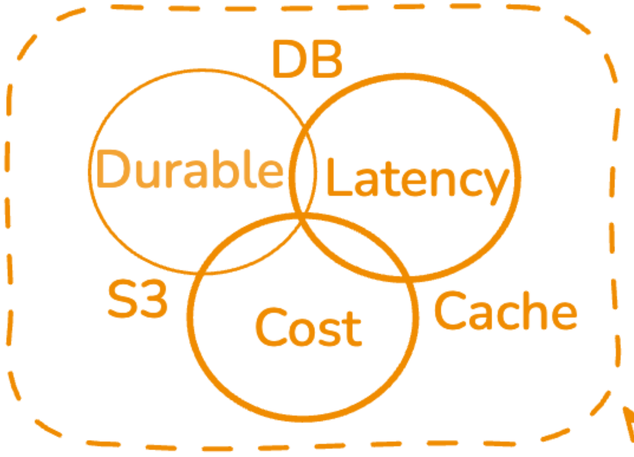

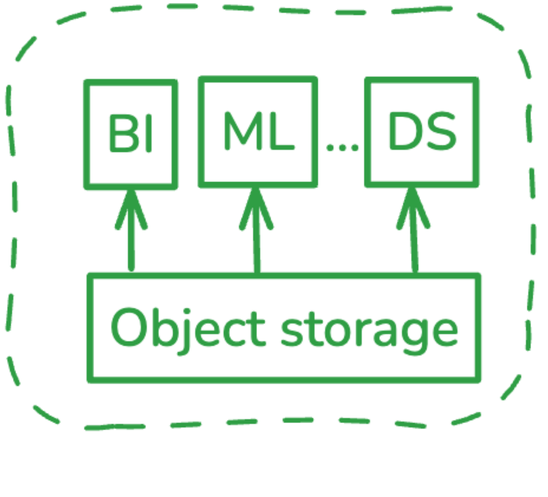

Data-Aware Caching for Cloud Analytics

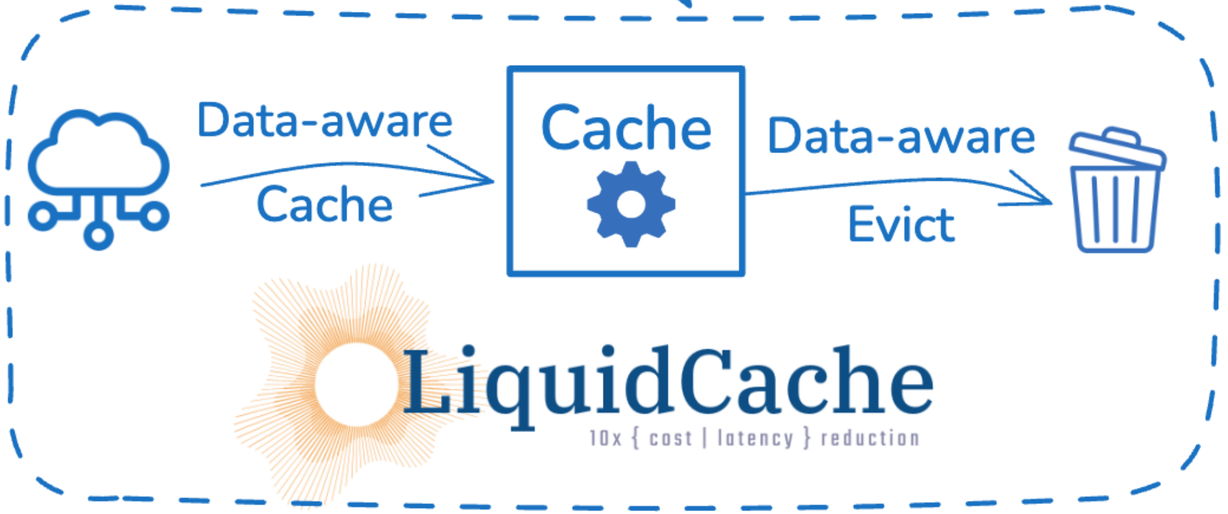

Xiangpeng Hao

Department of Computer Science
University of Wisconsin-Madison

May 19th, 1 PM CDT
CS2310 or Zoom

On premise → cloud (2010-2020)

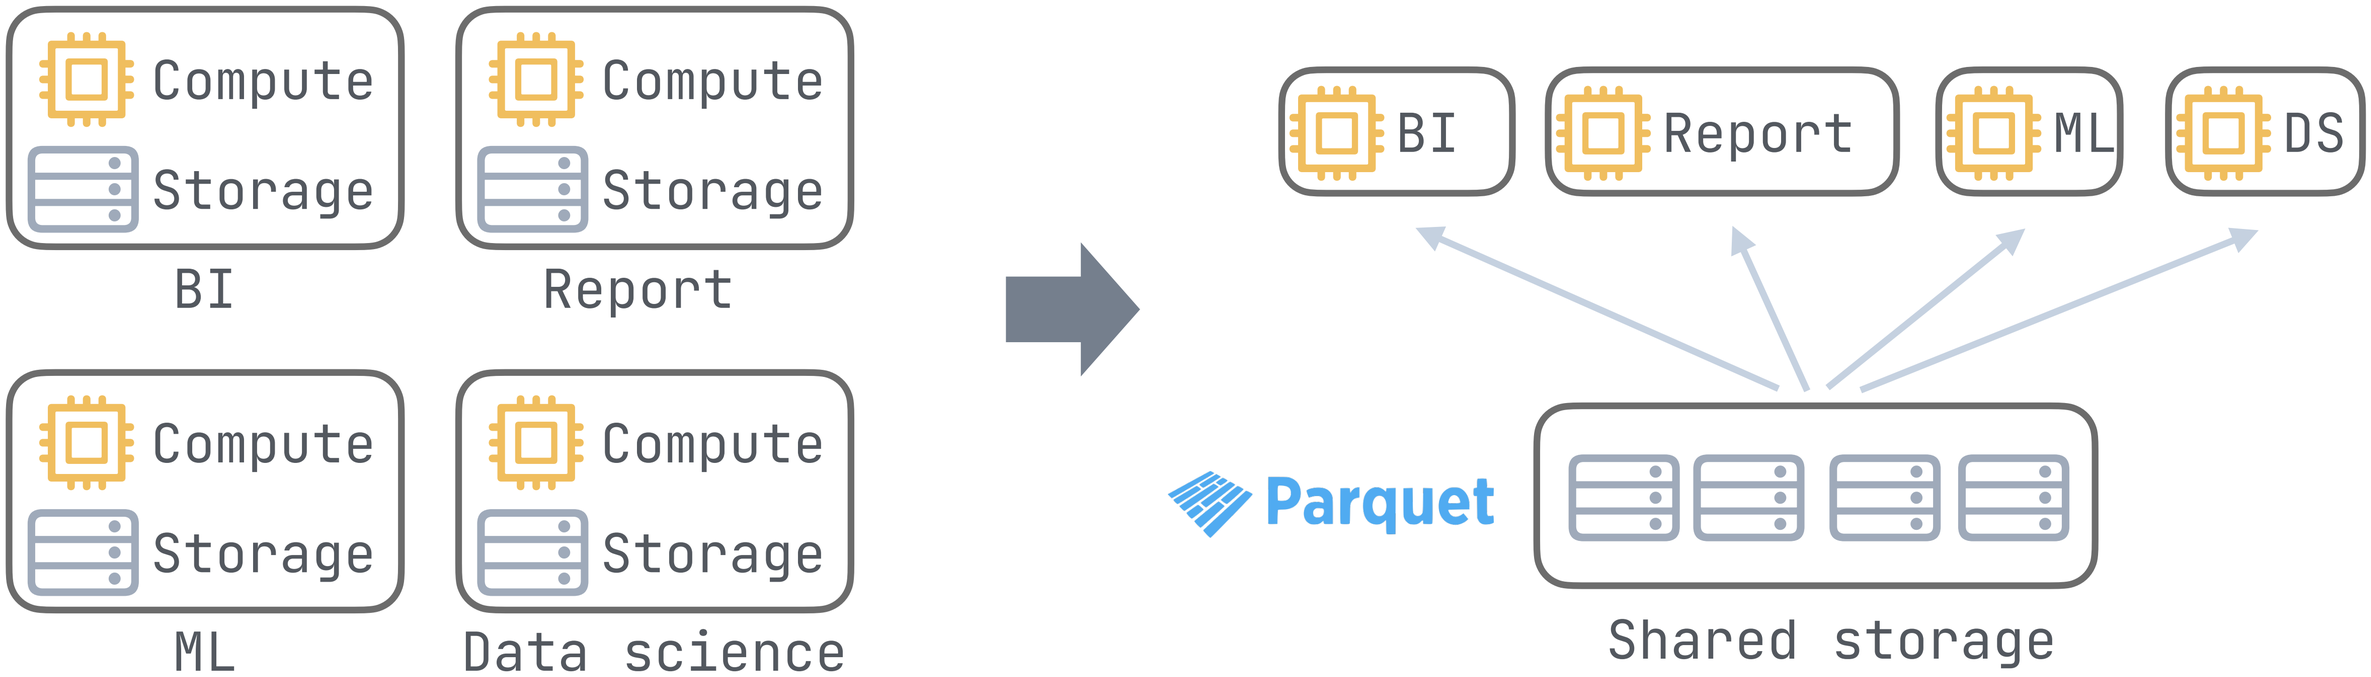

SAP databricks snowflake influxdata®

Every system has a cache (2020-2025)

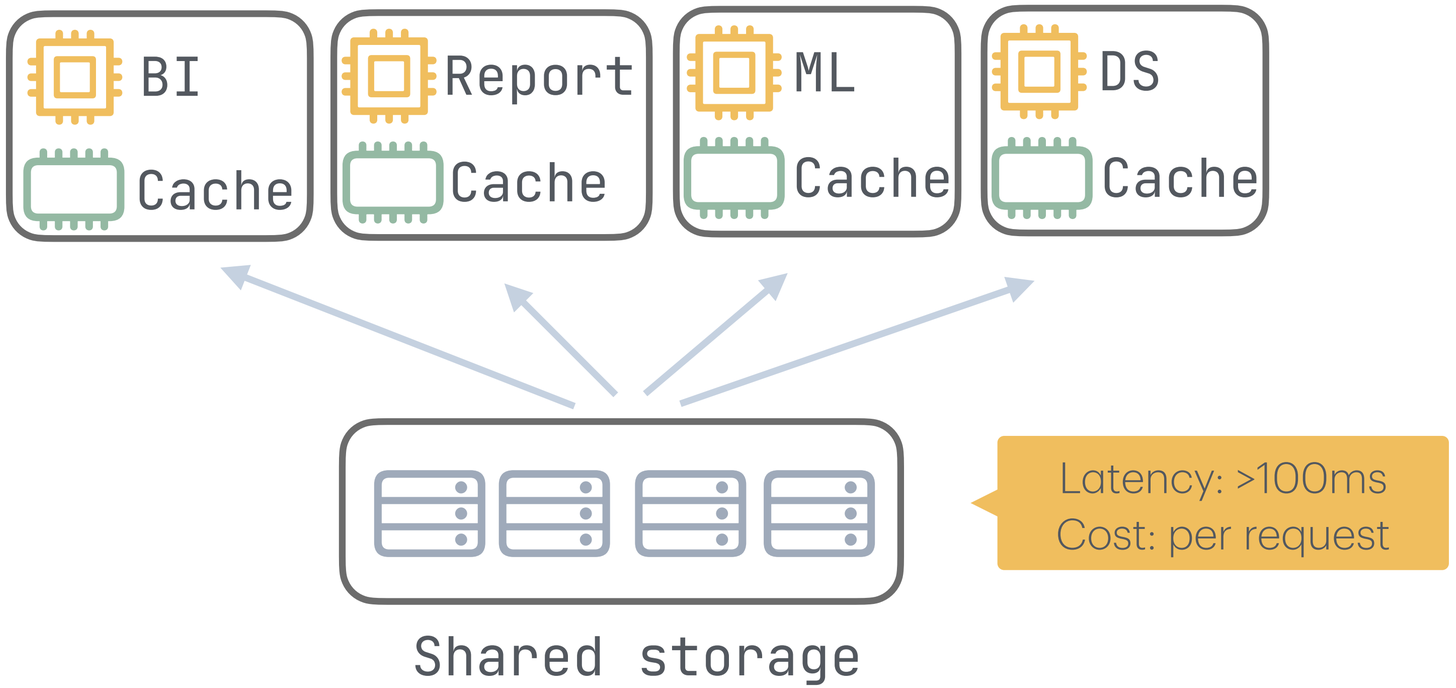

Vision: shared cache (2025+)

