



**L-STORE DATABASE - MILESTONE 2**  
**ECS 165A**  
**TEAM: CAN'T-STORE**

HANSON LAU, CHRISTOPHER PHAN, IVAN CVJETINOVIC, DANIEL MEDINA, NAMAN SHETH

# TOPICS

Durability and Bufferpool

Data Reorg

Indexing

# Durability and Bufferpool

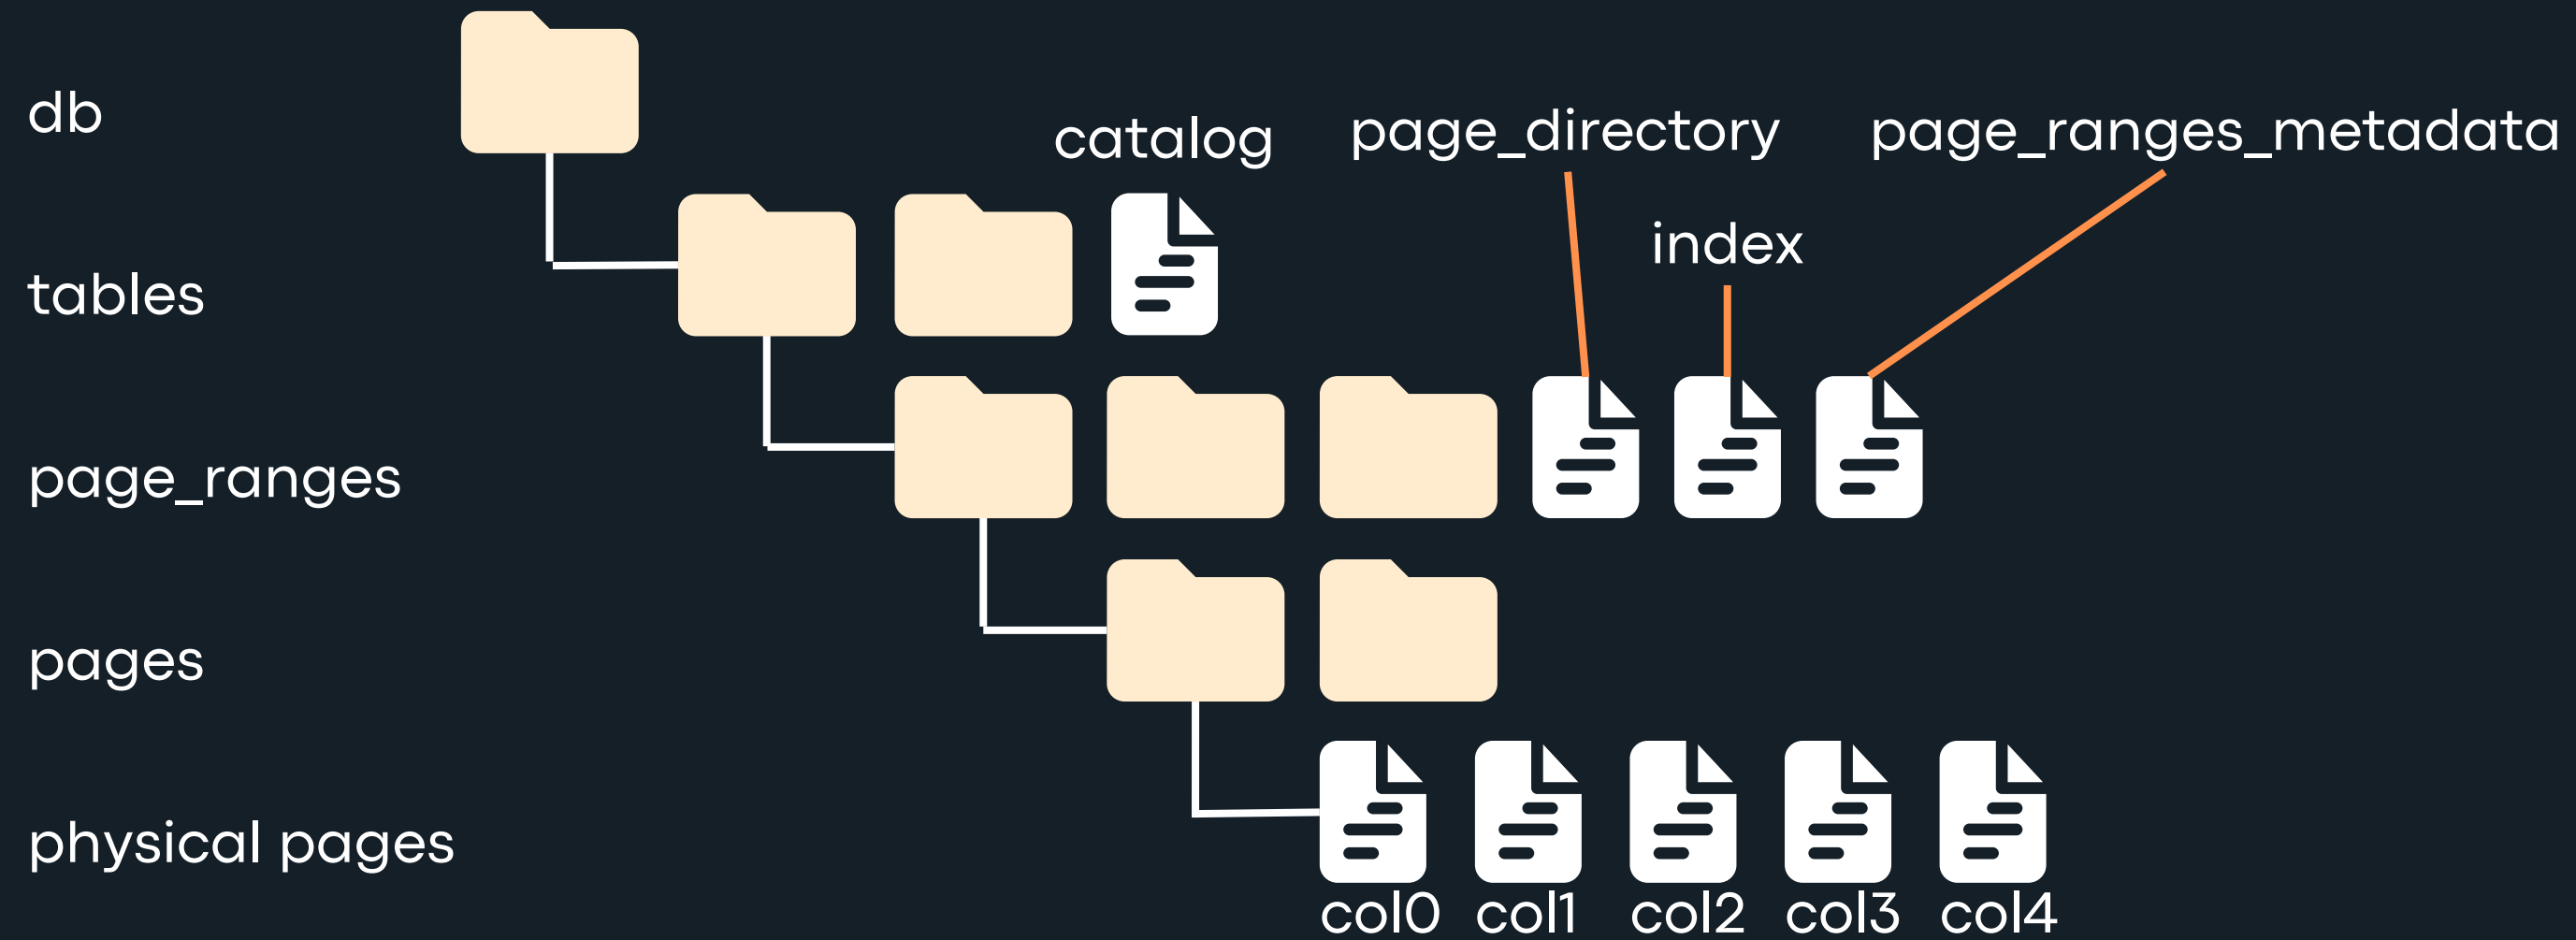
File Structure

Bufferpool

Eviction Strategy

Retrieving Pages

# File Structure



Path: <db\_name>/<table\_name>/page\_range<#>/page<#>/col<#>

# Helper Files



## catalog:

Holds "<table\_name, num\_columns, key\_column, num\_records>"

e.g.:

```
Grades 5 0 6000  
Students 4 0 991  
Staff 6 0 34
```



## page\_directory\_and\_index:

persisted using pickle

e.g.:

```
0 5 18  
1 15 16  
2 4 40
```



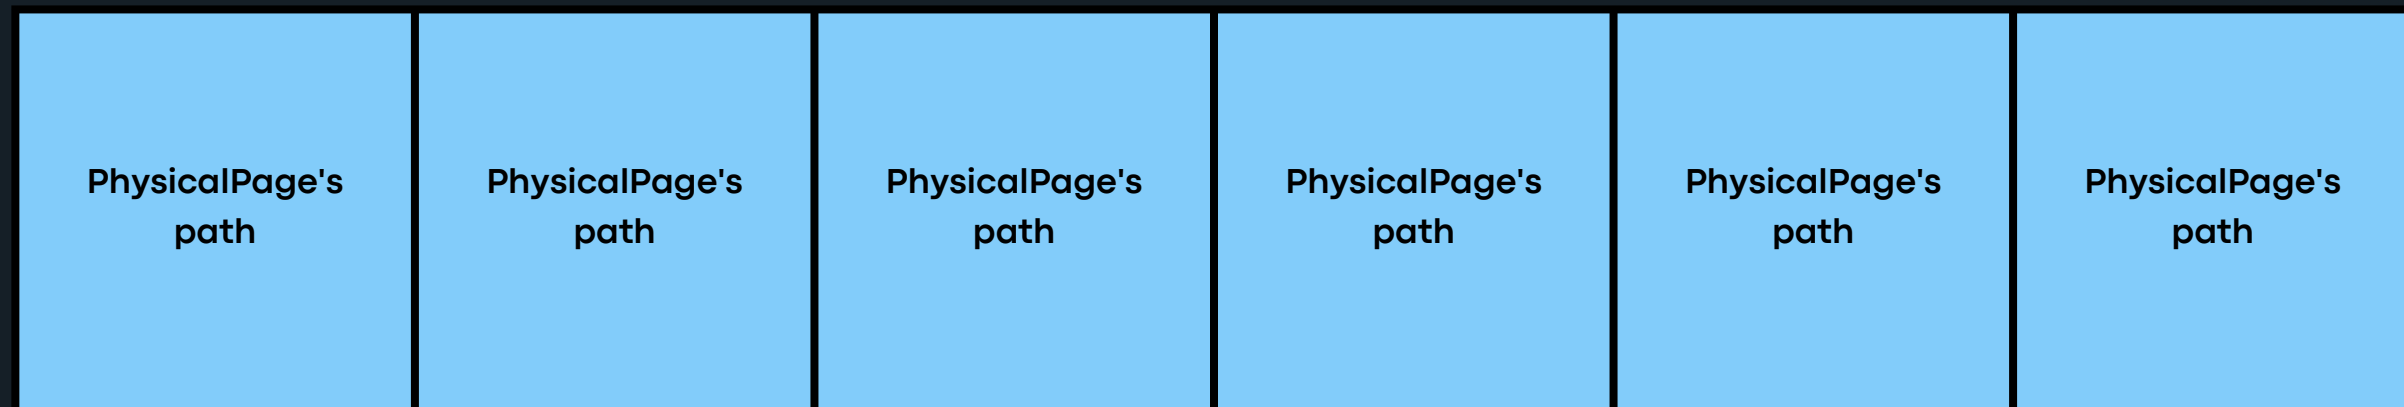
## page\_ranges\_metadata:

Holds "<page\_range\_num> <latest\_bp\_num> <latest\_tp\_num>"



