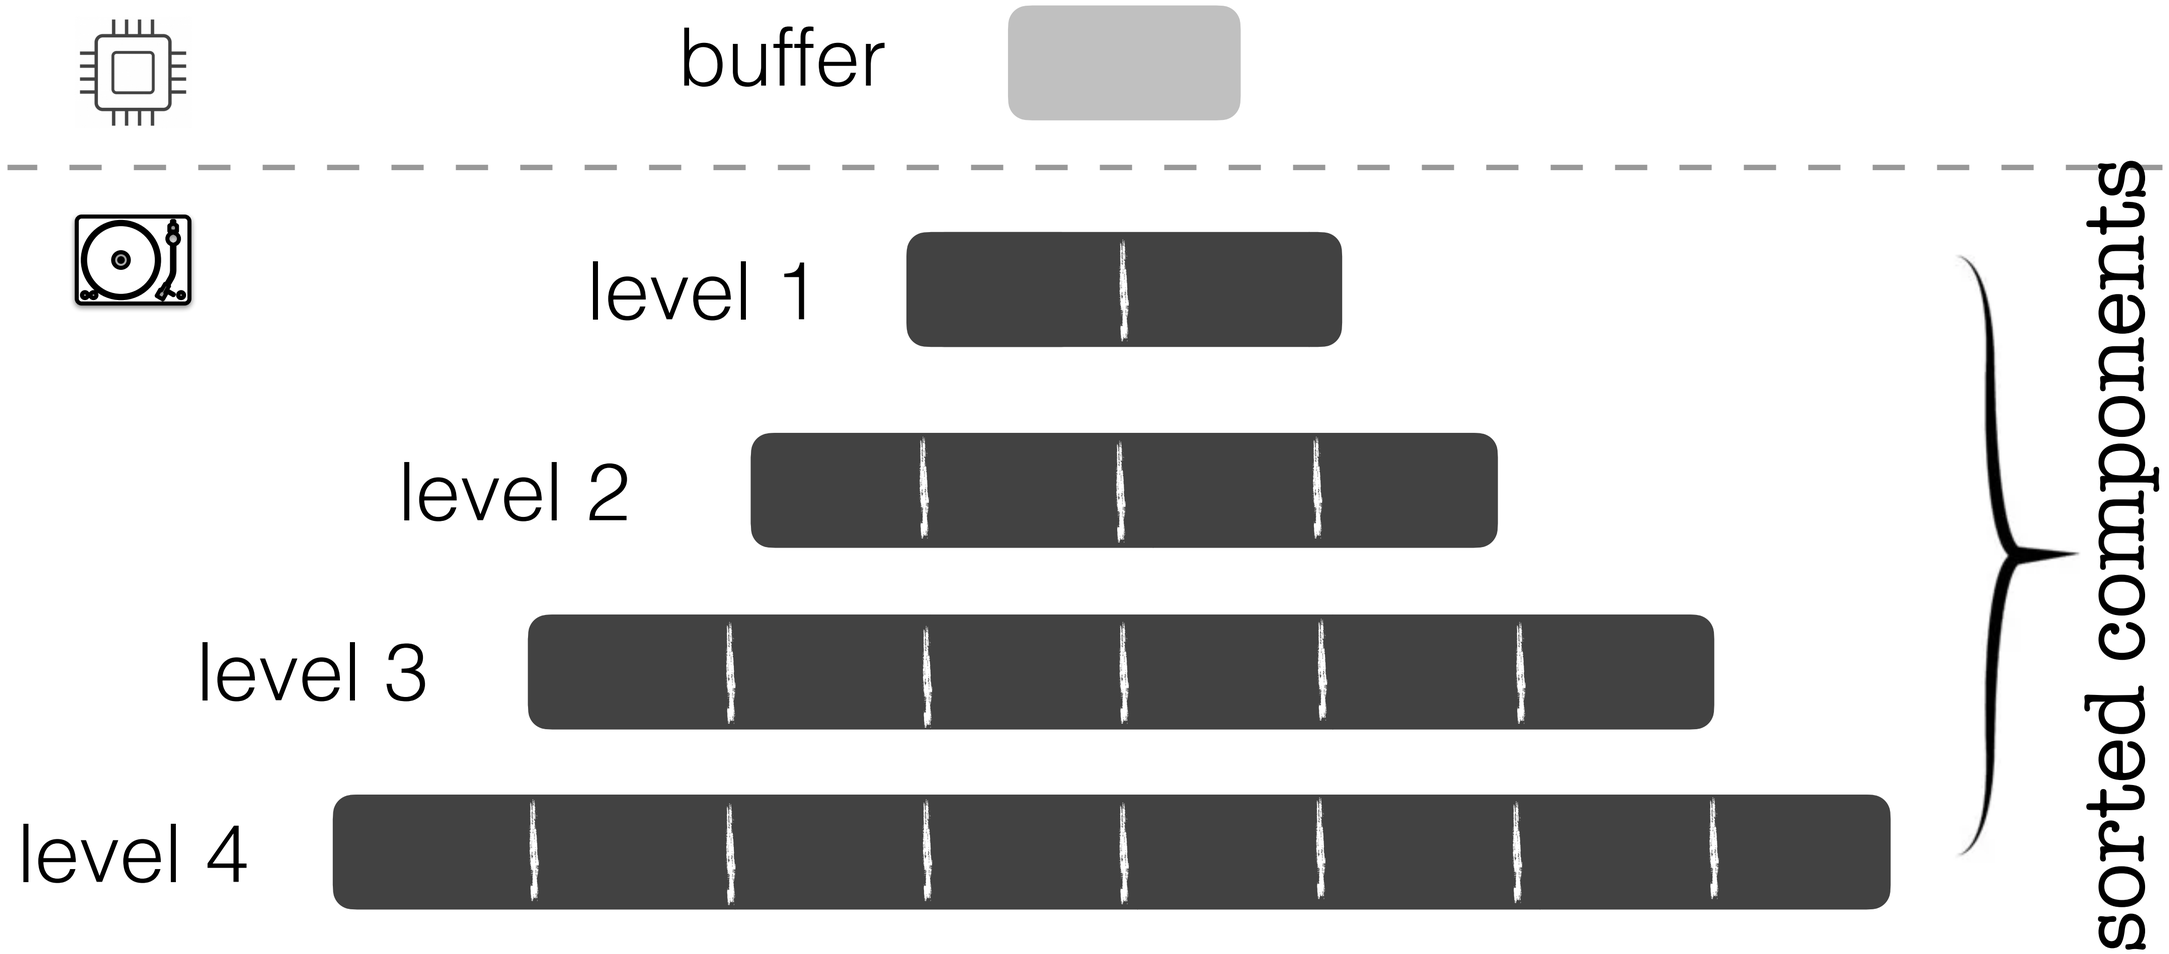
buffer
level 1
level 2
level 3
level 4
sorted components

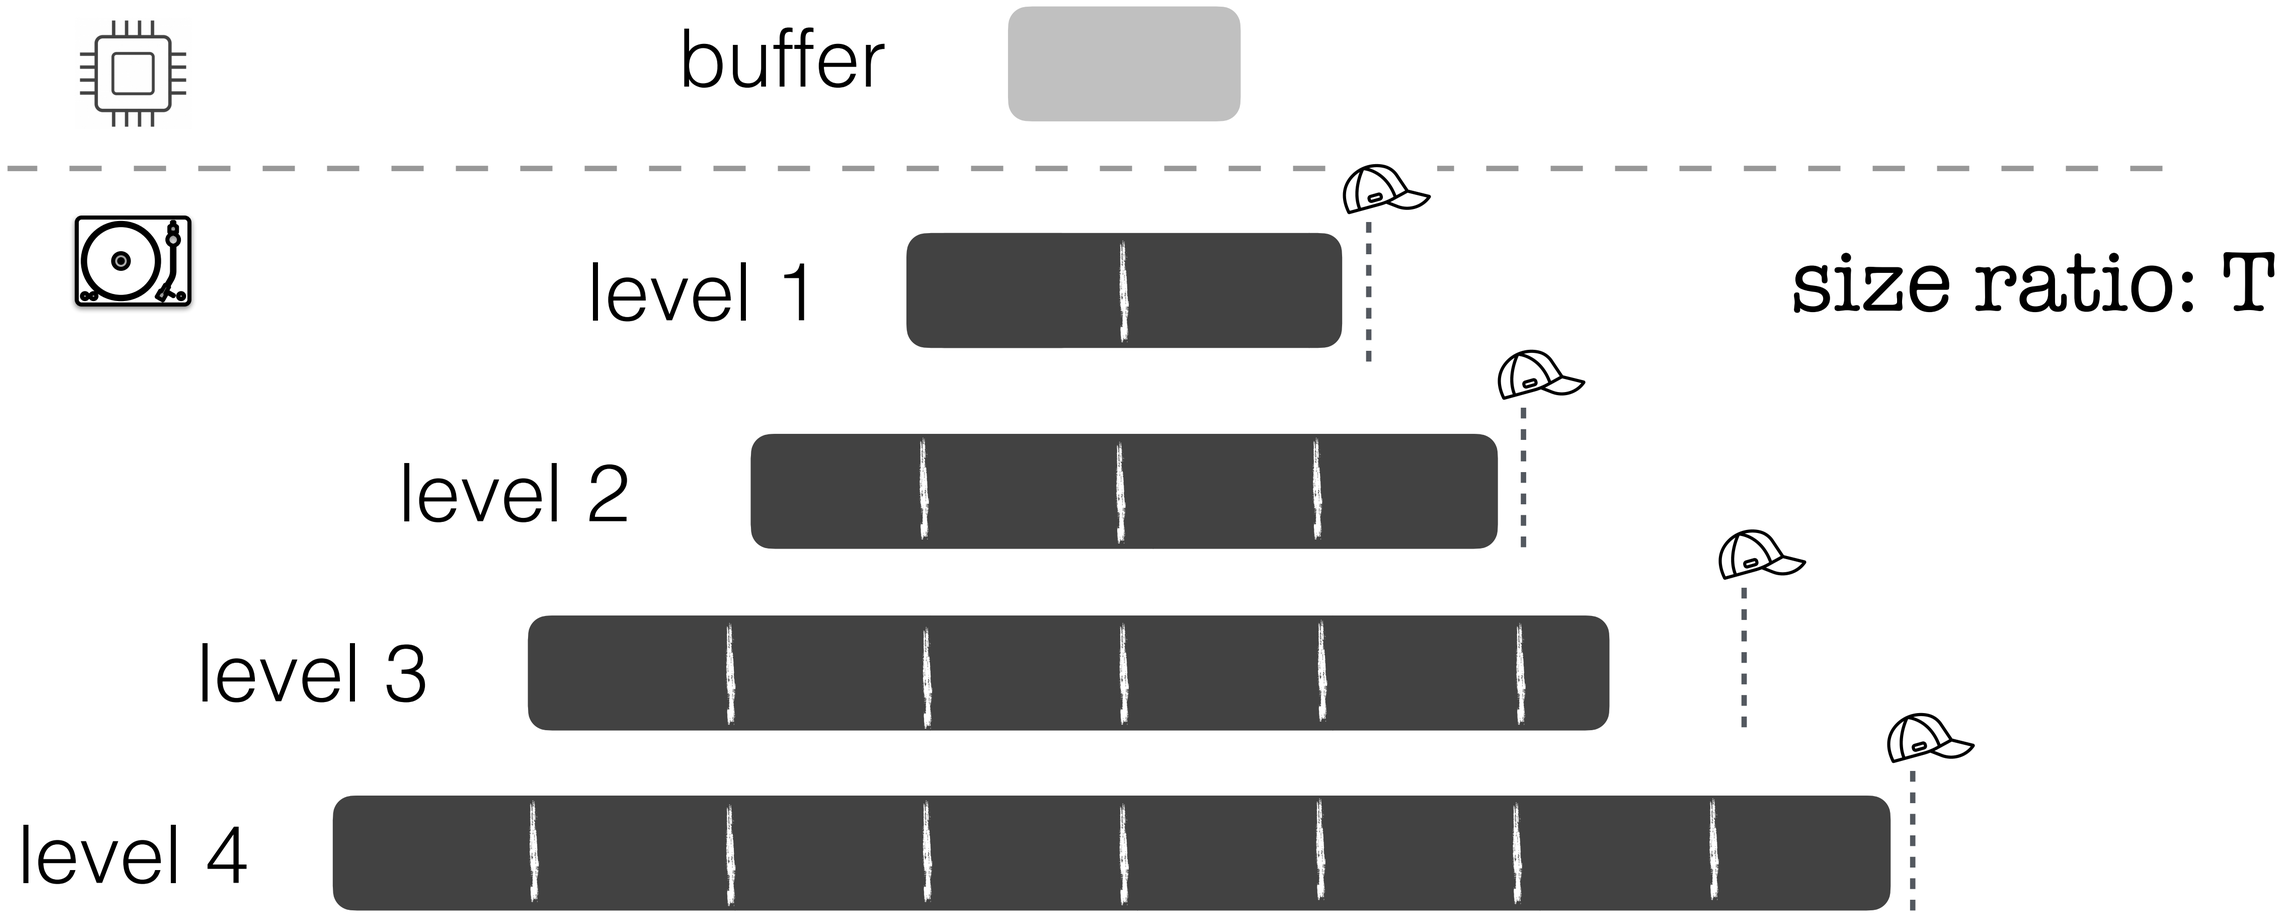
buffer
level 1
level 2
level 3
level 4
size ratio: T

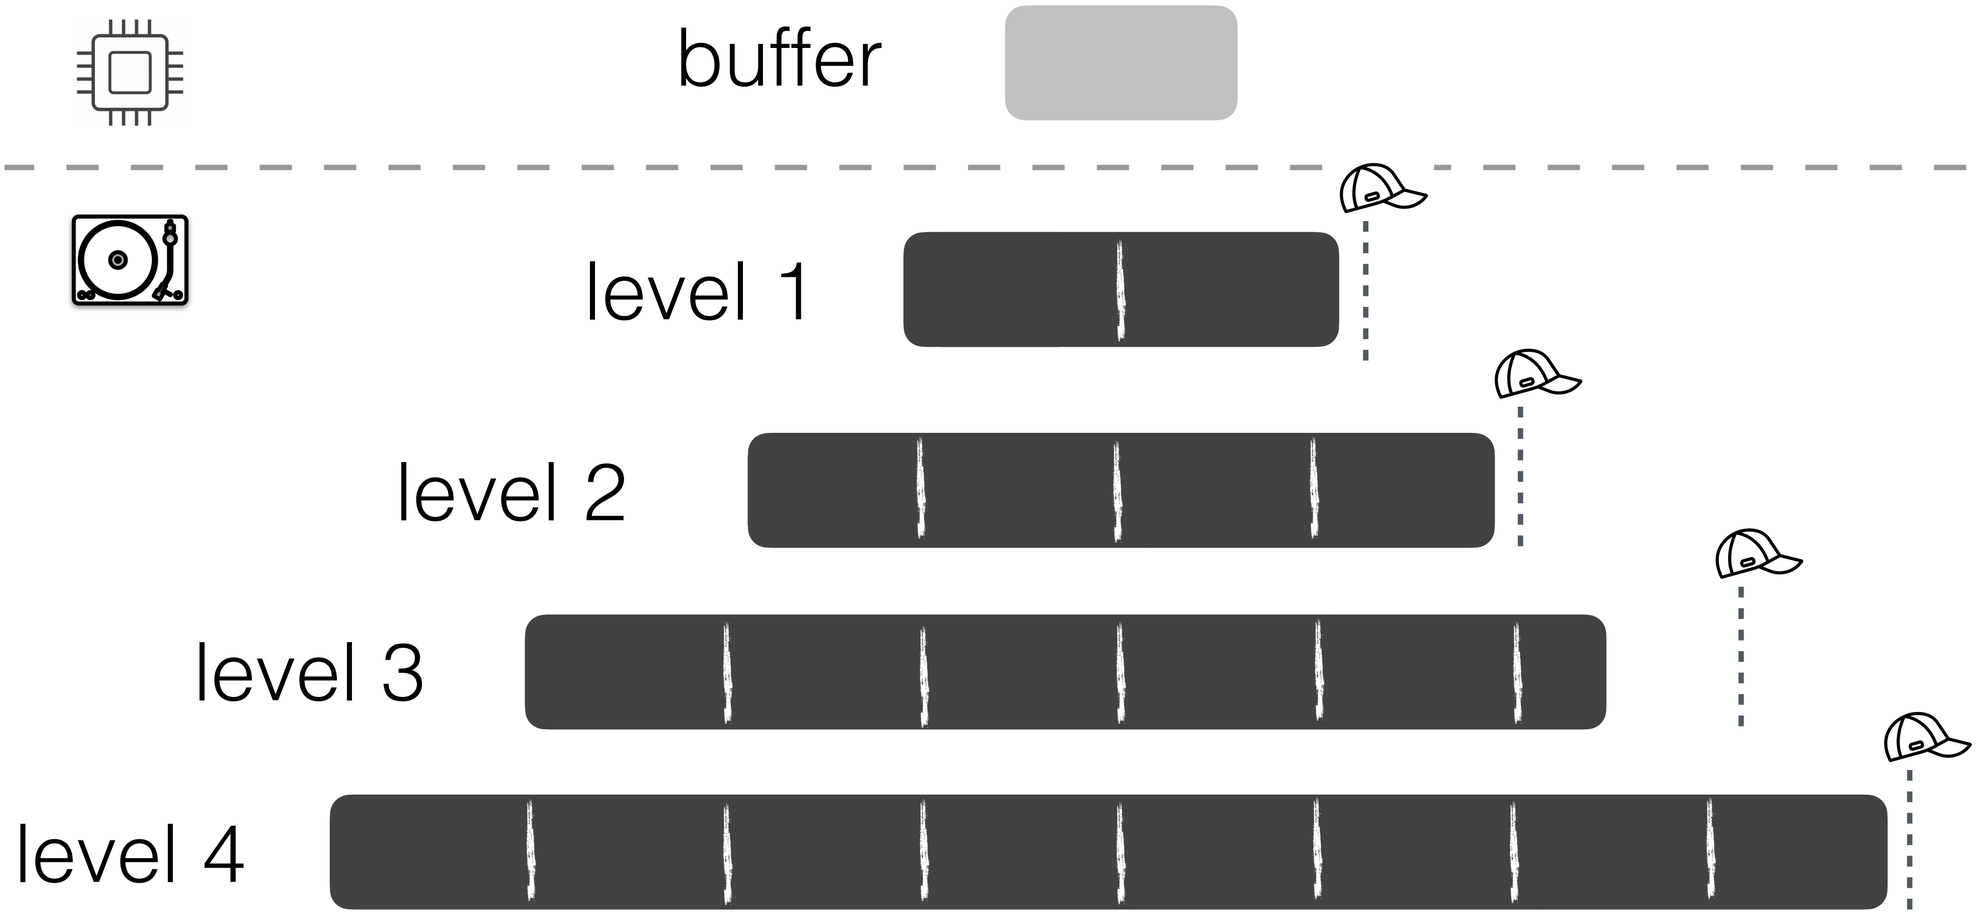
buffer
level 1
level 2
level 3
level 4

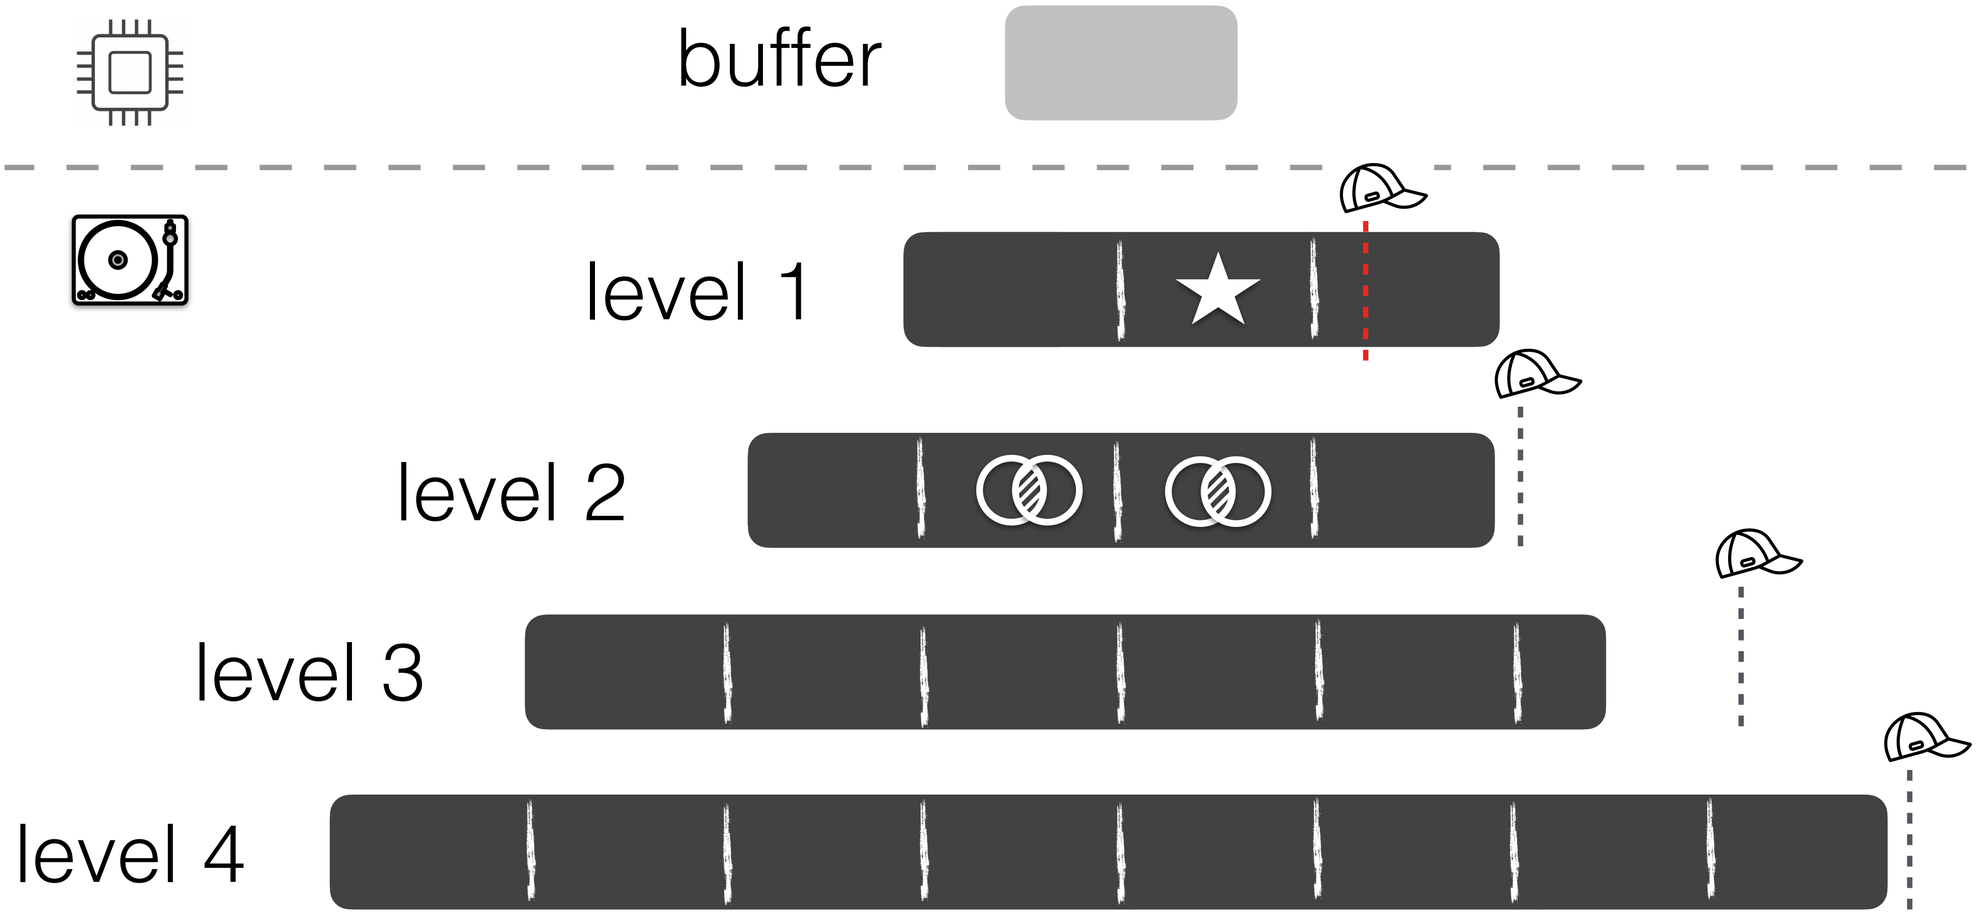
buffer
level 1
level 2
level 3
level 4

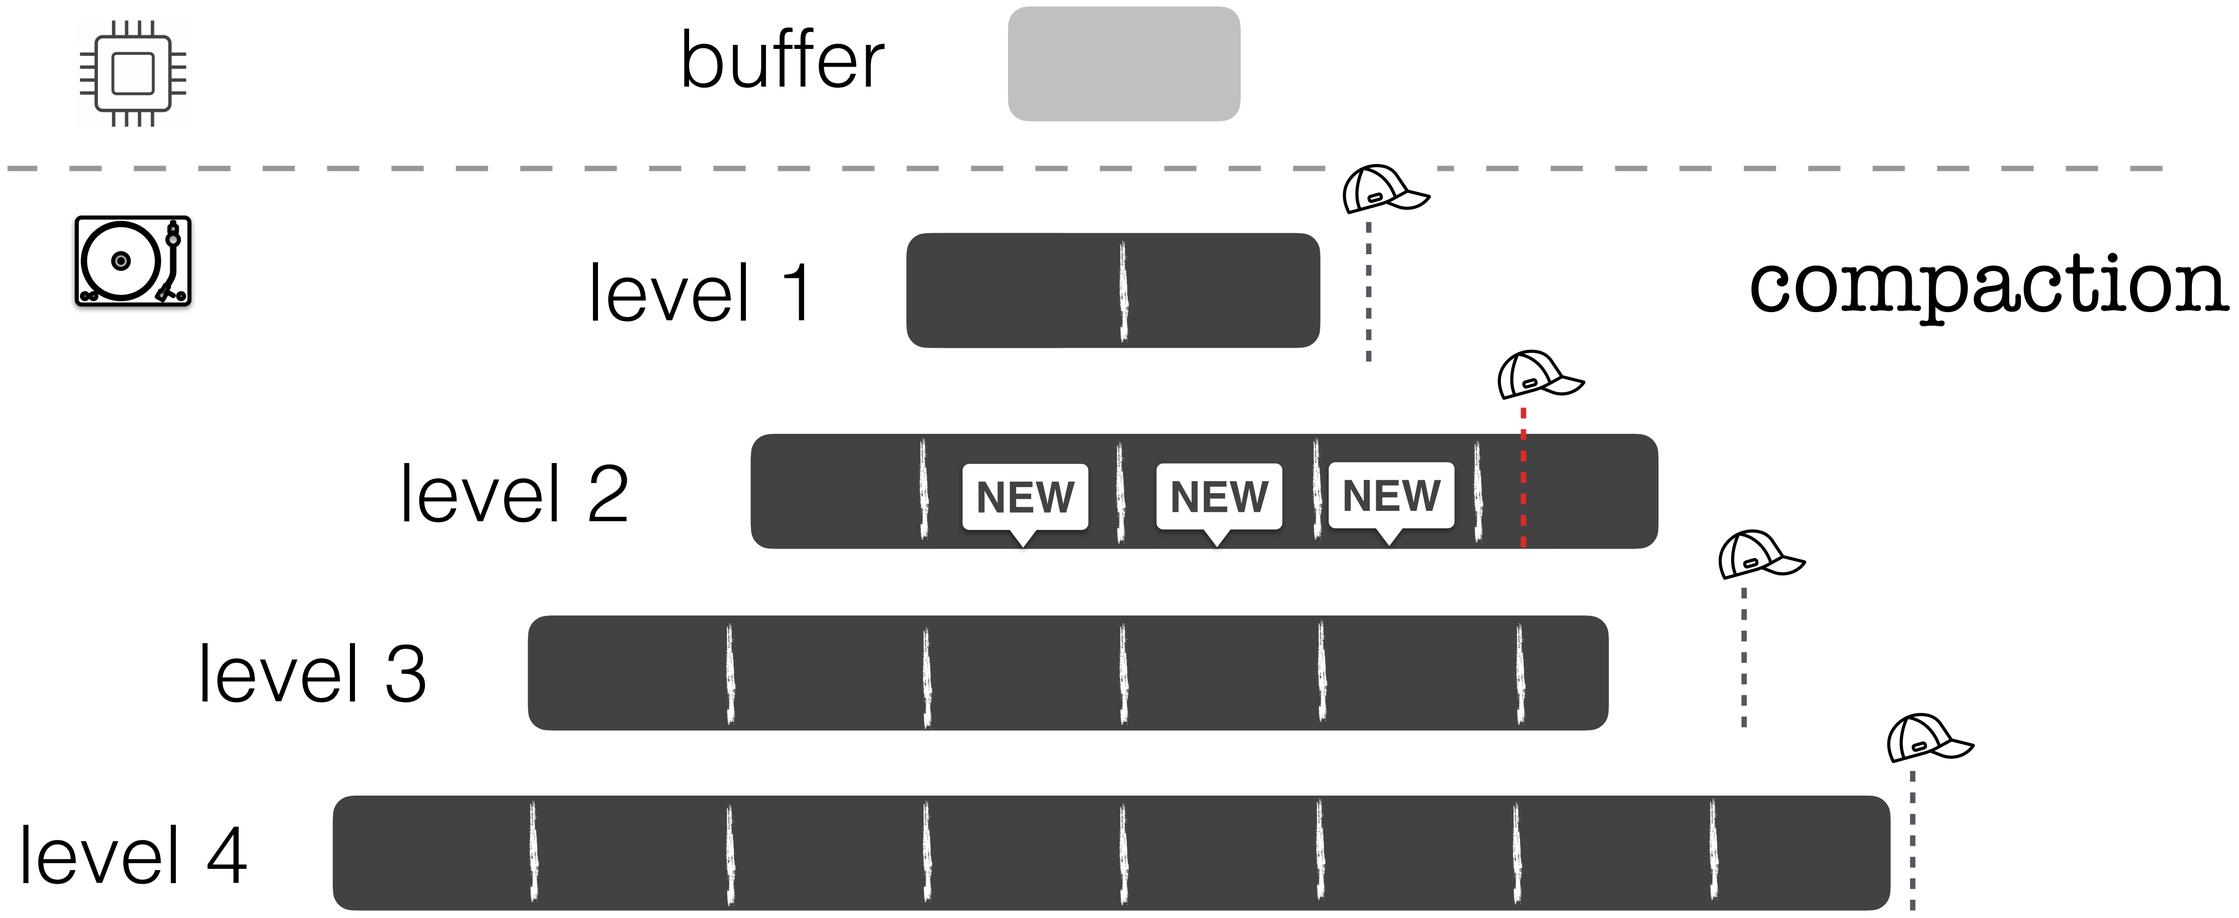
buffer
level 1
compaction
level 2
NEW
NEW
NEW
level 3
level 4

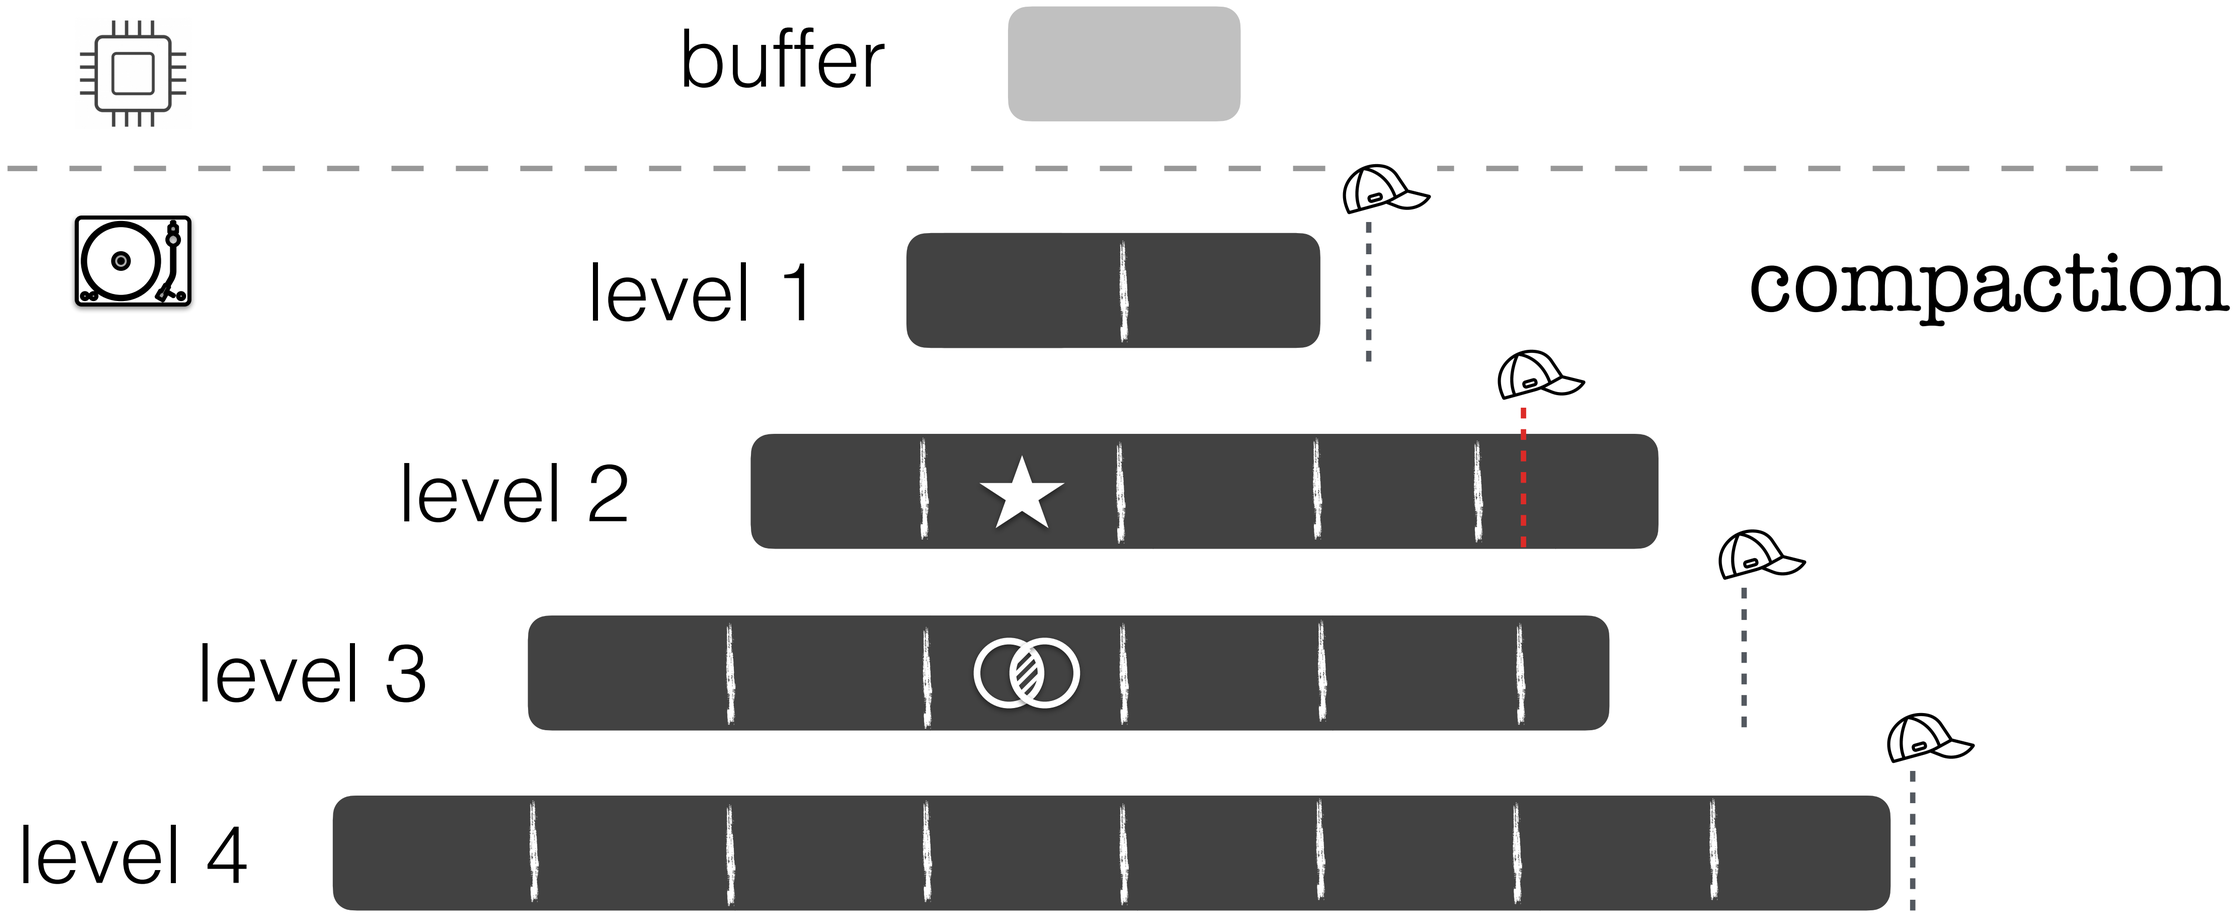
buffer
level 1
compaction
level 2
level 3
level 4