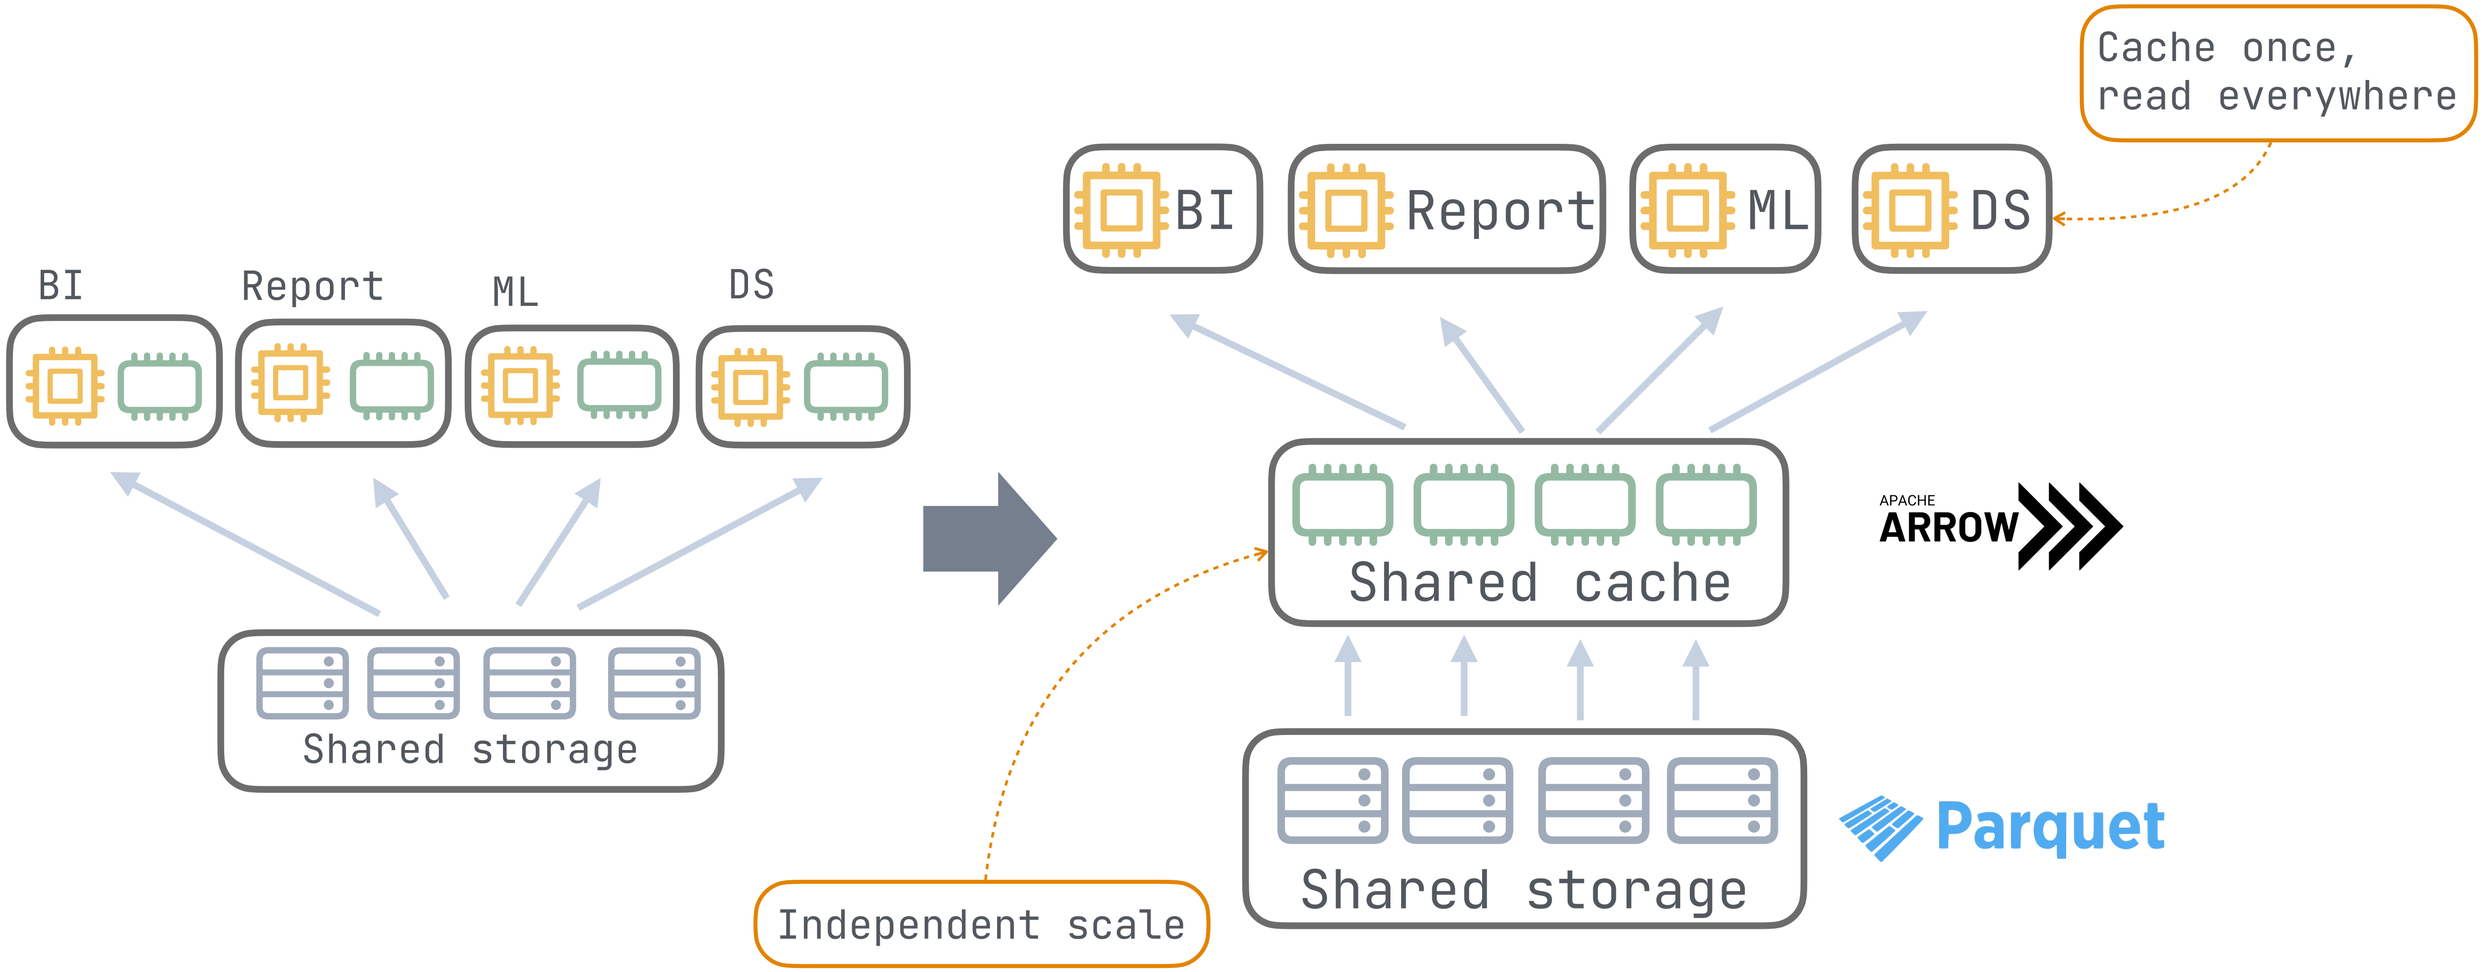

Thesis goal: LiquidCache

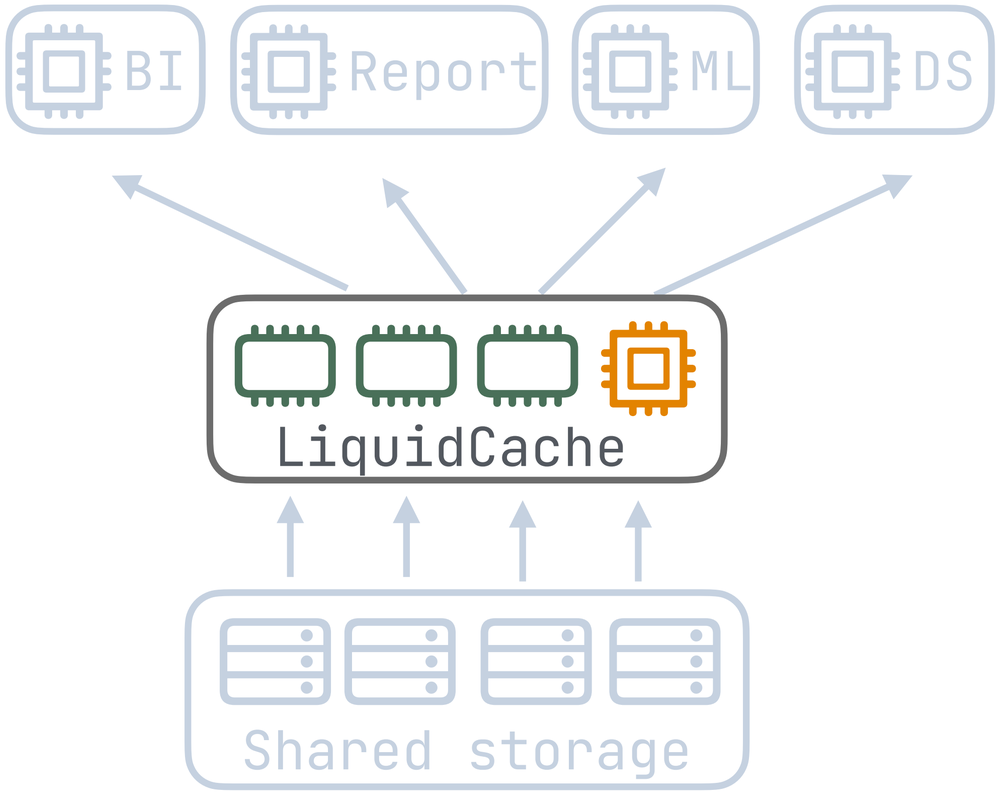

Thesis goal:

Design a cost-effective shared cache system by combining compute and data, while preserving ecosystem compatibility.

First attempt: byte cache

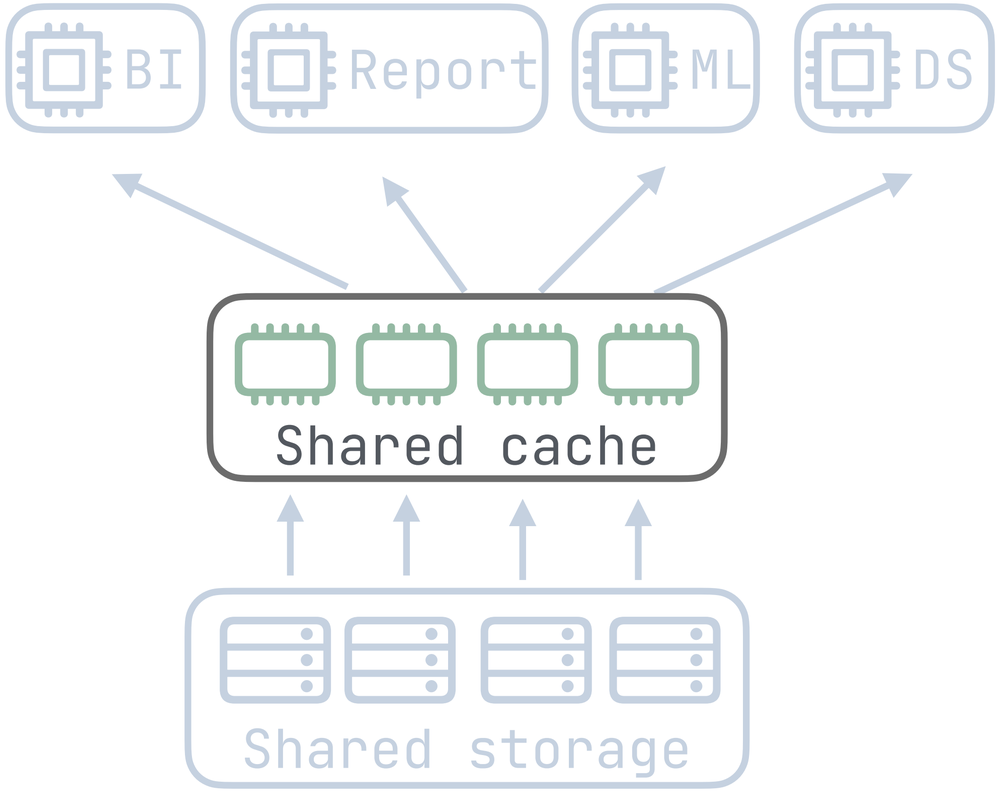

Challenges:

1. Network bottleneck
2. Inefficient cache eviction

LiquidCache = compute + data

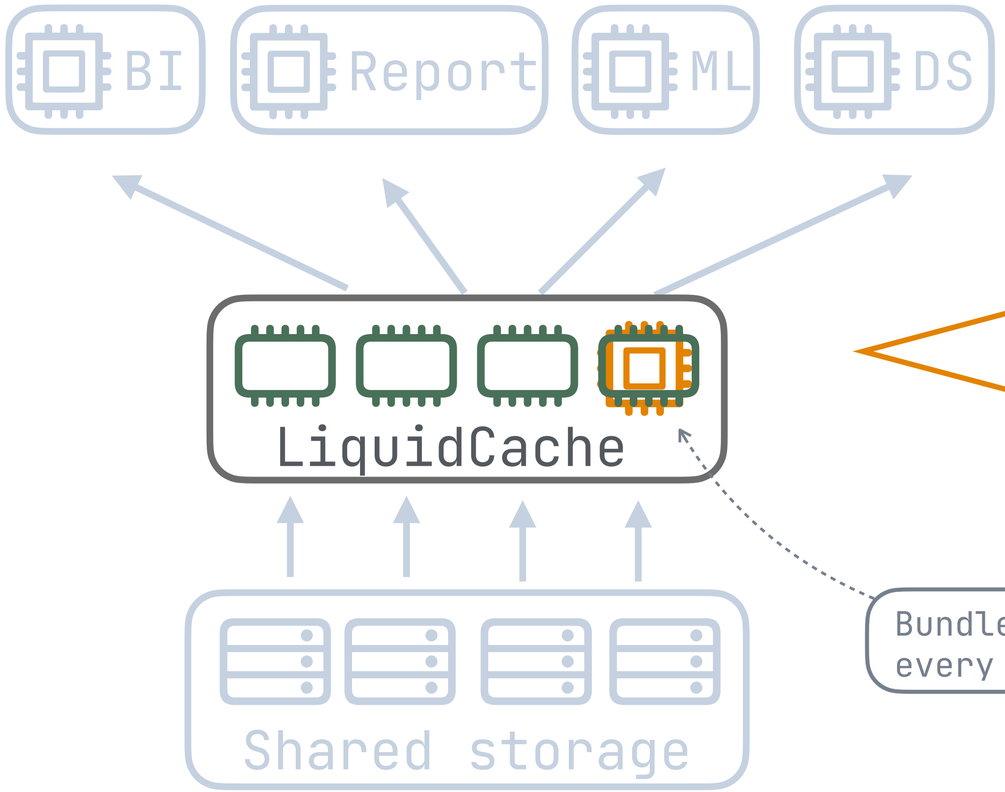

① Pushdown to reduce network

② Data-guided eviction

LiquidCache = compute + data

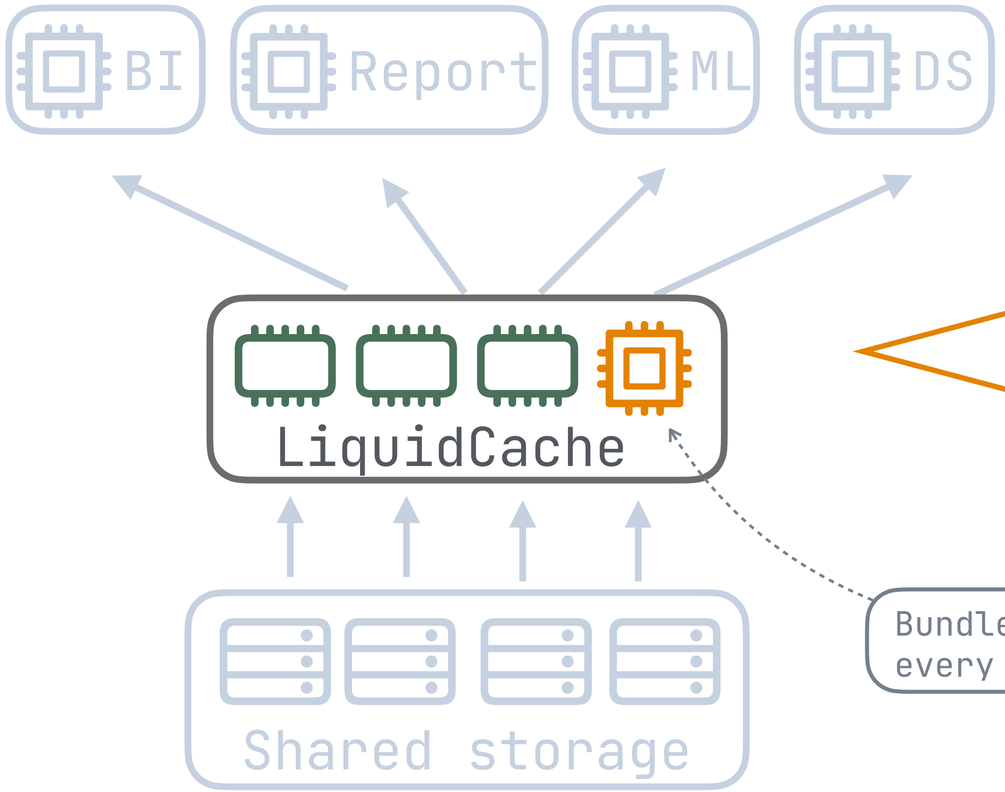

VLDB 2025
(hopefully)

