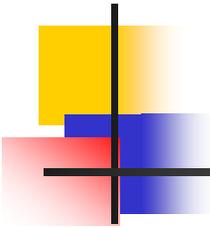
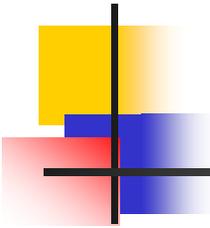


Query Processing and Optimization



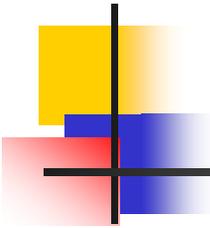
Objectives

- Introduction +
- Major Phases in Query Processing +
- Optimizing Queries +
 - Steps in Query Optimization +
 - Query Trees and Graphs +
 - Relational Algebra Transformation +
 - Query Processing and Optimization +

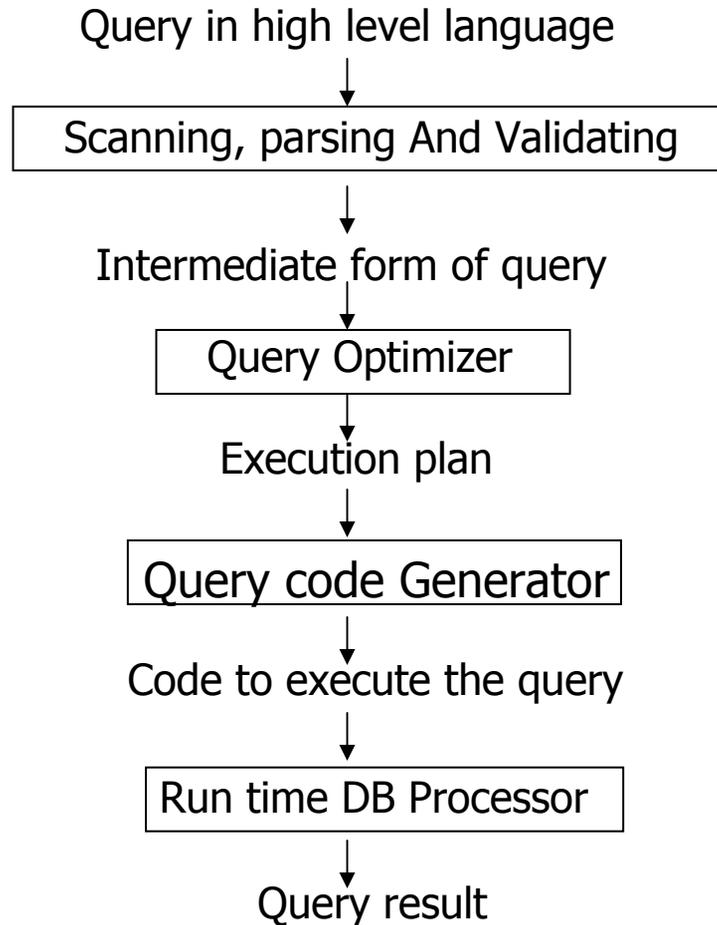


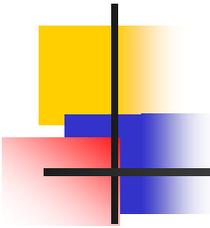
- Introduction

- With declarative languages such as SQL, the user specifies what data is required rather than how it is retrieved
- Relieves user of knowing what constitutes a good execution strategy
- Also gives the DBMS more control over system performance.



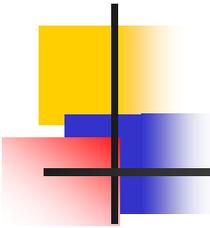
-- Processing a High-Level Query





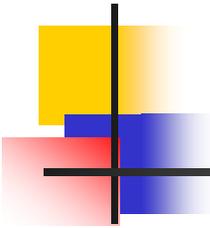
- Major Phases in Query Processing

- Scan, Parse, and Validate Query: Transform query written in high level language (e g. SQL), into correct efficient execution strategy expressed in low-level language (e. g. Relational Algebra)
- Optimize query and generate code: produce an execution strategy or plan and run-time code.
- Execute query: Implement the plan to retrieve required data.



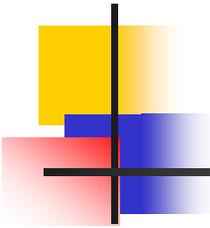
- Optimizing Queries

- The objective in query optimization is to select an efficient execution strategy.
- As there are many equivalent transformations from same high-level query, aim is to choose the one that minimizes resource usage.
- Generally, an efficient execution plan reduces the total execution time of a query; thereby reducing the response time of a query.
- The problem is computationally intractable with large number of relations, so the strategy adopted is reduced to finding a near optimum solution.
- Two main techniques for query optimization are:
 - Heuristic rules that order operations in a query
 - Comparing different strategies based on relative cost, and selecting one that minimizes resource usage.



