Milestone 2

Dulce Torres, Kyle Pickle, Pranav Kode, Sean Nguyen, Ruqayyah Siddique, Wen Wei Tan

Durability & Bufferpool Extension

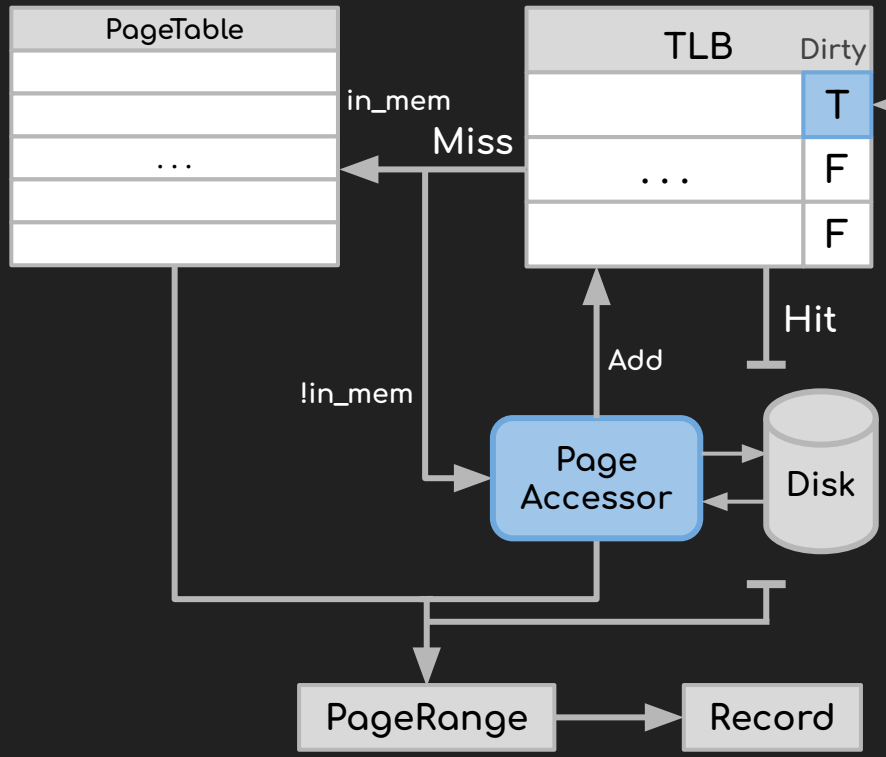

- New **PageDirectory** works at the granularity of **PageRanges**
- Two modes, in-memory and on-disk (automatic)
- **TLB** has been updated to work with **PageRanges**
  - **Dirty** and **Pinned** bits
- New **PageAccessor** works with **FileService**, pulling from and writing to disk

Bufferpool (TLB)

- New columns, **Dirty** and **Pinned**
- When an entry is deleted, updated, or added to, it is marked as **Dirty**
- When an entry is being merged, it is **Pinned**

