

Database Systems (資料庫系統)

October 4, 2003


Lecture #4

By Hao-hua Chu (朱浩華)

1

Ubicomp Project of the Week: Flexible Displays

- Displays using plastic substrate (rather than glass) are bendable and rollable.
 - Credit cards
- Flexible display has a wearable version – works on fabrics.
(www.visson.net)
 - Your cloth is a display.
 - Invisible Cloak, like in Harry Potter Movie (U. Tokyo)



2

Course Administration

- Homework #1: pick them up
- Homework #2: drop them off at the end of lecture or tomorrow at TA office hour
- Homework #3: post on the course homepage – due next Tuesday.
- Next Week Reading: Chapter 9

3

Overview of Storage & Indexing

Chapter 8

4

Outline

- Types of external Storage Devices
- File organizations
 - Question: how a file of records are stored on external storage device (e.g., a disk)?
 - Heap file, Sorted file
 - Indexing data structures
 - Question: how to speed up access to needed records on a disk?
 - Tree-based indexing, Hash-based indexing
- Comparison on file organizations
 - Question: which one is better/worse in performance?
- Indexes and Performance
 - Question: how to use indexing for better performance?

5

Data on External Storage

- External Storage: Offer persistent data storage
 - Unlike physical memory, data saved on a persistent storage is not lost when the system shutdowns or crashes.
- Magnetic Disks: Can retrieve random page at fixed cost
 - \$1 per Gigabyte
 - But reading several consecutive pages is much cheaper than reading them in random order
- Tapes: Can only read pages in sequence
 - \$0.3 per Gigabyte
 - Cheaper than disks; used for archival storage
- Other types of persistent storage devices:
 - Optical storage (CD-R, CD-RW, DVD-R, DVD-RW)



6

Definitions

- A [record](#) is a tuple or a row in a relation table.
 - Fixed-size records or variable-size records
- A [file](#) is a collection of records.
 - Store one table per file, or multiple tables in the same file
- A [page](#) is a fixed length block of data for disk I/O.
 - A file is consisted of pages.
 - A data page also contains a collection of records.
 - Typical page sizes are 4 and 8 KB.

7

File Organization

- [File organization](#): Method of arranging a file of records on external storage.
 - [Record id \(rid\)](#) is used to locate a record on a disk, e.g., (page id, slot number = i-th record in that page)
 - [Indexes](#) are data structures that allow us to efficiently search rids of given values in [index search key](#) fields
- [Architecture \(DB Storage and Indexing\)](#):
 - [Disk Space Manager](#) allocates/de-allocates spaces on disks.
 - [Buffer manager](#) moves pages between disks and main memory.
 - [File and index layers](#) organize records on files, and manage the indexing data structure.

8

Alternative File Organizations

Many alternatives exist, *each ideal for some situations, and not so good in others:*

- Heap files: Records are unsorted. Suitable when typical access is a file scan retrieving all records without any order.
 - Fast update (insertions / deletions)
- Sorted Files: Records are sorted. Best if records must be retrieved in some order, or only a `range` of records is needed.
 - Examples: employees are sorted by age.
 - Slow update in comparison to heap file.
- Indexes: Data structures to organize records via trees or hashing.
 - For example, create an index on employee age.
 - Like sorted files, they speed up searches for a subset of records that match values in certain ("search key") fields
 - Updates are much faster than in sorted files.