# Bufferpool: self.lru\_pages

Least recently used

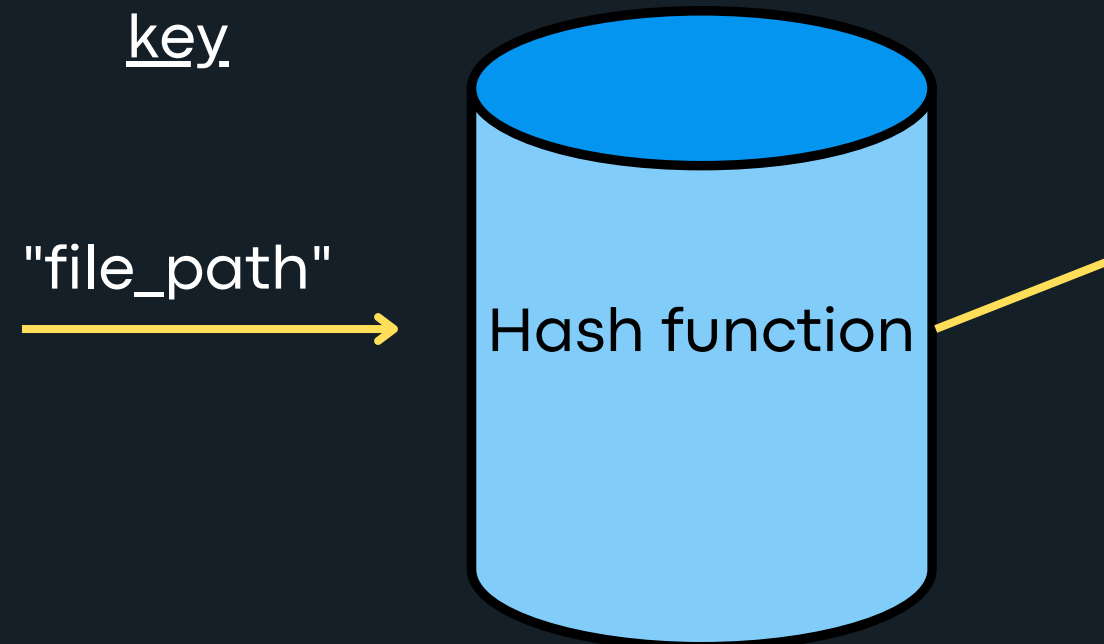


Most recently used

Implemented using a queue data structure.

The most recently used pages are at the end of the queue and the least recently used are at the beginning.

