L-Store Milestone 2

Alana Rufer, Eseosa Omorogieva, Nina Gopal,
Riddhi Barbhaiya, Kushaal Rao

Review of logical memory organization

RID → [page_range_offset, page_index, offset]

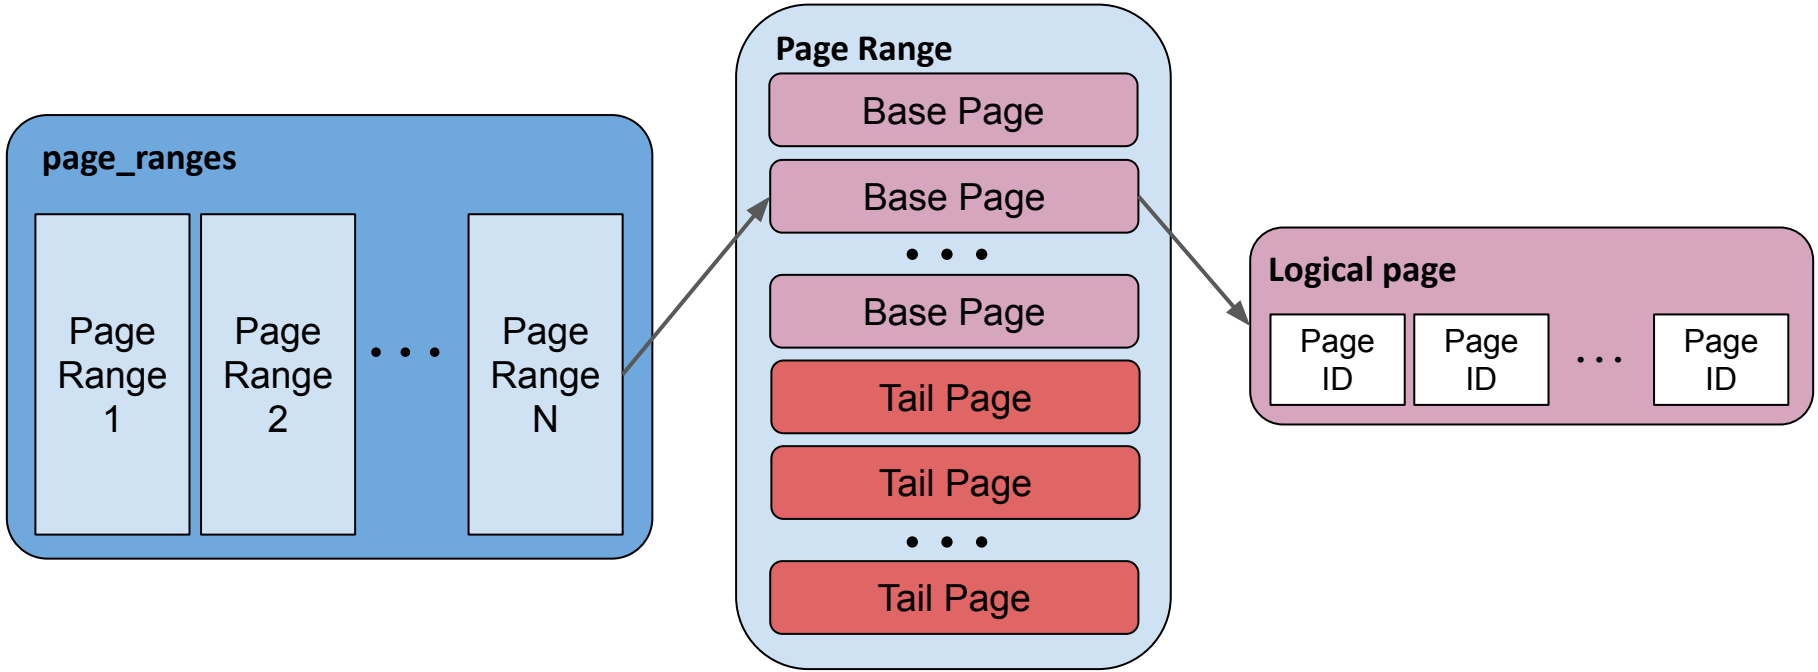

Indexing

Index Structure

Array of dictionaries, with the index of each dictionary corresponding to the column number

Columns without indexes have a None placeholder

Value are mapped to a list of RIDs of records that contain the value in that column