① Pushdown to reduce network

② Data-guided eviction

Bundled CPU in
every server

Pushdown to reduce network

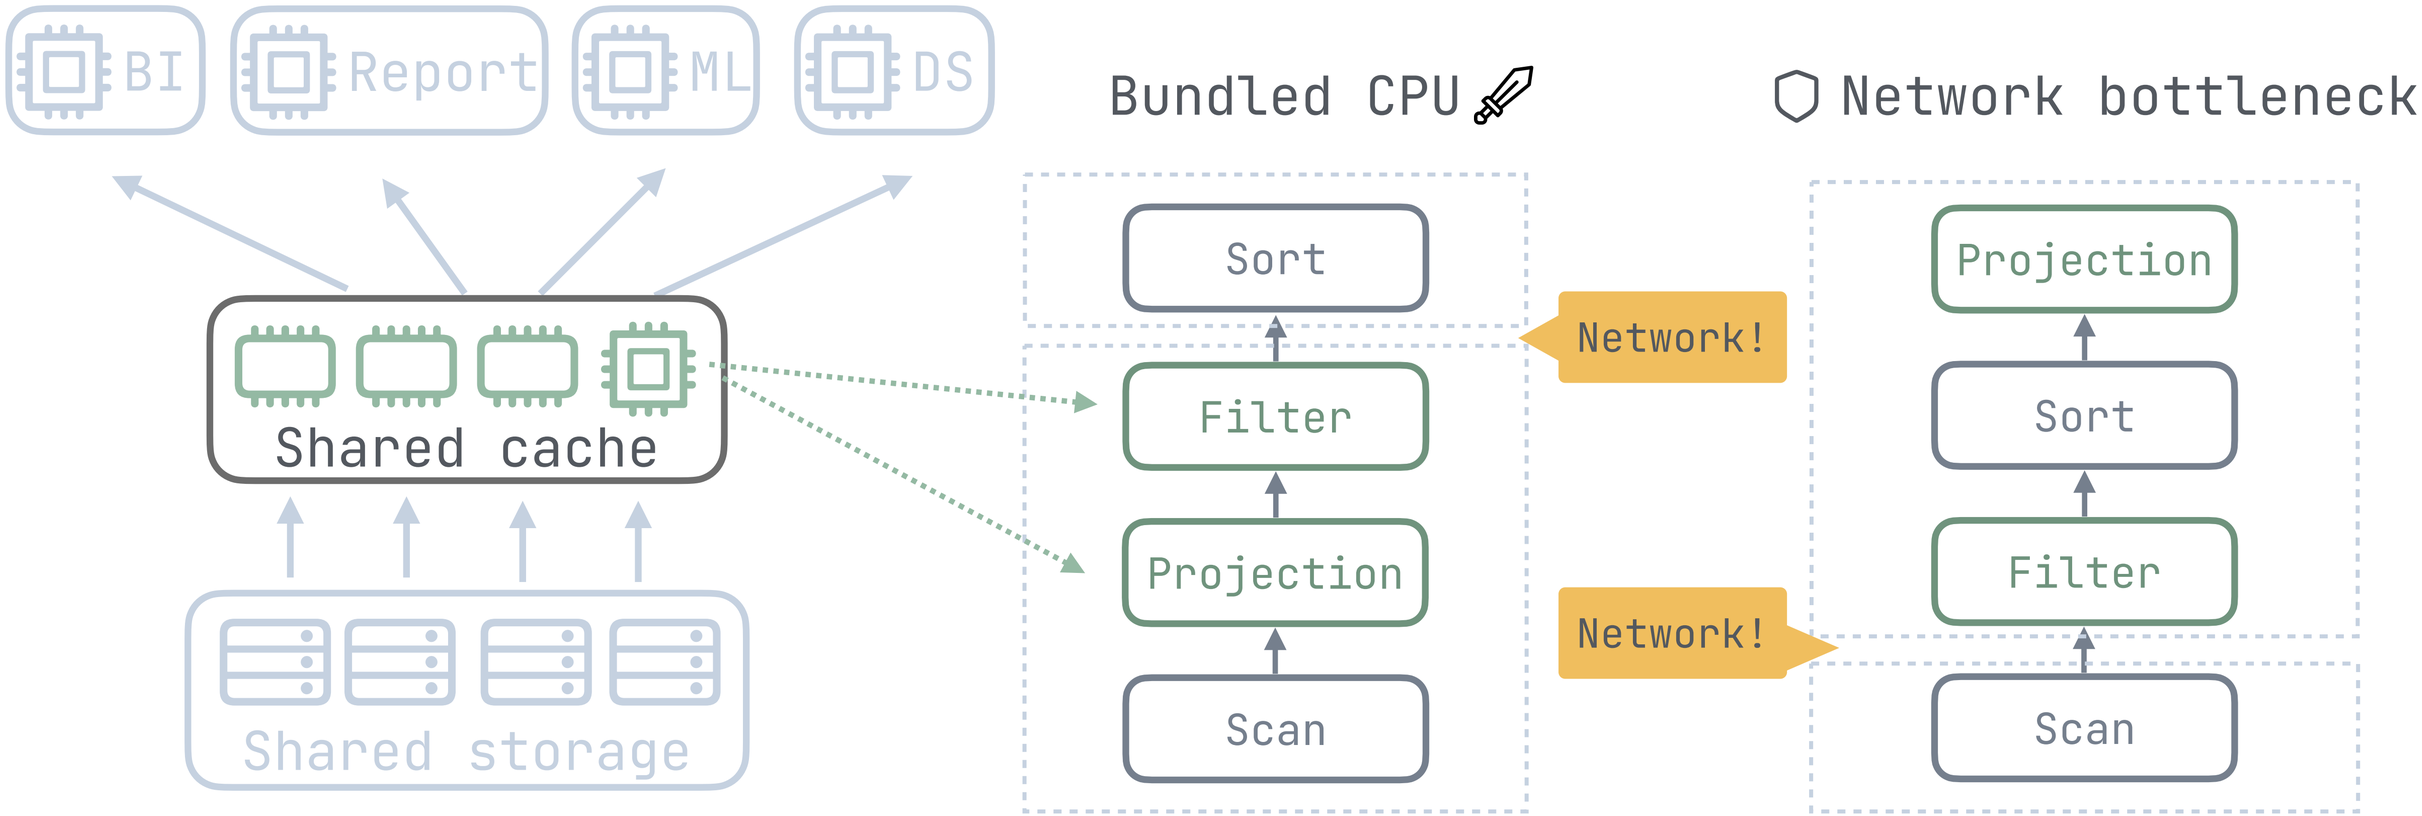

Pushdown overwhelms cache CPU

Previously believed bottleneck:
Filter evaluation

Our findings:
Data decoding

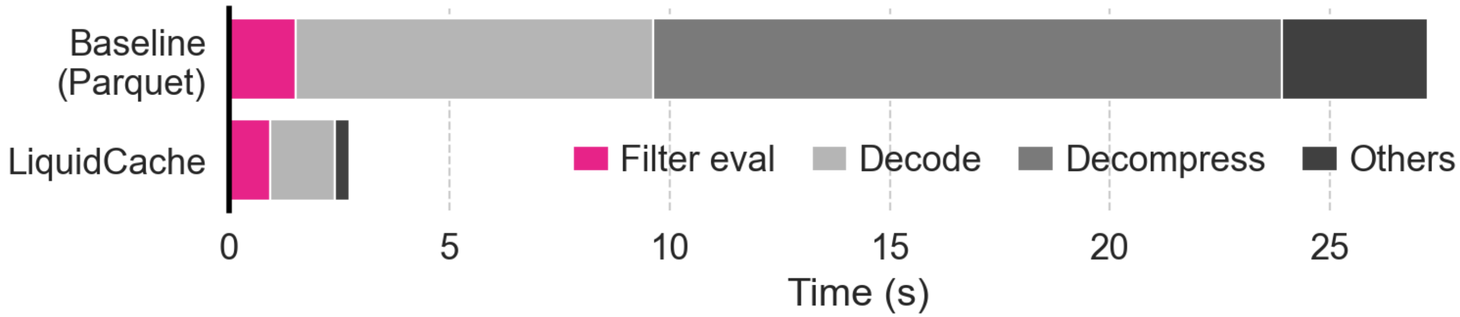

LiquidCache

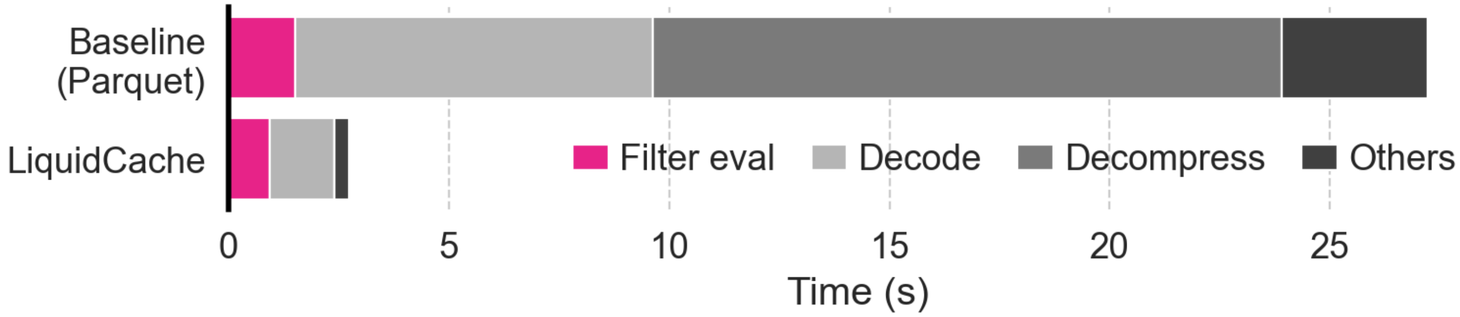

Part 1:
co-designed format to skip decoding

Part 2:
progressive, selective, asynchronous transcoding

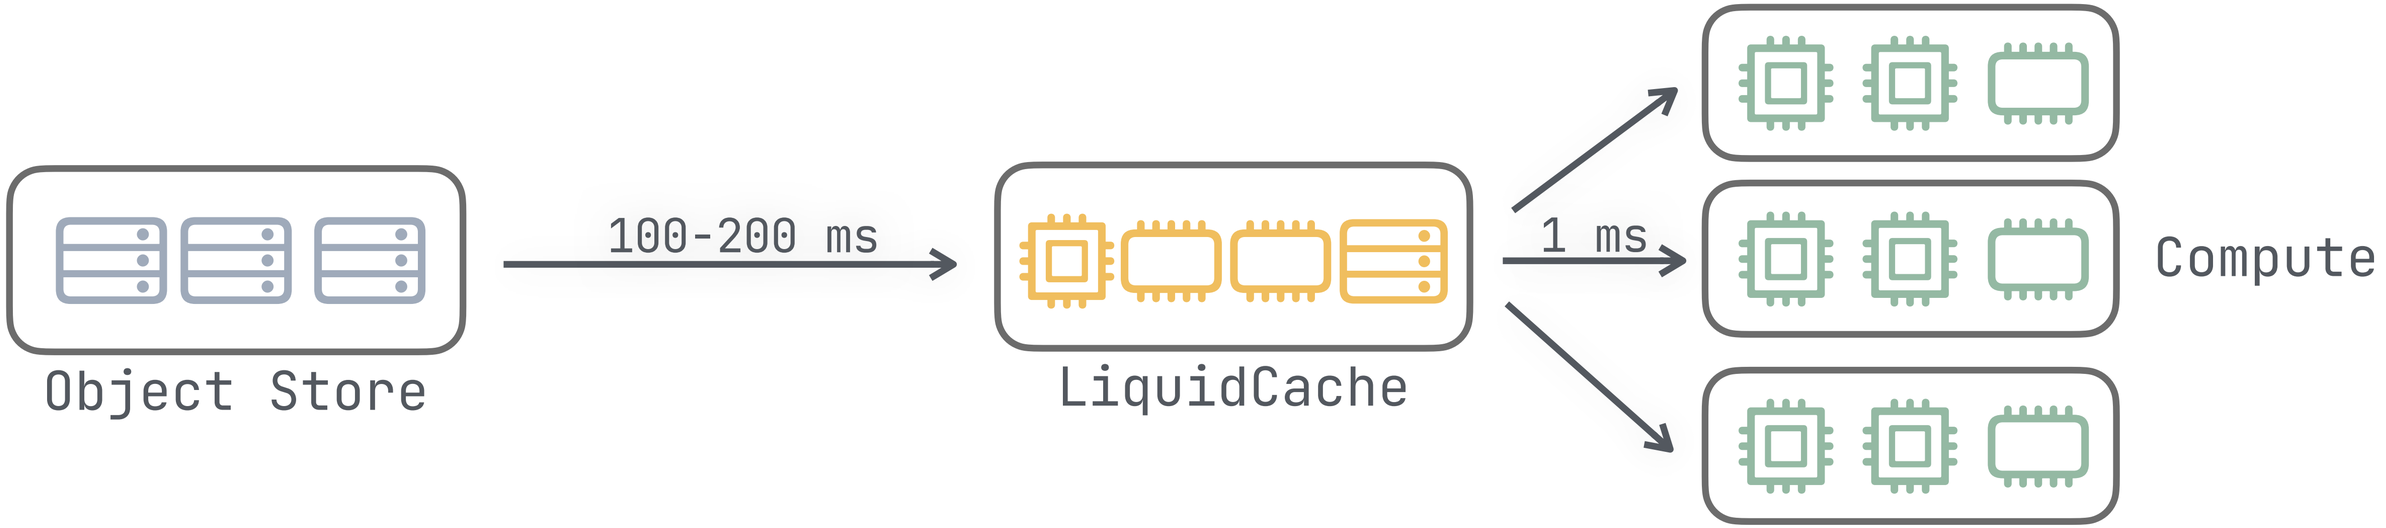

Part 1:
co-designed format to skip decoding

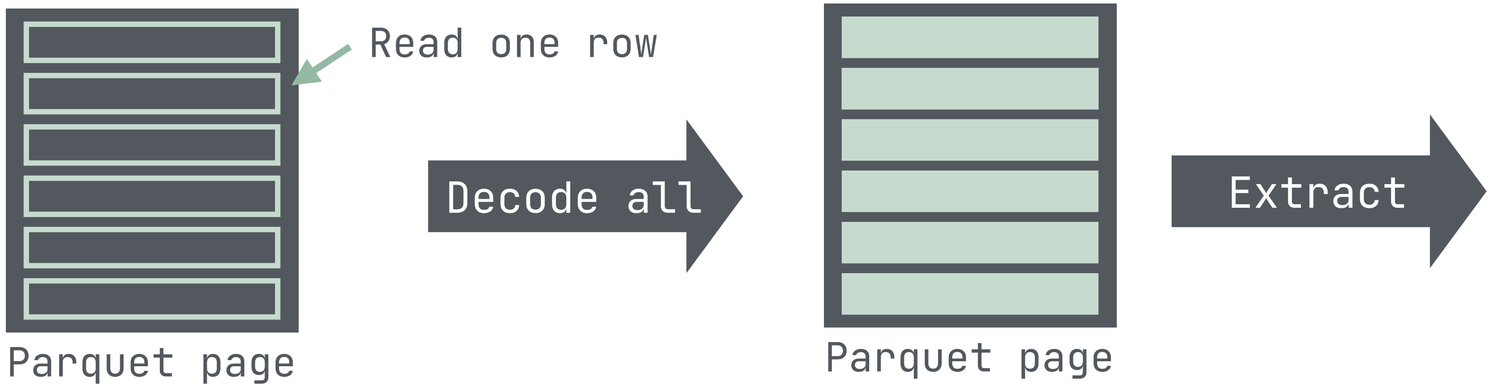

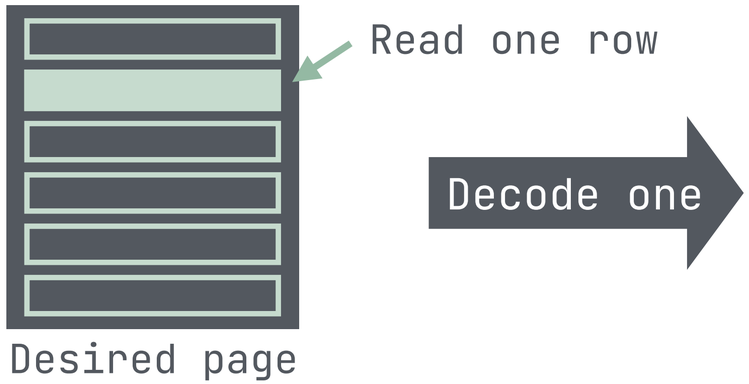

Co-design principle:
Each row must be independently decodable

Each row must be independently decodable (string example)