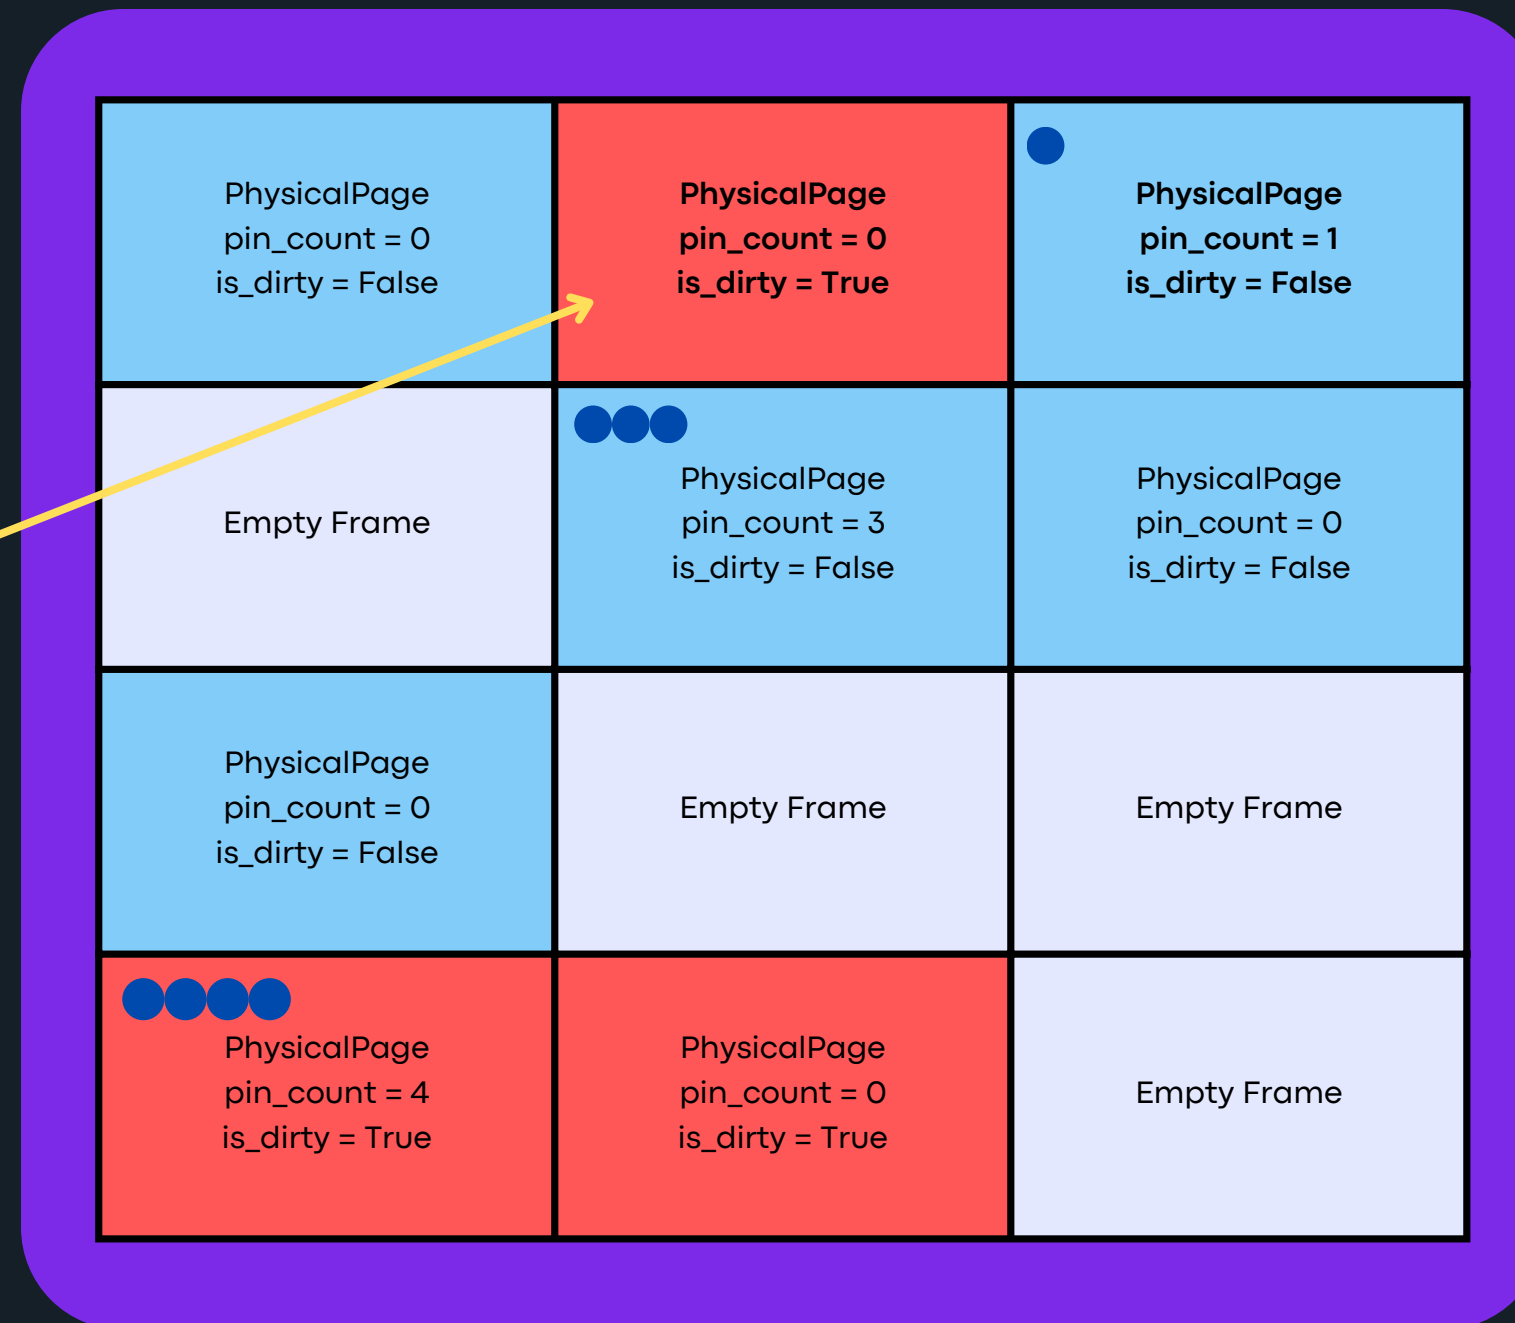
# Bufferpool: self.pages



Holds a configurable number of PhysicalPage objects

Implemented using a hash map

- key: string of path to file
- value: PhysicalPage object

