
ECS165A

Milestone 2 Overview

Haley Raizes, Brittany Bates,
Jim McKerney, Nicholas Chen, Chris Bried

Objectives

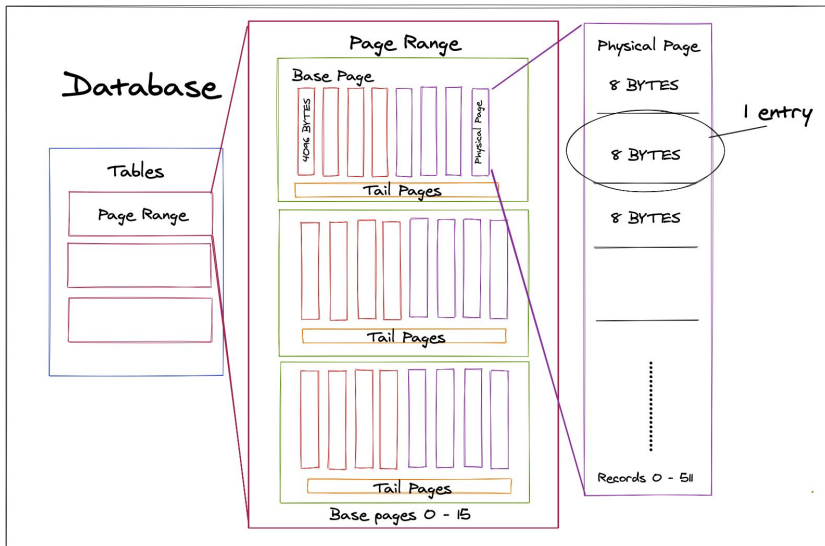
- Disk Structure
 - Bufferpool Design
 - Eviction Policy
 - Indexing
 - Merge
 - Performance
 - Questions
 - Demo
-