9

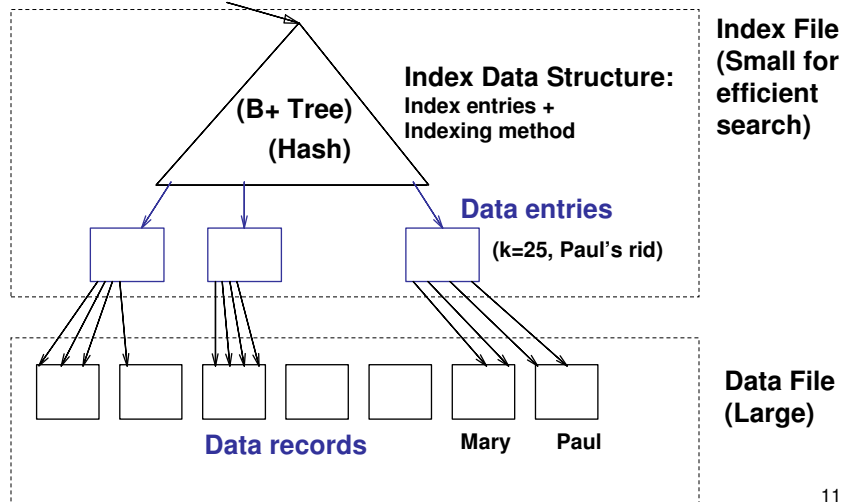
Indexes

- An index on a file speeds up selections on the search key fields for the index.
 - Any subset of the attributes of a table can be the search key for an index on the relation.
 - Search key is not the same as candidate key (minimal set of fields that uniquely identify a record in a relation).
 - Example: employee age is not a candidate key.
- An index file contains a collection of data entries (called **k***).
 - Quickly search an index to locate a data entry with a key value k.
 - Example of a data entry: <age, rid>
 - Use the data entry to find the data record.
 - Example of a data record: <name, age, salary>
 - Create multiple indexes on the same data records.
 - Example indexes: age, salary, name

10

Indexing Example

Search key value: find employees with age = 25



11

Alternatives for Data Entry k^*

- Three alternatives for what to store in a data entry:
 - (Alternative 1): Data record with key value k
 - Example: <age, name, salary>
 - (Alternative 2): < k , rid of data record with search key value k >
 - Example: <age, rid>
 - (Alternative 3): < k , list of rids of data records with search key k >
 - Example: <age, rid_1, rid_2, ...>
- Choice of alternative for data entries is independent of the indexing method used to locate data entries with a given key value k .
 - Indexing method takes a search key and finds the data entries matching the search key.
 - Examples of indexing methods: B+ trees or hashing.

12

Alternatives for Data Entries (Contd.)

- **Alternative 1: data record with key value k**
 - Data entries are also the data records.
 - At most one index on a given collection of data records can use Alternative 1.
 - Otherwise, data records are duplicated, why?
 - If data records are very large, # of pages containing data entries is high.
 - This may lead to less efficient search.

13

Alternatives for Data Entries (Contd.)

- **Alternatives 2 and 3: $\langle k, \text{rid or rid-list of data record(s)} \rangle$ with search key value k**
 - Data entries typically much smaller than data records.
 - More efficient search than Alternative 1. Why?
 - Alternative 3 more compact than Alternative 2,
 - Lead to variable sized data entries (size of rid-list is not fixed)

14

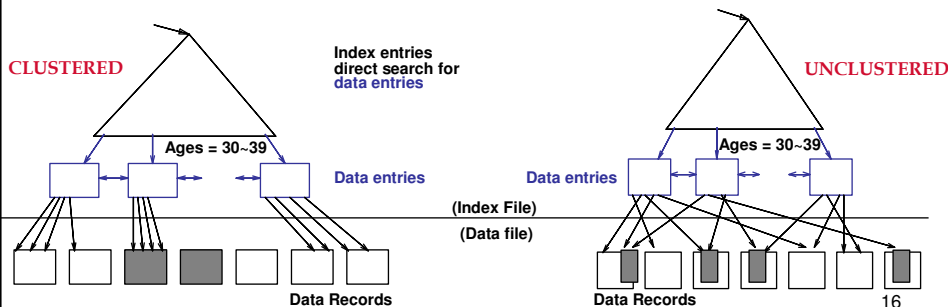
Index Classification

- **Primary vs. secondary:** If search key contains primary key, then called primary index; otherwise it is called secondary index.
- **Clustered vs. unclustered:** If order of data records is the same as, or close to the order of data entries, then it is called clustered index.
 - Alternative 1 implies clustered; in practice, clustered also implies Alternative 1.
 - One clustered index and multiple unclustered indexes
 - Why is this important?
 - Consider the cost of range search query (e.g., find all records $30 < \text{age} < 39$)

