


# Chapter 19

## Algorithms for Query Processing and Optimization



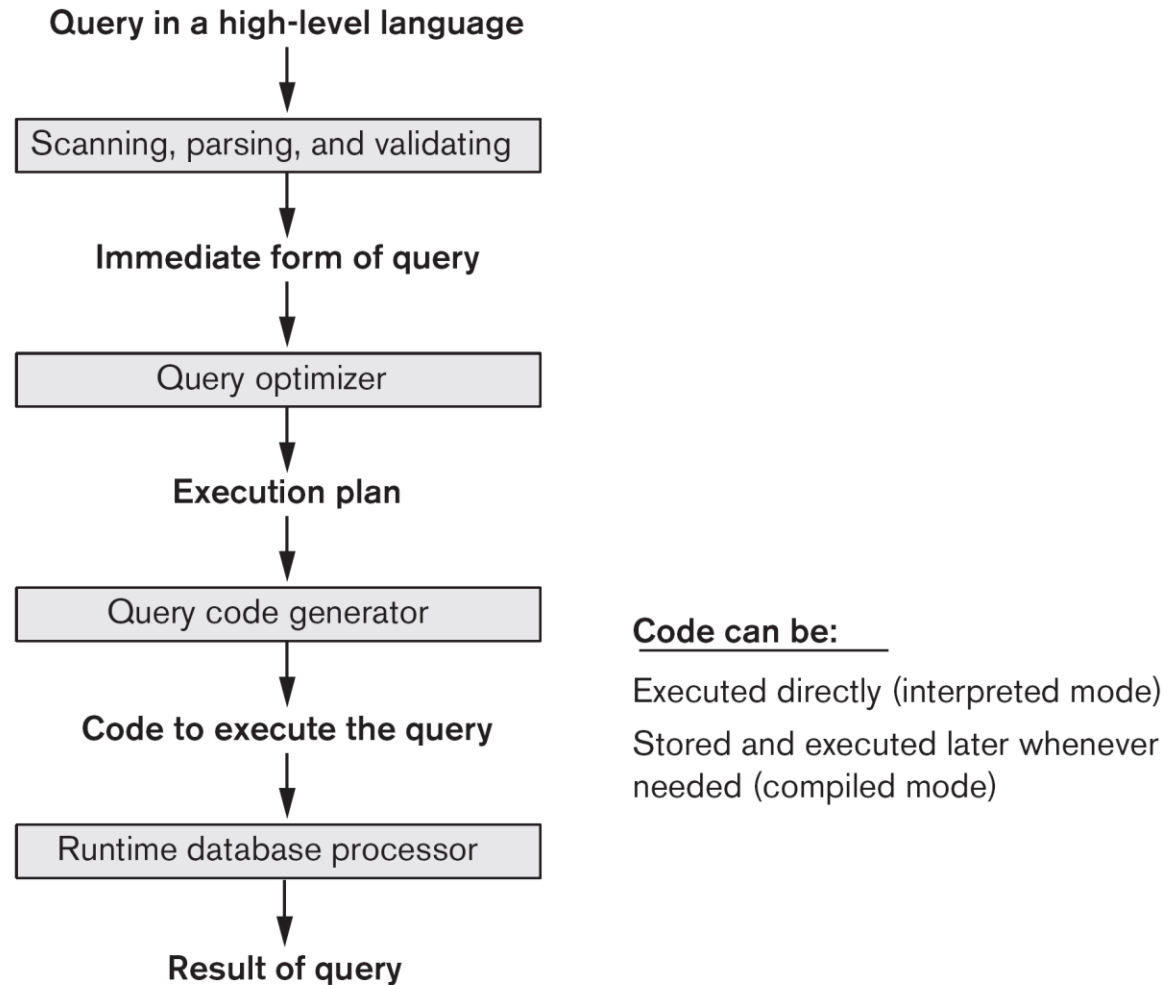
Addison-Wesley  
is an imprint of

PEARSON

# 0. Introduction to Query Processing (1)

- **Query optimization:**
  - The process of choosing a suitable execution strategy for processing a query.
- Two internal representations of a query:
  - **Query Tree**
  - **Query Graph**

# Introduction to Query Processing (2)



**Figure 19.1**

Typical steps when processing a high-level query.

# 1. Translating SQL Queries into Relational Algebra (1)

- **Query block:**
  - The basic unit that can be translated into the algebraic operators and optimized.
- A query block contains a single SELECT-FROM-WHERE expression, as well as GROUP BY and HAVING clause if these are part of the block.
- **Nested queries** within a query are identified as separate query blocks.
- Aggregate operators in SQL must be included in the extended algebra.

