CSE462/562: Database Systems (Spring 22) Lecture 7: Index 2/22/2022



Index

- Sometimes, we want to retrieve records by specifying the values in one or more fields
 - Find all students in CSE
 - Find all students admitted in year 2021
- Not very efficient to handle with heap file/sorted file
 - Heap file: always need to linear scan
 - Sorted file: only (somewhat) efficient for the sorted column
 - i.e., can't use binary search on a file sorted on major for specified adm_year



Index

sid

100

101

102

103

- An index: a data structure that speeds up search on a few fields on a relation
 - Maps index key k to data entry k*
 - Any subset of the columns of a relation can be the index key k
 - Index key is not (candidate/primary) key; doesn't have to be unique
 - Data entry k^*
 - e.g., the data record itself
 - Store the key k with the data entry k^* ?
 - Sometimes we do, sometimes we don't
 - Essentially an associative container, but more with more functionalities



Index classification

- Representation of data entries in index
 - i.e., what kind of info is the index actually storing?
 - 3 alternatives
- What selections does it support
- Indexing techniques: tree/hash/other
- Primary vs. Secondary Indexes
 - Unique indexes
- Clustered vs. Unclustered Indexes
- Single Key vs. Composite Indexes

Alternatives for the data entry k^* in index

- Three alternatives:
 - Alternative 1: the record itself (with its key k)
 - Alternative 2: <*k*, record ID of a matching record>
 - Alternative 3: <*k*, list of record IDs of matching records>
- Choice of the alternative is orthogonal to the indexing technique
 - Example of indexing techniques: B+-Tree, hash index, R-Tree, KD-Tree, and etc...
- A heap/sorted file can have multiple indexes
 - e.g., a B-tree index on adm_year and a hash index on major for the heap file of student relation
 - each usually stored as a separate file
 - usually at most one alternative-1 index per file (why?)

More on the alternatives of the data entries in index

- Alternative 1: actual data record (with its key k^*)
 - If this is used, it is another file organization for data records (just like heap file/sorted file)
 - At most 1 alt-1 index
 - Good: avoids record id/pointer lookups
 - Bad: less efficient to maintain for insertion/deletion/update
- Alternative 2 & 3

<*k*, record id of a matching record> or <*k*, list of record ids of matching records>

- Good: Can have multiple alt-2/alt-3 indexes
- Good: more efficient to maintain than alternative 1
- Bad: additional record id/pointer lookup (usually random I/O)
 - How to work around it? Include non-key columns.
- Alt-3 is more compact than alt-2, but the variation in data entry size can be much larger
 - Harder to deal with when they need to be split/merged
- Alt-3: key skew could lead to extremely long record id lists
 - Workaround: split them into shorter alt-3 data entries that fit into individual data pages

Index operations

- Inserts a data entry the index
- Deletes a data entry from the index
- Updates the value of a data entry
 - Can you change the index key of a data entry?
- Search and scan
 - Point lookup: find the data entry (entries) of a *search key*
 - Range scan: enumerate all the data entries in a range of *search keys*
 - e.g., $adm_year \in [2020, 2021]$, $adm_year > 2020$, $adm_year \le 2015$
 - sometimes the search key is a subset of the index key
 - Full index scan: enumerate all data entries in an index
 - Might be useful for ordering/efficiency
 - Other search operations:
 - String prefix matching
 - 2-D, 3-D, or higher dimensional range search
 - ...

Index Types

- Tree and hash indexes are the two most common categories of indexes
 - More details in the next 3-4 lectures
 - Example: B-Tree and static hash index

Tree-based indexes



- Leaf pages contain data entries, and are chained (prev & next page ids)
- Internal pages have index entries; only used to direct searches
- Good for equality and range selection
 - Results are ordered by index key

Example: B-Tree index



- Technically, this is the B+-Tree index, not the original B-Tree
 - Difference: B+-Tree only stores keys rather than data entries in internal nodes
 - But most DBMS uses B+-Tree, but use the term B-Tree...

Hash-based indexes

- Good for equality selections.
- Index is a collection of *buckets*.
 - Bucket = *primary* page plus zero or more *overflow* pages.
 - Buckets contain data entries.
- Hashing function h: h(r) = bucket in which (data entry for) record r belongs.
 h looks at the *index key* fields of r.
 - No need for "index entries" in this scheme.

Example: static hashing index

- Fixed number of primary pages = # of buckets (denoted as M)
 - allocated sequentially; never de-allocated
 - allocate overflow pages if needed
- h(k) % M = the bucket id for a data entry with index key k.



Clustered vs unclustered index

- Clustered index
 - An index over a file such that the order of the data records is the same as, or "close to" that of the index data entries
 - A file can only be clustered on one index key
 - Sorted file can be used for clustering, but may be expensive to maintain
 - Can we use heap file? Yes, but with some tricks.
 - Using Alternative 1 in a B+-tree implies clustered, *but not vice-versa*.
 - aka clustered file

Clustered vs unclustered index

- Assume alternative 2 for data entries, and data records are stored in a heap file.
 - To build clustered index
 - first sort the heap file, with some free space on each block for future updates/inserts.
 - The percentage of free space in the initial sort/append is called *fill factor*
 - Overflow pages may be needed for inserts/updates.
 - Thus, the order of data records is "close to", if not not identical to, the sort order.



Access cost of clustered vs unclustered index

- Cost of accessing data records through index varies *greatly* based on whether index is clustered!
 - e.g. range scan with n matching data records in a B-Tree
 - assuming we ignore the buffer pool's effect
 - clustered: $H + \left[\frac{n}{M}\right]$ I/Os

• unclustered:
$$H + \left[\frac{n}{B}\right] - 1 + n I/Os$$



Tradeoffs between clustered and unclustered indexes

- What are the tradeoffs?
- Clustered Pros
 - Efficient for range searches for records: sequential access in a sorted file
 - May be able to do some types of compression
 - Locality benefits
- Clustered Cons
 - Expensive to maintain (on the fly or sloppy with reorganization)
- Unclustered
 - Pros: easy and efficient to maintain, allow multiple indexes
 - Cons: expensive for range scans for records: 1 random IO for each matching record.

Primary, secondary and unique index

- Primary index: index key contains the primary key
 - e.g., for student table, an index over (sid) is its primary index
 - at most one per relation
- Unique index: index key contains a candidate key
 - Primary index is a unique index, but not vice versa
 - Can be clustered or unclustered.
- Secondary index (not well-defined but often used)
 - It may have different meanings
 - an index that is not indexed over the primary key
 - unclustered
 - or both

Unconventional "index"

- There might be alternative file organization also considered/called as "index"
 - e.g., columnstore index in MS SQL Server
 - Good compression, fast scan, but more expensive to update in general



- What it really means:
 - It may be used as the primary storage format (aka clustered columnstore)
 - i.e., may be thought of as a clustered file or a file organization
 - It may also be used as a copy of the (subset of) data (aka unclustered columnstore)
 - i.e., may be thought of as a secondary and unclustered index
- Takeaway: always read the fine prints of DBMS documentation CSE462/562 (Spring 2022): Lecture 7

Summary

- This lecture
 - Index classification and types
 - Index operations
 - HW2 released today
 - No submission needed
 - Solution will be released in one week
- Next time
 - Hash indexes