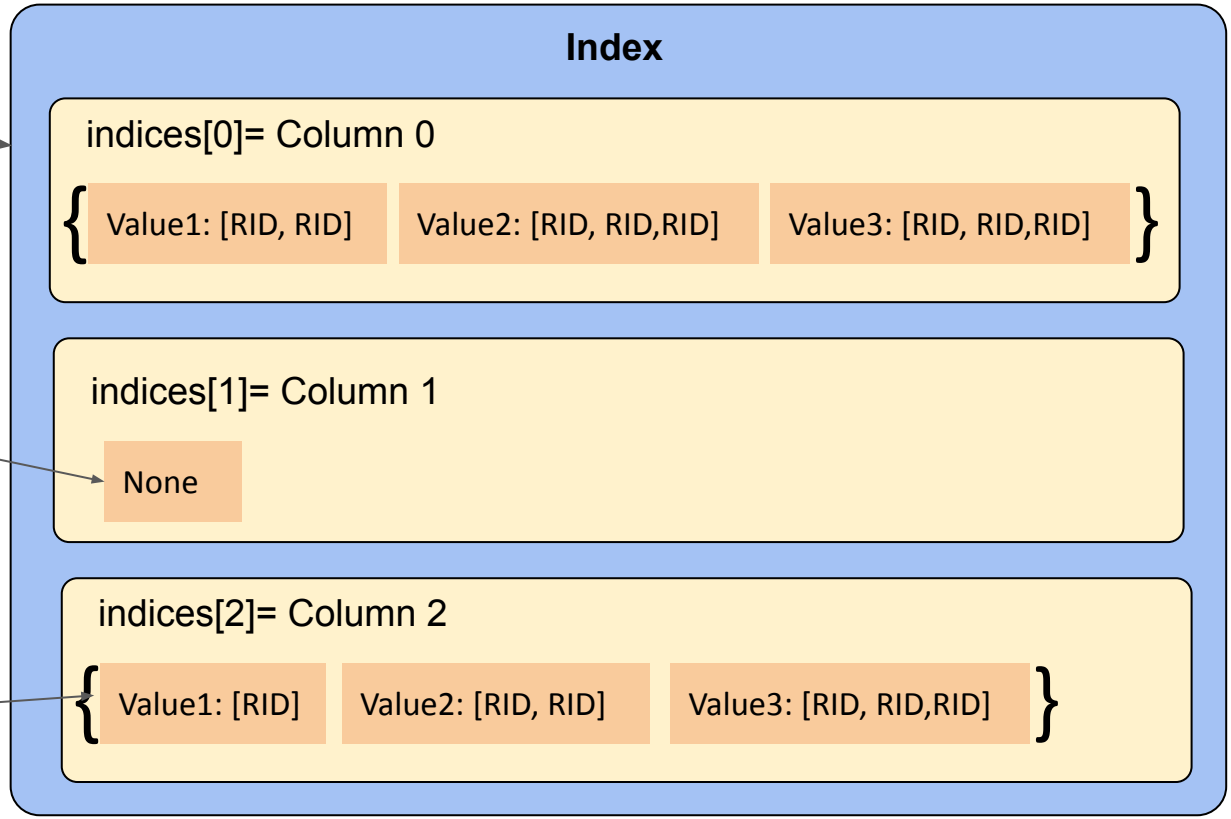

Index Create and Drop Column

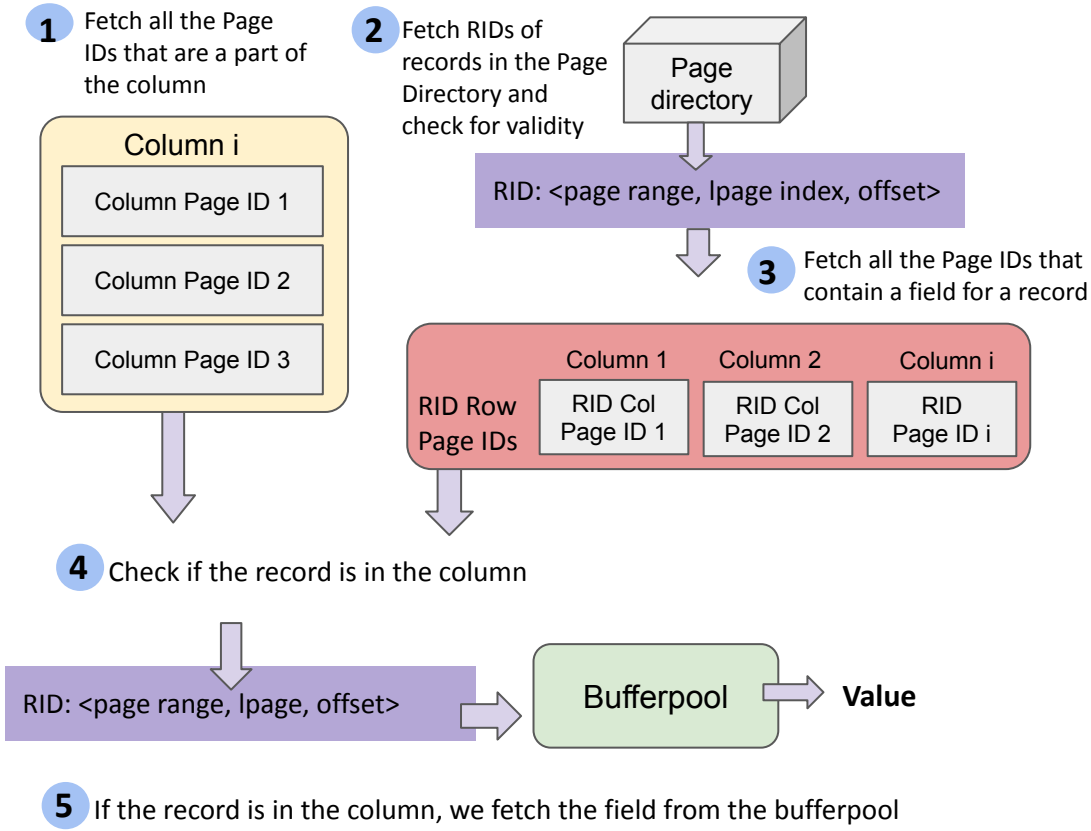

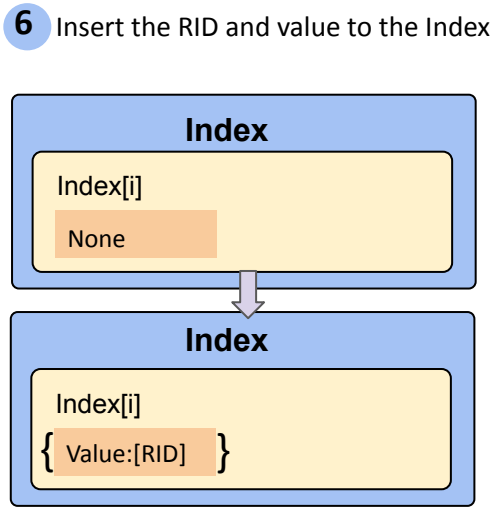

Durability

File Organization & Structure

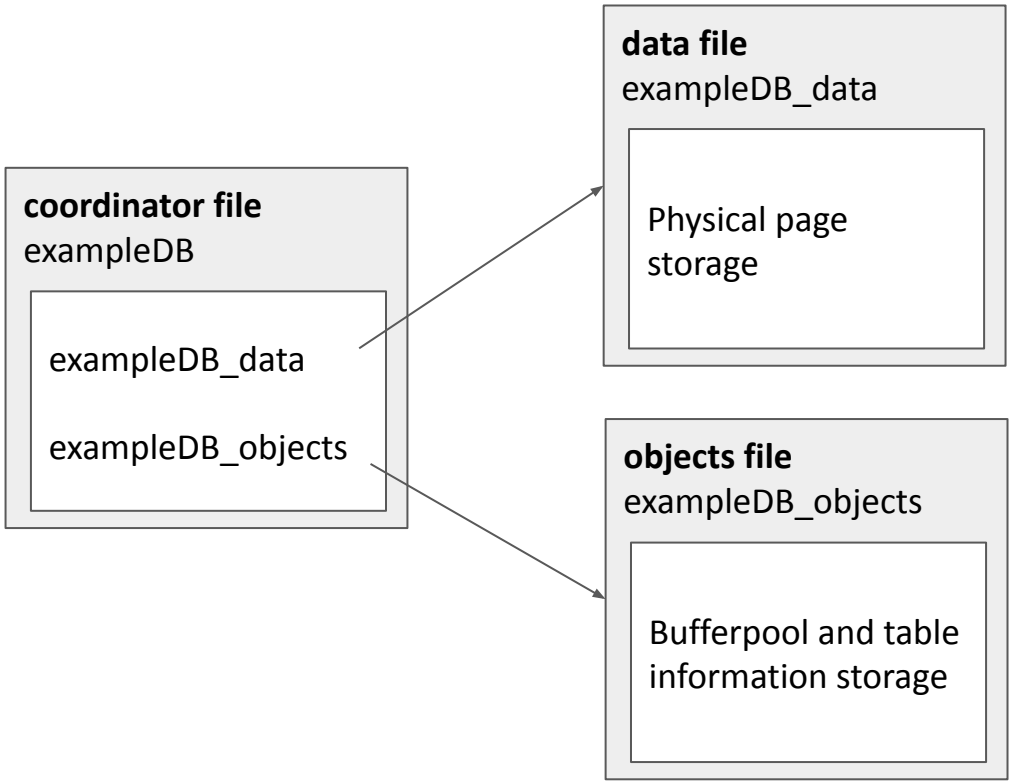

3 files

Separate files for data and for table & bufferpool information

db.open(): restore from files

db.close(): flush dirty pages to data file and store other information in objects file

