


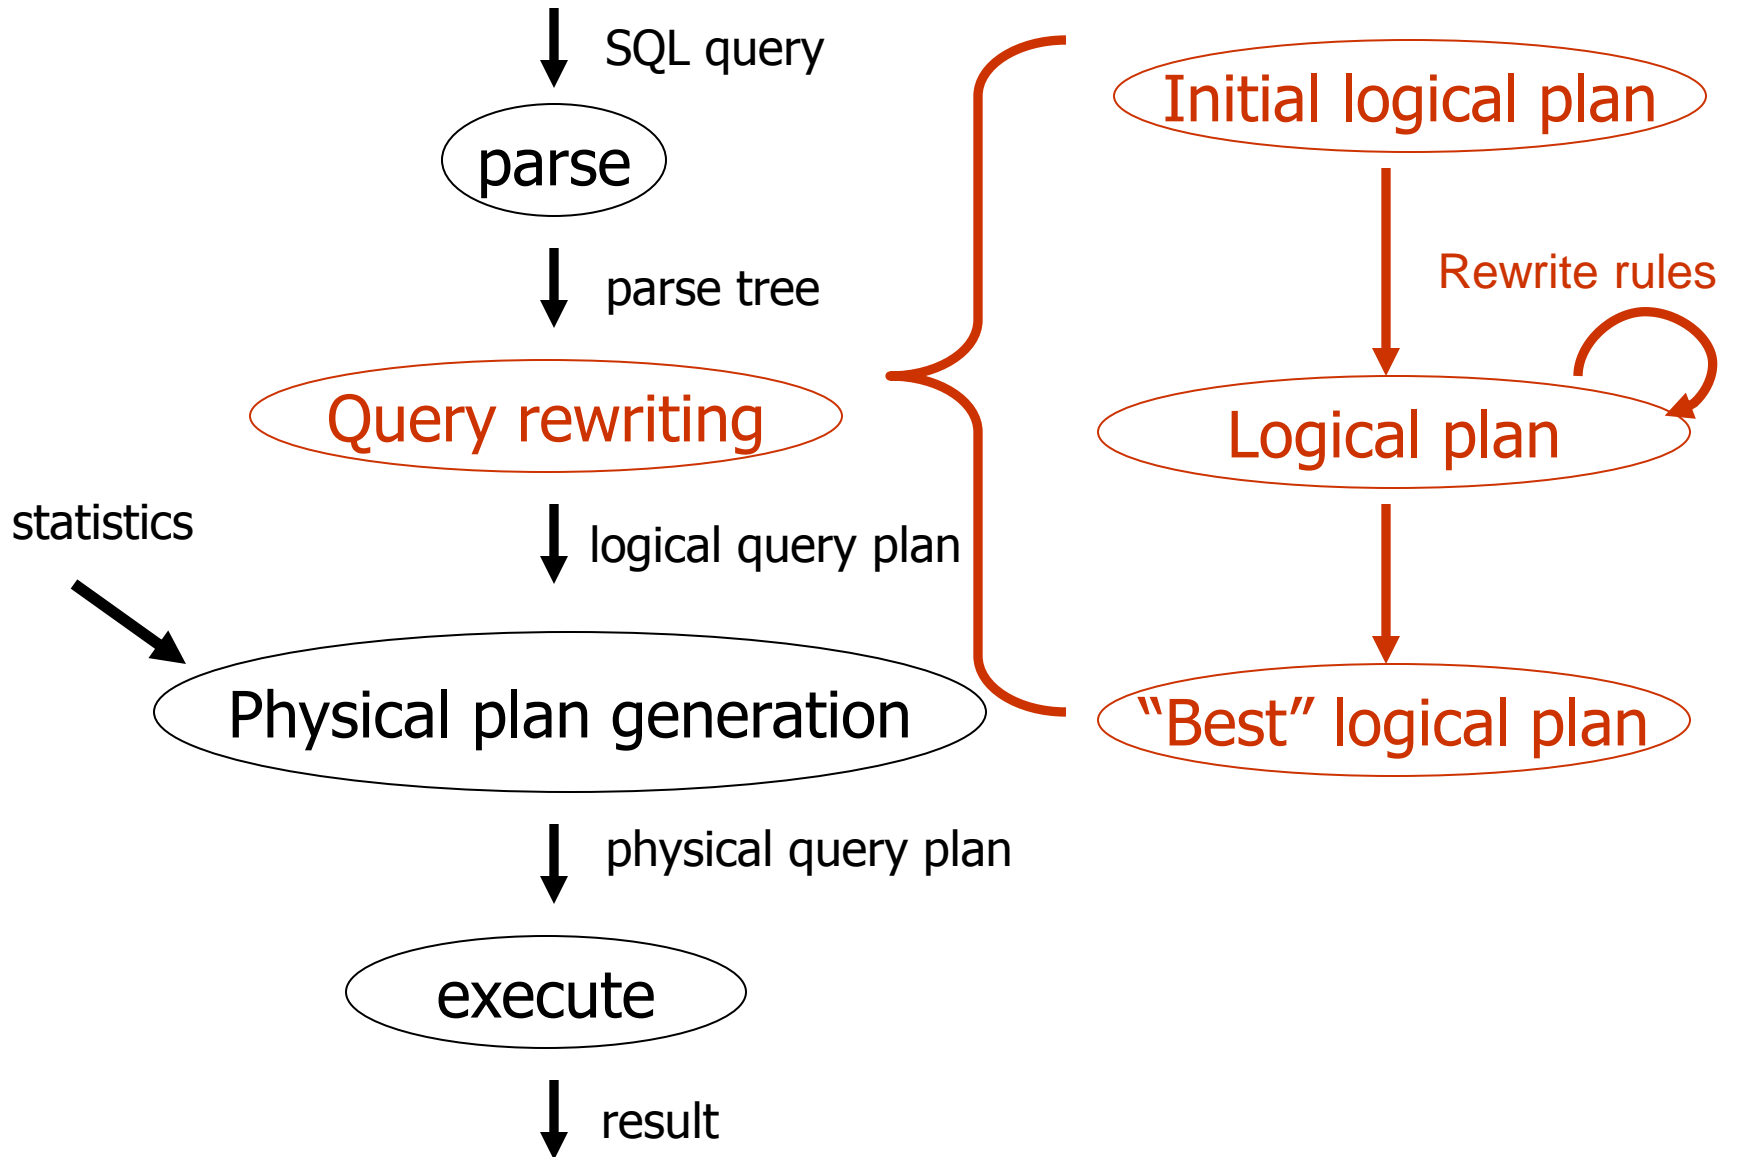
# CPS216: Data-Intensive Computing Systems

