


# Data Storage and Querying

Dr G.N.Wikramanayake  
 Dr Jeevani Goonetillake  
 University of Colombo School of Computing



© 2008, University of Colombo School of Computing



© 2008, University of Colombo School of Computing



## Physical View

- The DBMS must know
  - exact physical location
  - precise physical structure



### Employee record

A.B.C. De Silva		222, Galle Road, Colombo	
-----------------	--	--------------------------	--

Name (20 characters)      Address (40 characters)

650370690V Senior Lecturer
----------------------------

NID (10 char)      Designation (15 char)



© 2008, University of Colombo School of Computing



## File



- A collection of data records with similar characteristics
- Student list, course list at the university
- price list, stock records at a supermarket

Item#	Description	Shelf#	Max Stock	Reorder	Balance
1152	Milk	D01	100	25	70
1167	Bread	D01	70	25	30
1175	Sugar	G02	300	50	120
1172	Rice	G04	50	20	35




© 2008, University of Colombo School of Computing




## Physical Design

- Provide good performance
  - Fast response time
  - Minimum disk accesses



© 2008, University of Colombo School of Computing

## Disk Assembly



$$\text{Access time} = \text{seek time} + \text{rotational delay} + \text{transfer time}$$



© 2008, University of Colombo School of Computing



## Disks


- 6 disks (platters); 12 surfaces; 2 outer protected surfaces; 10 inner data surfaces (coated with a magnetic substance to record data)
- Each surface 200-400 concentric tracks
- Read/write heads placed over specific track; with one active at a time
- Set of corresponding tracks is a cylinder, i.e. track 1 of all 10 surfaces




© 2008, University of Colombo School of Computing



## Disks



© 2008, University of Colombo School of Computing



## Data Block

- A data block (sector) is the smallest unit of data defined within the database.
- block size may be defined by the `DB_BLOCK_SIZE`



© 2008, University of Colombo School of Computing



## Definitions

- Seek time
  - Average time to move the read-write head to the correct cylinder
- Rotational delay
  - Average time for the sector to move under the read-write head
- Transfer time
  - Time to read a sector and transfer the data to memory



© 2008, University of Colombo School of Computing



## More Definitions

- Logical Record
  - The data about an entity (a row in a table)
- Physical Record
  - A sector, page or block on the storage medium
- Typically several logical records can be stored in one physical record.



© 2008, University of Colombo School of Computing



## File Organization Techniques

- Three techniques
  - Serial or Heap (unordered)
  - Sorted
    - Sequential (SAM)
    - Indexed Sequential (ISAM)
  - Hashed or Direct or Random



© 2008, University of Colombo School of Computing



## Serial

- Used for temporary files such as transaction files, dump files

<u>Transaction#</u>	<u>Item#</u>	<u>Description</u>	<u>Quantity</u>
101	1152	Milk	01
101	1167	Bread	02
102	1167	Bread	01
102	1175	Sugar	01
103	1172	Rice	01
103	1152	Milk	01



© 2008, University of Colombo School of Computing



## Heap File

Tony Lama was the first record added,  
Digital was the last.

<u>ID</u>	<u>Company</u>	<u>Industry</u>	<u>SymbL</u>	<u>Price</u>	<u>Earns.</u>	<u>Dividnd.</u>
1767	Tony Lama	Apparel	TONY	45.00	1.50	0.25
1152	Lockheed	Aero	LCH	112.00	1.25	0.50
1175	Ford	Auto	F	88.00	1.70	0.20
1122	Exxon	Oil	XON	46.00	2.50	0.75
1231	Intel	Comp.	INTL	30.00	2.00	0.00
1323	GM	Auto	GM	158.00	2.10	0.30
1378	Texaco	Oil	TX	230.00	2.80	1.00
1245	Digital	Comp.	DEC	120.00	1.80	0.10



© 2008, University of Colombo School of Computing



## Heap File Characteristics

- Insertion
  - Fast: New records added at the end of the file
- Retrieval
  - Slow: A sequential search is required
- Update - Delete
  - Slow:
    - Sequential search to find the page
    - Make the update or mark for deletion
    - Re-write the page