Physical Page Metadata

Moved metadata inside the 4096 block allocated to a physical page

Added **TPS** (Tail Page Sequence Number): RID of last merged tail record

TPS is used to help determine whether a field is up to date

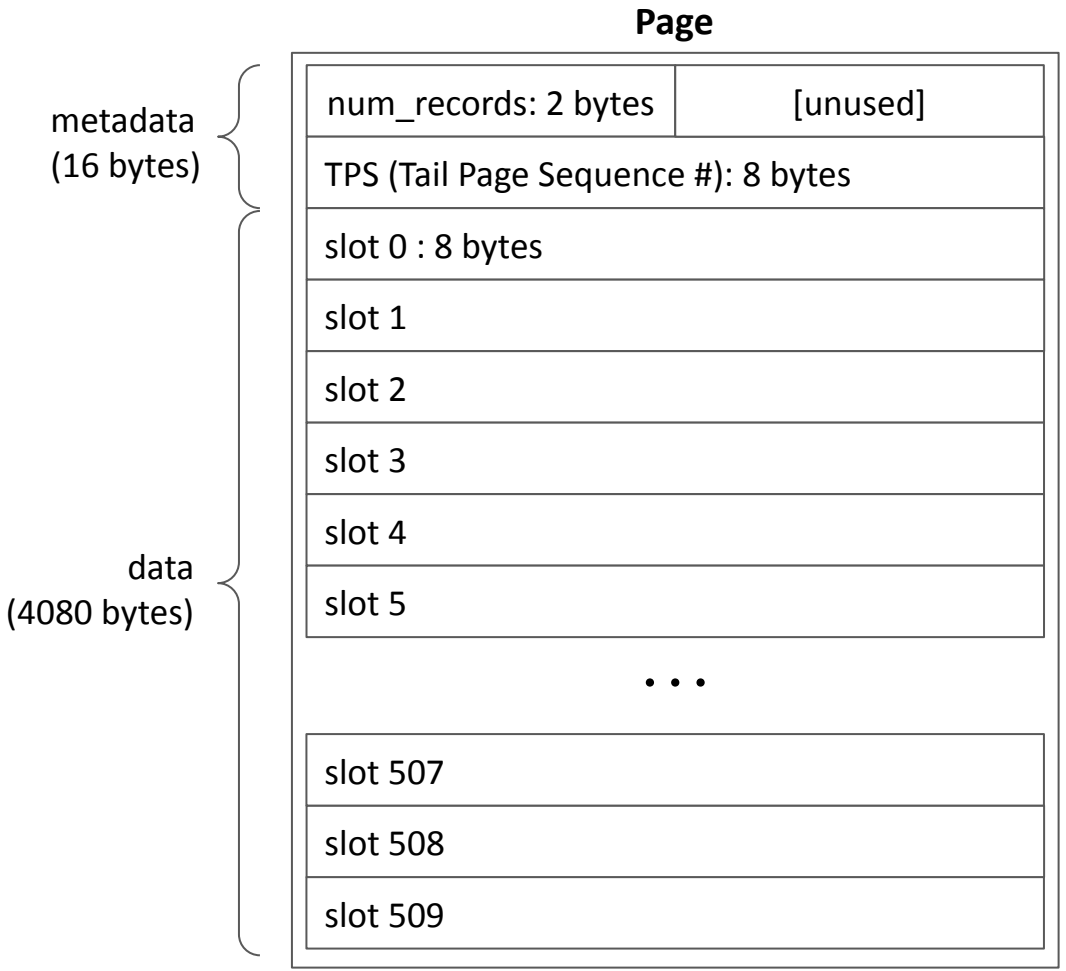

Bufferpool

Abstracts disk accesses

Shared by all database tables

Granularity of physical pages

Uses locks to prevent conflict between merge and user threads

Deallocates disk space when a table is dropped

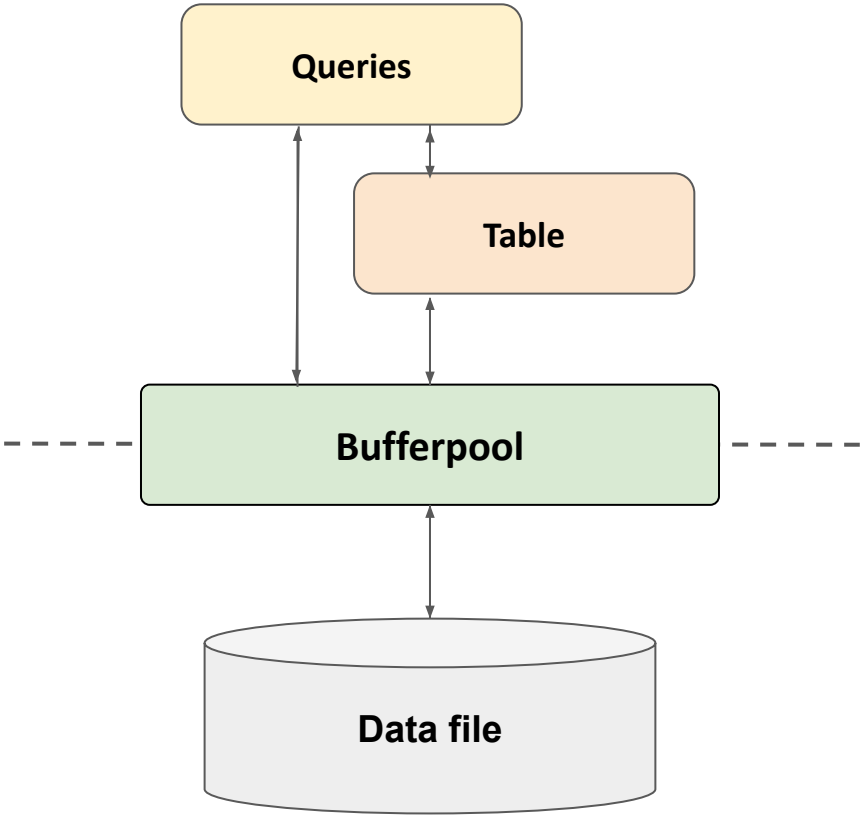

Bufferpool

pool_index: index of physical page in bufferpool and pins array

pageid: virtual page identifier

frameid: disk location identifier

Page Replacement (LRU)