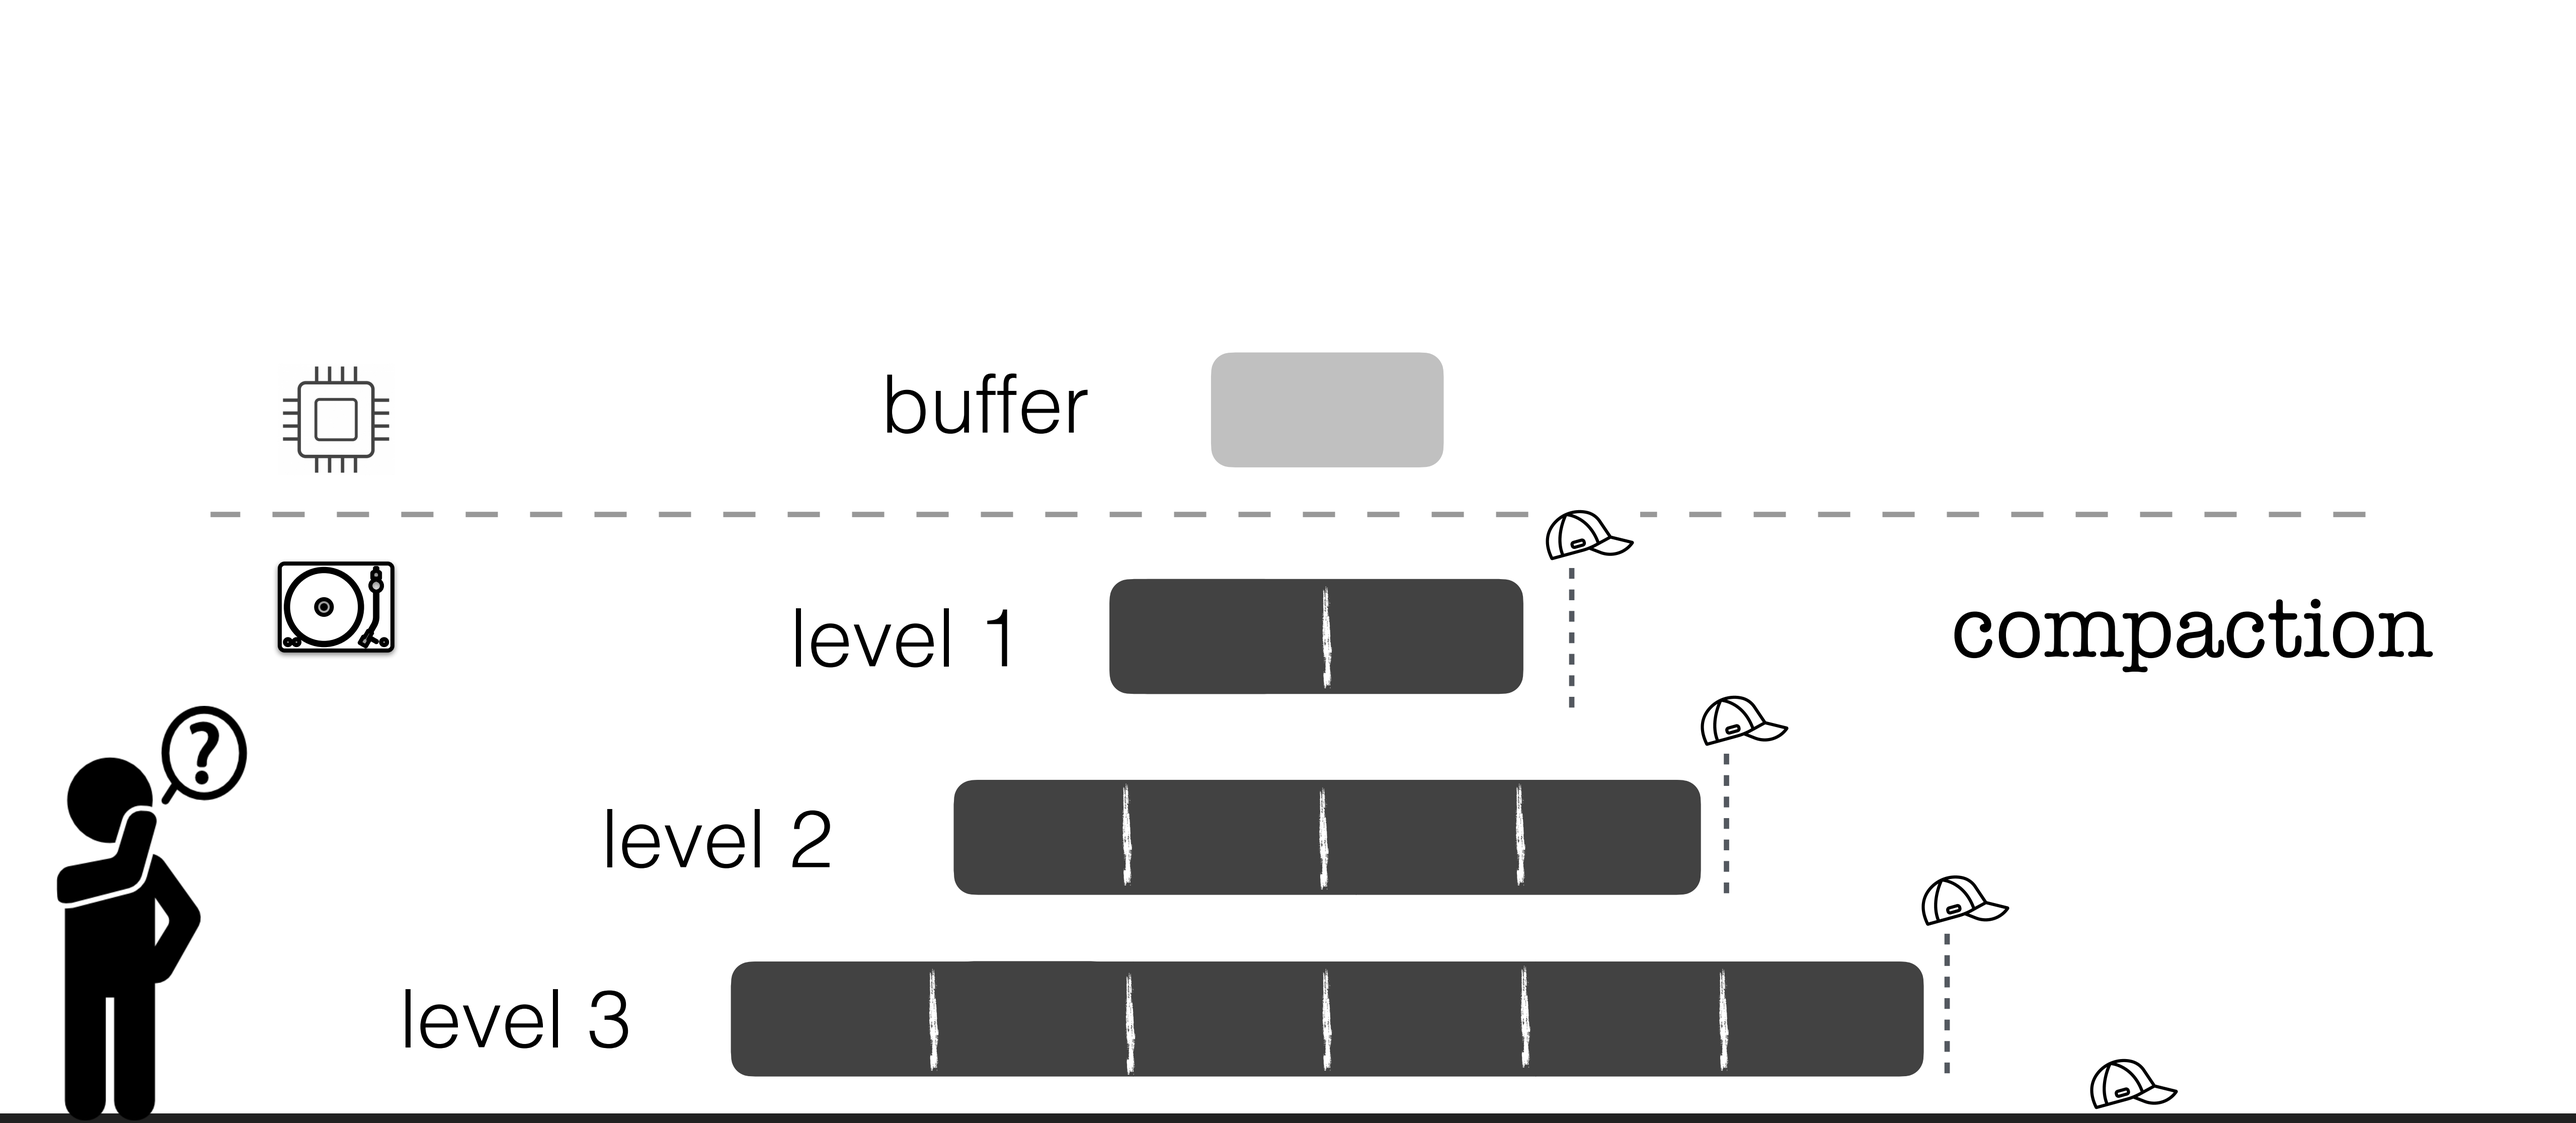

So, what about deletes?

Deletes in LSMs

delete := insert tombstone

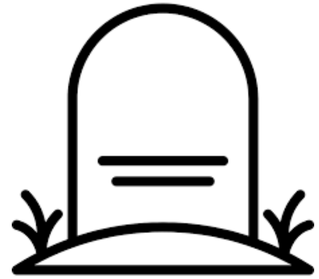

Deletes in LSMs

delete := insert tombstone

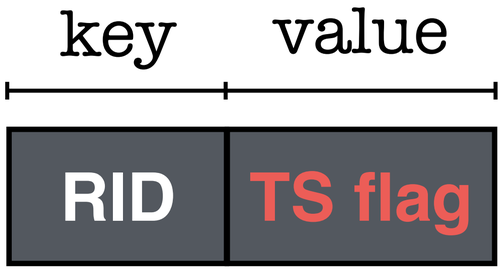

Deletes in LSMs

delete(5)

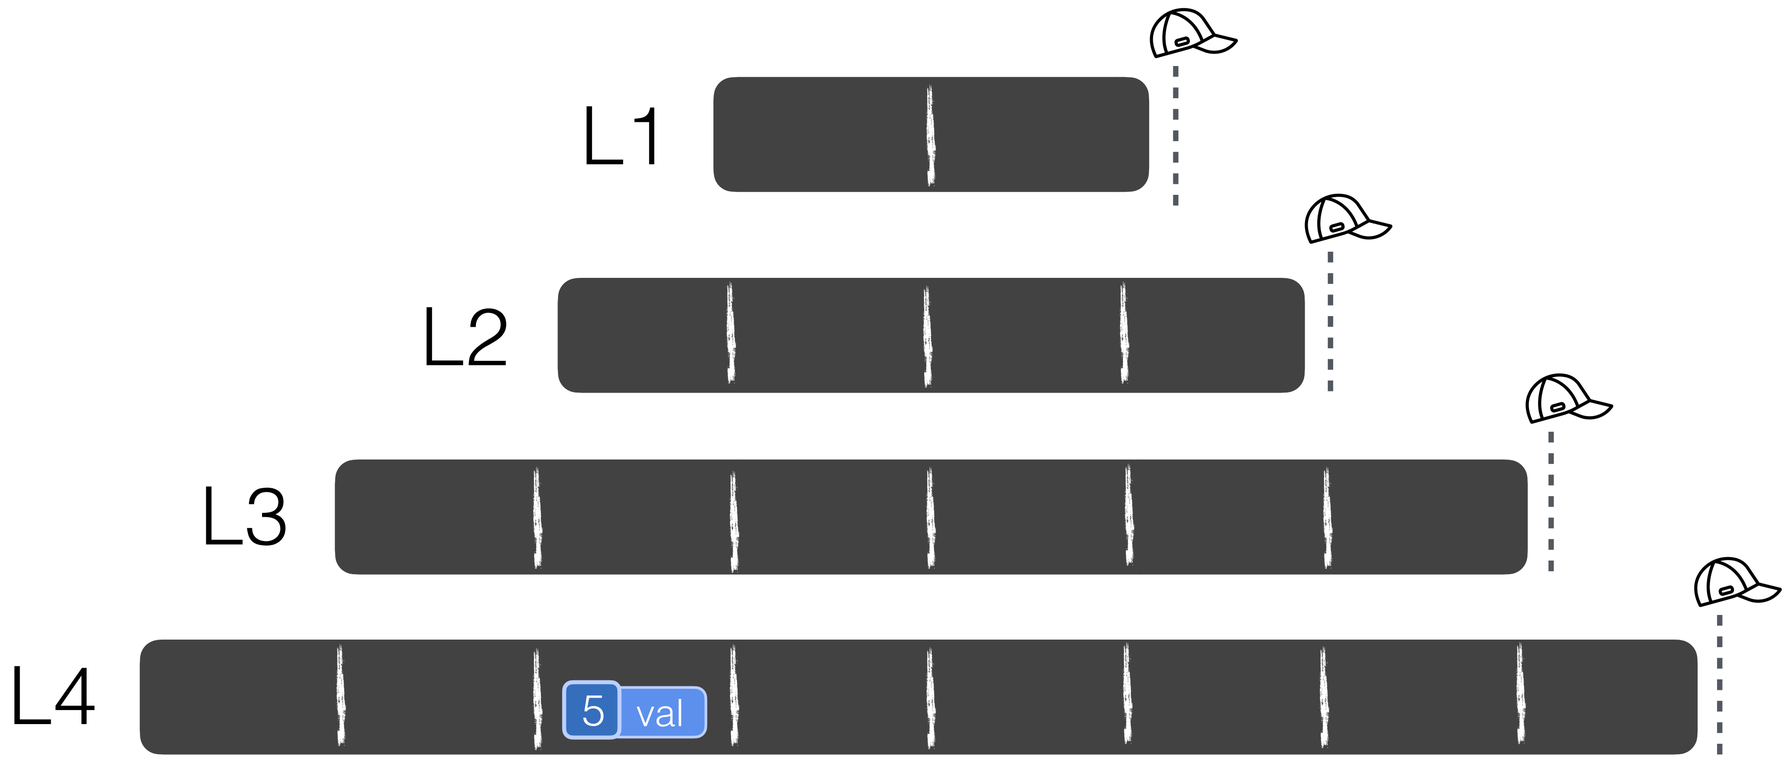

Deletes in LSMs

delete(5)

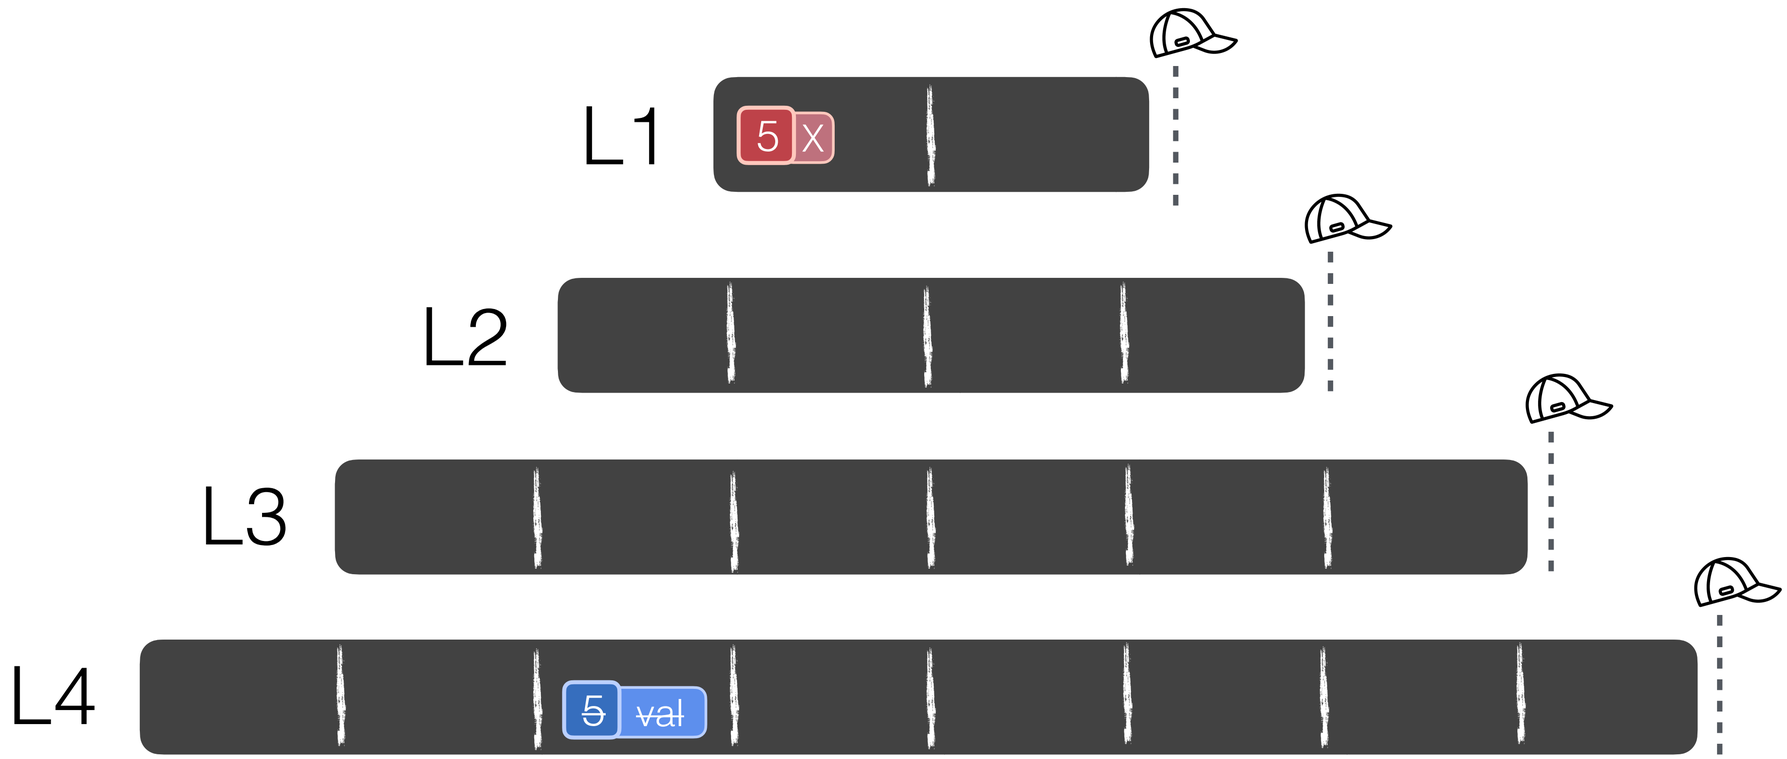

Deletes in LSMs

delete(5)