# Translating SQL Queries into Relational Algebra (2)

```
SELECT      LNAME, FNAME
FROM        EMPLOYEE
WHERE       SALARY > ( SELECT      MAX (SALARY)
                        FROM        EMPLOYEE
                        WHERE       DNO = 5);
```

```
SELECT      LNAME, FNAME
FROM        EMPLOYEE
WHERE       SALARY > C
```

$\pi_{LNAME, FNAME} (\sigma_{SALARY > C}(EMPLOYEE))$

```
SELECT      MAX (SALARY)
FROM        EMPLOYEE
WHERE       DNO = 5
```

$\mathcal{F}_{MAX\ SALARY} (\sigma_{DNO=5}(EMPLOYEE))$

# 7. Using Heuristics in Query Optimization (1)

- Process for heuristics optimization
  1. The parser of a high-level query generates an initial internal representation;
  2. Apply heuristics rules to optimize the internal representation.
  3. A query execution plan is generated to execute groups of operations based on the access paths available on the files involved in the query.
  
- The main heuristic is to apply first the operations that reduce the size of intermediate results.
  - E.g., Apply `SELECT` and `PROJECT` operations before applying the `JOIN` or other binary operations.

# Using Heuristics in Query Optimization (2)

## ■ Query tree:

- A tree data structure that corresponds to a relational algebra expression. It represents the input relations of the query as **leaf nodes** of the **tree**, and represents the relational algebra operations as internal nodes.
- An execution of the query tree consists of executing an internal node operation whenever its operands are available and then replacing that internal node by the relation that results from executing the operation.

## ■ Query graph:

- A graph data structure that corresponds to a relational calculus expression. It does *not* indicate an order on which operations to perform first. There is only a *single* graph corresponding to each query.

# Using Heuristics in Query Optimization (3)

- Example:
  - For every project located in 'Stafford', retrieve the project number, the controlling department number and the department manager's last name, address and birthdate.
- Relation algebra:

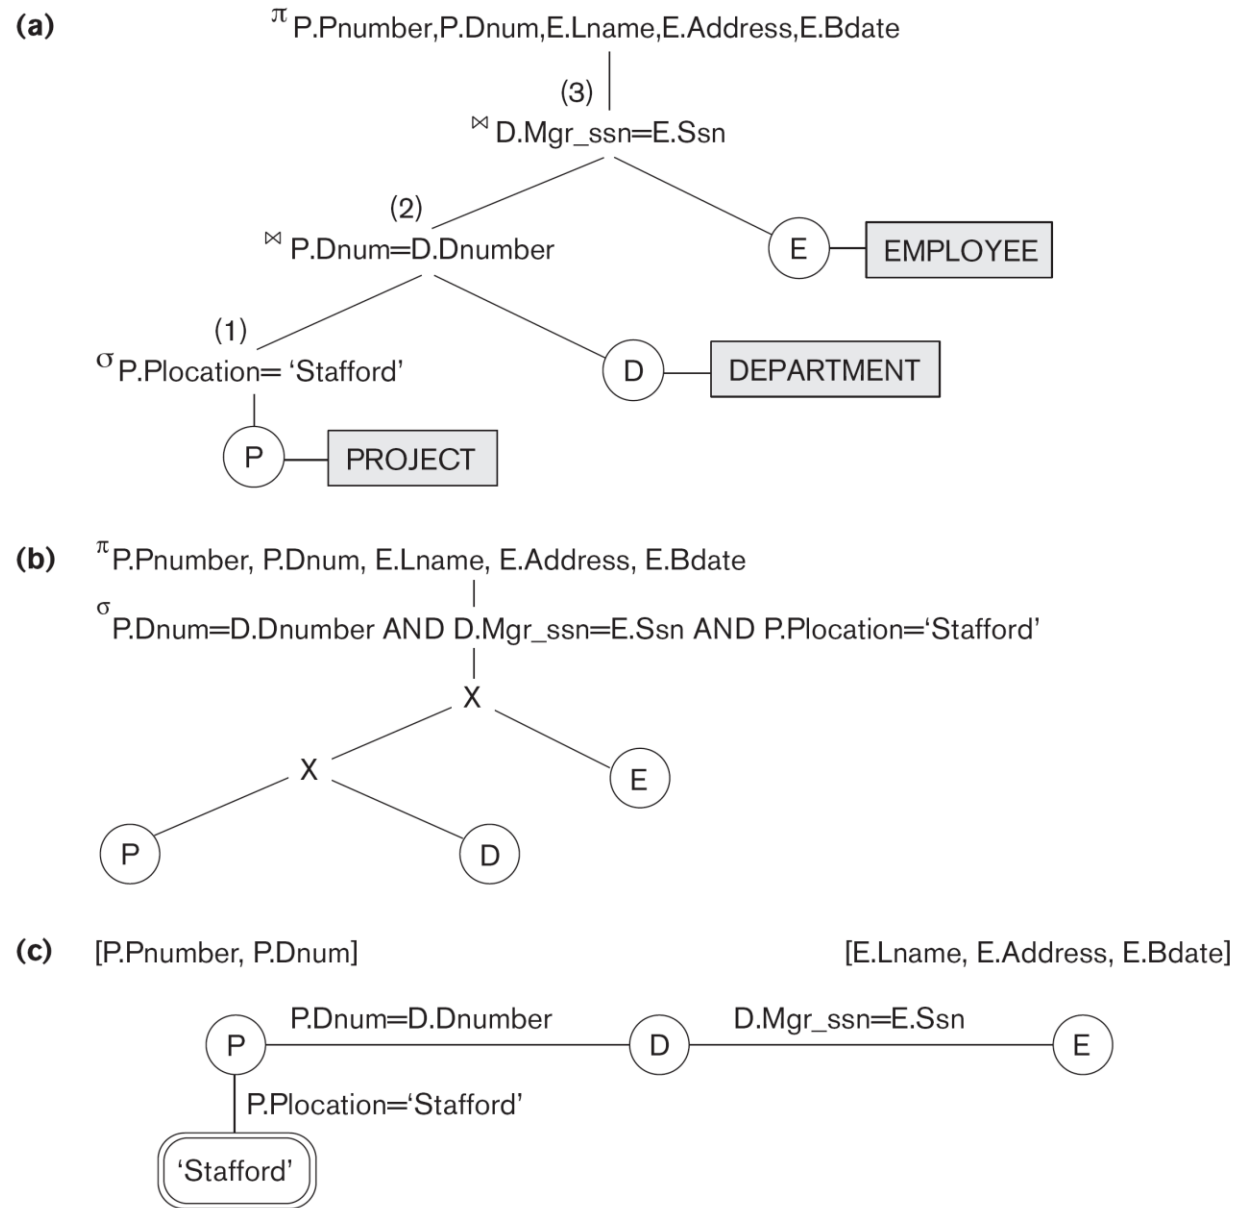
$$\pi_{\text{PNUMBER, DNUM, LNAME, ADDRESS, BDATE}} \left( \left( \left( \sigma_{\text{PLOCATION='STAFFORD'}}(\text{PROJECT}) \right) \right) \bowtie_{\text{DNUM=DNUMBER}} (\text{DEPARTMENT}) \right) \bowtie_{\text{MGRSSN=SSN}} (\text{EMPLOYEE})$$