© 2008, University of Colombo School of Computing



## Sequential

- Records are recoded in key sequence, but have no index
- Used for master files in a normal batch processing

<u>Item#</u>	<u>Description</u>	<u>Price</u>
1152	Milk	35.00
1167	Bread	30.00
1172	Rice	60.00
1175	Sugar	50.00



© 2008, University of Colombo School of Computing



## Sequential (Ordered) File


<u>ID</u>	<u>Company</u>	<u>Industry</u>	<u>SymbL</u>	<u>Price</u>	<u>Earns.</u>	<u>Dividnd.</u>
1122	Exxon	Oil	XON	46.00	2.50	0.75
1152	Lockheed	Aero	LCH	112.00	1.25	0.50
1175	Ford	Auto	F	88.00	1.70	0.20
1231	Intel	Comp.	INTL	30.00	2.00	0.00
1245	Digital	Comp.	DEC	120.00	1.80	0.10
1323	GM	Auto	GM	158.00	2.10	0.30
1378	Texaco	Oil	TX	230.00	2.80	1.00
1480	Conoco	Oil	CON	150.00	2.00	0.50
1767	Tony Lama	Apparel	TONY	45.00	1.50	0.25



© 2008, University of Colombo School of Computing



## Sequential Access



© 2008, University of Colombo School of Computing



## Sequential File Characteristics

- Older media (cards, tapes)
- Records physically ordered by primary key
- Use when direct access to individual records is not required
- Accessing records
  - Sequential search until record is found
- Binary search can speed up access
  - Must know file size and how to determine mid-point,



© 2008, University of Colombo School of Computing




## Search Sugar

- Sequence of file based on Description
- 4 sequential accesses to reach sugar
- 2 binary search accesses to reach sugar



© 2008, University of Colombo School of Computing



© 2008, University of Colombo School of Computing



## Inserting Records in SAM files

- Insertion
  - Slow:
    - Sequential search to find where the record goes
    - If sufficient space in that page, then rewrite
    - If insufficient space, move some records to next page
    - If no space there, keep bumping down until space is found
  - May use an “overflow” file to decrease time



© 2008, University of Colombo School of Computing



## Deletions and Updates to SAM

- Deletion
  - Slow:
    - Find the record
    - Either mark for deletion or free up the space
    - Rewrite
- Updates
  - Slow:
    - Find the record
    - Make the change
    - Rewrite



© 2008, University of Colombo School of Computing



## Binary Search to Find GM (1323)

	ID	Company	Industry	Symb.	Price	Earns.	Dividnd.
	1122	Exxon	Oil	XON	46.00	2.50	0.75
	1152	Lockheed	Aero	LCH	112.00	1.25	0.50
	1175	Ford	Auto	F	88.00	1.70	0.20
	1231	Intel	Comp.	INTL	30.00	2.00	0.00
1.	1245	Digital	Comp.	DEC	120.00	1.80	0.10
3.	1323	GM	Auto	GM	158.00	2.10	0.30
2.	1378	Texaco	Oil	TX	230.00	2.80	1.00
	1480	Conoco	Oil	CON	150.00	2.00	0.50
	1767	Tony Lama Apparel	TONY	45.00	1.50	0.25	

- Takes 3 accesses as opposed to 6 for linear search.



© 2008, University of Colombo School of Computing



## Indexed Sequential

- Disk (usually)
- Records physically ordered by primary key
- Index gives physical location of each record
- Records accessed sequentially or directly via the index
- The index is stored in a file and read into memory when the file is opened.
- Indexes must be maintained



© 2008, University of Colombo School of Computing



## Indexed Sequential Access

- Given a value for the key
  - search the index for the record address
  - issue a read instruction for that address
  - Fast: Possibly just one disk access



© 2008, University of Colombo School of Computing



## Indexed Sequential Access: Fast

Find record with key 777-13-1212

Key	Cyl.	Trck	Sect.
279-66-7549	3	10	2
452-75-6301	3	10	3
789-12-3456	3	10	4

222-66-7634
255-75-5531
279-66-7549
333-88-9876
382-32-0658
452-75-6301
701-43-5634
777-13-1212
789-12-3456

777-13-1212 < 789-12-3456  
Search Cyl. 3, Trck 10, Sect. 4 sequentially.