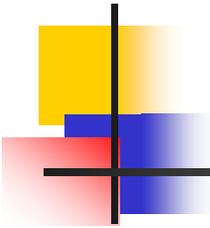
- Steps in Query Optimization

- In using heuristics during query optimization, the following steps must be followed:
 1. Translate queries into query trees or query graphs
 2. Apply transformation rules for relational algebra expression.
 3. Perform heuristic optimization



-- Query Trees and Graphs

- **A query tree** is 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.
- The execution terminates when the root node is executed and produces the result relation for the query.



--- Query Tree Example ...

- For every project located in 'Stafford', list the project number, the controlling department number and the department manager's last name, address, and birth date.

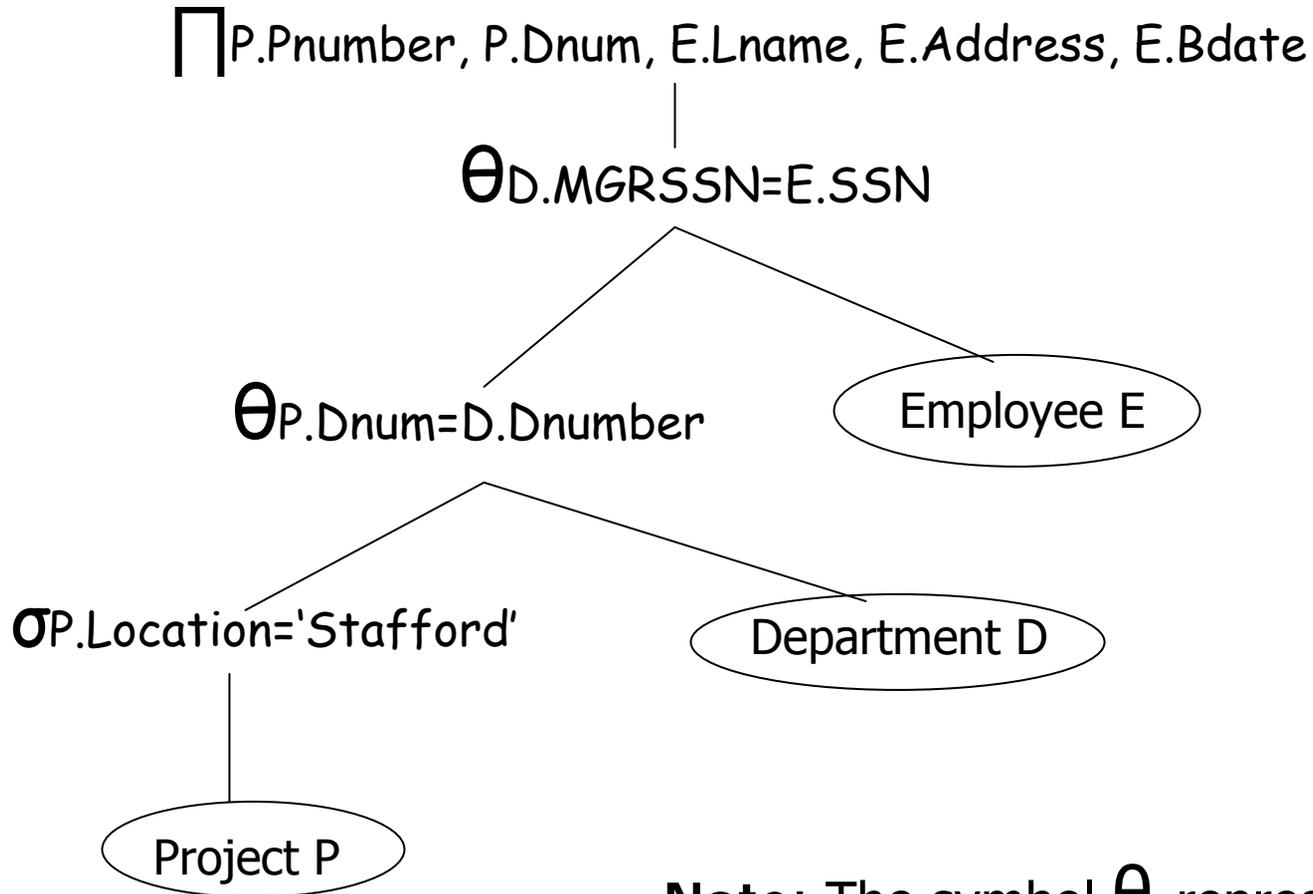
```
SELECT Pnumber, Dnum, Lname, Address, Bdate  
FROM Project, Department, Employee  
WHERE Dnum=Dnumber  
AND MgrSSN = SSN  
AND Plocation = 'Stafford'
```

