

L-Store Milestone 1

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

Workflow

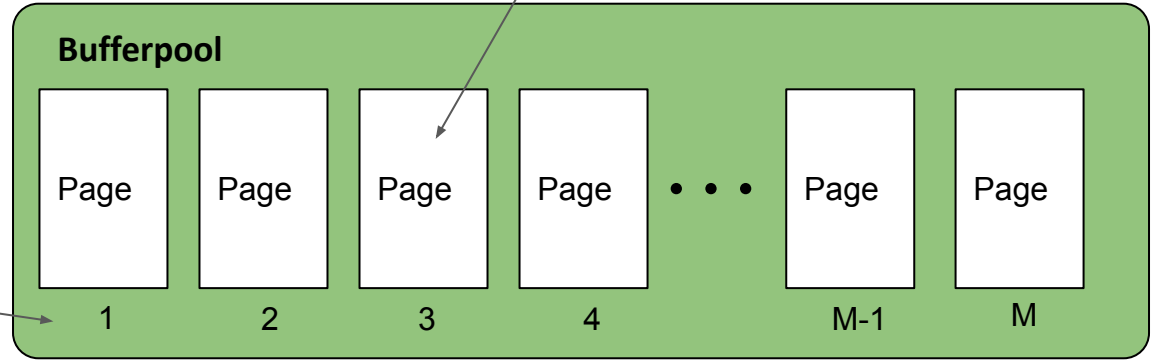
- Worked from two ends, then integrated:
 - Page-up (lower-level)
 - User-down (higher-level)
- Agreed on API to interface between higher level and lower level components, enabling us to work from both ends somewhat independently
- Agreed on using cumulative tail records

Database Design

Pages & Logical Pages


Page: object representing a physical page

Page ID: index of physical page in bufferpool



metadata page IDs

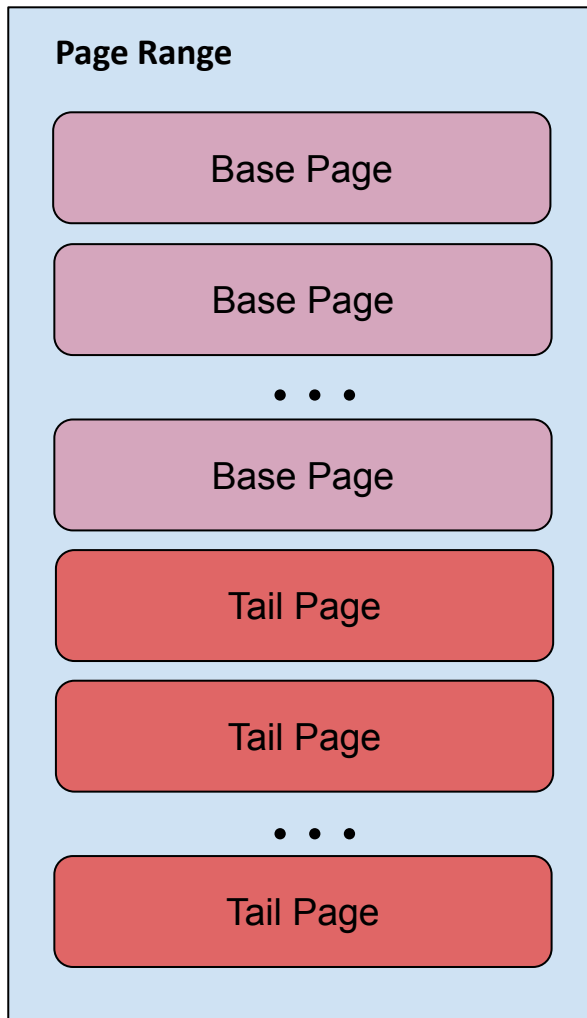
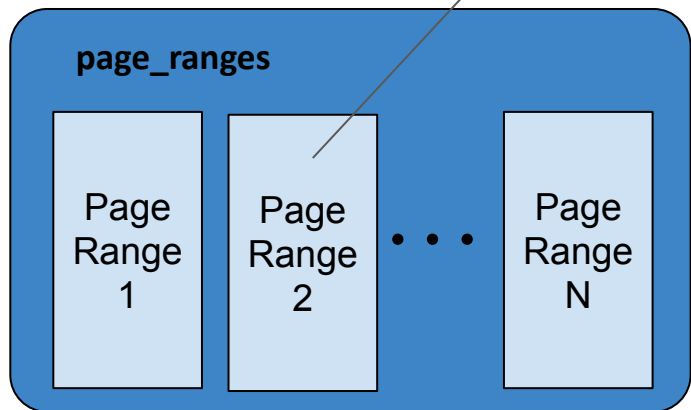
data page IDs



