ECS165A Milestone 2

Winter 2022

Michael Pitts, Jericho Mata, Ethan Yu, Ryan Cen and Yuyi Li

Goals & Objectives

There are three objectives we hope to achieve in this project:

- The first concern is balancing data in memory and disk.
- The second objective is to expand select index capabilities.
- The last objective is data reorganization through a contention-free merge.

Updated Architecture

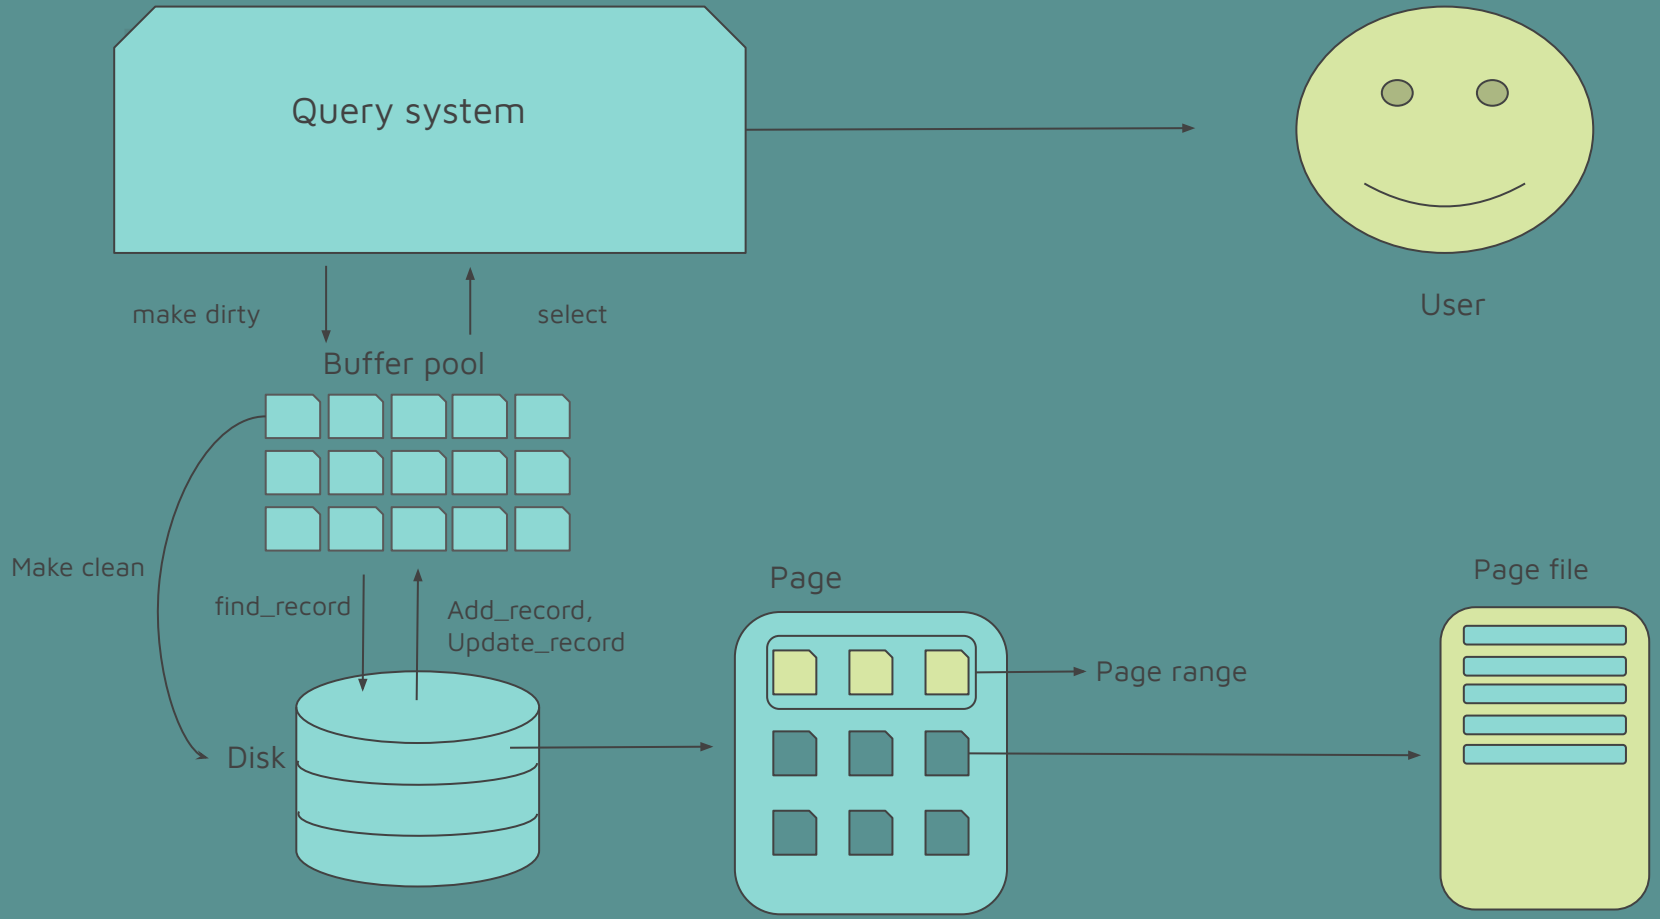

Bufferpool and Extension

- Due to limited memory space, we would like to save the more frequently used data(most recently used pages) in memory while keeping the less frequent ones in the disk.