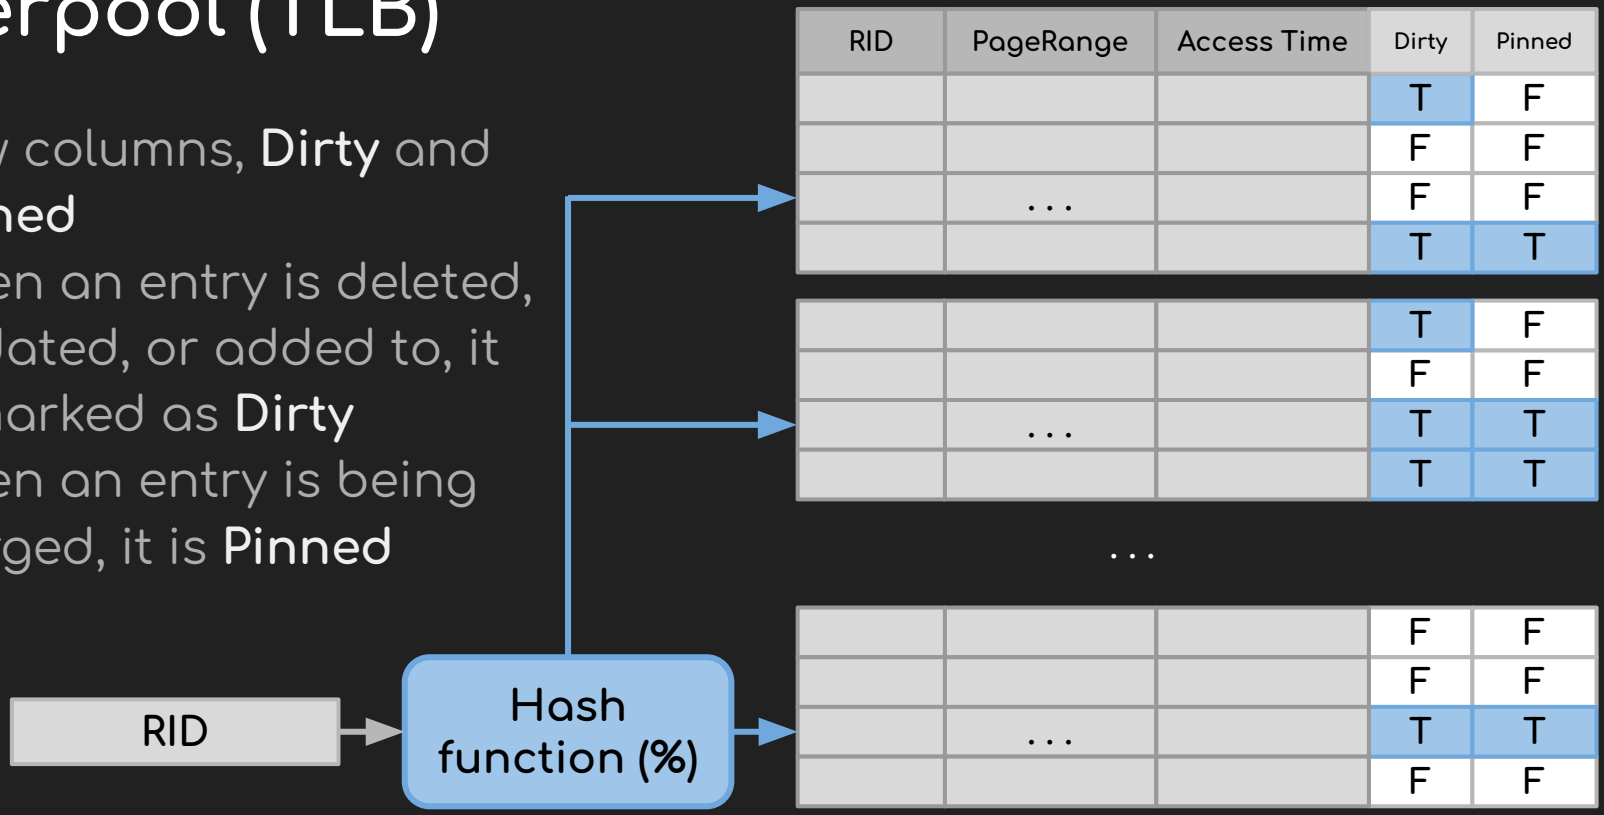

FileService

- **DataBase** is read from / written to a single .db file
  - No **Pickle** used; organized and parsed through from scratch
- **FileService** has 3 Functions:
  - load_tables
  - pull_page_range
  - merge_tables