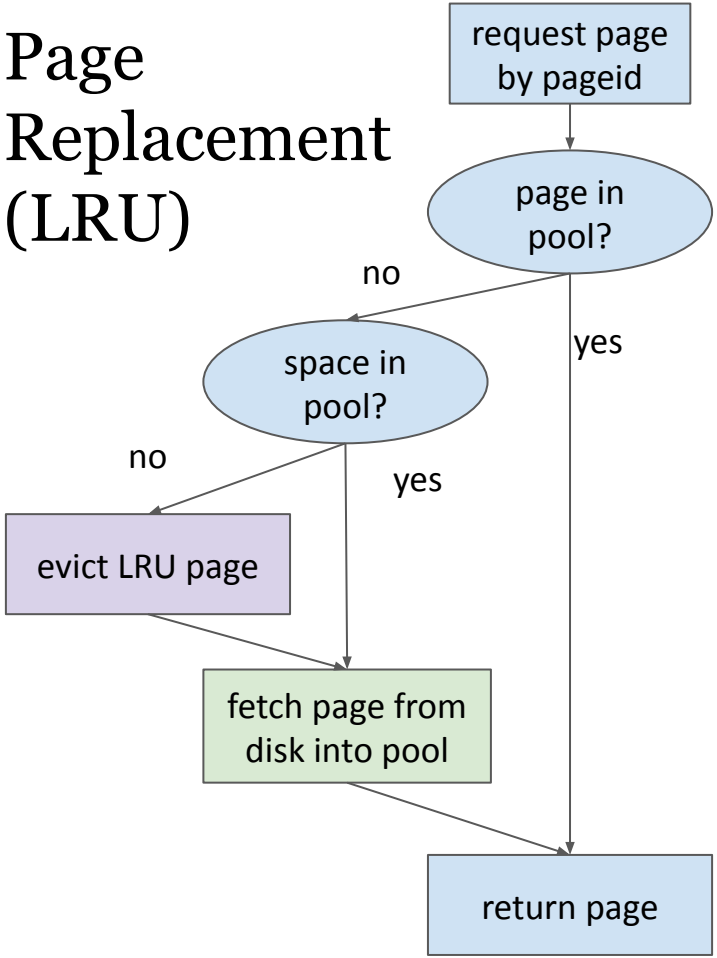

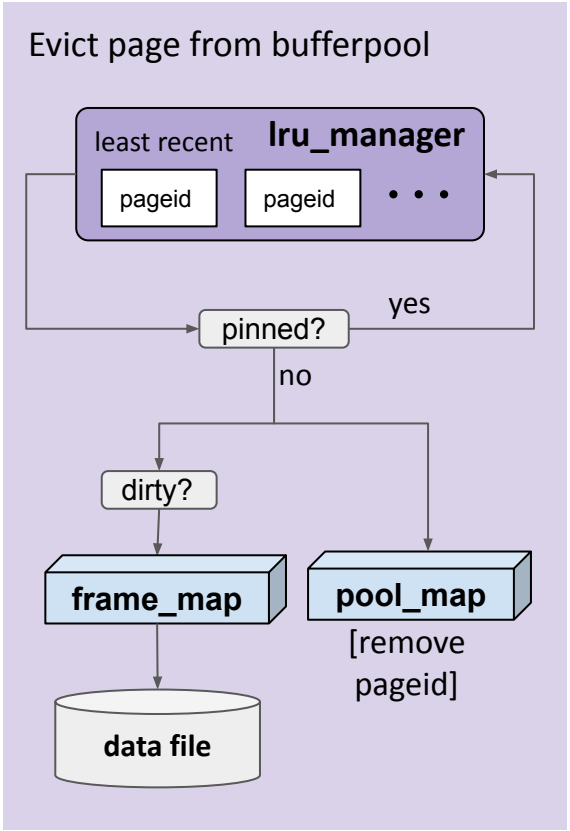

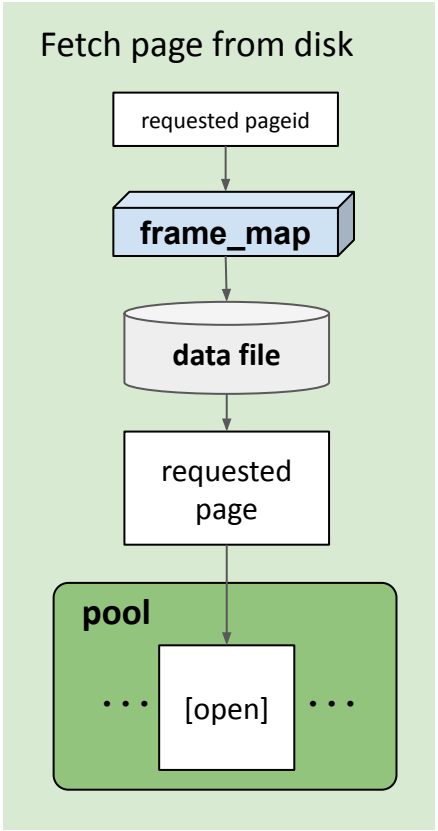

Merging

Initiating Merge

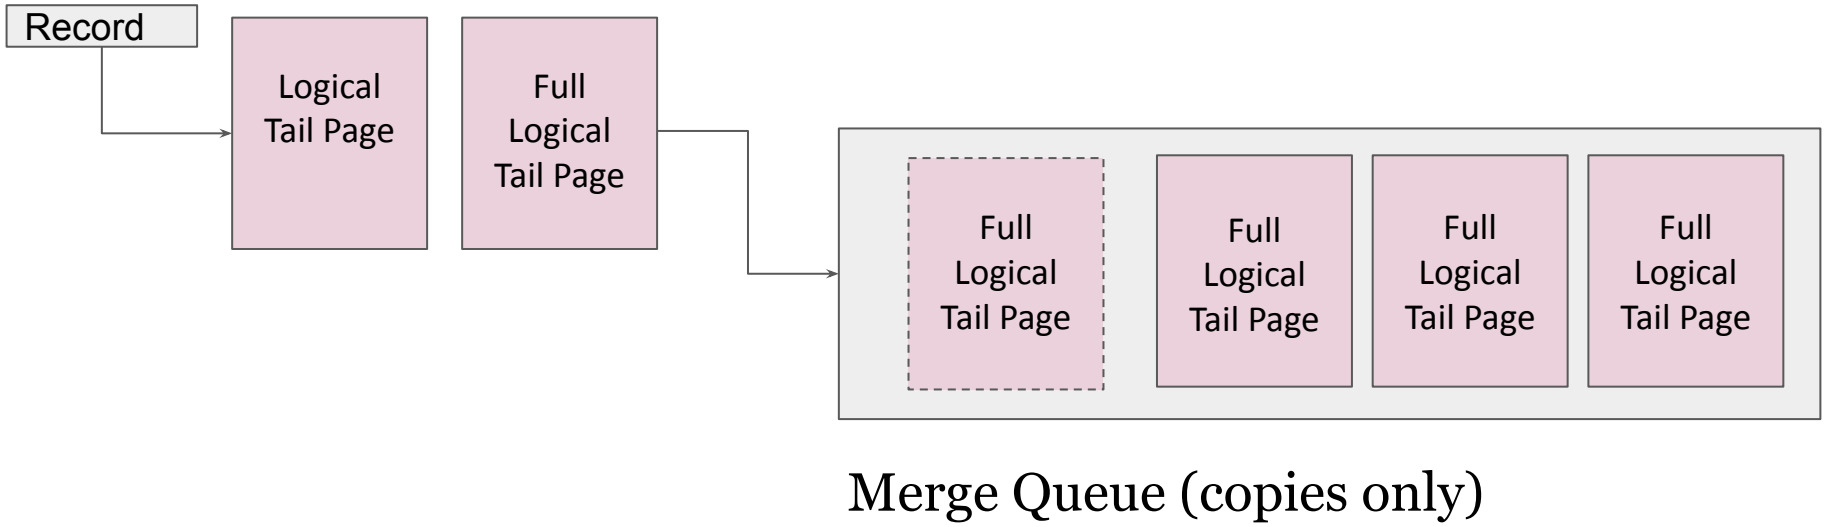

Merge after reaching threshold value of filled tail pages