## **Query Processing (contd.)**


Shivnath Babu

# Overview of Query Processing

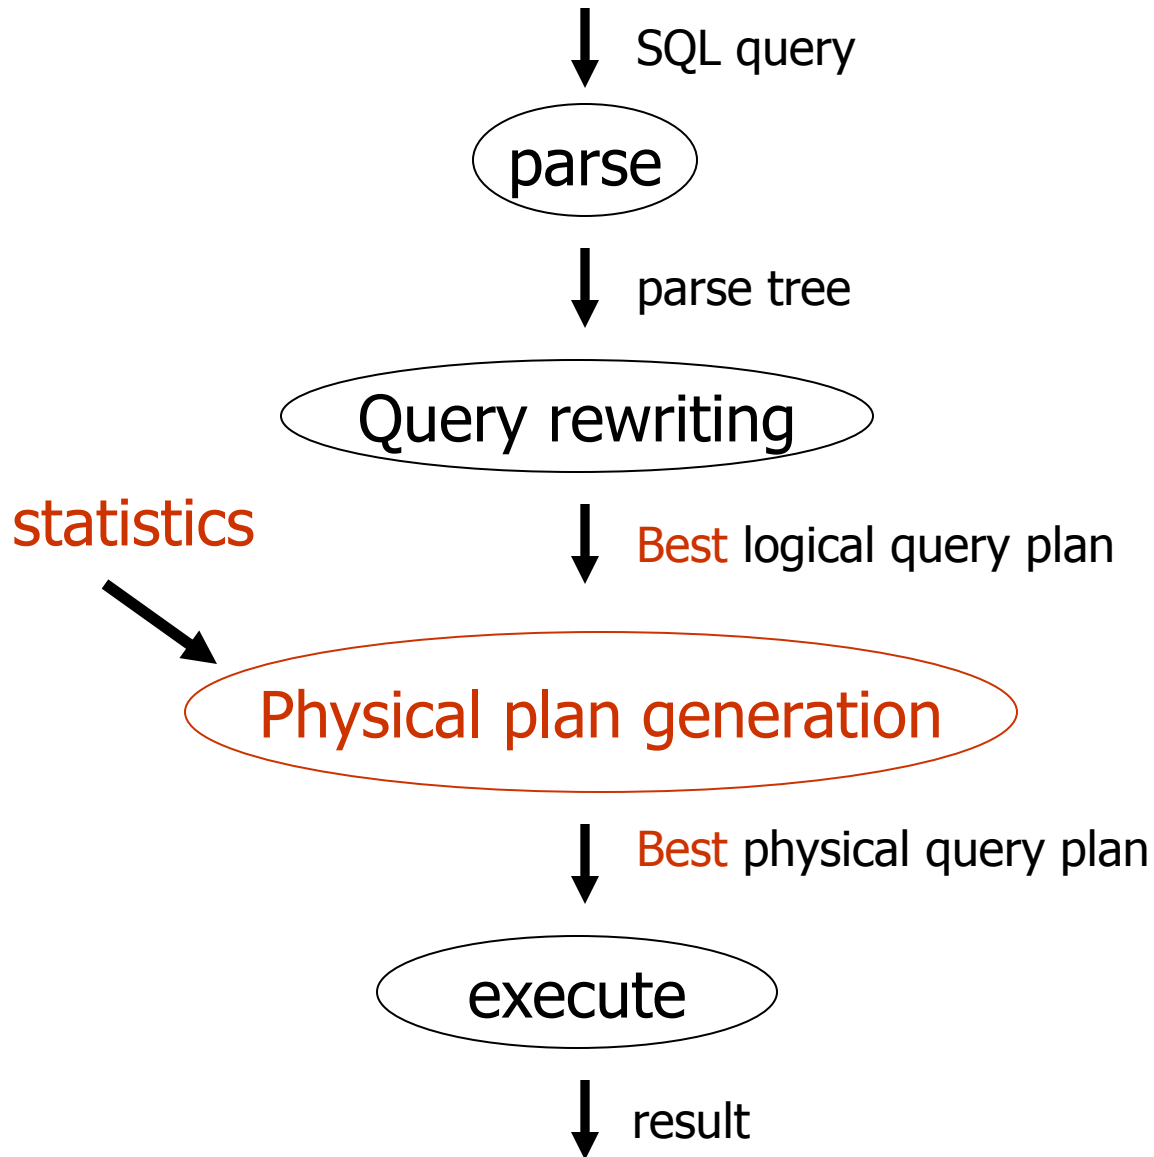