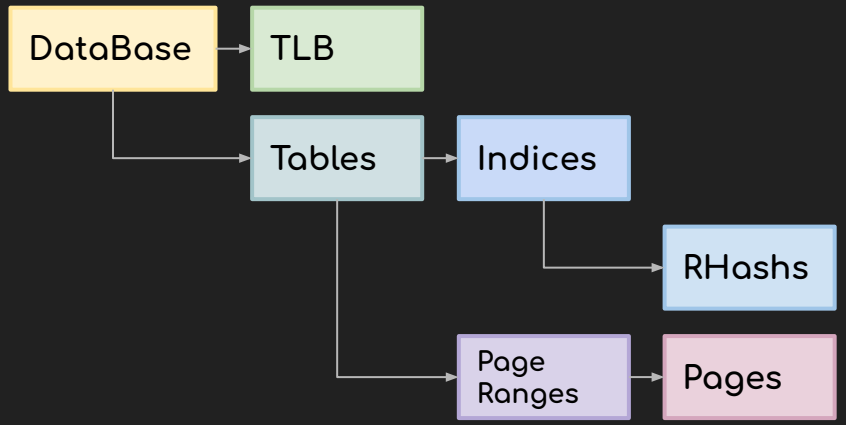

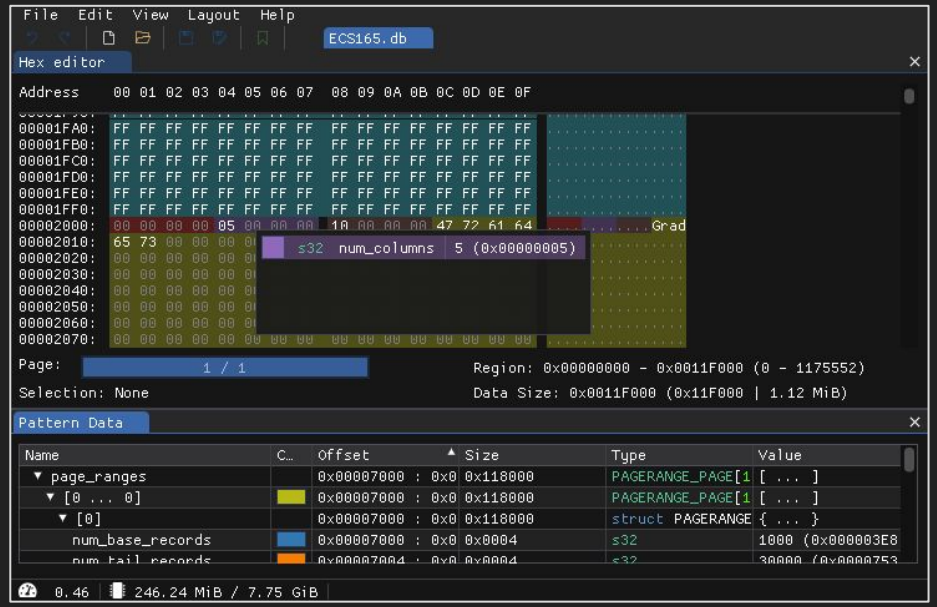

Merge

- Creates copy of new base page
- Updates the records of the old base page
- All prior tail records marked as deleted
- Old base page also marked as deleted
- Appends the new copy of the base page

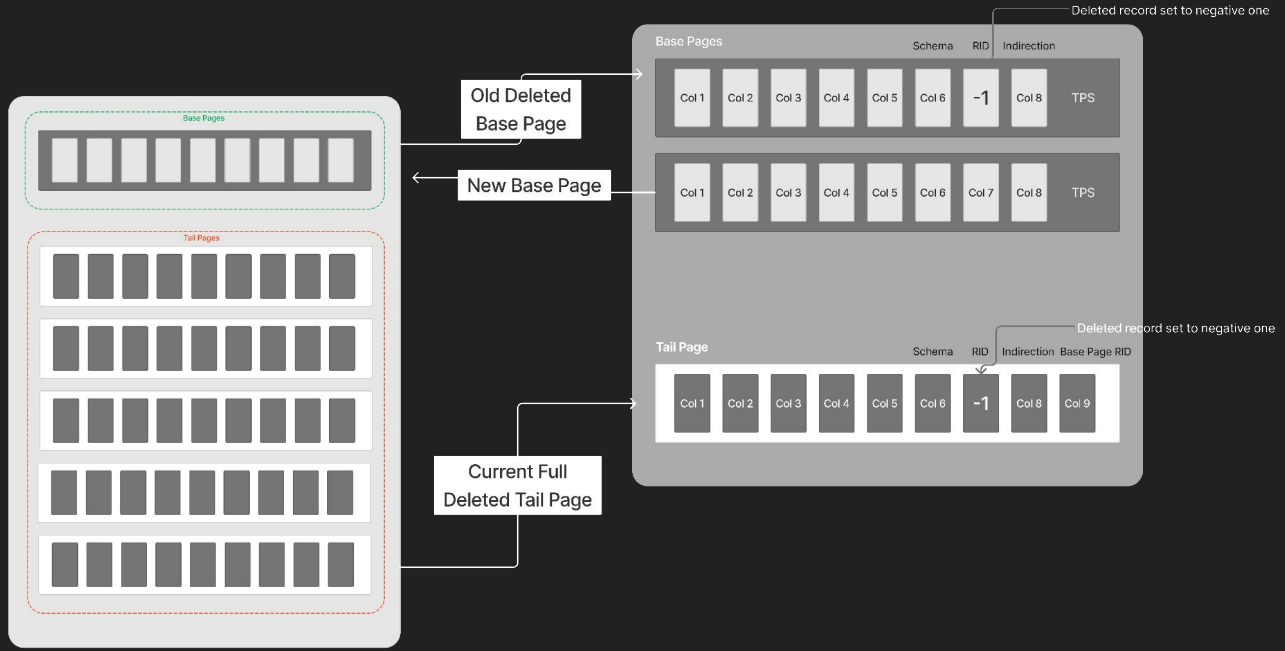

Indexing

- Create_index
  - Create a new RHash object
  - Scan through all table records and call the index's insert function for the corresponding column
- Drop_index
  - Deletes the index of the specified column from memory
- New seeding scheme
  - 1% chance to add a seed after an index insert
  - Smallest value is always one of the seeds
  - Max number of seeds is 25% of the total number of records
- Indices are persisted in the .db file