no support for timely delete persistence

great for inserts

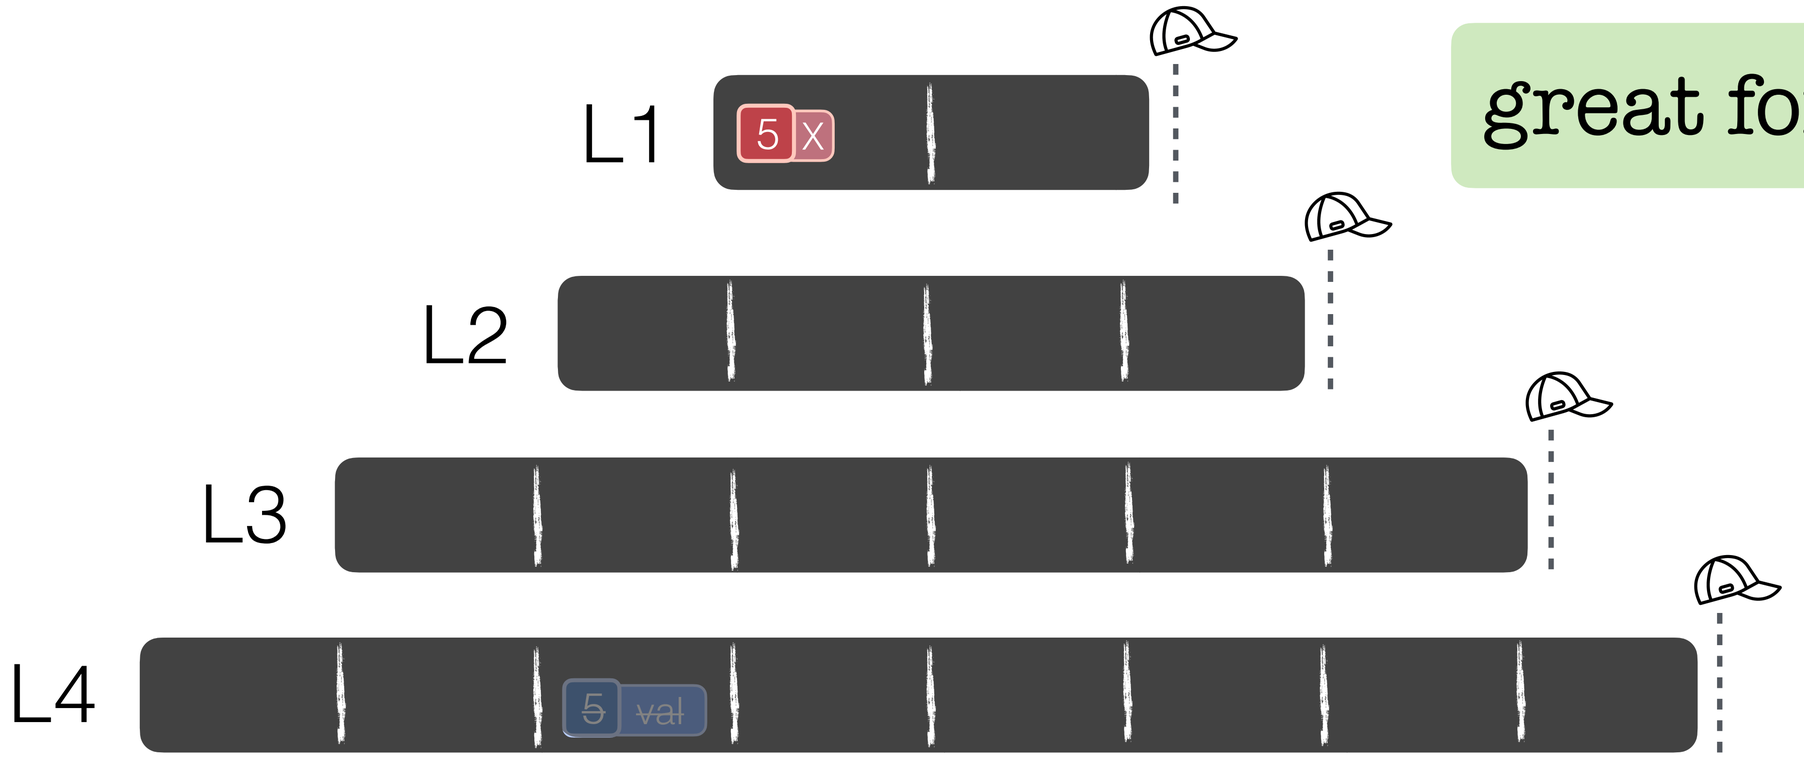

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

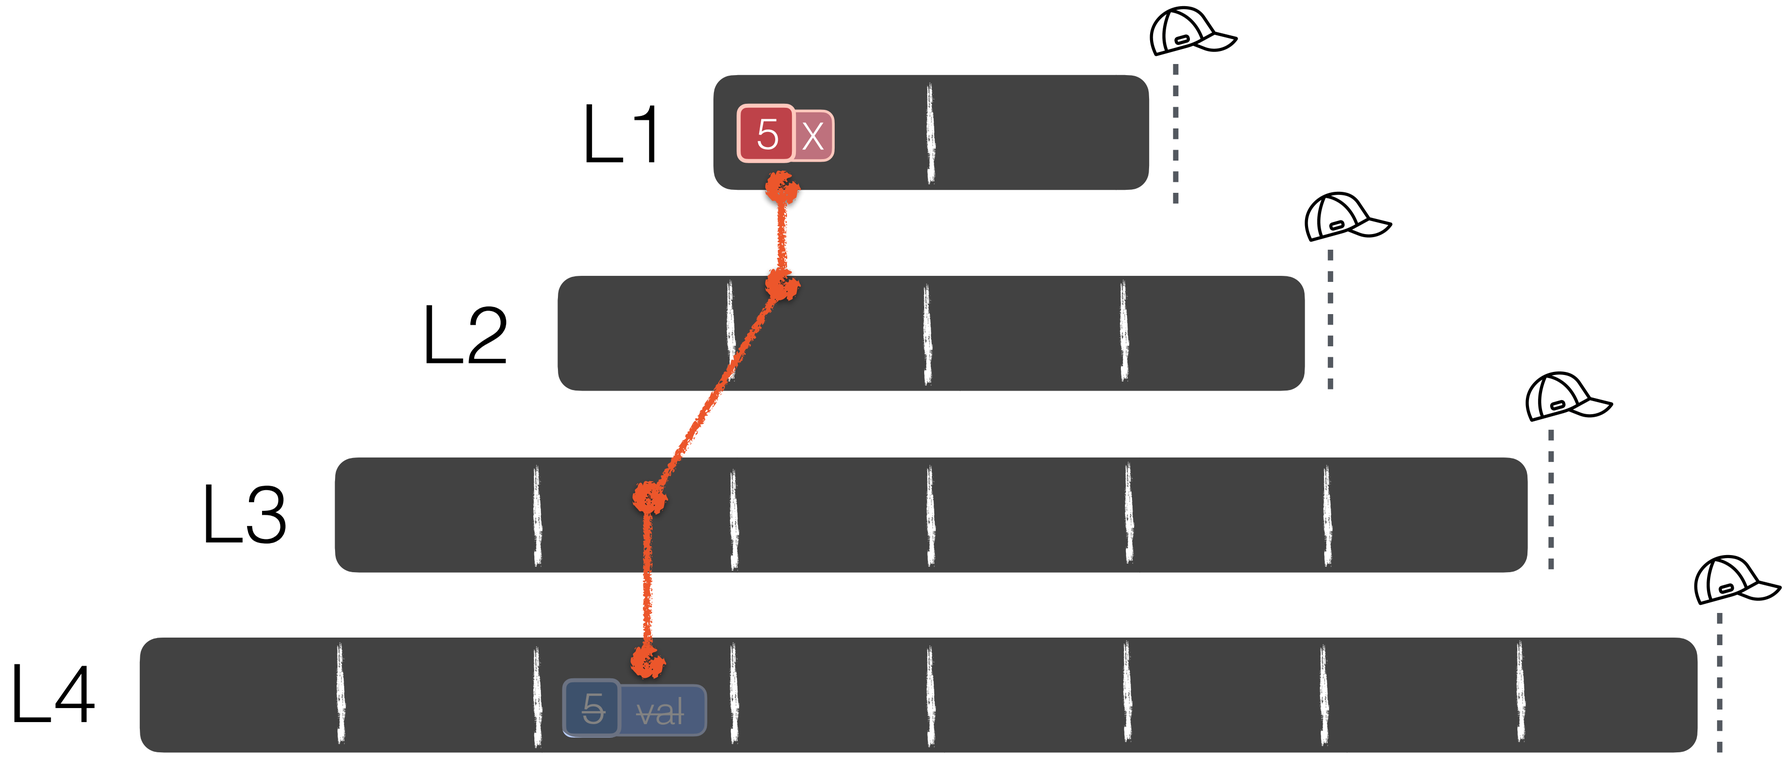

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

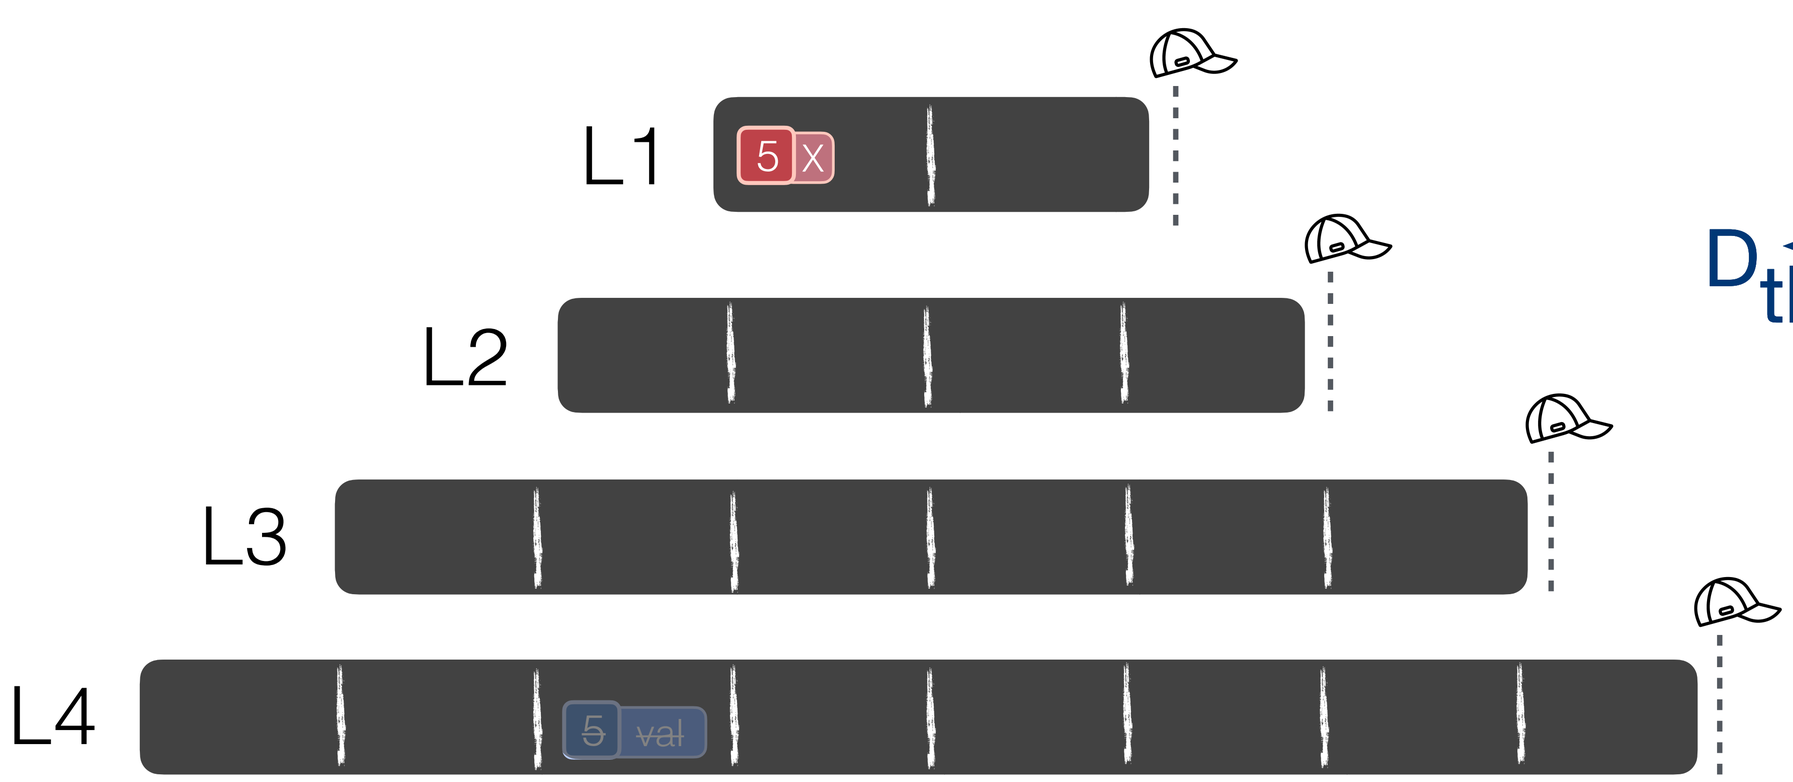

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

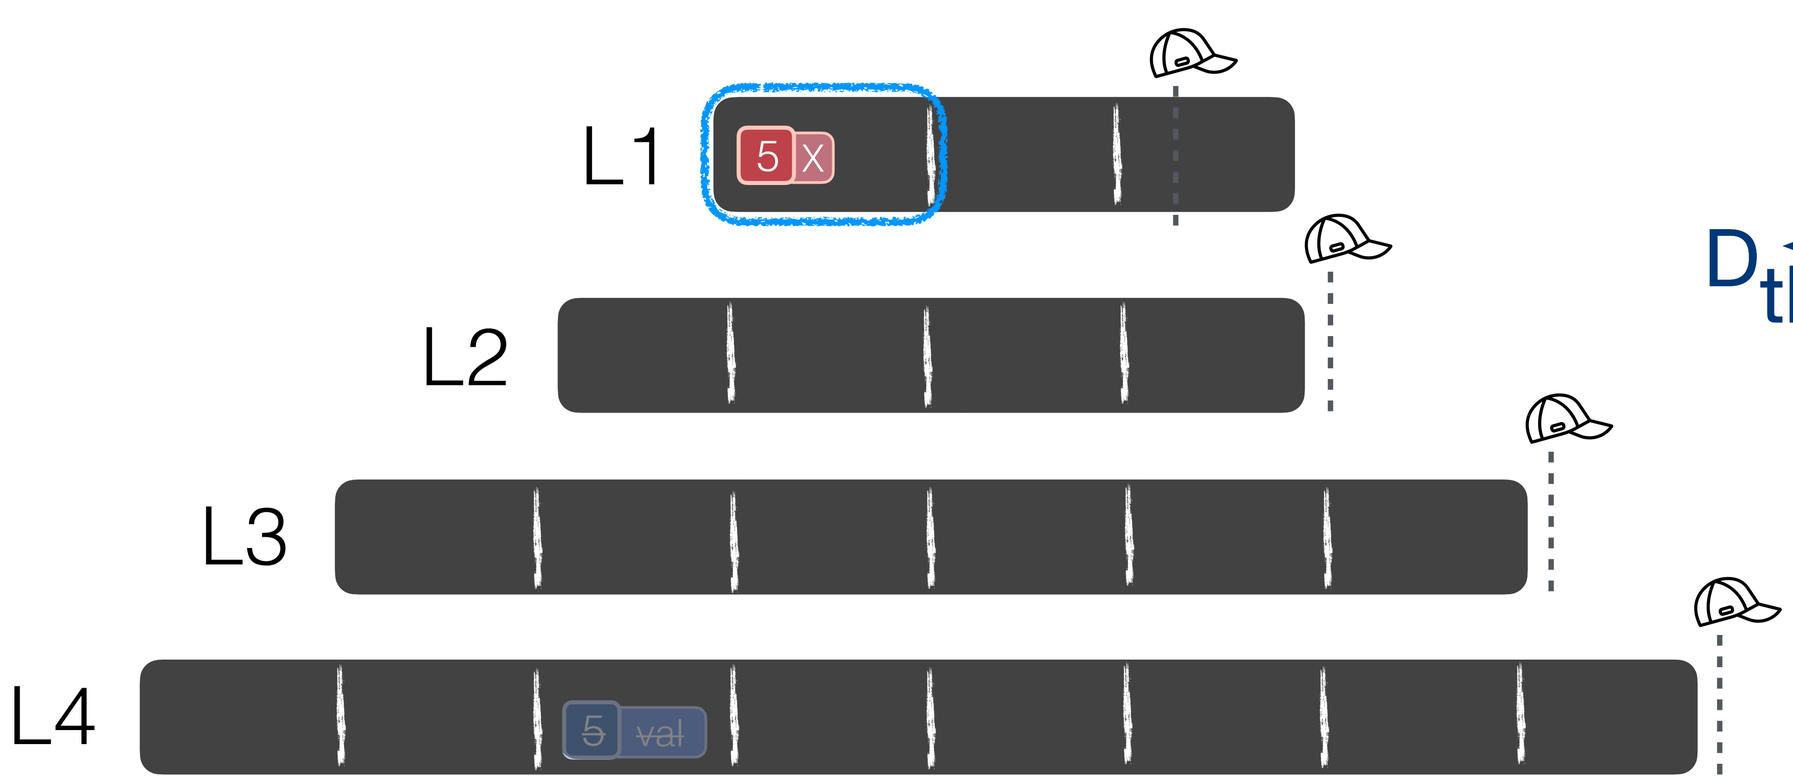

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

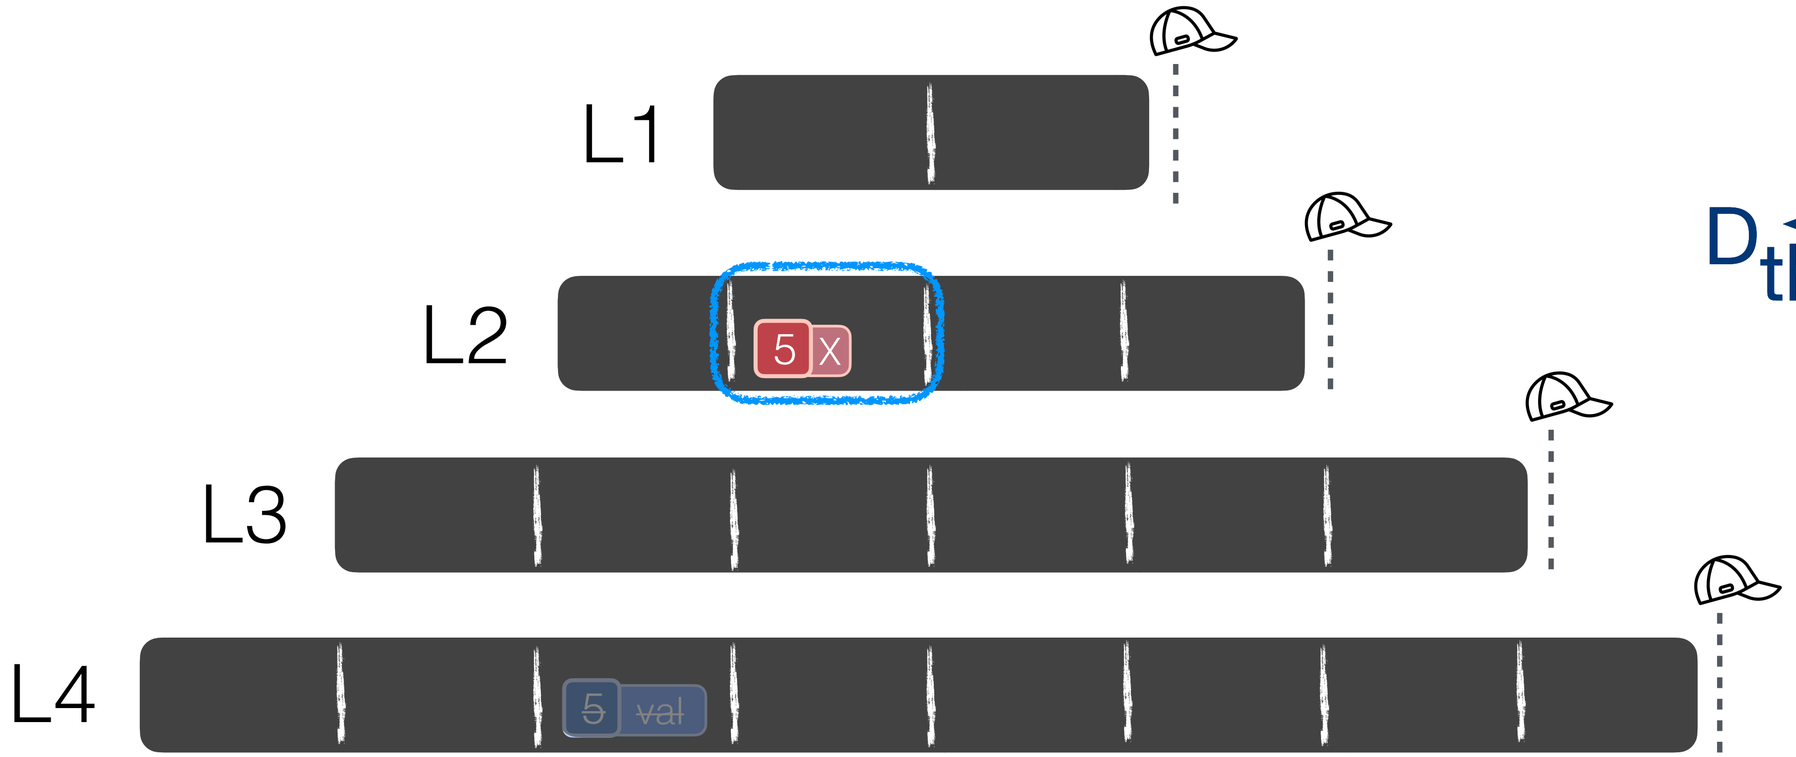

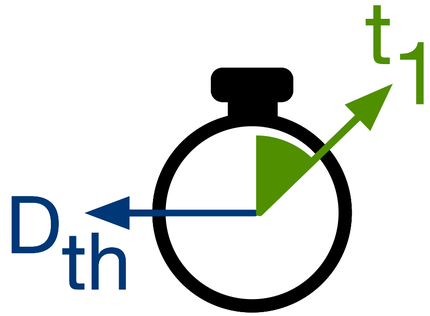

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

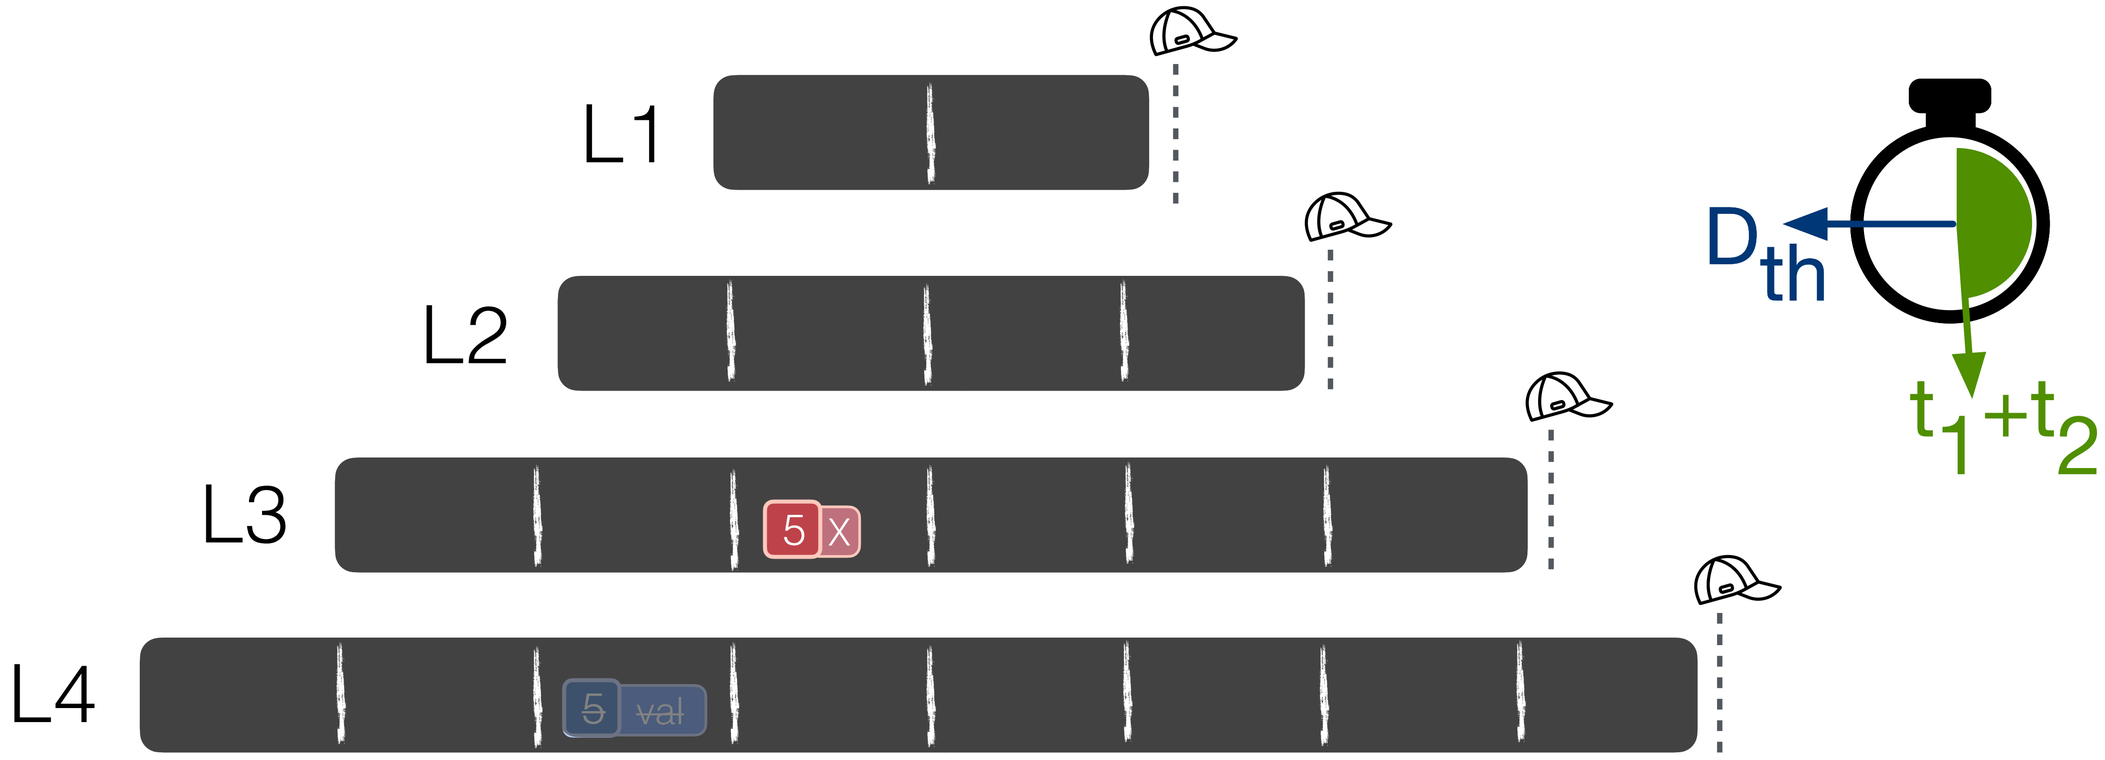

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

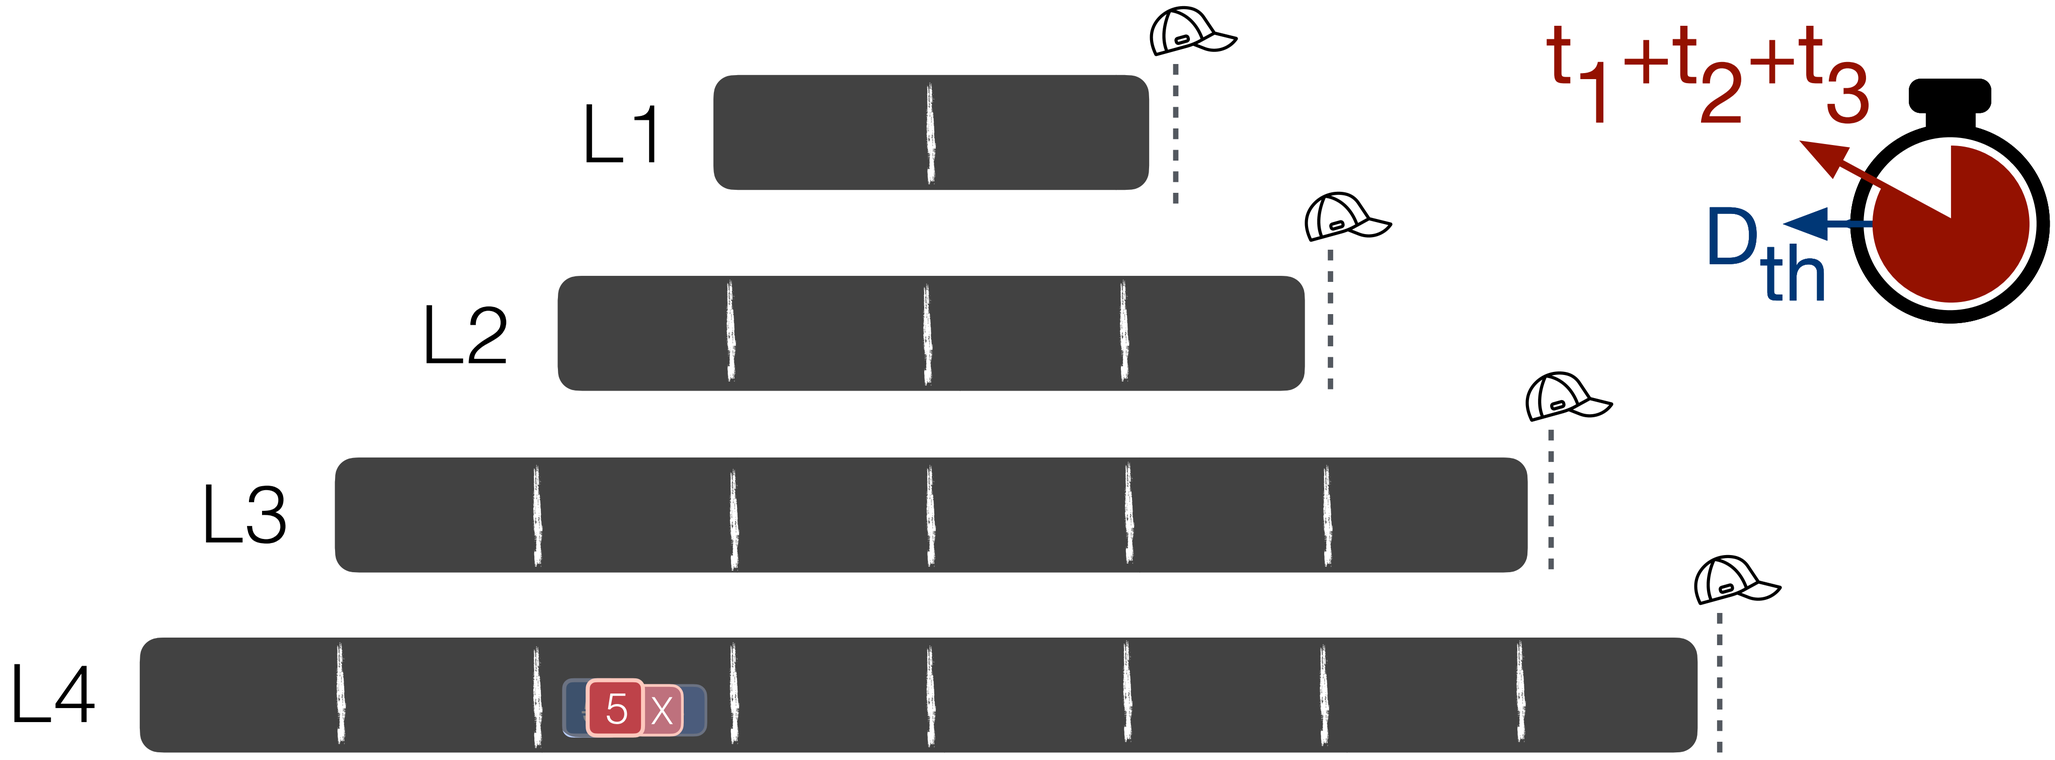

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

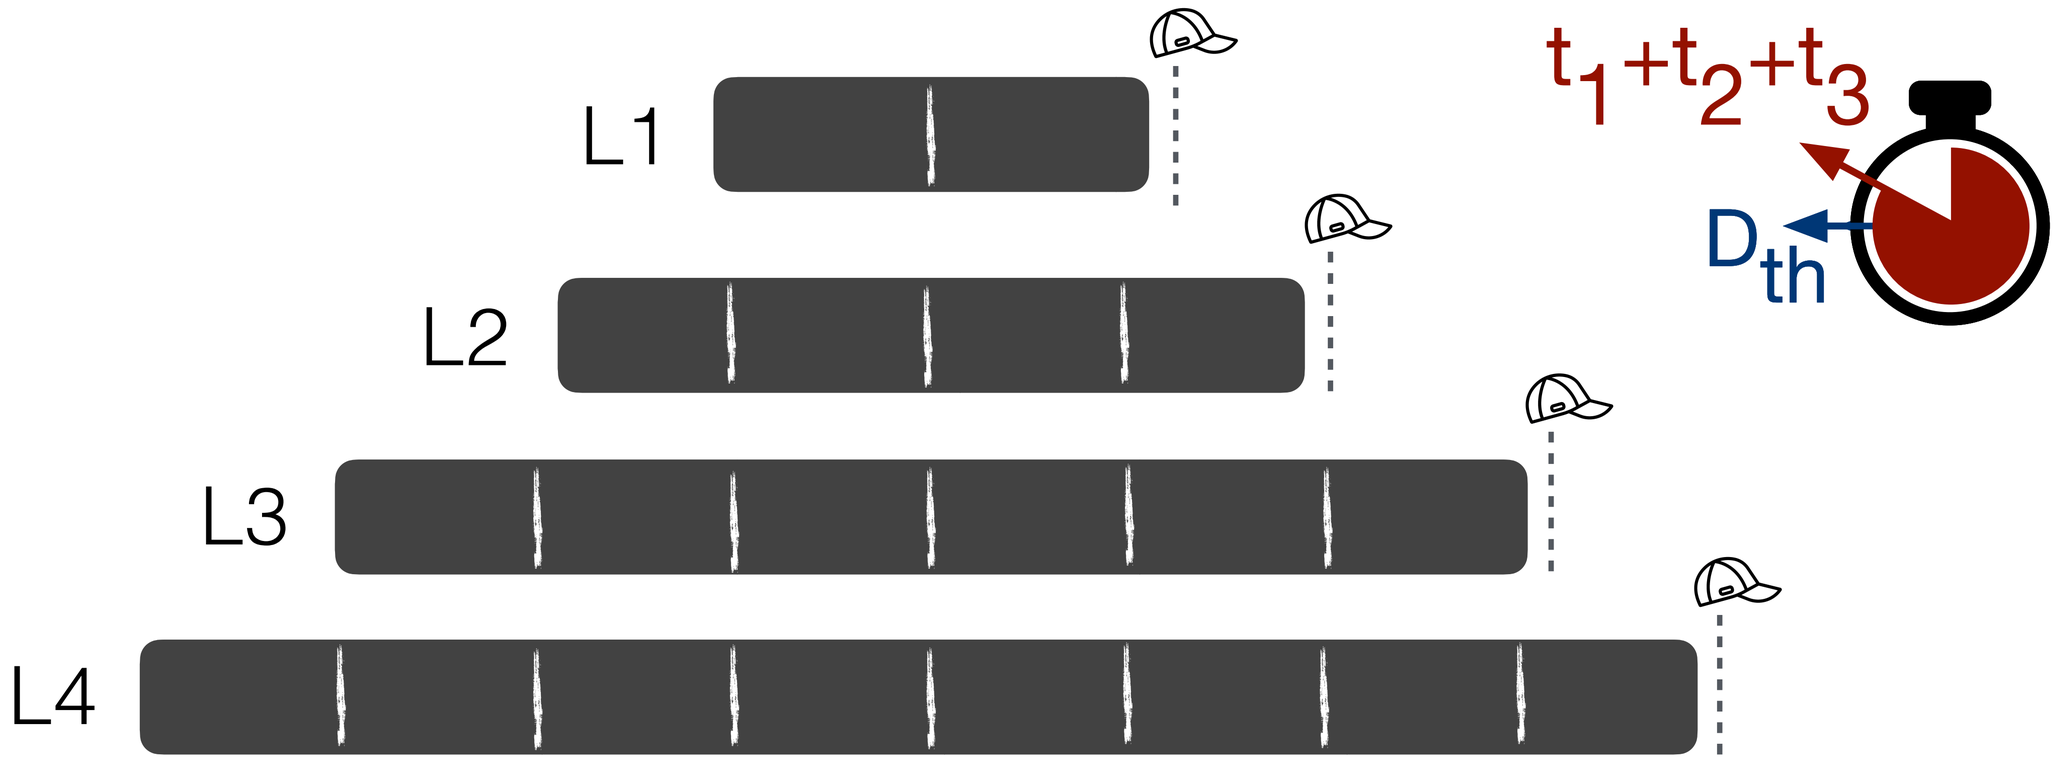

Problem: Persisting Deletes Timely

delete(5) within threshold: D_{th}

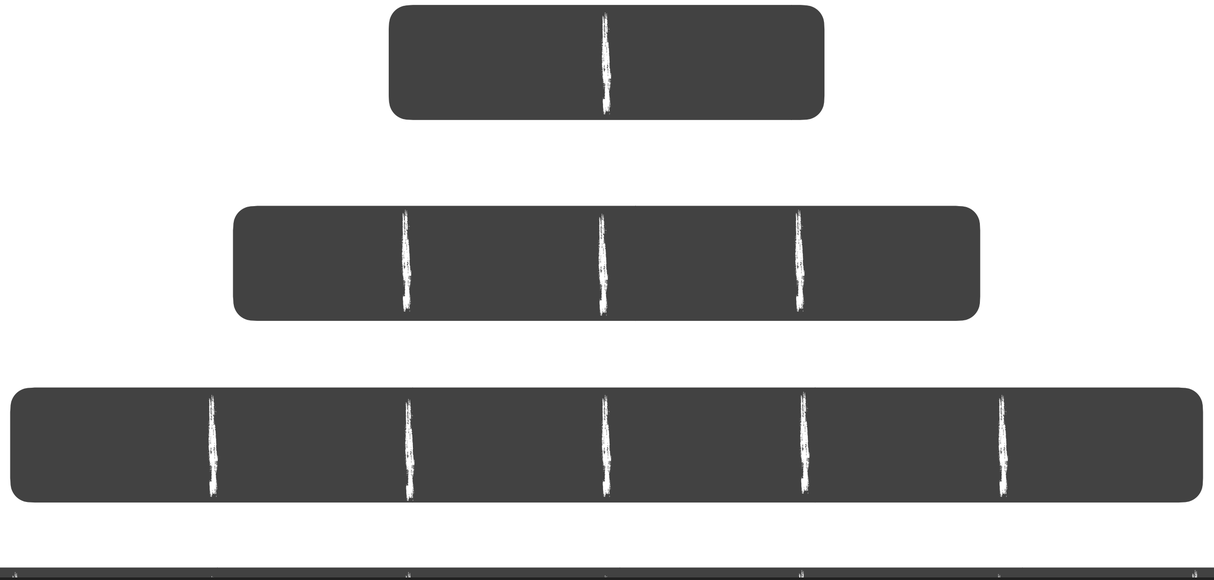

unbounded delete persistence latency

- File picking policy
- Tree shape
- Ingestion rate

Intuition: Compaction holds the key!

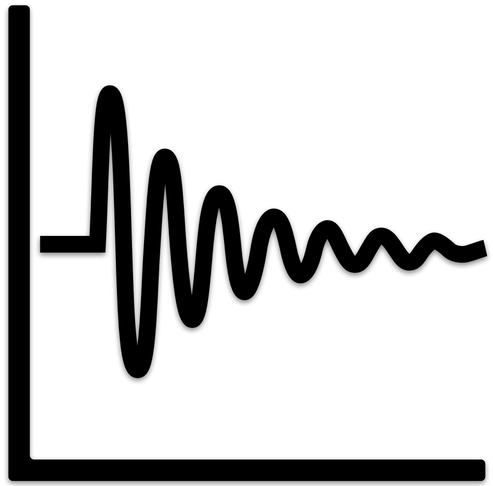

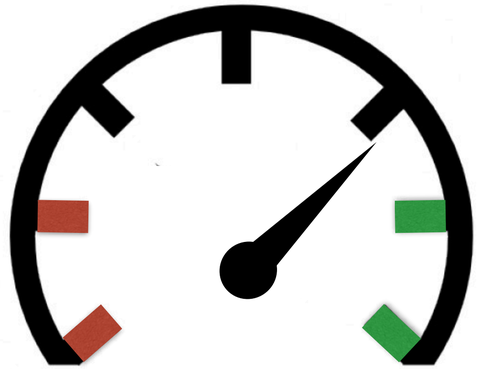

workload

Compaction

performance

What are the **design choices**?

How does a choice **affect performance**?

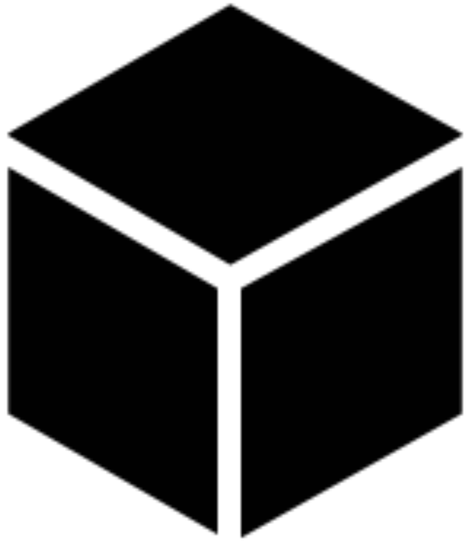

Compaction

VLDB 2021

SIGMOD 2022-a

VLDB 2021

SIGMOD 2022-a

How to organize the data on device?

How much data to move at-a-time?

Which block of data to be moved?

When to re-organize the data layout?

VLDB 2021

SIGMOD 2022-a

Data Layout

How to organize the data on device?

Compaction granularity

How much data to move at-a-time?

Data movement policy

Which block of data to be moved?

Trigger

When to re-organize the data layout?

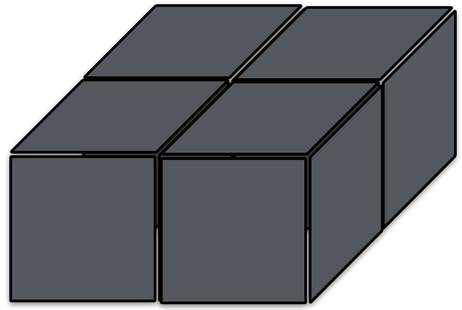

any existing + completely new
compaction strategies

Database	Data layout	Compaction Trigger					Compaction Granularity				Data Movement Policy							
		Level saturation	#Sorted runs	File staleness	Space amp.	Tombstone-TTL	Level	Sorted run	File (single)	File (multiple)	Round-robin	Least overlap (+1)	Least overlap (+2)	Coldest file	Oldest file	Tombstone density	Expired TS-TTL	N/A (entire level)
RocksDB [30], Monkey [22]	Leveling / 1-Leveling	✓		✓					✓	✓		✓		✓	✓	✓		
	Tiering		✓		✓	✓		✓										✓
LevelDB [32], Monkey (J.) [21]	Leveling	✓							✓		✓	✓	✓					
SlimDB [47]	Tiering	✓							✓	✓								✓
Dostoevsky [23]	L-leveling	$✓^L$	$✓^T$				$✓^L$	$✓^T$				$✓^L$						$✓^T$
LSM-Bush [24]	Hybrid leveling	$✓^L$	$✓^T$				$✓^L$	$✓^T$				$✓^L$						$✓^T$
Silk [11], Silk+ [12]	Leveling	✓							✓	✓	✓							
HyperLevelDB [35]	Leveling	✓							✓		✓	✓	✓					
PebblesDB [46]	Hybrid leveling	✓							✓	✓								✓
Cassandra [8]	Tiering		✓	✓		✓		✓										✓
	Leveling	✓				✓			✓	✓		✓				✓	✓	
WiredTiger [62]	Leveling	✓					✓											✓
X-Engine [34], Leaper [63]	Hybrid leveling	✓							✓	✓		✓				✓		
HBase [7]	Tiering		✓					✓										✓
AsterixDB [3]	Leveling	✓					✓											✓
	Tiering		✓					✓										✓
Tarantool [57]	L-leveling	$✓^L$	$✓^T$				$✓^L$	$✓^T$										✓