- Pinning pages happens when we are currently using it.

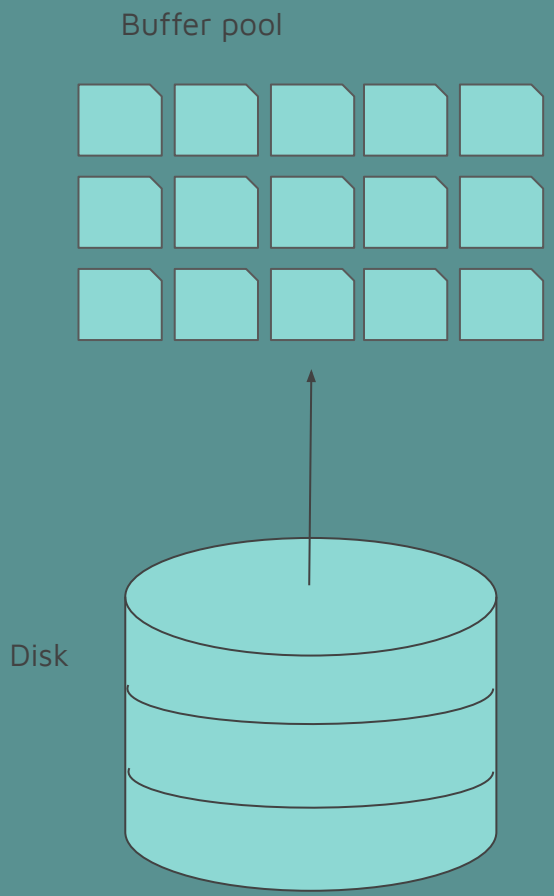

Bufferpool Eviction policy

- When it comes to evicting pages, we get rid of the least recently used ones. In the function evict_page(), we first clean the page first with make_clean(), then iterate through the bufferpool to find the least used, unpinned pages to evict.

- Lock up the disk whenever committing to it so we utilized the lock function alike to multithreading in the OS.