- SQL query:

```
Q2:      SELECT      P.NUMBER,P.DNUM,E.LNAME,
           E.ADDRESS, E.BDATE
           FROM        PROJECT AS P,DEPARTMENT AS D,
                    EMPLOYEE AS E
           WHERE       P.DNUM=D.DNUMBER AND
                    D.MGRSSN=E.SSN AND
                    P.PLOCATION='STAFFORD';
```



# Using Heuristics in Query Optimization (4)



**Figure 19.4**

Two query trees for the query Q2. (a) Query tree corresponding to the relational algebra expression for Q2. (b) Initial (canonical) query tree for SQL query Q2. (c) Query graph for Q2.

# Using Heuristics in Query Optimization (5)

- **Heuristic Optimization of Query Trees:**
  - The same query could correspond to many different relational algebra expressions — and hence many different query trees.
  - The task of heuristic optimization of query trees is to find a **final query tree** that is efficient to execute.

- **Example:**

```
Q: SELECT      LNAME
      FROM      EMPLOYEE, WORKS_ON, PROJECT
      WHERE     PNAME = 'AQUARIUS' AND
                PNMUBER=PNO AND ESSN=SSN
                AND BDATE > '1957-12-31';
```

# Using Heuristics in Query Optimization (6)

**Figure 19.5**

Steps in converting a query tree during heuristic optimization.