- Equivalent relational algebra expression is:

```
 $\Pi_{Pnumber, Dnum, Lname, Address, Bdate}$   
 $((\sigma_{Plocation = 'Stafford'}(Project))$   
 $\theta_{Dnum=Dnumber}(Department))$   
 $\theta_{MgrSSN=SSN}(Employee))$ 
```

Note: θ stands for JOIN

... --- Query Tree: Example

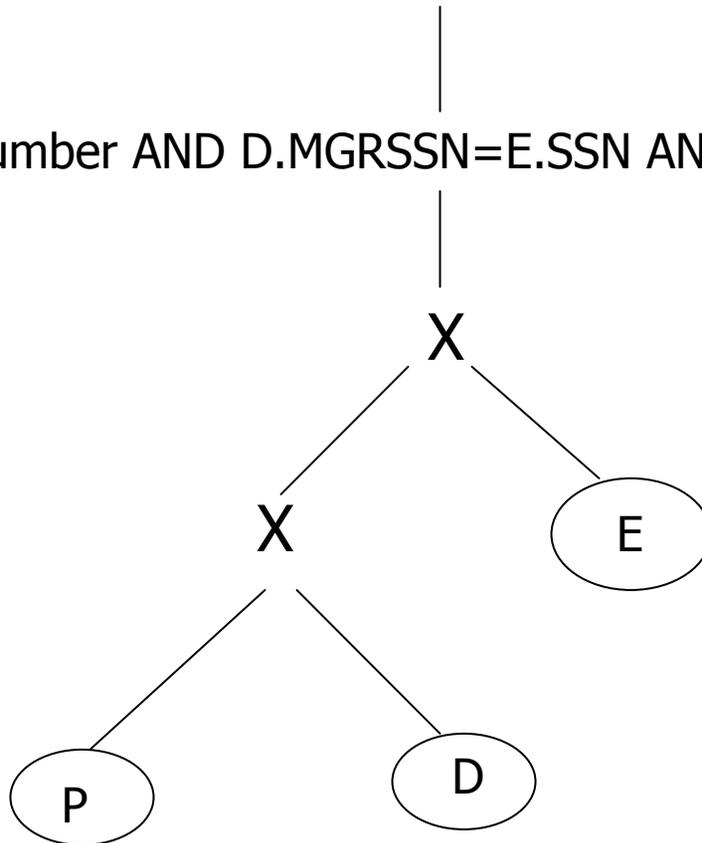


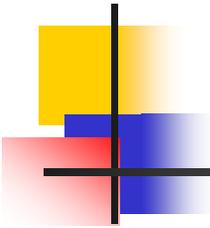
Note: The symbol Θ represents JOIN

... --- Query Tree: Example

Π P.Pnumber, P.Dnum, E.Lname, E.Address, E.Bdate

σ P.Dnum=D.Dnumber AND D.MGRSSN=E.SSN AND P.Location='Stafford'