Page Ranges

Structure used to manage
base and tail pages

Allocated on as-needed
basis



Fixed maximum
number of base pages
(positions reserved on
page range creation)

Unlimited tail pages

Table

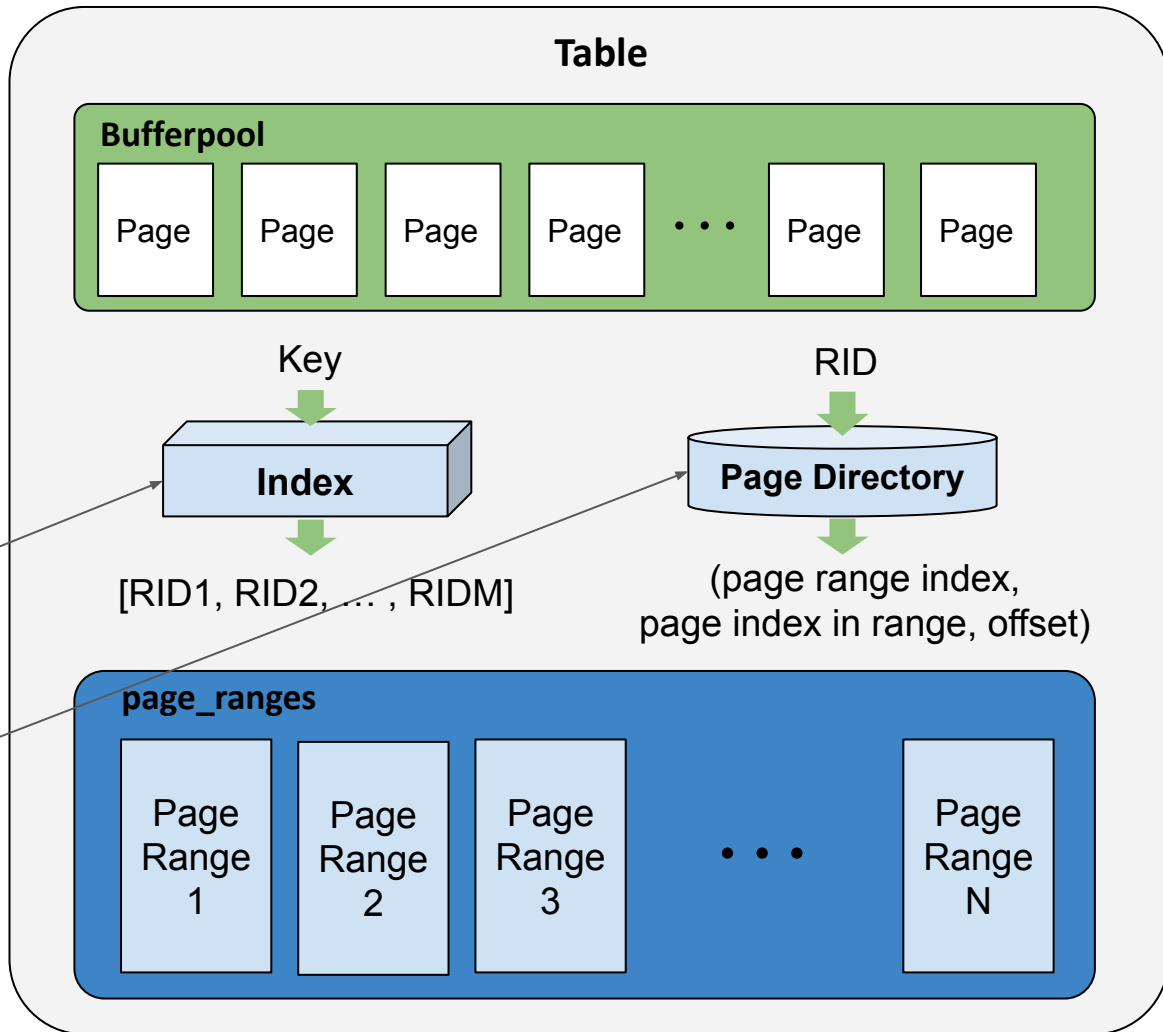
Connects pages and queries

Manages page ranges and allocates pages as needed