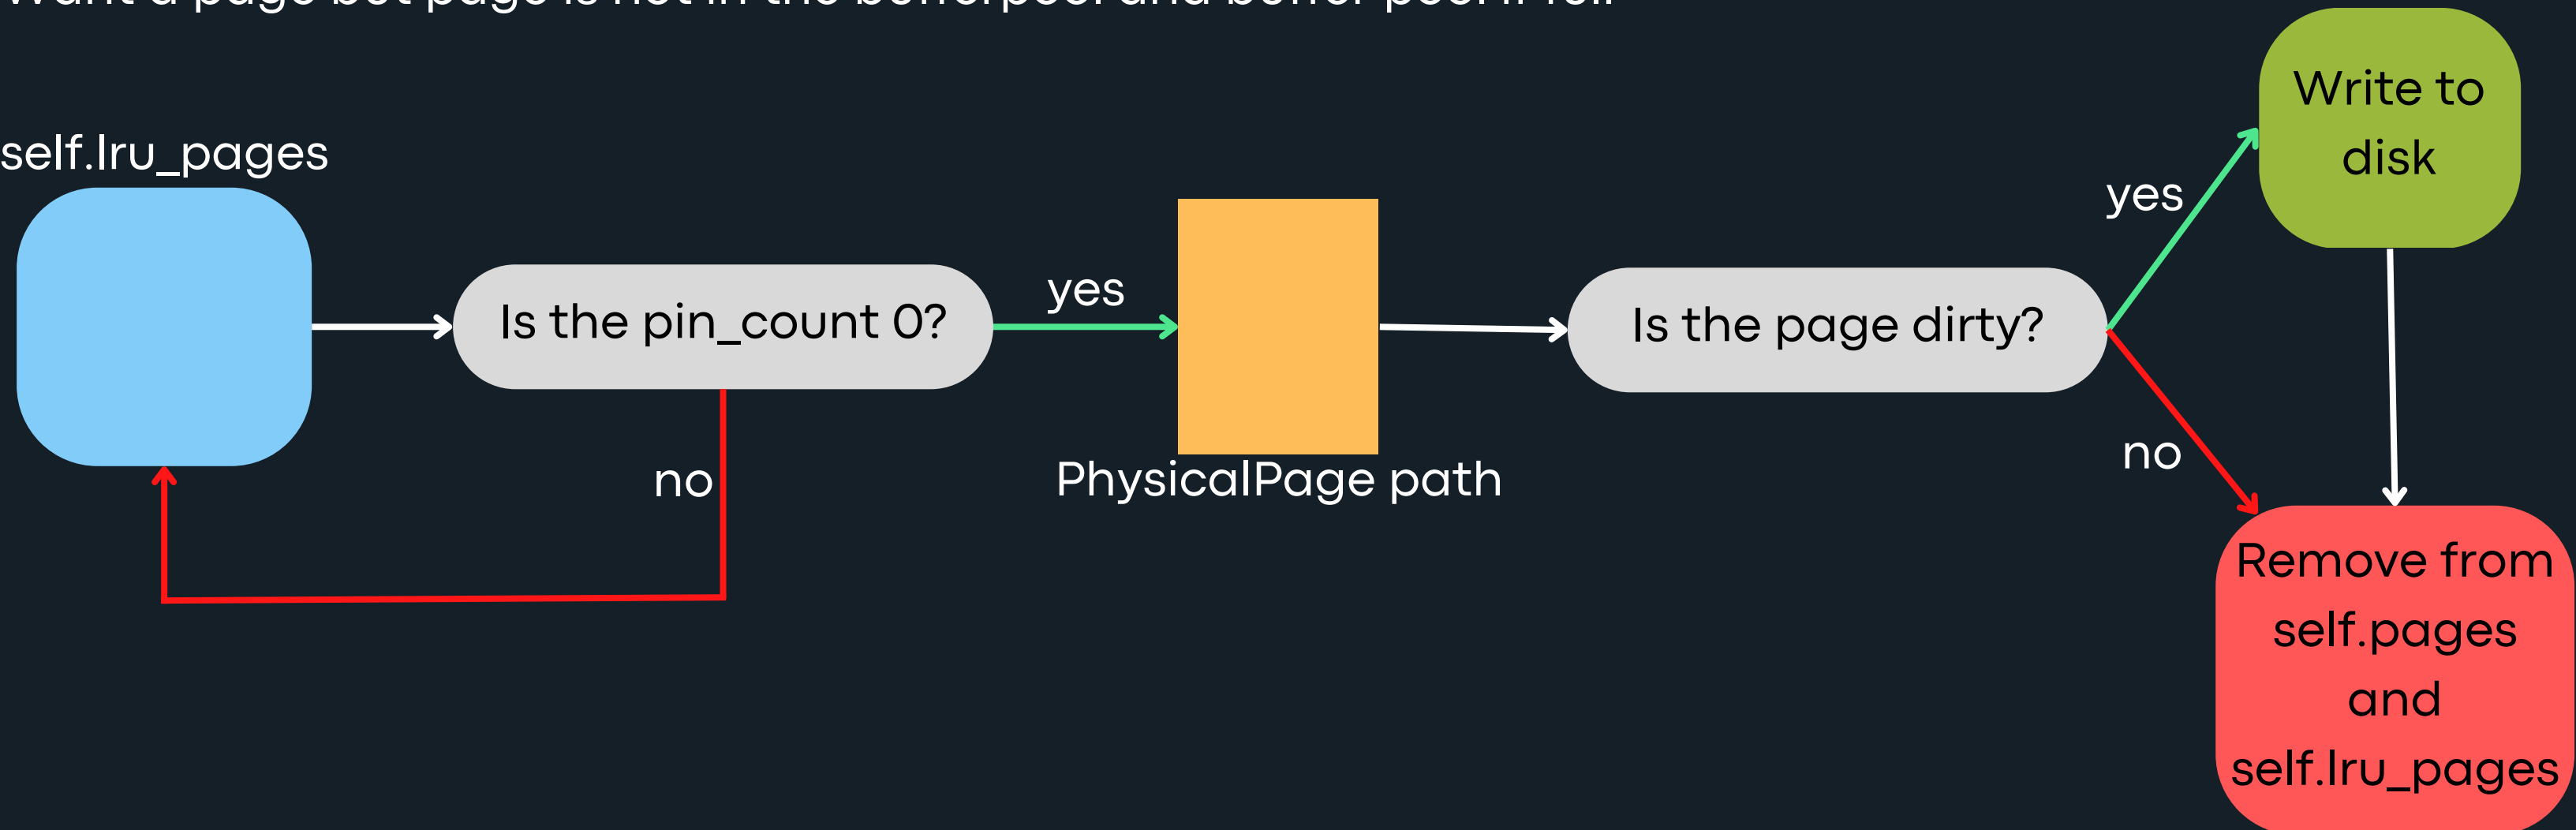


# Bufferpool: Evicting a Page

Situation:

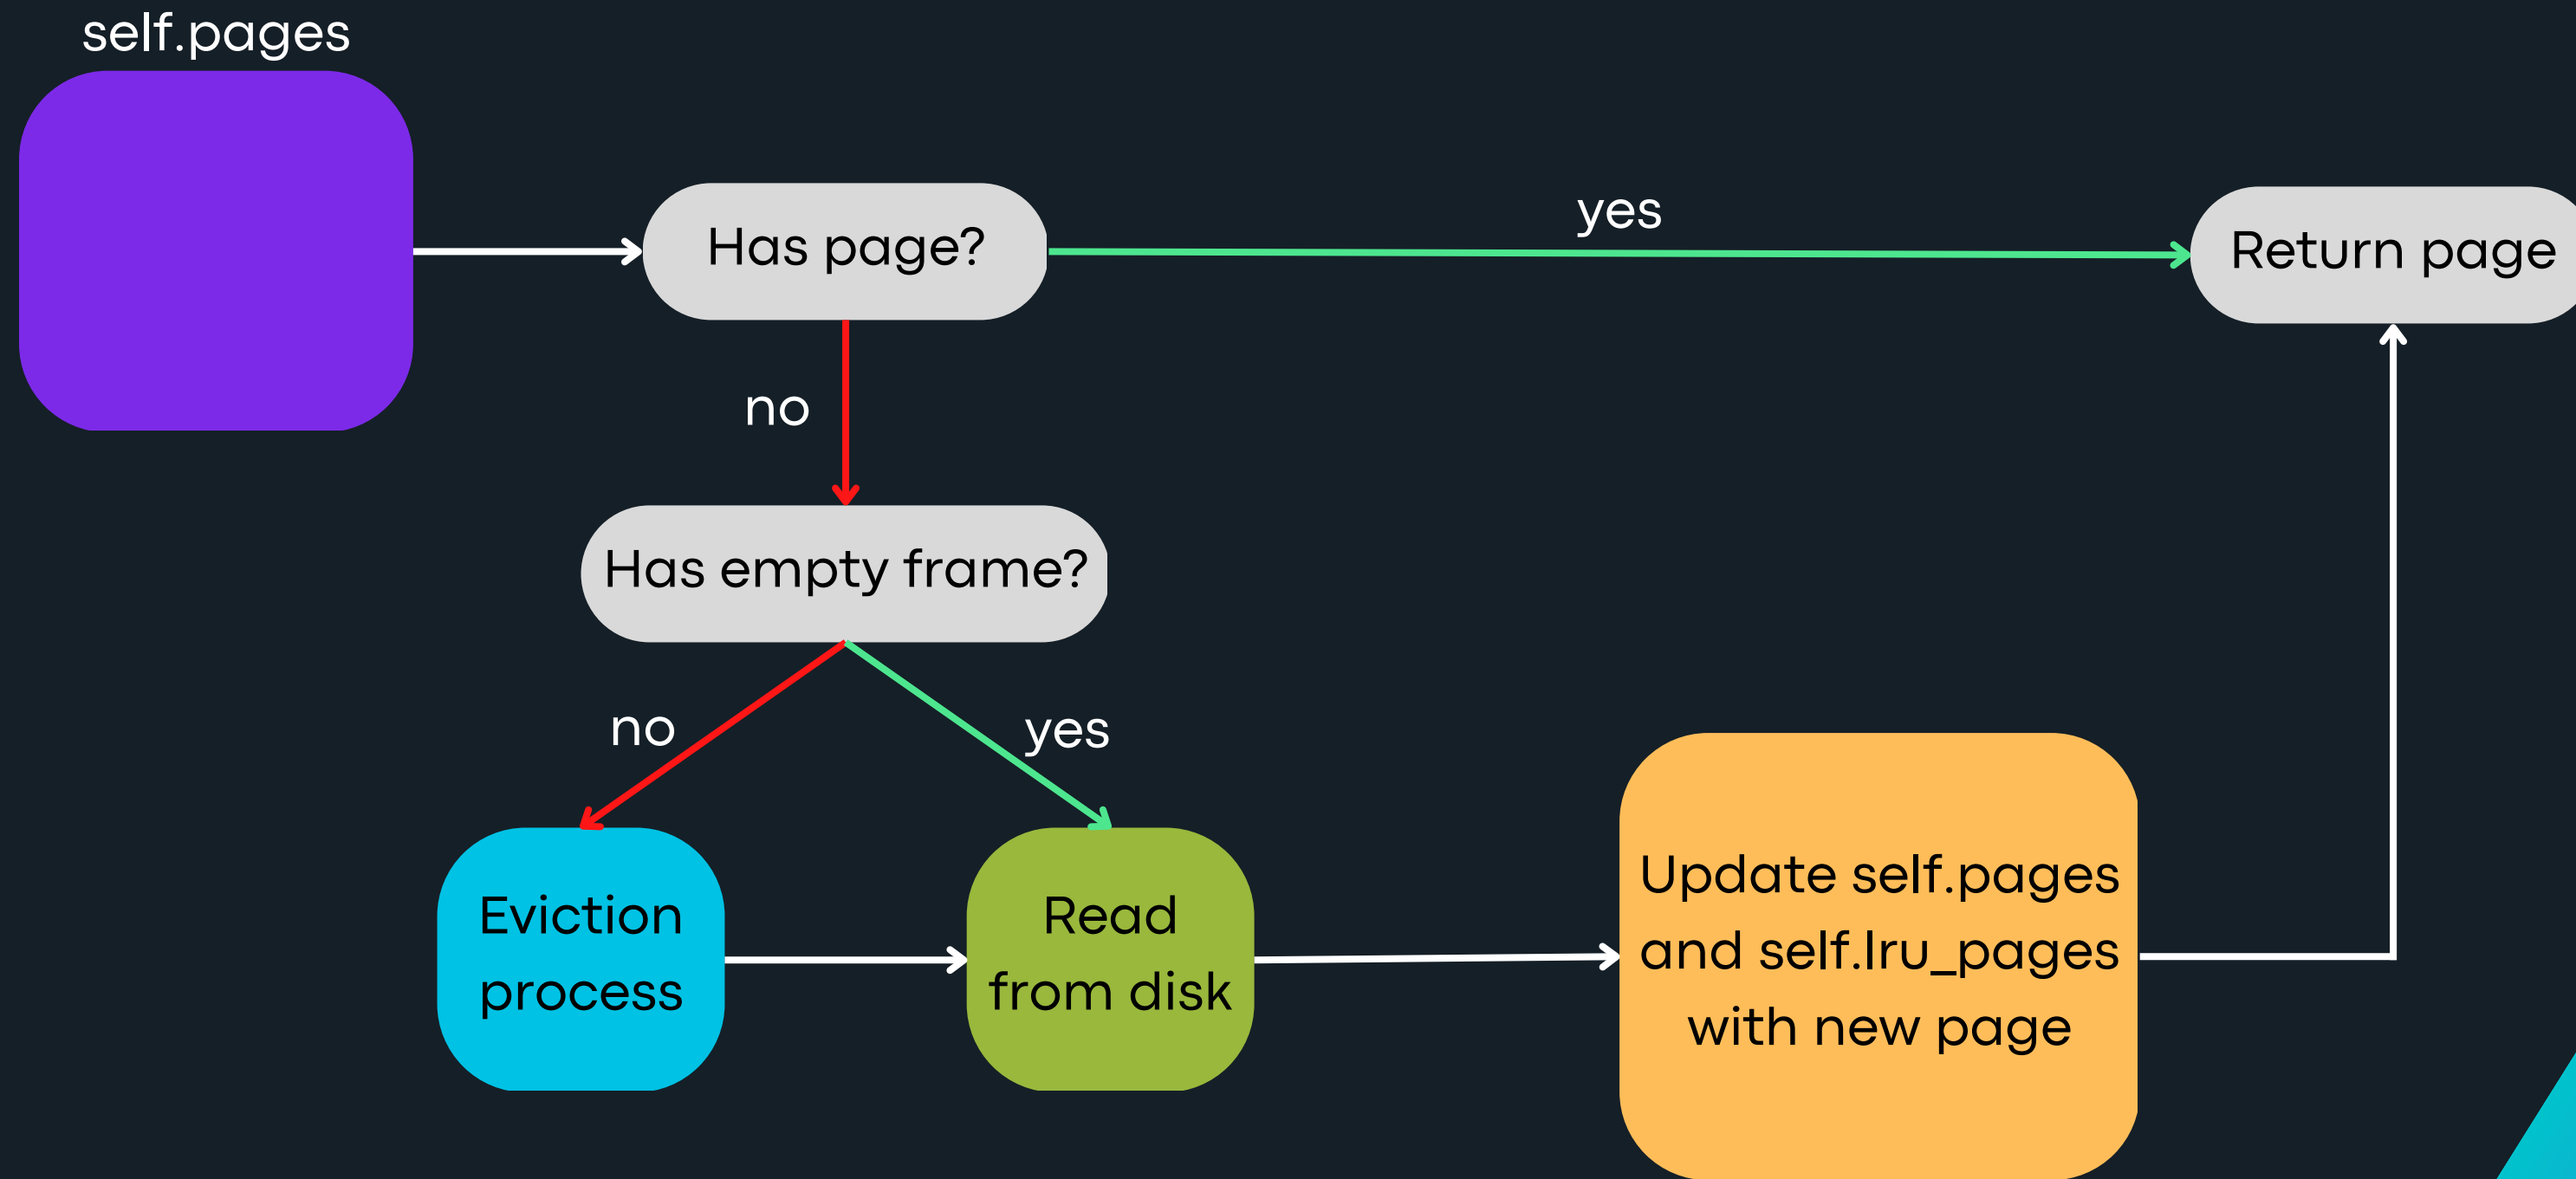
Want a page but page is not in the bufferpool and buffer pool is full