The Path to Durability

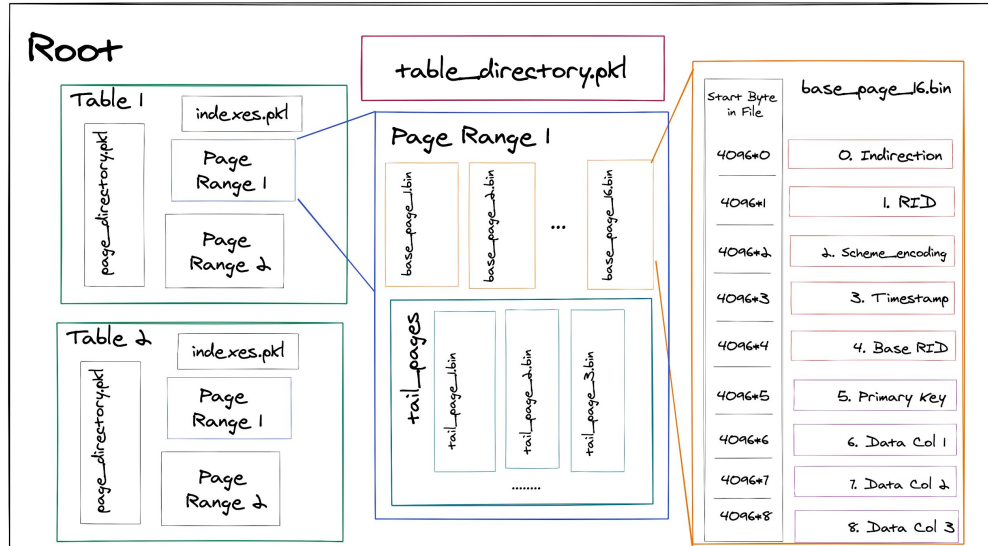
Milestone 1



Milestone 2




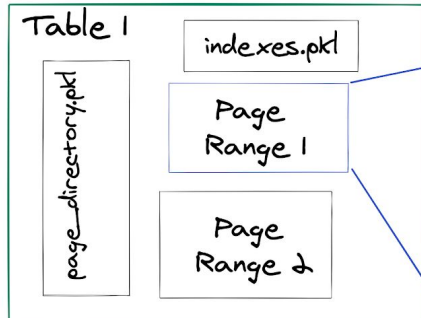
Volatile Memory



Non-Volatile Memory

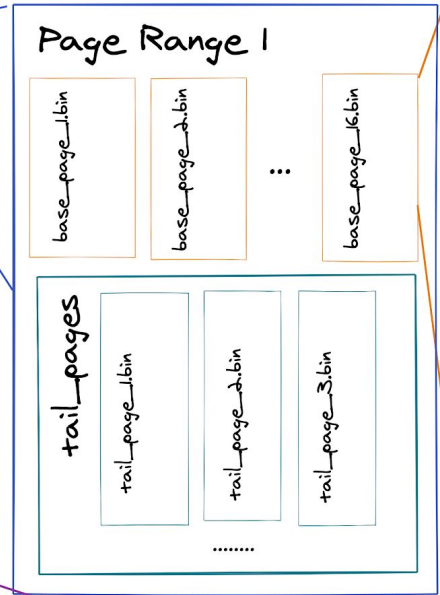
Disk Structure

Root



```
index.es.pkl
"all_indexes = {
"primary_key": {...},
"column_1": {...}
}"
```

table_directory.pkl



```
"page_directory" = {
"table_info": {num_records: 87, ...},
0: {"base_page" = true, "page_range" = 0, "base_page" = 0, "page_index" = 0},
...
}"
```

page_directory.pkl

Start Byte in File	base_page_16.bin
4096*0	0. Indirection
4096*1	1. RID
4096*2	2. Scheme_encoding
4096*3	3. Timestamp
4096*4	4. Base RID
4096*5	5. Primary Key
4096*6	6. Data Col 1
4096*7	7. Data Col 2
4096*8	8. Data Col 3

Indirection Col	Size
0	8 BYTES
1	8 BYTES
2	8 BYTES
3	8 BYTES
.....
511	8 BYTES

Bufferpool Design


Bufferpool

*A frame loads 1 base or tail page

```
frame_directory = {  
  (+table_name, page_range, base/tail page index, is_base_record) : frame_index  
}
```

Main Methods
evict_page()
load_page()
commit_page()
at_capacity()
is_record_in_pool()
add_frame_to_directory()
commit_all_frames()