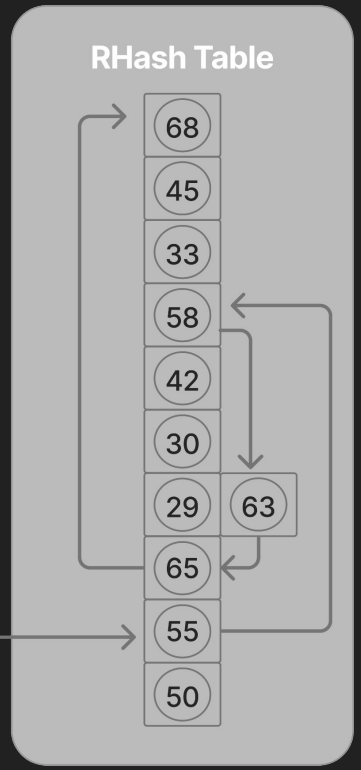

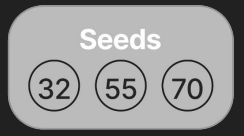

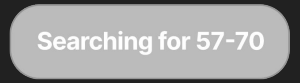

Milestone 1 vs Milestone 2

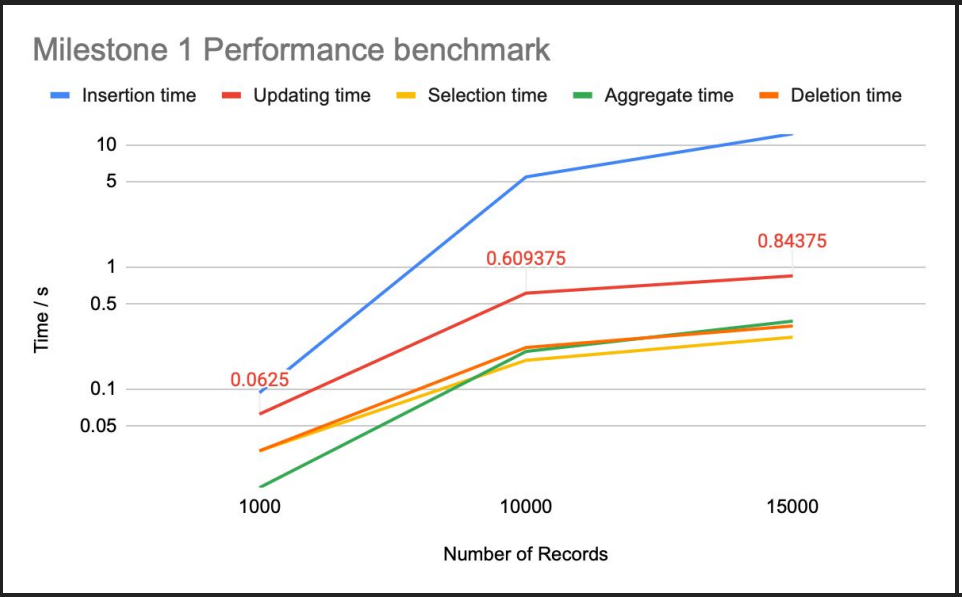

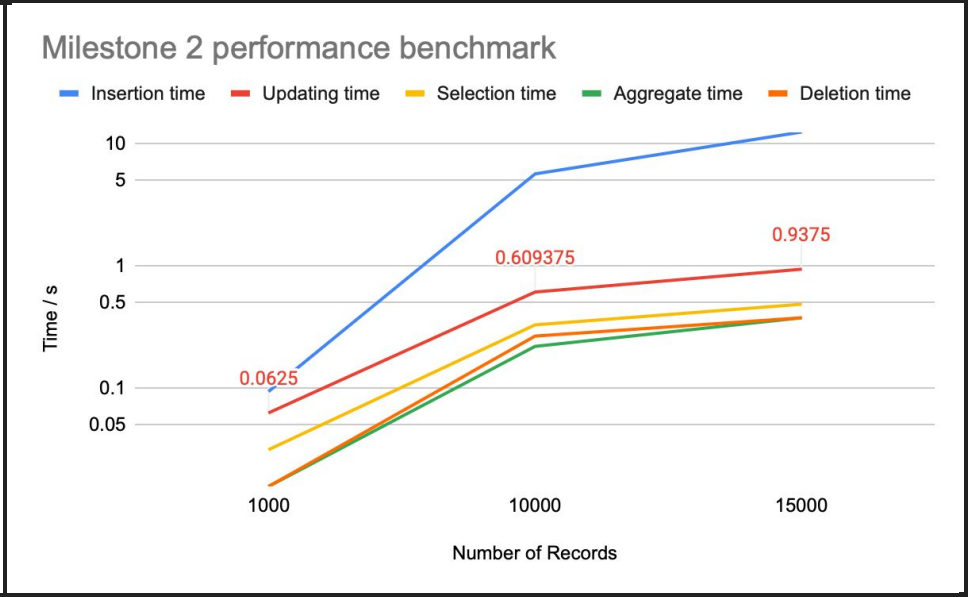

System: Ryzen 7 2.9 Ghz with 8MB L3 cache, 16GB ram
Workload: __main__.py

Thank You!