Site of physical page storage

Index: maps column values to RIDs
O(1) access time

Page Directory (hash table) :
maps RID to physical location
O(1) access time

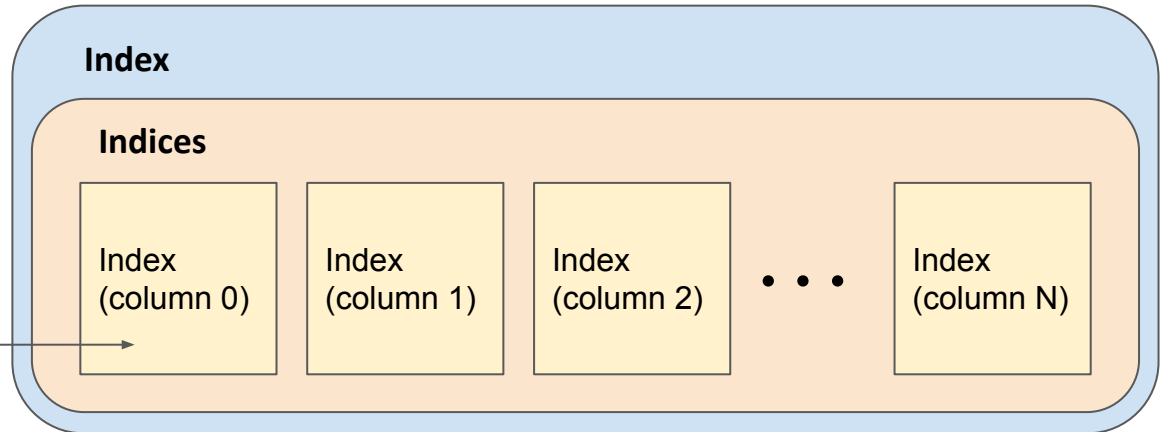


Index

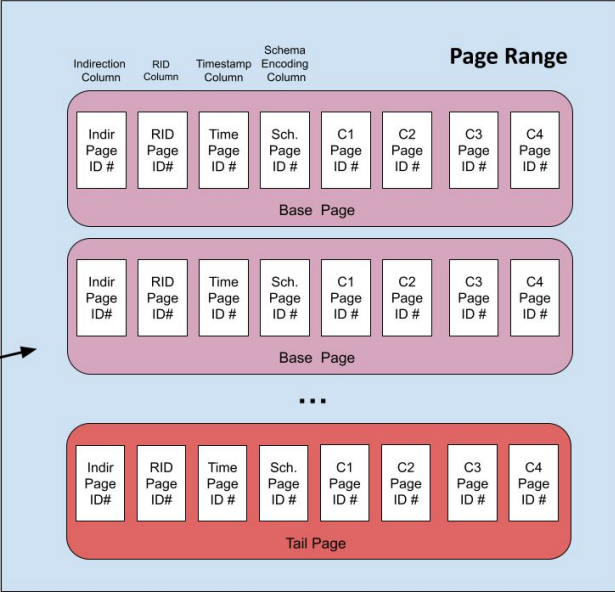
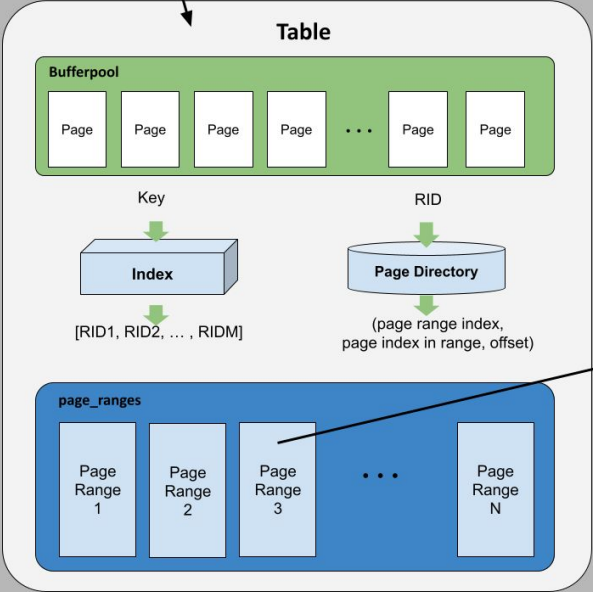
Hash tables for each data column index: $O(1)$ access