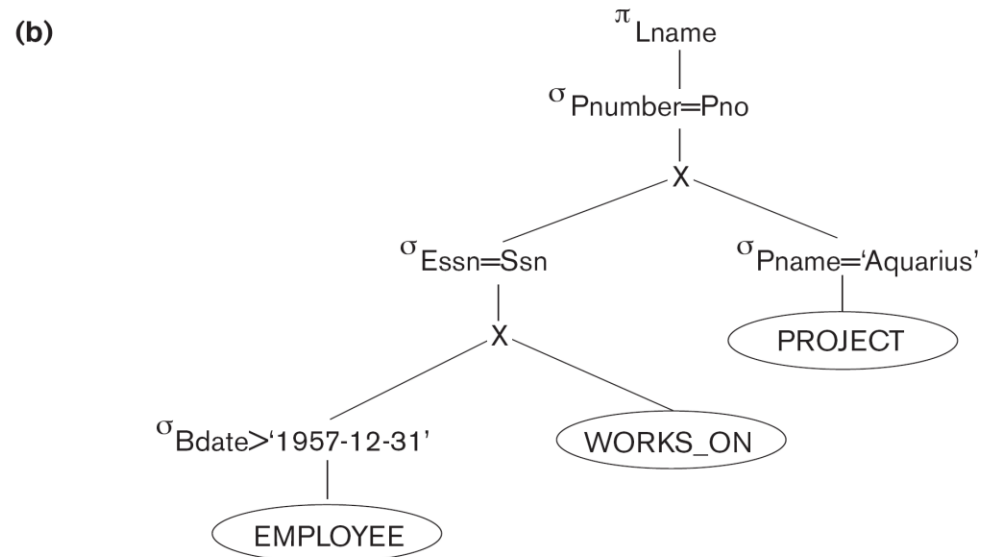
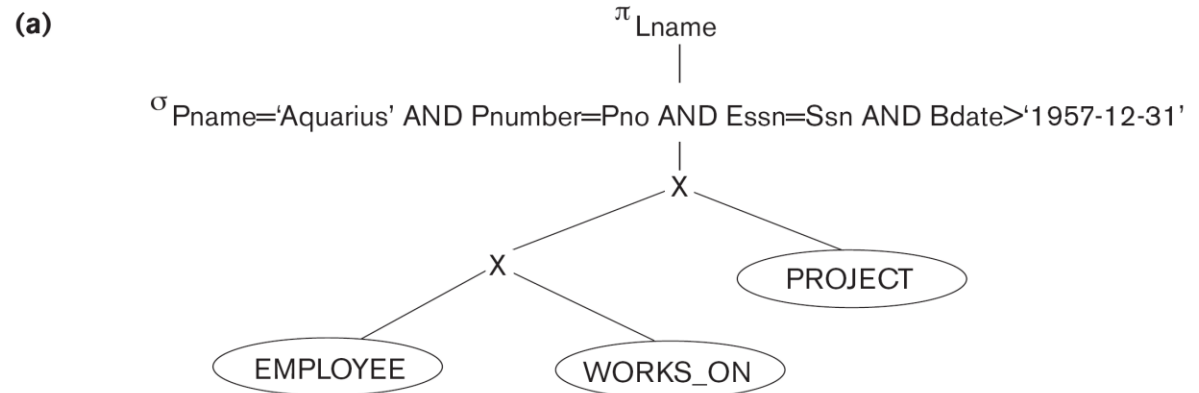
(a) Initial (canonical) query tree for SQL query Q.

(b) Moving SELECT operations down the query tree.

(c) Applying the more restrictive SELECT operation first.

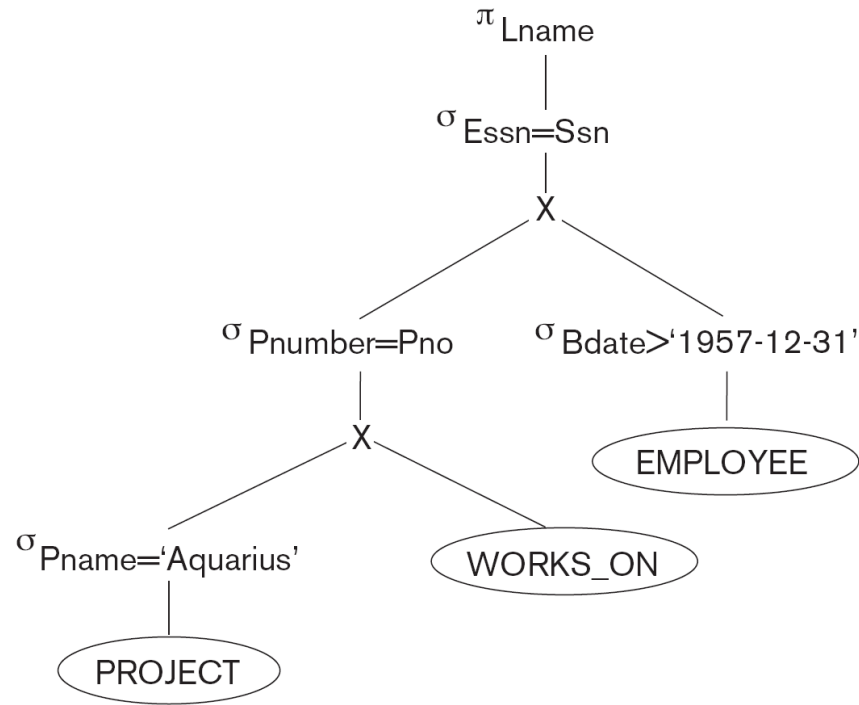
(d) Replacing CARTESIAN PRODUCT and SELECT with JOIN operations.

