

Bharathidasan University

Centre for Differently Abled Persons

Khajamalai Campus, Tiruchirappalli
Tamil Nadu, India



Bachelor of Computer Applications
(For Students with Speech and Hearing Impairment)

Course: Database Management System
Unit - 3

Compiled By
Dr.M.Prabavathy
(Assistant Professor)
Dr.R.Nandhakumar

Modification of the Database

- ▶ Deletion of tuples from a given relation.
- ▶ Insertion of new tuples into a given relation
- ▶ Updating of values in some tuples in a given relation

File Organization

- ▶ The database is stored as a collection of *files*. Each file is a sequence of *records*. A record is a sequence of fields.
- ▶ One approach:
 - ▶ assume record size is fixed
 - ▶ each file has records of one particular type only
 - ▶ different files are used for different relations

This case is easiest to implement; will consider variable length records later.

Fixed-Length Records

- ▶ Simple approach:
 - ▶ Store record i starting from byte $n * (i - 1)$, where n is the size of each record.
 - ▶ Record access is simple but records may cross blocks
 - ▶ Modification: do not allow records to cross block boundaries

- ▶ Deletion of record i :
alternatives:

- ▶ move records $i + 1, \dots, n$
to $i, \dots, n - 1$
- ▶ move record n to i
- ▶ do not move records, but
link all free records on a
free list

record 0	10101	Srinivasan	Comp. Sci.	65000
record 1	12121	Wu	Finance	90000
record 2	15151	Mozart	Music	40000
record 3	22222	Einstein	Physics	95000
record 4	32343	El Said	History	60000
record 5	33456	Gold	Physics	87000
record 6	45565	Katz	Comp. Sci.	75000
record 7	58583	Califieri	History	62000
record 8	76543	Singh	Finance	80000
record 9	76766	Crick	Biology	72000
record 10	83821	Brandt	Comp. Sci.	92000
record 11	98345	Kim	Elec. Eng.	80000

Free Lists


- ▶ Store the address of the first deleted record in the file header.
- ▶ Use this first record to store the address of the second deleted record, and so on
- ▶ Can think of these stored addresses as **pointers** since they “point” to the location of a record.
- ▶ More space efficient representation: reuse space for normal attributes of free records to store pointers. (No pointers stored in in-use records.)

header				
record 0	10101	Srinivasan	Comp. Sci.	65000
record 1				
record 2	15151	Mozart	Music	40000
record 3	22222	Einstein	Physics	95000
record 4				
record 5	33456	Gold	Physics	87000
record 6				
record 7	58583	Califieri	History	62000
record 8	76543	Singh	Finance	80000
record 9	76766	Crick	Biology	72000
record 10	83821	Brandt	Comp. Sci.	92000
record 11	98345	Kim	Elec. Eng.	80000




Variable-Length Records

- ▶ Variable-length records arise in database systems in several ways:
 - ▶ Storage of multiple record types in a file.
 - ▶ Record types that allow variable lengths for one or more fields such as strings (**varchar**)
 - ▶ Record types that allow repeating fields (used in some older data models).
- ▶ Attributes are stored in order
- ▶ Variable length attributes represented by fixed size (offset, length), with actual data stored after all fixed length attributes
- ▶ Null values represented by null-value bitmap



Variable-Length Records: Slotted Page Structure



- ▶ **Slotted page** header contains:
 - ▶ number of record entries
 - ▶ end of free space in the block
 - ▶ location and size of each record
- ▶ Records can be moved around within a page to keep them contiguous with no empty space between them; entry in the header must be updated.
- ▶ Pointers should not point directly to record — instead they should point to the entry for the record in header.

Organization of Records in Files

- ▶ **Heap** – a record can be placed anywhere in the file where there is space
- ▶ **Sequential** – store records in sequential order, based on the value of the search key of each record
- ▶ **Hashing** – a hash function computed on some attribute of each record; the result specifies in which block of the file the record should be placed
- ▶ Records of each relation may be stored in a separate file. In a **multiple clustering file organization** records of several different relations can be stored in the same file
 - ▶ Motivation: store related records on the same block to minimize I/O

Sequential File Organization

- ▶ Suitable for applications that require sequential processing of the entire file
- ▶ The records in the file are ordered by a **search-key**

10101	Srinivasan	Comp. Sci.	65000	
12121	Wu	Finance	90000	
15151	Mozart	Music	40000	
22222	Einstein	Physics	95000	
32343	El Said	History	60000	
33456	Gold	Physics	87000	
45565	Katz	Comp. Sci.	75000	
58583	Califieri	History	62000	
76543	Singh	Finance	80000	
76766	Crick	Biology	72000	
83821	Brandt	Comp. Sci.	92000	
98345	Kim	Elec. Eng.	80000	




Sequential File Organization (Cont.)

