

Data-Intensive Computing Systems

Introduction to Query Processing

Shivnath Babu

Query Processing

Declarative **SQL** Query → Query Plan

NOTE: You will not be tested on how well you know SQL. Understanding the SQL introduced in class will be sufficient (a primer follows). SQL is described in Chapter 6, GMUW.

Focus: Relational System (i.e., data is organized as tables, or relations)

SQL Primer

We will focus on SPJ, or Select-Project-Join Queries

Select <attribute list>

From <relation list>

Where <condition list>

Example Filter Query over R(A,B,C):

Select B

From R

Where R.A = "c" \wedge R.C > 10

SQL Primer (contd.)

We will focus on SPJ, or Select-Project-Join-Queries

Select <attribute list>

From <relation list>

Where <condition list>

Example Join Query over R(A,B,C) and S(C,D,E):

Select B, D

From R, S

Where R.A = "c" \wedge S.E = 2 \wedge R.C = S.C

R	A	B	C	S	C	D	E
a	1	10	10	10	x	2	
b	1	20	20	20	y	2	
c	2	10	30	30	z	2	
d	2	35	40	40	x	1	
e	3	45	50	50	y	3	

Select B,D

From R,S

Where R.A = "c" \wedge

S.E = 2 \wedge R.C=S.C

Answer

B	D
2	x

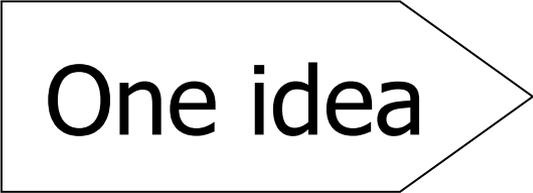
- How do we execute this query?

Select B,D

From R,S

Where R.A = "c" \wedge S.E = 2 \wedge

R.C=S.C



One idea

- Do Cartesian product
- Select tuples
- Do projection

R X S

R.A	R.B	R.C	S.C	S.D	S.E
-----	-----	-----	-----	-----	-----

Select B,D

From R,S

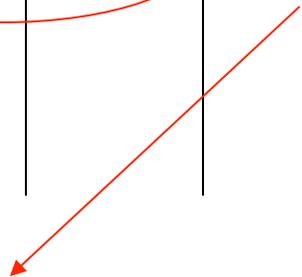
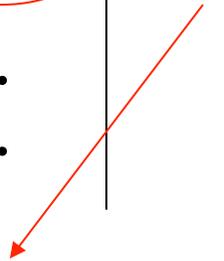
Where R.A = "c"

\wedge S.E = 2 \wedge

R.C=S.C

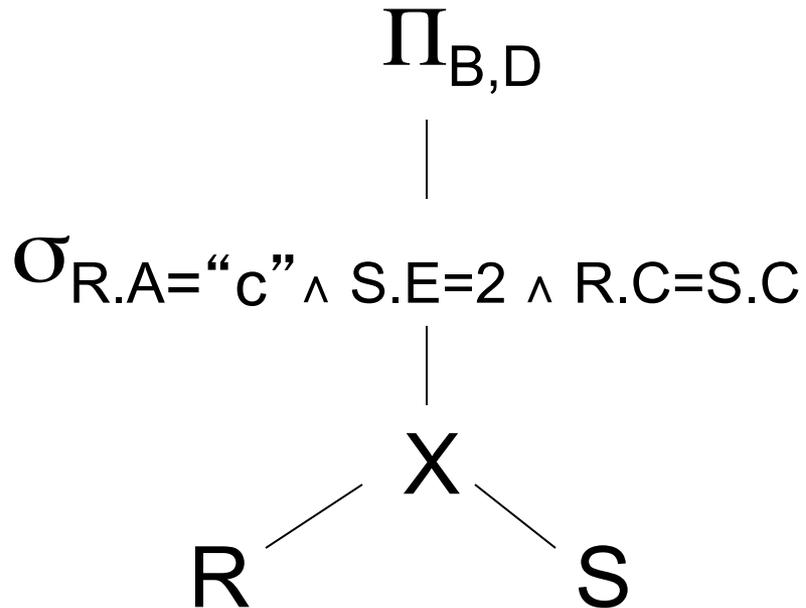
a	1	10	10	x	2
a	1	10	20	y	2
.					
.					
c	2	10	10	x	2
.					
.					