(e) Moving PROJECT operations down the query tree.

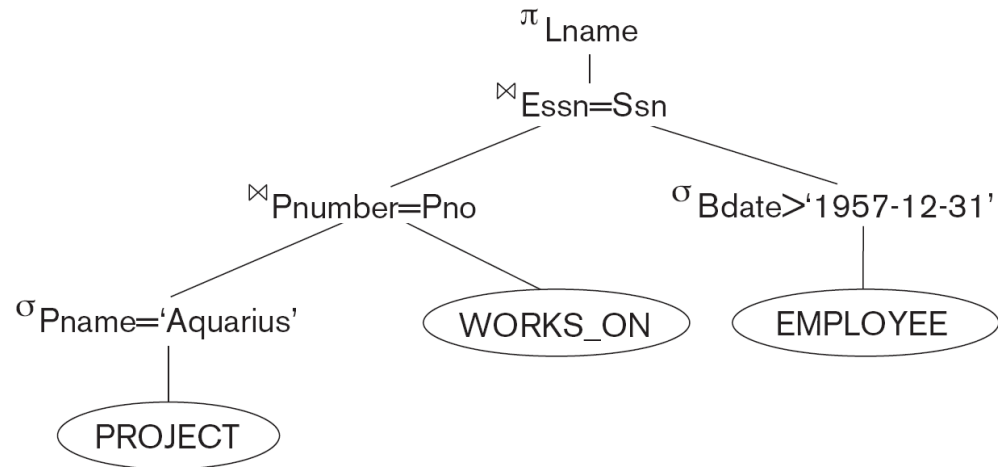


# Using Heuristics in Query Optimization (7)

(c)

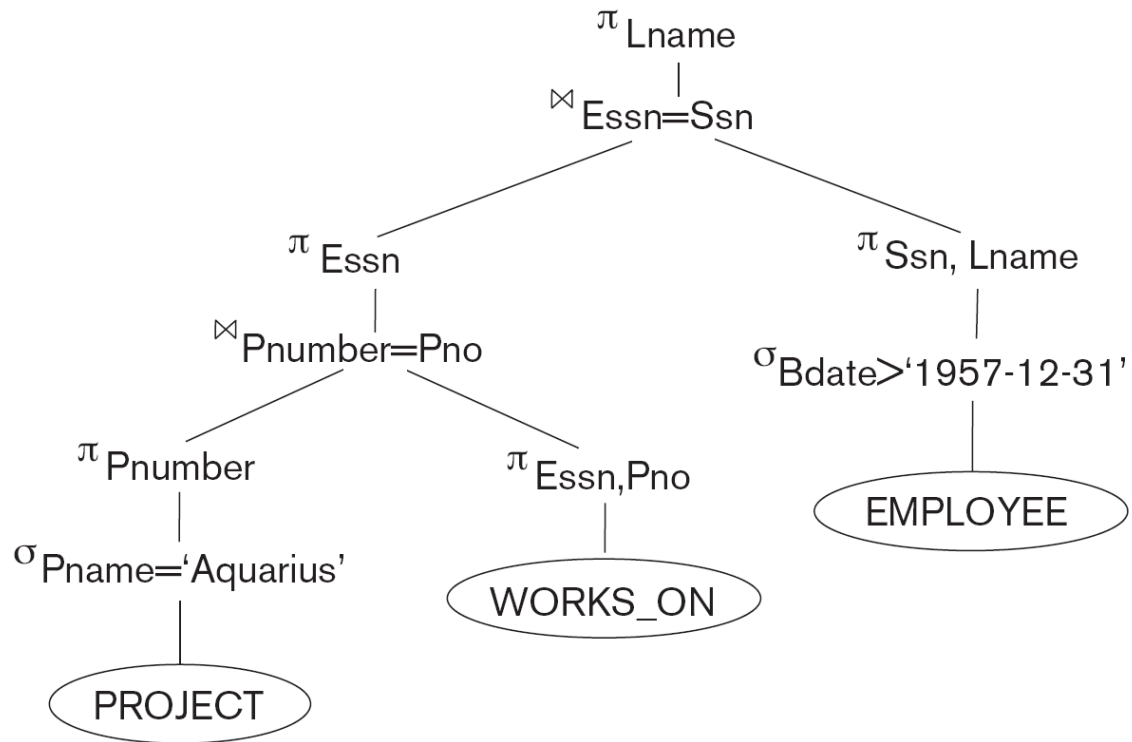


(d)



# Using Heuristics in Query Optimization (8)

(e)