FADE

FAst **DE**lete

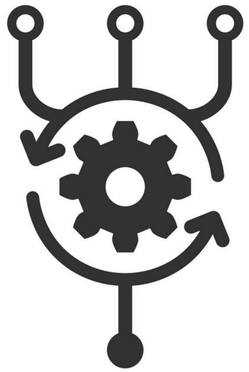

family of
compaction strategies

FAst **DE**lete

compaction trigger	compaction file picking policy
	

decompose tasks

piggyback

FAst **DE**lete

delete(5) within threshold: D_{th}

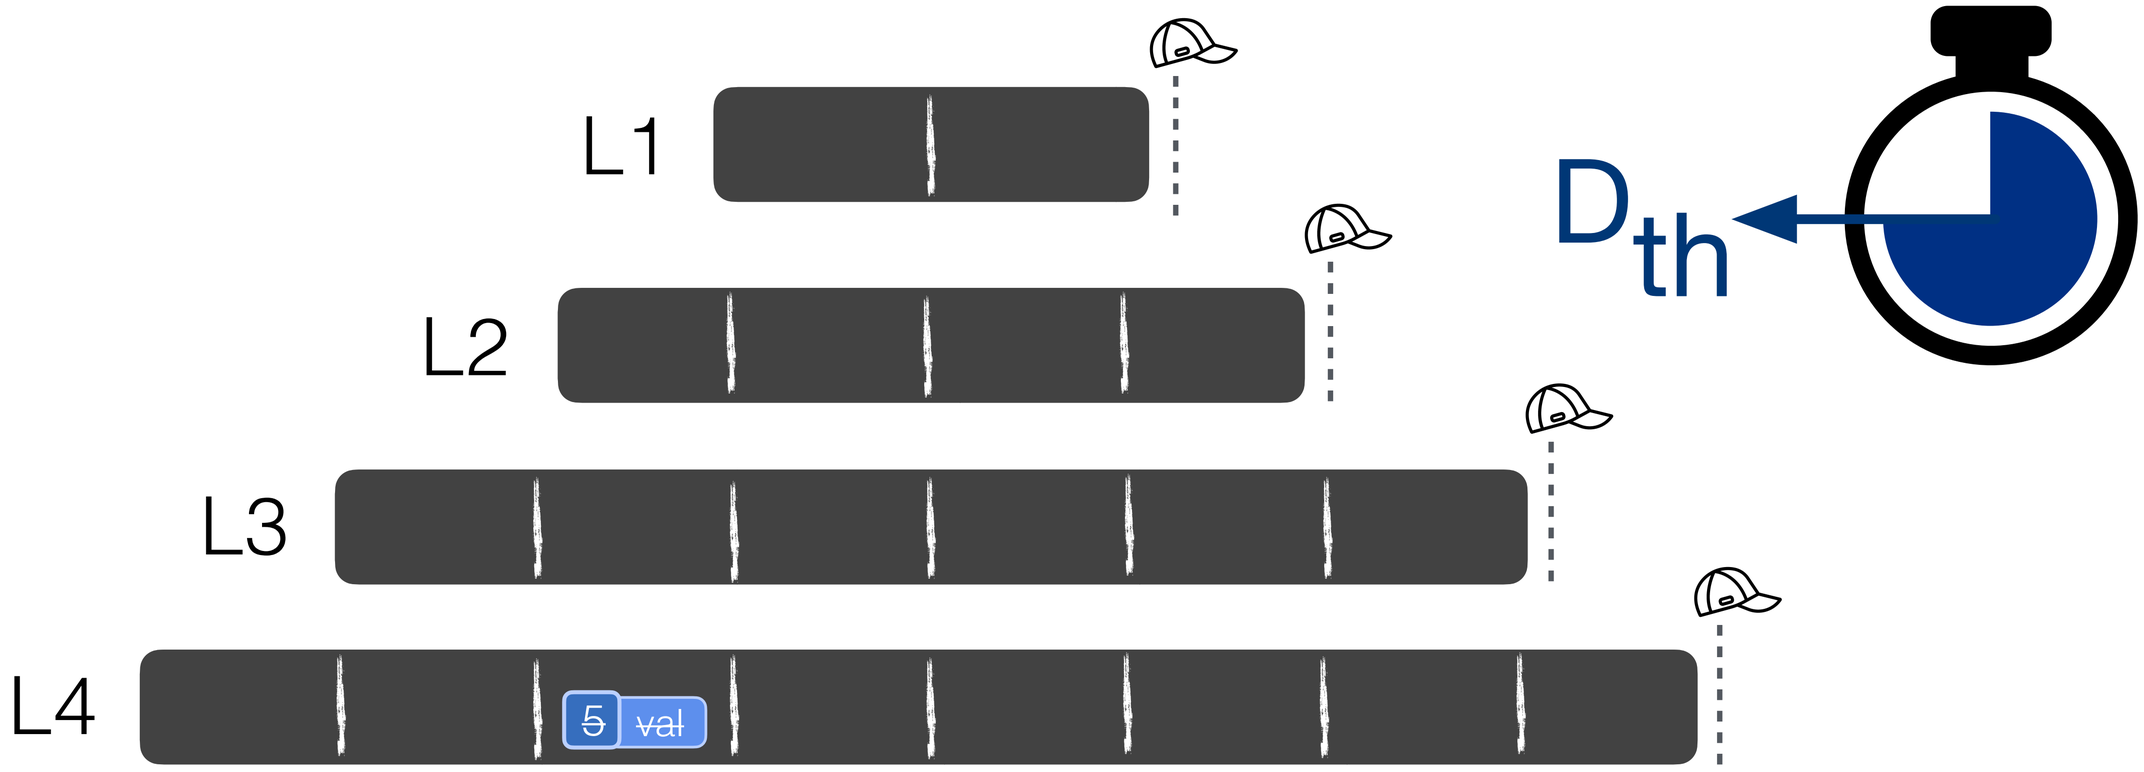

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

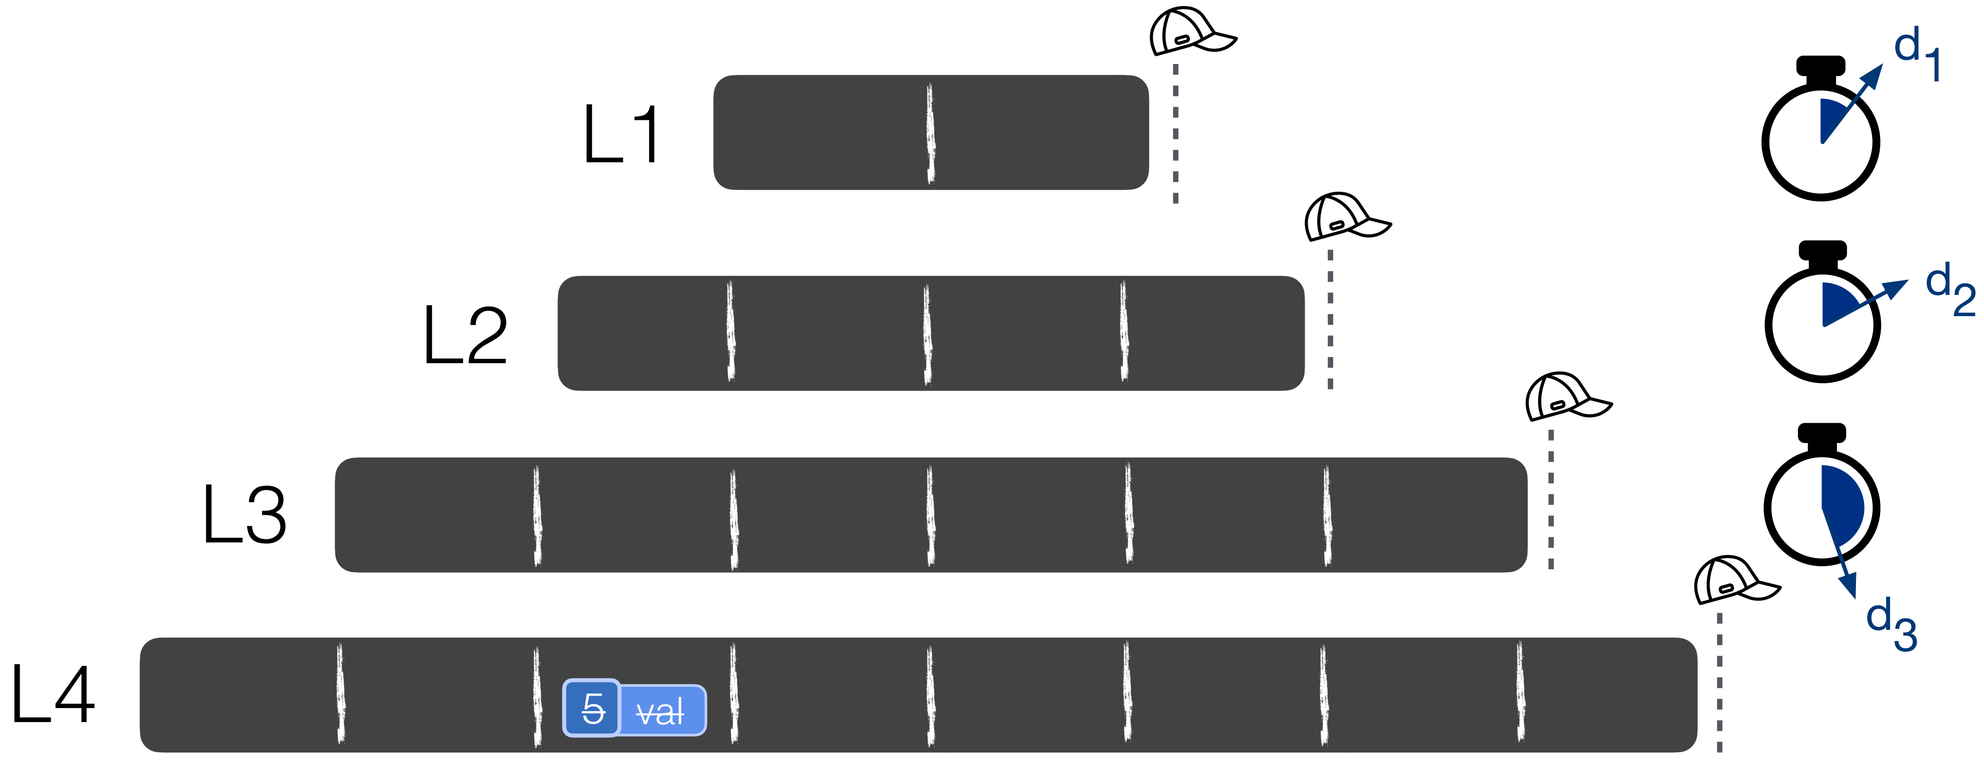

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

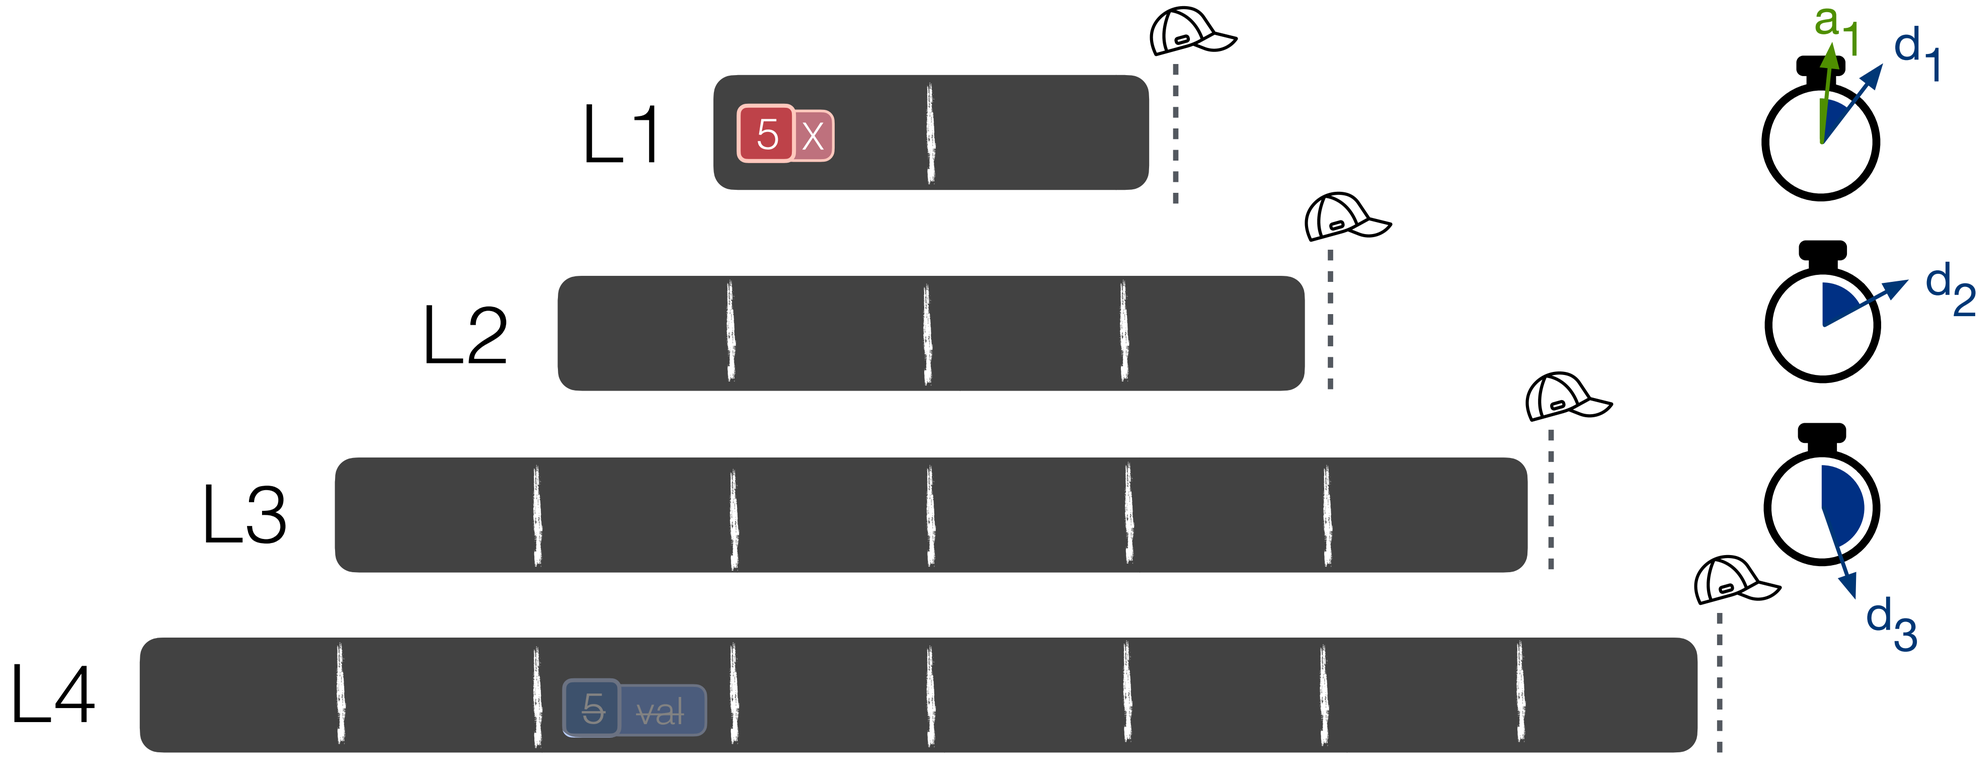

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

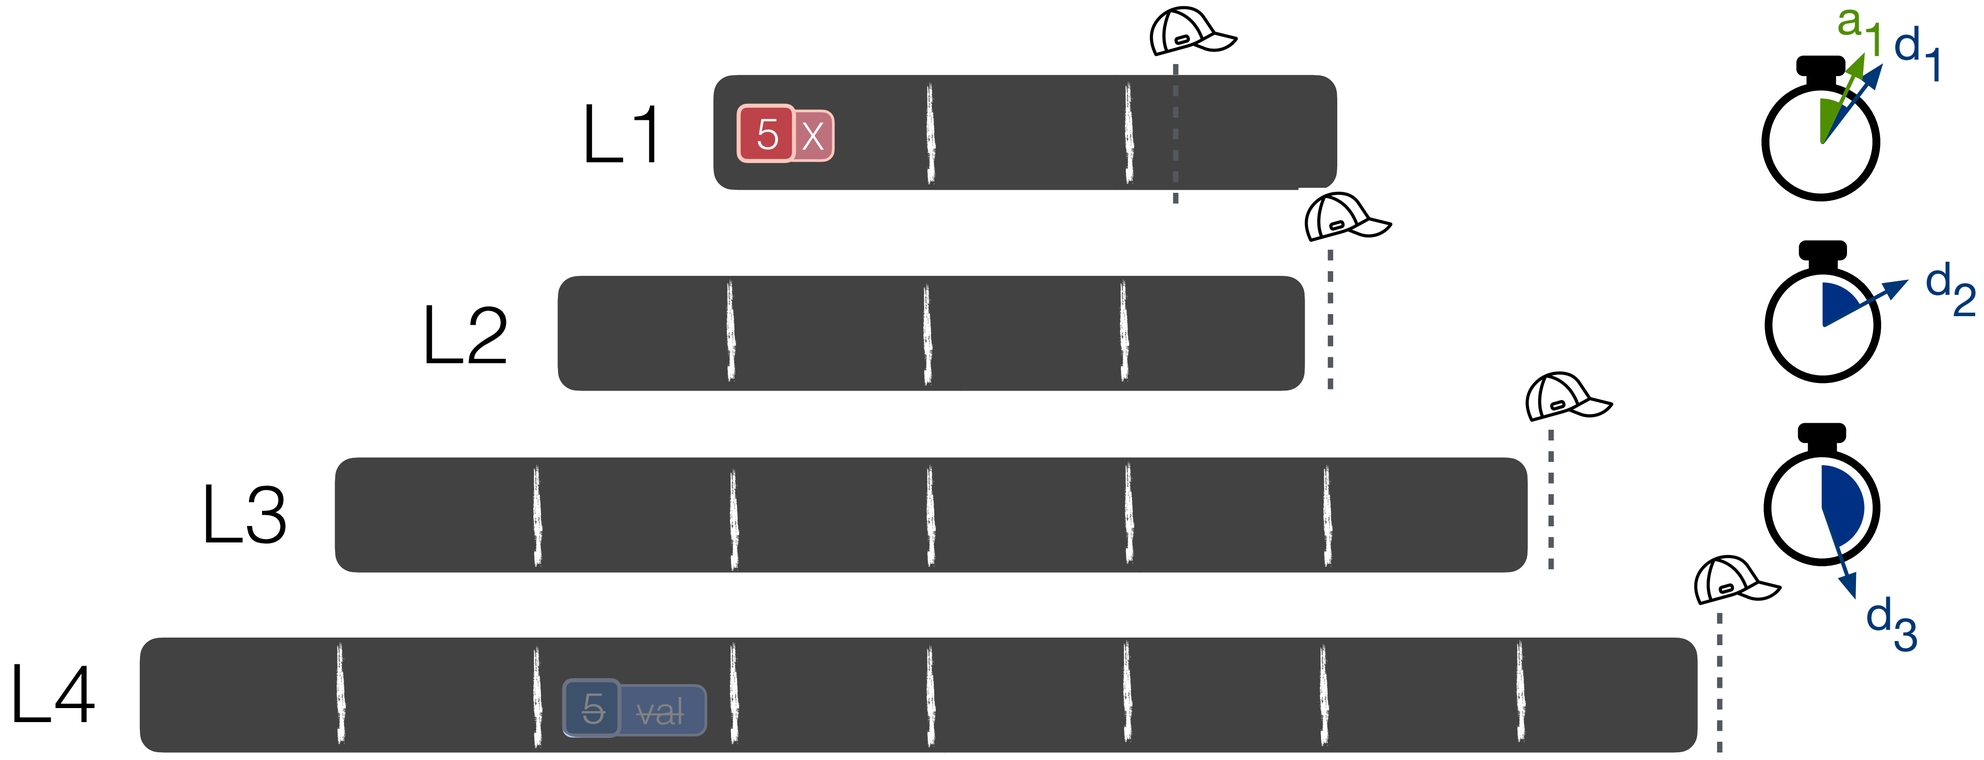

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

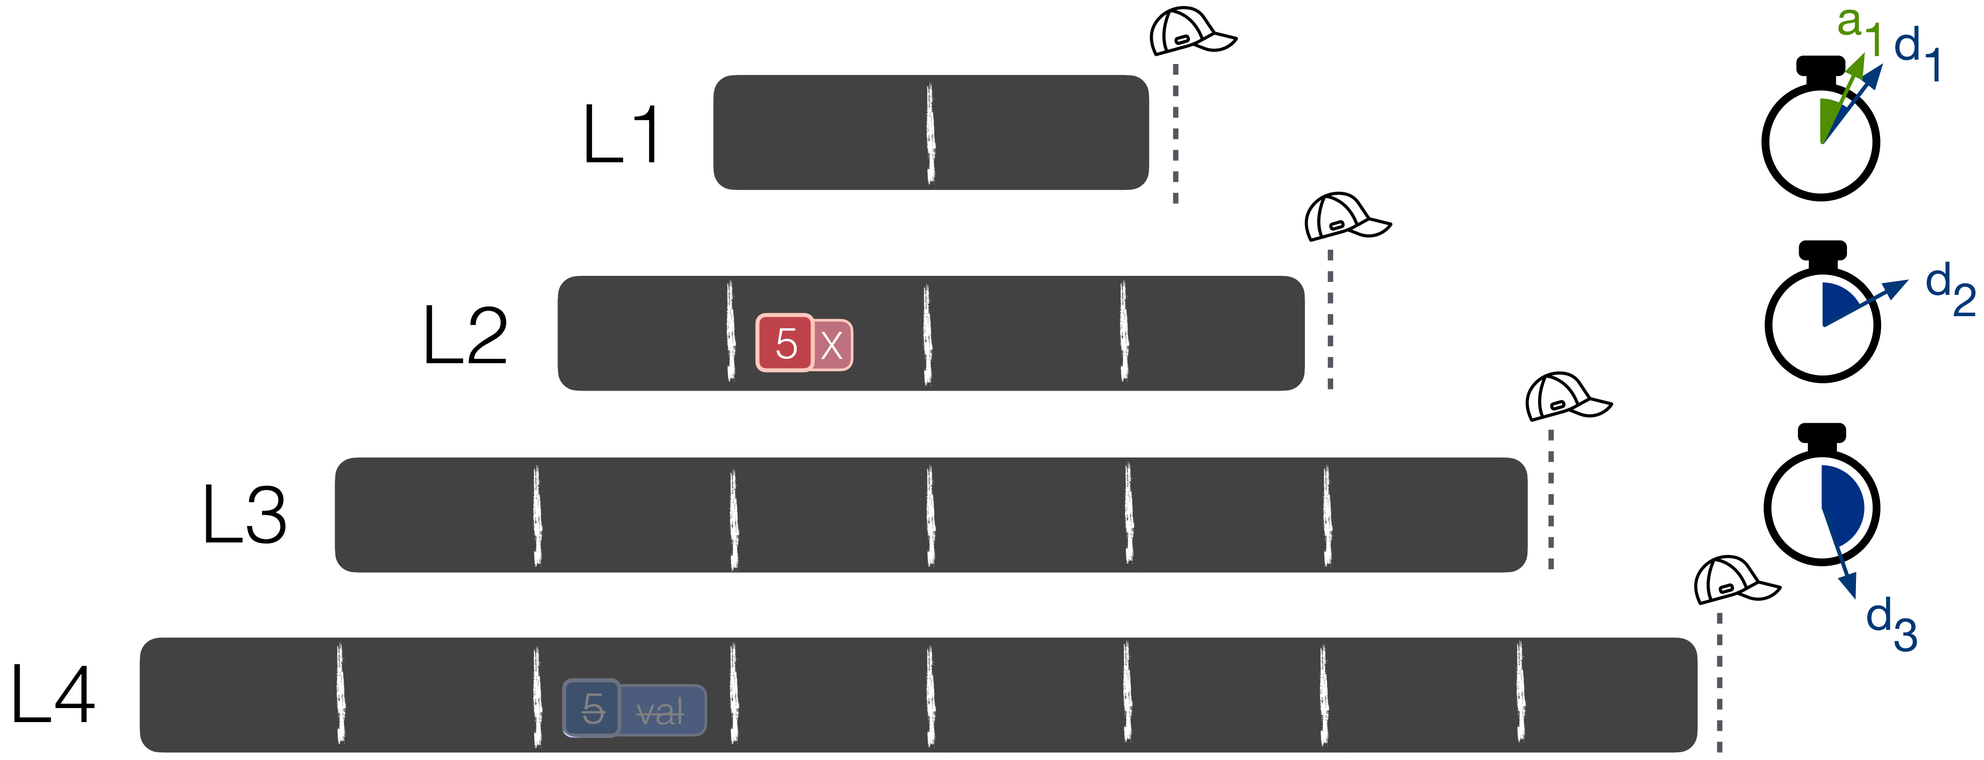

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

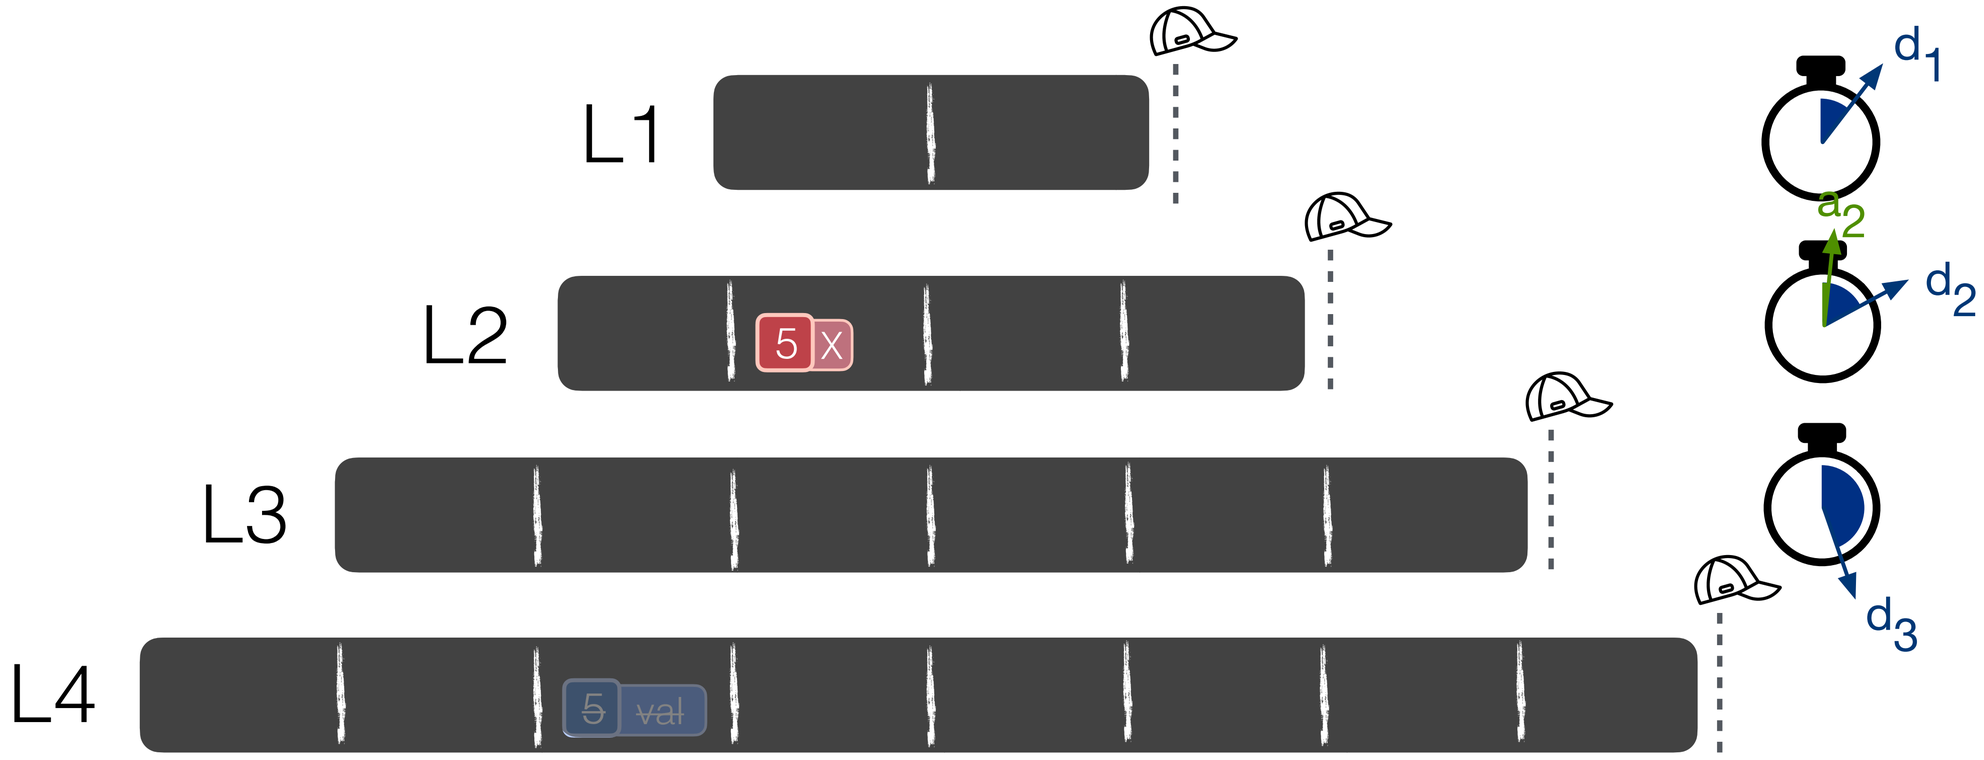

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

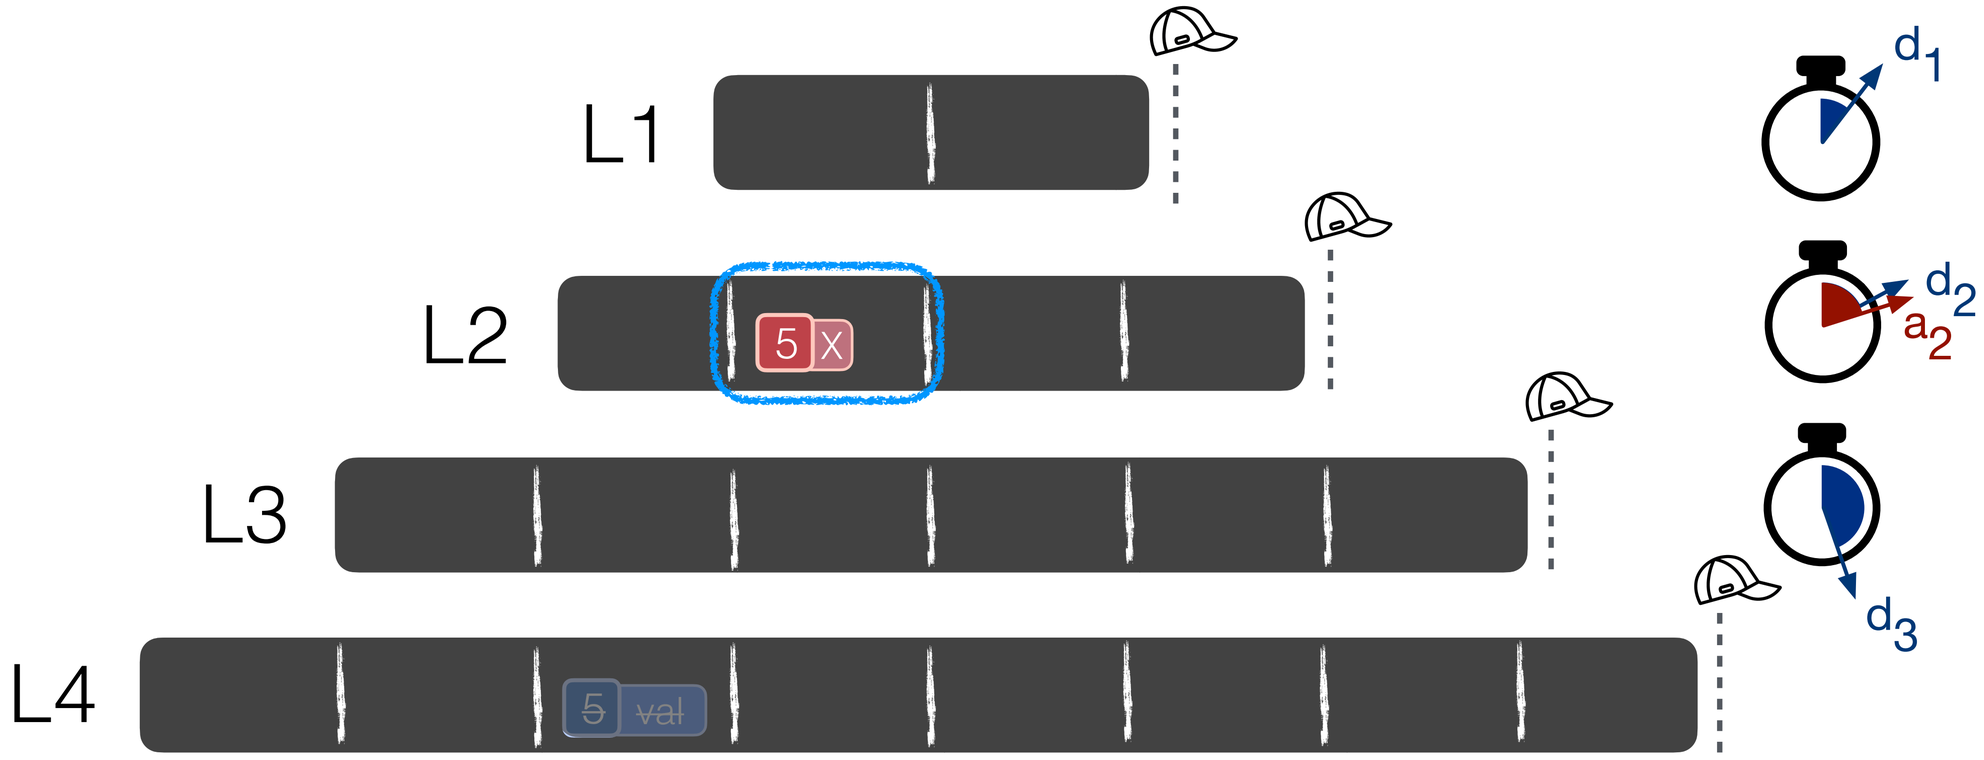

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

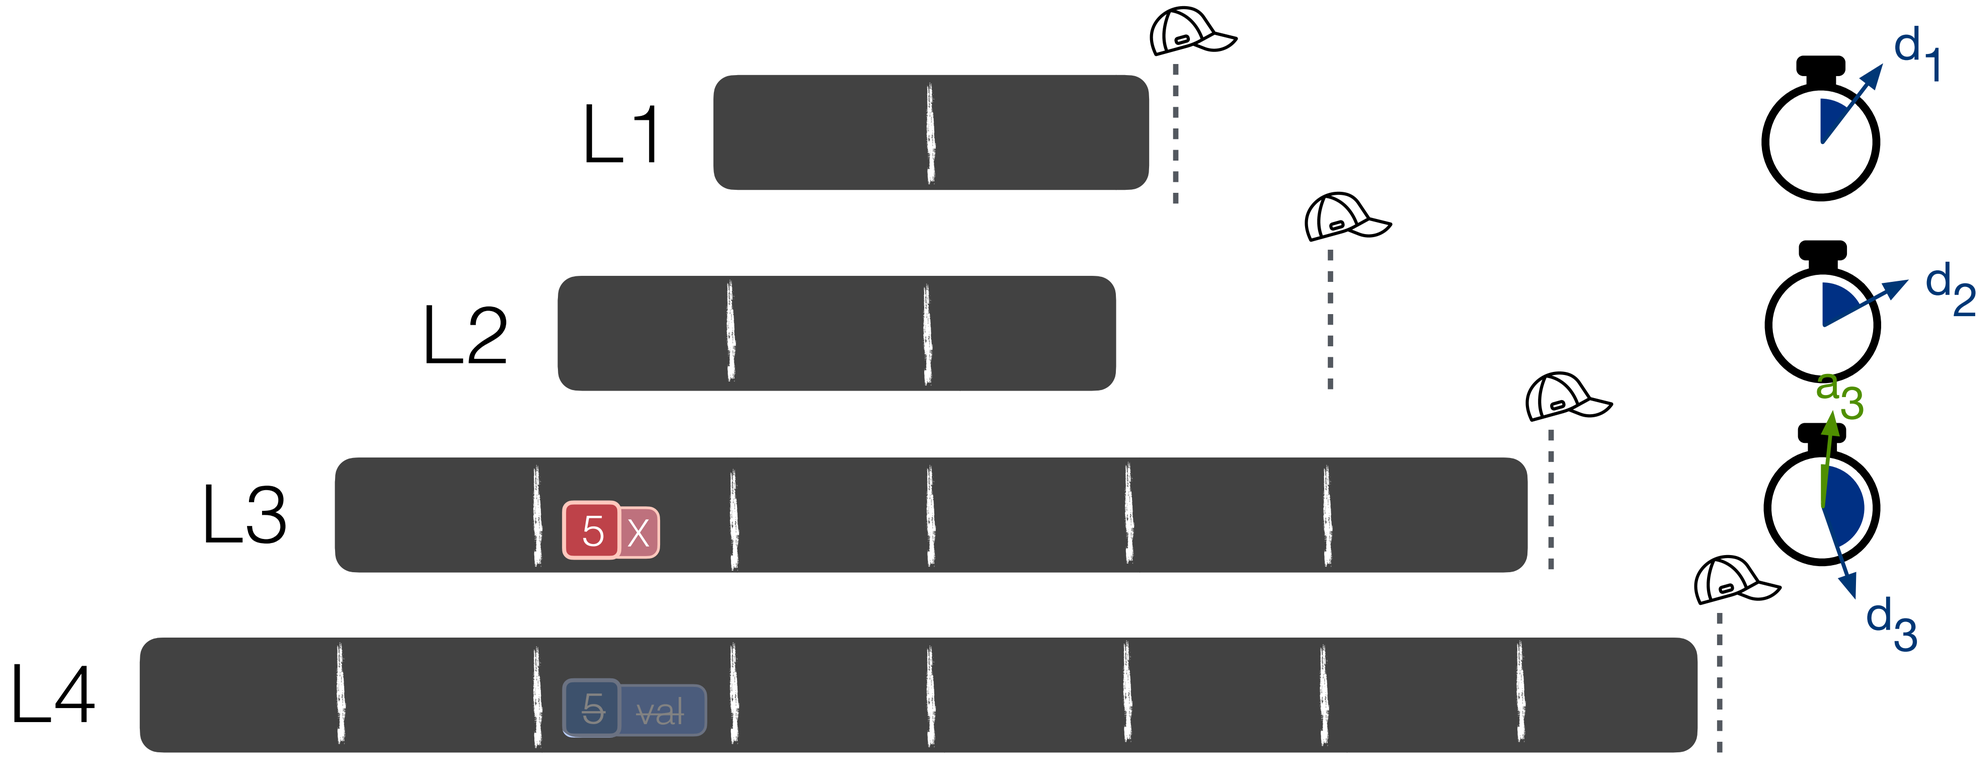

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

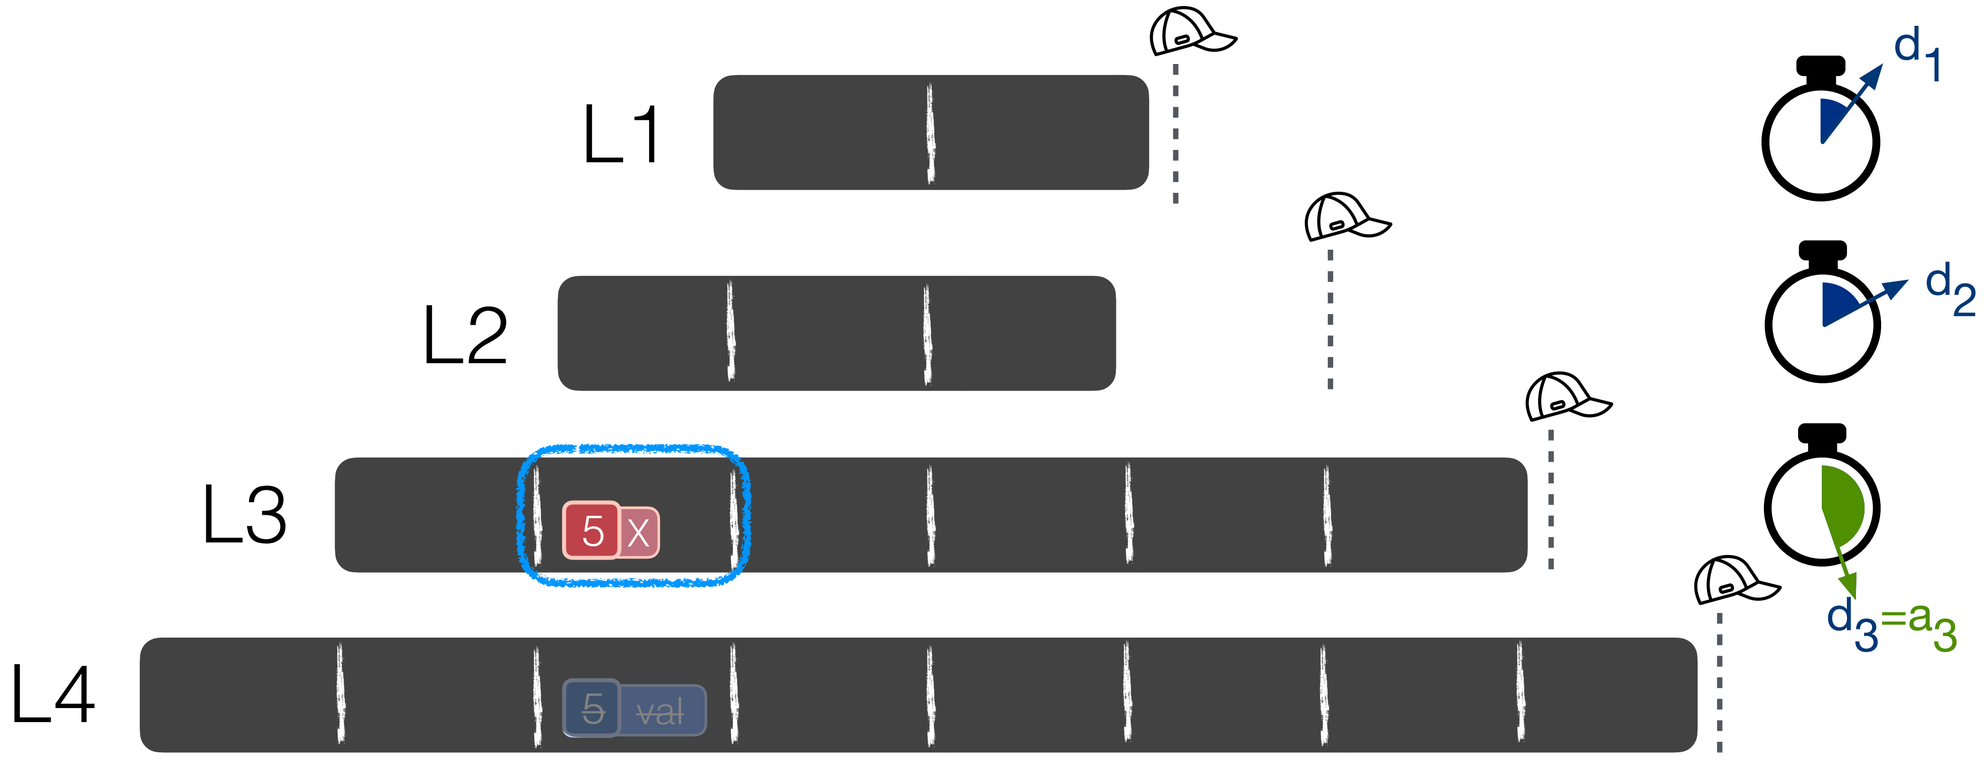

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

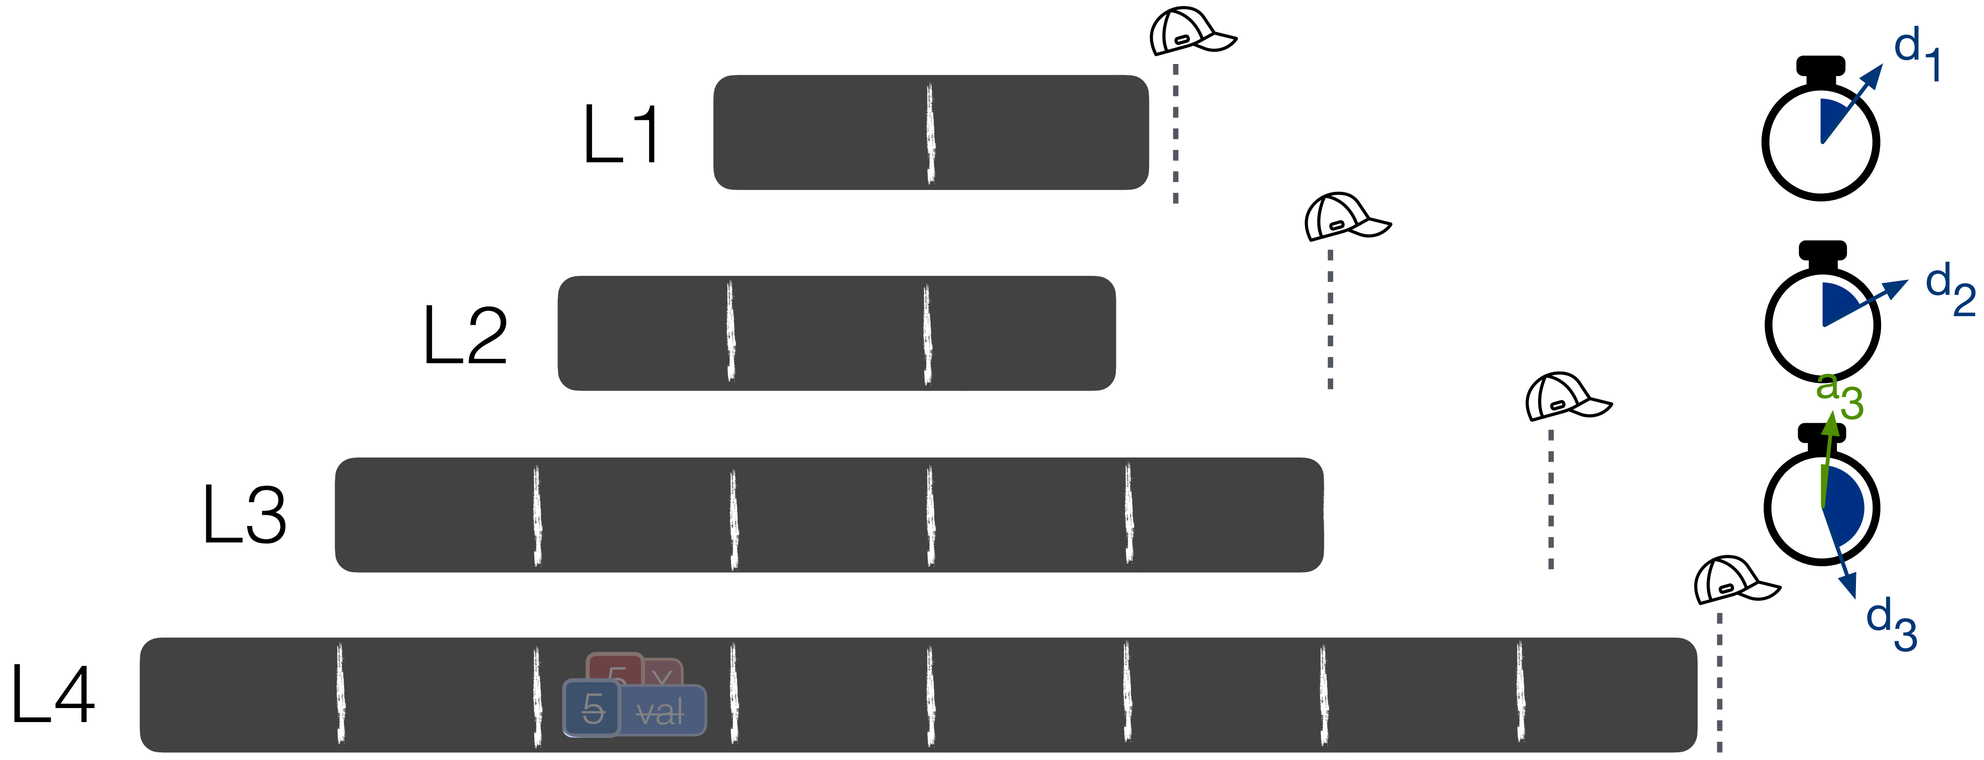

FAst **DE**lete

$$\sum_{i=1}^{L-1} d_i \leq D_{th}$$

$$d_i = T \cdot d_{i-1}$$

delete(5) within threshold: D_{th}

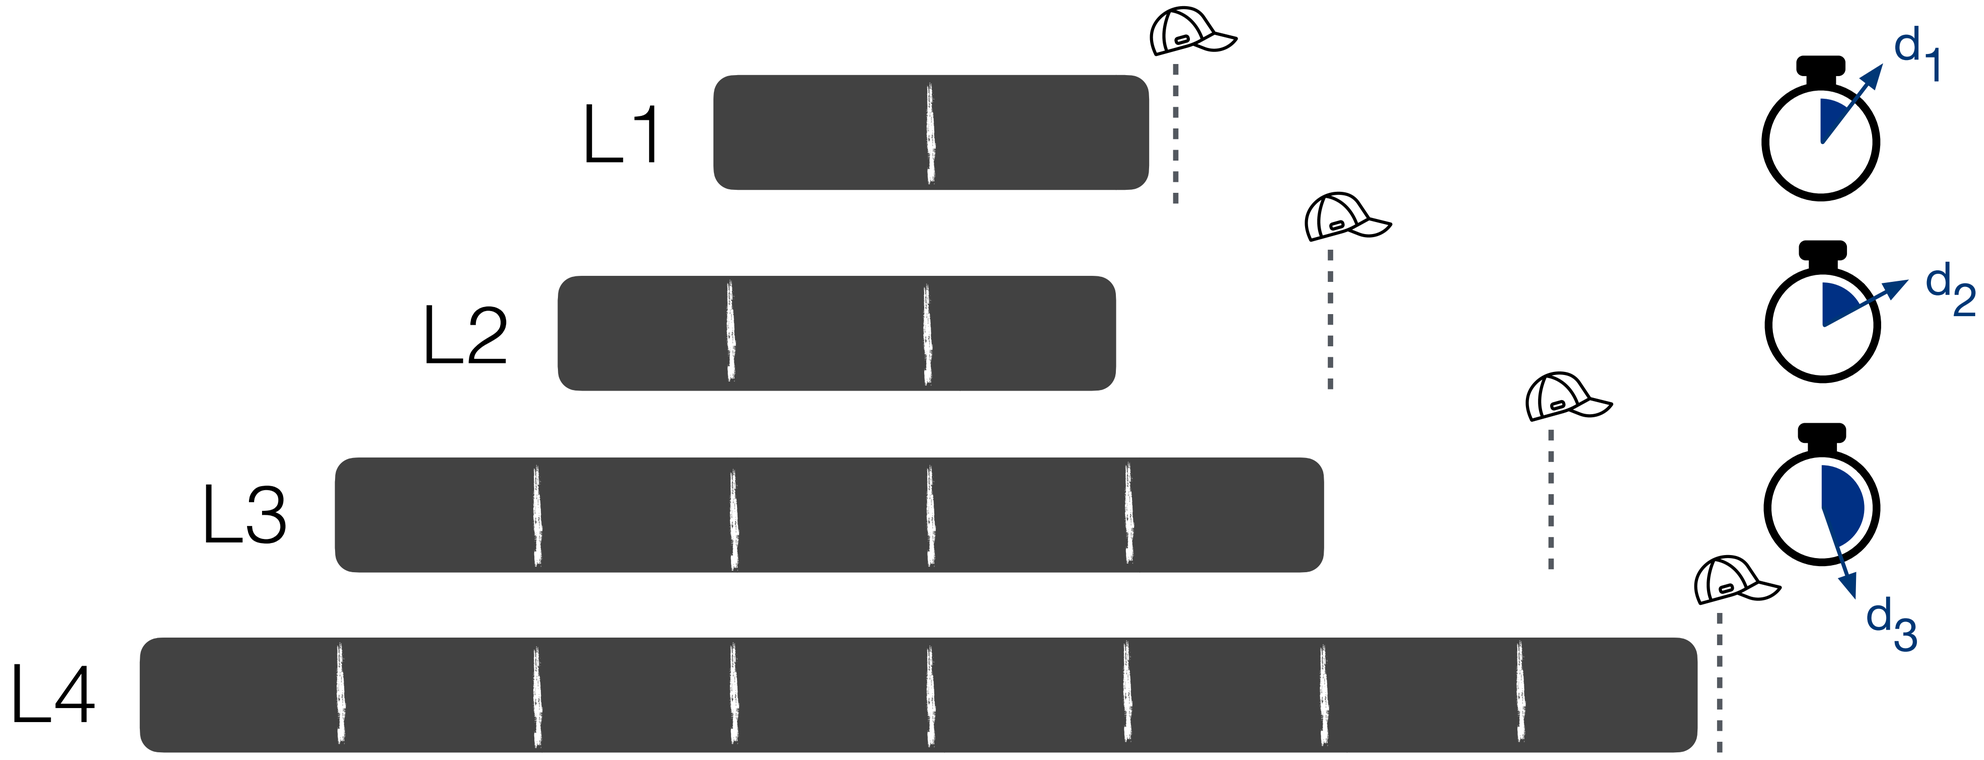

persist all deletes within:

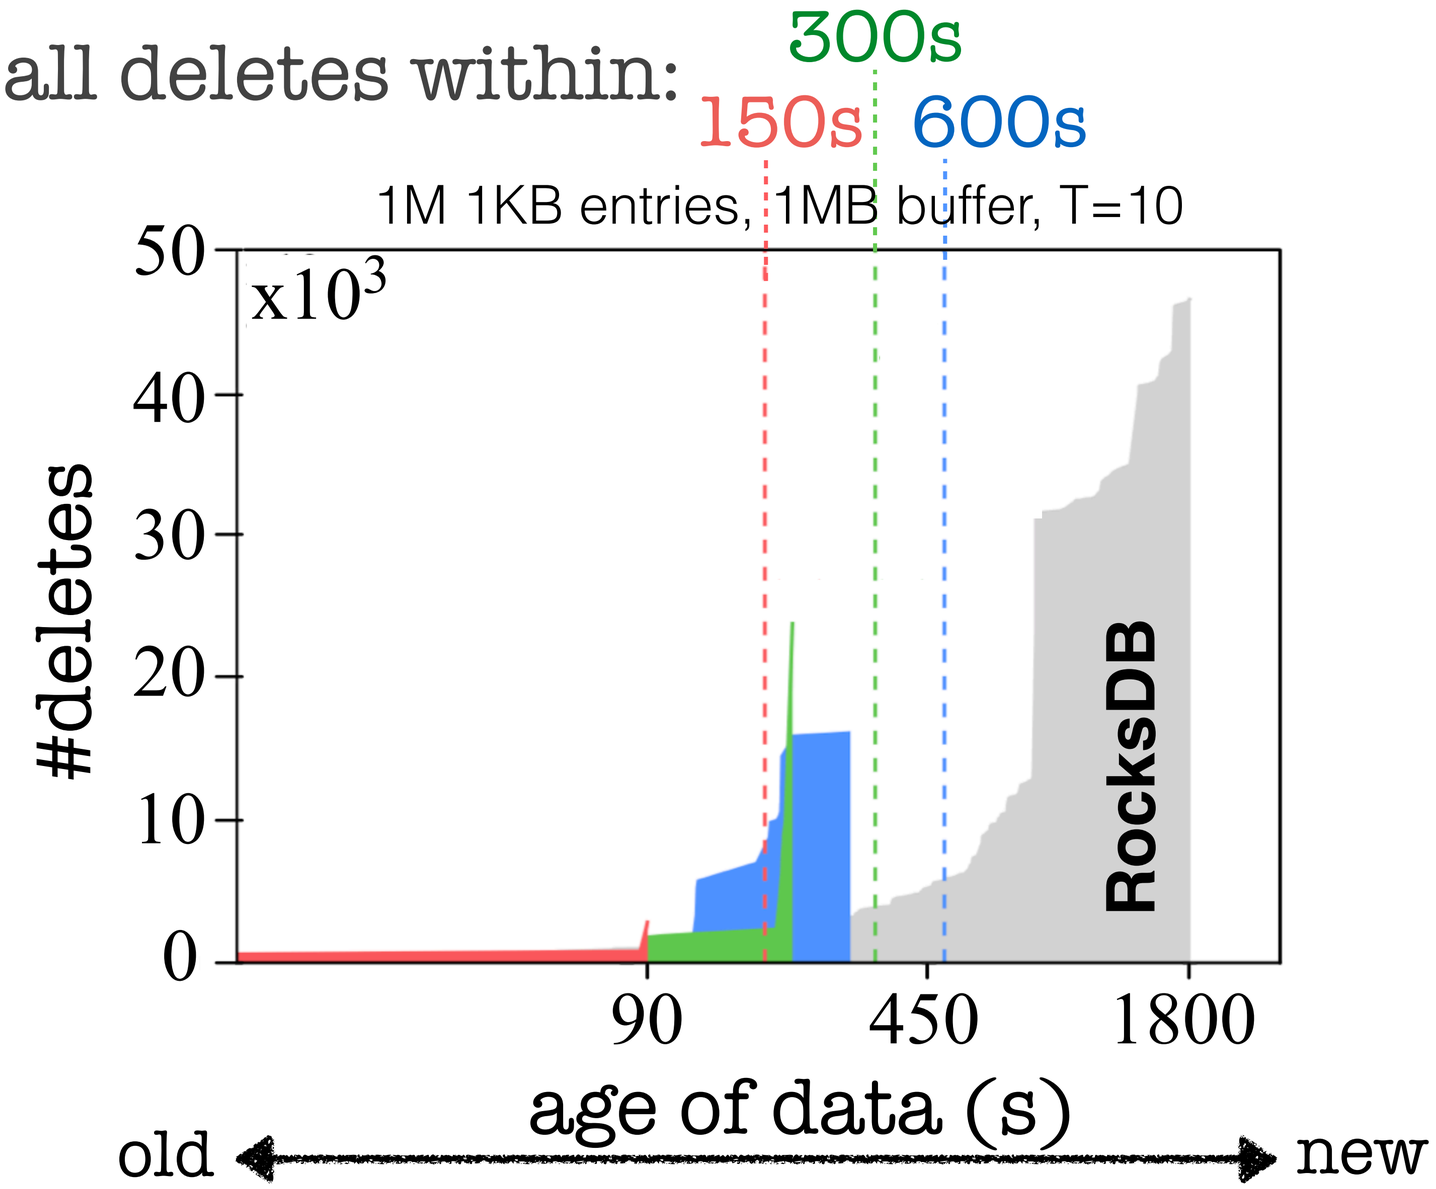

persists deletes timely

within threshold

reduced space amplification

2.1x - 9.8x

persists deletes timely

within threshold

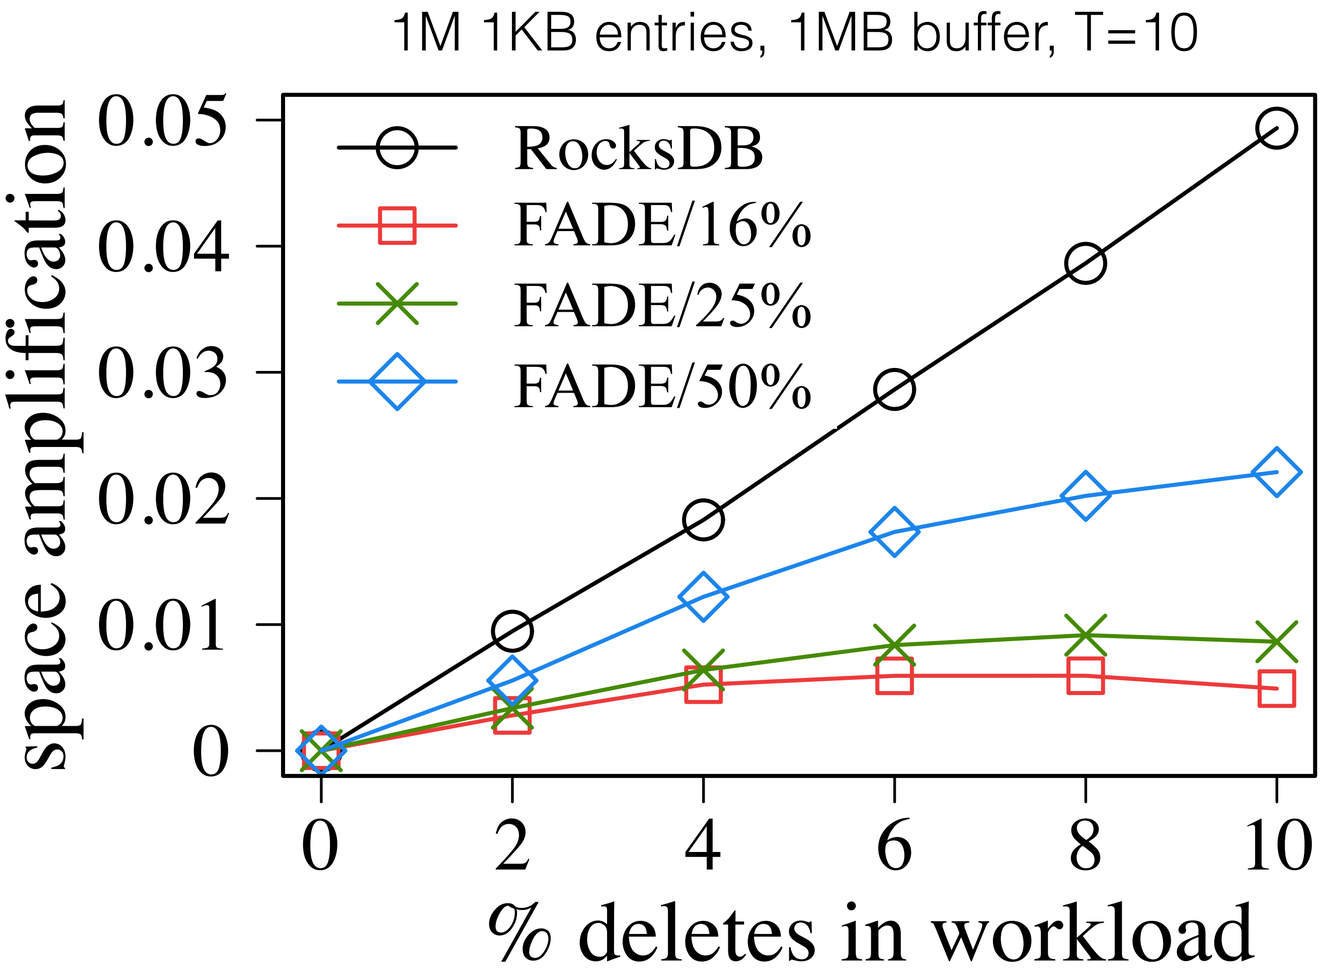

improved read performance

1.2x - 1.4x

reduced space amplification

2.1x - 9.8x

persists deletes timely

within threshold

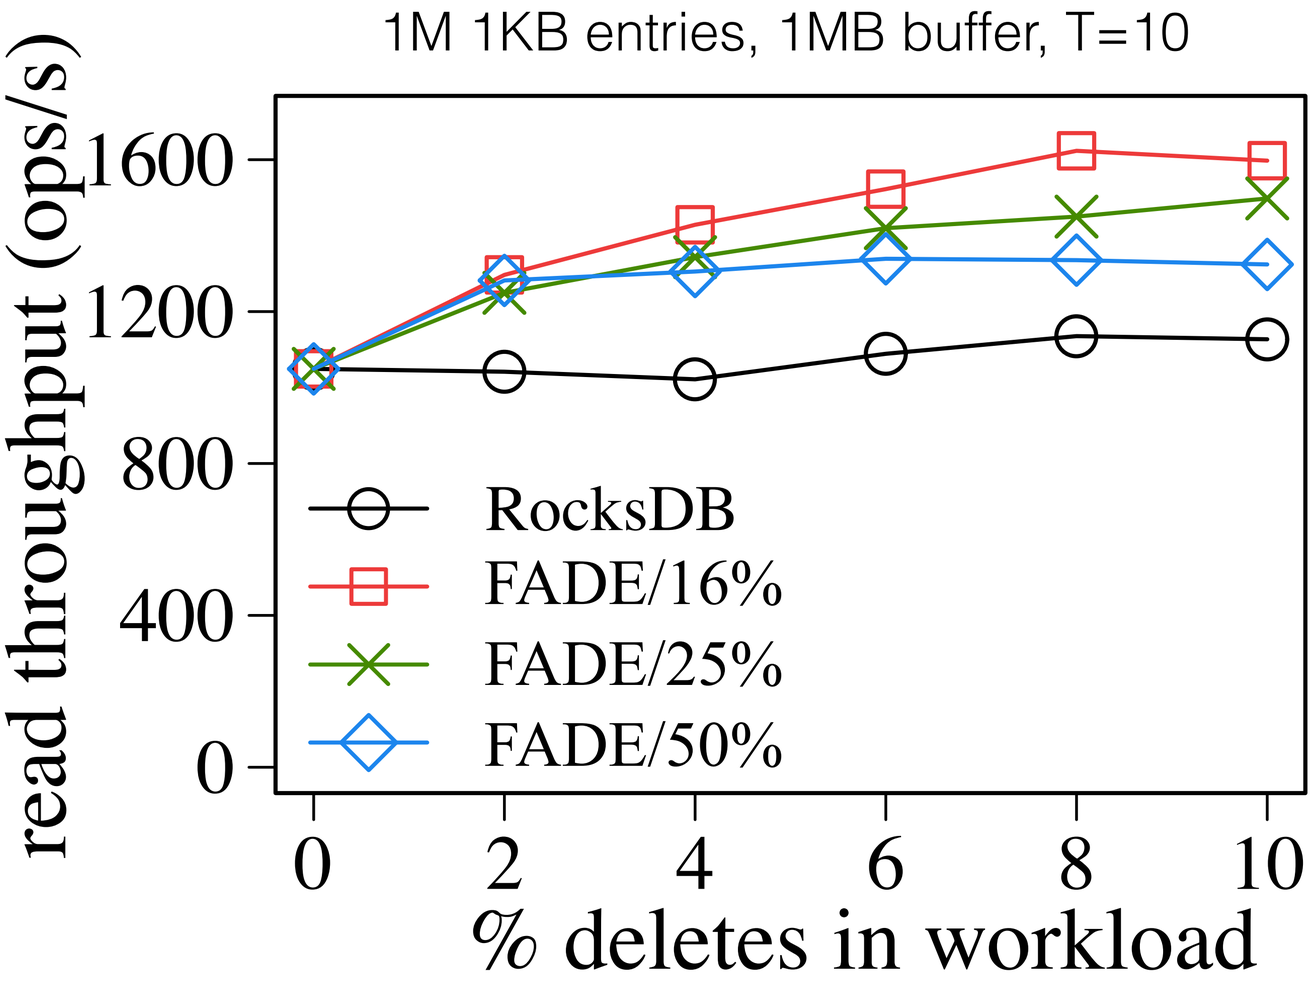

higher write amplification

4% - 25%

improved read performance

1.2x - 1.4x

reduced space amplification

2.1x - 9.8x

persists deletes timely

within threshold

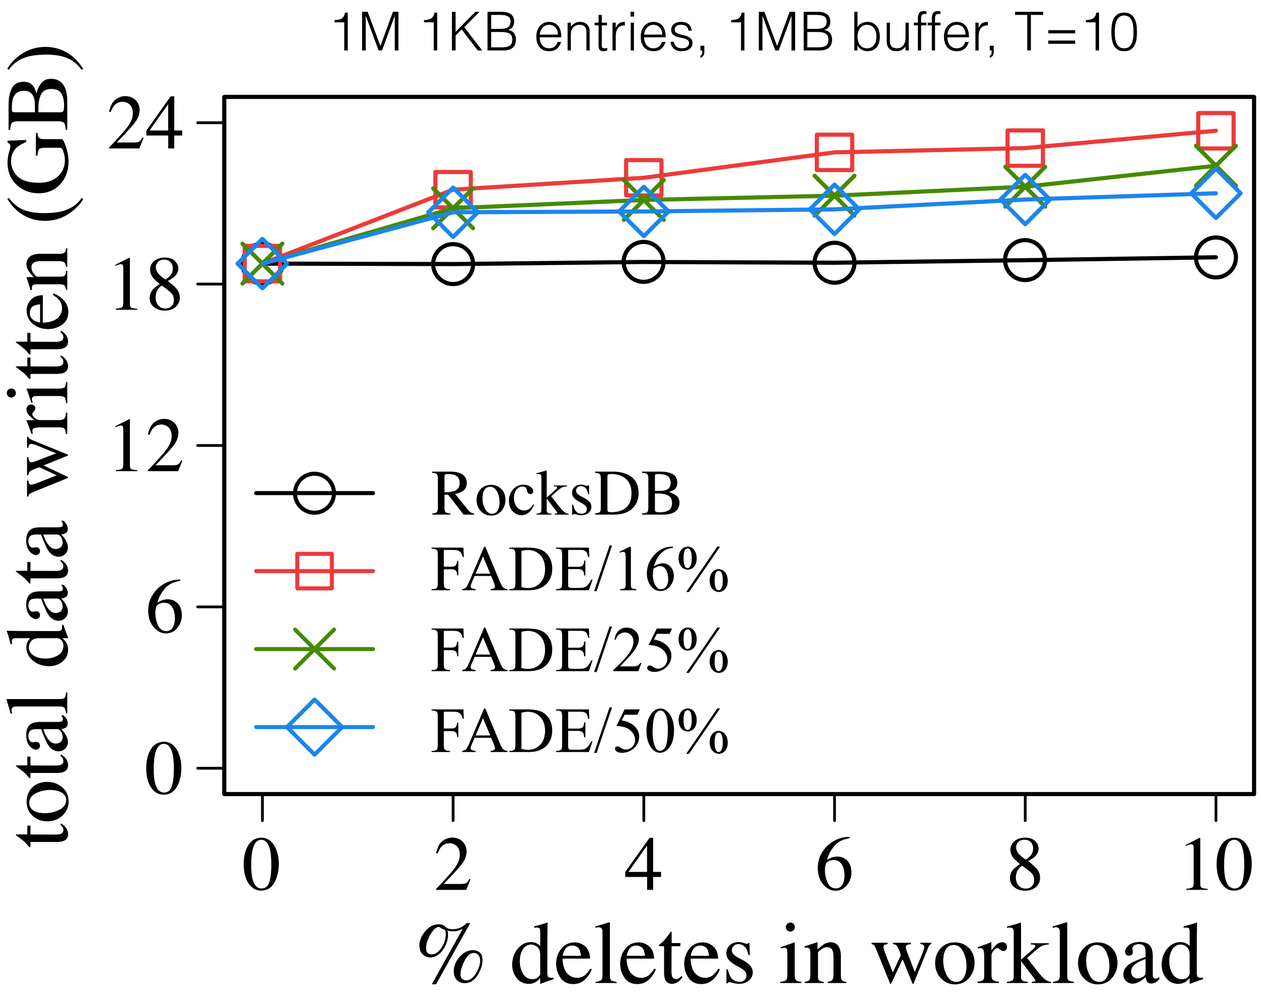

higher write amplification

4% - 25%

improved read performance

1.2x - 1.4x

reduced space amplification

2.1x - 9.8x

persists deletes timely

within threshold

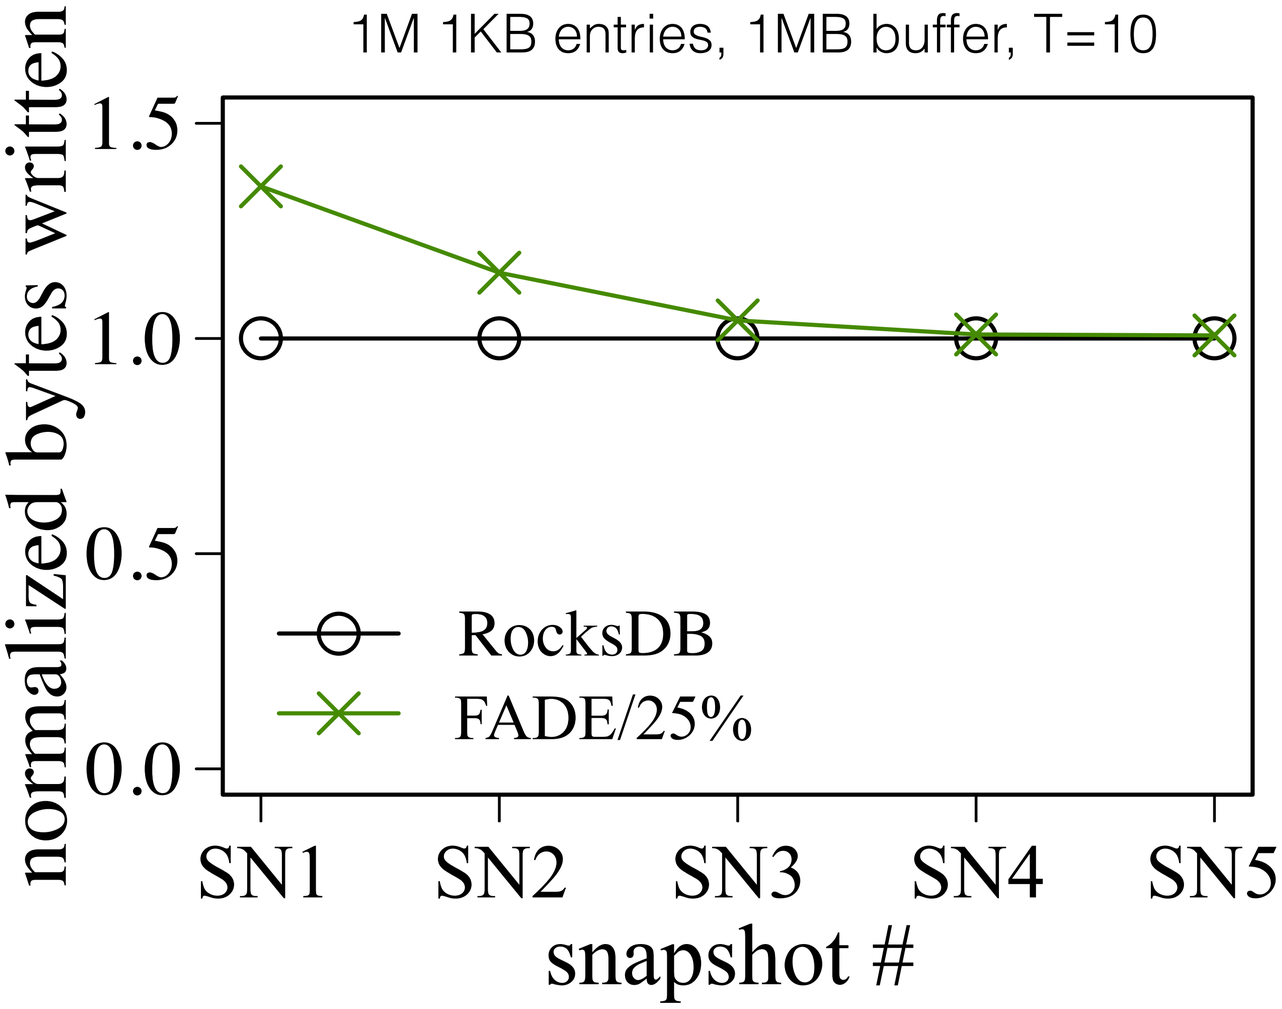

higher write amplification

0.7%

improved read performance

1.2x - 1.4x

reduced space amplification

2.1x - 9.8x

persists deletes timely

within threshold

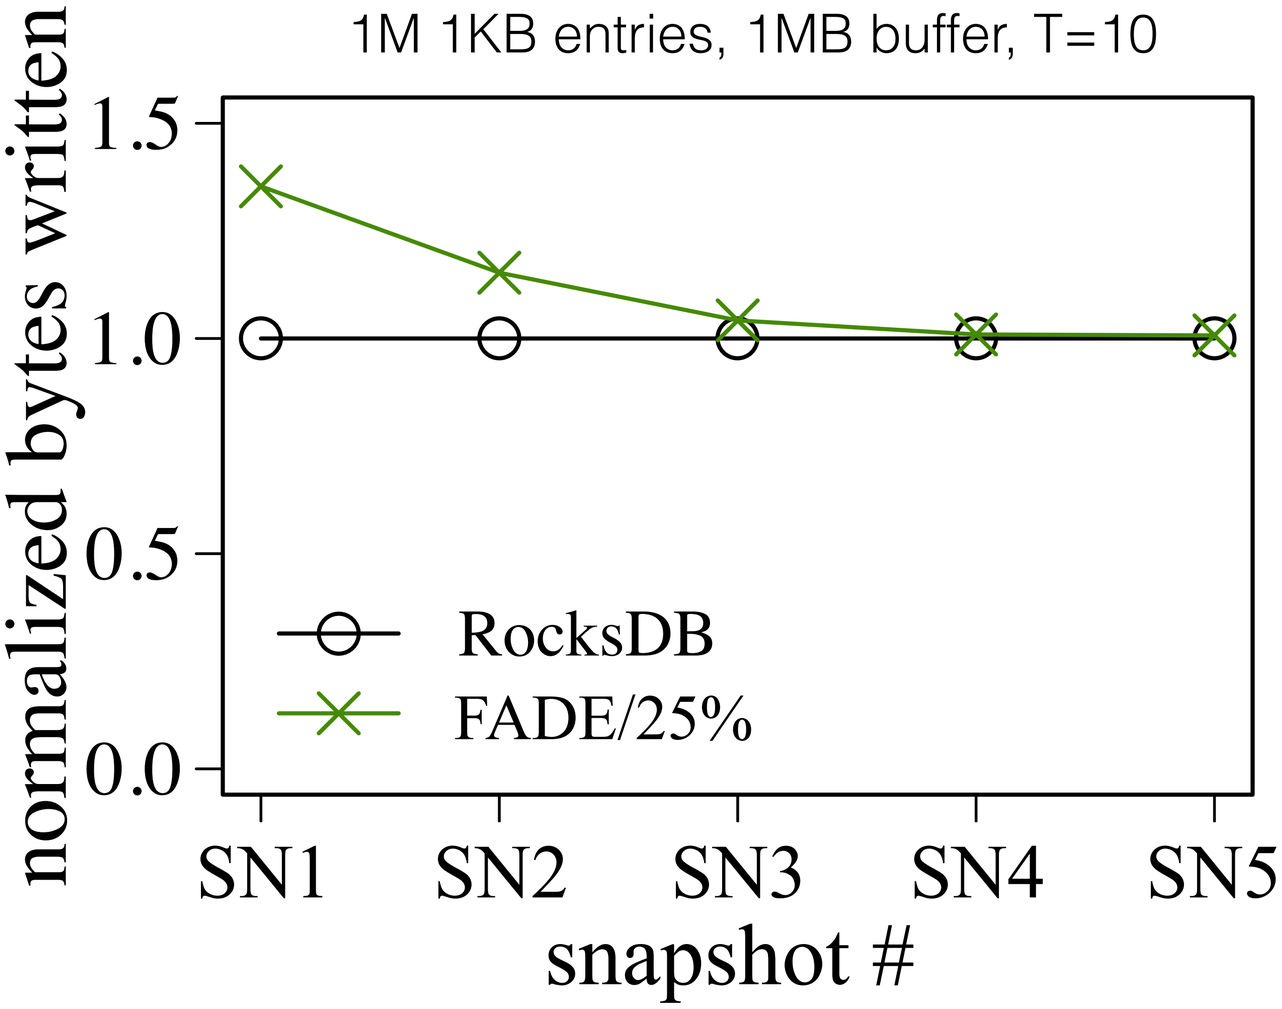

EDBT 2022

DEBull 2022

FADE

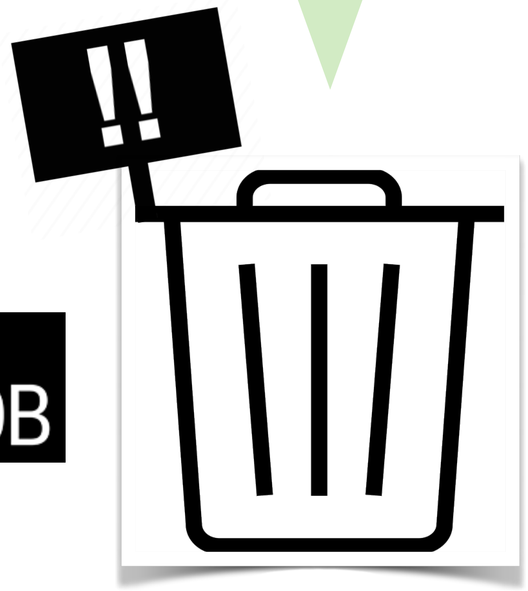

RocksDB

on-demand

persist all logical
deletes within D days

FADE

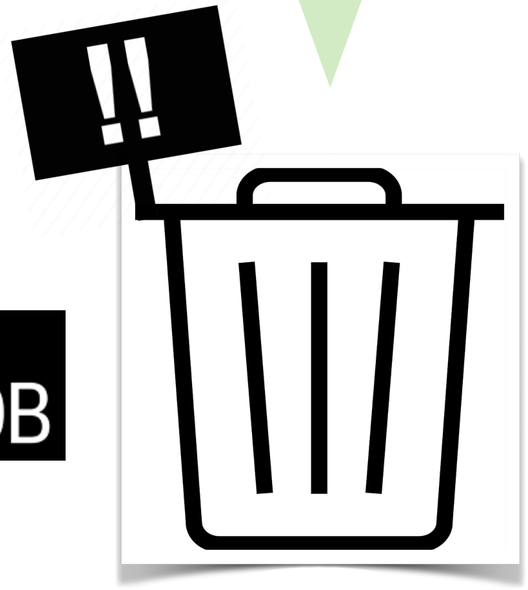

RocksDB

on-demand

EDBT 2022

DEBull 2022

retention-based

persist all logical
deletes within D days

delete **all data older**
than T days

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

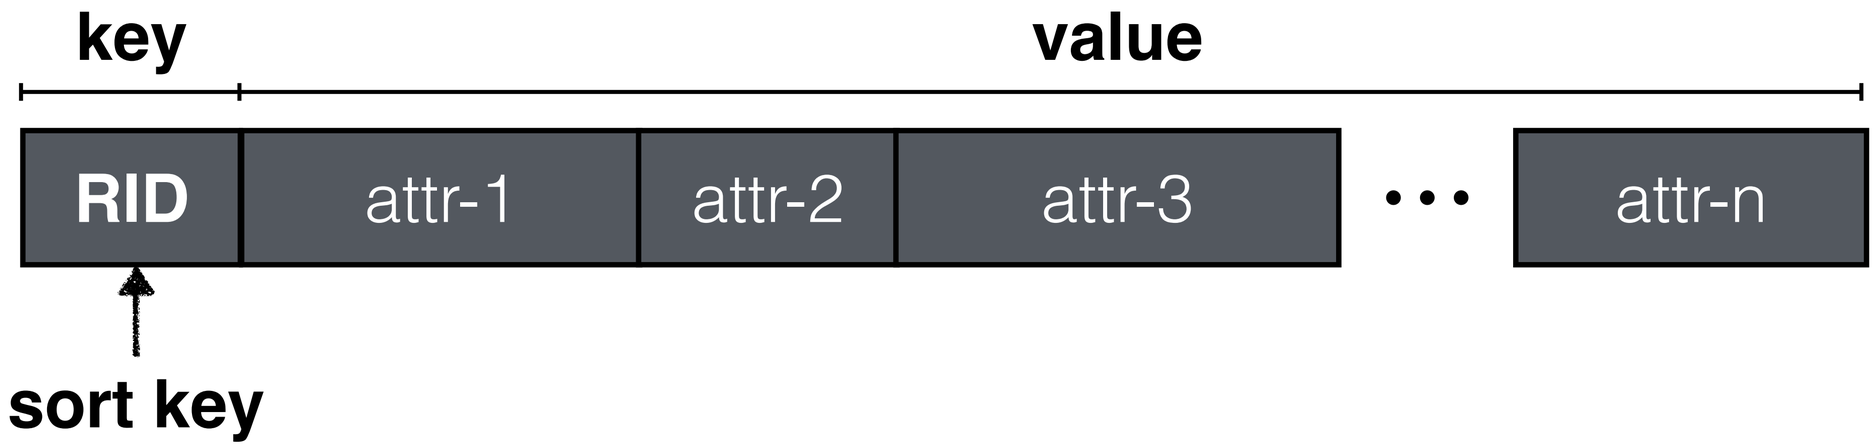

sort key = delete key

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

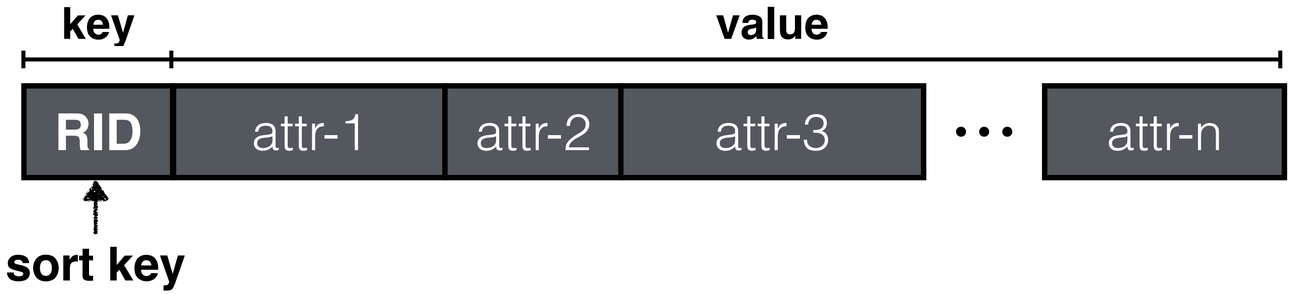