15

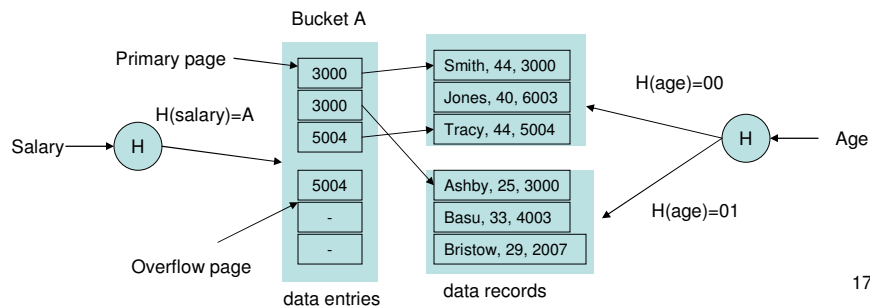
Clustered vs. Unclustered Index

- Cost of retrieving data records through index varies *greatly* based on whether index is clustered or not!
- Examples: retrieve all the employees of ages 30~39.
 - What is the cost (# disk page I/Os) of clustered index?
 - What is the cost of unclustered index?



Hash-Based Indexes

- Good for equality selections.
 - Data entries (key, rid) are grouped into buckets.
 - Bucket = *primary page* plus zero or more *overflow pages*.
 - *Hashing function h*: $h(r)$ = bucket in which record r belongs. h looks at the *search key* fields of r .
 - If Alternative (1) is used, the buckets contain the data records.




17

Hash-based Indexes (Cont)


- Search on key value:
 - Apply key value to the hash function -> bucket number
 - Retrieve the primary page of the bucket. Search records in the primary page. If not found, search the overflow pages.
 - Cost of locating rids: # pages in bucket (small)
- Insert a record:
 - Apply key value to the hash function -> bucket number
 - If all (primary & overflow) pages in that bucket are full, allocate a new overflow page.
 - Cost: similar to search.
- Delete a record
 - Cost: Similar to search.
- More details about hash-based index in Chapter 11.

18

B+ Tree Indexes

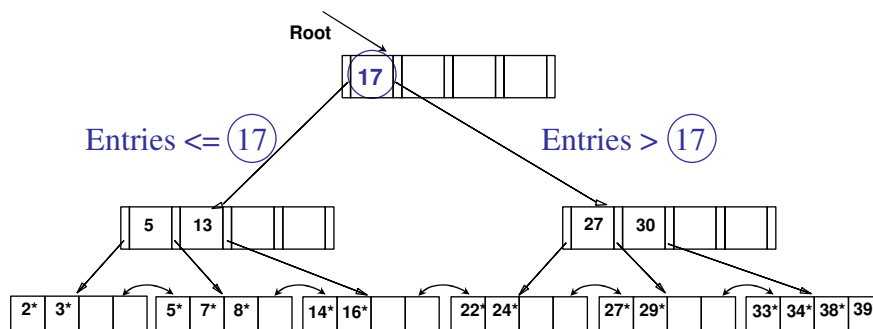


- ❖ Leaf pages contain *data entries*, and are chained (prev & next)
- ❖ Non-leaf pages contain *index entries* and direct searches:



19

Example B+ Tree



- Find 7*, 29*? 15* < age < 30*
- Insert/delete: Find data entry in leaf, then change it. Need to adjust parent sometimes.
 - And change sometimes bubbles up the tree (keep the tree balance)
- More details about tree-based index in Chapter 10.

20

Cost Model for Our Analysis

- Ignore CPU costs, for simplicity.
- Measuring IO costs: number of page I/O's
 - Ignores gains of pre-fetching a sequence of pages; thus, even I/O cost is only approximated.
- Average-case analysis; based on several simplistic assumptions.

B: The number of data pages

R: Number of records per page

D: (Average) time to read or write disk page

** Good enough to show the overall trends!*

21

Comparing File Organizations

- Heap files (random order; insert at eof)
- Sorted files, sorted on *<age, salary>*
- Clustered B+ tree file, Alternative (1), search key *<age, salary>*
- Heap file with unclustered B + tree index on search key *<age, salary>*
- Heap file with unclustered hash index on search key *<age, salary>*

22

Operations to Compare

- Scan: Fetch all records from disk
- Equality search
 - Example: find all employees with age = 23 and salary = 5000.
- Range selection
 - Example: find all employees with age > 35.
- Insert a record
 - Identify the page for inserting the record, fetch it, modify it, and write it back.
- Delete a record
 - Similar to insert.