- ▶ Deletion – use pointer chains
- ▶ Insertion – locate the position where the record is to be inserted
 - ▶ if there is free space insert there
 - ▶ if no free space, insert the record in an **overflow block**
 - ▶ In either case, pointer chain must be updated
- ▶ Need to reorganize the file from time to time to restore sequential order

10101	Srinivasan	Comp. Sci.	65000	
12121	Wu	Finance	90000	
15151	Mozart	Music	40000	
22222	Einstein	Physics	95000	
32343	El Said	History	60000	
33456	Gold	Physics	87000	
45565	Katz	Comp. Sci.	75000	
58583	Califieri	History	62000	
76543	Singh	Finance	80000	
76766	Crick	Biology	72000	
83821	Brandt	Comp. Sci.	92000	
98345	Kim	Elec. Eng.	80000	

32222	Verdi	Music	48000	
-------	-------	-------	-------	--



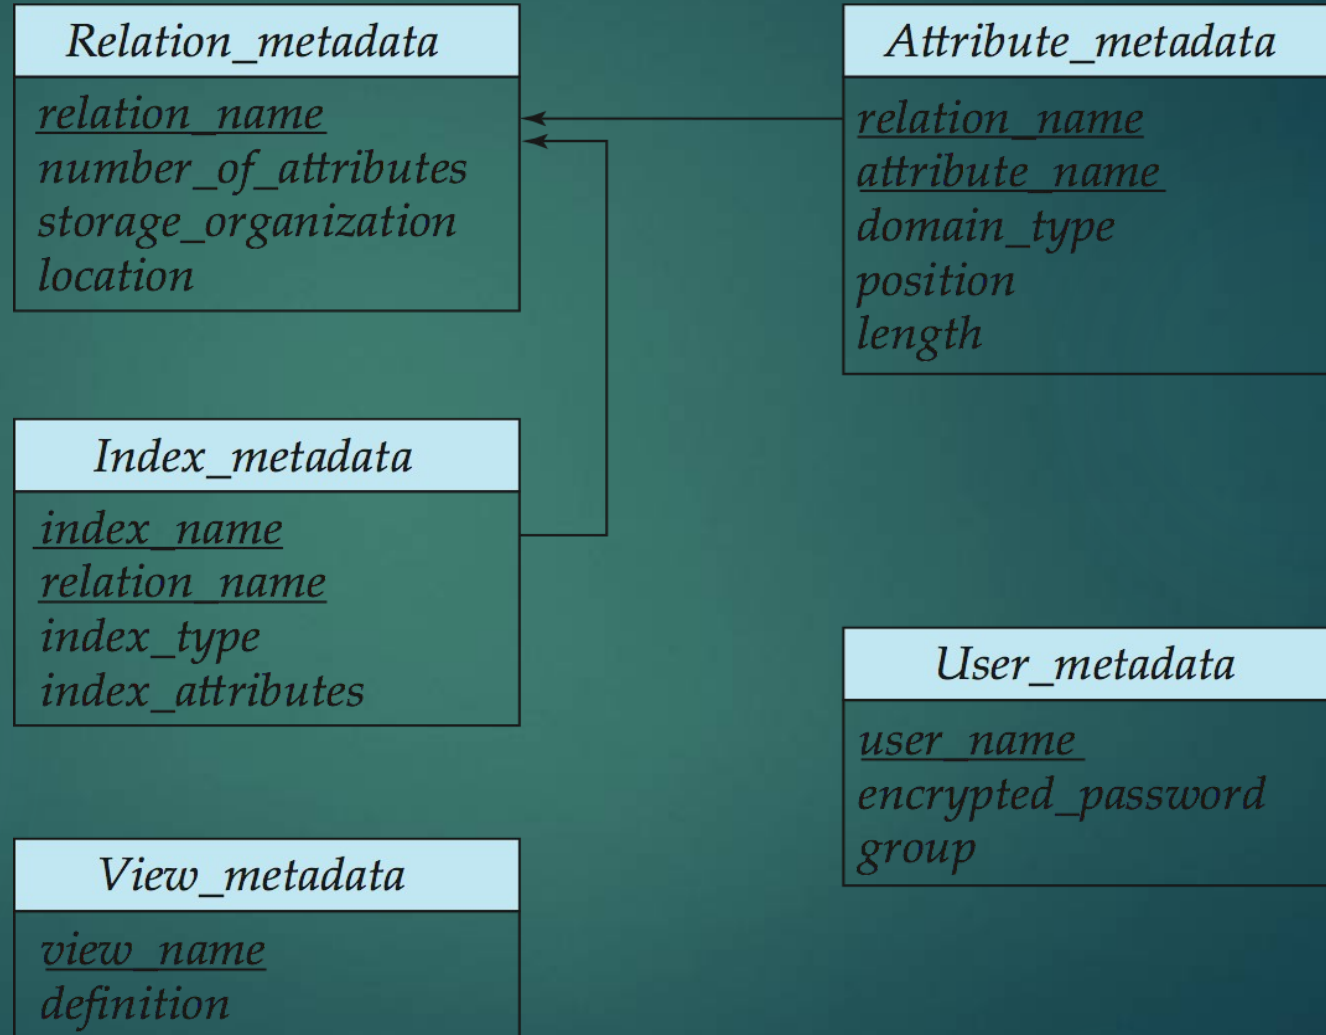
Data Dictionary Storage

The **Data dictionary** (also called **system catalog**) stores **metadata**; that is, data about data, such as