# Query Rewriting




We will revisit it towards the end of this lecture

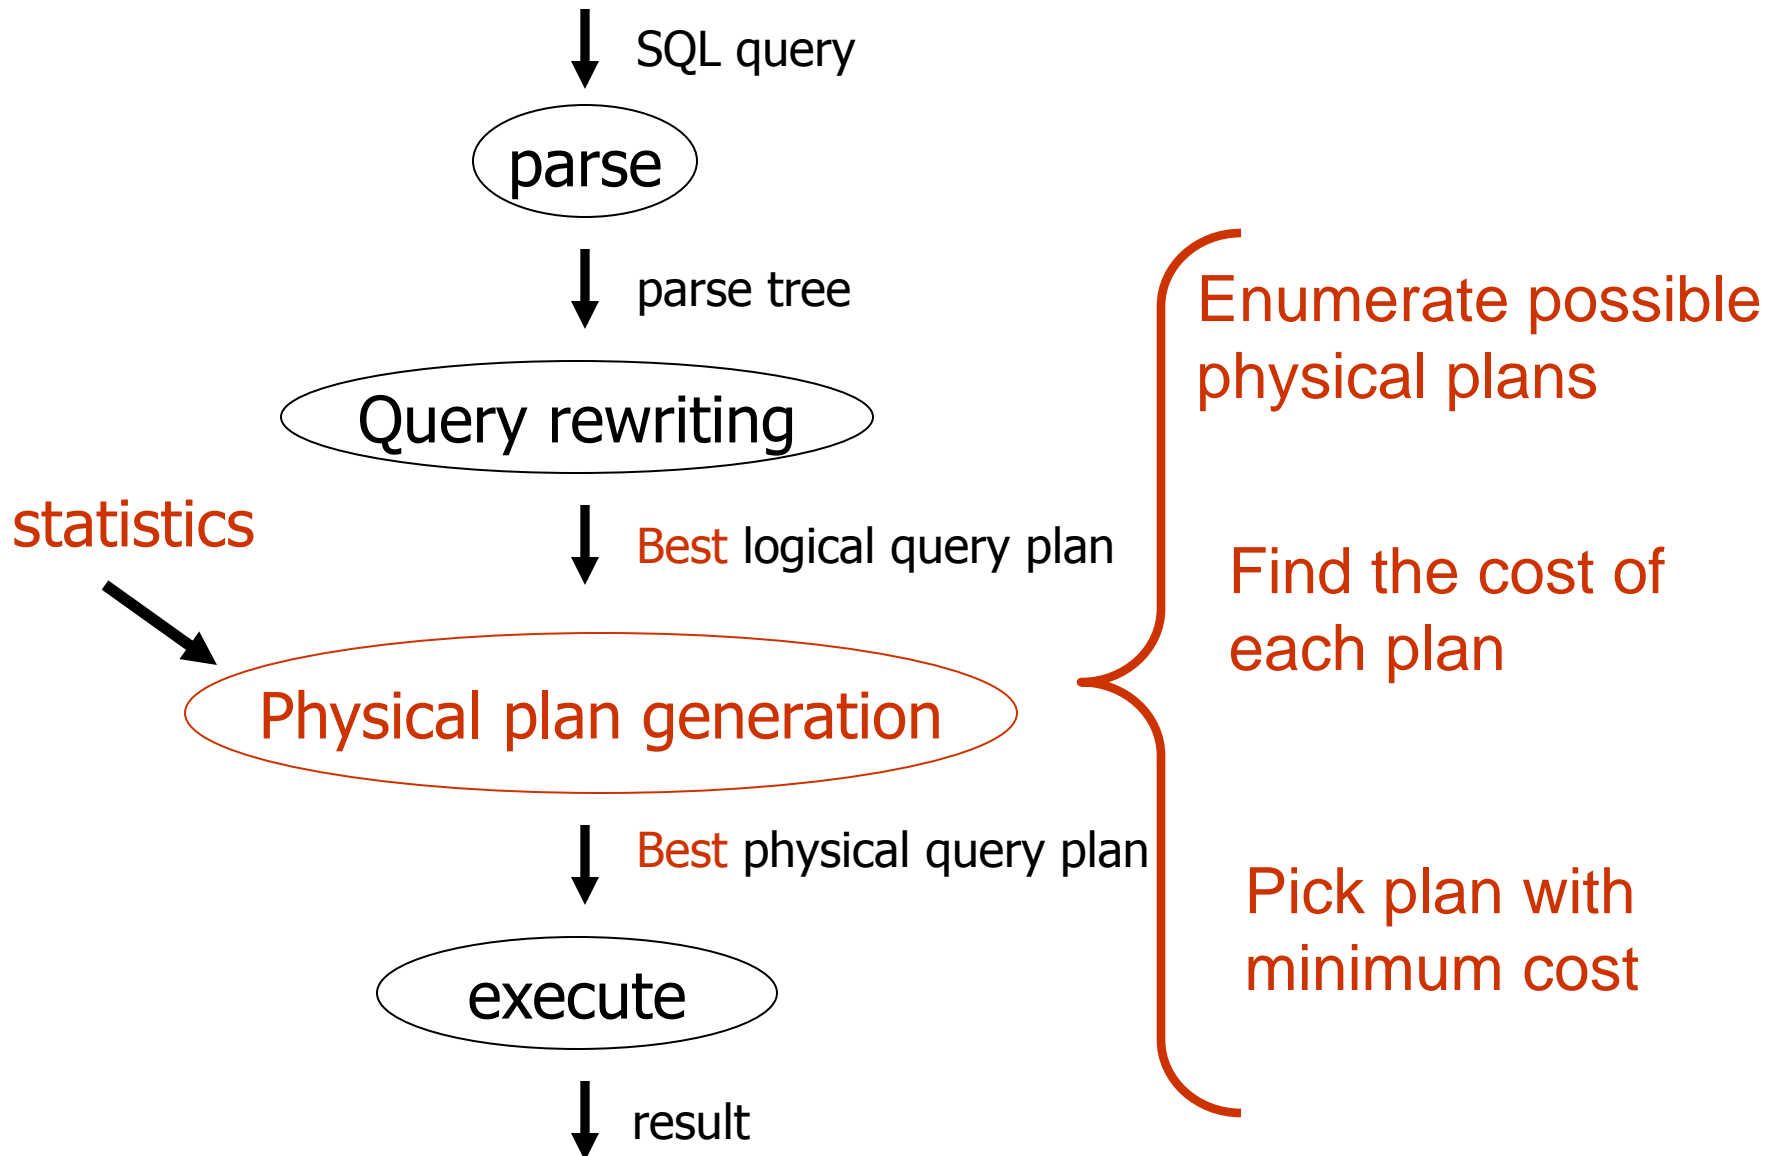


# Physical Plan