23

Assumptions in Our Analysis

- Heap Files:
 - Equality selection on key; exactly one match.
- Sorted Files:
 - Files compacted after deletions.
- Indexes:
 - Alt (2), (3): data entry size = 10% size of data record

24

Heap Files

B: The number of data pages
R: Number of records per page
D: (Average) time to read or write one disk page

- Scan: $B \cdot D$
- Search with equality selection: $\frac{1}{2} B \cdot D$
- Search with range selection: $B \cdot D$
- Insert: $2D$
 - New record is inserted at the end of the file. Read/write out last page.
- Delete a record: $\text{search cost} + D$ (no compacting)
- Delete a record (with rid): $2 \cdot D$
 - search cost = D

25

Sorted Files

B: The number of data pages
R: Number of records per page
D: (Average) time to read or write one disk page

- Scan: BD
- Search with equality selection: $D \log_2(B)$
 - Binary search
- Search with range selection: $D \log_2(B) + \# \text{ matched tuples}$
 - Search for first matching record / page, then find all the qualifying records / pages in sequential order.
- Insert: $2D$: $\text{search} + BD$
 - Find the right position/page (search), make space for the inserting record by shifting all subsequent records by one slot, then insert.
- Delete a record: $\text{search} + BD$
 - Search the record, delete it, shift all subsequent records by one slot.

26

Clustered B+ Tree File

Tree: 2/3 occupancy (this is typical). # physical data pages = $1.5B$

B: The number of data pages
R: Number of records per page
D: (Average) time to read or write one disk page

- Scan: $1.5 BD$
- Search with equality selection: $D \log_F(1.5 B)$
 - F is number of children per B+ tree node
- Search with range selection: $D \log_F(B) + \# \text{ matched tuples}$
 - Search for first matching record (page), then find all the qualifying records / pages in sequential order.
- Insert: $2D$: search + D
 - Find the right position (page), insert + write out the modified page. No need to shift records -> empty data entries in modified page.
- Delete a record: search + D
 - Search record, delete record, and write back modified page. No need to shift.

27

Heap File with Unclustered Tree Index

Tree: 67% occupancy (this is typical). Data entry is 1/10 of data record. # data entry pages = $0.15B$

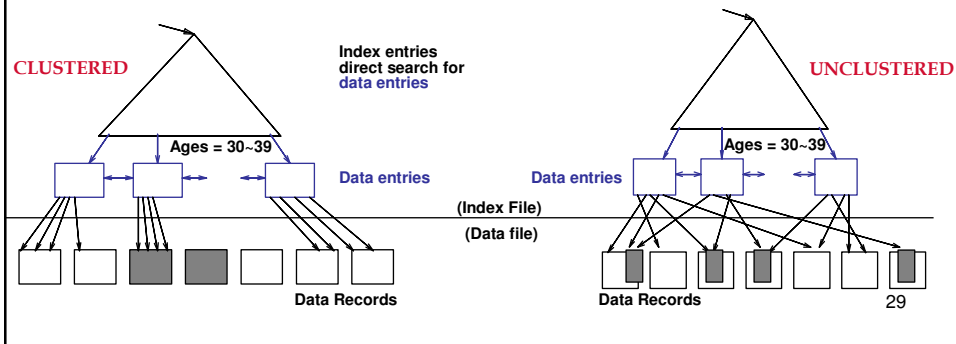
B: The number of data pages
R: Number of records per page
D: (Average) time to read or write one disk page

- Scan (in-order): $0.15 BD + RBD = BD (0.15 + R)$, why?
 - Unorder scan: BD
- Search with equality selection: $D (\log_F(0.15 B) + 1)$
 - Search for data entry takes $\log_F(0.15 B)$ pages + one read on data record page.
- Search with range selection: $D (\log_F(0.15B) + \# \text{ matched tuples})$
 - Search for first matching data entry, then find all the qualifying entries in sequential order. But each data entry may point to a data record on a different data page.
- Insert: search + $2D$
 - One read/write to heap file page + search + one write to data entry page.
- Delete a record: search + $2D$

28

Clustered vs. Unclustered Index

- Cost of retrieving data records through index varies *greatly* based on whether index is clustered or not!
- Examples: retrieve all the employees of ages 30~39.
 - What is the cost (# disk page I/Os) of clustered index?
 - What is the cost of unclustered index?