Vital for query performance: Selecting 10,000 records went from finishing on the order of minutes (using scanning) to on the order of seconds

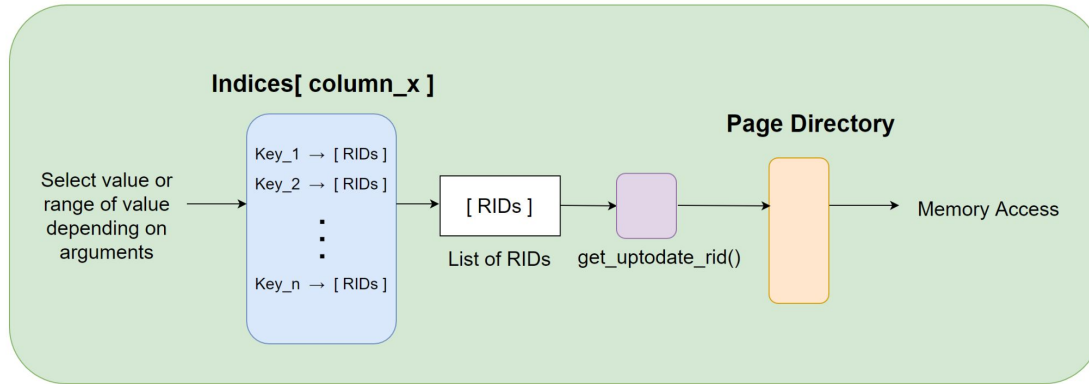
indices[i]: maps a value to an array of RIDs of the records with that value in the i 'th column