Generation

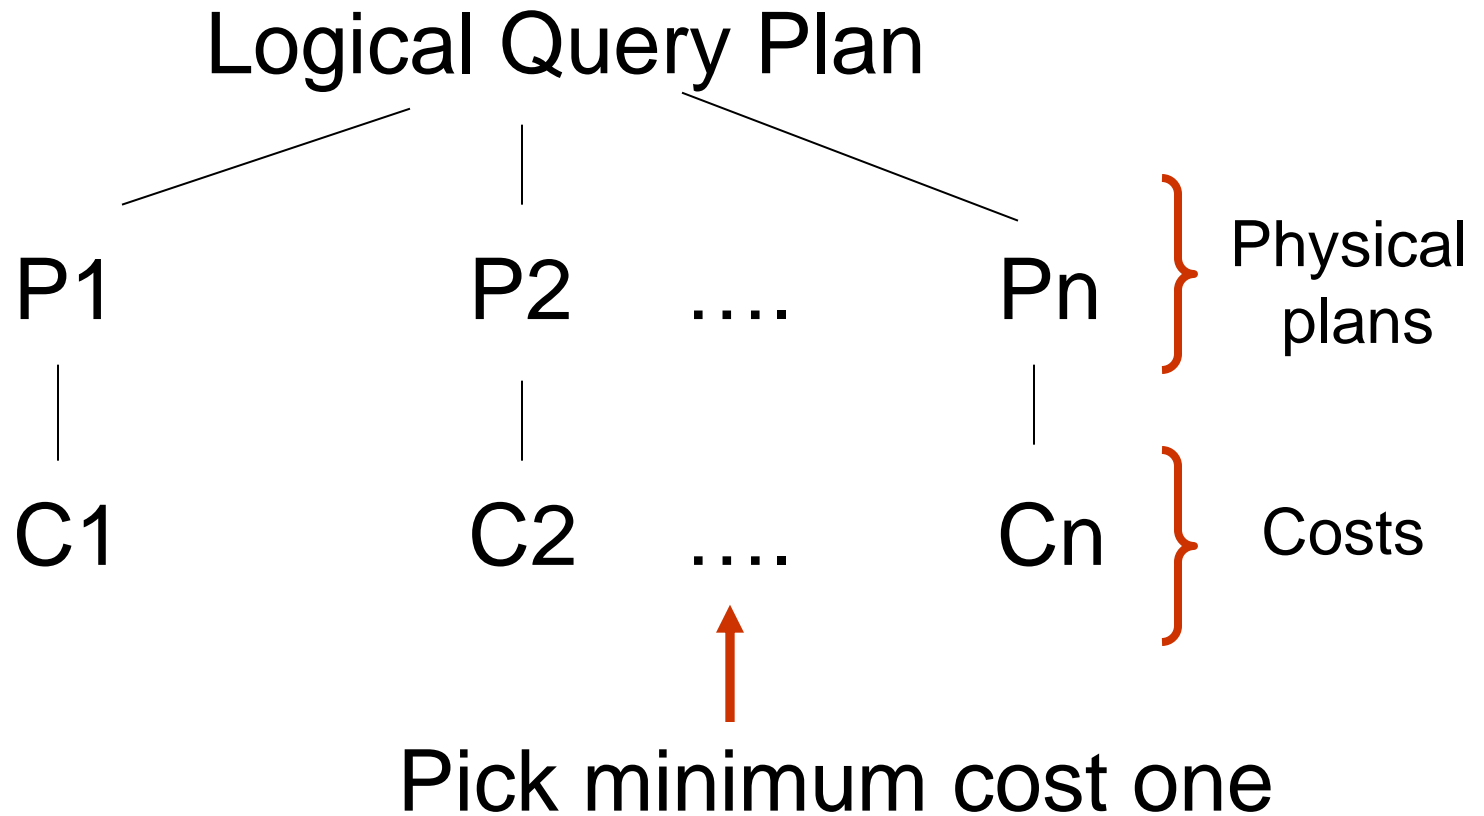


Best logical plan





# Physical Plan Generation







# Plans for Query Execution

- Roadmap
  - Path of a SQL query
  - Operator trees
  - Physical Vs Logical plans
  - Plumbing: Materialization Vs pipelining