self.lru\_pages





# Bufferpool: Retrieving a Page

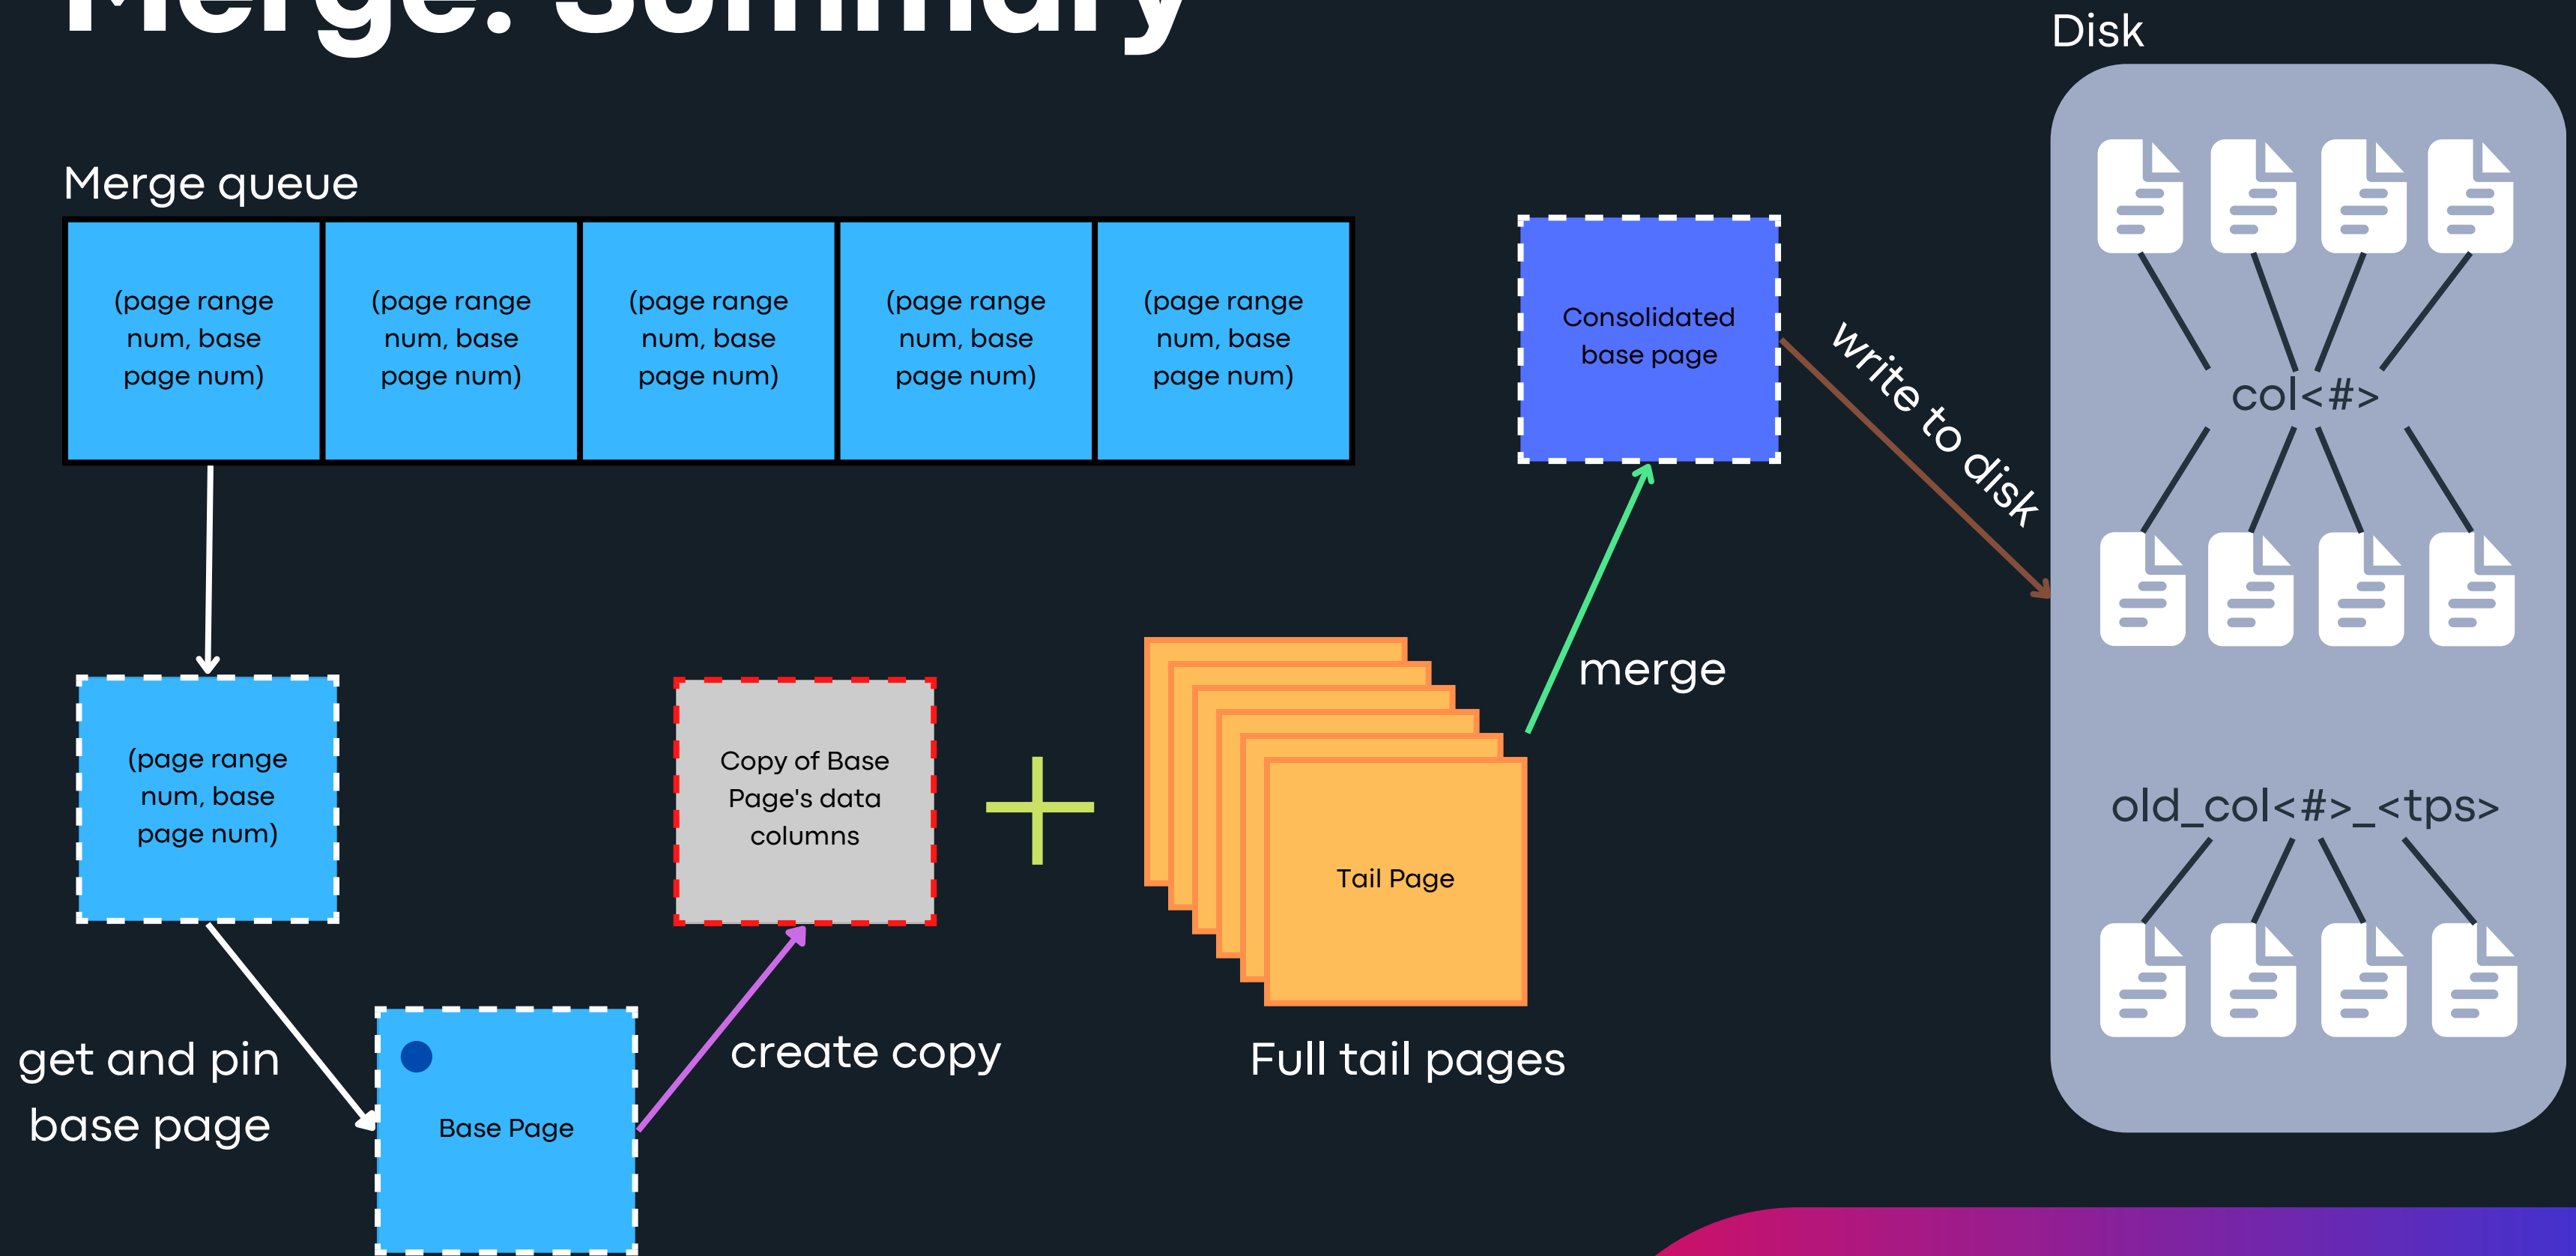


# Data Reorg

Merging

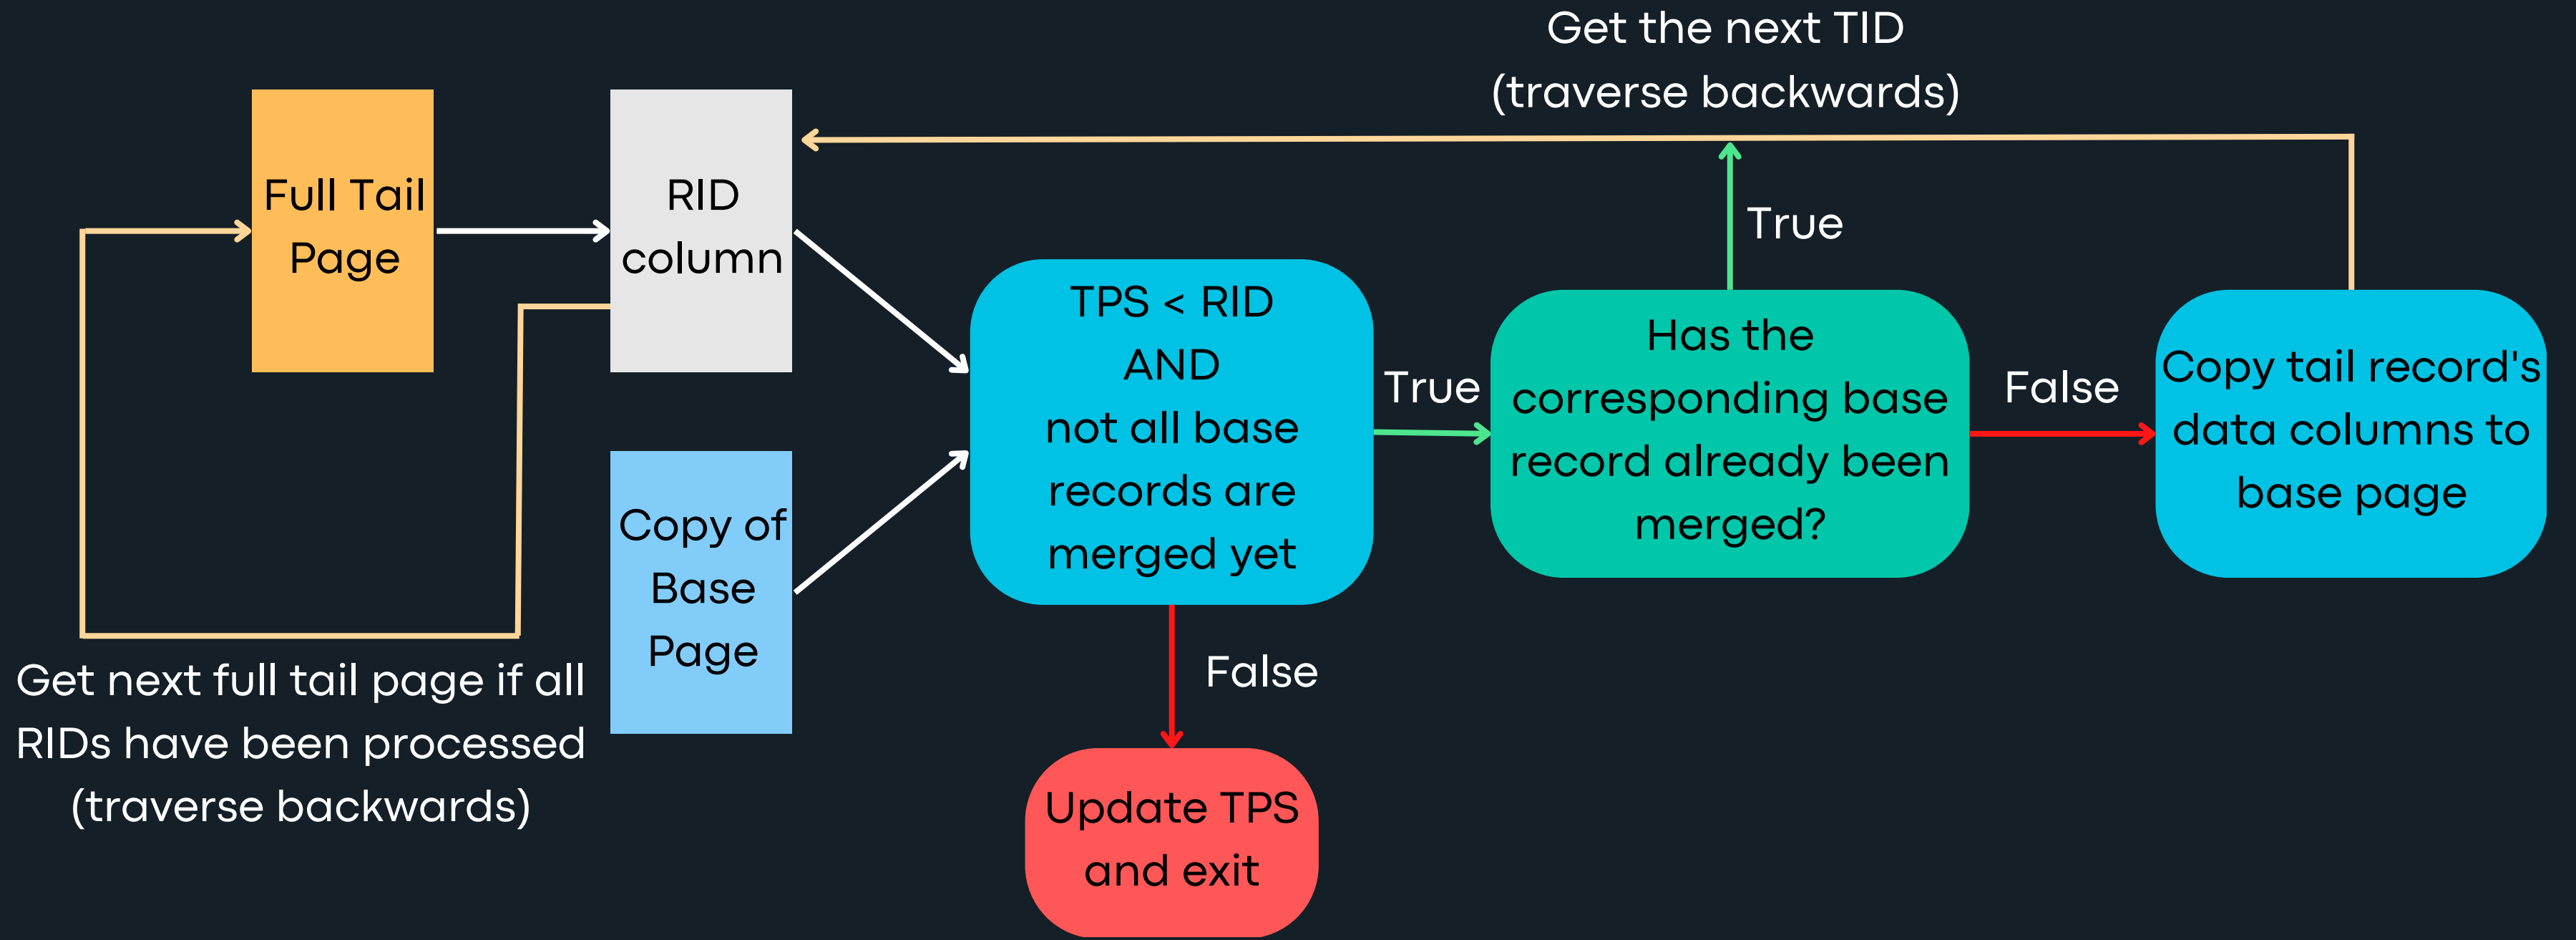
Result of Merging

# Merge: Summary

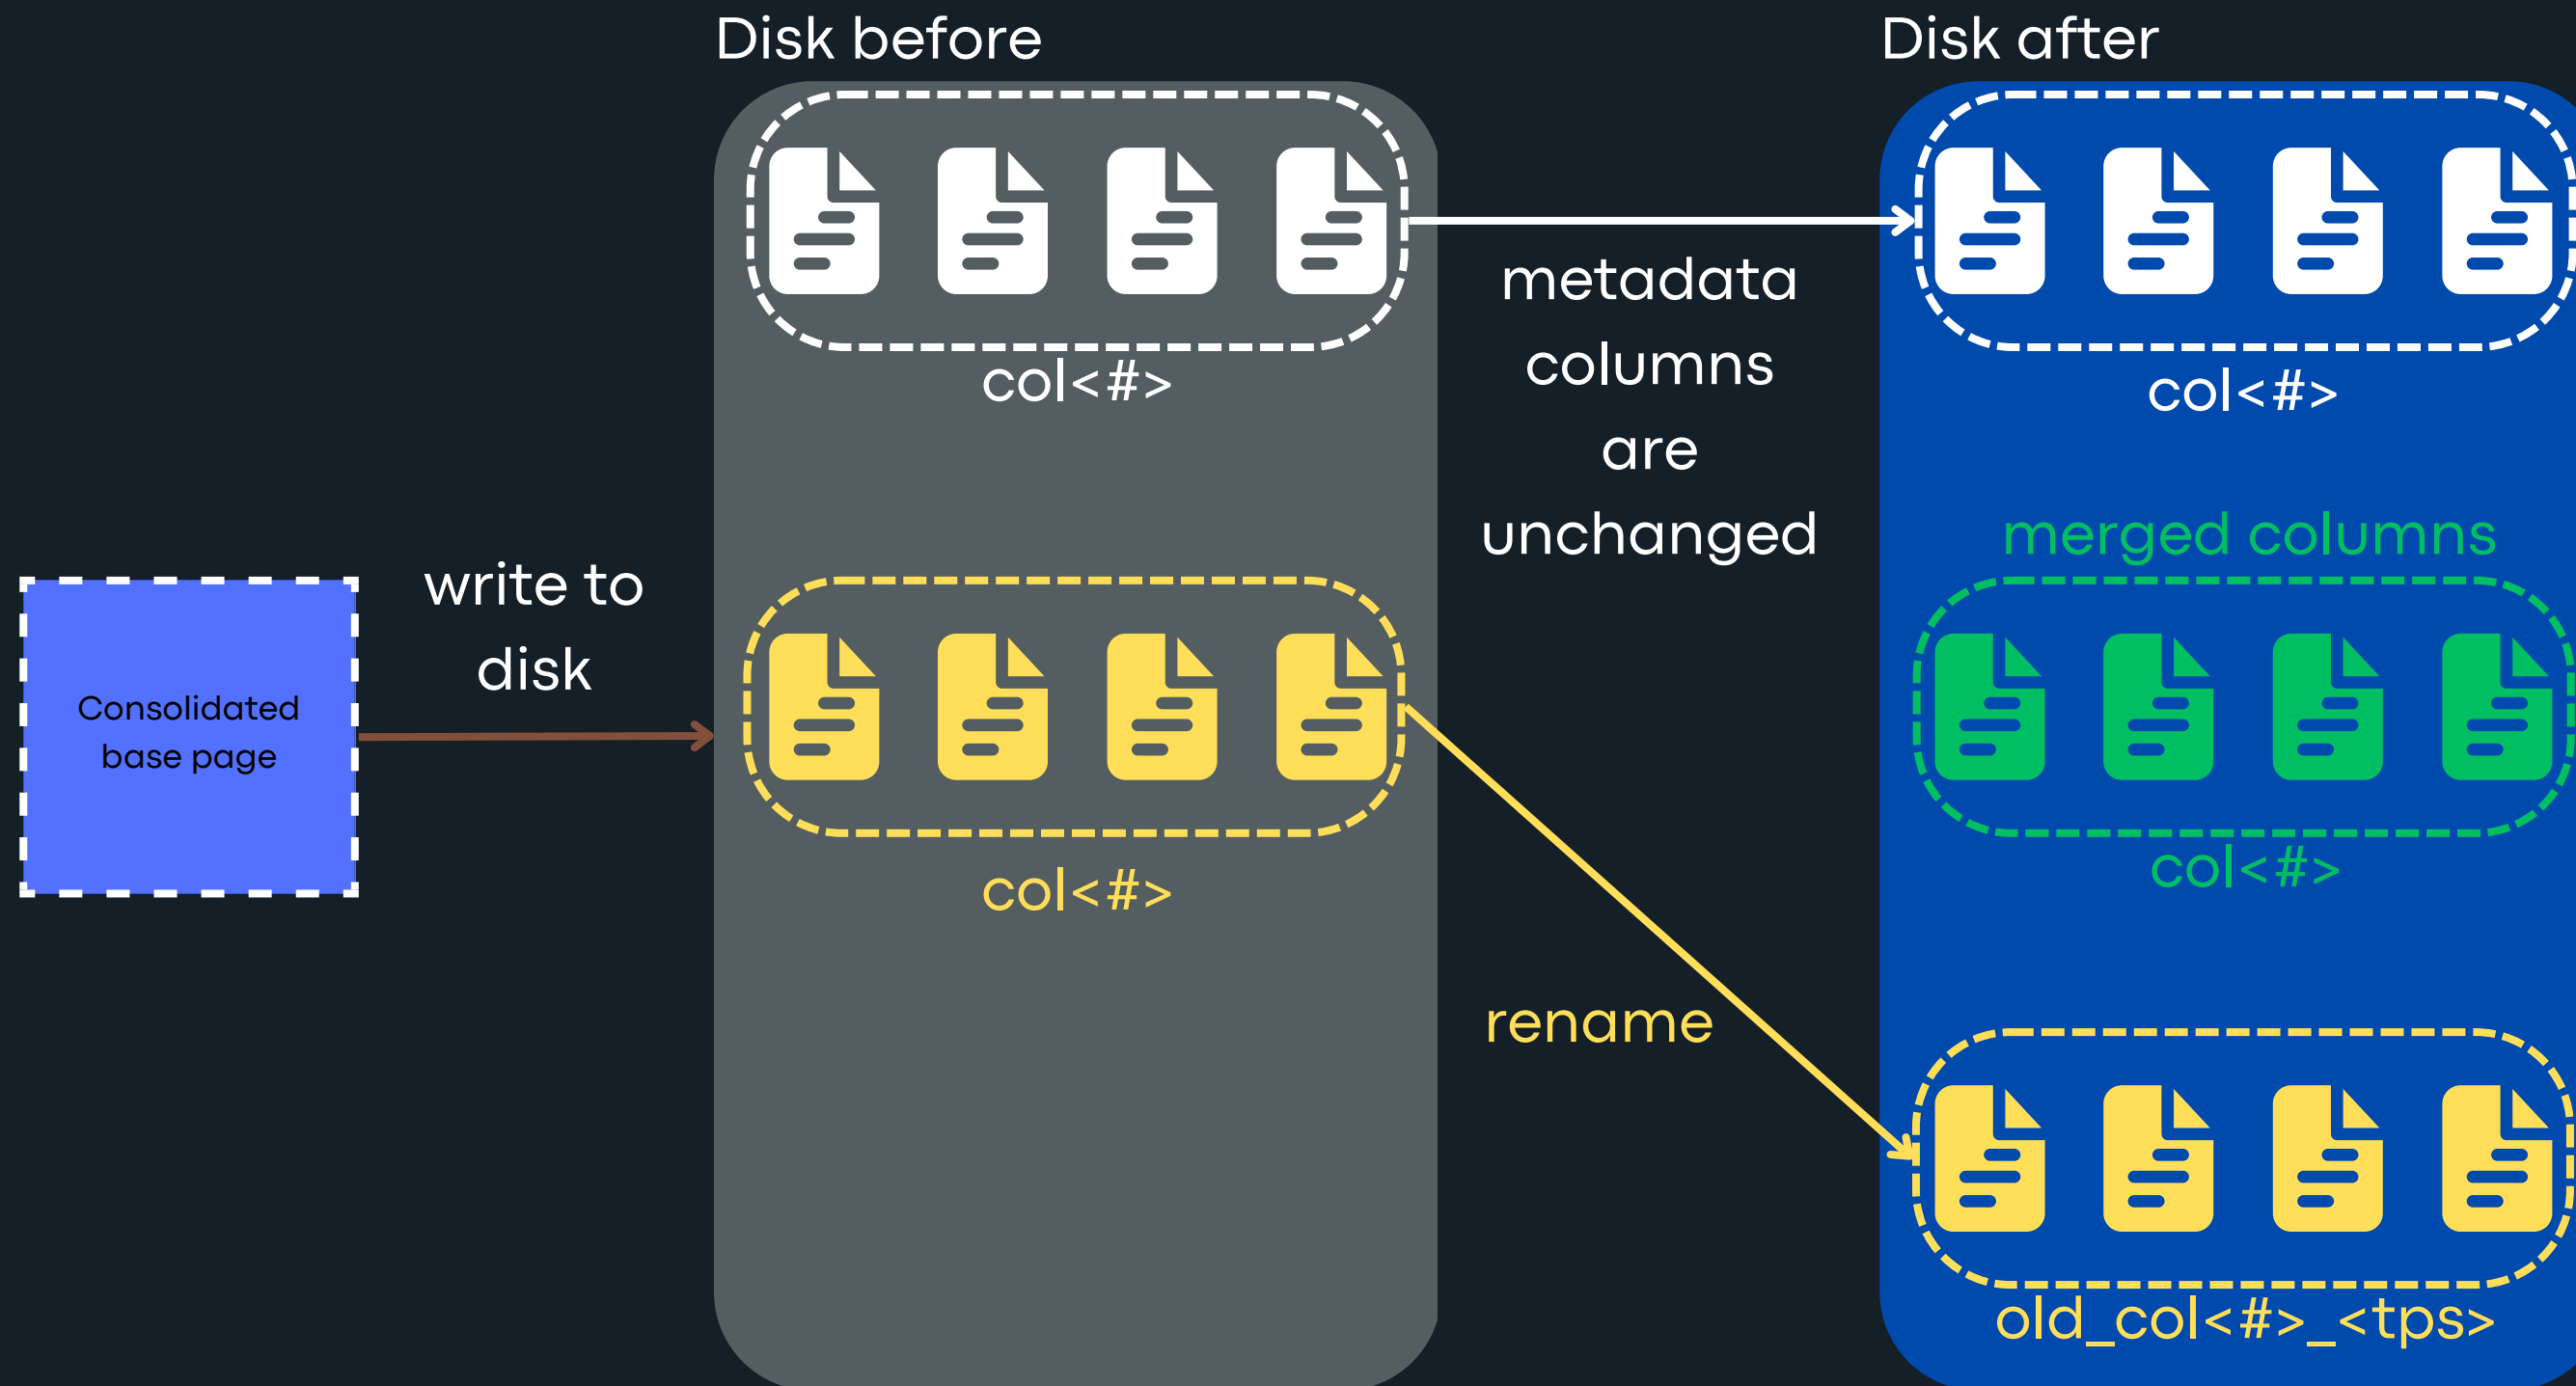


# Merging

By default, merge happens after 1024 updates on a full base page



# Result of Merging



Files are renamed instead of updating the page directory

The same purpose of rerouting all RIDs to the consolidated base page is achieved since if the old base pages are being used, they will still be in memory and every new transaction will use the consolidated base pages.