Copy To Merge Pool

Separate storage area for pages used in the merge

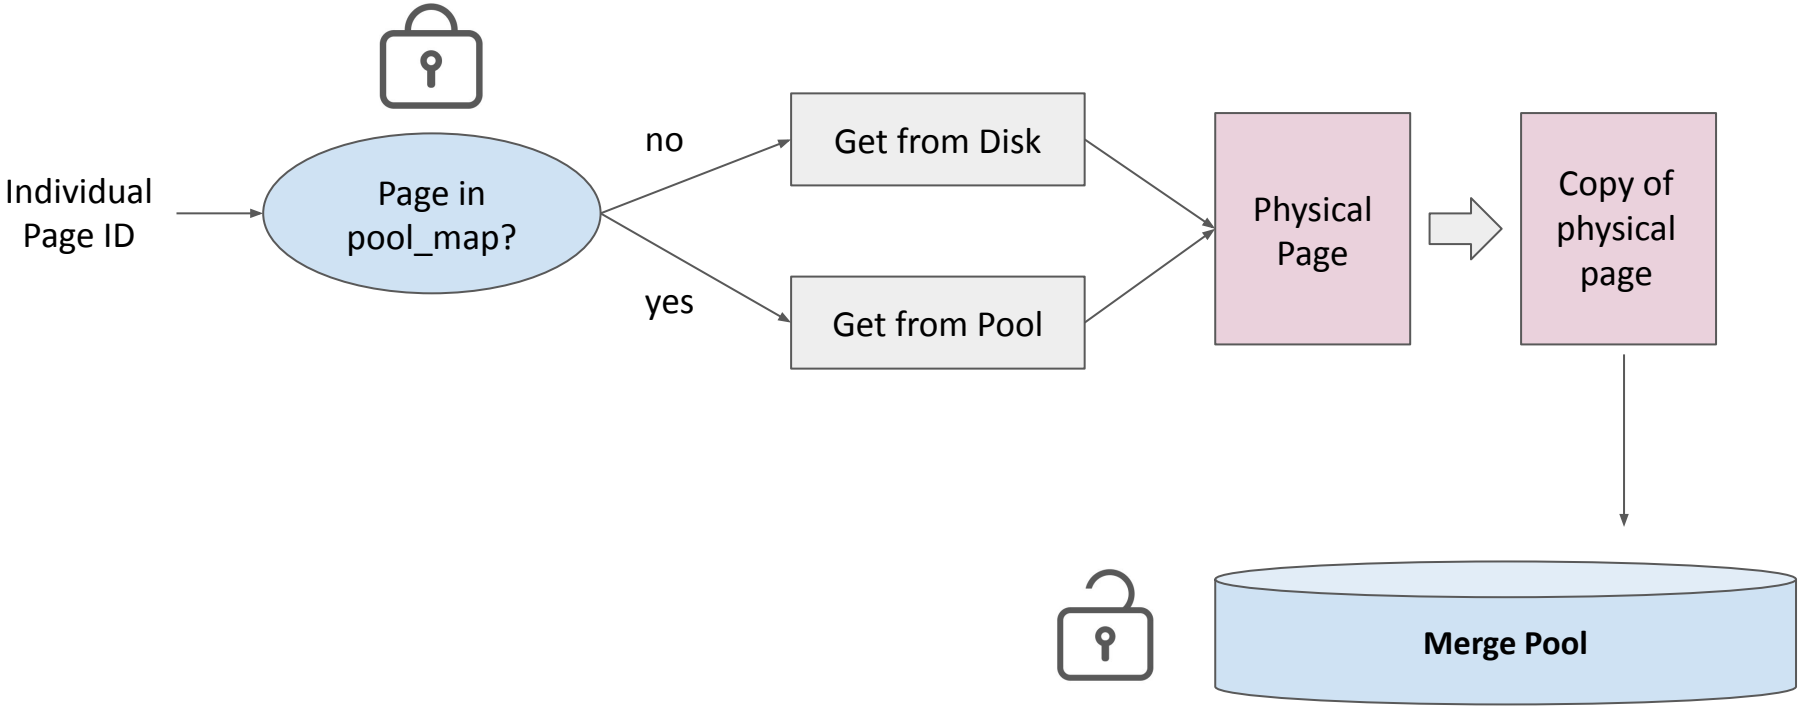

Merging

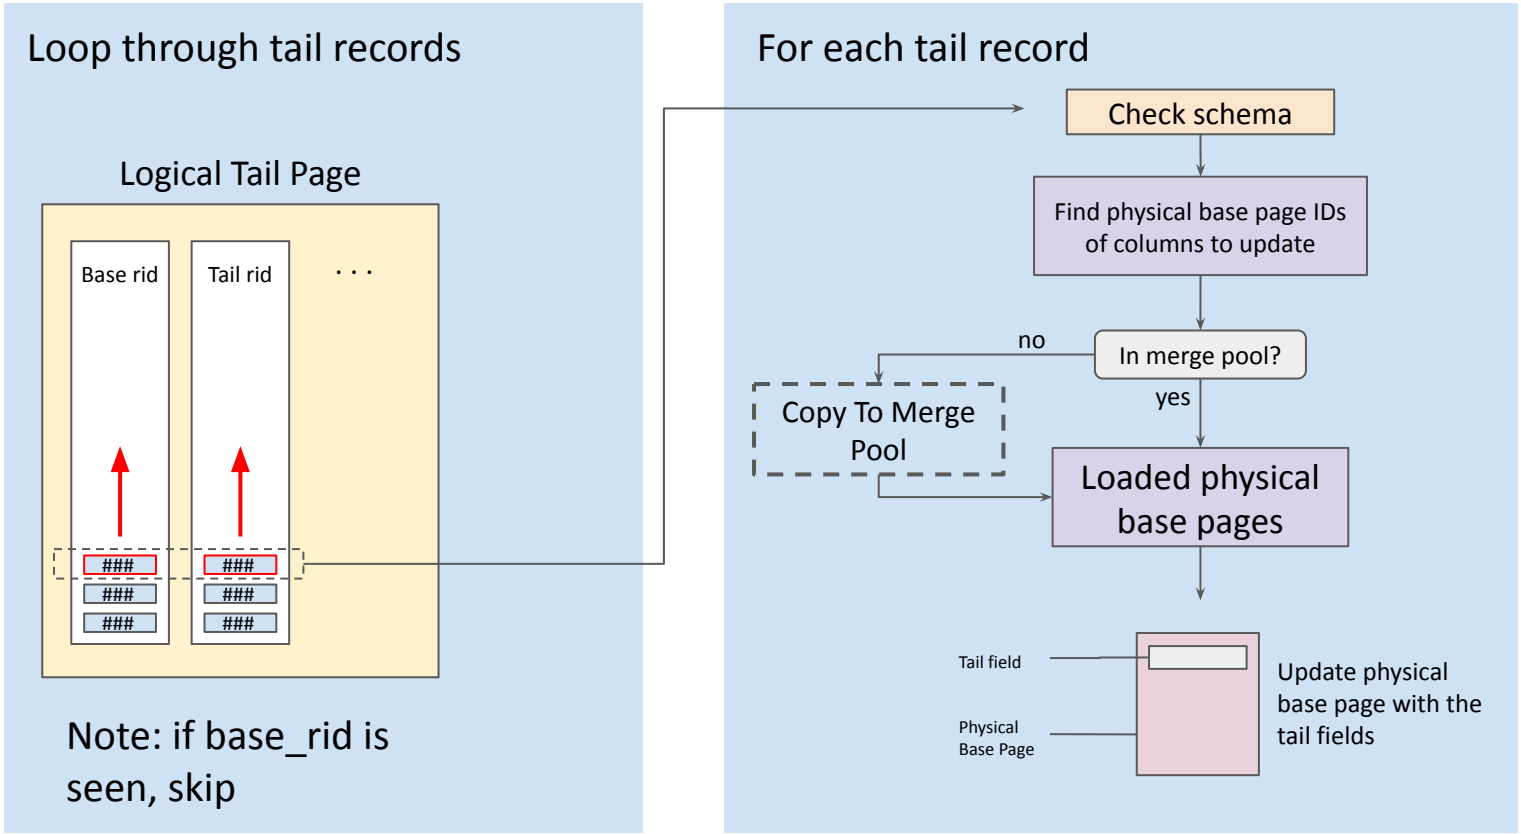
Loop through tail records
Logical Tail Page
Base rid
Tail rid
. . .
###
Note: if base_rid is seen, skip
For each tail record
Check schema
Find physical base page IDs of columns to update
In merge pool?
no
yes
Copy To Merge Pool
Loaded physical base pages
Tail field
Physical Base Page
Update physical base page with the tail fields