# Logical Plans Vs. Physical Plans



Best logical plan

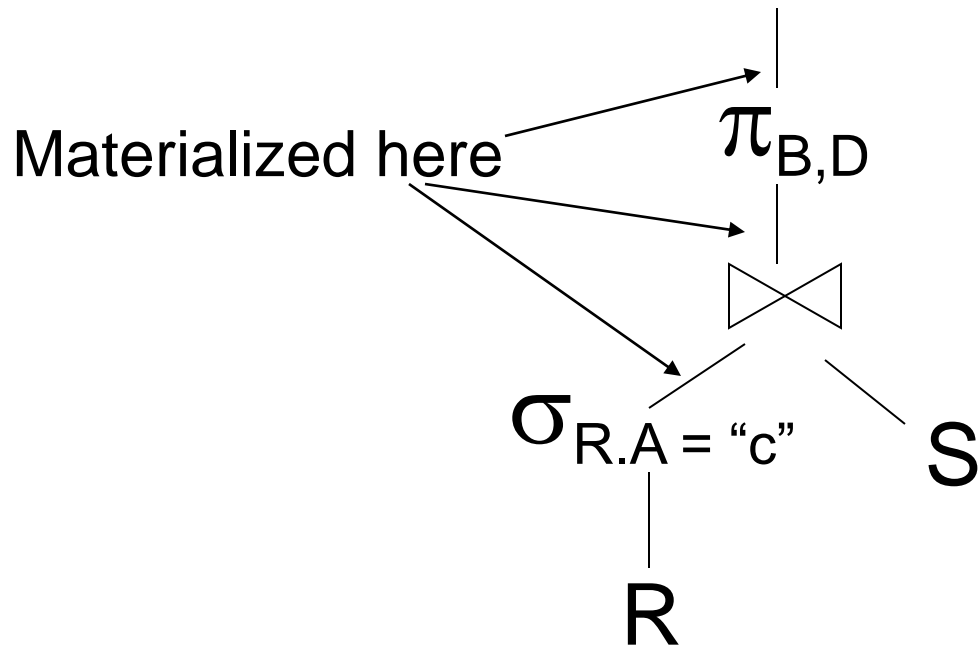


# Operator Plumbing




- **Materialization:** output of one operator written to disk, next operator reads from the disk
- **Pipelining:** output of one operator directly fed to next operator

# Materialization



# Iterators: Pipelining



→ Each operator supports:

- `Open()`
- `GetNext()`
- `Close()`

# Iterator for Table Scan (R)


```
Open() {  
    /** initialize variables */  
    b = first block of R;  
    t = first tuple in block b;  
}
```

```
Close() {  
    /** nothing to be done */  
}
```

```
GetNext() {  
    IF (t is past last tuple in block b) {  
        set b to next block;  
        IF (there is no next block)  
            /** no more tuples */  
            RETURN EOT;  
        ELSE t = first tuple in b;  
    }  
    /** return current tuple */  
    oldt = t;  
    set t to next tuple in block b;  
    RETURN oldt;  
}
```

# Iterator for Select

$\sigma_{R.A = "c"}$




```
Open() {  
    /** initialize child */  
    Child.Open();  
}
```

```
Close() {  
    /** inform child */  
    Child.Close();  
}
```

```
GetNext() {  
    LOOP:  
        t = Child.GetNext();  
        IF (t == EOT) {  
            /** no more tuples */  
            RETURN EOT;  
        }  
        ELSE IF (t.A == "c")  
            RETURN t;  
    ENDLOOP:  
}
```

# Iterator for Sort




```
getNext() {  
    IF (more tuples)  
        RETURN next tuple in order;  
    ELSE RETURN EOT;  
}
```

```
Open() {  
    /** Bulk of the work is here */  
    Child.Open();  
    Read all tuples from Child  
    and sort them  
}
```

```
Close() {  
    /** inform child */  
    Child.Close();  
}
```




# Iterator for Tuple Nested Loop Join




- TNLJ (conceptually)
  - for each  $r \in L_{exp}$  do
    - for each  $s \in R_{exp}$  do
      - if  $L_{exp}.C = R_{exp}.C$ , output  $r,s$

# Example 1: Left-Deep Plan



Question: What is the sequence of getNext() calls?

# Example 2: Right-Deep Plan



Question: What is the sequence of getNext() calls?

# Cost Measure for a Physical Plan

- There are many cost measures
  - Time to completion
  - Number of I/Os (we will see a lot of this)
  - Number of getNext() calls
- Tradeoff: Simplicity of estimation Vs. Accurate estimation of performance as seen by user

# Why do we need Query Rewriting?

- **Pruning** the HUGE space of physical plans
    - Eliminating redundant conditions/operators
    - Rules that will improve performance with very high probability
  - **Preprocessing**
    - Getting queries into a form that we know how to handle best
- ➔ Reduces optimization time drastically without noticeably affecting quality

# Some Query Rewrite Rules

- Transform one **logical plan** into another
  - Do not use statistics
- Equivalences in relational algebra
- Push-down predicates
- Do projects early
- Avoid cross-products if possible