sort key = delete key

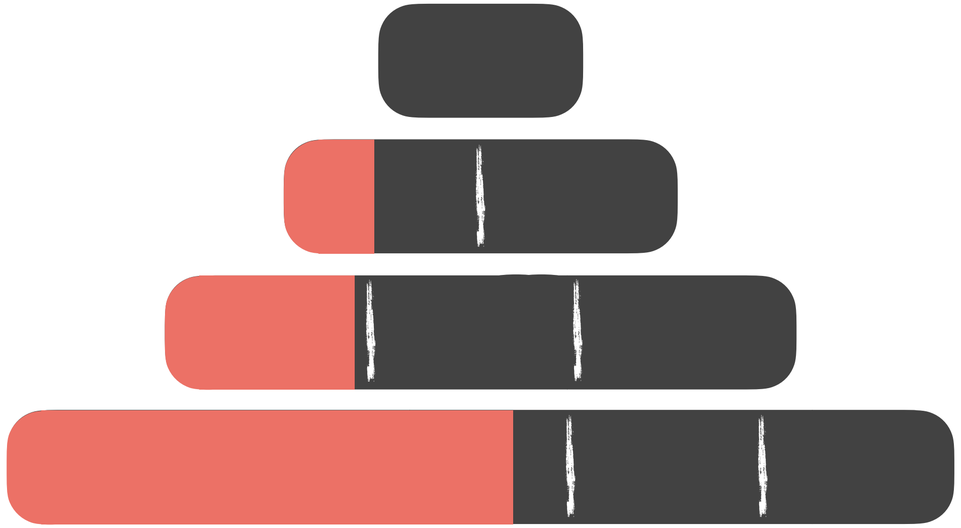

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

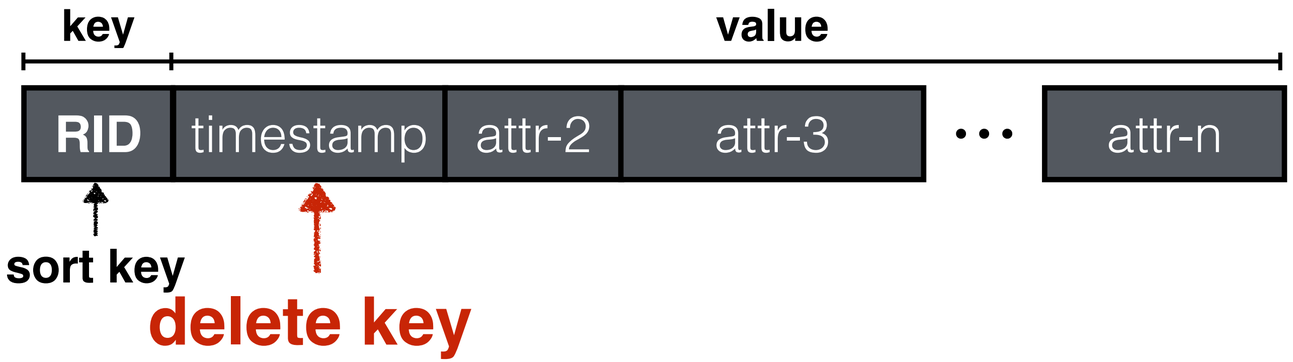

sort key ≠ delete key

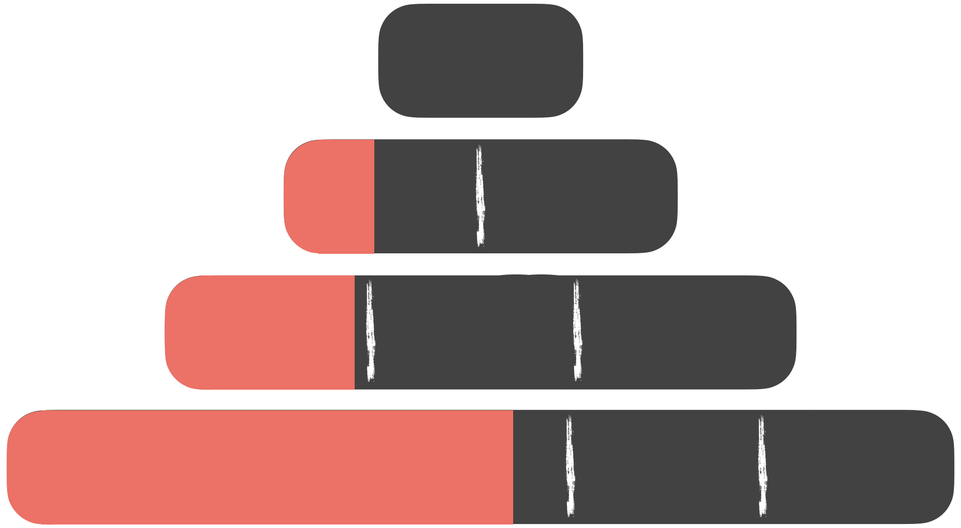

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

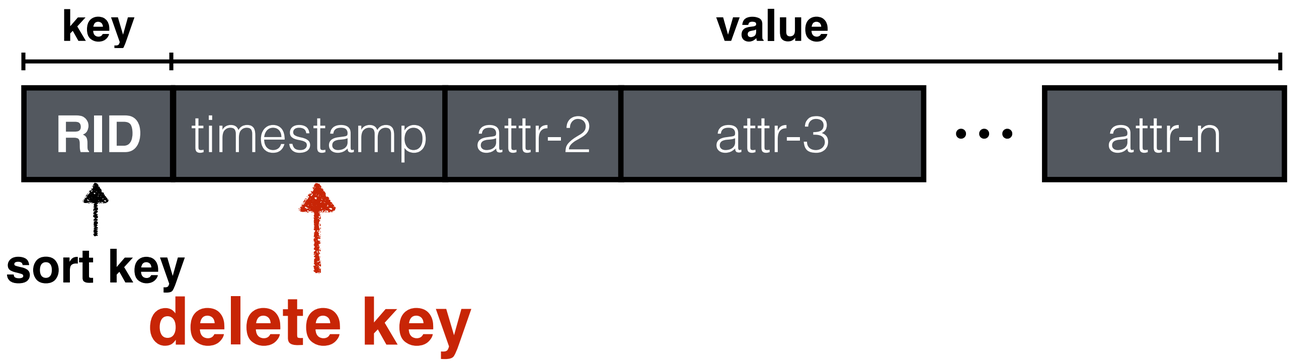

sort key ≠ delete key

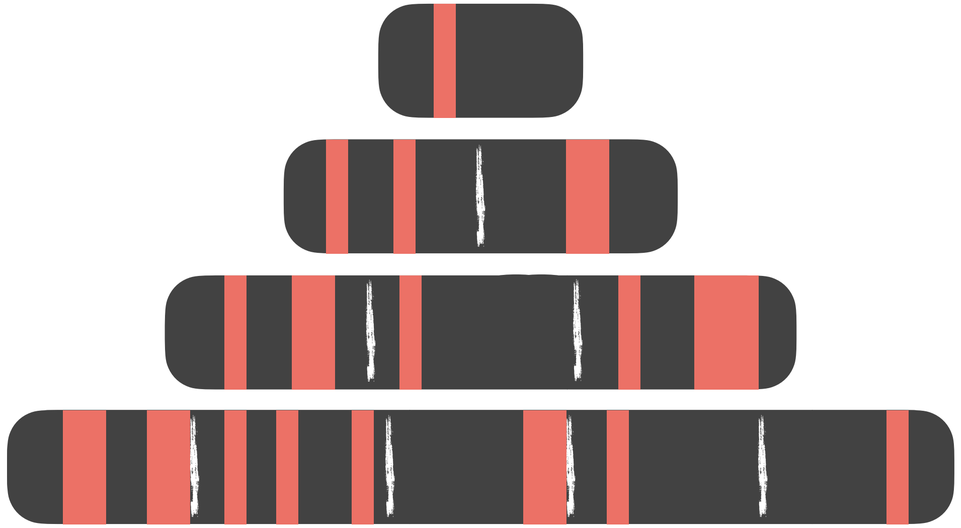

scattered occurrences

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

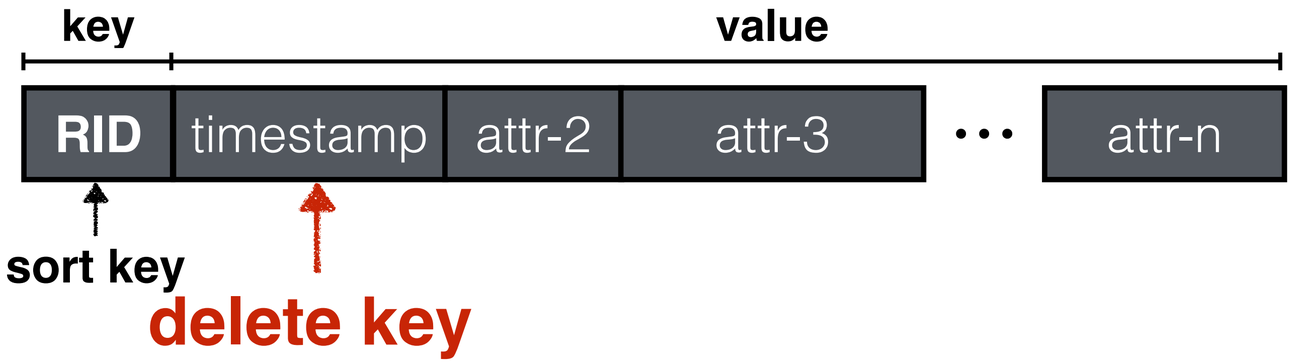

sort key ≠ delete key

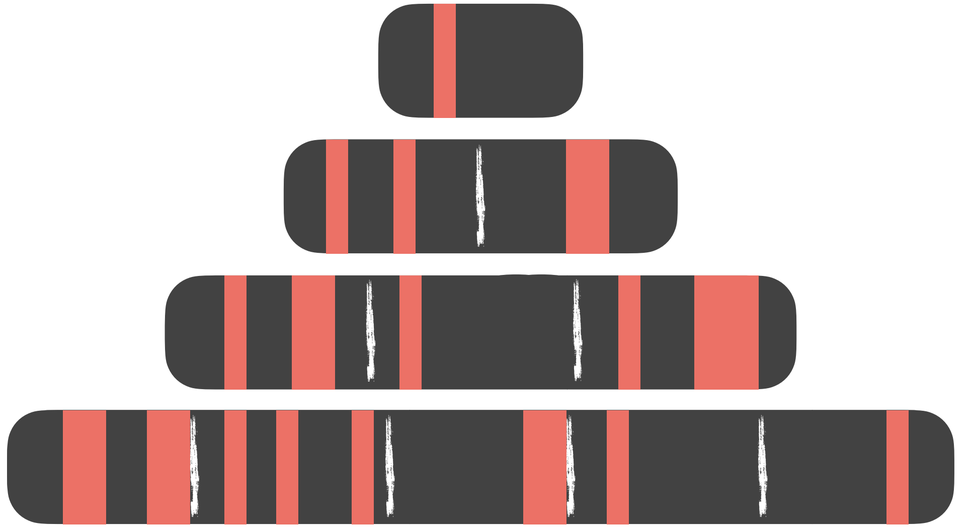

scattered occurrences

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

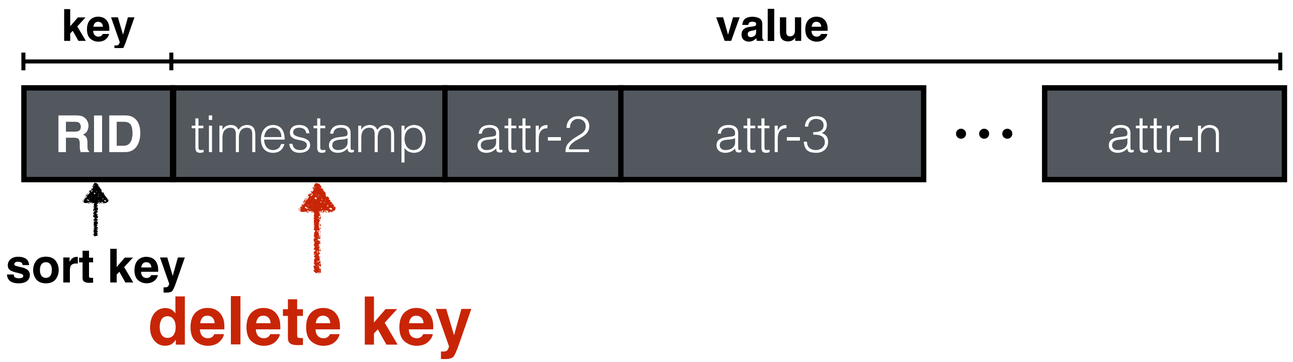

sort key ≠ delete key

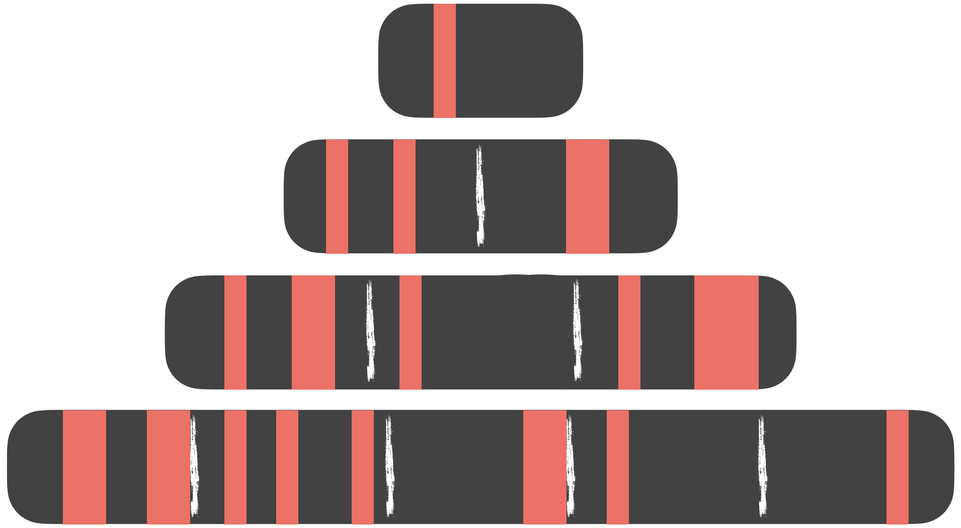

scattered occurrences

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

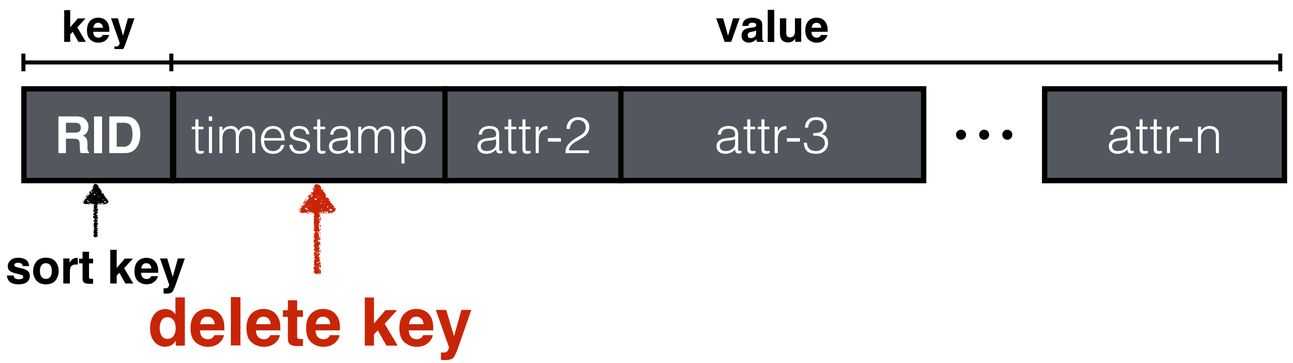

sort key ≠ delete key

latency spikes

superfluous I/Os

Realizing **Retention-Based Deletes**

delete all entries older than: TS_x

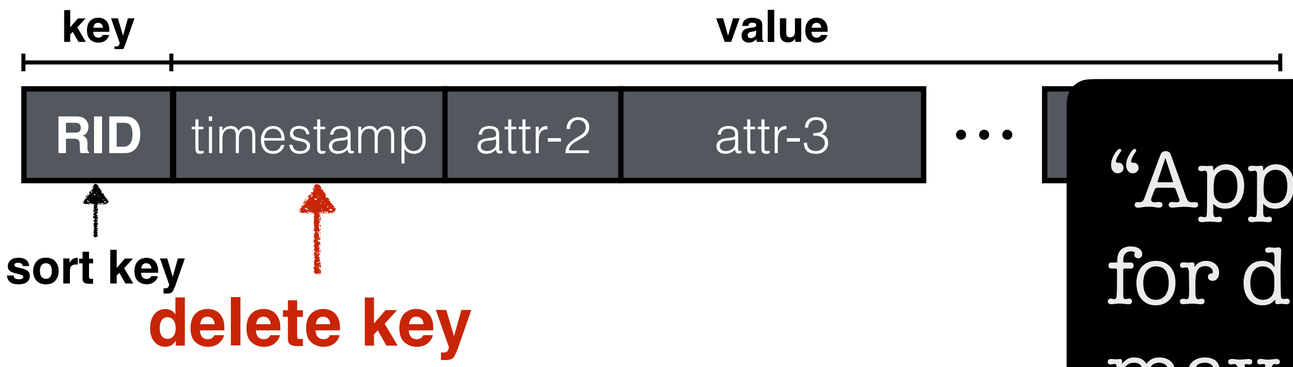

sort key ≠ delete ke

"Applications have requirements for deletes every day. E.g., they may keep data for 30 days, … effectively purging 1/30 of the database every day.

This induces performance pains!"

Realizing **Retention-Based Deletes**

delete all entries older than <= 65_D

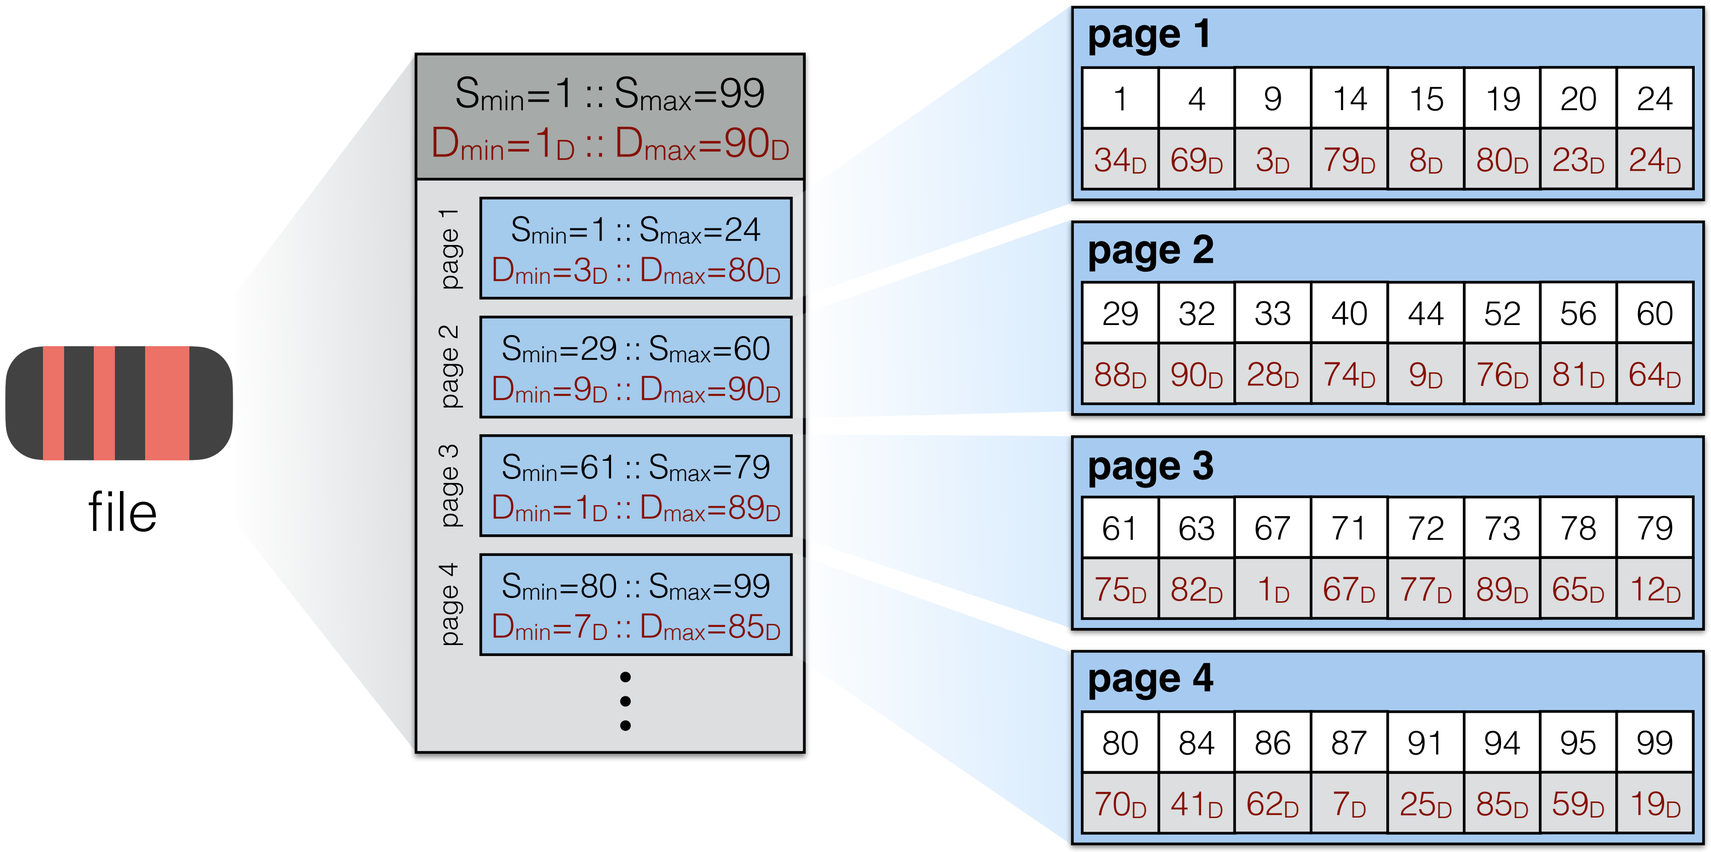

Realizing **Retention-Based Deletes**

delete all entries older than <= 65_D

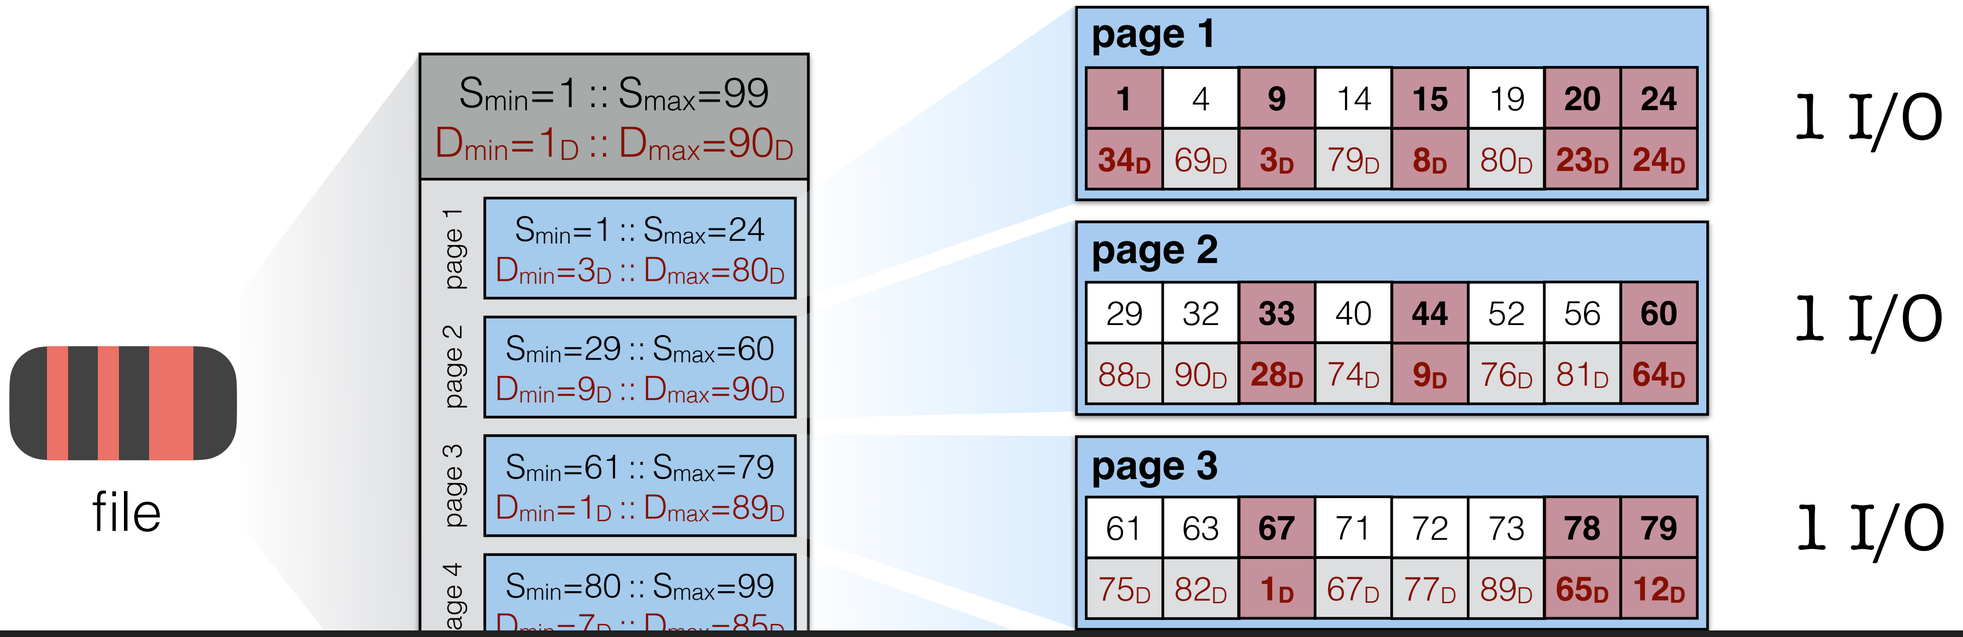

Intuition: Data Layout holds the key!