Concluding Merge

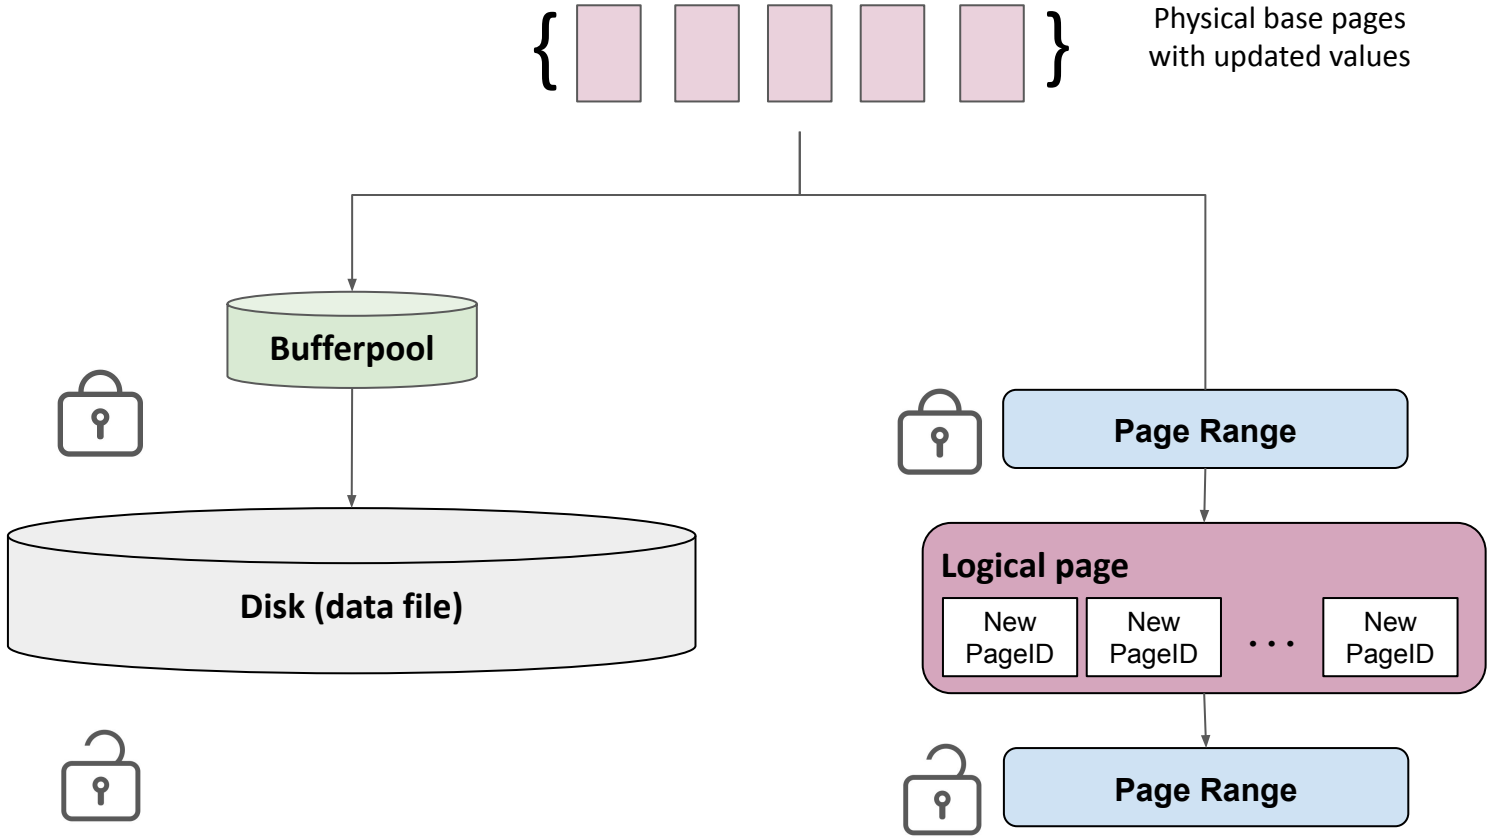

Performance

Tuning Bufferpool and the Threshold for Merging

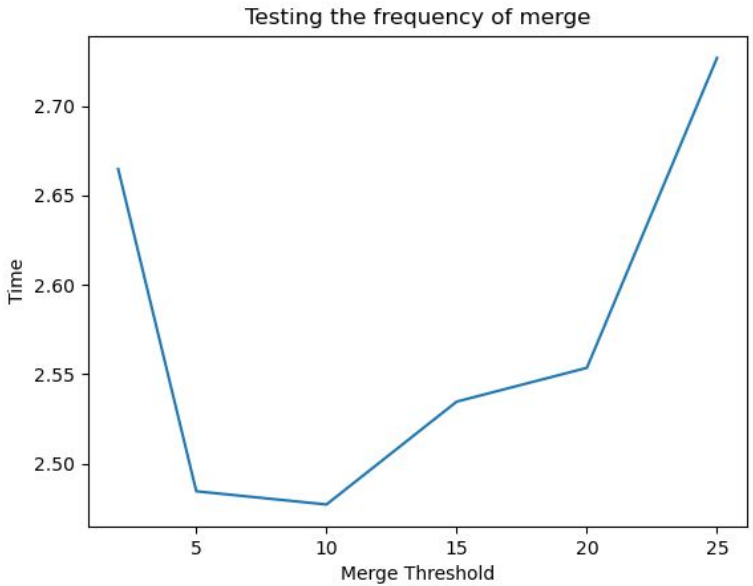

Finalized Merge Threshold = 10

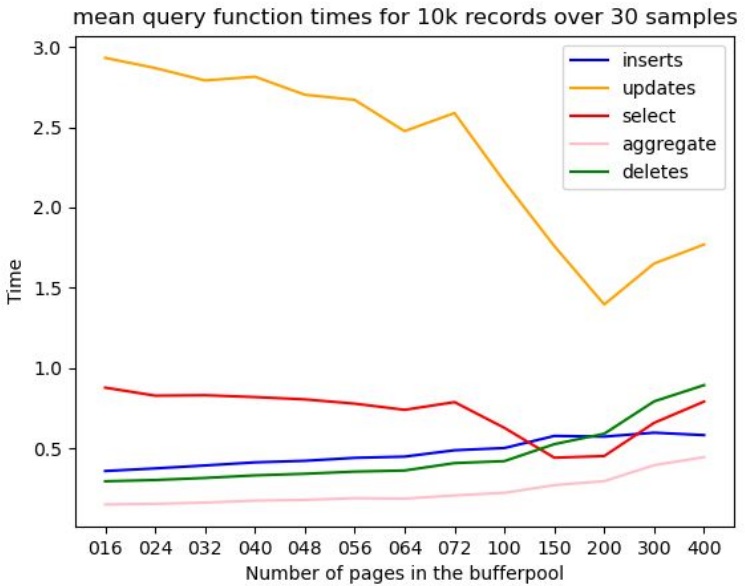

Finalized Bufferpool = 200

Workload: tuning_merge.py, tuning_buffer.py