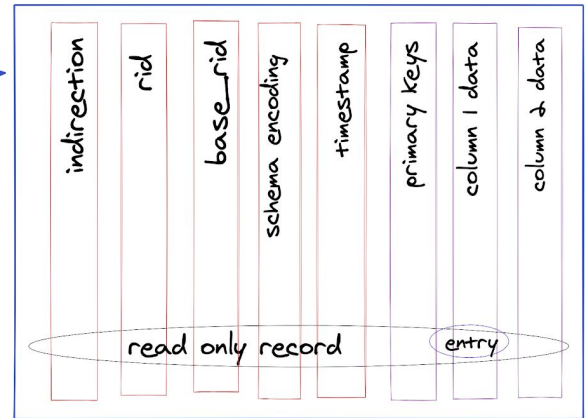
Frame Array



*Bufferpool size is set in the configuration file

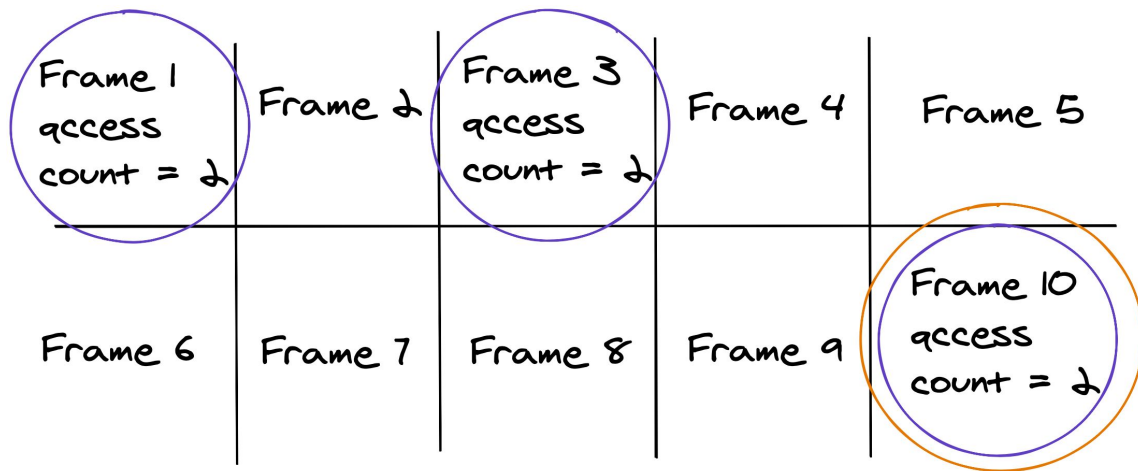
Frame

- all_columns
- pin
- dirty_bit
- time_in_bufferpool
- access_count
- path_to_page_on_disk
- frame_key



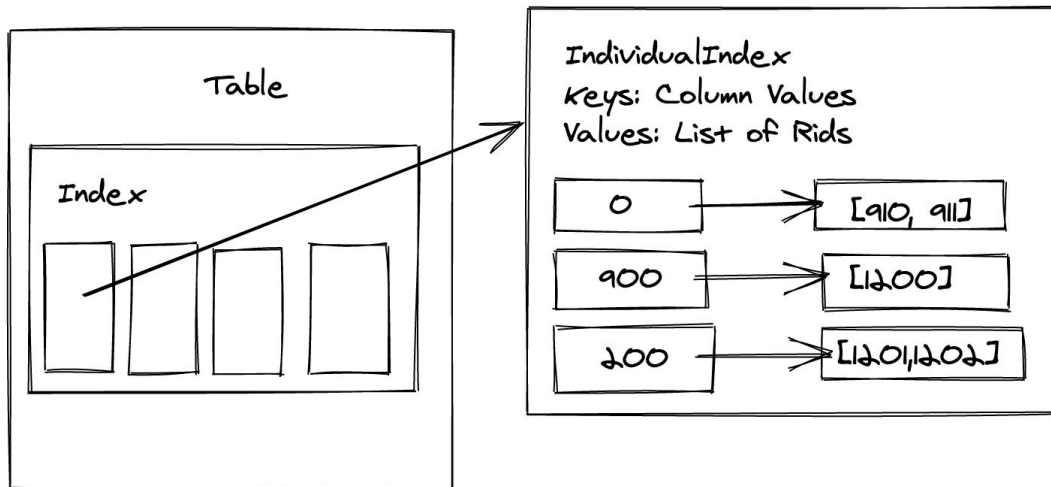
Eviction Policy

- Synthesis of LRU and LFU policies
 - Among the least frequently used records, evict the least recently used
- Chose this method primarily for speed & simplicity
- Does not distinguish between privileged and unprivileged data



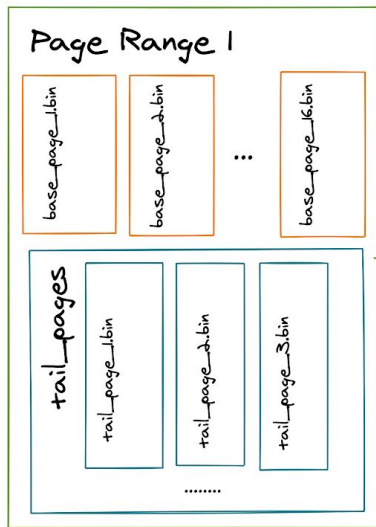
Indexing

- Hash index
- Maps column values to list of RIDs
- Used for select and update
- Index created automatically for primary key
- Bufferpool frames are committed before index creation
- Index persisted as .pkl file



Merge - General Flow


①. Call to update instantiates merge after x number of updates, where x is defined in config.py