Realizing **Retention-Based Deletes**

KiWi

Key Weaving storage layout

Key Weaving storage layout

delete all entries older than <= 65_D

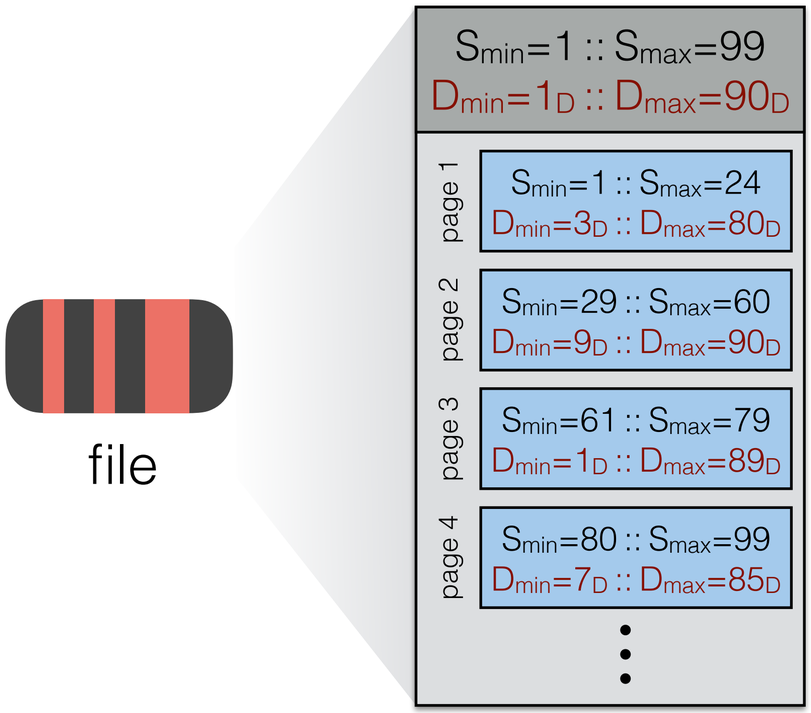

Key Weaving storage layout

delete all entries older than <= 65_D

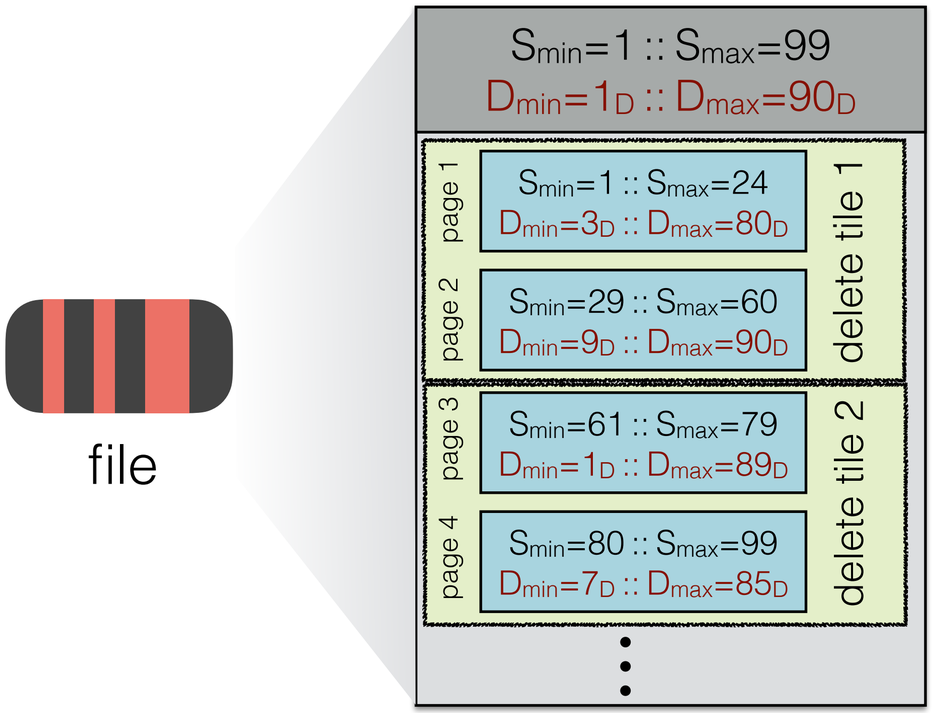

Key Weaving storage layout

delete all entries older than <= 65_D

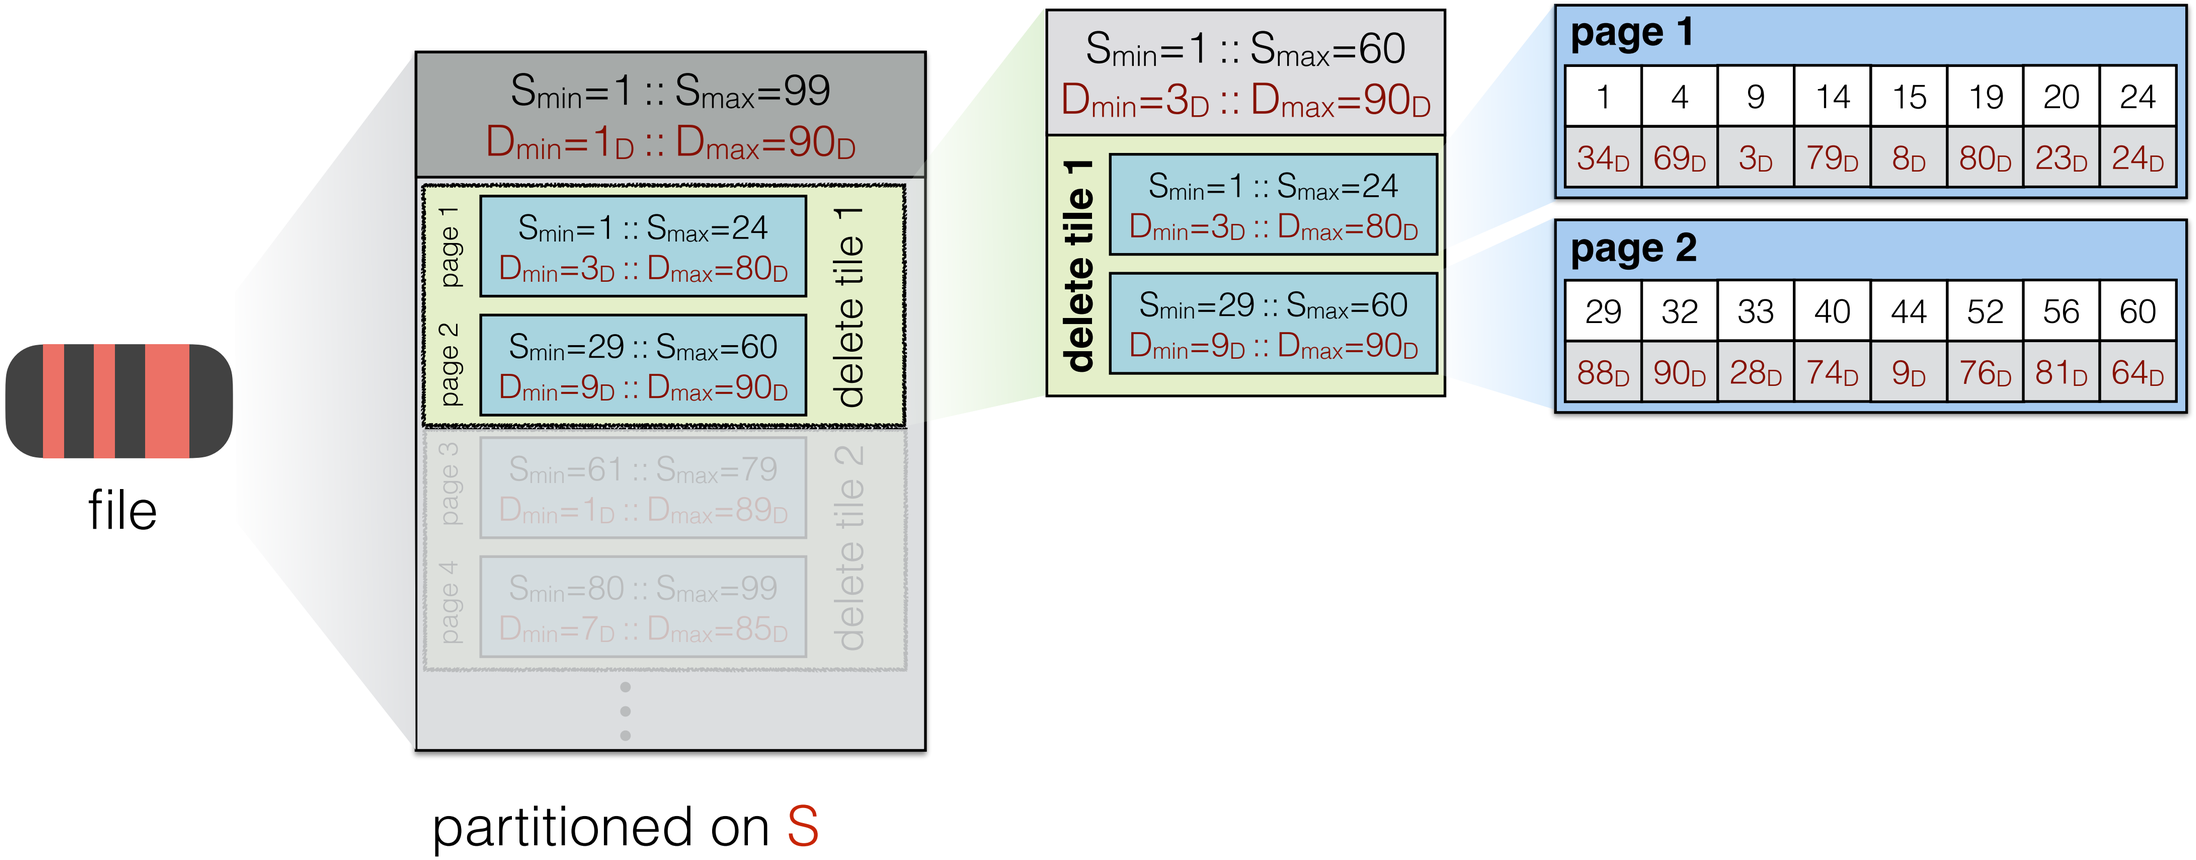

Key Weaving storage layout

delete all entries older than <= 65_D

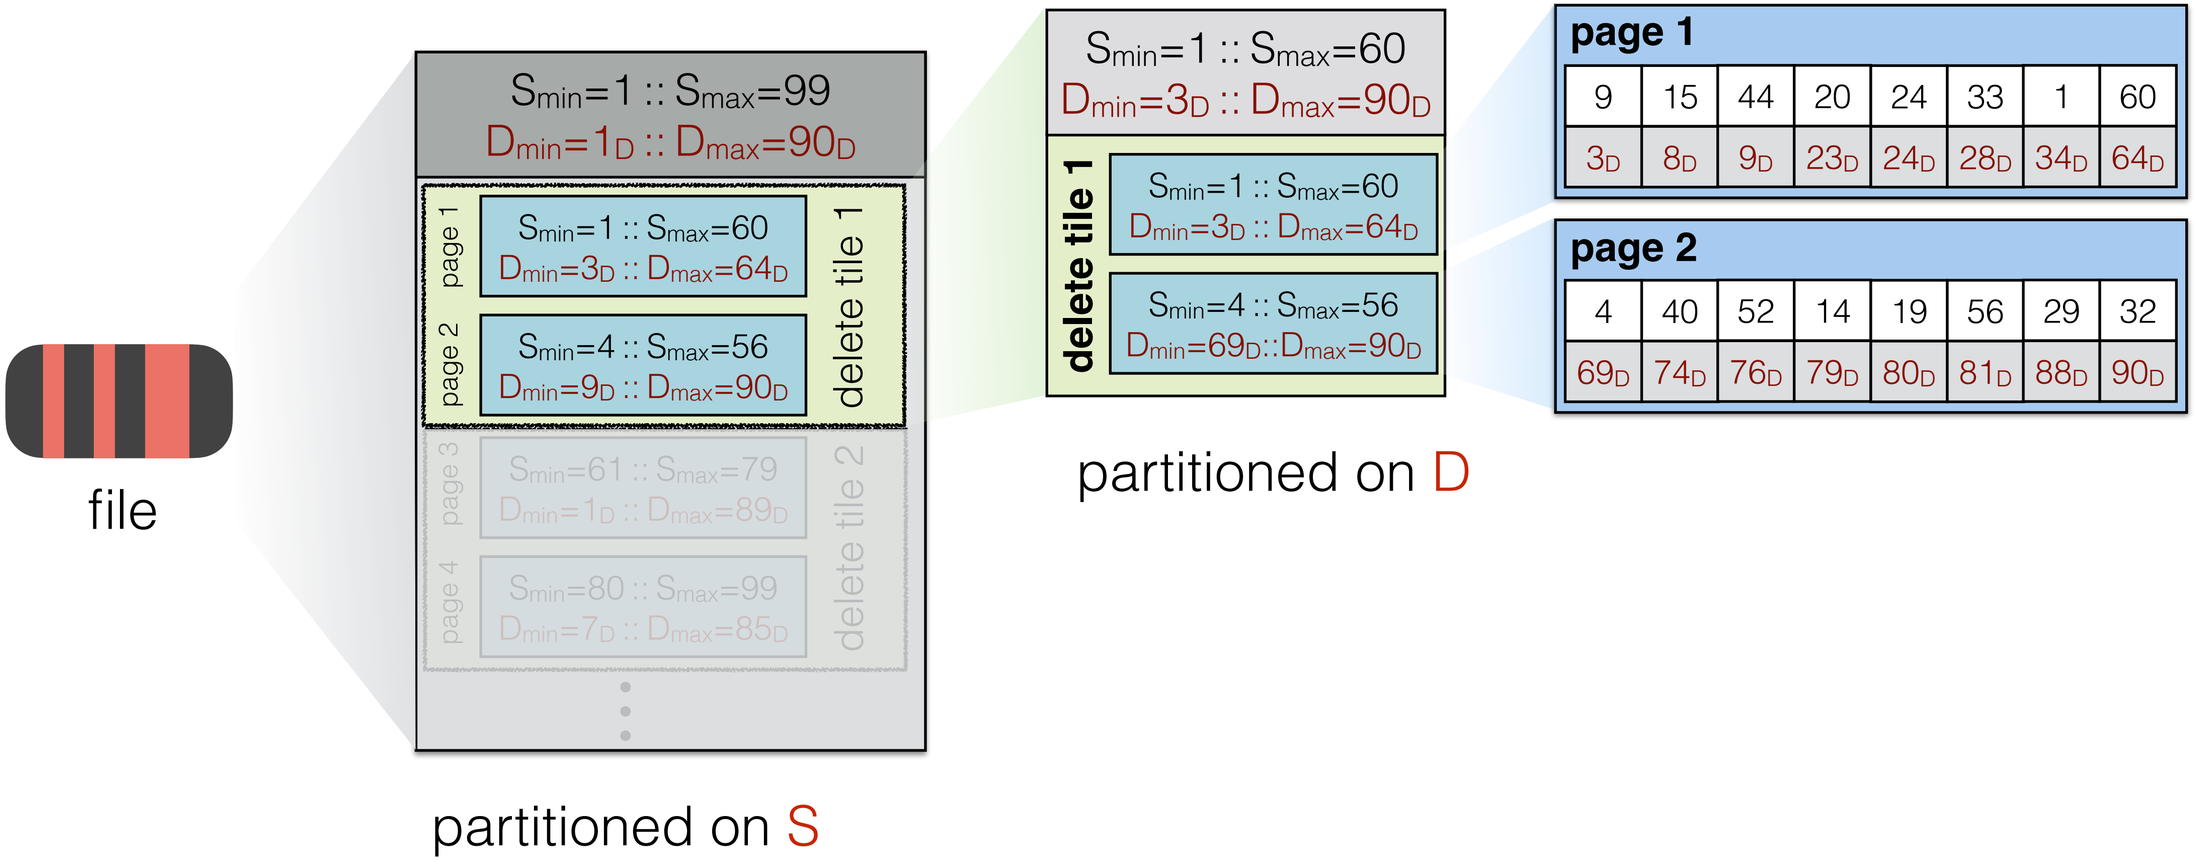

Key Weaving storage layout

delete all entries older than <= 65_D

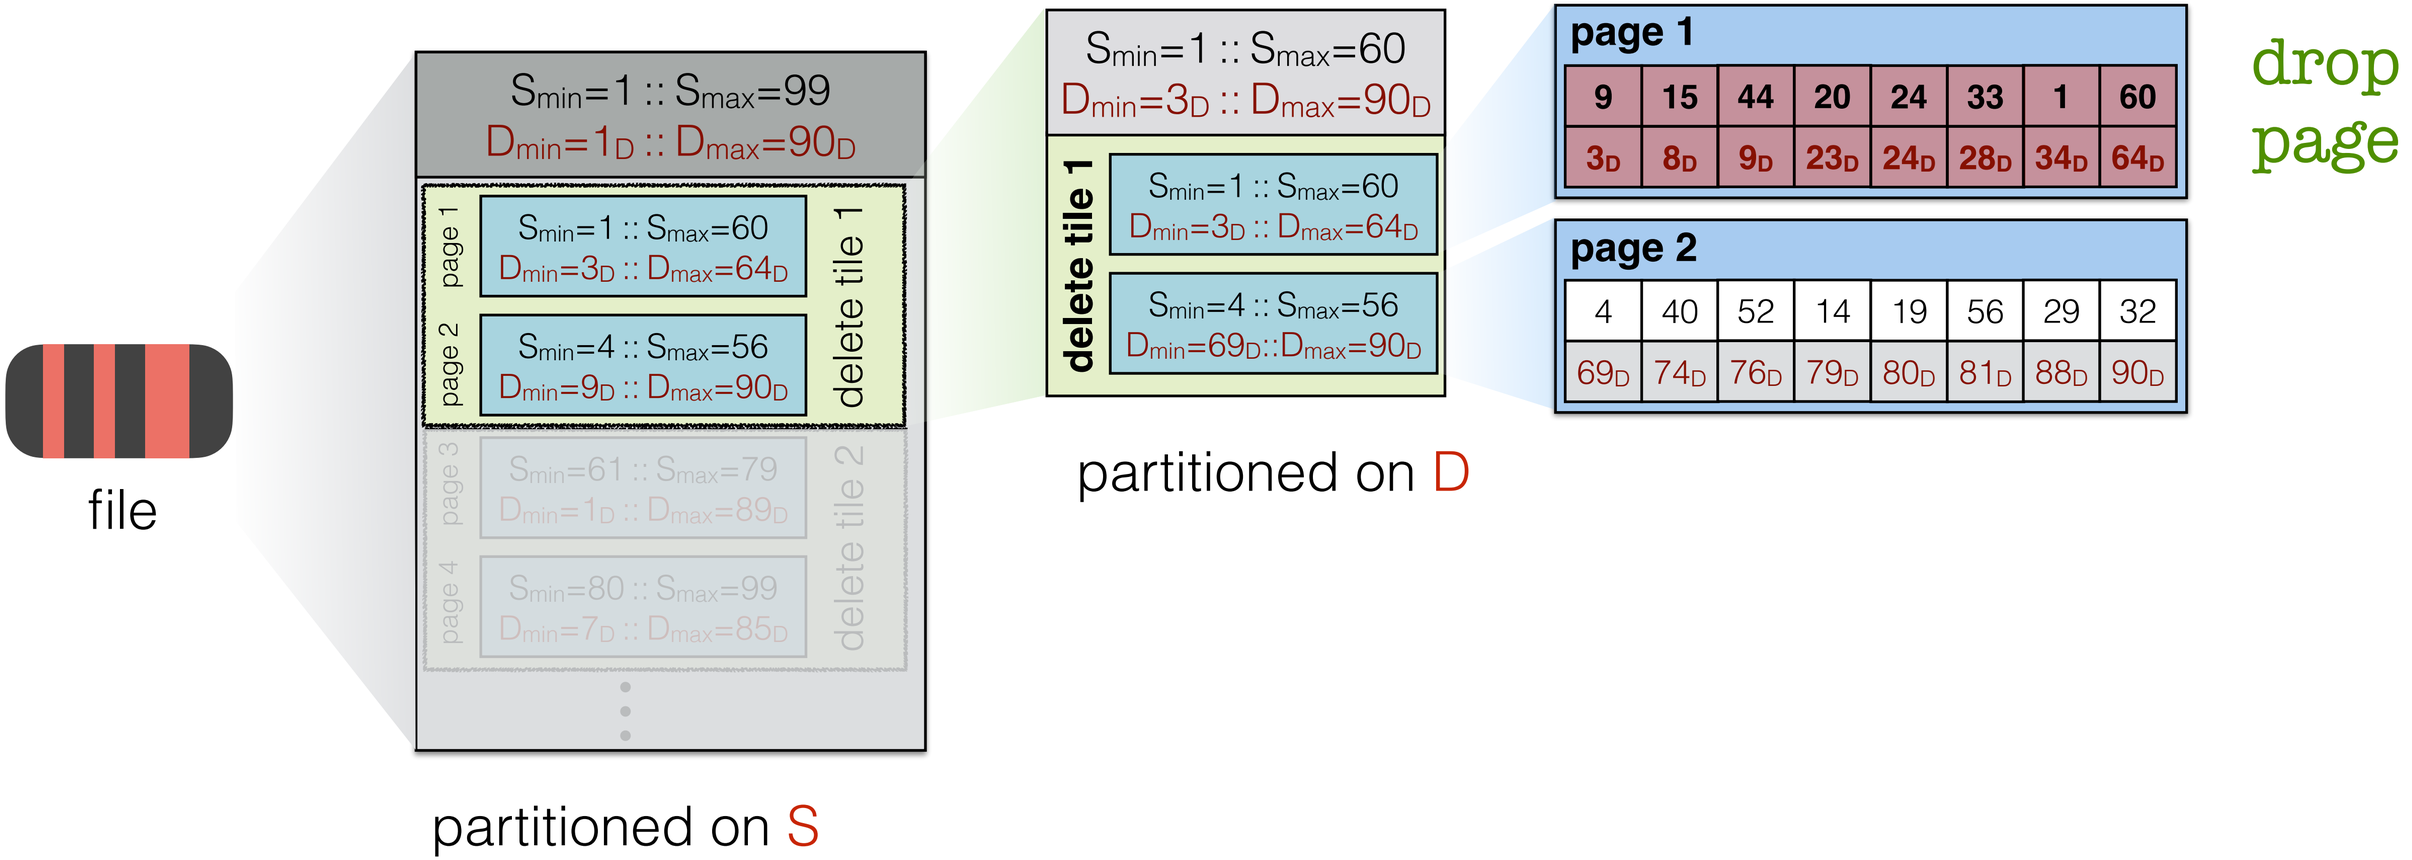

Key Weaving storage layout

delete all entries older than <= 65_D

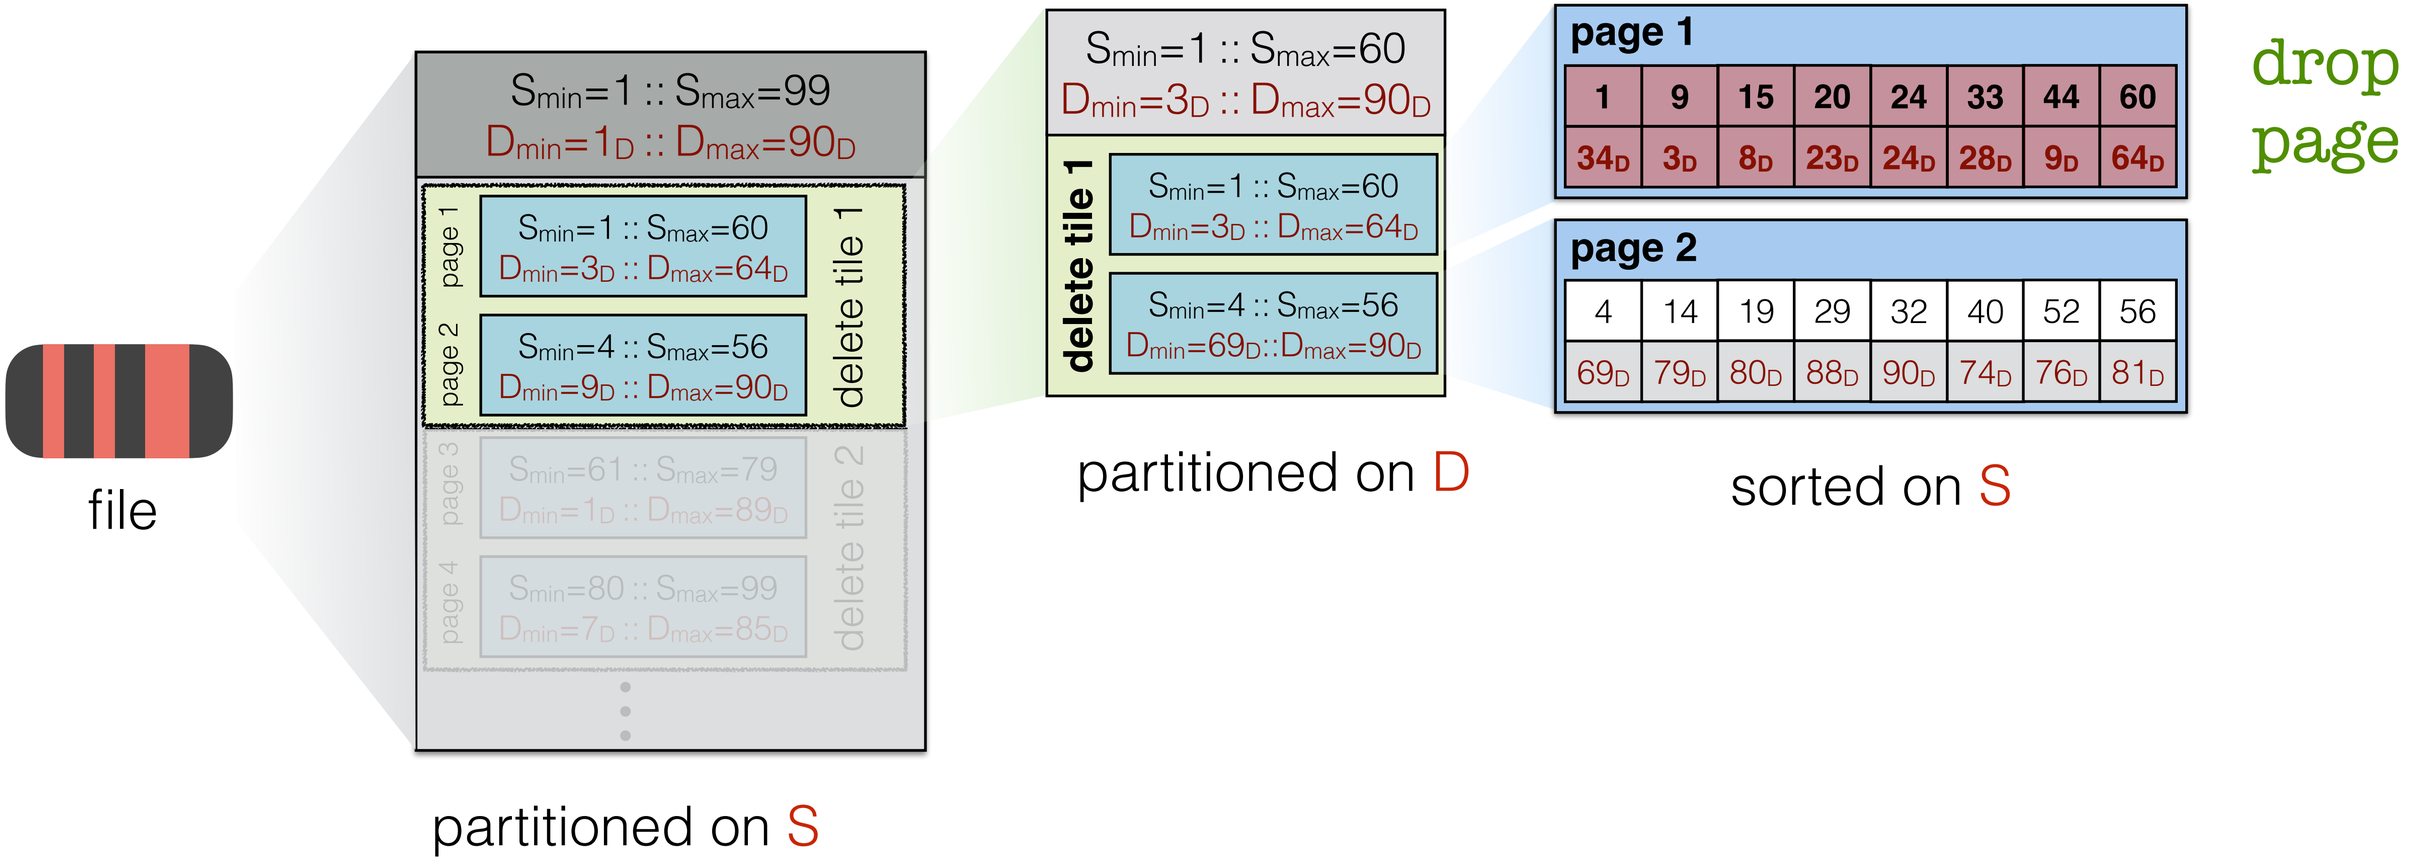

Key Weaving storage layout

delete all entries older than <= 65_D

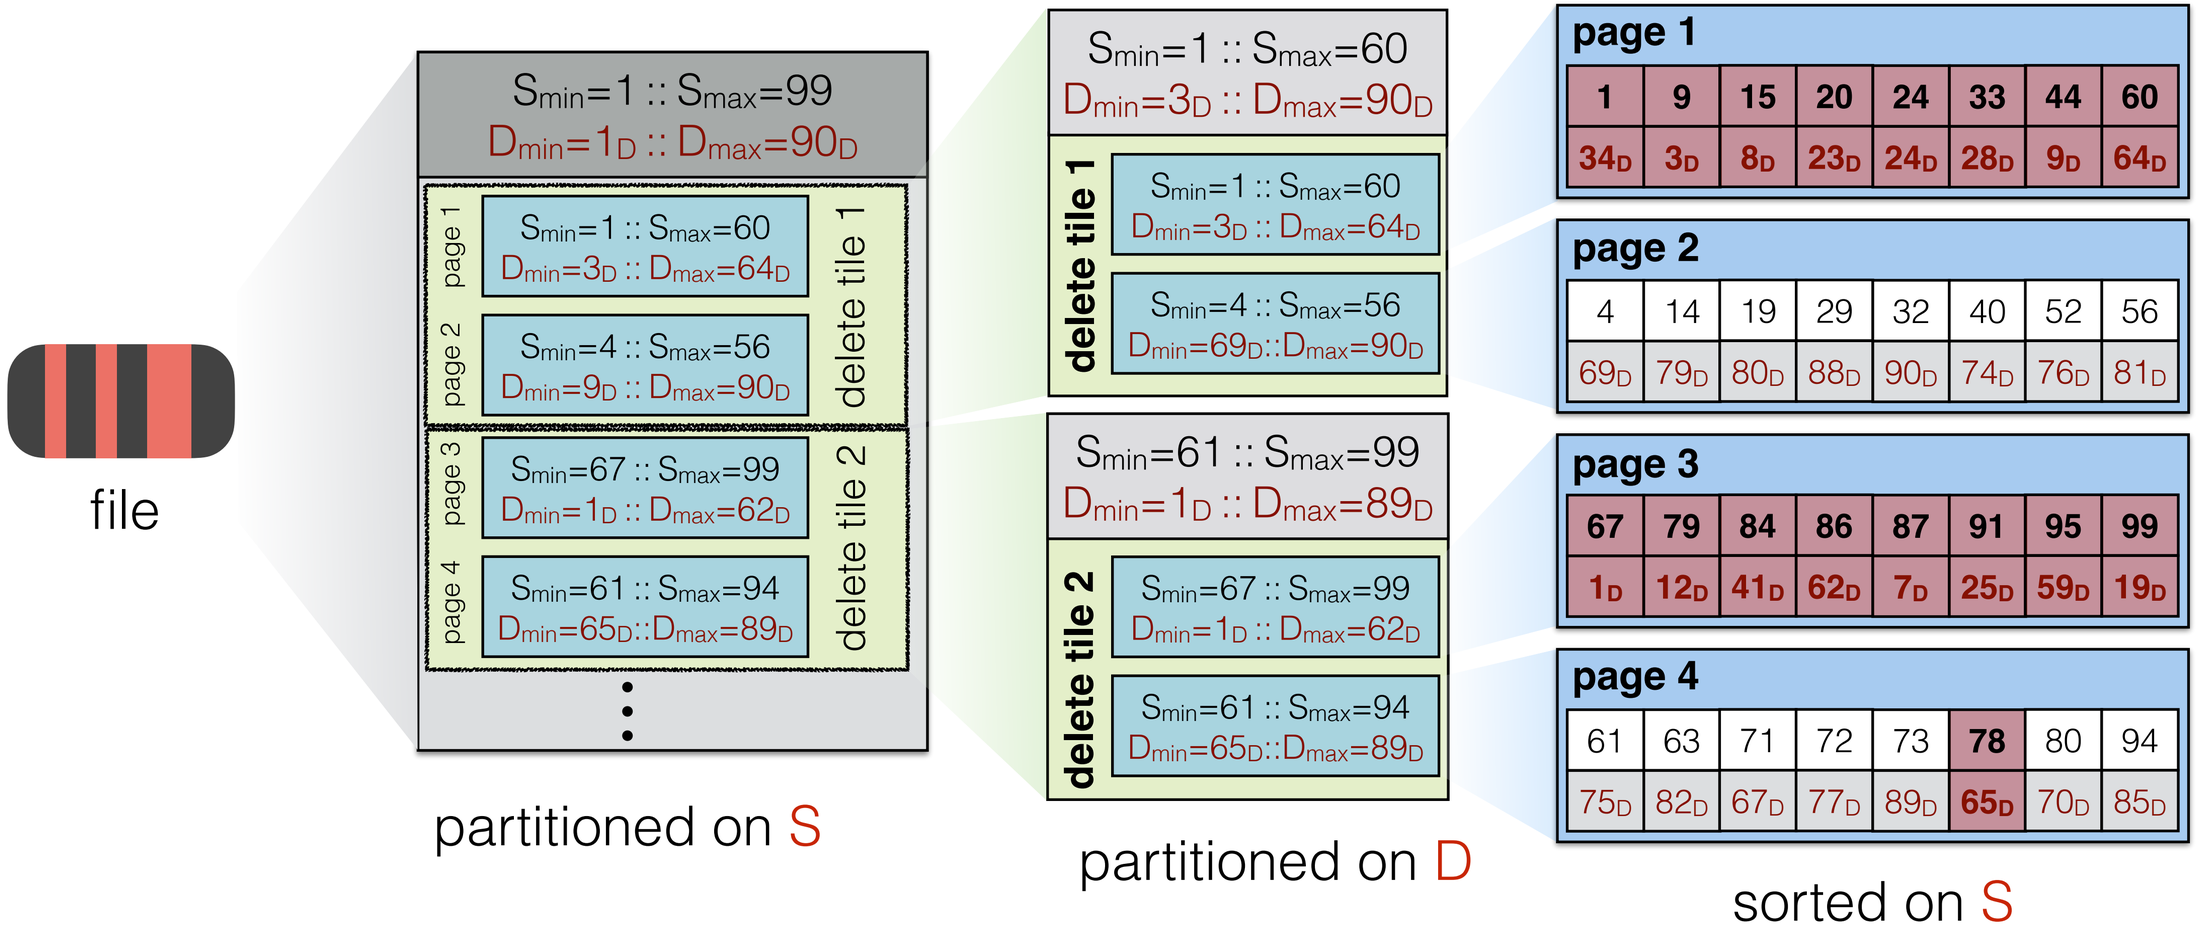

superior delete performance

up to 2.5x

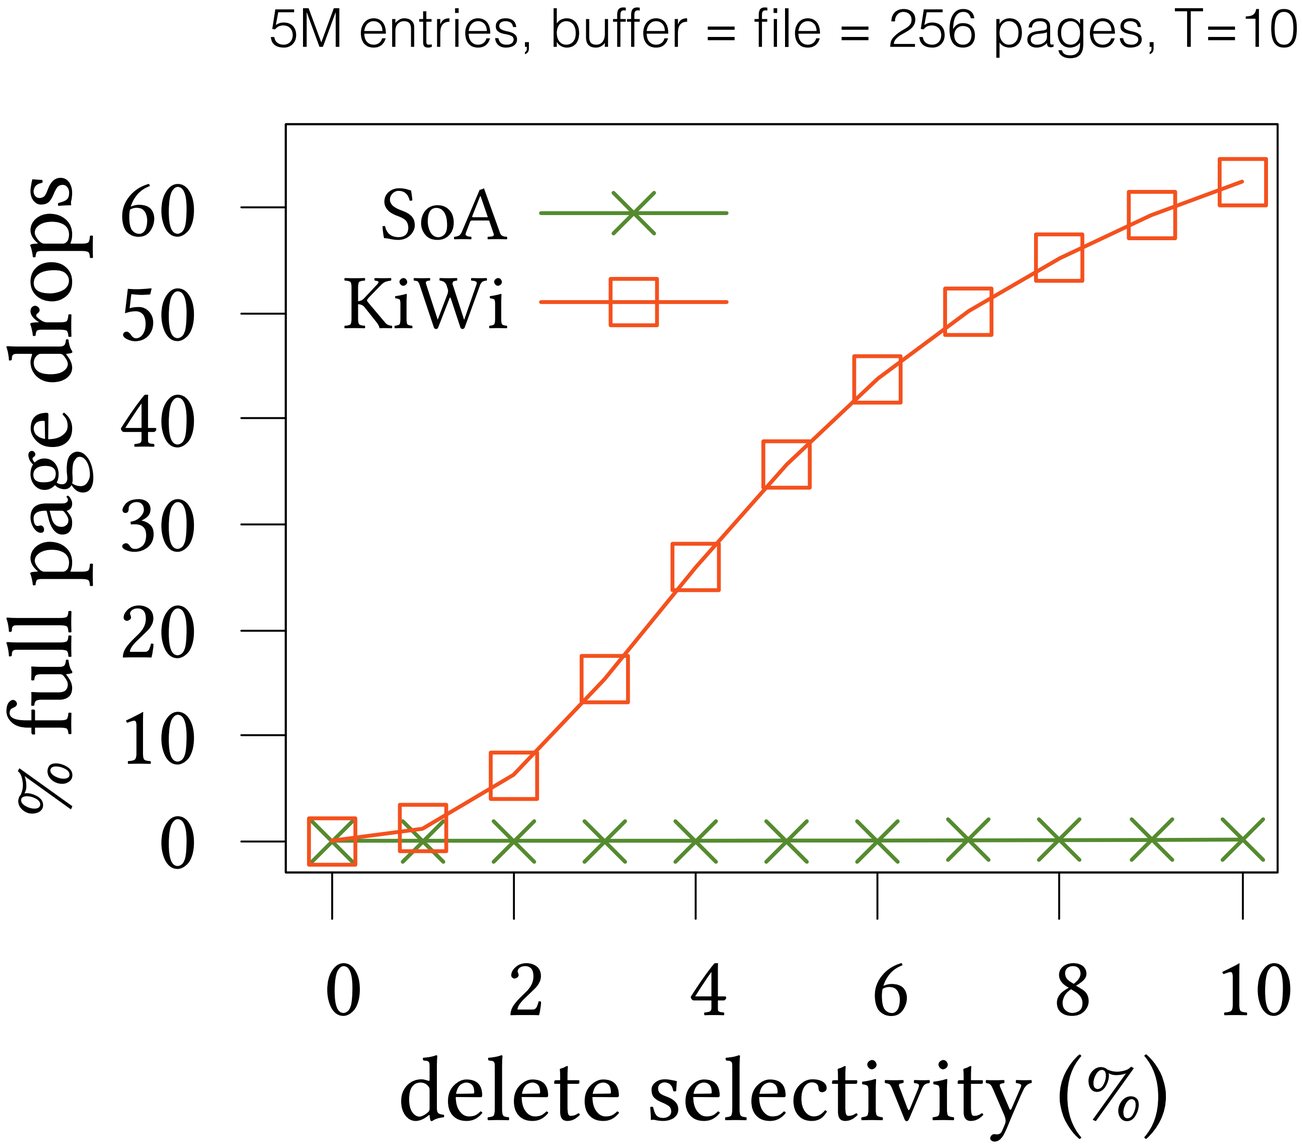

Key Weaving storage layout

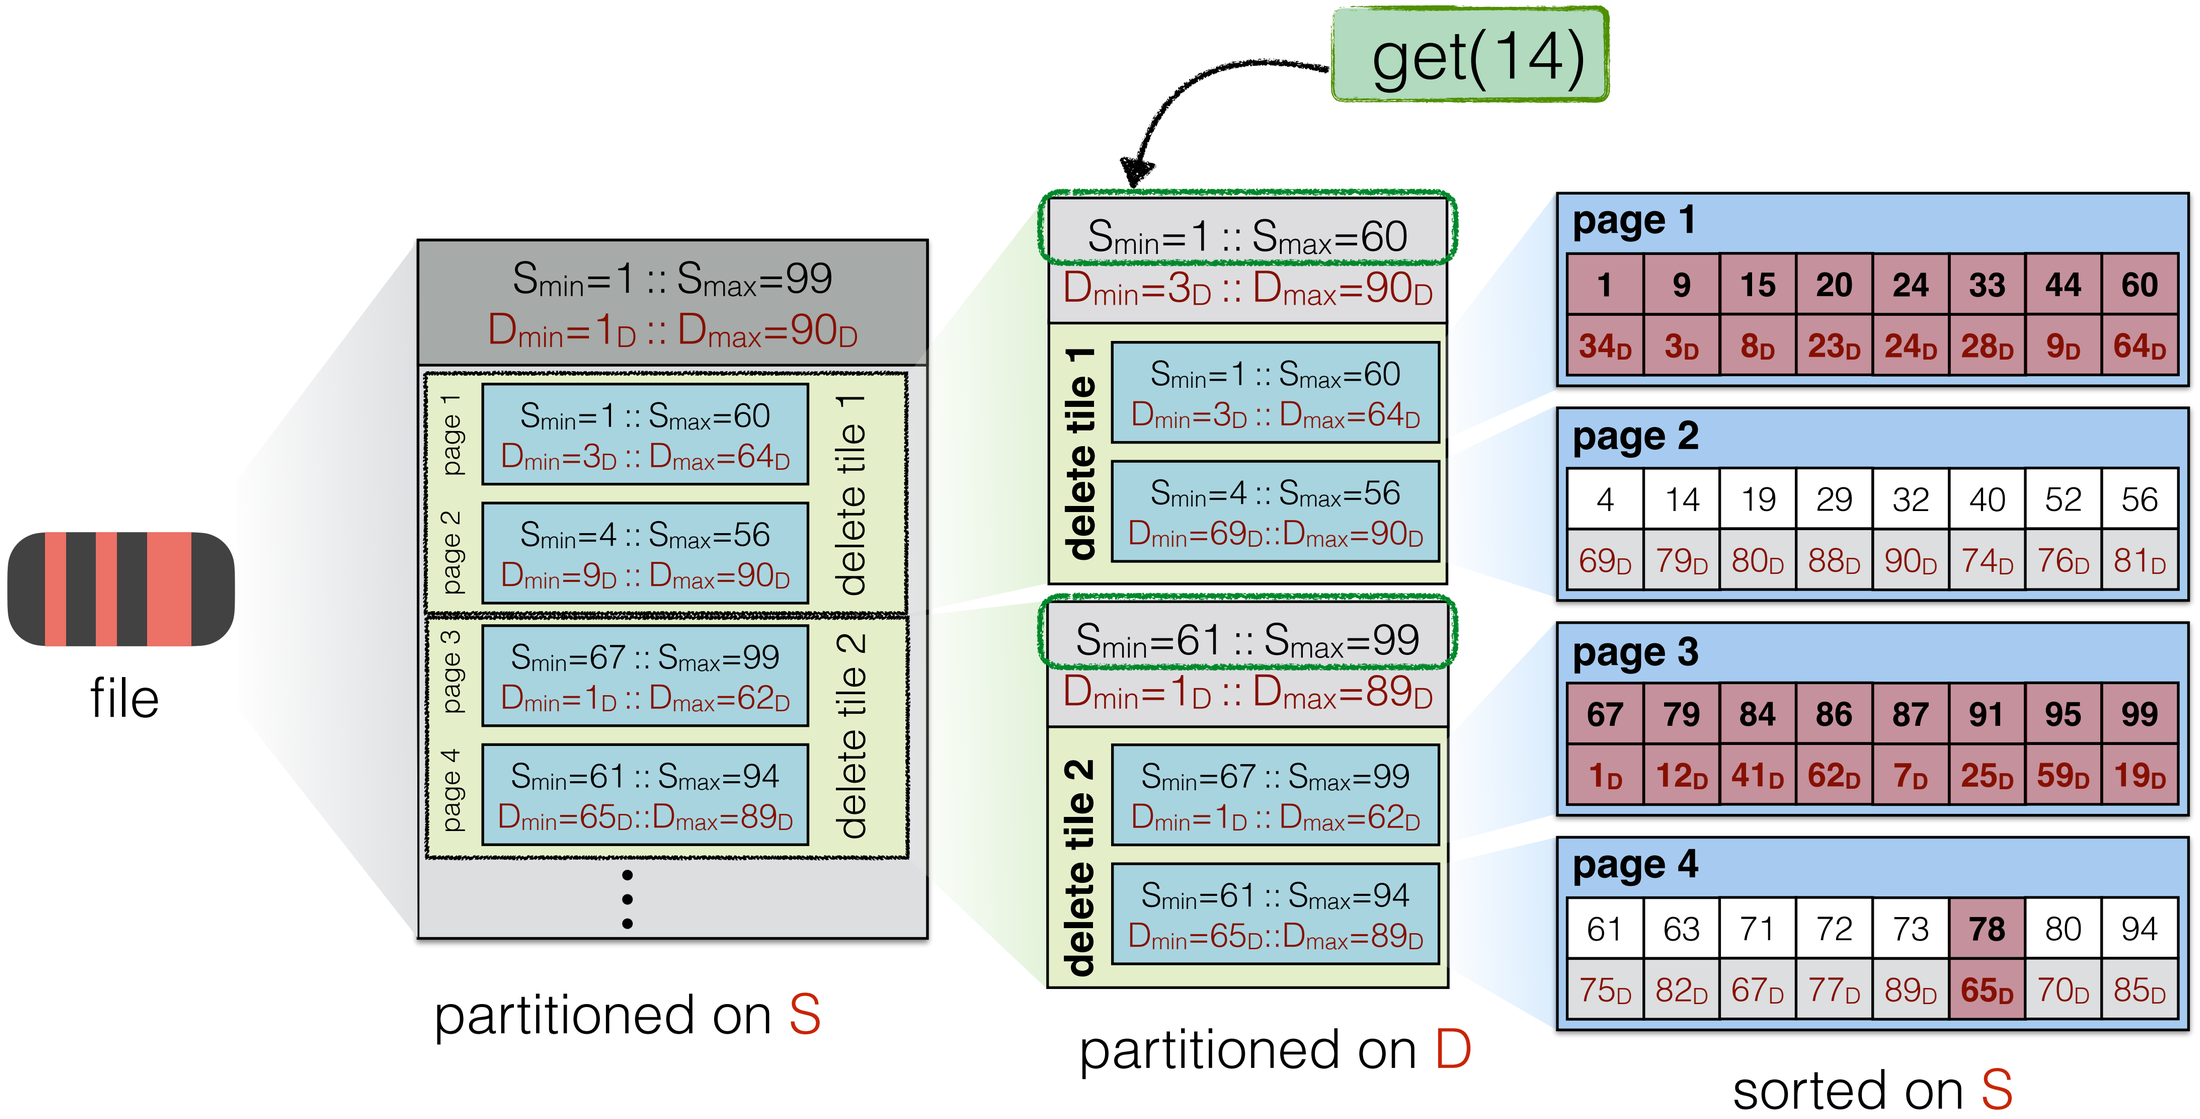

Key Weaving storage layout

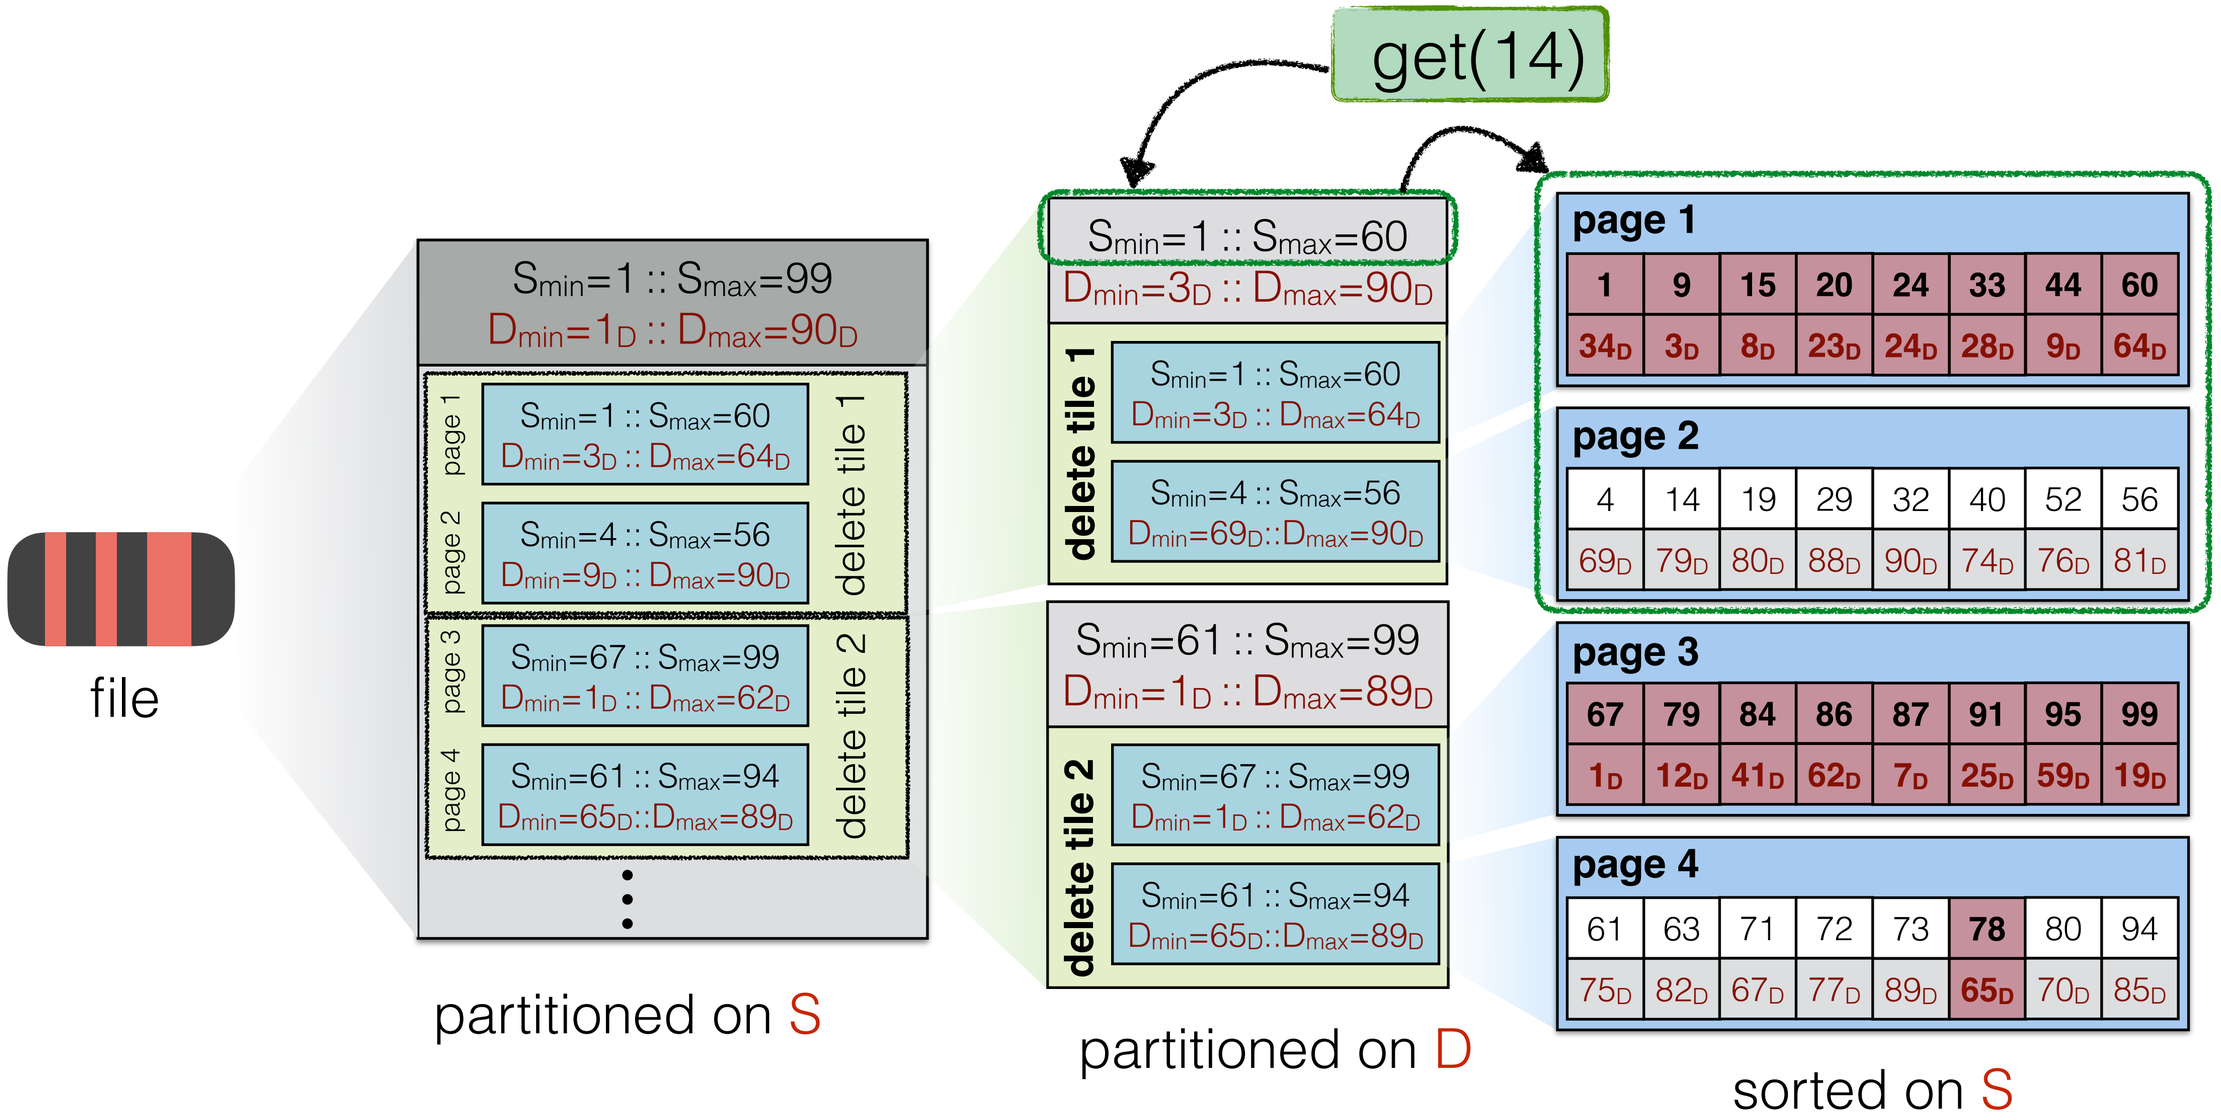

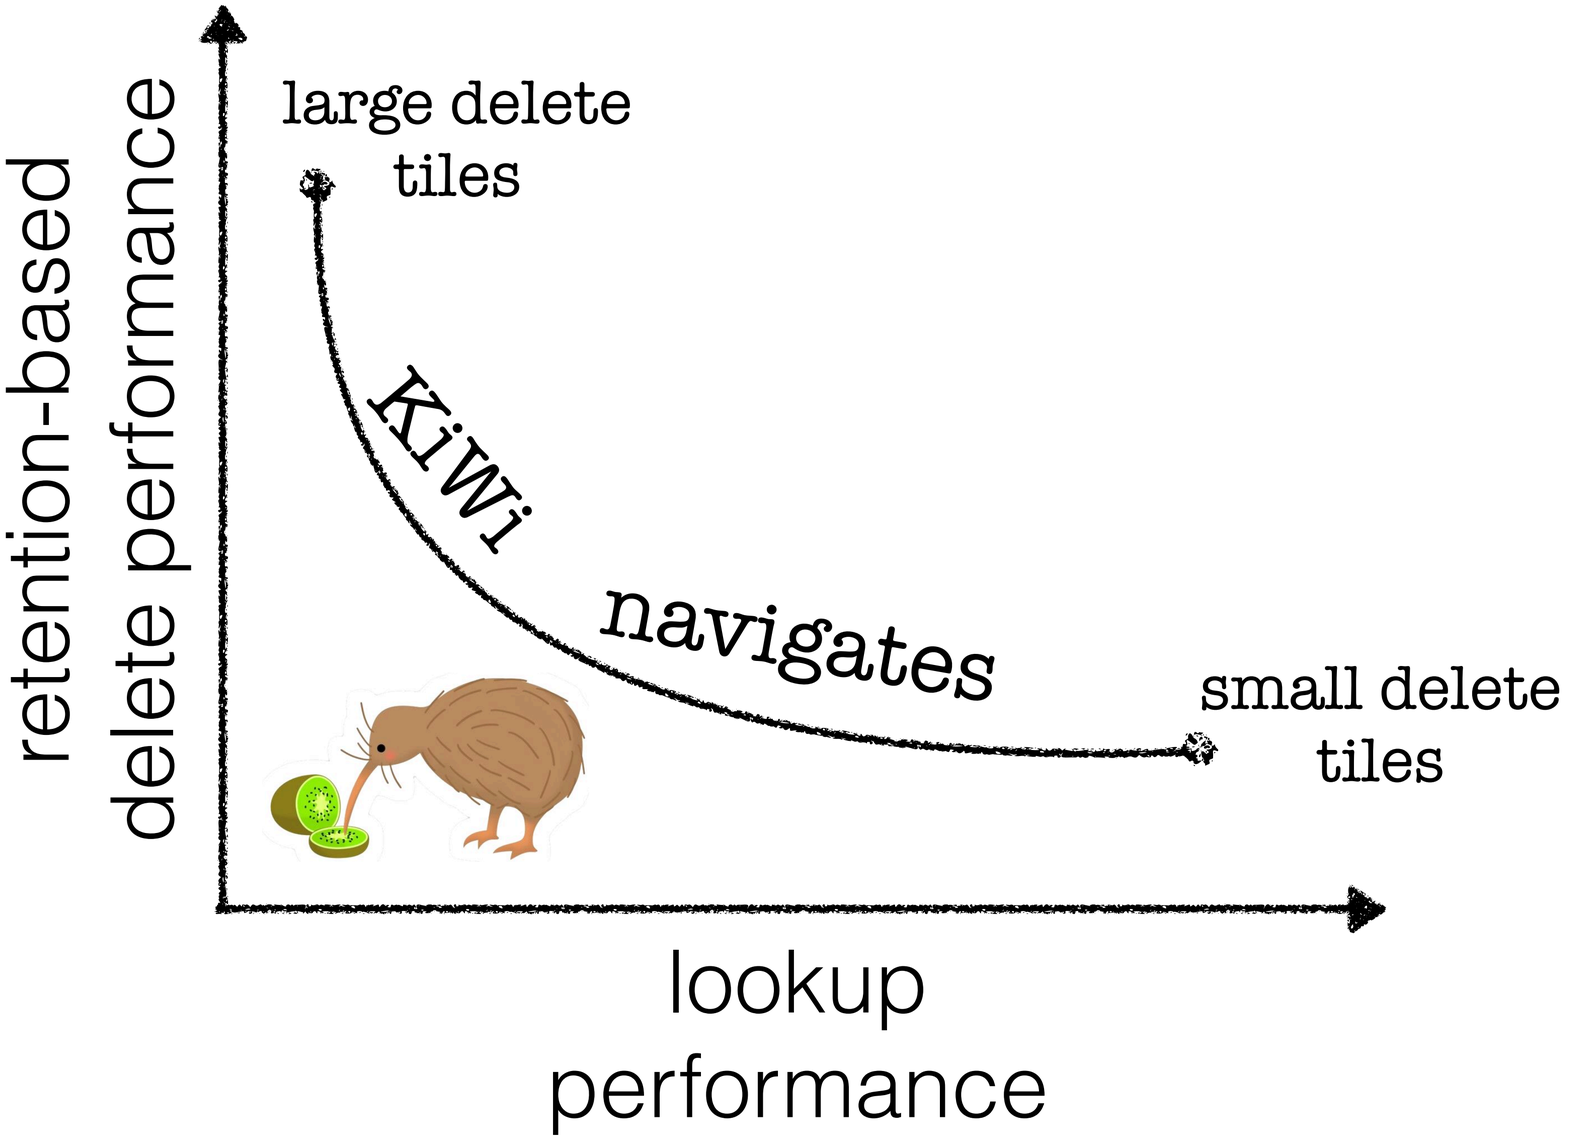
retention-based
delete performance
large delete
tiles
KiWi
navigates
small delete
tiles
lookup
performance

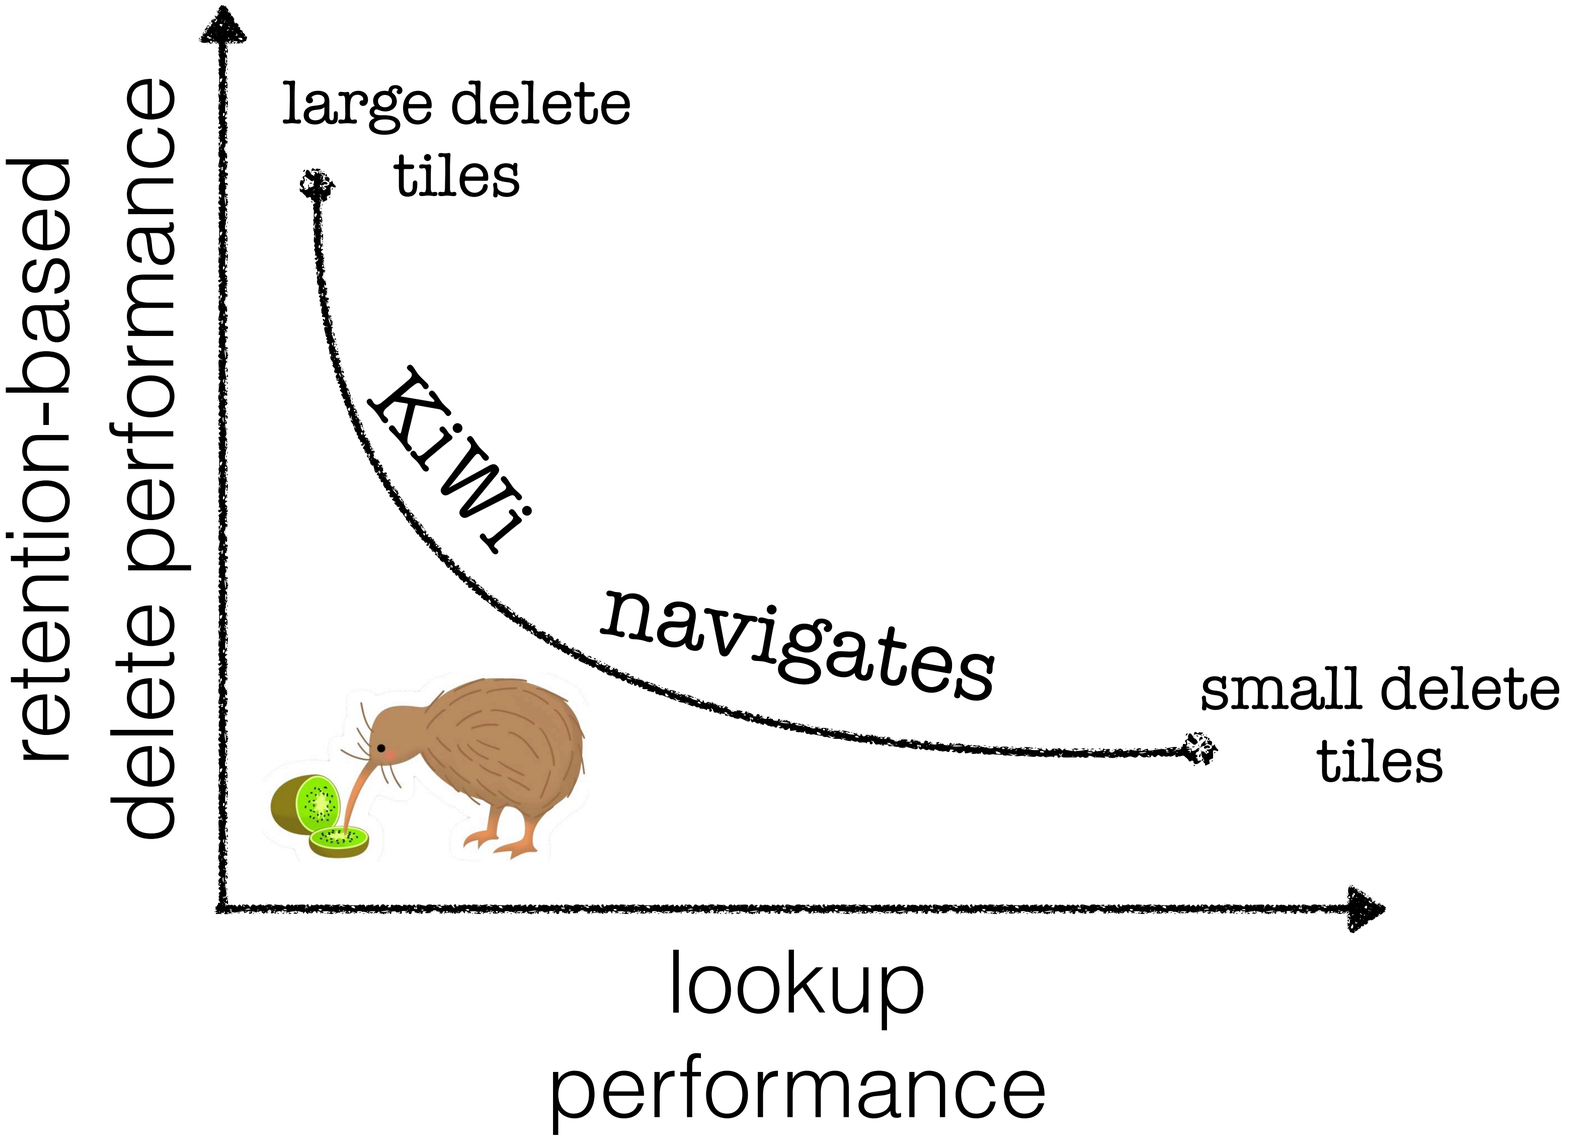
retention-based
delete performance
large delete
tiles
KiWi
navigates
small delete
tiles
lookup
performance

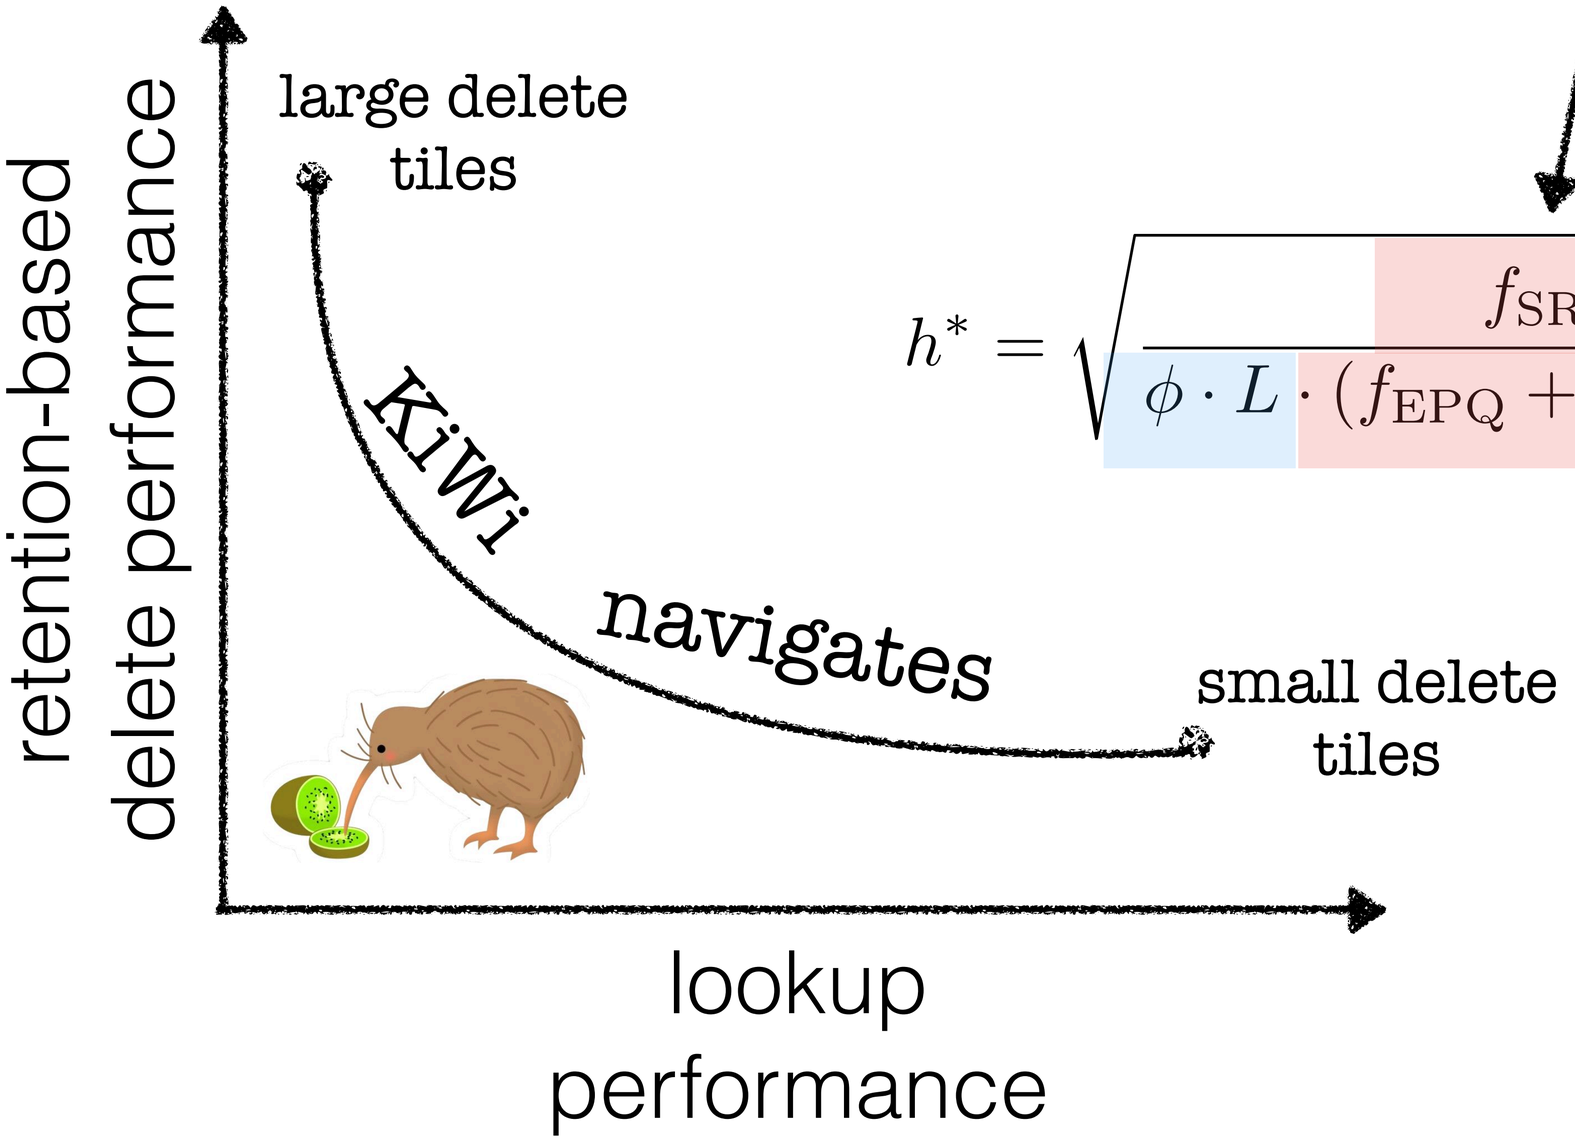