# Using Heuristics in Query Optimization (9)

- General Transformation Rules for Relational Algebra Operations:
  1. Cascade of  $\sigma$ : A conjunctive selection condition can be broken up into a cascade (sequence) of individual  $\sigma$  operations:
    - $\sigma_{c_1 \text{ AND } c_2 \text{ AND } \dots \text{ AND } c_n}(R) = \sigma_{c_1} (\sigma_{c_2} (\dots (\sigma_{c_n}(R)) \dots))$
  2. Commutativity of  $\sigma$ : The  $\sigma$  operation is commutative:
    - $\sigma_{c_1} (\sigma_{c_2}(R)) = \sigma_{c_2} (\sigma_{c_1}(R))$
  3. Cascade of  $\pi$ : In a cascade (sequence) of  $\pi$  operations, all but the last one can be ignored:
    - $\pi_{List1} (\pi_{List2} (\dots (\pi_{Listn}(R)) \dots)) = \pi_{List1}(R)$
  4. Commuting  $\sigma$  with  $\pi$ : If the selection condition  $c$  involves only the attributes  $A_1, \dots, A_n$  in the projection list, the two operations can be commuted:
    - $\pi_{A_1, A_2, \dots, A_n} (\sigma_c (R)) = \sigma_c (\pi_{A_1, A_2, \dots, A_n} (R))$

# Using Heuristics in Query Optimization (10)

- General Transformation Rules for Relational Algebra Operations (contd.):
5. Commutativity of  $\bowtie$  ( and  $\times$  ): The  $\bowtie$  operation is commutative as is the  $\times$  operation:
    - $R \bowtie_c S = S \bowtie_c R$ ;  $R \times S = S \times R$
  6. Commuting  $\sigma$  with  $\bowtie$  (or  $\times$  ): 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 follows:
    - $\sigma_c ( R \bowtie S ) = ( \sigma_c ( R ) ) \bowtie S$
  - Alternatively, if the selection condition  $c$  can be written as  $(c_1 \text{ and } c_2)$ , where condition  $c_1$  involves only the attributes of  $R$  and condition  $c_2$  involves only the attributes of  $S$ , the operations commute as follows:
    - $\sigma_c ( R \bowtie S ) = ( \sigma_{c_1} ( R ) ) \bowtie ( \sigma_{c_2} ( S ) )$

# Using Heuristics in Query Optimization (11)

- General Transformation Rules for Relational Algebra Operations (contd.):
7. Commuting  $\pi$  with  $\bowtie$  (or  $\times$ ): Suppose that the projection list is  $L = \{A_1, \dots, A_n, B_1, \dots, B_m\}$ , where  $A_1, \dots, A_n$  are attributes of  $R$  and  $B_1, \dots, B_m$  are attributes of  $S$ . If the join condition  $c$  involves only attributes in  $L$ , the two operations can be commuted as follows:
- $\pi_L ( R \bowtie_C S ) = ( \pi_{A_1, \dots, A_n} ( R ) ) \bowtie_C ( \pi_{B_1, \dots, B_m} ( S ) )$
  - If the join condition  $C$  contains additional attributes not in  $L$ , these must be added to the projection list, and a final  $\pi$  operation is needed.



# Using Heuristics in Query Optimization (12)

- General Transformation Rules for Relational Algebra Operations (contd.):
- 8. Commutativity of set operations: The set operations  $\cup$  and  $\cap$  are commutative but “ $-$ ” is not.
- 9. Associativity of  $\bowtie$ ,  $\times$ ,  $\cup$ , and  $\cap$  : These four operations are individually associative; that is, if  $\theta$  stands for any one of these four operations (throughout the expression), we have
  - $(R \theta S) \theta T = R \theta (S \theta T)$
- 10. Commuting  $\sigma$  with set operations: The  $\sigma$  operation commutes with  $\cup$  ,  $\cap$  , and  $-$ . If  $\theta$  stands for any one of these three operations, we have
  - $\sigma_c (R \theta S) = (\sigma_c (R)) \theta (\sigma_c (S))$

# Using Heuristics in Query Optimization (13)

- General Transformation Rules for Relational Algebra Operations (contd.):

- The  $\pi$  operation commutes with  $\cup$ .

$$\pi_L ( R \cup S ) = ( \pi_L ( R ) ) \cup ( \pi_L ( S ) )$$

- Converting a  $(\sigma, x)$  sequence into  $\bowtie$ : If the condition  $c$  of a  $\sigma$  that follows a  $x$  corresponds to a join condition, convert the  $(\sigma, x)$  sequence into a  $\bowtie$  as follows:

$$(\sigma_C ( R \times S )) = ( R \bowtie_C S )$$

- Other transformations

# Using Heuristics in Query Optimization (14)

- Outline of a Heuristic Algebraic Optimization Algorithm:
  1. Using rule 1, break up any select operations with conjunctive conditions into a cascade of select operations.
  2. Using rules 2, 4, 6, and 10 concerning the commutativity 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.
  3. Using rule 9 concerning associativity of binary operations, rearrange the leaf nodes of the tree so that the leaf node relations with the most restrictive select operations are executed first in the query tree representation.
  4. Using Rule 12, combine a Cartesian product operation with a subsequent select operation in the tree into a join operation.
  5. Using rules 3, 4, 7, and 11 concerning the cascading of project and the commuting of project with other operations, break down and move lists of projection attributes down the tree as far as possible by creating new project operations as needed.
  6. Identify subtrees that represent groups of operations that can be executed by a single algorithm.