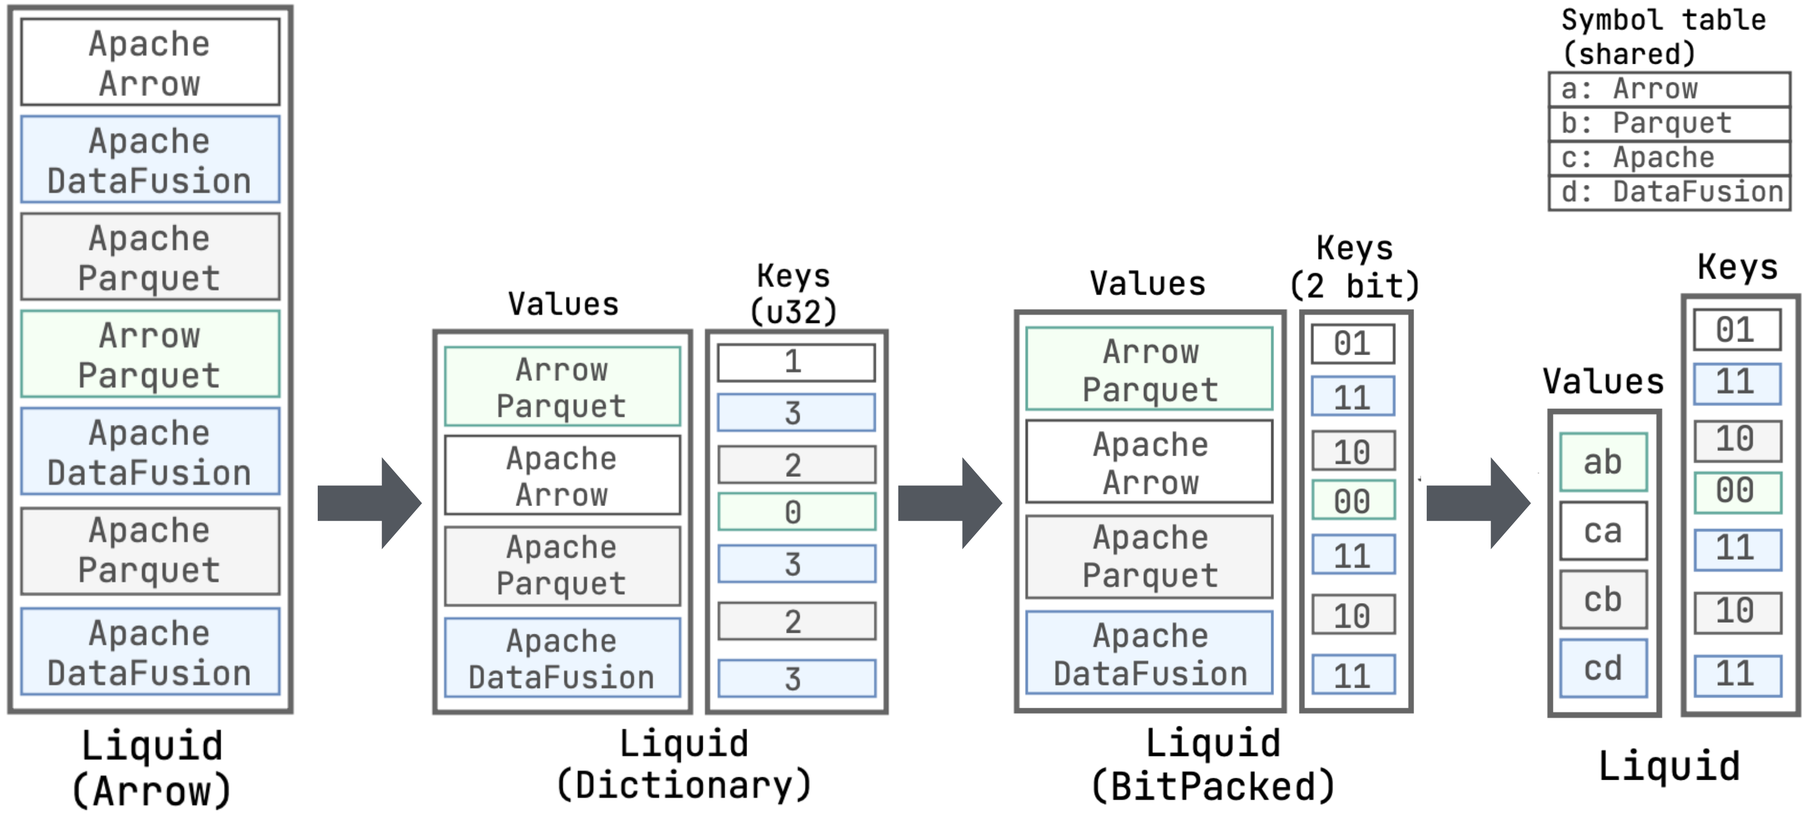

No general purpose compression

Leverages state-of-the-art encoding schemes

Carefully designed encoding/layout for each data types

Co-design with filter pushdown (selective decoding)

```sql
SELECT val, location
FROM sensor_data
WHERE location = 'office';
```

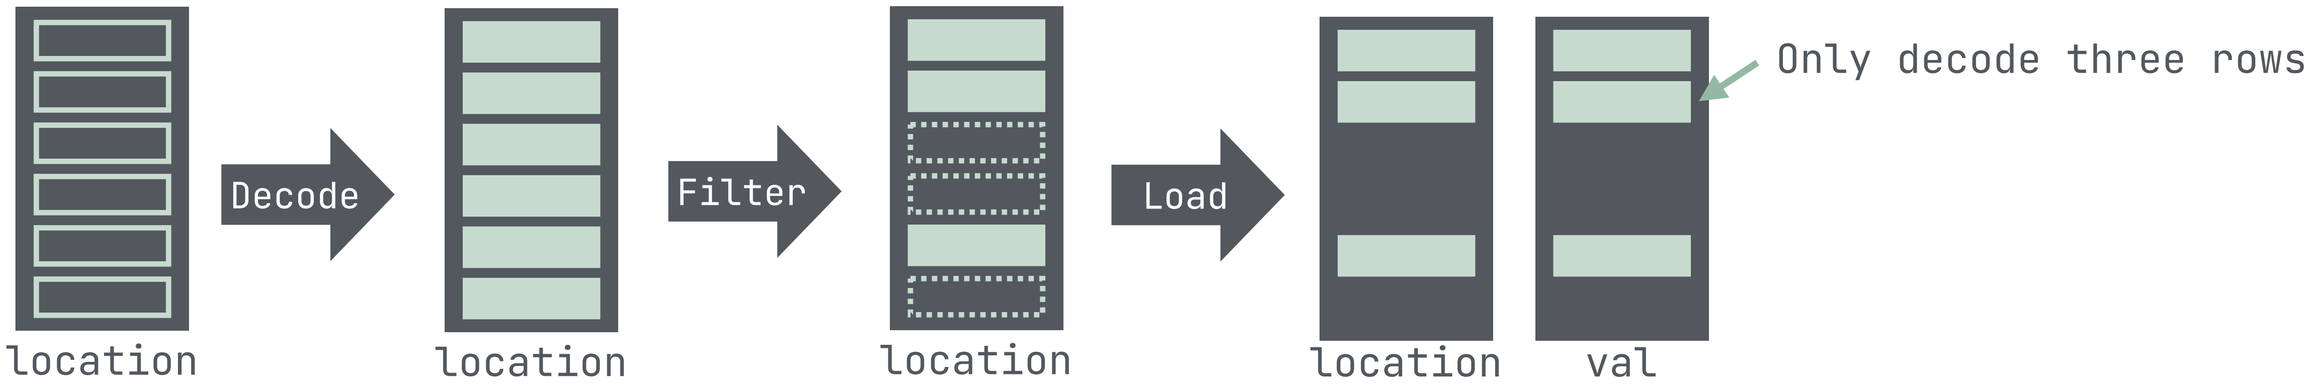

Co-design with filter pushdown (filter late materialization)

```sql
SELECT val, location
FROM sensor_data
WHERE location = 'office', date > '2025-03-21';
```

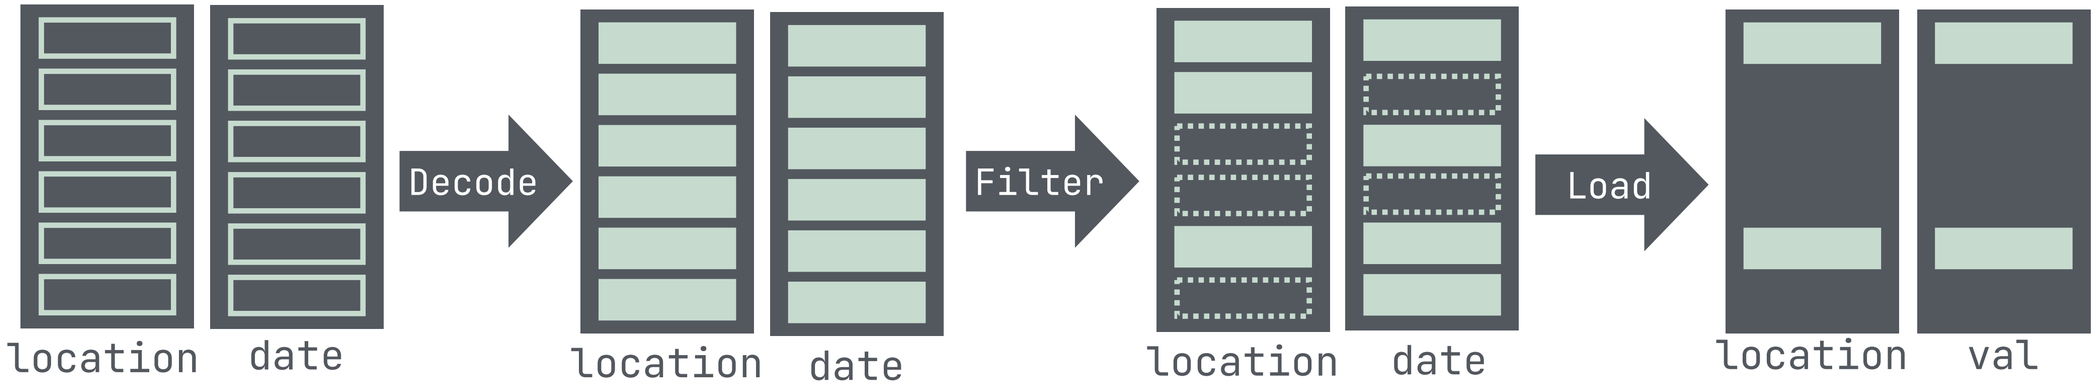

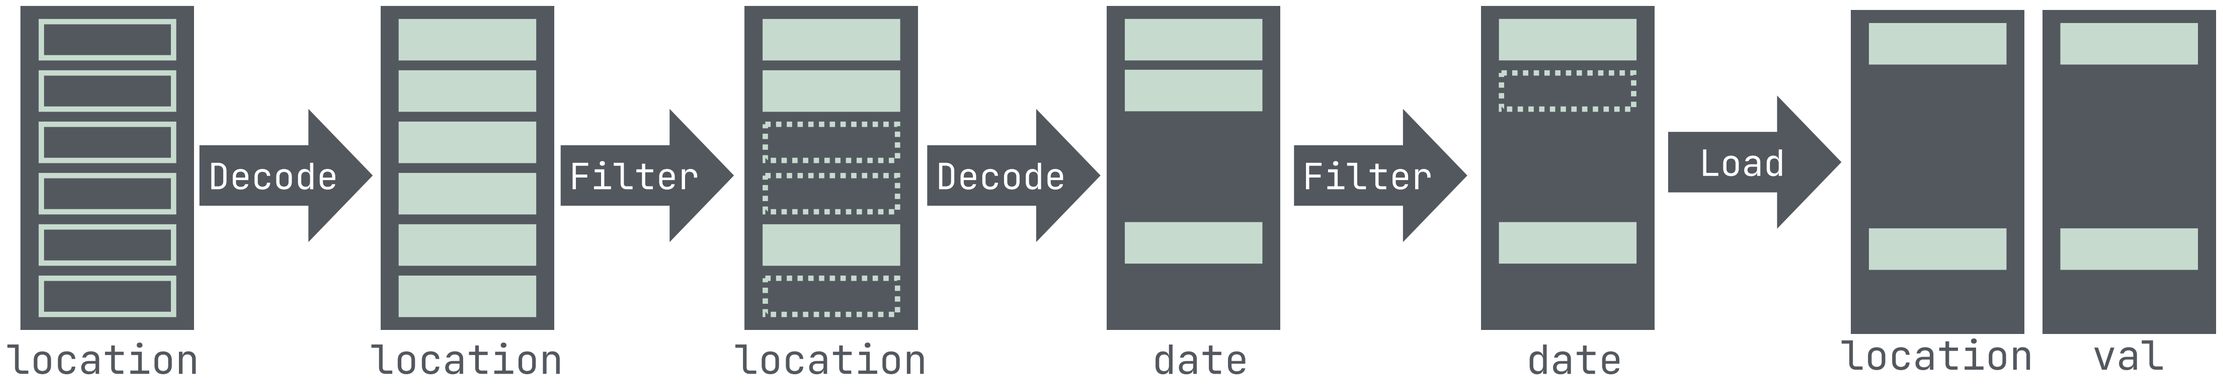

Co-design with filter pushdown (evaluate on encoded data)

```
SELECT val, location
FROM sensor_data
WHERE name = 'Apache Arrow'
```