workload

data
structure

$$h^* = \sqrt{\frac{f_{\text{SRD}} \cdot \frac{N}{B}}{\phi \cdot L \cdot (f_{\text{EPQ}} + f_{\text{PQ}}) \cdot + L \cdot f_{\text{SRQ}}}}$$

h^* = optimal delete tile size

f_{SRD} = prop. of retention-based deletes

f_{EPQ} = prop. of empty point queries

f_{PQ} = prop. of non-empty point queries

f_{SRQ} = prop. of short range queries

L = levels in tree

N = entries in tree

B = entries in a page

ϕ = false positive rate of query filter

better overall performance

up to 4x

superior delete performance

up to 2.5x

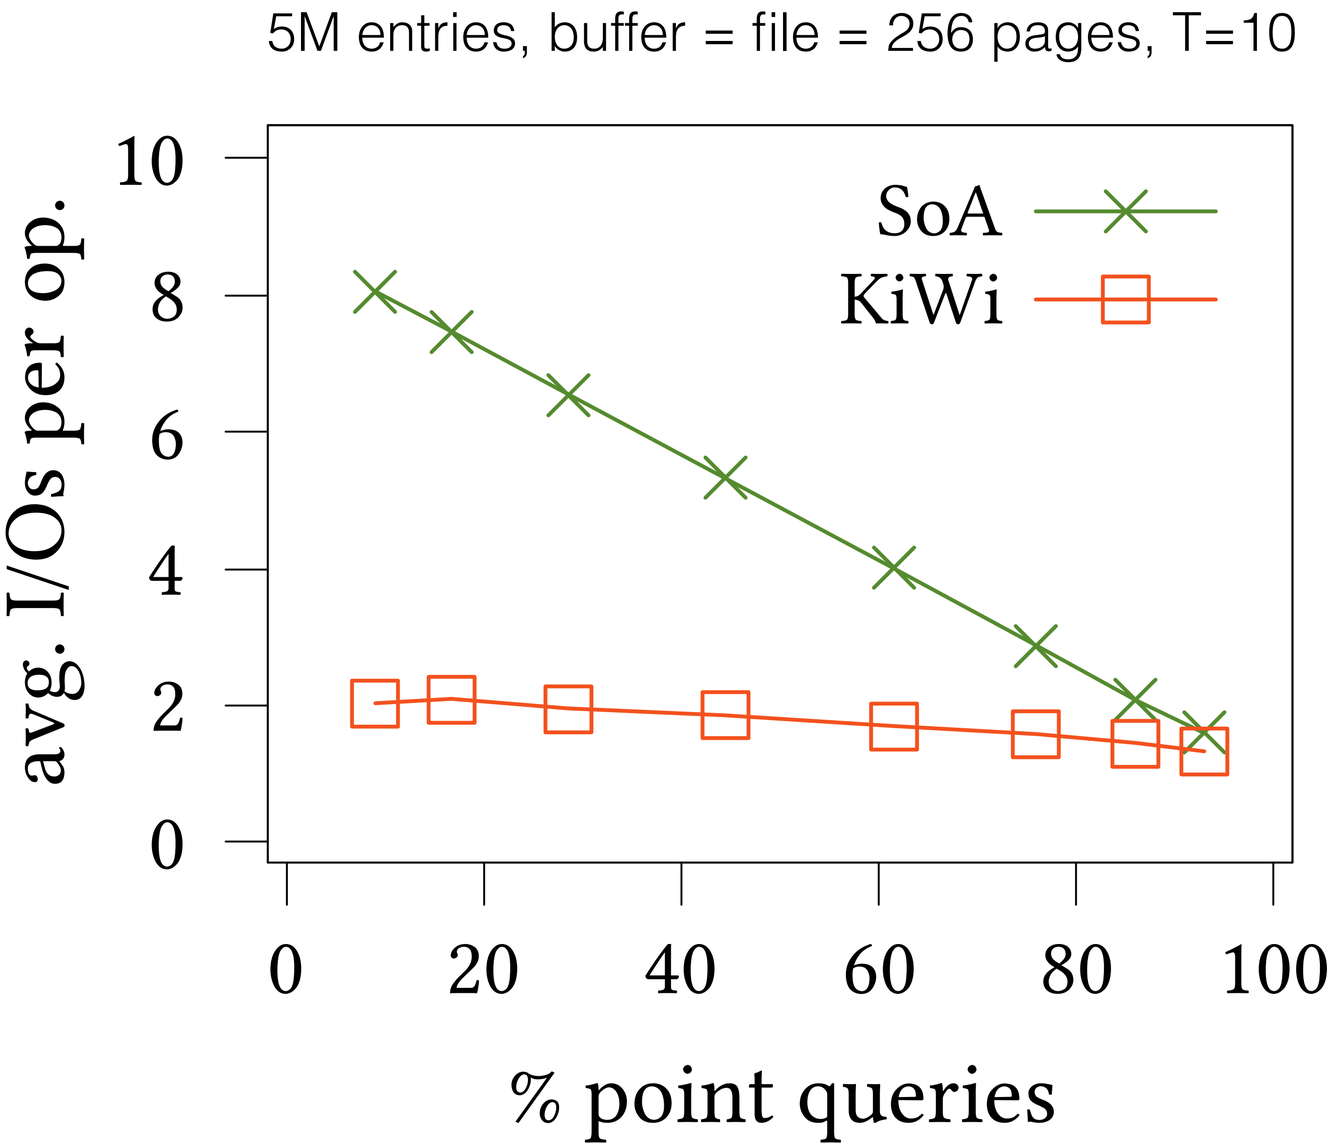

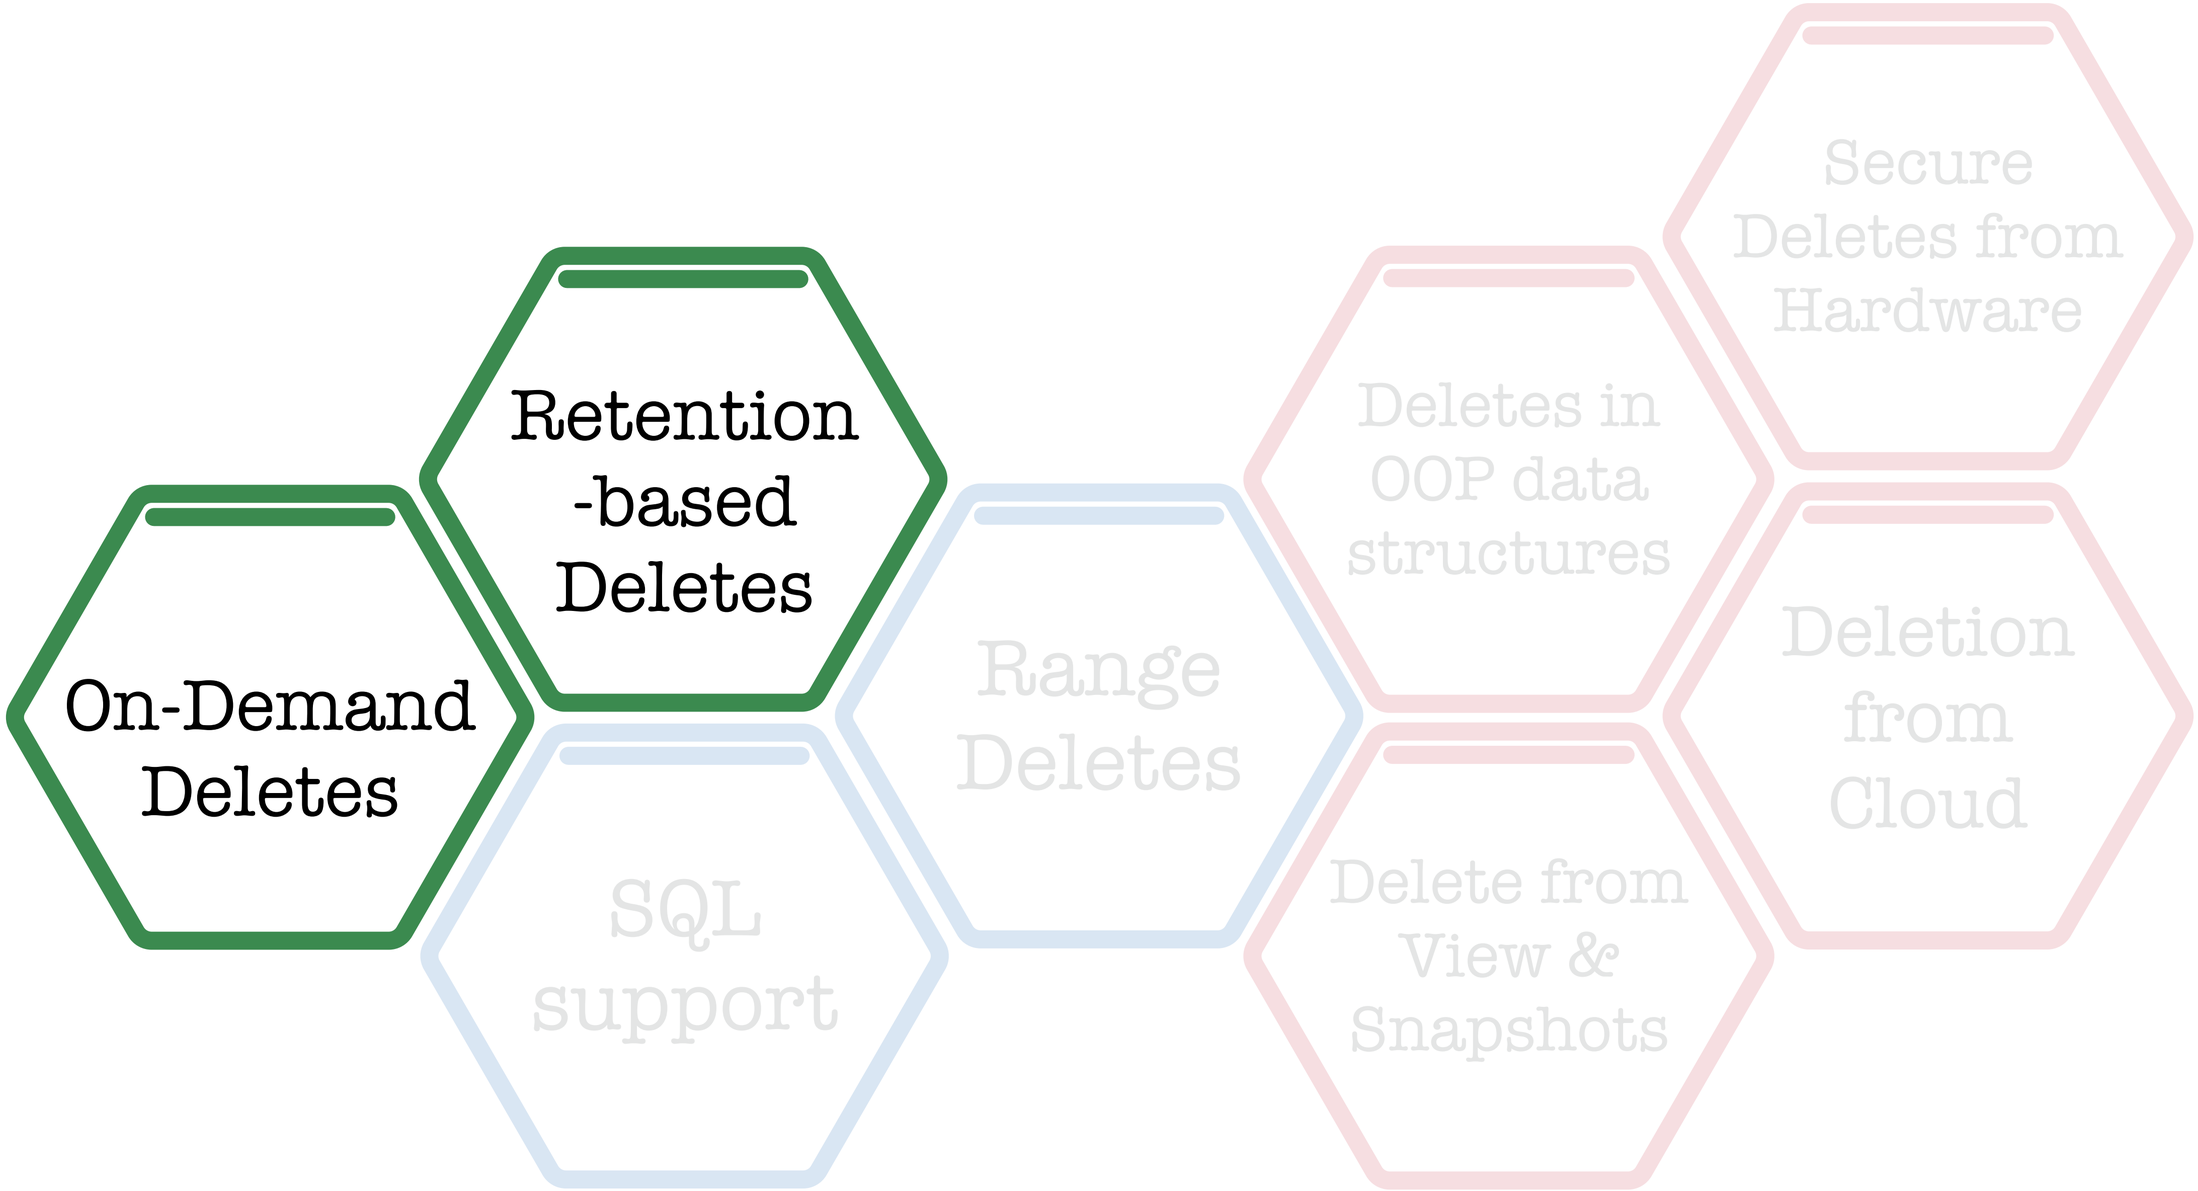
On-Demand Deletes
Retention -based Deletes
SQL support
Range Deletes
Delete from View & Snapshots
Deletes in OOP data structures
Deletion from Cloud
Secure Deletes from Hardware

FADE

Range Deletes

SQL support

Deletes in OOP data structures

Secure Deletes from Hardware

Deletion from Cloud

Delete from View & Snapshots

FADE

SQL support

Range Deletes

Deletes in OOP data structures

Secure Deletes from Hardware

Deletion from Cloud

Delete from View & Snapshots

FADE

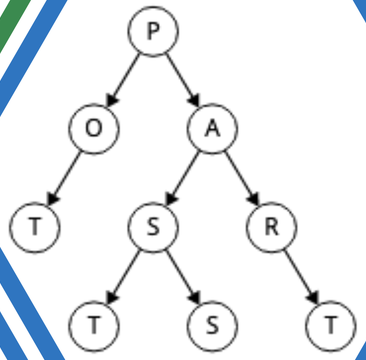

Deletes in OOP data structures

Secure Deletes from Hardware

Deletion from Cloud

Delete from View & Snapshots

THANK YOU!

Actively seeking students and collaborators!

Subhadeep Sarkar

Brandeis UNIVERSITY

https://subhadeep.net/

lab