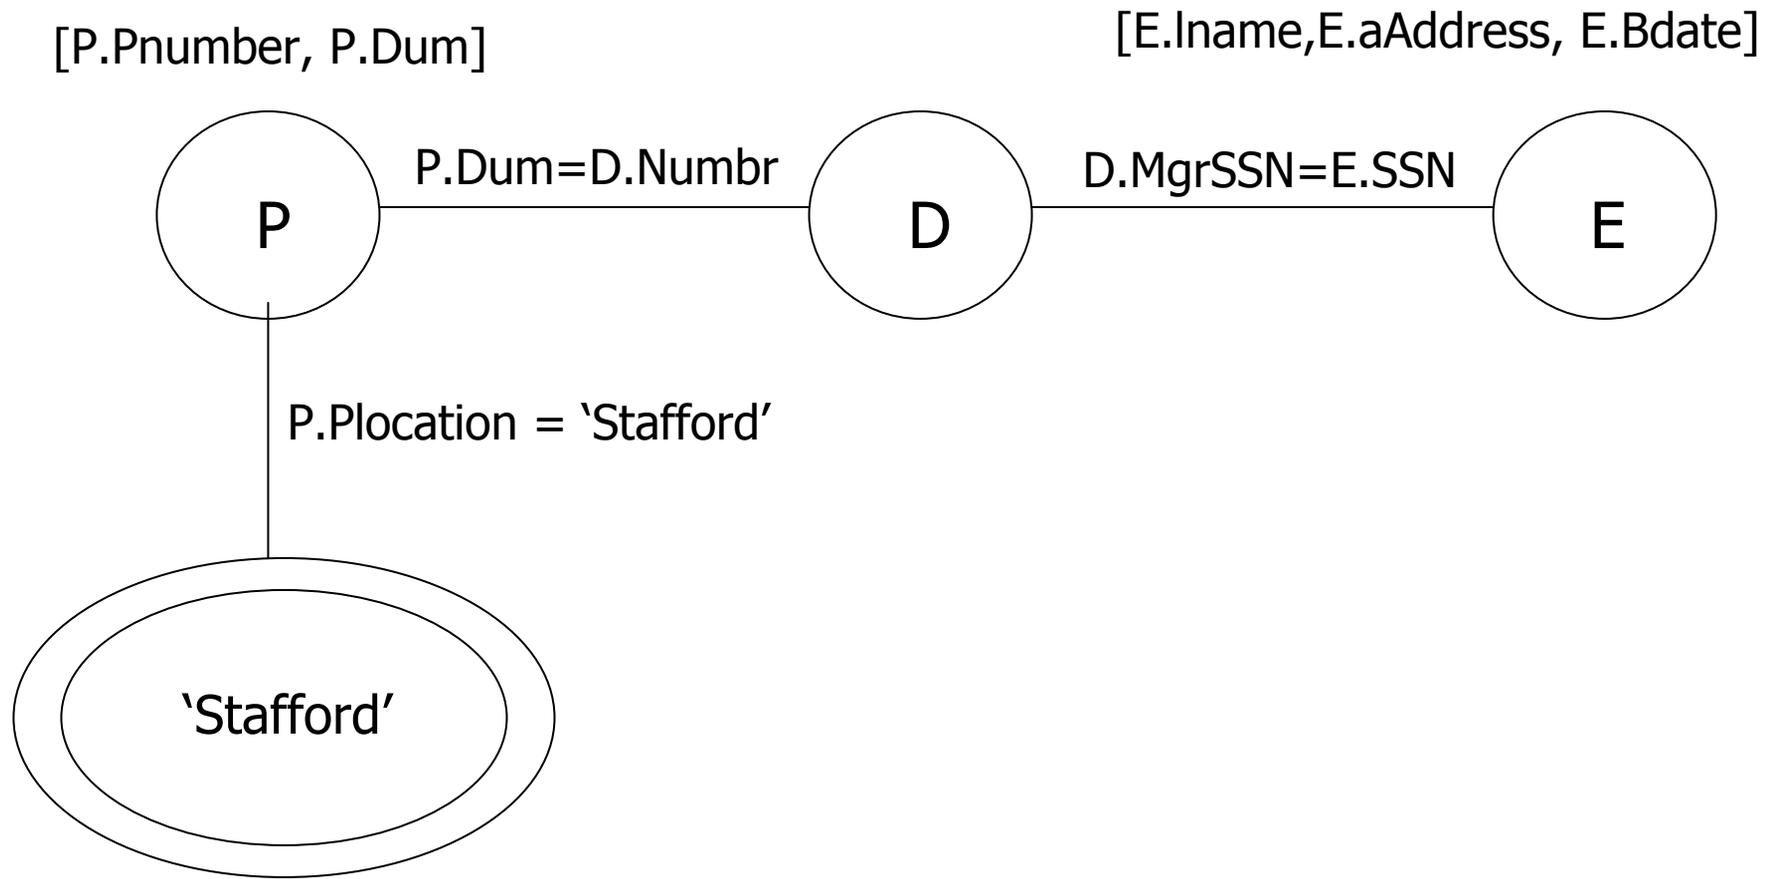


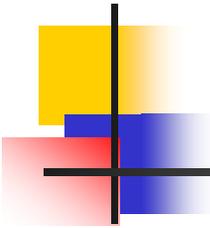
--- Query Graph Example ...

- For every project located in 'Stafford', list the project number, the controlling department number, and the department manager's last name, address, and birth date.

```
SELECT Pnumber, Dnum, Lname, Address, Bdate  
FROM Project, Department, Employee  
WHERE Dnum = Dnumber  
AND MgrSSN = SSN  
AND Plocation = 'Stafford'
```

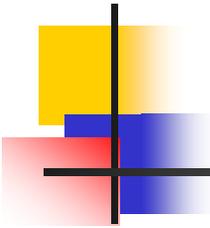
... --- Query Graph Example





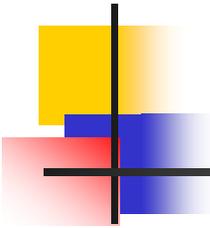
- Relational Algebra Transformation ...

- The following list gives a basic selection of relational algebra transformation:
 - **Cascade of selection:** Conjunctive selection operations can cascade into individual selection operations
$$\sigma_{p \text{ AND } q \text{ AND } r}(\mathbf{R}) = \sigma_p(\sigma_q(\sigma_r(\mathbf{R})))$$
 - **Cascade of projection:** In a sequence of projection operations, only the last in sequence is required.
$$\Pi_L(\Pi_M(\dots(\Pi_N(\mathbf{R}))) = \Pi_L(\mathbf{R})$$
 - **Commutativity of selection:** A sequence of selection operations are commutative
$$\sigma_q(\sigma_p(\mathbf{R})) = \sigma_p(\sigma_q(\mathbf{R}))$$
 - **Associativity of Natural Join and Cross Product:** Natural join (*) and Cross Product (X) are associative:
$$(R * S) * T = R * (S * T)$$
$$(R X S) X T = R X (S X T)$$



-- Heuristic Optimization

- Perform selection operations as early as possible
 - Keep predicates on same relation together
- Combine Cartesian product with subsequent selection whose predicate represents join condition into a join operation.
- Use associative of binary operations to rearrange leaf nodes with most restrictive selection operations first.
- Perform projection as early as possible
 - Keep projection attributes on same relations together.
- Compute common expressions once.
 - If common expression appears more than once, and result not too large, store result and reuse it when required.
 - Useful when querying views, as same expression is used to construct view each time.