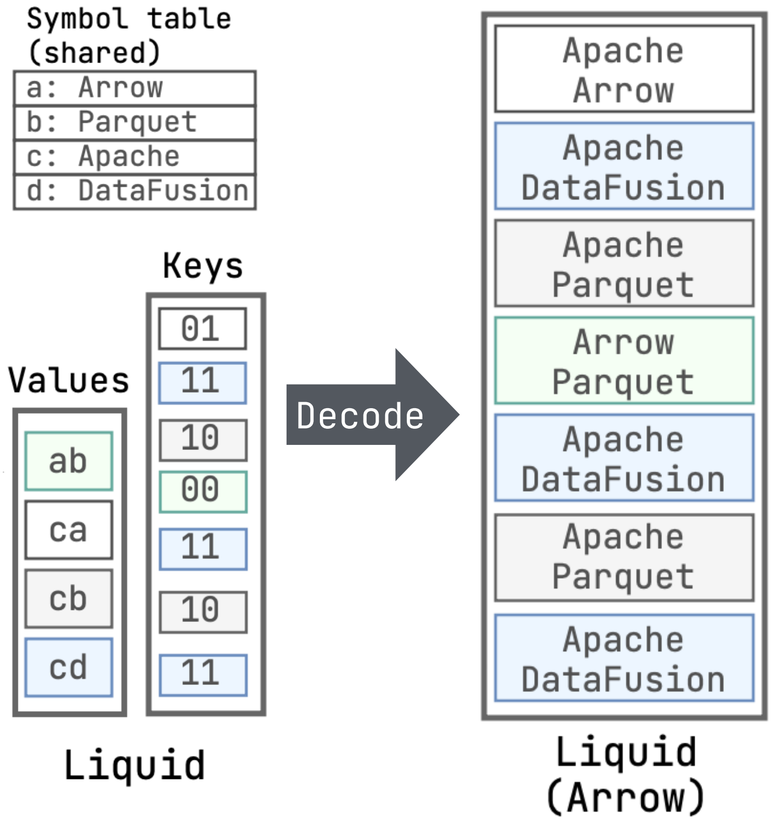

Baseline

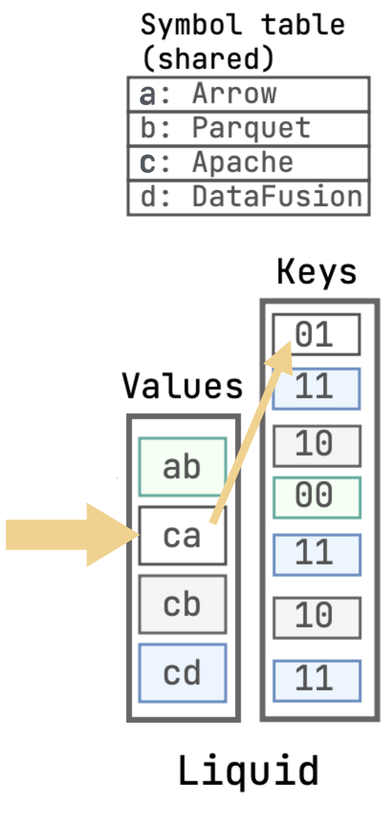

Evaluate on encoded

In paper:
Evaluate on **partially** encoded data

LiquidCache

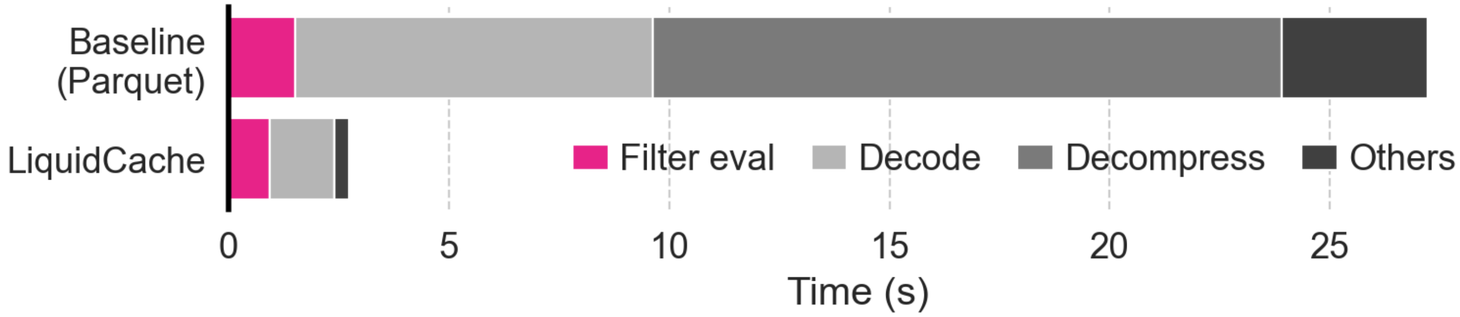

Part 1:
co-designed format to skip decoding

Part 2:
progressive, selective, asynchronous transcoding

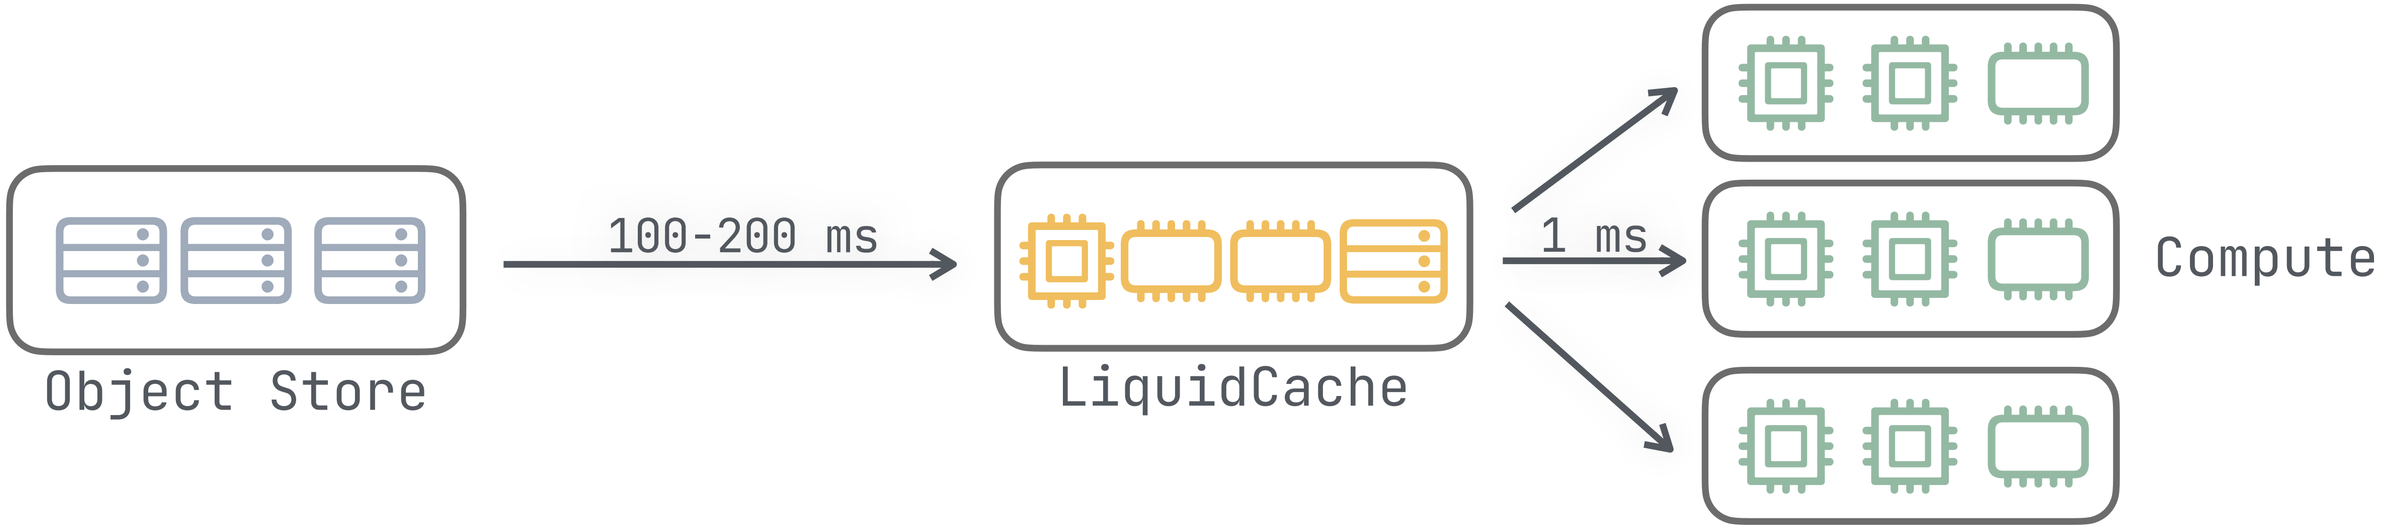

Part 2:
progressive, selective, asynchronous transcoding

Yet another file format?

Nimble
LanceDB
Vortex
ClickHouse
DuckDB

No!

Technical: they are not much different
Organizational: license, governance, ecosystem
LiquidCache: progressively bend the world