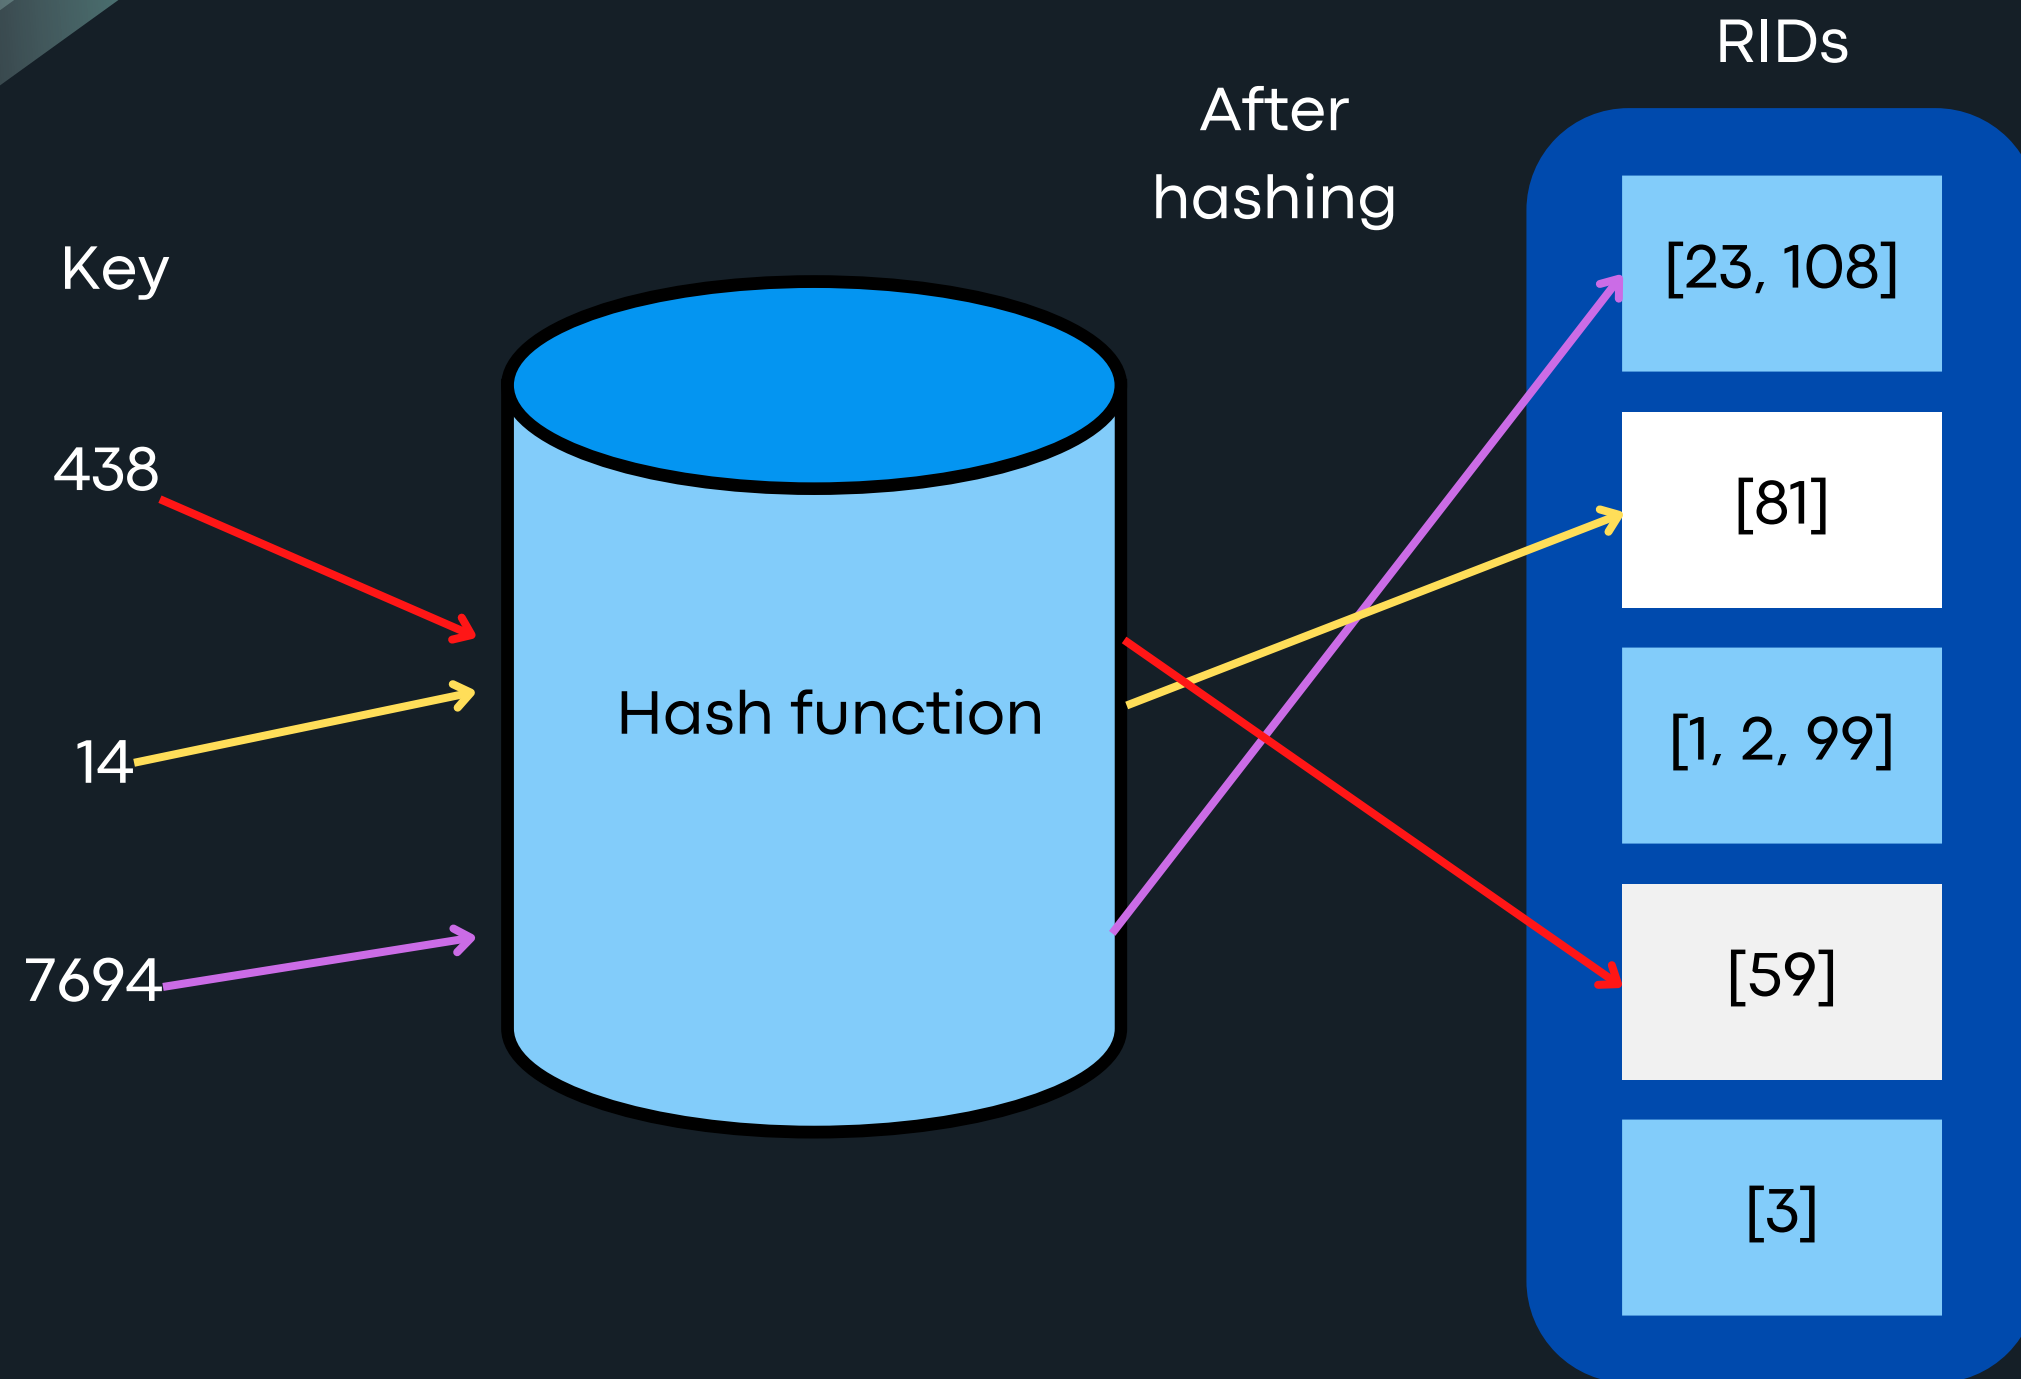
# Index

Index Structure

Creating Indices



# Index



Given a key, the index returns a list of RIDs of base pages whose latest record contains that key

The primary key column is indexed by default. Secondary indices can be created and will need to be maintained.

Limitation:

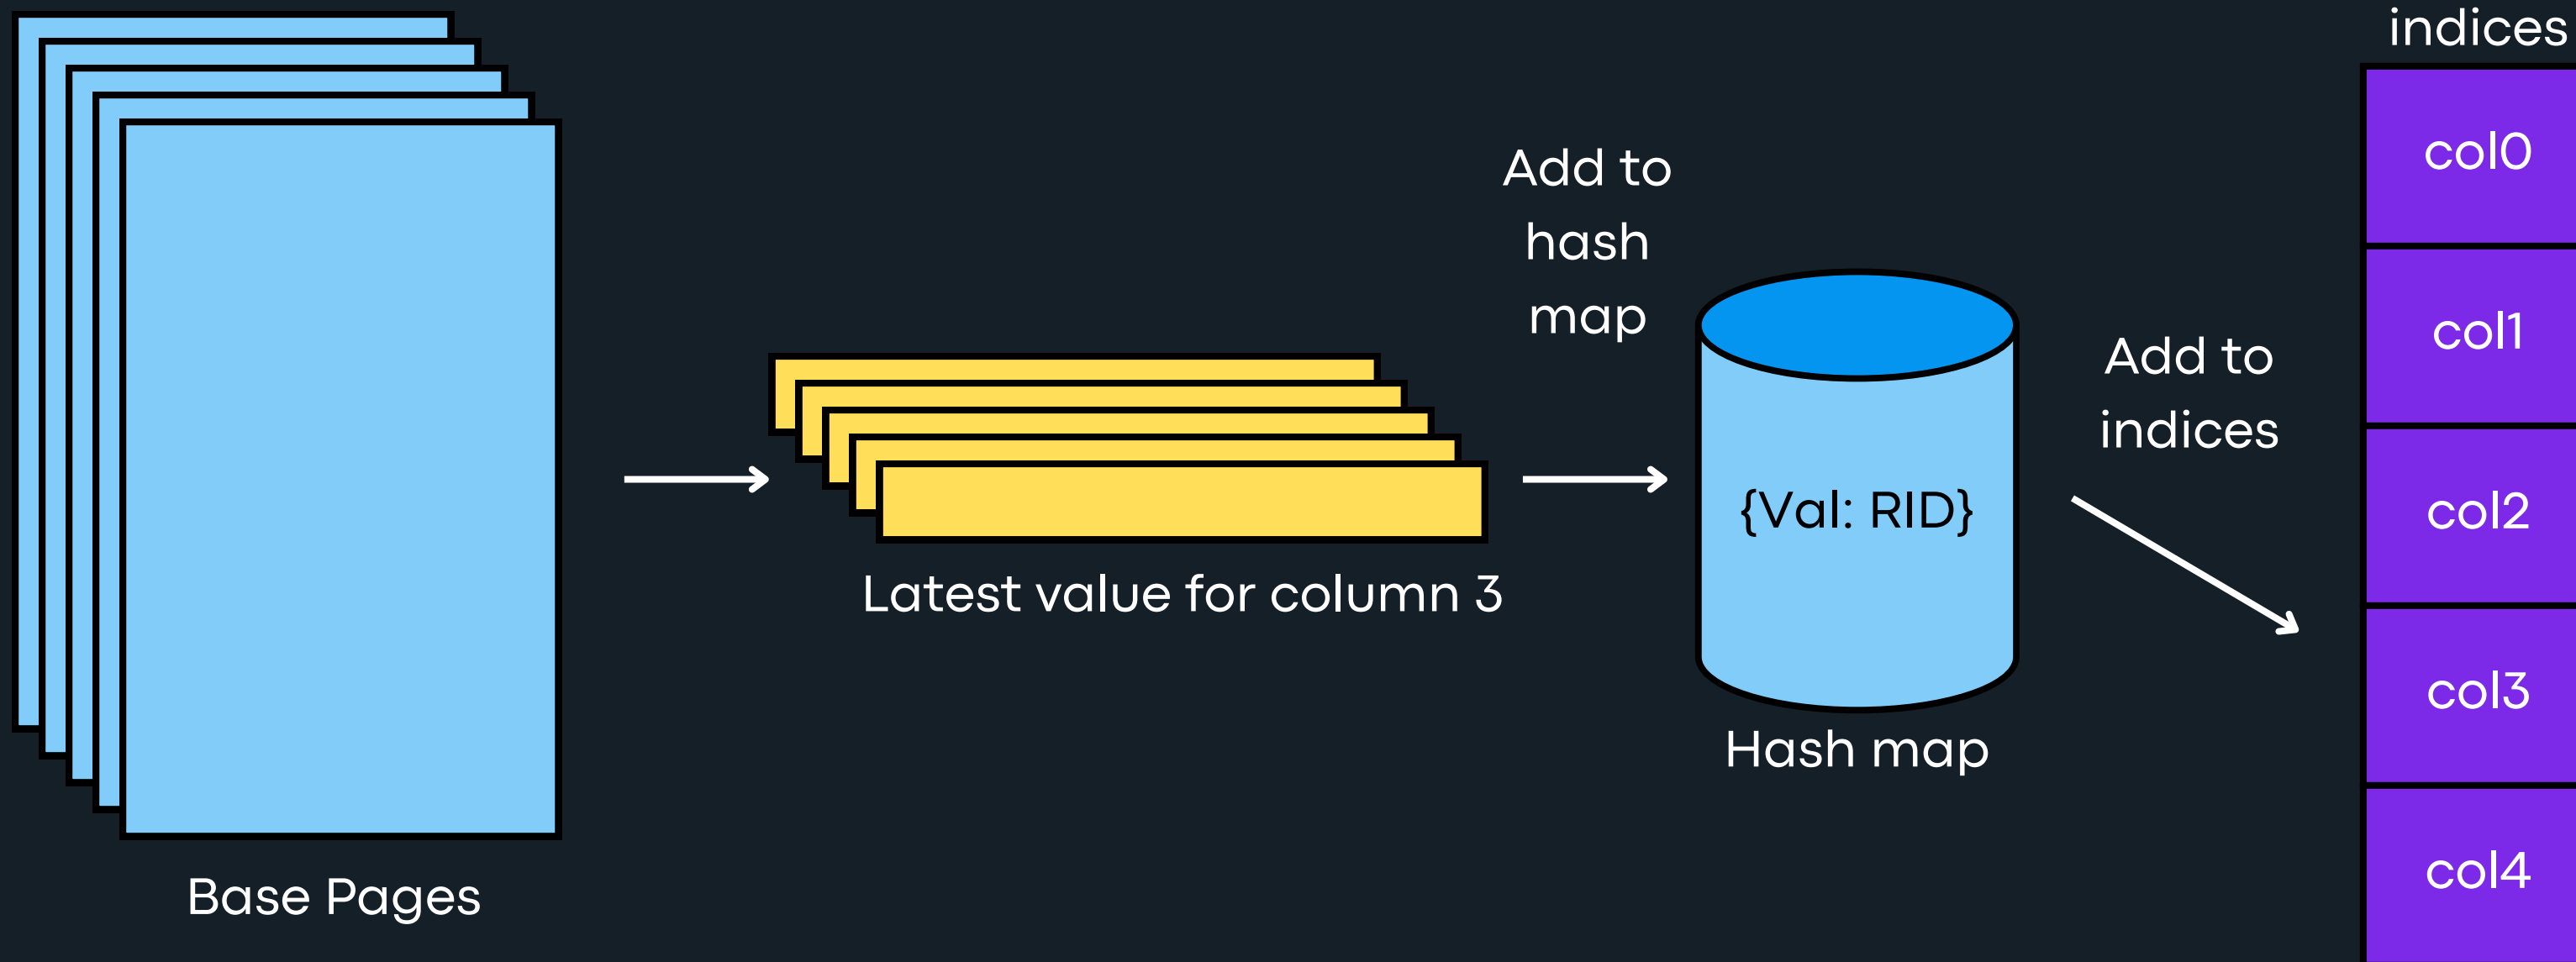
Optimized for equality

Needs multiple accesses for a range query

Hash map allows for  $O(1)$  access time

# Creating Indices

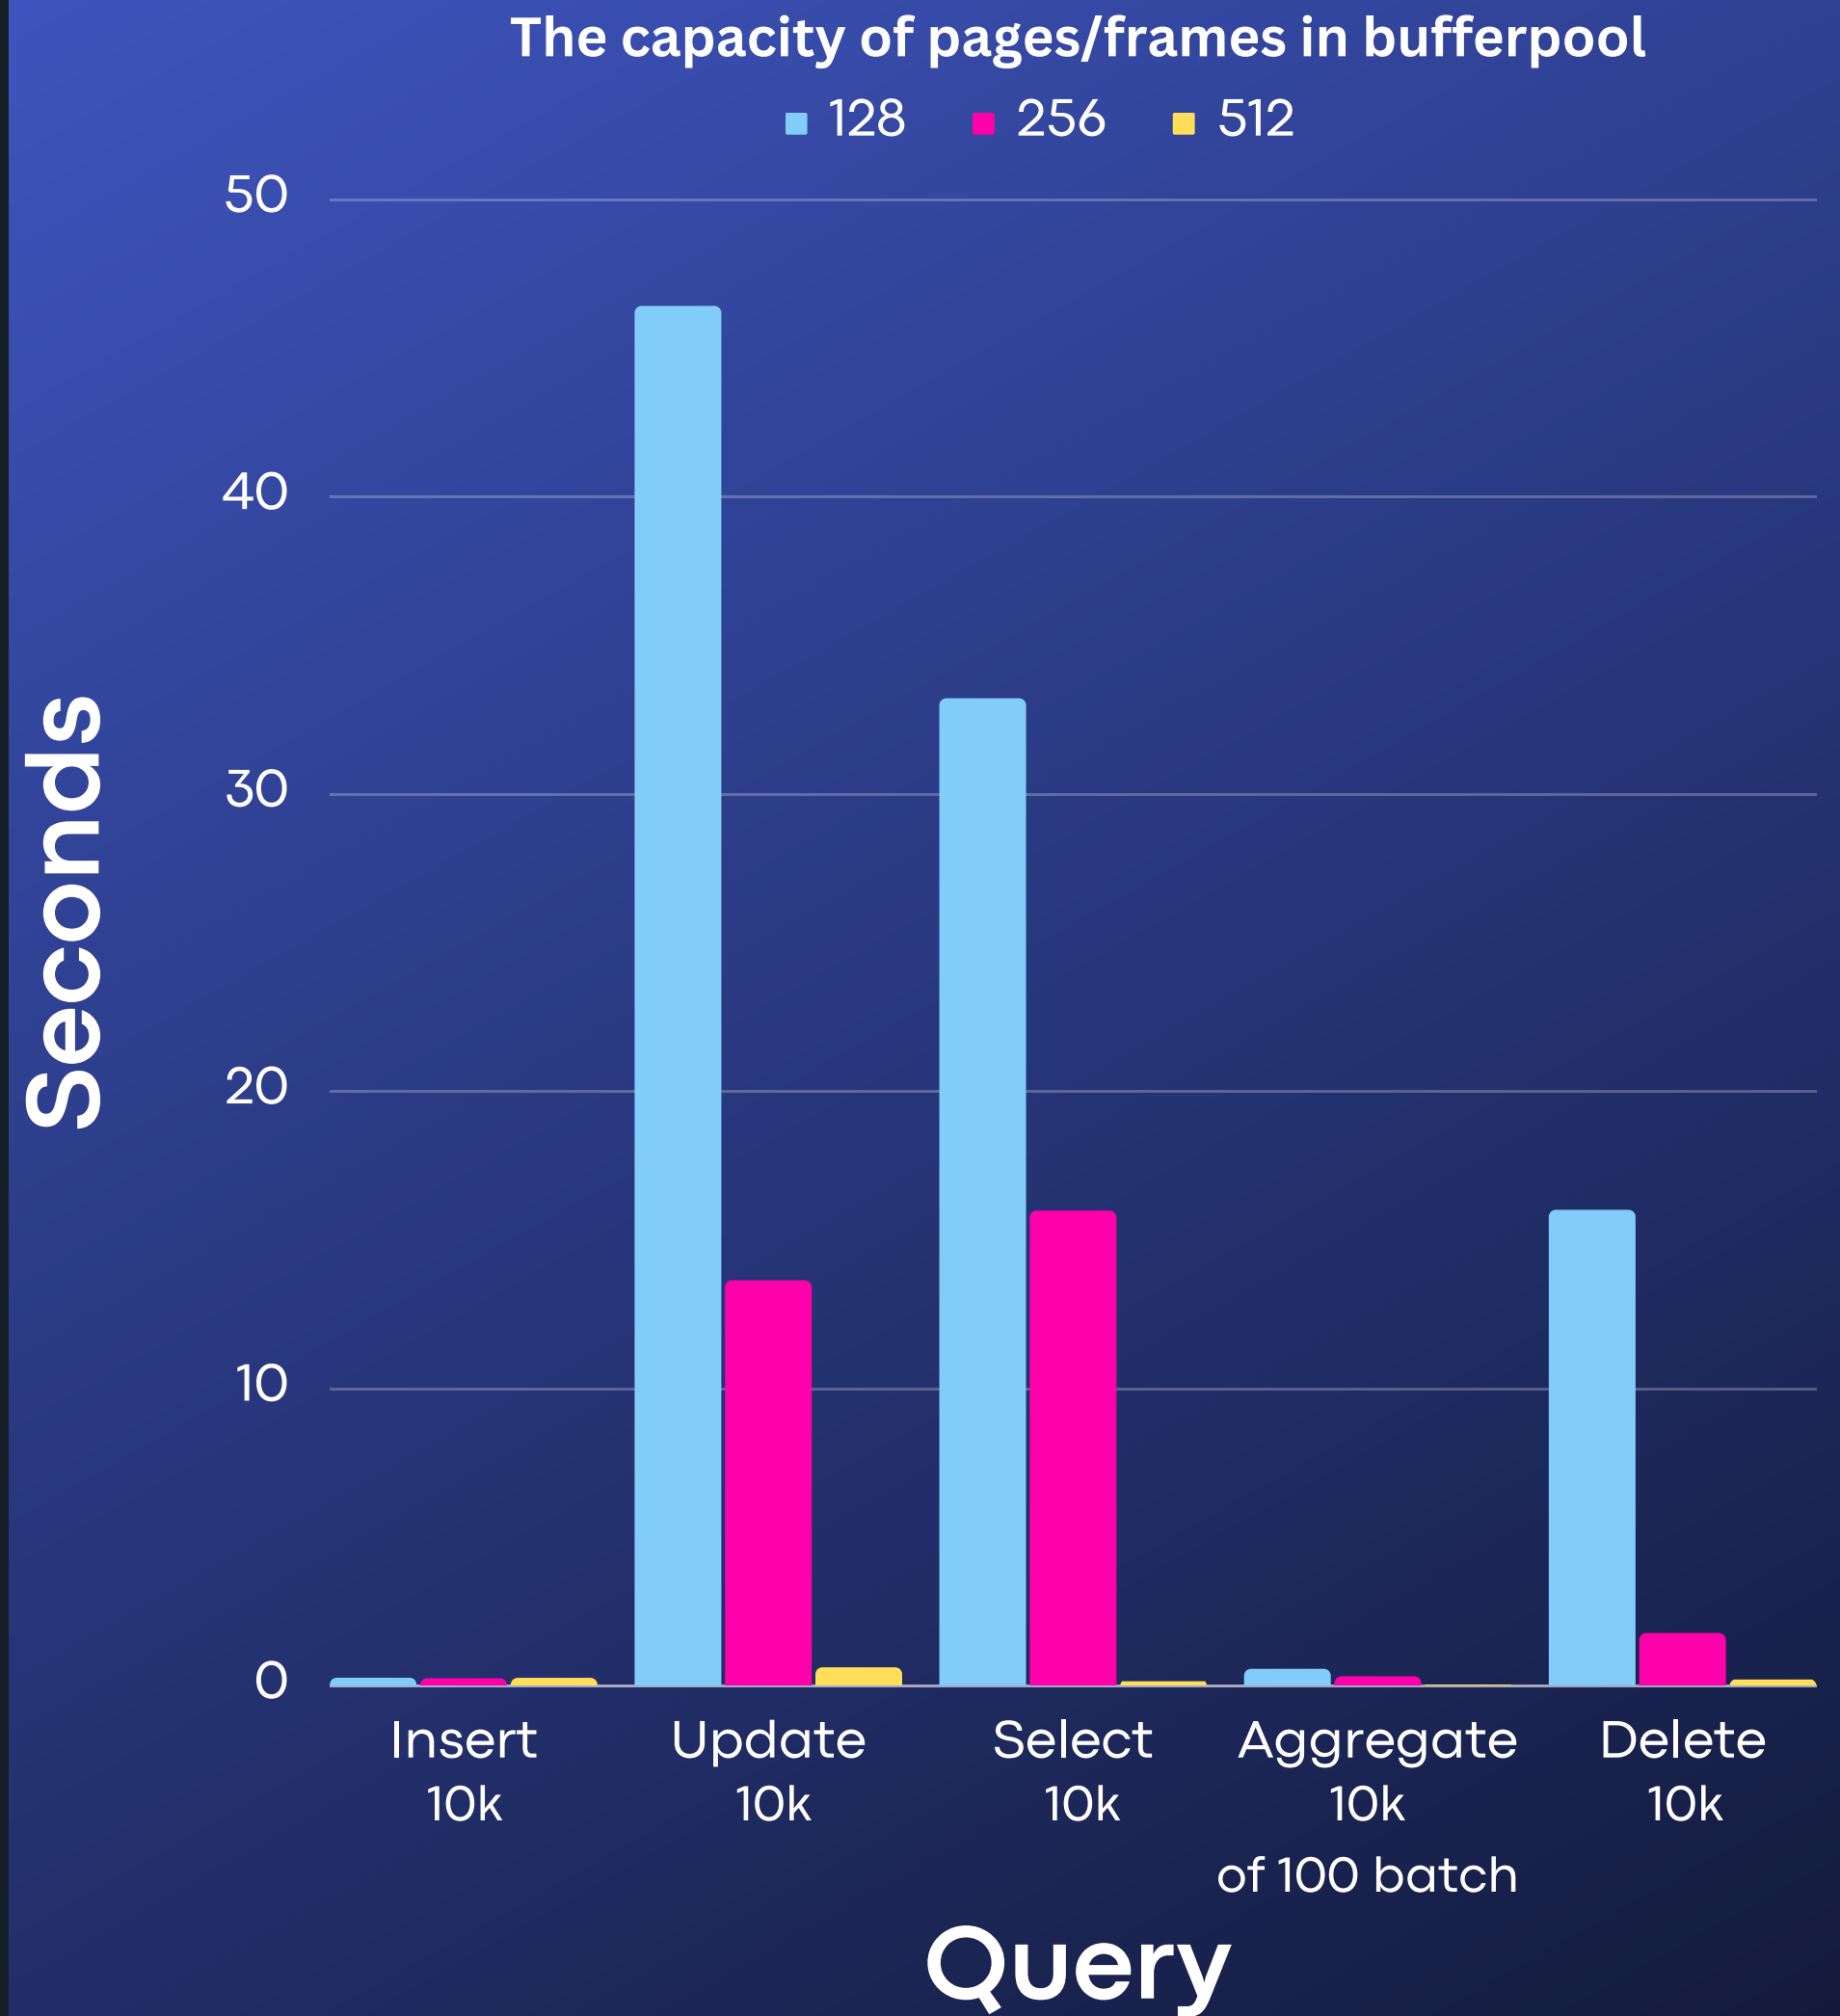
Creating an index for column 3



# QUERY RUNTIMES (\_\_MAIN\_\_.PY)

## Specs

Mac: M1 pro 16 GB RAM



# Questions