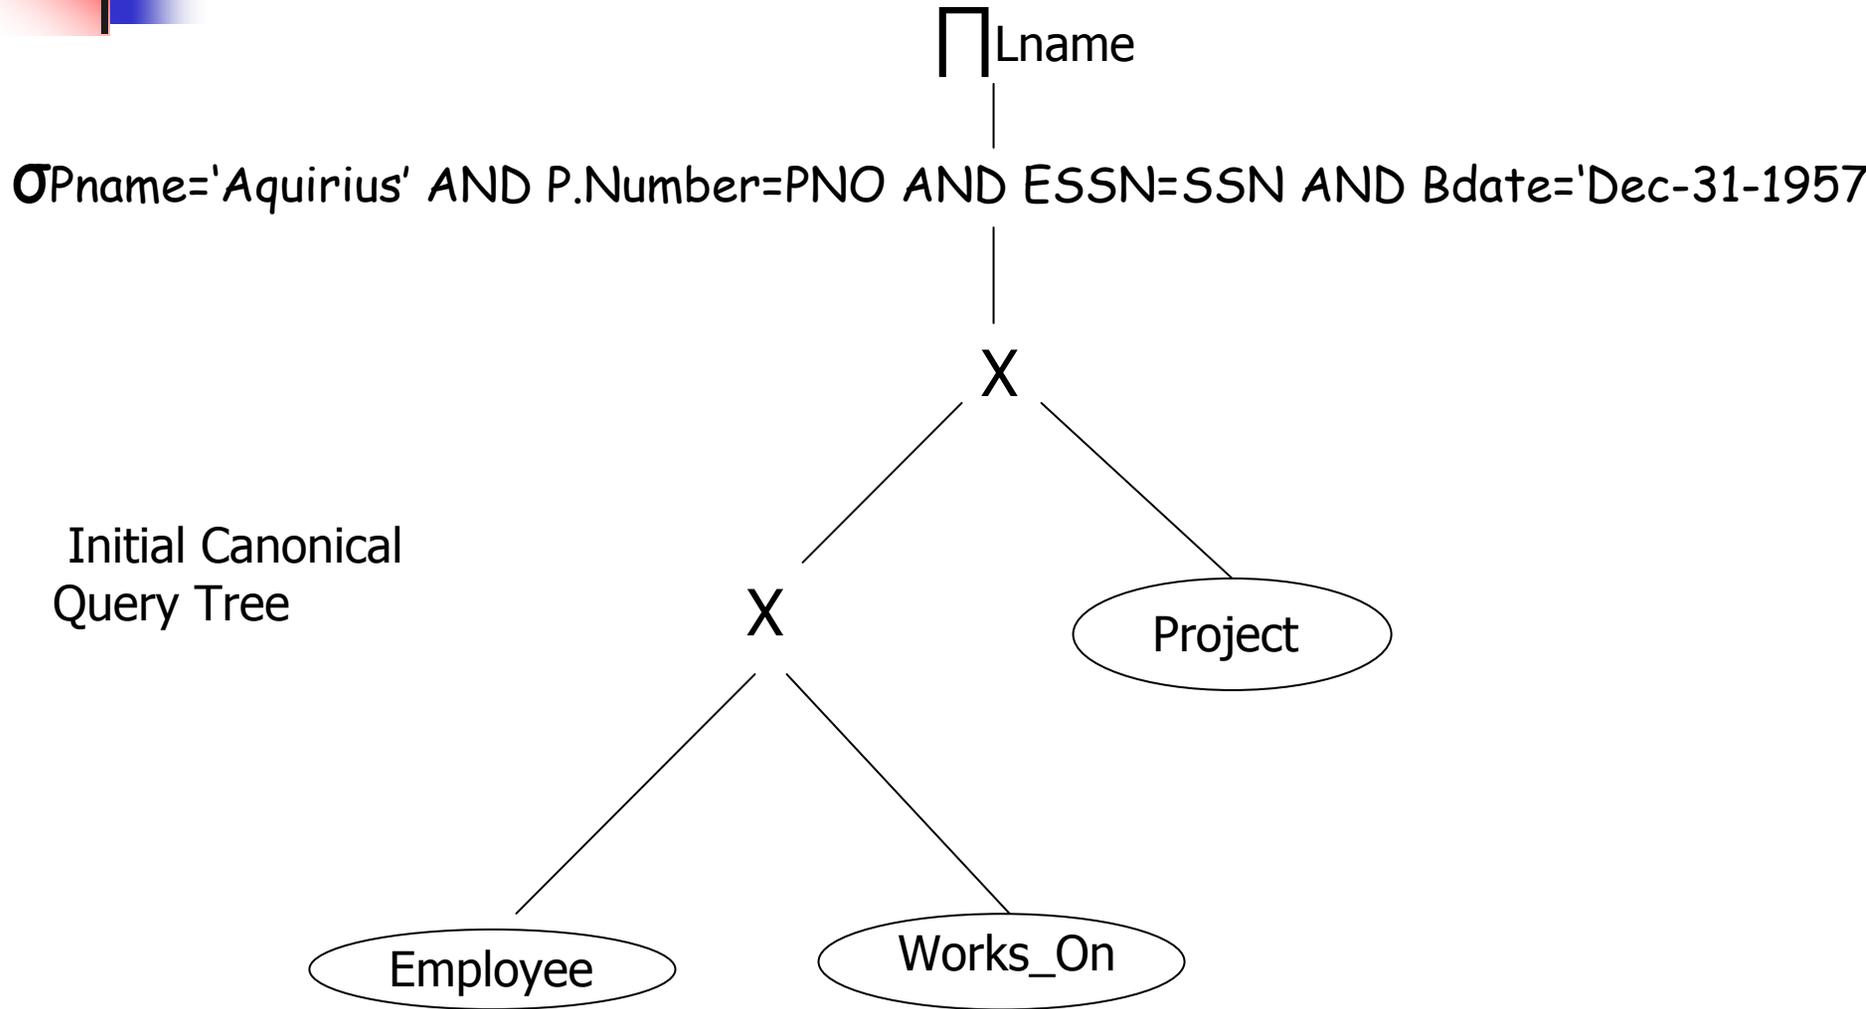


-- Query Optimization: Example ...

- Query: Find the last name of employees born after 1975 who work on a project named 'Aquarius'

```
SELECT Lname  