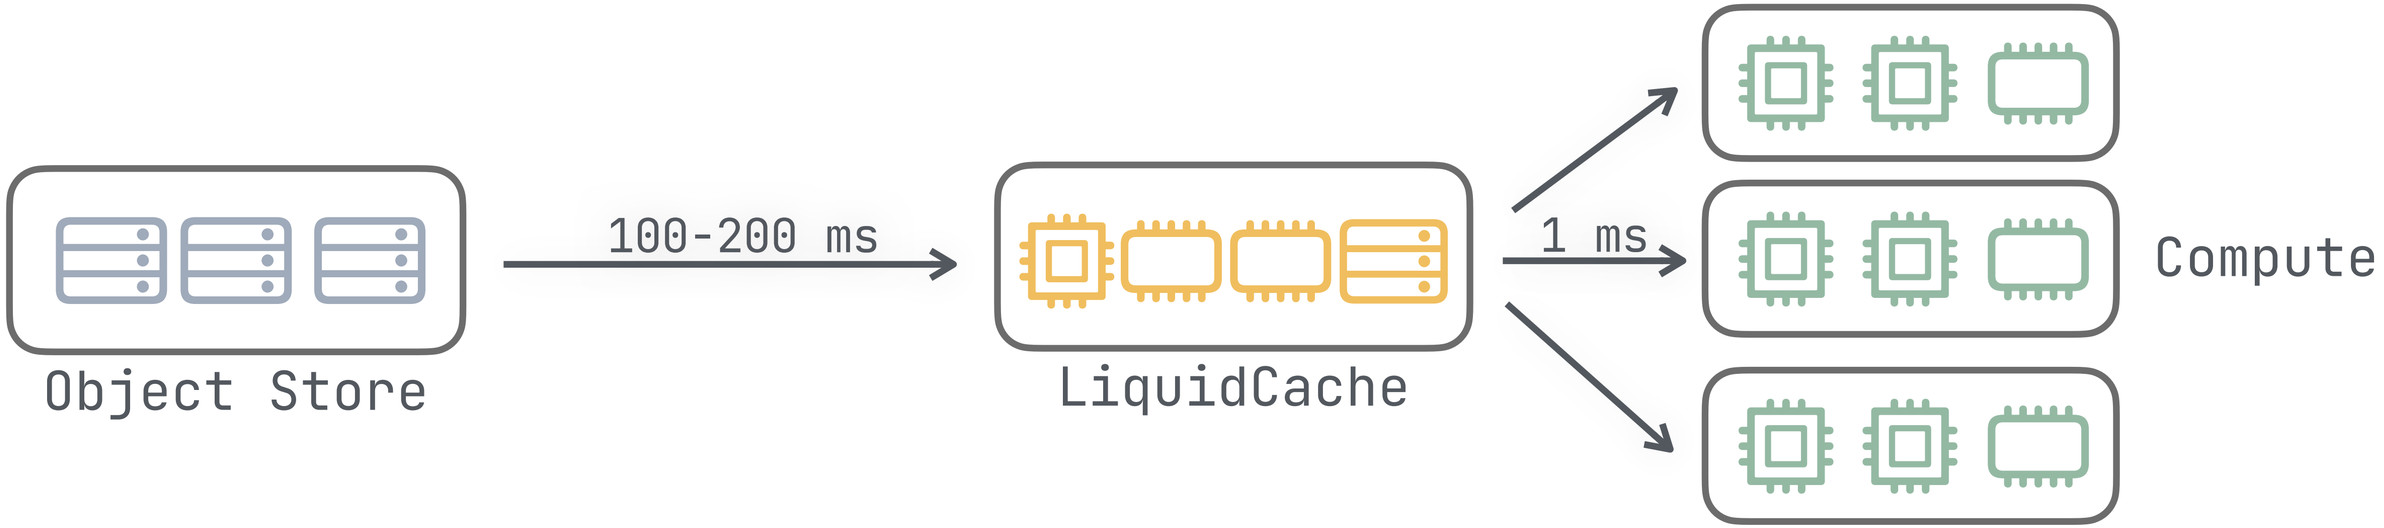

Progressive transcoding

Progressive:
Transcode as needed,
no upfront cost

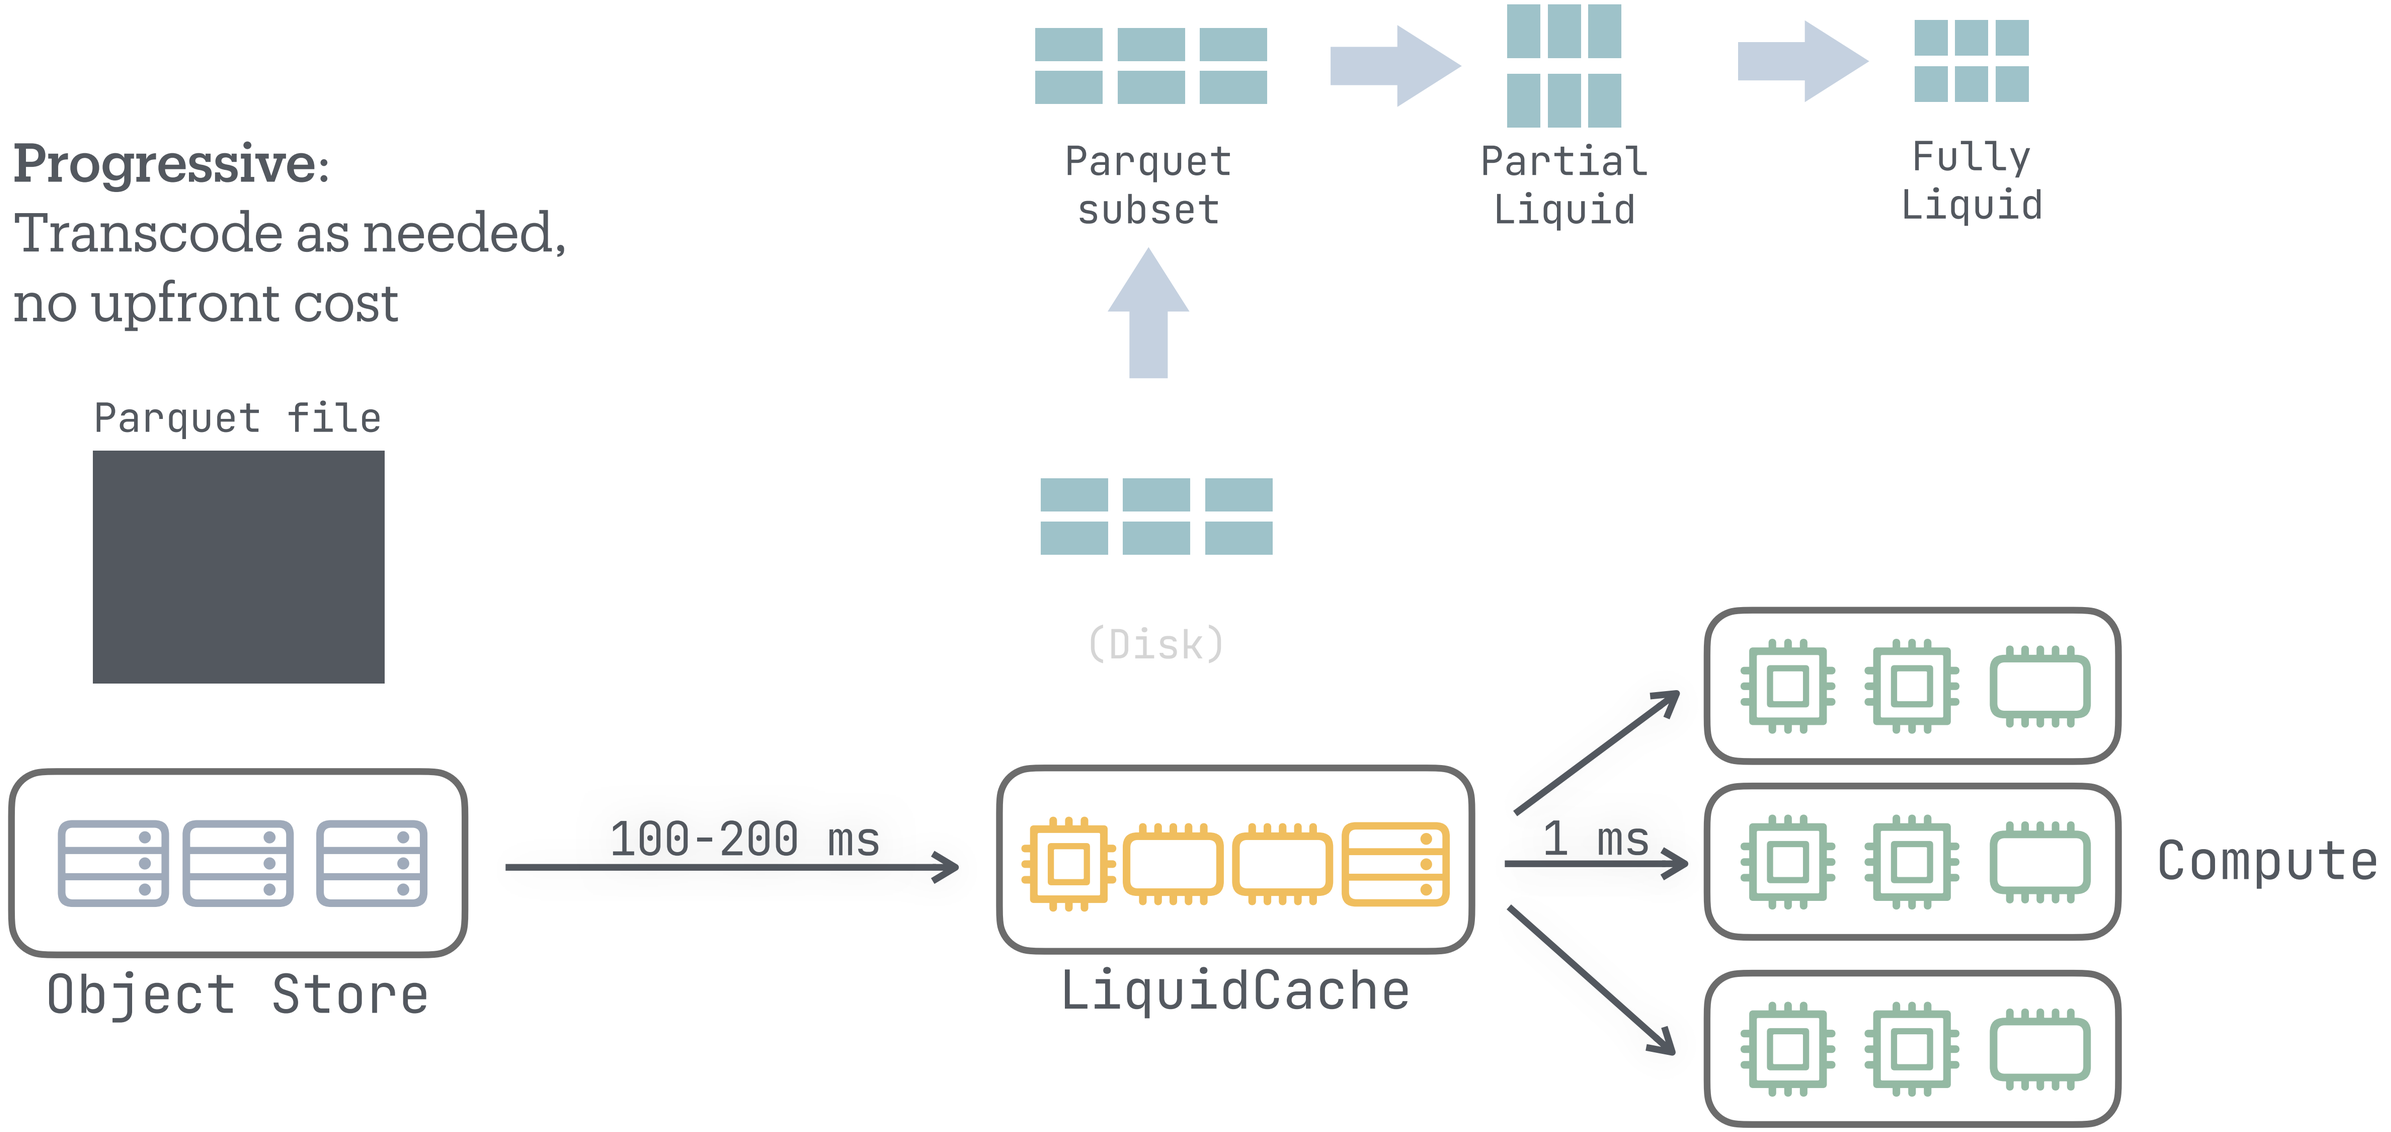

Selective transcoding

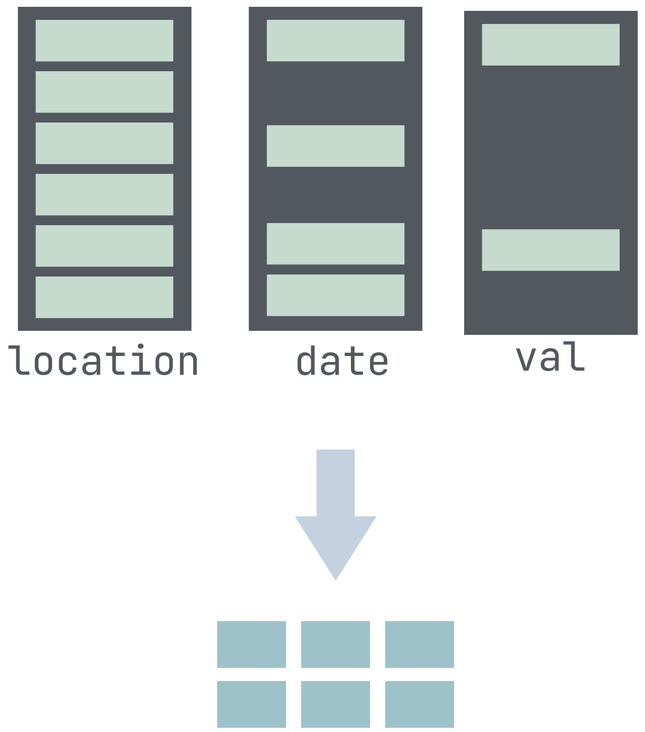

Selective:
Transcode only touched data

Asynchronous transcoding

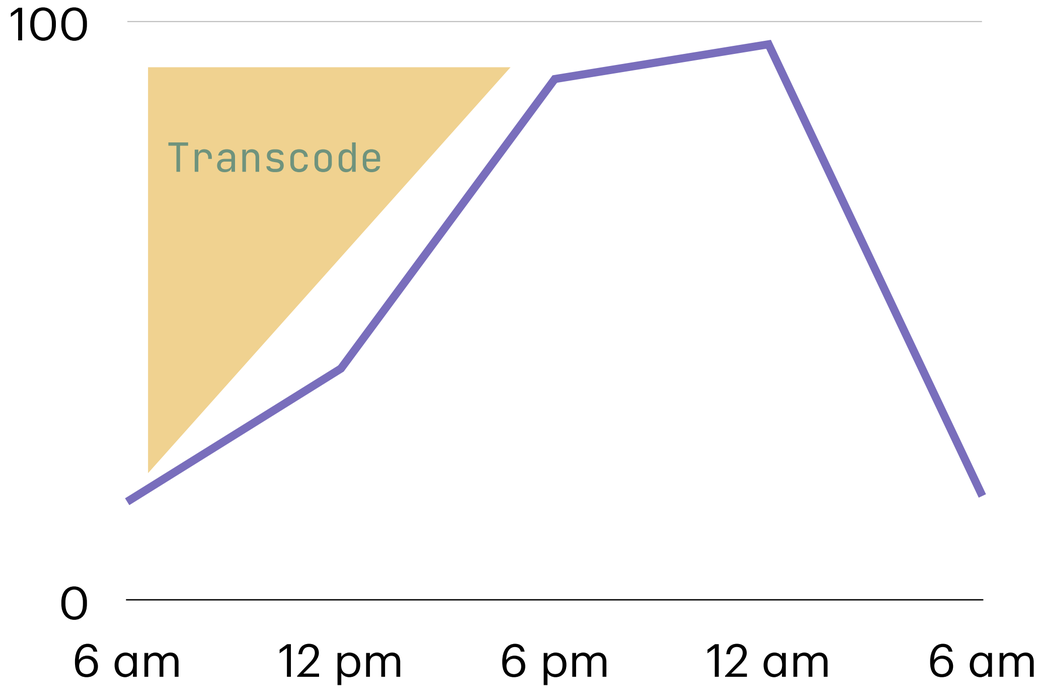

Asynchronous:
Transcode when less busy

LiquidCache

10x { cost | latency } reduction

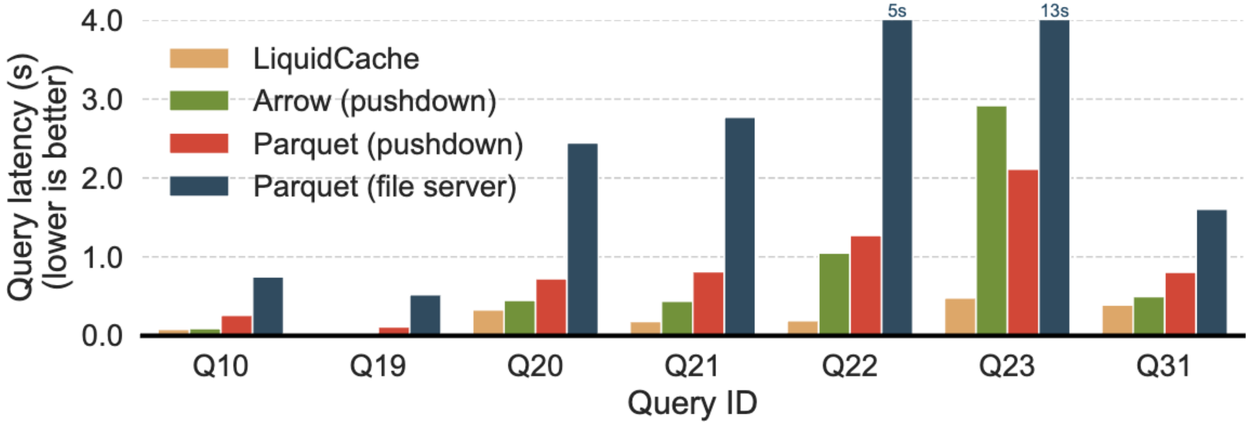

With same memory: 10x lower latency

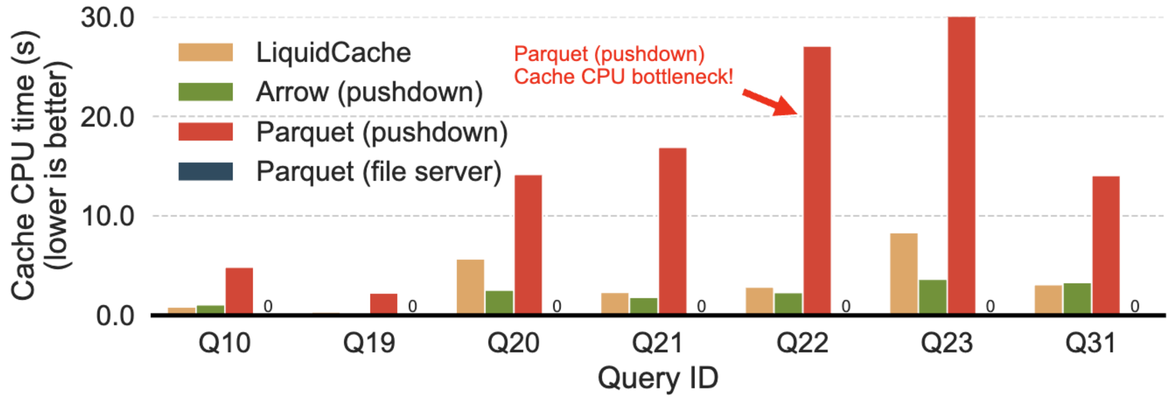

With same CPU: 10x lower CPU time

LiquidCache

10x { cost | latency } reduction

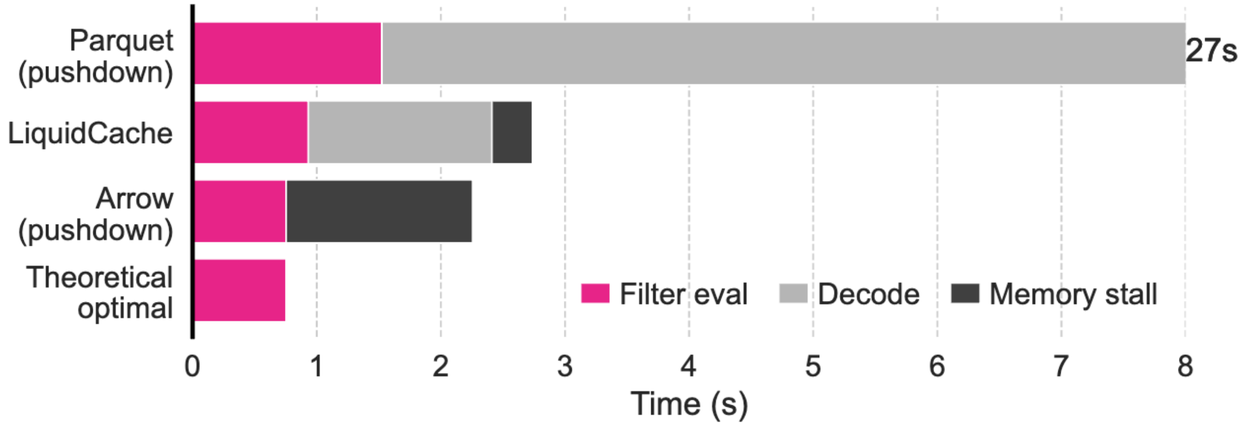

Decoding cost: close to theoretical optimal