# Equivalences in Relational Algebra

$$R \bowtie S = S \bowtie R \quad \text{Commutativity}$$

$$(R \bowtie S) \bowtie T = R \bowtie (S \bowtie T) \quad \text{Associativity}$$

Also holds for: Cross Products, Union, Intersection


$$R \times S = S \times R$$

$$(R \times S) \times T = R \times (S \times T)$$

$$R \cup S = S \cup R$$

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

# Apply Rewrite Rule (1)



$$\Pi_{B,D} [\sigma_{R.C=S.C} [\sigma_{R.A=\text{"c"}}(R \ X \ S)]]$$

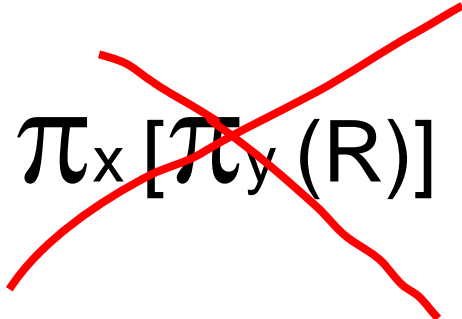


# Rules: Project

Let:  $X$  = set of attributes

$Y$  = set of attributes

$XY = X \cup Y$

$$\pi_{xy}(R) = \pi_x[\pi_y(R)]$$


## Rules: $\sigma + \bowtie$ combined

Let  $p$  = predicate with only R attribs


$q$  = predicate with only S attribs

$m$  = predicate with only R,S attribs

$$\sigma_p (R \bowtie S) = [\sigma_p (R)] \bowtie S$$


$$\sigma_q (R \bowtie S) = R \bowtie [\sigma_q (S)]$$

## Apply Rewrite Rule (2)



$\Pi_{B,D} [ \sigma_{R.C=S.C} [ \sigma_{R.A="c"}(R) ] X S ]$

# Apply Rewrite Rule (3)



$$\Pi_{B,D} [[\sigma_{R.A='c'}(R)] \bowtie S]$$

Rules:  $\sigma + \bowtie$  combined (continued)

$$\sigma_{p \wedge q} (R \bowtie S) = [\sigma_p (R)] \bowtie [\sigma_q (S)]$$

$$\sigma_{p \wedge q \wedge m} (R \bowtie S) =$$
$$\sigma_m [(\sigma_p R) \bowtie (\sigma_q S)]$$

$$\sigma_{p \vee q} (R \bowtie S) =$$
$$[(\sigma_p R) \bowtie S] \cup [R \bowtie (\sigma_q S)]$$

Which are “good” transformations?

$$\square \sigma_{p_1 \wedge p_2} (R) \rightarrow \sigma_{p_1} [\sigma_{p_2} (R)]$$

$$\square \sigma_p (R \bowtie S) \rightarrow [\sigma_p (R)] \bowtie S$$

$$\square R \bowtie S \rightarrow S \bowtie R$$

$$\square \pi_x [\sigma_p (R)] \rightarrow \pi_x \{ \sigma_p [\pi_{xz} (R)] \}$$


Conventional wisdom: do projects early

Example:  $R(A,B,C,D,E)$

$P: (A=3) \wedge (B=\text{"cat"})$

$\pi_E \{ \sigma_P (R) \}$  vs.  $\pi_E \{ \sigma_P \{ \pi_{ABE}(R) \} \}$

# But: What if we have A, B indexes?





## Bottom line:

- No transformation is always good
- Some are usually good:
  - Push selections down
  - Avoid cross-products if possible
  - Subqueries → Joins

# Avoid Cross Products (if possible)

Select B,D

From R,S,T,U

Where  $R.A = S.B \wedge$

$R.C = T.C \wedge R.D = U.D$

- Which join trees avoid cross-products?
- If you can't avoid cross products, perform them as late as possible

# More Query Rewrite Rules

- Transform one **logical plan** into another
  - Do not use statistics
- Equivalences in relational algebra
- Push-down predicates
- Do projects early
- Avoid cross-products if possible
- **Use left-deep trees**
- **Subqueries → Joins**
- **Use of constraints, e.g., uniqueness**