②. Merge gets its own thread

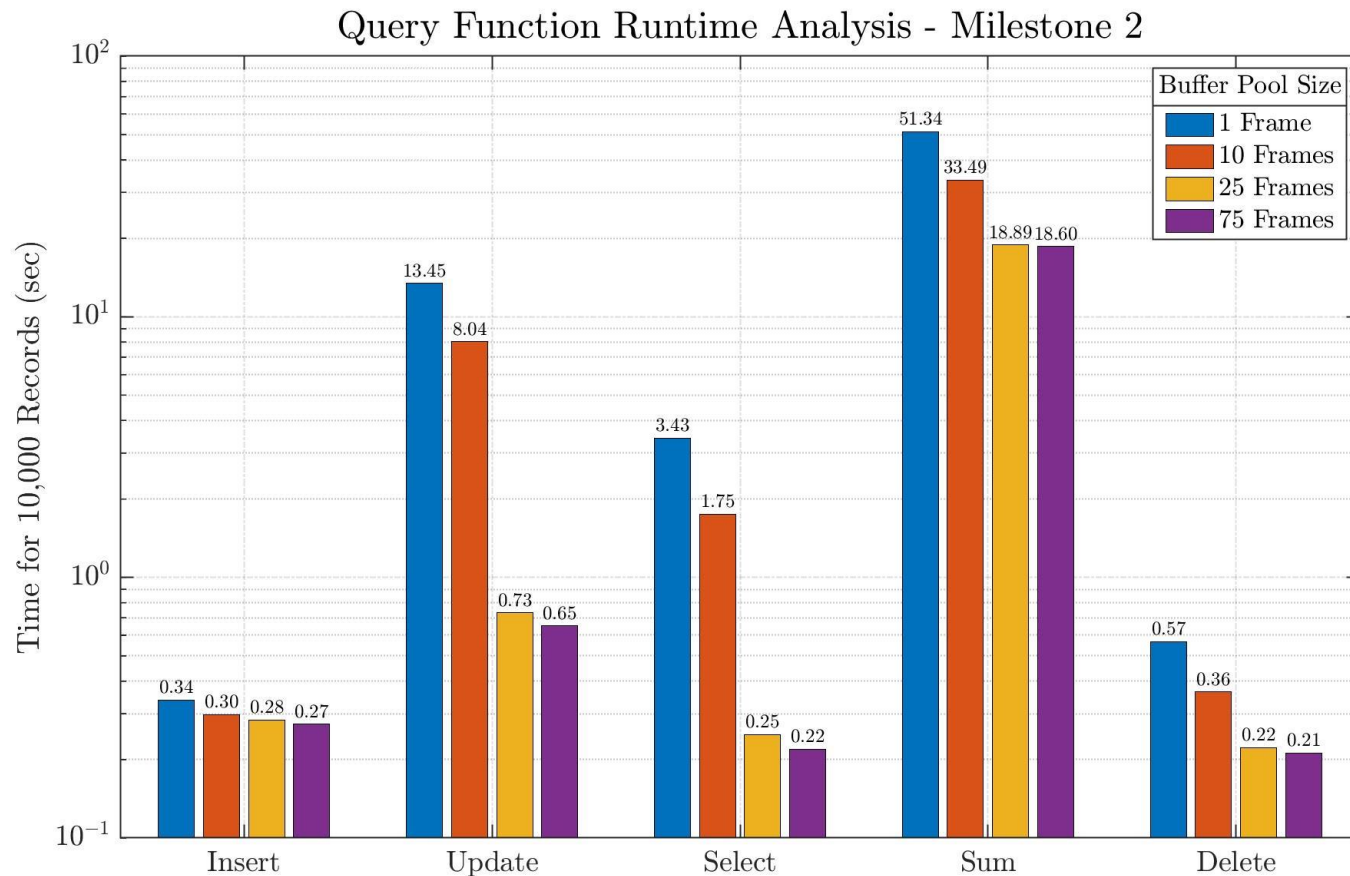
③. Load in base and tail pages per page range for merge

Bufferpool Instance



⑤. Write updated records back to disk

Optimizations & Performance



*These times are based on 10 run averages using the provided `__main__.py`

Specs: 6 core Intel Core i7, 2.6 GHz, 256KB L2 cache per core, 12 MB L3 Cache, 16GB Memory

Questions

Demo