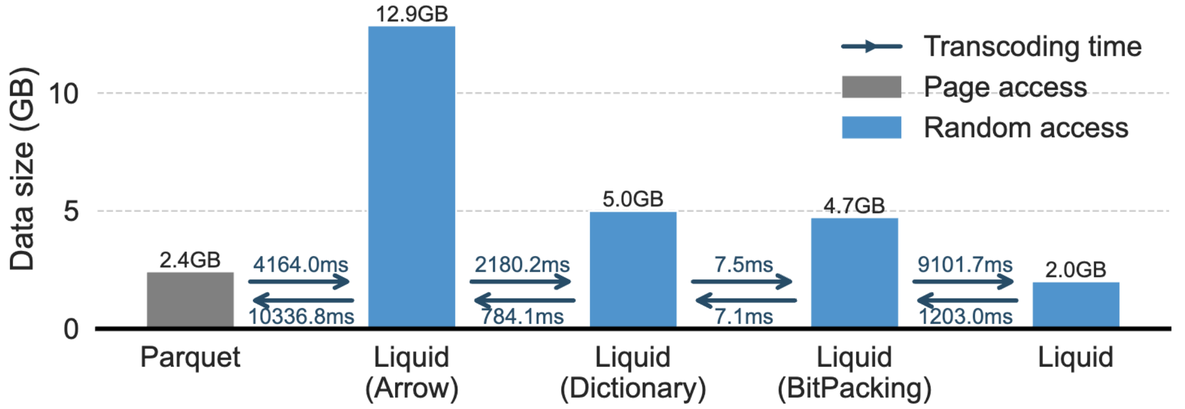

Compression ratio: comparable to Parquet

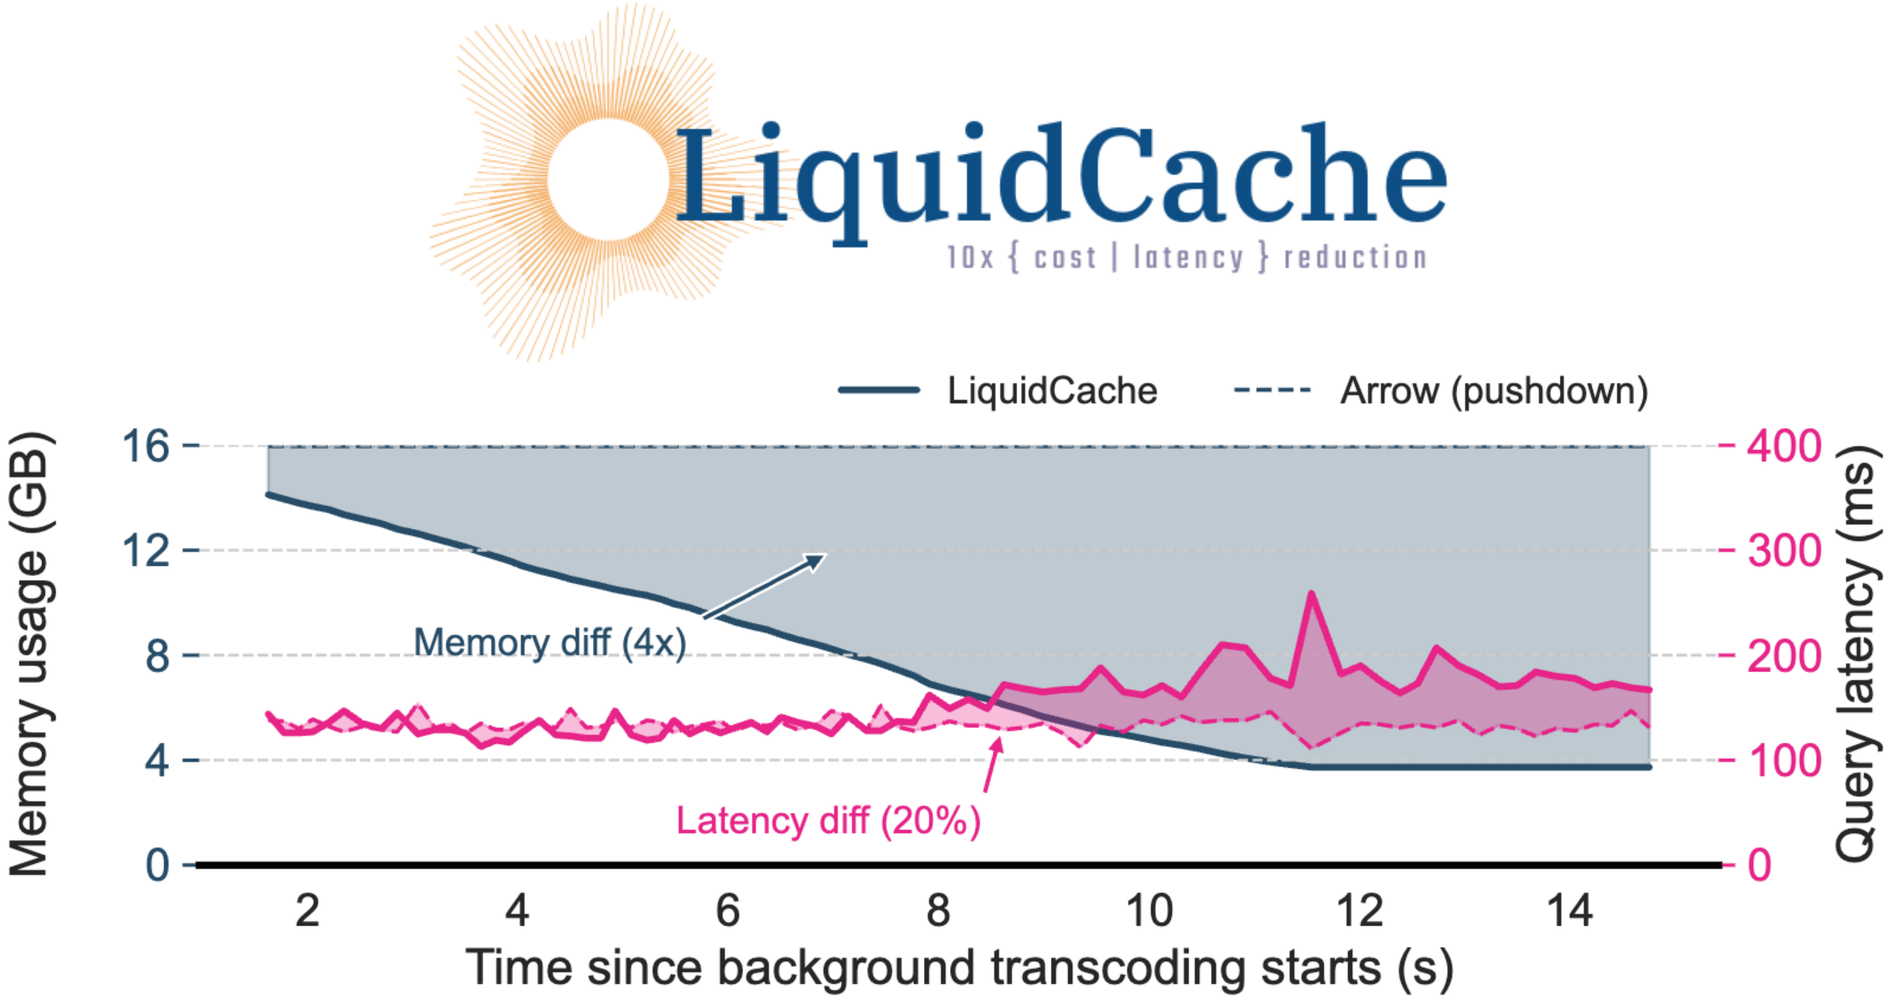

Transcoding cost: negligible, no latency spike

LiquidCache = compute + data

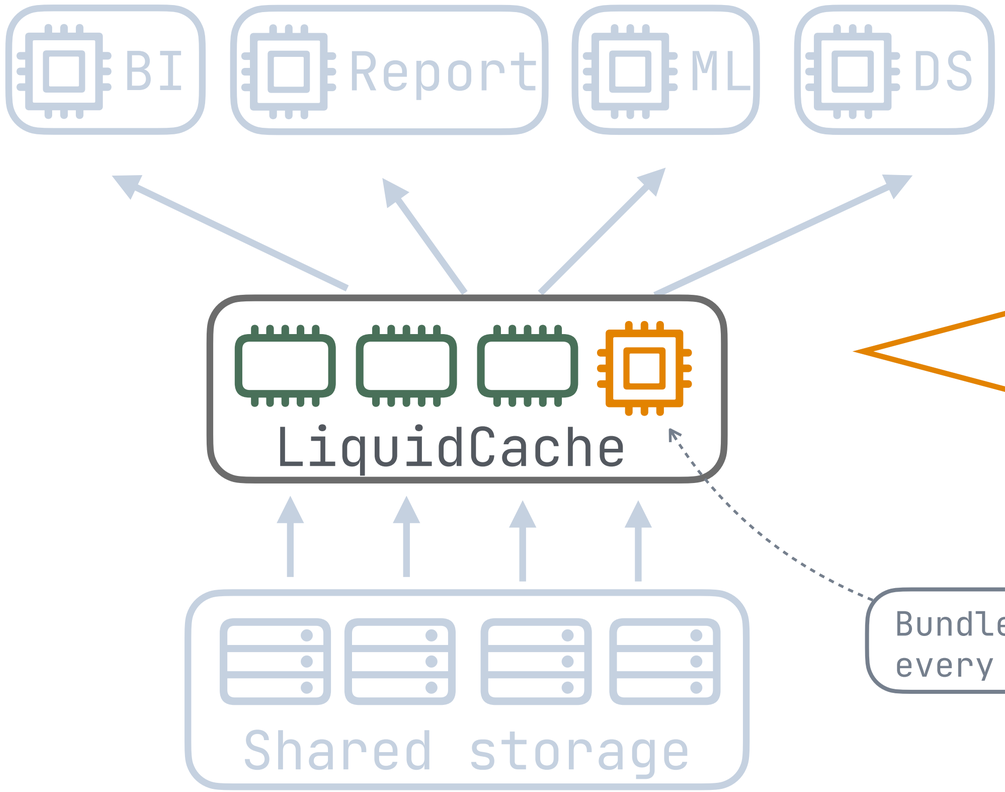

① Pushdown to reduce network

② Data-guided eviction

Proposed work

Motivating example

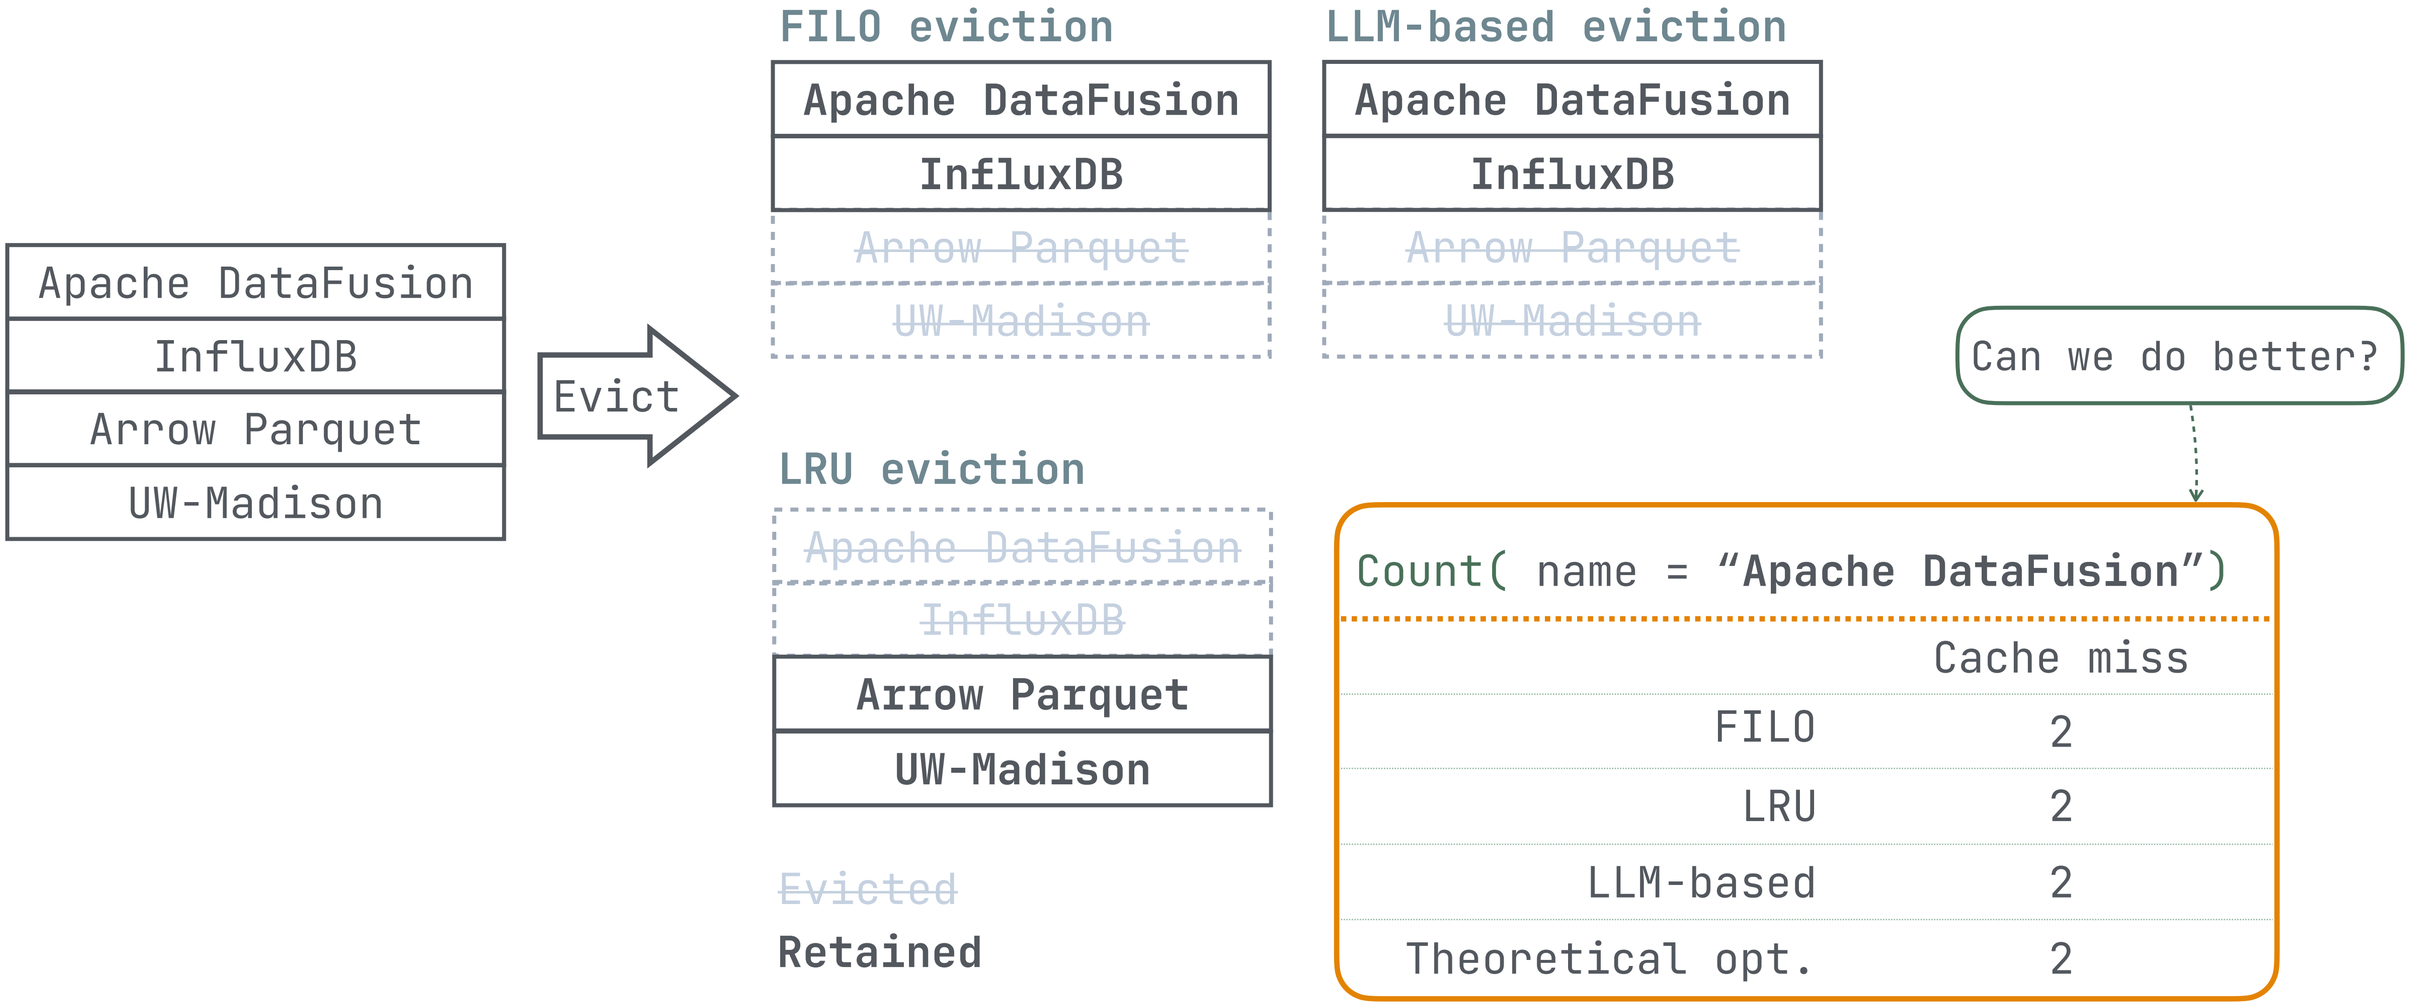

Data-guided eviction

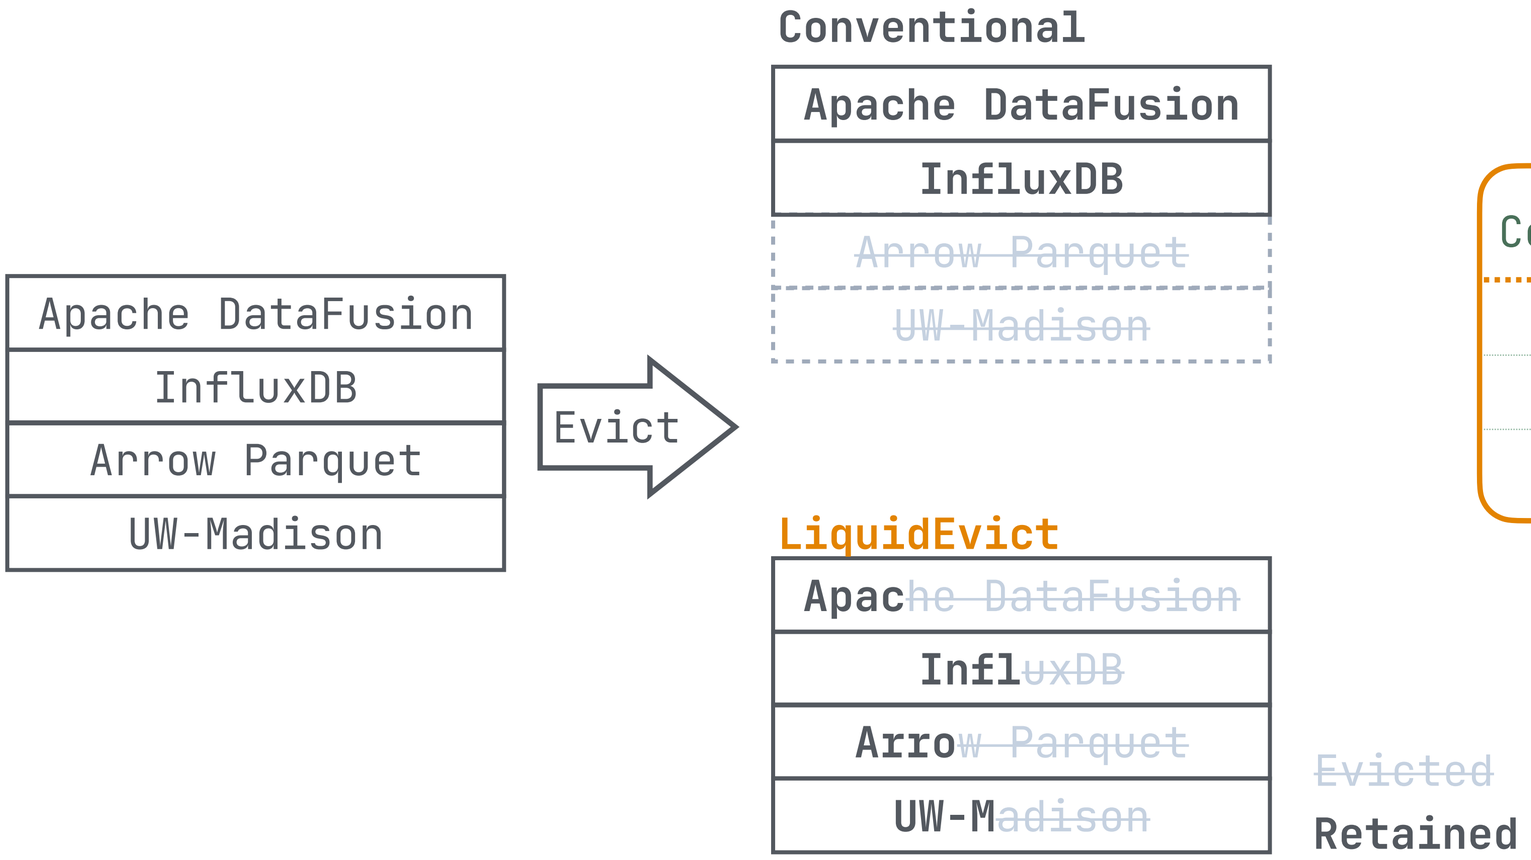

`Count( name = “Apache DataFusion”)`

	Cache miss
Conventional	2
LiquidEvict	**1**

Insight: evict partial data

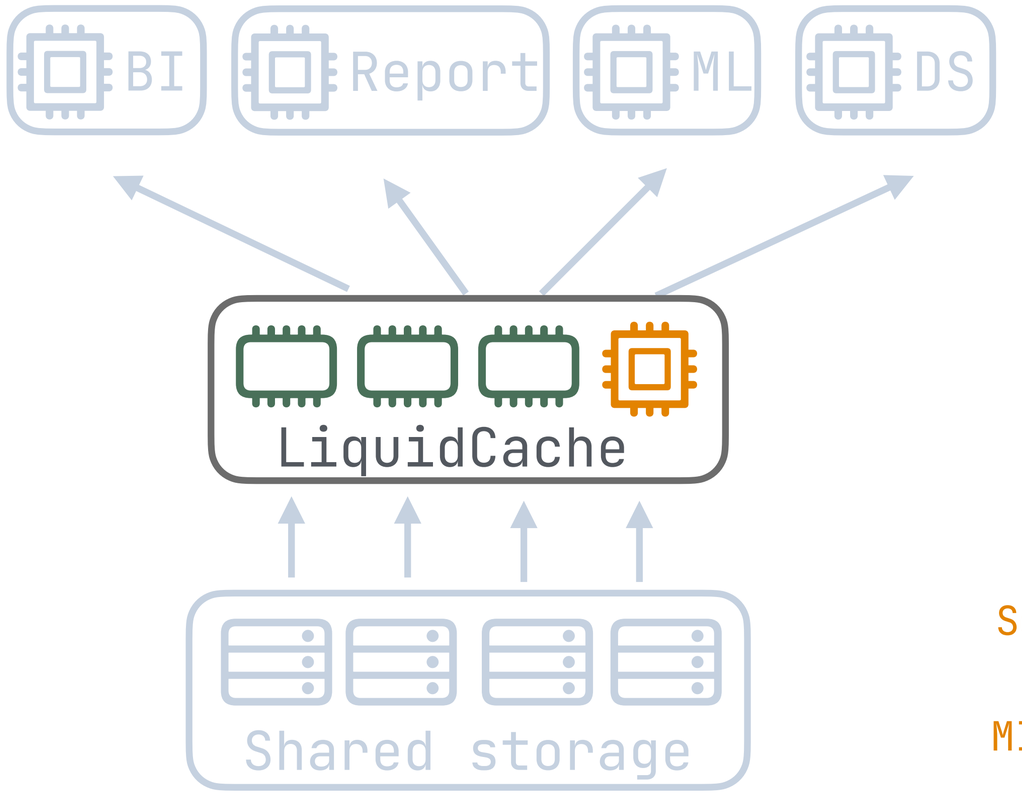

Some parts of data are more important than others

String prefix	Bit width	Dictionary
String length	Min/max/avg	Nullable mask

r_name = 'EUROPE'