# Using Heuristics in Query Optimization

## (15)

- **Summary of Heuristics for Algebraic Optimization:**
  1. The main heuristic is to apply first the operations that reduce the size of intermediate results.
  2. Perform select operations as early as possible to reduce the number of tuples and perform project operations as early as possible to reduce the number of attributes. (This is done by moving select and project operations as far down the tree as possible.)
  3. The select and join operations that are most restrictive should be executed before other similar operations. (This is done by reordering the leaf nodes of the tree among themselves and adjusting the rest of the tree appropriately.)

# Using Heuristics in Query Optimization (16)

## ■ Query Execution Plans

- An execution plan for a relational algebra query consists of a combination of the relational algebra query tree and information about the access methods to be used for each relation as well as the methods to be used in computing the relational operators stored in the tree.
- **Materialized evaluation:** the result of an operation is stored as a temporary relation.
- **Pipelined evaluation:** as the result of an operator is produced, it is forwarded to the next operator in sequence.

# 8. Using Selectivity and Cost Estimates in Query Optimization (1)

## ■ Cost-based query optimization:

- Estimate and compare the costs of executing a query using different execution strategies and choose the strategy with the lowest cost estimate.
- (Compare to heuristic query optimization)

## ■ Issues

- Cost function
- Number of execution strategies to be considered

# Using Selectivity and Cost Estimates in Query Optimization (2)

- Cost Components for Query Execution
  1. Access cost to secondary storage
  2. Storage cost
  3. Computation cost
  4. Memory usage cost
  5. Communication cost
  
- Note: Different database systems may focus on different cost components.

# Using Selectivity and Cost Estimates in Query Optimization (3)

- Catalog Information Used in Cost Functions
  - Information about the size of a file
    - number of records (tuples) ( $r$ ),
    - record size ( $R$ ),
    - number of blocks ( $b$ )
    - blocking factor ( $bfr$ )
  - Information about indexes and indexing attributes of a file
    - Number of levels ( $x$ ) of each multilevel index
    - Number of first-level index blocks ( $bl_1$ )
    - Number of distinct values ( $d$ ) of an attribute
    - Selectivity ( $sl$ ) of an attribute
    - Selection cardinality ( $s$ ) of an attribute. ( $s = sl * r$ )



# Using Selectivity and Cost Estimates in Query Optimization (4)

- Examples of Cost Functions for SELECT
- S1. Linear search (brute force) approach
  - $C_{S1a} = b$ ;
  - For an equality condition on a key,  $C_{S1a} = (b/2)$  if the record is found; otherwise  $C_{S1a} = b$ .
- S2. Binary search:
  - $C_{S2} = \log_2 b + (s/bfr) \lceil -1$
  - For an equality condition on a unique (key) attribute,  $C_{S2} = \log_2 b$
- S3. Using a primary index (S3a) or hash key (S3b) to retrieve a single record
  - $C_{S3a} = x + 1$ ;  $C_{S3b} = 1$  for static or linear hashing;
  - $C_{S3b} = 1$  for extendible hashing;

# Using Selectivity and Cost Estimates in Query Optimization (5)

- Examples of Cost Functions for SELECT (contd.)
- S4. Using an ordering index to retrieve multiple records:
  - For the comparison condition on a key field with an ordering index,  $C_{S4} = x + (b/2)$
- S5. Using a clustering index to retrieve multiple records:
  - $C_{S5} = x + \lceil (s/bfr) \rceil$
- S6. Using a secondary (B+-tree) index:
  - For an equality comparison,  $C_{S6a} = x + s$ ;
  - For an comparison condition such as  $>$ ,  $<$ ,  $>=$ , or  $<=$ ,
  - $C_{S6a} = x + (b_{l1}/2) + (r/2)$

# Using Selectivity and Cost Estimates in Query Optimization (6)

- Examples of Cost Functions for SELECT (contd.)
- S7. Conjunctive selection:
  - Use either S1 or one of the methods S2 to S6 to solve.
  - For the latter case, use one condition to retrieve the records and then check in the memory buffer whether each retrieved record satisfies the remaining conditions in the conjunction.
- S8. Conjunctive selection using a composite index:
  - Same as S3a, S5 or S6a, depending on the type of index.
- Examples of using the cost functions.