FROM Employee, Works_on, Project  
WHERE Pname = 'Aquarius'  
AND ESSN = SSN  
AND Pnumber = PNO  
AND Bdate > '1975-12-31'
```

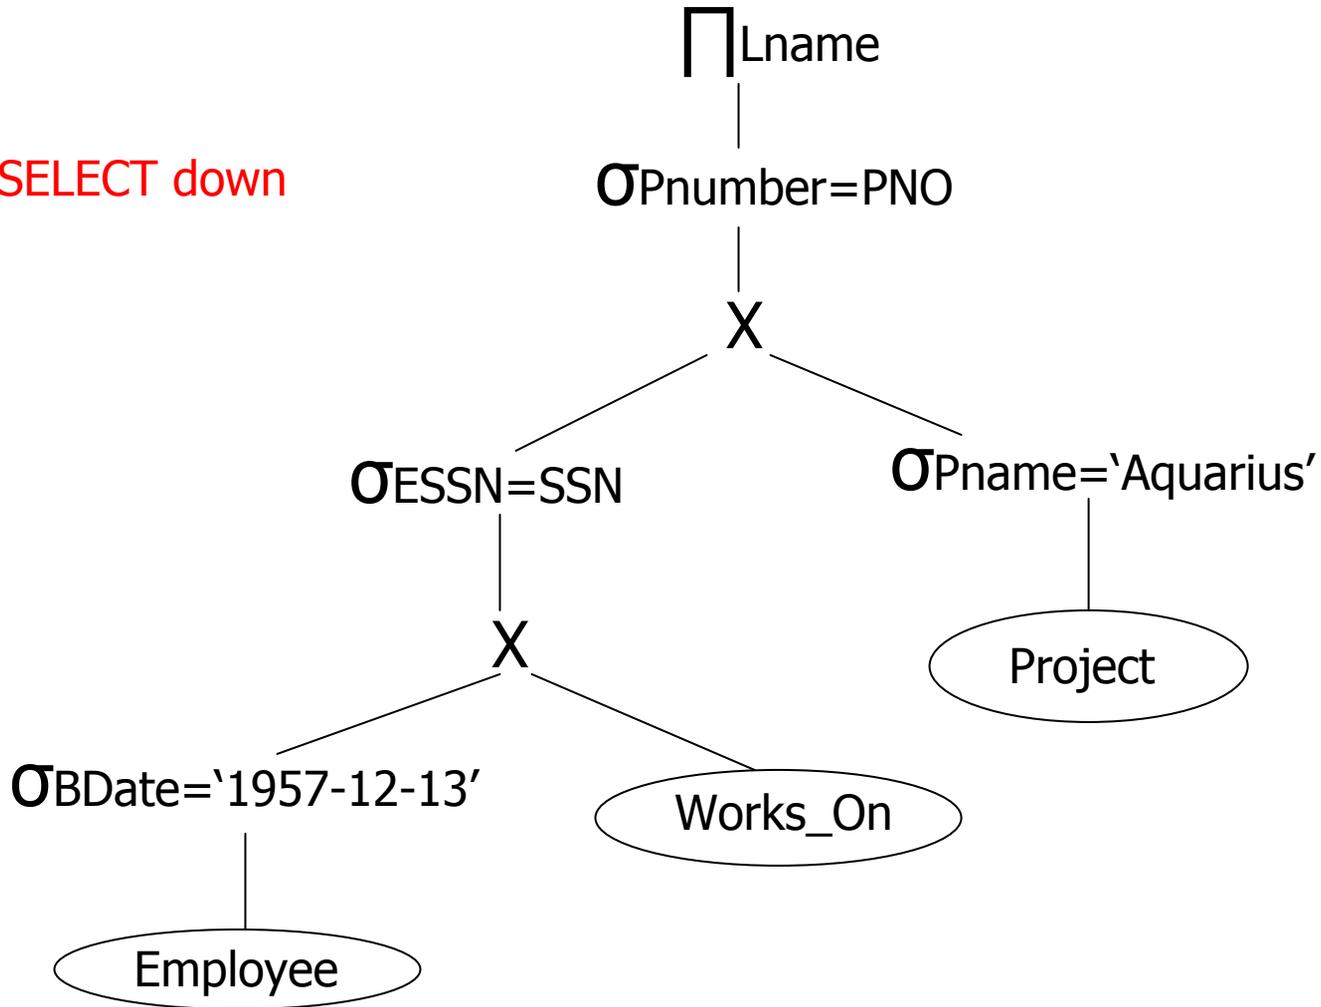
... -- Query Optimization: Example ...