- ▶ Information about relations
 - ▶ names of relations
 - ▶ names, types and lengths of attributes of each relation
 - ▶ names and definitions of views
 - ▶ integrity constraints
- ▶ User and accounting information, including passwords
- ▶ Statistical and descriptive data
 - ▶ number of tuples in each relation
- ▶ Physical file organization information
 - ▶ How relation is stored (sequential/hash/...)
 - ▶ Physical location of relation
- ▶ Information about indices (Chapter 11)

Relational Representation of System Metadata

- ▶ Relational representation on disk
- ▶ Specialized data structures designed for efficient access, in memory



Storage Access

- ▶ A database file is partitioned into fixed-length storage units called **blocks**. Blocks are units of both storage allocation and data transfer.
- ▶ Database system seeks to minimize the number of block transfers between the disk and memory. We can reduce the number of disk accesses by keeping as many blocks as possible in main memory.
- ▶ **Buffer** – portion of main memory available to store copies of disk blocks.
- ▶ **Buffer manager** – subsystem responsible for allocating buffer space in main memory.

Buffer Manager

- ▶ Programs call on the buffer manager when they need a block from disk.
 1. If the block is already in the buffer, buffer manager returns the address of the block in main memory
 2. If the block is not in the buffer, the buffer manager
 1. Allocates space in the buffer for the block
 1. Replacing (throwing out) some other block, if required, to make space for the new block.
 2. Replaced block written back to disk only if it was modified since the most recent time that it was written to/fetched from the disk.
 2. Reads the block from the disk to the buffer, and returns the address of the block in main memory to requester.

Buffer-Replacement Policies

- ▶ Most operating systems replace the block **least recently used** (LRU strategy)
- ▶ Idea behind LRU – use past pattern of block references as a predictor of future references
- ▶ Queries have well-defined access patterns (such as sequential scans), and a database system can use the information in a user's query to predict future references
 - ▶ LRU can be a bad strategy for certain access patterns involving repeated scans of data
 - ▶ For example: when computing the join of 2 relations r and s by a nested loops
 - for each tuple tr of r do
 - for each tuple ts of s do
 - if the tuples tr and ts match ...
 - ▶ Mixed strategy with hints on replacement strategy provided by the query optimizer is preferable

Buffer-Replacement Policies (Cont.)

- ▶ **Pinned block** – memory block that is not allowed to be written back to disk.
- ▶ **Toss-immediate** strategy – frees the space occupied by a block as soon as the final tuple of that block has been processed
- ▶ **Most recently used (MRU) strategy** – system must pin the block currently being processed. After the final tuple of that block has been processed, the block is unpinned, and it becomes the most recently used block.
- ▶ Buffer manager can use statistical information regarding the probability that a request will reference a particular relation
 - ▶ E.g., the data dictionary is frequently accessed. Heuristic: keep data-dictionary blocks in main memory buffer
- ▶ Buffer managers also support **forced output** of blocks for the purpose of recovery (more in Chapter 16)

Deletion

- ▶ Delete all instructors

```
delete from instructor
```

- ▶ Delete all instructors from the Finance department

```
delete from instructor  
where dept_name = 'Finance';
```

- ▶ Delete all tuples in the *instructor* relation for those instructors associated with a department located in the Watson building.

```
delete from instructor  
where dept name in (select dept name  
                        from department  
                        where building = 'Watson');
```

Deletion (Cont.)

- ▶ Delete all instructors whose salary is less than the average salary of instructors

```
delete from instructor  
where salary < (select avg (salary)  
                  from instructor);
```

- Problem: as we delete tuples from deposit, the average salary changes
- Solution used in SQL:
 1. First, compute **avg** (*salary*) and find all tuples to delete
 2. Next, delete all tuples found above (without recomputing **avg** or retesting the tuples)

Insertion

- ▶ Add a new tuple to *course*

```
insert into course  
  values ('CA-04', 'Database Systems', 'BCA.', 4);
```

- ▶ or equivalently

```
insert into course (course_id, title, dept_name, credits)  
  values ('CA-04', 'Database Systems', 'BCA.', 4);
```

- ▶ Add a new tuple to *student* with *tot_creds* set to null

```
insert into student  
  values ('3003', 'Green', 'Finance', null);
```

Insertion (Cont.)

- ▶ Add all instructors to the *student* relation with *tot_creds* set to 0

```
insert into student  
    select ID, name, dept_name, 0  
from instructor
```

- ▶ The **select from where** statement is evaluated fully before any of its results are inserted into the relation.

Otherwise queries like

```
insert into table1 select * from table1
```

would cause problem

Case Statement for Conditional Updates

- ▶ Same query as before but with case statement

update *instructor*

```
    set salary = case  
    when salary <= 100000 then salary * 1.05  
    else salary * 1.03  
    end
```



▶ Thank You