SearchPharse ≠ '' p_brand = 'Brand#23' MIN("URL")

MIN("TITLE") **Real-world queries**

MobilePhoneModel ≠ '' n_name = 'SAUDI ARABIA'

UserID = 4350909328996404 49 l_quantity < 24

Evicts data, keep summaries

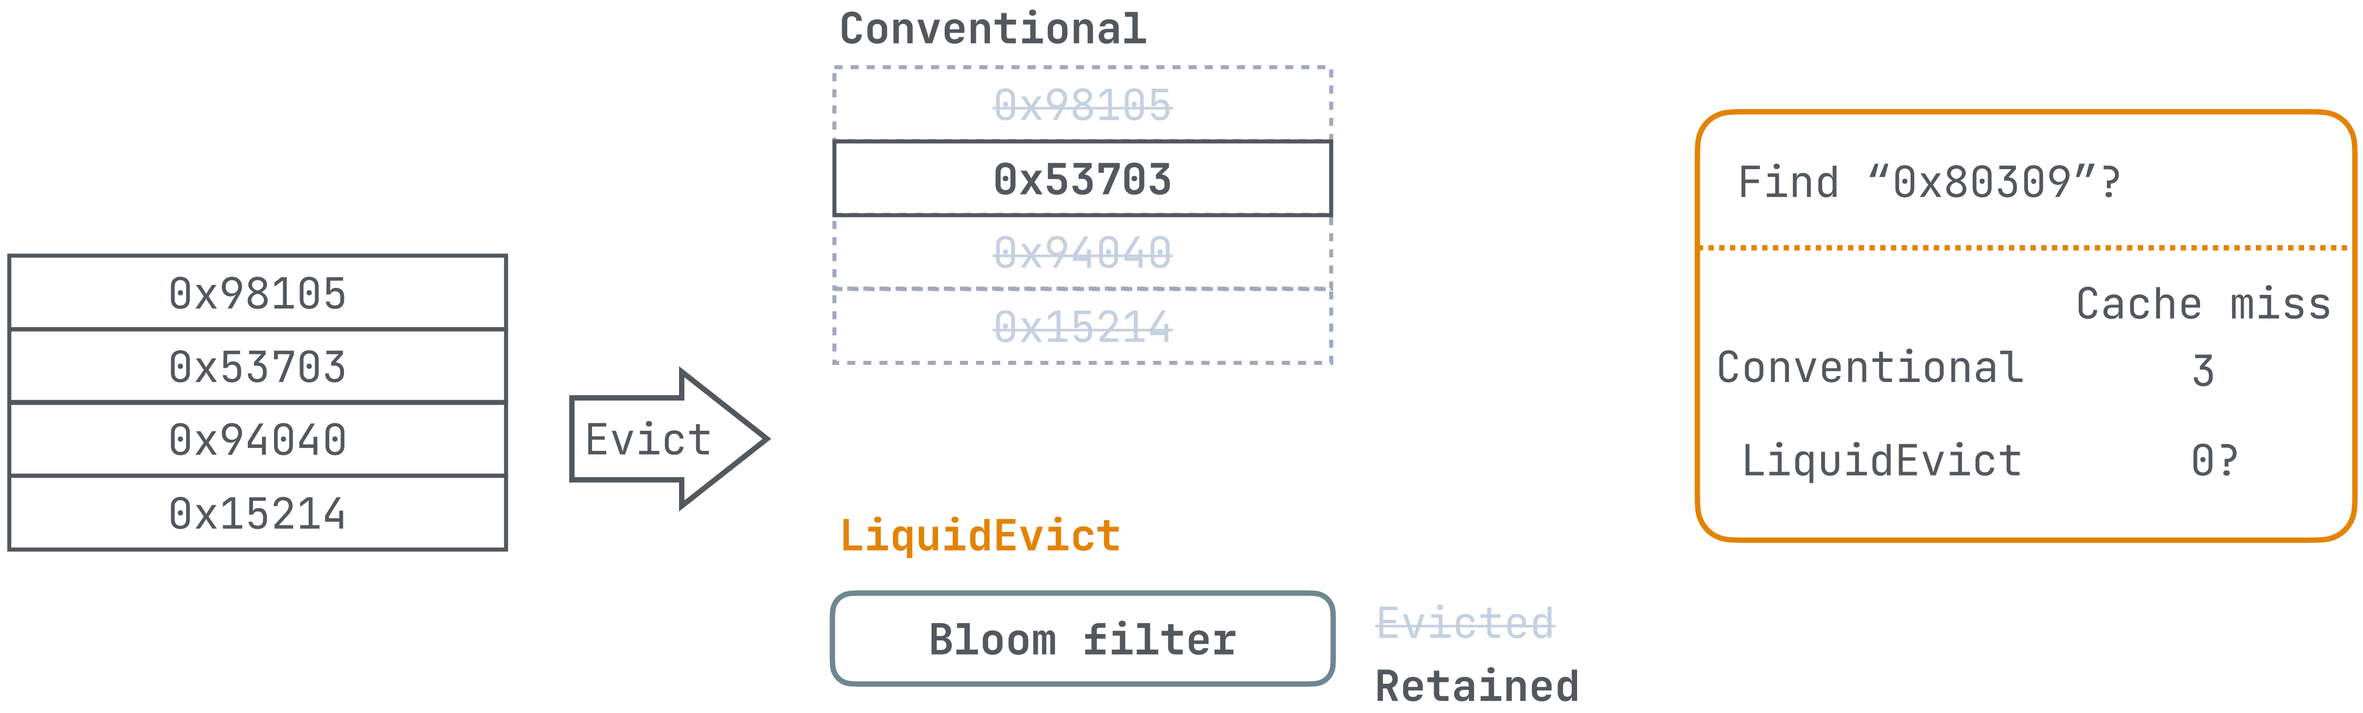

Real world example: StringView eviction

Logical content

Apache DataFusion
InfluxDB
Arrow Rust Impl
Parquet pushdown
Apache DataFusion

Physical representation

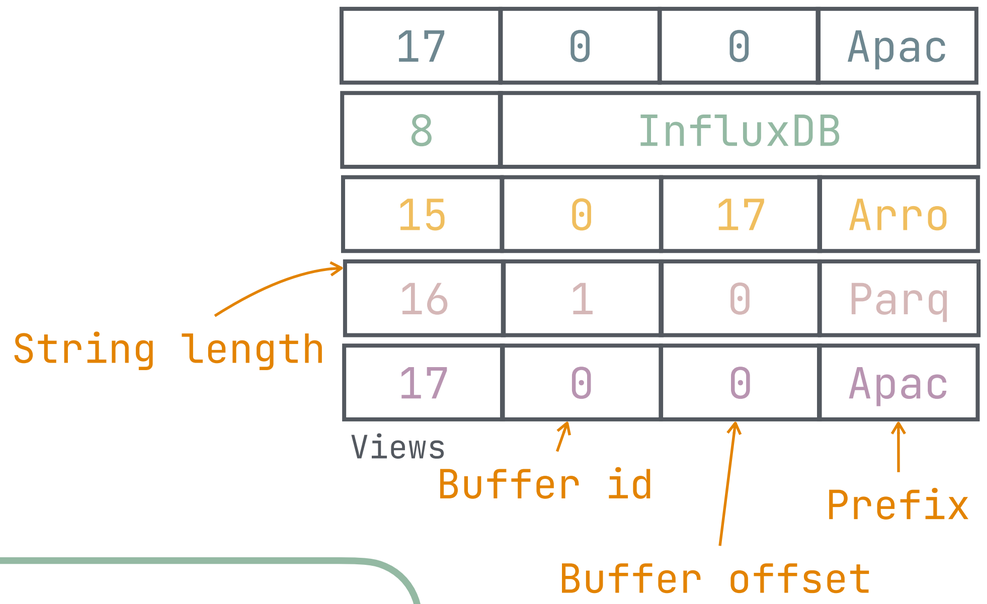

A	p	a	c	h	e		D
a	t	a	F	u	s	i	o
n	A	r	r	o	w		R
u	s	t		I	m	p	l

Buffer 0

P	a	r	q	u	e	t	
p	u	s	h	d	o	w	n

Buffer 1

Conventional: evict the entire array

LiquidEvict (structure-aware):

1. Evict buffers
2. Retain only prefix and str len
3. Retain only prefix

LiquidCache + LiquidEvict

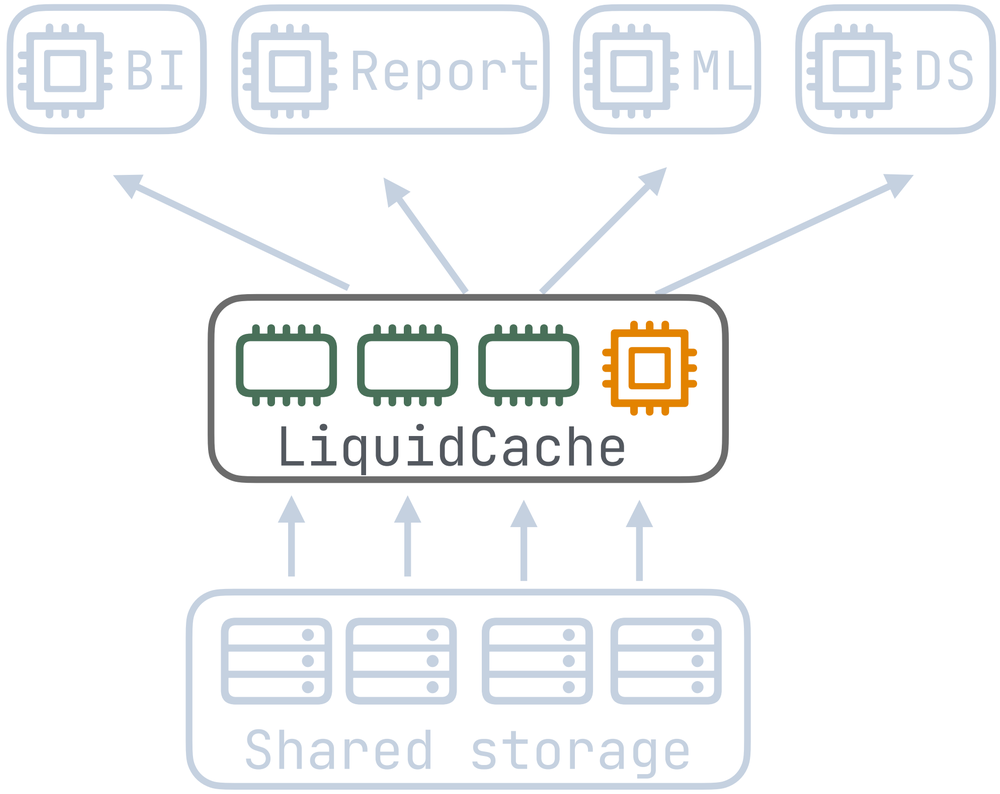

Data-aware cache

1. Pushdown to reduce traffic
2. Efficient decoding

Data-aware eviction

1. Evicts unimportant parts of data
2. Evicts data, keep summaries

Shared, pushdown cache system

10x lower latency, 10x lower cost

Timeline

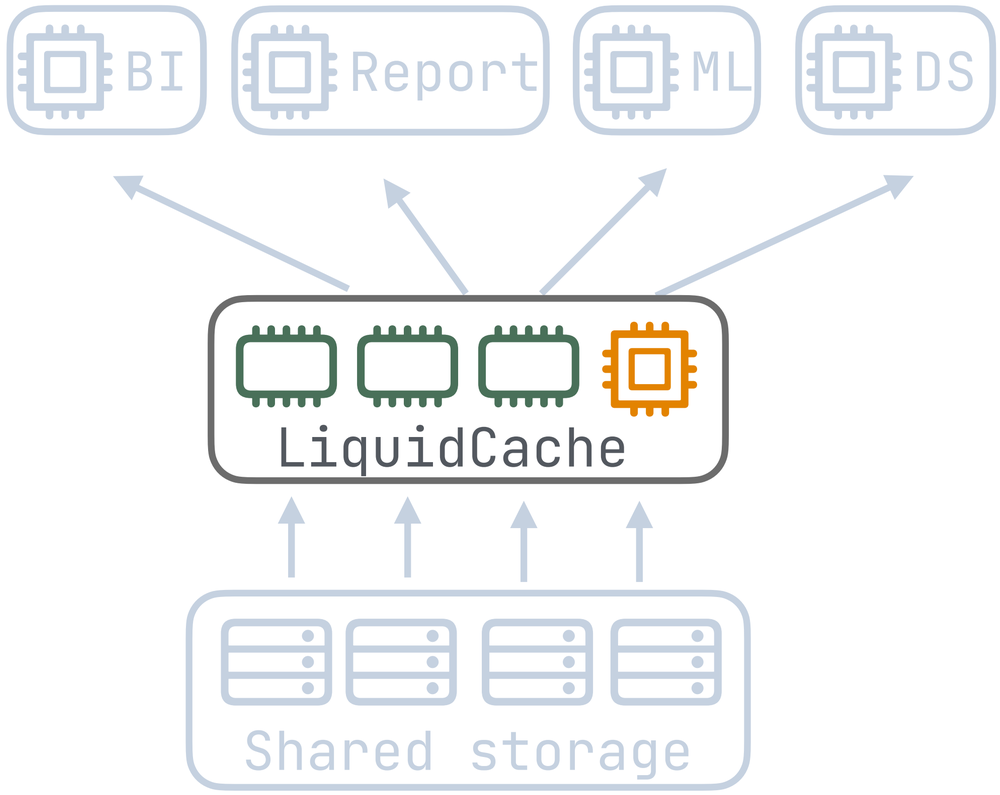

Jun. 1 - LiquidCache VLDB revision

July 1 - LiquidEvict prototype

Aug 1 - LiquidEvict bench & polish

Sep 1 - LiquidEvict VLDB submission

Dec 1 - LiquidEvict revision

Jan - Mar 2026 - Polish & present

May - Defense + graduate

Other first-author papers:

Bf-Tree: A Modern Read-Write-Optimized Concurrent Larger-Than-Memory Range Index.
Xiangpeng Hao, Badrish Chandramouli. (VLDB 2024)

Towards Buffer Management with Tiered Main Memory.
Xiangpeng Hao, Xinjing Zhou, Xiangyao Yu, Michael Stonebraker. (SIGMOD 2024)

https://github.com/XiangpengHao/liquid-cache

Huge thanks to InfluxData for supporting the project in 2024-25

influxdata®

Need your help:

1. Make LiquidEvict real
2. Make LiquidCache ready for your company
3. Advance science and serve the public good

$50k charitable giving supports 1 PhD/year

20% of a FTE cost!

Better than most chatbots!