Database

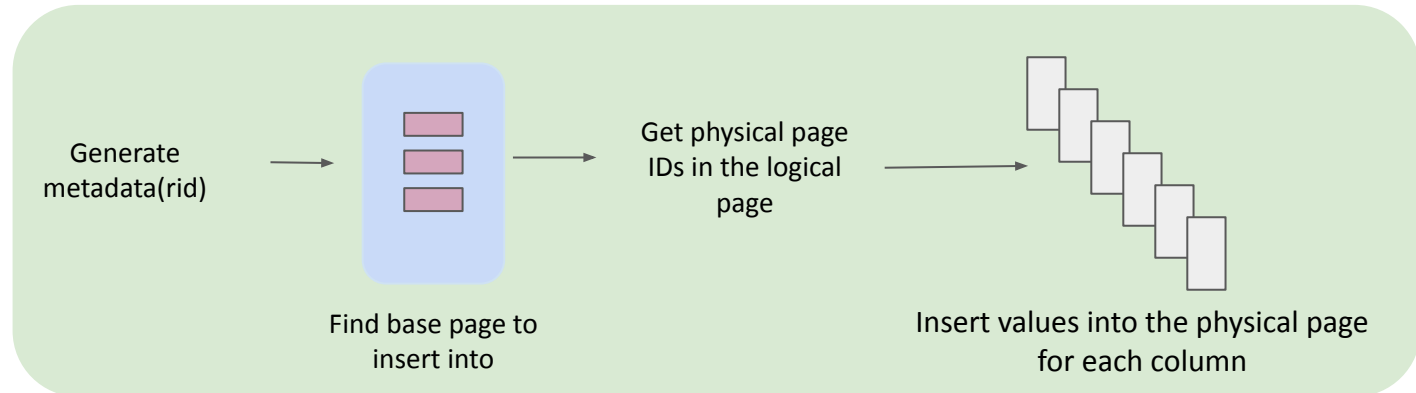


Query Logic: Select & Sum

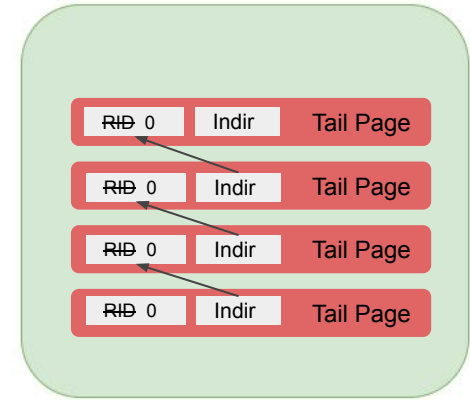
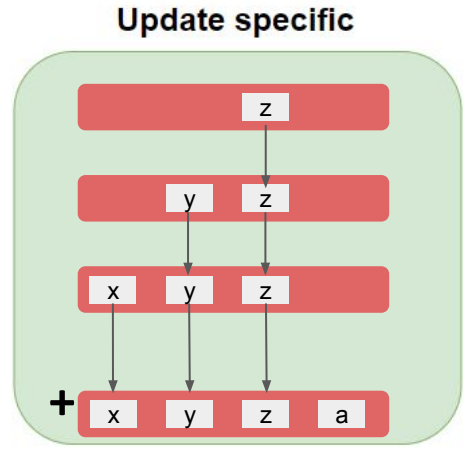
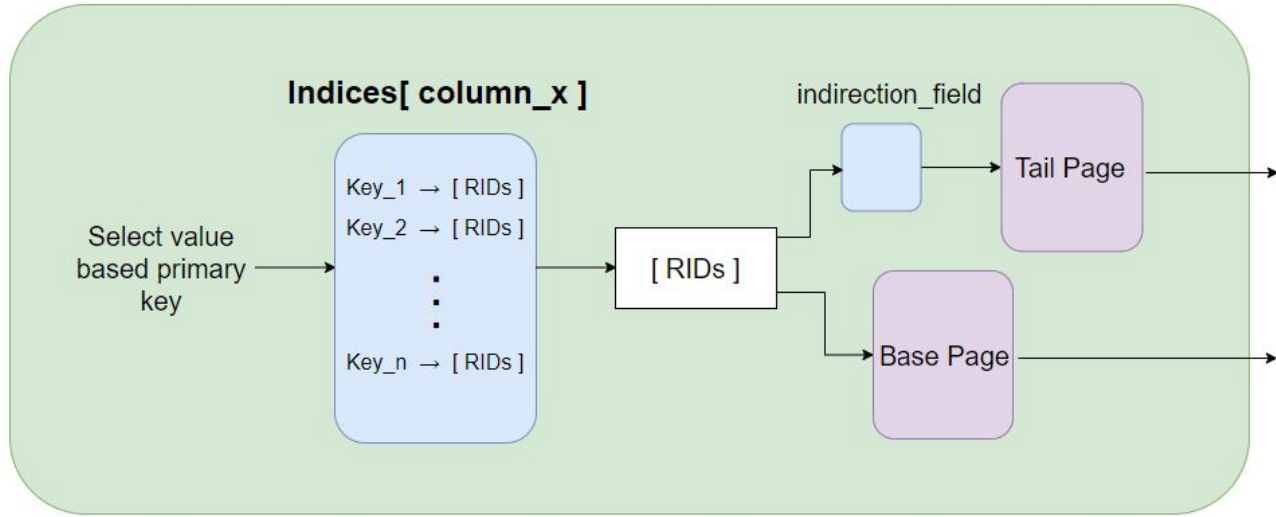