Heap File with Unclustered Hash Index

Hash: No overflow buckets.
 80% page occupancy
 => # data entry pages
 = $0.125B$

B: The number of data pages
R: Number of records per page
D: (Average) time to read or write one disk page

- Scan (in-order): $0.125 BD + RBD = BD (0.125 + R)$
 - Unorder scan: BD
- Search with equality selection: $2D$
 - Hash function + read data entry page + read data record page
- Search with range selection: BD
 - Hash function is useless, do unorder scan.
- Insert: $4D$
 - Read/write on heap file page + read/write the data entry page.
- Delete a record: $4D$
 - Same as Insert

Cost of Operations

	(a) Scan	(b) Equality	(c) Range	(d) Insert	(e) delete
(1) Heap	BD	0.5BD ☹	BD ☹	2D ☺	Search + D ☹
(2) Sorted	BD	$\log_2(B)$	$\log_2(B) + \#matches$ ☺	Search + BD ☹	Search + BD ☹
(3) Clustered Tree Index	1.5 BD	D ☺ ($\log_F(1.5B)$)	$\log_F(1.5B) + \#matches$ ☺	Search + D ☺	Search + D ☺
(4) Unclustered Tree index	BD(R+0.15) ☹	D (1 + $\log_F(0.15B)$) ☺	D ($\log_F(0.15B) + \#matches$) ☹	Search + 2D ☺	Search + 2D ☺
(5) Unclustered Hash Index	BD(R+0.125) ☹	2D ☺	BD ☹	4D ☺	4D ☺

**No one file organizations is uniformly superior*

31

General Guidelines

- An index supports efficient retrieval of data entries satisfying a selection condition:
 - There are two types of selections: equality and range
 - Hash-based indexing is only optimized for equality selection, useless for range selection.
 - Tree-based indexing is better for both.
 - Tree-based clustering index is best for range selection.
- Clustered index can be expensive than unclustered index:
 - When inserting a new record into a full page, shift existing records into other pages => change data entries for these records => expensive.
 - Consider this as a tradeoff for more efficient range selection.

32

Understanding the Workload

- How to decide the best indexing for a table?
 - Need to understand the workload, how queries are evaluation, and concurrency control.
- For each query in the workload:
 - Which tables does it access?
 - Which fields are retrieved?
 - Which fields are involved in selection/join conditions? How selective are these conditions likely to be?
- For each update in the workload:
 - Which fields are involved in selection/join conditions? How selective are these conditions likely to be?
 - The type of update (INSERT/DELETE/UPDATE), and the fields that are affected.

33

Choice of Indexes

- What indexes should we create?
 - Which tables should have indexes?
 - What field(s) should be the search key?
 - Should we build several indexes?
- For each index, what kind of an index should it be?
 - Clustered or unclustered?
 - Hash index or Tree index?

34

Choice of Indexes (Contd.)

- **One approach:** Consider the most important queries in turn. Consider the best plan using the current indexes, and see if a better plan is possible with an additional index. If so, create it.
 - Obviously, this implies that we must understand how a DBMS evaluates queries and creates **query evaluation plans!**
 - For now, we discuss simple 1-table queries.
- Before creating an index, must also consider the impact on updates in the workload!
 - **Trade-off:** Indexes can make queries go faster, updates slower (because also have to update the indexes). Indexes also require disk space, too.

35

Index Selection Guidelines

- Attributes in WHERE clause are candidates for index keys.
 - Exact match condition suggests hash index.
 - Range query suggests tree index.
 - Clustering is especially useful for range queries; can also help on equality queries if there are many duplicates.
- Multi-attribute search keys should be considered when a WHERE clause contains several conditions.
 - Order of attributes is important for range queries.
 - Such indexes can sometimes enable **index-only** strategies for important queries.
 - For index-only strategies, clustering is not important!
- Try to choose indexes that benefit as many queries as possible. Since only one index can be clustered per relation, choose it based on important queries that would benefit the most from clustering.