© 2008, University of Colombo School of Computing



## Inserting into ISAM files

- Not very efficient
  - Indexes must be updated
  - Must locate where the record should go
  - If there is space, insert the new record and rewrite
  - If no space, use an overflow area
  - Periodically merge overflow records into file



© 2008, University of Colombo School of Computing



## Deletion and Updates for ISA

- Fairly efficient
  - Find the record
  - Make the change or mark for deletion
  - Rewrite
  - Periodically remove records marked for deletion



© 2008, University of Colombo School of Computing



## Use ISAM files when:

- Both sequential and direct access is needed.
- Say we have a retail application like Foley's.
- Customer balances are updated daily.
- Usually sequential access is more efficient for batch updates.
- But we may need direct access to answer customer questions about balances.



© 2008, University of Colombo School of Computing



## Random

- Randomly organized file contains records stored without regard to the sequence of their control fields.
- Records are stored in some convenient order establishing a direct link between the key of the record and the physical address of that record




© 2008, University of Colombo School of Computing



## Direct or Hashed Access

- A portion of disk space is reserved
- A "hashing" algorithm computes record address



© 2008, University of Colombo School of Computing



## Hashed Access Characteristics

- No indexes to search or maintain
- Very fast direct access
- Inefficient sequential access
- Use when direct access is needed, but sequential access is not
- Data cannot be sorted easily



© 2008, University of Colombo School of Computing



## Query Optimization

A query typically has many possible execution strategies and the process of choosing a suitable one for processing a query is known as query optimisation.

The job of the heuristic query optimiser is to transform this initial query tree into a final query tree that is efficient to execute.

The optimiser must include rules for equivalence among relational algebra expressions that can be applied to the initial tree, guided by the heuristic query optimisation rules to produce the final optimised query tree.



© 2008, University of Colombo School of Computing



**Execution strategy** - The DBMS must then devise an execution strategy for retrieving the result of the query from the internal database files.

**Query code generator** - According to the chosen execution plan, the code generator generates the code to execute the plan.

**Runtime database processor** - The runtime database processor has the task of running the query code (compiled or interpreted) to produce the query result. If a runtime error occurs the runtime database processor generates an error message.



© 2008, University of Colombo School of Computing



## Translating SQL Queries into Relational Algebra

```
SELECT p.pno, d.dno, e.ename
FROM Project as p, Department as d, Employee as e
WHERE d.dno=p.dept and d.mgr=e.empno and
p.location='Colombo';
```

$T1 \leftarrow \text{Project} \bowtie_{\text{dno=dept}} \text{Department}$

$T2 \leftarrow T1 \bowtie_{\text{mgr=empno}} \text{Employee}$

$T3 \leftarrow \sigma_{\text{location='Colombo'}}(T2)$

$\text{Result} \leftarrow \pi_{\text{pno, dno, ename}}(T3)$

JOIN Project and Department over dno=dept giving T1

JOIN T1 and Employee over mgr=empno giving T2

RESTRICT T2 where location='Colombo' giving T3

PROJECT T3 over pno, dno, ename giving Result



© 2008, University of Colombo School of Computing



## Queries in Relational Algebra

### Equivalent Queries

(a)  $\text{Result} \leftarrow \pi_{\text{pno, dno, ename}}(\sigma_{\text{location='Colombo'}}((\text{Project} \bowtie_{\text{dno=dept}} \text{Department}) \bowtie_{\text{mgr=empno}} \text{Employee}))$

(b)  $T1 \leftarrow \sigma_{\text{location='Colombo'}}(\text{Project})$   
 $T2 \leftarrow T1 \bowtie_{\text{dno=dept}} \text{Department}$   
 $T3 \leftarrow T2 \bowtie_{\text{mgr=empno}} \text{Employee}$   
 $\text{Result} \leftarrow \pi_{\text{pno, dno, ename}}(T3)$



© 2008, University of Colombo School of Computing



## Basic algorithms for executing query operations

- Each DBMS typically has a number of general database access algorithms that implement relational operations such as SELECT or JOIN or combinations of these operations.
- The query optimisation module will consider only execution strategies that can be implemented by the DBMS access algorithms (i.e. storage structures and access paths).