Initial Canonical
Query Tree

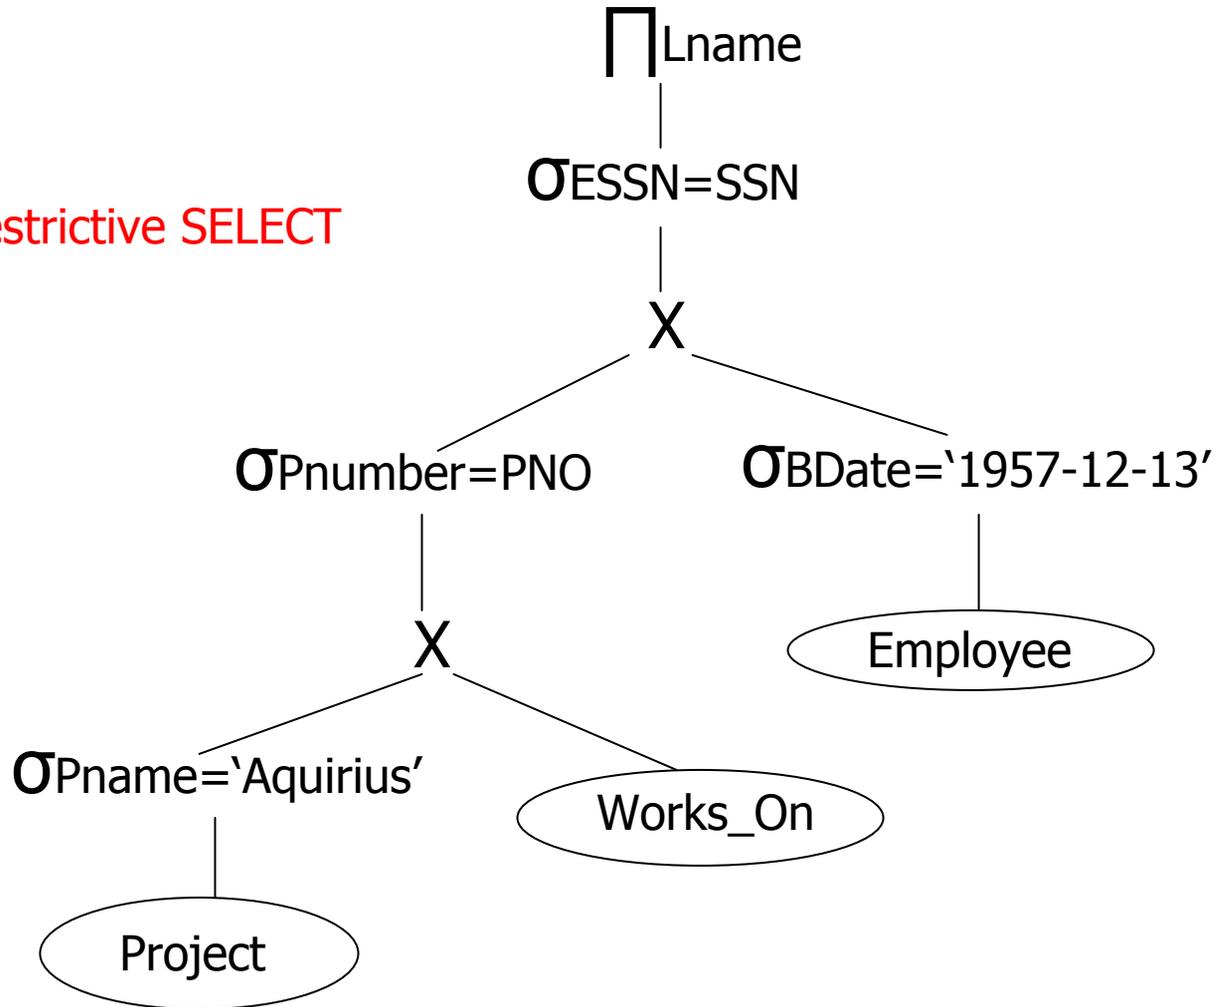
... -- Query Optimization: Example ...