Eviction process

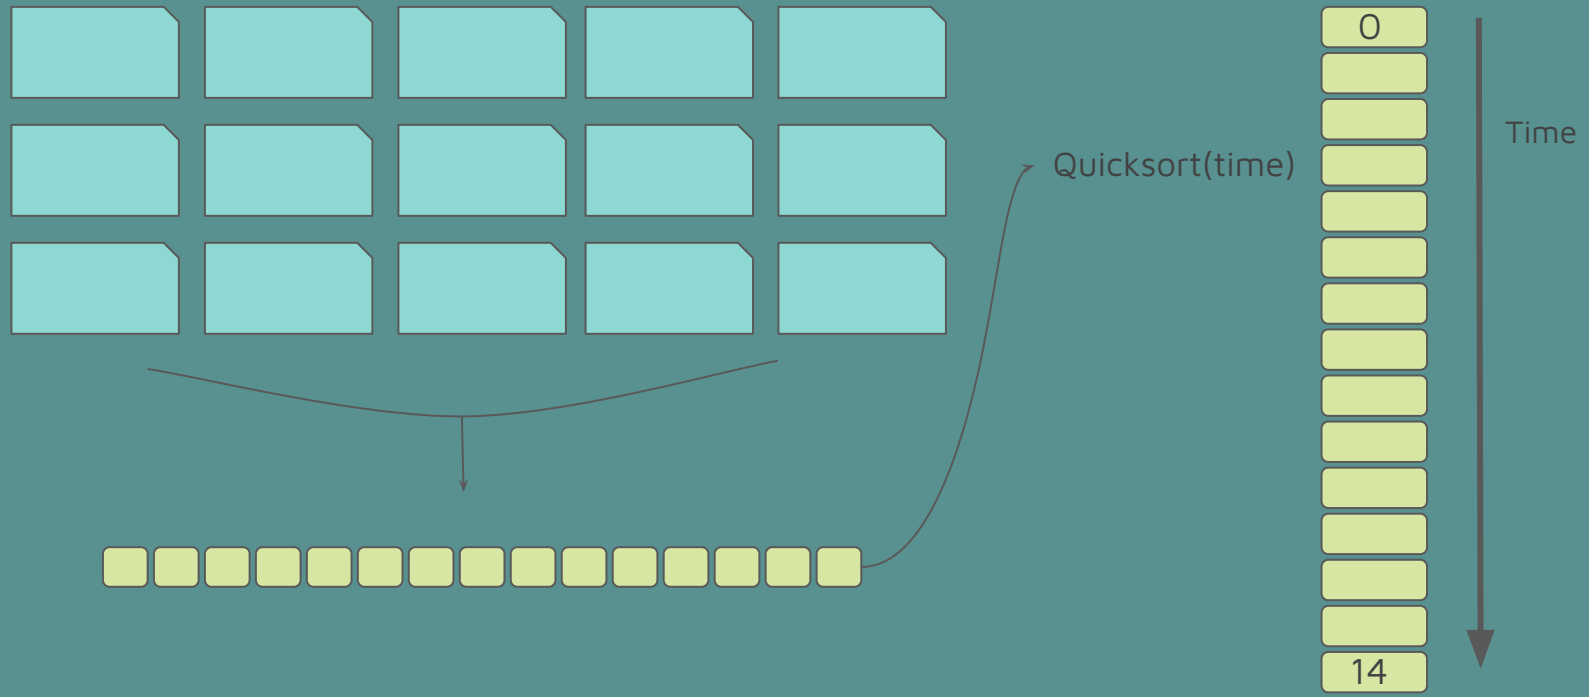

Eviction process

What's happening

Once the sorted array is extracted the function the runs down the list of pages from least recently used to most looking for an unpinned page. Once it is found that page is evicted.

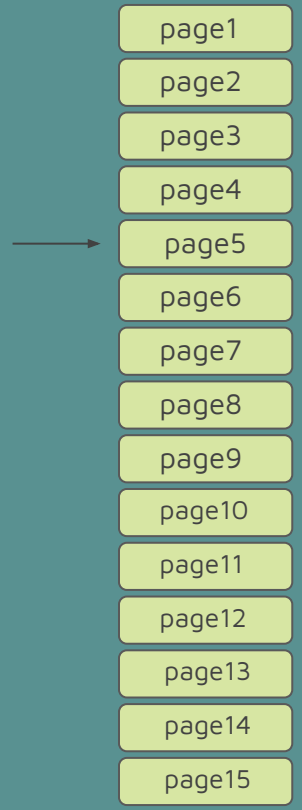


```
For page in sorted_pages:
        if(page not in pinned_pages):
                bpool.remove(page)
```

The Pseudo-disk:

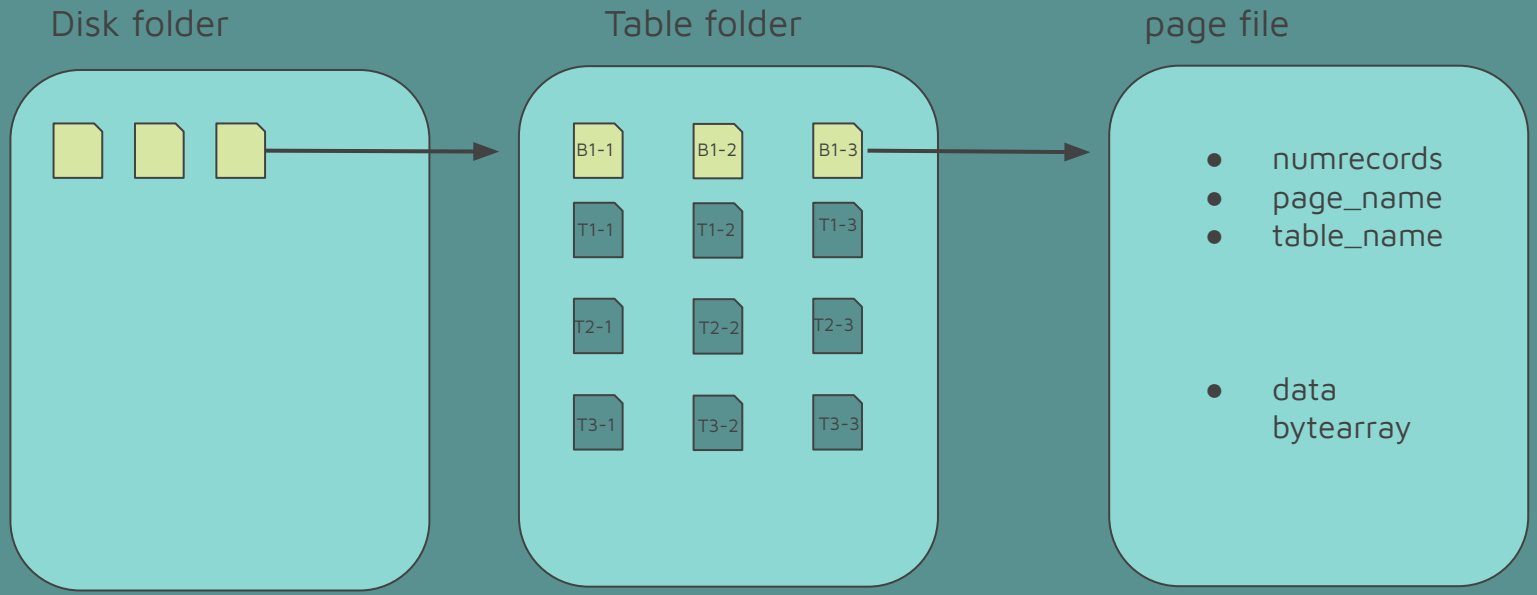

Merge function

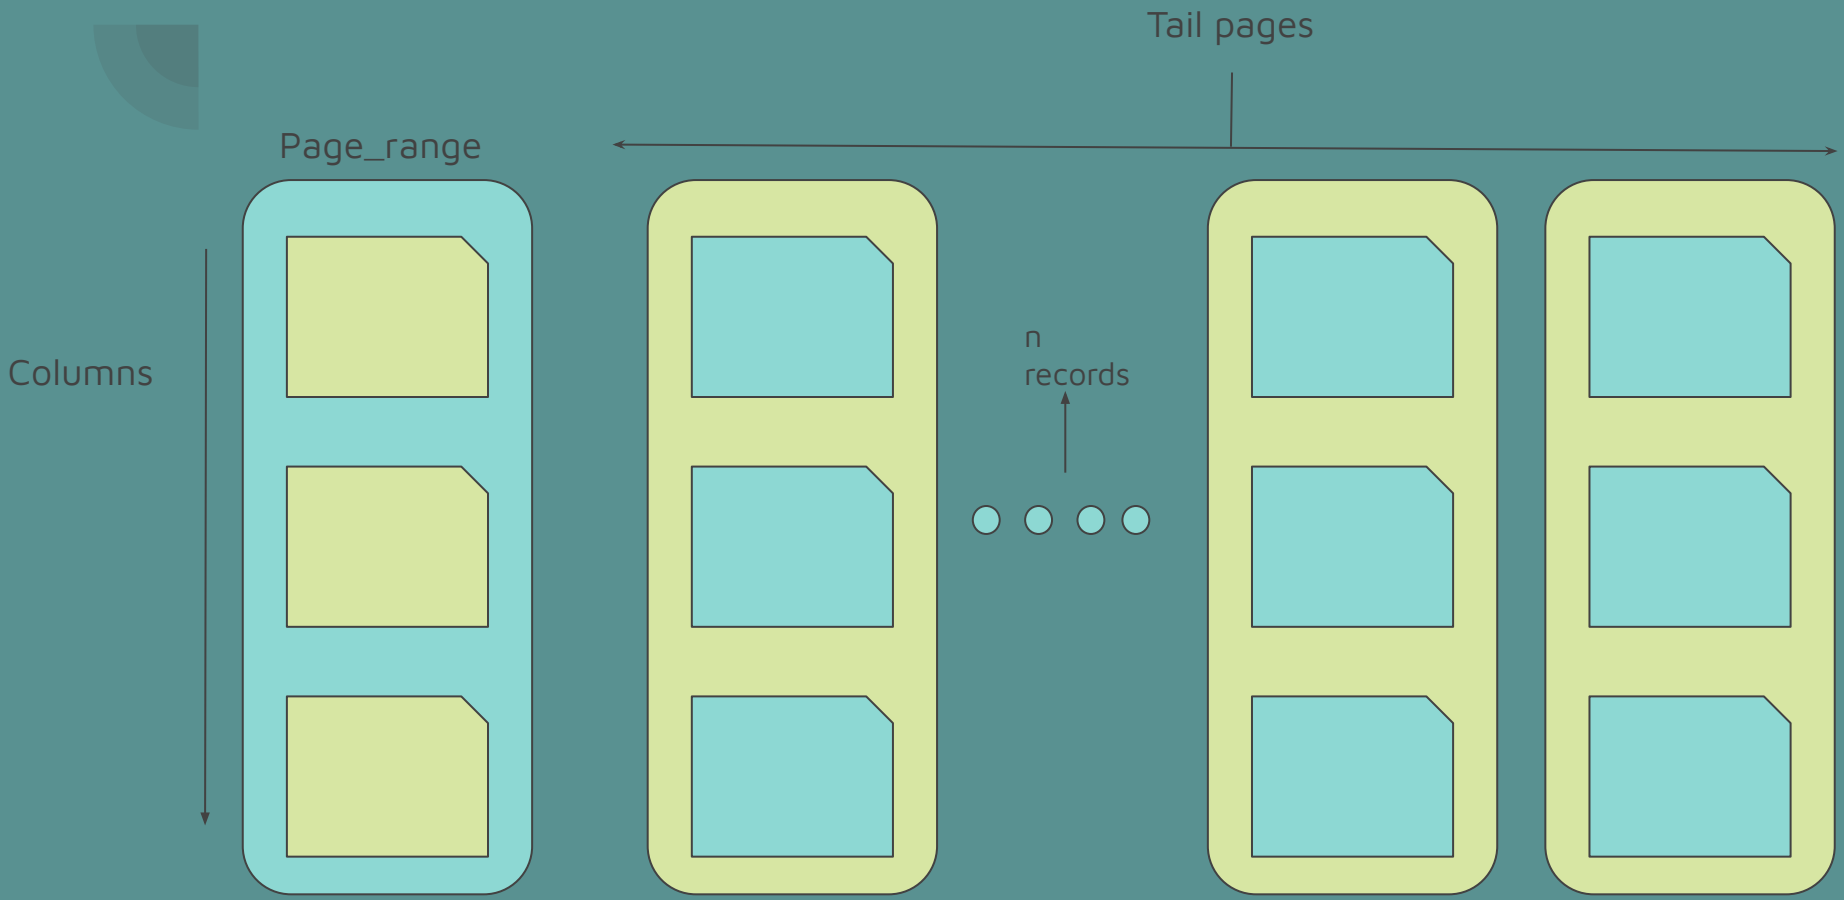

Merge function Cont'l

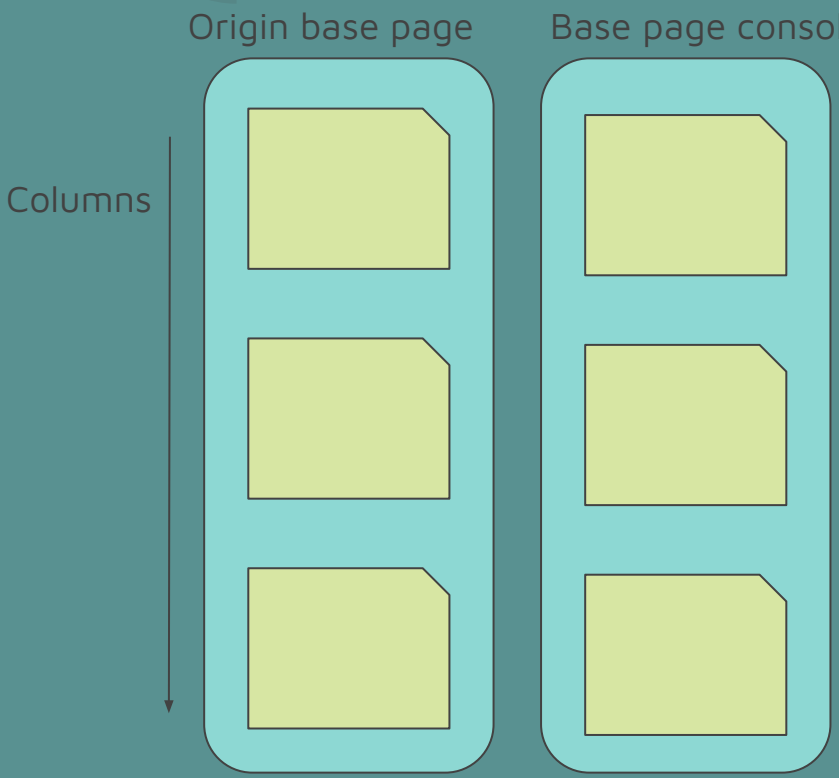

The 999 tail pages are the compressed for each column in the page range.

Result of Merge:

Because of this merge the update and select function no longer have to wade through a sea of tail pages to find their target. The result of this is greatly improved performance.

Indexing

Table: Company

Age	job title	salary	YOE
Base rid	Base rid	Base rid	Base rid
Base rid	Base rid	Base rid	Base rid
Base rid	Base rid	Base rid	Base rid
Base rid	Base rid	Base rid	Base rid
Base rid	Base rid	Base rid	Base rid

Primary Key

What's happening

All the columns are filled with base rids that are sorted by their get_newest_value which returns the most updated tail page value.