© 2008, University of Colombo School of Computing



Sorting is one of the primary algorithms used in query processing.

For example, whenever an SQL query specifies an ORDER BY clause, the query result must be sorted.

Sorting is also a key component in sort-merge algorithms used for JOIN and other operations (such as UNION and INTERSECTION), and in duplicate elimination algorithms for the PROJECT operation (when an SQL query specifies the DISTINCT option in the SELECT clause).



© 2008, University of Colombo School of Computing



There are many options for executing a SELECT operation.

A number of searching algorithms are possible for selecting records from a file.

Linear search (brute force), binary search, using a primary index or hash key, clustering index and using secondary index (B-tree) on an equality comparison are examples of searching.



© 2008, University of Colombo School of Computing



## Selecting Records, e.g.

- Primary index** (records ordered on a key field)

$T1 \leftarrow \sigma_{\text{empno}=12345}(\text{Employee})$  (single record)

$T1 \leftarrow \sigma_{\text{dno}=5}(\text{Department})$  (multiple records)

- Clustering Index** (records ordered on a non key field)

$T1 \leftarrow \sigma_{\text{dno}=5}(\text{Employee})$  (multiple records)

- B<sup>+</sup>-Tree Index** (secondary index on equality comparison)

$T1 \leftarrow \sigma_{\text{salary} \geq 30000 \text{ and } \text{salary} \leq 35000}(\text{Employee})$  (multiple records)



© 2008, University of Colombo School of Computing



## Using Heuristics in Query Optimisation

- There are two main techniques for query optimisation: heuristic rules and systematic estimating.
- Heuristic rules are used to order the operations in a query execution strategy. The rules typically reorder the operations in a query tree or determine an order for executing the operations specified by a query graph.
- Systematic estimating is used to cost the different execution strategies and to choose the execution plan with the lowest cost estimate.
- Both strategies are usually combined in a query optimiser.



© 2008, University of Colombo School of Computing



## Transformation rules for relational algebra operations

- There are many rules for transforming relational algebra operations into equivalent ones.
- These are in addition to those discussed under relational algebra.
- These rules are used in heuristic optimisation.
- Algorithms that utilise these rules are used to transform an initial query tree into an optimised tree that is more efficient to execute.
- Here we look at some examples that demonstrate such transformations.



© 2008, University of Colombo School of Computing



## Rule 1 (cascade of $\sigma$ )

Break up any SELECT operations ( $\sigma$ ) with conjunctive conditions (AND) into a cascade (sequence) of individual SELECT operations.

$$\sigma_{c_1 \text{ AND } c_2 \text{ AND } \dots \text{ AND } c_n} (R) \equiv \sigma_{c_1} (\sigma_{c_2} (\dots (\sigma_{c_n} (R)) \dots))$$

This permits a greater degree of freedom in moving SELECT operations down different branches of the tree.



© 2008, University of Colombo School of Computing



## Rule 2 (commutative of $\sigma$ )

The SELECT operation is commutative

$$\sigma_{c_1} (\sigma_{c_2} (R)) \equiv \sigma_{c_2} (\sigma_{c_1} (R))$$

## Rule 3 (commutative of $\sigma$ with $\pi$ )

If the SELECT condition  $c$  involves only attributes  $a_1, a_2, \dots, a_n$  in the PROJECTION list, the two operations can be commuted.

$$\pi_{a_1, a_2, \dots, a_n} (\sigma_c (R)) \equiv \sigma_c (\pi_{a_1, a_2, \dots, a_n} (R))$$



© 2008, University of Colombo School of Computing



## Rule 4 (commutative of $\sigma$ with $X$ or $\infty$ )

If all the attributes in the selection condition  $c$  involve only the attributes of one of the relations being joined (say  $R$ ) the two operations can be commuted as

$$\sigma_c (R \infty S) \equiv (\sigma_c (R)) \infty S$$

Alternatively if the selection condition  $c$  can be written as  $c_1$  and  $c_2$ , where  $c$  involves only the attributes of  $S$ , the operations commute as

$$\sigma_c (R \infty S) \equiv (\sigma_{c_1} (R)) \infty (\sigma_{c_2} (S))$$

The same rules apply if the  $\infty$  is replaced by a  $X$  operation.