**Select and Sum
(Read Focused)**

**Insert
(Write Focused)**



Query Logic: Updates & Deletes




Delete specific

Performance


Sum Workloads

Without Updates




Insert Workloads

Without Updates




Select Workloads

Without Updates



Delete Workloads

Without Updates



Hardware

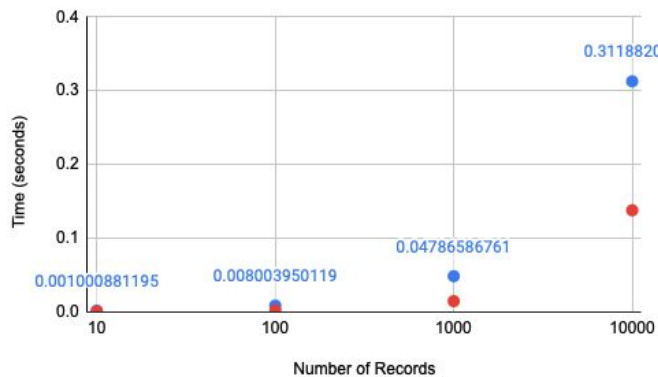
● (Windows)
Intel Core i7,
1.8GHz, 8GB,
8MB L3 Cache

● (Mac)
Apple M1,
3.2 GHz,
16GB, 12 MB
L2 Cache

Workloads:
custom


Sum Workloads

With Updates




Insert Workloads

With Updates



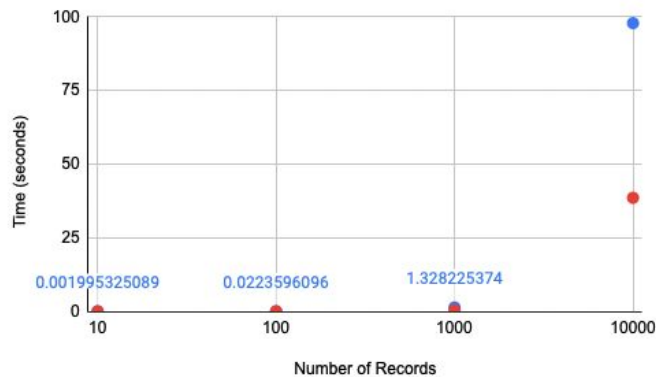
Select Workloads

With Updates



Delete Workloads

With Updates




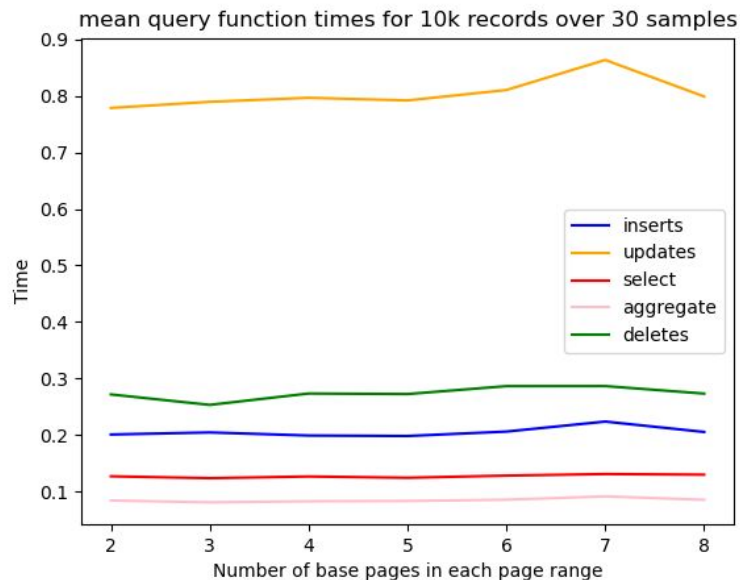
Hardware

● (Windows)
Intel Core i7,
1.8GHz, 8GB,
8MB L3 Cache

● (Mac)
Apple M1,
3.2 GHz,
16GB, 12 MB
L2 Cache

Workloads:
custom

Testing: Page Range Size



Workload: `__main__.py`

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

Q & A