Moving SELECT down



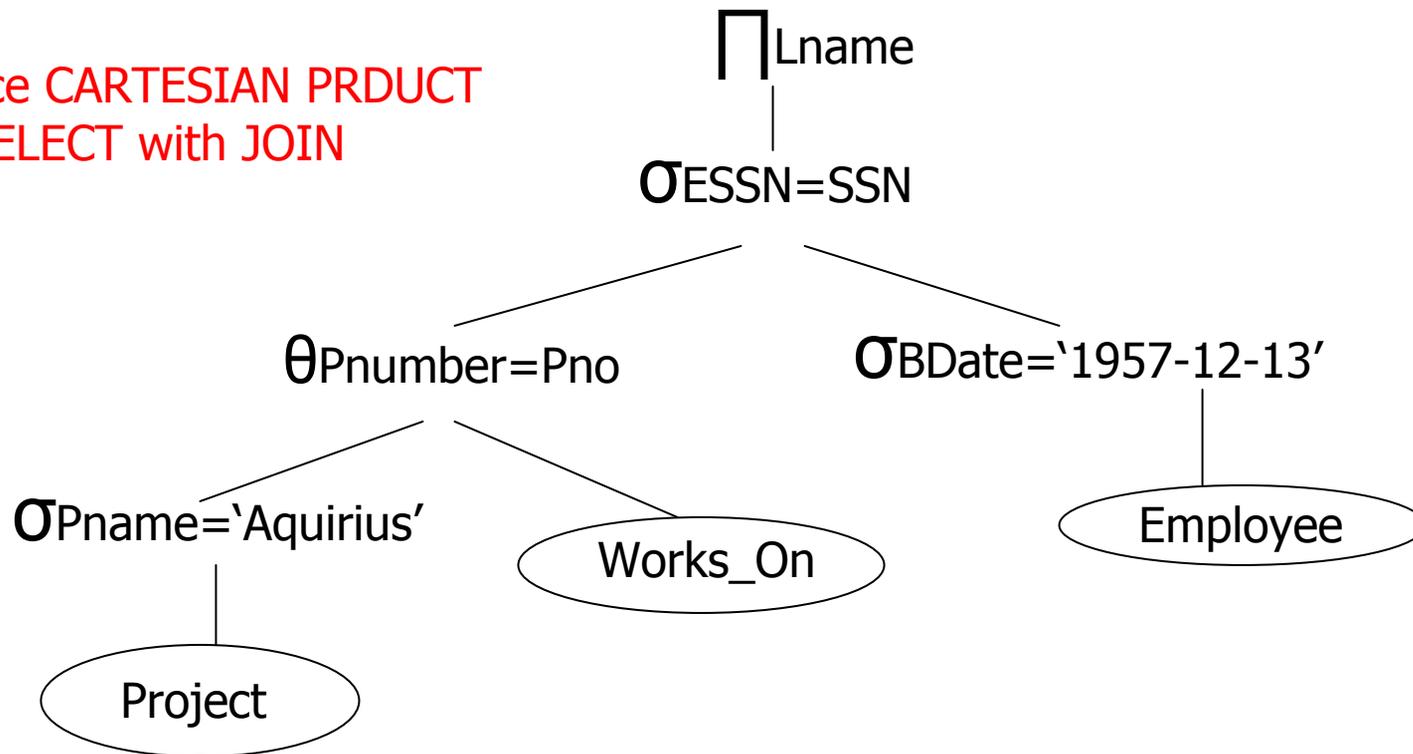
... -- Query Optimization: Example ...

Apply more restrictive SELECT



... -- Query Optimization: Example ...

Replace CARTESIAN PRODUCT
And SELECT with JOIN



... -- Query Optimization: Example

Move PROJECTION down