© 2008, University of Colombo School of Computing



## Rule 5 (commuting $\sigma$ with set operations)

The  $\sigma$  operation commutes with  $\cup$ ,  $\cap$  and  $-$ . If  $\theta$  stands for any one of these 3 operations then

$$\sigma_c (R \theta S) \equiv (\sigma_c (R)) \theta (\sigma_c (S))$$

Using rules 2, 3, 4, and 5 concerning the commutative of SELECT with other operations, move each SELECT operation as far down the query tree as is permitted by the attributes involved in the select condition.

The objective is to reduce the number of tuples that appear in the Cartesian product.




© 2008, University of Colombo School of Computing




SELECT Iname  
 FROM Employee, WorksOn, Project  
 WHERE pname='Aquarius' and pnumber=pno and essn=ssn  
 and bdate>'31/12/1957';

**(a) initial query tree**


Depending on the capacity of the memory intermediate records may have to be written to disk. Thus increase total Read/Write operations.



**(b) Moving σ operation down the query tree**




**(c) Applying the more restrictive σ operation first**




Use the information that pname is a unique attribute of Project relation.

**(d) Replacing Cartesian product and Select with Join operations**



**(e) Moving Project operation (π) down the query tree**



**Using Selectivity and Cost Estimates in Query Optimisation**

A query optimiser should not depend solely on heuristic rules, it should also estimate and compare the costs of executing a query using different execution strategies and should choose the strategy with the lowest cost estimate.

- Cost components
- Cost functions
- Examples

## Cost Components for Query Execution

- Access cost to secondary storage
  - Cost of searching for, reading and writing data blocks that reside on secondary storage.
  - Cost of searching depends on access file structures such as ordering, hashing, indexes. Also based on how data are allocated on disk.
- Storage cost
  - Cost of storing any intermediate files that are generated by an execution strategy for the query



© 2008, University of Colombo School of Computing



## Cost Components ...

- Computation cost
  - Cost of performing in-memory operations on the data buffers during query execution. E.g. searching, sorting, merging for join, performing computations on field values.
- Memory usage cost
  - Cost pertaining to the number of memory buffers needed during query execution.
- Communication cost
  - Cost of sending the query and its results from database server to client.



© 2008, University of Colombo School of Computing



## Catalog Information used in Cost Functions

- $n_r$  - number of tuples in relation  $r$
- $s_r$  - size of tuple in relation  $r$
- $b_r$  - number of blocks containing tuples of  $r$
- $f_r$  - blocking factor – number of tuples of relation  $r$  that fit into one block
- $x_a$  - number of levels of each multilevel index on attribute  $a$ .
- $d_{a,r}$  - number of distinct values in relation  $r$  for attribute  $a$ .  
 $d_{a,r} = n_r / \pi_a(r)$  if  $a$  is unique
- $s_{a,r}$  - Selection cardinality of attribute  $a$  of relation  $r$  – average number records satisfying equality condition. If  $a$  is unique  $d_{a,r} = n_r$ ,  $s_{a,r} = 1$  else (if uniformly distributed),  $s_{a,r} = n_r / d_{a,r}$



© 2008, University of Colombo School of Computing



## Cost Functions for SELECT

- Linear Search
  - Retrieve all file blocks; cost =  $b_r$
  - For equal condition on average cost =  $b_r/2$  if found else cost =  $b_r$
- Binary Search
  - Search access cost =  $\log_2 b_r + (s_{a,r}/f_r) - 1$
  - If unique attribute average cost =  $\log_2 b_r$
- Using a primary / secondary key index
  - Cost =  $x_a + 1$
- Using a hash key
  - Cost  $\approx 1$



© 2008, University of Colombo School of Computing