Hardware: Dual-Core Intel Core i7, 2.5GHz, 16GB, 4 MB L3 Cache

Query performance

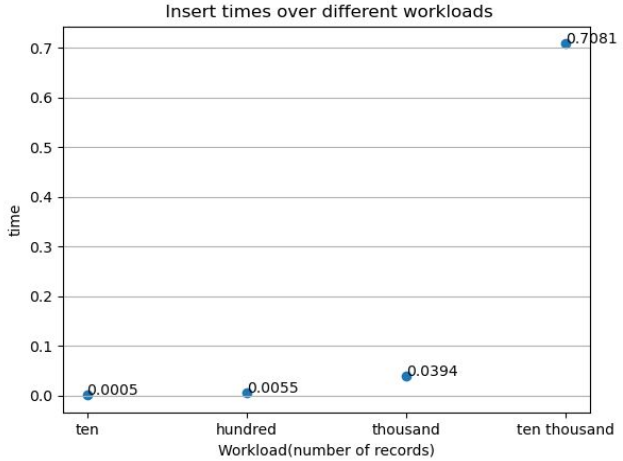

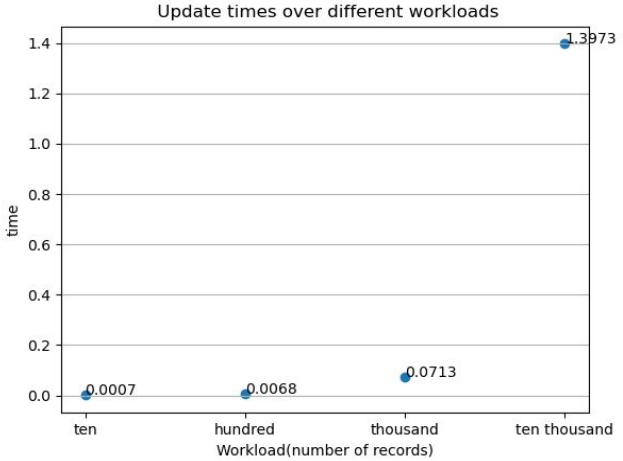

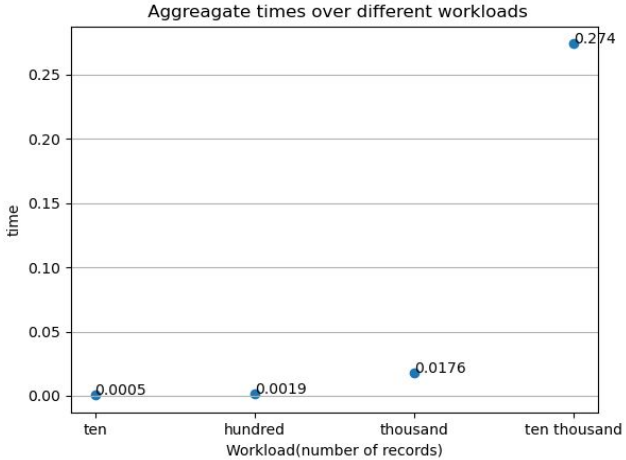

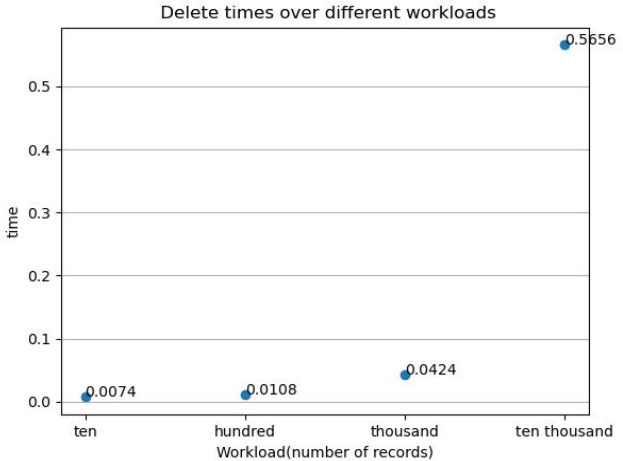

Workload: __main__.py

Hardware: Dual-Core Intel Core i7, 2.5GHz, 16GB, 4 MB L3 Cache

Q&A