Indexing Cont'l

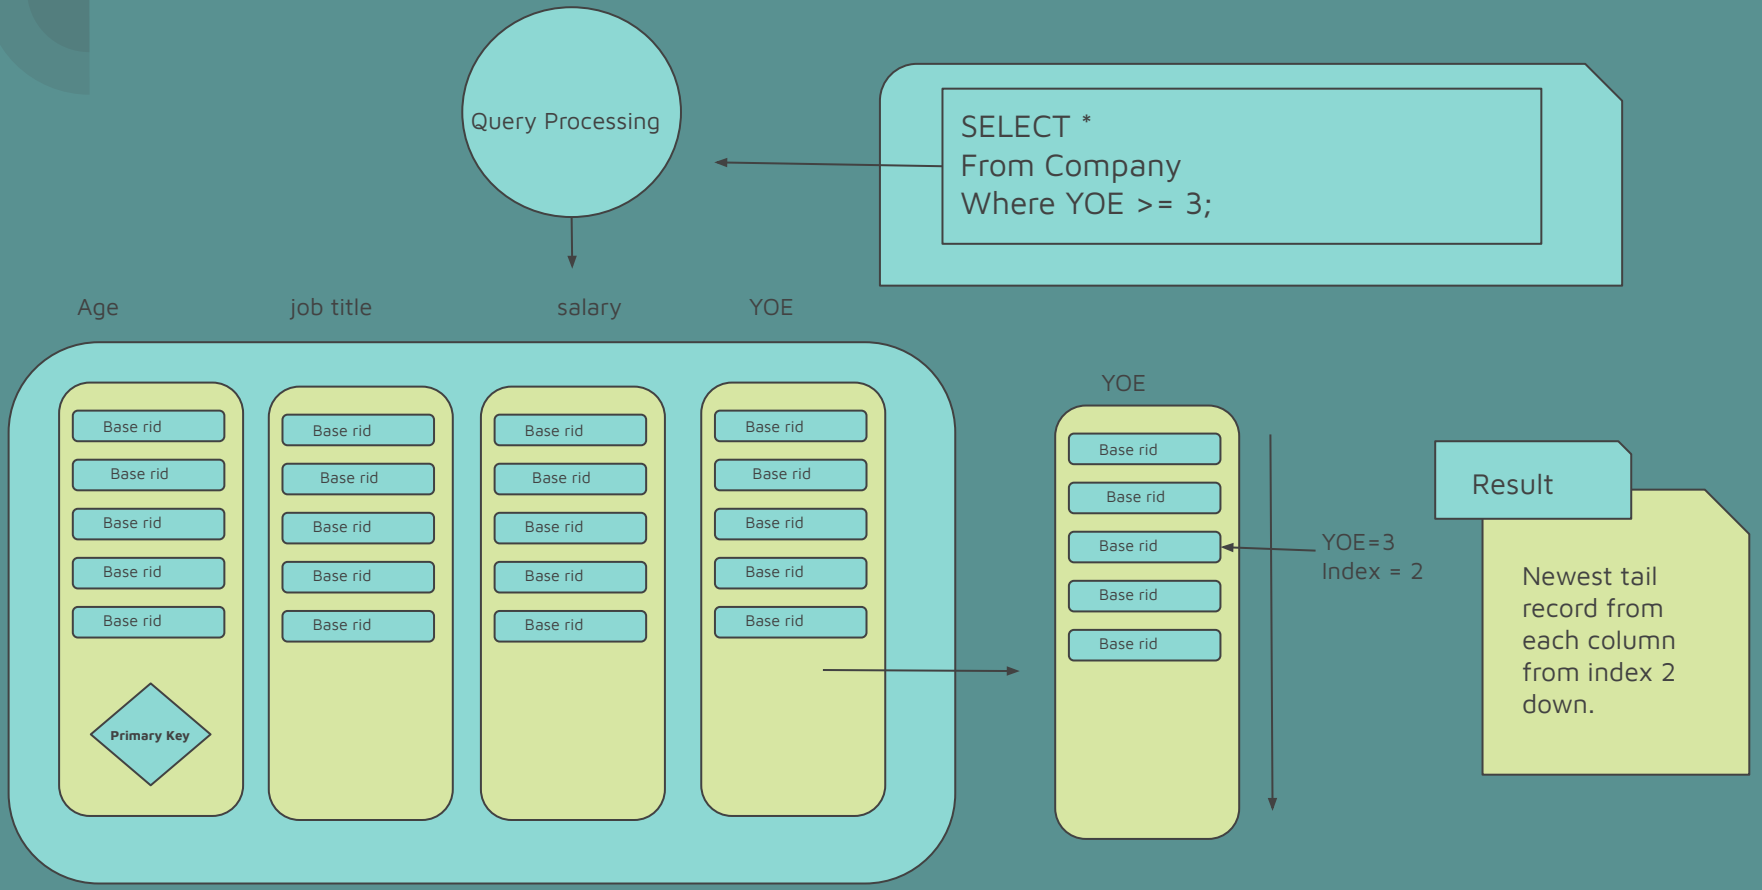

Things to improve

- Speed, this can be improved through optimizing what data-structure we use throughout our system.

- Switch from a cumulative database to a non-cumulative database

Q & A Time