36

Examples of Clustered Indexes

- B+ tree index on E.age can be used to get qualifying records.
 - How selective is the condition? (selective means % of qualified records)
 - Is this index useful?
- Consider the GROUP BY query.
 - Is E.age index good? Why not?
 - Clustered E.dno index may be better.
- Equality queries and duplicates:
 - Unclustering is bad in case of many qualified records.
 - Clustering on E.hobby helps!

```
SELECT E.dno
FROM Emp E
WHERE E.age>40
```


```
SELECT E.dno, COUNT (*)
FROM Emp E
WHERE E.age>10
GROUP BY E.dno
```

```
SELECT E.dno
FROM Emp E
WHERE E.hobby=Stamps
```

Composite Search Keys

- Search on a combination of fields.
 - **Equality query:** Every field value is equal to a constant value.
 - age=12 and sal =10
 - **Range query:** Some field value is not a constant.
 - age =13; or sal=10 and age > 5
 - The order of fields in composite key is important!
 - <sal, age>: data entries are sorted by sal first, then age.

Examples of composite key indexes using lexicographic order.



Composite Search Keys

- To retrieve Emp records with *age*=30 AND *sal*=4000,
 - an index on *<age,sal>* would be better than an index on *age* or an index on *sal*.
- If condition is: *20<age<30* AND *3000<sal<5000*:
 - Clustered tree index on *<age,sal>* or *<sal,age>*.
- If condition is: *age*=30 AND *3000<sal<5000*:
 - Clustered *<age,sal>* index much better than *<sal,age>* index. Why?
 - For *<age, sal>*, find the first data entry with (*age*=30, *sal*=3000) and the qualified entries are likely to be qualified. However, for *<sal, age>*, find the first data entry with (*sal*=3000, *age*=anything), subsequent entries can have any ages.
 - The order of fields in composite key is important!
- Composite indexes are larger, updated more often.

39

Index-Only Plans

- A number of queries can be answered without retrieving any records from one or more of the relations involved if a suitable index is available.

<E.dno>

```
SELECT D.mgr
FROM Dept D, Emp E
WHERE D.dno=E.dno
```

<E.dno,E.eid>
Tree index!

```
SELECT D.mgr, E.eid
FROM Dept D, Emp E
WHERE D.dno=E.dno
```

<E.dno>

```
SELECT E.dno, COUNT(*)
FROM Emp E
GROUP BY E.dno
```

<E.dno,E.sal>
Tree index!

```
SELECT E.dno, MIN(E.sal)
FROM Emp E
GROUP BY E.dno
```

<E.age,E.sal>
or
<E.sal, E.age>
Tree!

```
SELECT AVG(E.sal)
FROM Emp E
WHERE E.age=25 AND
E.sal BETWEEN 3000 AND 5000
```

Index-Only Evaluation (Contd.)

- Index-only evaluations are possible if the key is <dno,age> or we have a tree index with key <age,dno>
 - Which is better?
 - What if we consider the second query?
- consider selective-ness of condition vs. cost of sorting
 - Selective: <age, dno>
 - Not selective: <dno, age>

```
SELECT E.dno, COUNT (*)  
FROM Emp E  
WHERE E.age=30  
GROUP BY E.dno
```

```
SELECT E.dno, COUNT (*)  
FROM Emp E  
WHERE E.age>30  
GROUP BY E.dno
```

41

Summary

- Many alternative file organizations exist, each appropriate in some situation.
- If selection queries are frequent, sorting the file or building an *index* is important.
 - Hash-based indexes only good for equality search.
 - Sorted files and tree-based indexes best for range search; also good for equality search.
- Index is a collection of data entries plus a way to quickly find entries with given key values.

42

Summary (Contd.)

- Data entries can be actual data records, <key, rid> pairs, or <key, rid-list> pairs.
 - Choice orthogonal to *indexing technique* used to locate data entries with a given key value.
- Can have several indexes on a given file of data records, each with a different search key.
- Indexes can be classified as clustered vs. unclustered. Differences have important consequences for utility/performance.

43

Summary (Contd.)

- Understanding the nature of the *workload* for the application, and the performance goals, is essential to developing a good design.
 - What are the important queries and updates? What fields/relations are involved?
- Indexes must be chosen to speed up important queries (and perhaps some updates!).
 - Index maintenance overhead on updates to key fields.
 - Build indexes to support index-only strategies.
 - Clustering is an important decision; only one index on a given relation can be clustered!
 - Order of fields in composite index key can be important.

44