# Using Selectivity and Cost Estimates in Query Optimization (7)

## ■ Examples of Cost Functions for JOIN

### ■ Join selectivity ( $js$ )

$$■ \quad js = \frac{|(R \bowtie_C S)|}{|R \times S|} = \frac{|(R \bowtie_C S)|}{(|R| * |S|)}$$

■ If condition  $C$  does not exist,  $js = 1$ ;

■ If no tuples from the relations satisfy condition  $C$ ,  $js = 0$ ;

■ Usually,  $0 \leq js \leq 1$ ;

## ■ Size of the result file after join operation

$$■ \quad |(R \bowtie_C S)| = js * |R| * |S|$$

# Using Selectivity and Cost Estimates in Query Optimization (8)

- Examples of Cost Functions for JOIN (contd.)
- J1. Nested-loop join:
  - $C_{J1} = b_R + (b_R * b_S) + ((js * |R| * |S|) / bfr_{RS})$
  - (Use R for outer loop)
- J2. Single-loop join (using an access structure to retrieve the matching record(s))
  - If an index exists for the join attribute B of S with index levels  $x_B$ , we can retrieve each record s in R and then use the index to retrieve all the matching records t from S that satisfy  $t[B] = s[A]$ .
  - The cost depends on the type of index.

# Using Selectivity and Cost Estimates in Query Optimization (9)

- Examples of Cost Functions for JOIN (contd.)
- J2. Single-loop join (contd.)
  - For a secondary index,
    - $C_{J2a} = b_R + (|R| * (x_B + s_B)) + ((js * |R| * |S|)/bfr_{RS});$
  - For a clustering index,
    - $C_{J2b} = b_R + (|R| * (x_B + (s_B/bfr_B))) + ((js * |R| * |S|)/bfr_{RS});$
  - For a primary index,
    - $C_{J2c} = b_R + (|R| * (x_B + 1)) + ((js * |R| * |S|)/bfr_{RS});$
  - If a hash key exists for one of the two join attributes — B of S
    - $C_{J2d} = b_R + (|R| * h) + ((js * |R| * |S|)/bfr_{RS});$
- J3. Sort-merge join:
  - $C_{J3a} = C_S + b_R + b_S + ((js * |R| * |S|)/bfr_{RS});$
  - (CS: Cost for sorting files)

# Using Selectivity and Cost Estimates in Query Optimization (10)

## ■ Multiple Relation Queries and Join Ordering

- A query joining  $n$  relations will have  $n-1$  join operations, and hence can have a large number of different join orders when we apply the algebraic transformation rules.
- Current query optimizers typically limit the structure of a (join) query tree to that of left-deep (or right-deep) trees.

## ■ Left-deep tree:

- A binary tree where the right child of each non-leaf node is always a base relation.
  - Amenable to pipelining
  - Could utilize any access paths on the base relation (the right child) when executing the join.

# 9. Overview of Query Optimization in Oracle

- Oracle DBMS V8
  - **Rule-based query optimization:** the optimizer chooses execution plans based on heuristically ranked operations.
    - (Currently it is being phased out)
  - **Cost-based query optimization:** the optimizer examines alternative access paths and operator algorithms and chooses the execution plan with lowest estimate cost.
    - The query cost is calculated based on the estimated usage of resources such as I/O, CPU and memory needed.
  - Application developers could specify hints to the ORACLE query optimizer.
  - The idea is that an application developer might know more information about the data.



# 10. Semantic Query Optimization

- **Semantic Query Optimization:**

- Uses constraints specified on the database schema in order to modify one query into another query that is more efficient to execute.

- Consider the following SQL query,

```
SELECT      E.LNAME, M.LNAME
FROM        EMPLOYEE E M
WHERE       E.SUPERSSN=M.SSN AND E.SALARY>M.SALARY
```

- **Explanation:**

- Suppose that we had a constraint on the database schema that stated that no employee can earn more than his or her direct supervisor. If the semantic query optimizer checks for the existence of this constraint, it need not execute the query at all because it knows that the result of the query will be empty. Techniques known as theorem proving can be used for this purpose.

# Summary

0. Introduction to Query Processing
1. Translating SQL Queries into Relational Algebra
2. Algorithms for External Sorting
3. Algorithms for SELECT and JOIN Operations
4. Algorithms for PROJECT and SET Operations
5. Implementing Aggregate Operations and Outer Joins
6. Combining Operations using Pipelining
7. Using Heuristics in Query Optimization
8. Using Selectivity and Cost Estimates in Query Optimization
9. Overview of Query Optimization in Oracle
10. Semantic Query Optimization