Bingo! →
Got one...



Relational Algebra - can be used to describe plans

Ex: Plan 1



Relational Algebra Primer (Chapter 5, GMUW)

Select: $\sigma_{R.A="c" \wedge R.C=10}$

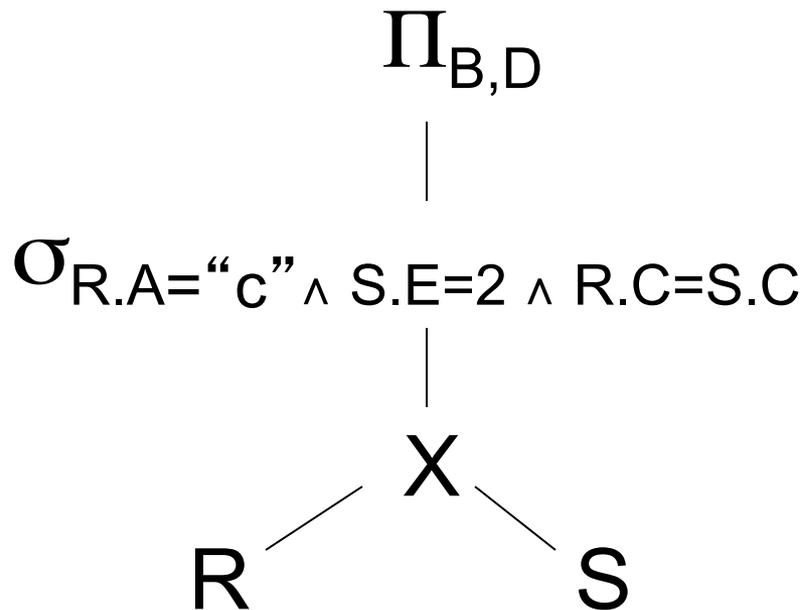
Project: $\Pi_{B,D}$

Cartesian Product: $R \times S$

Natural Join: $R \bowtie S$

Relational Algebra - can be used to describe plans

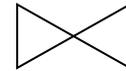
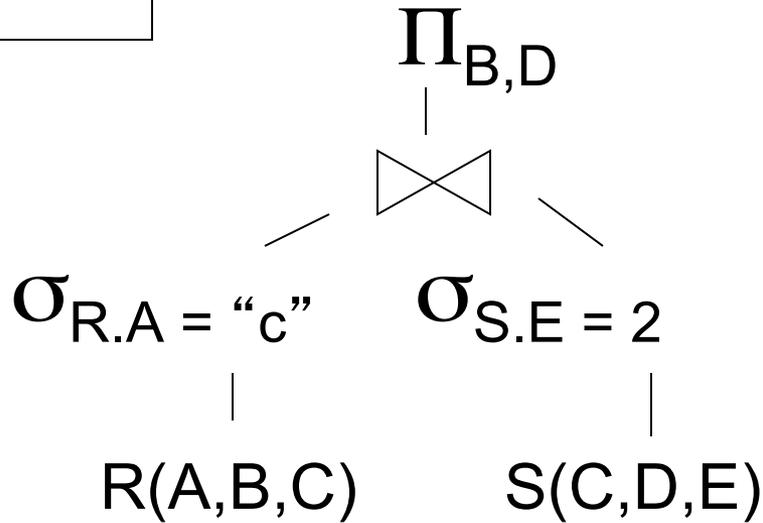
Ex: Plan 1



OR: $\Pi_{B,D} [\sigma_{R.A="c" \wedge S.E=2 \wedge R.C=S.C} (RXS)]$

Another idea:

Plan II



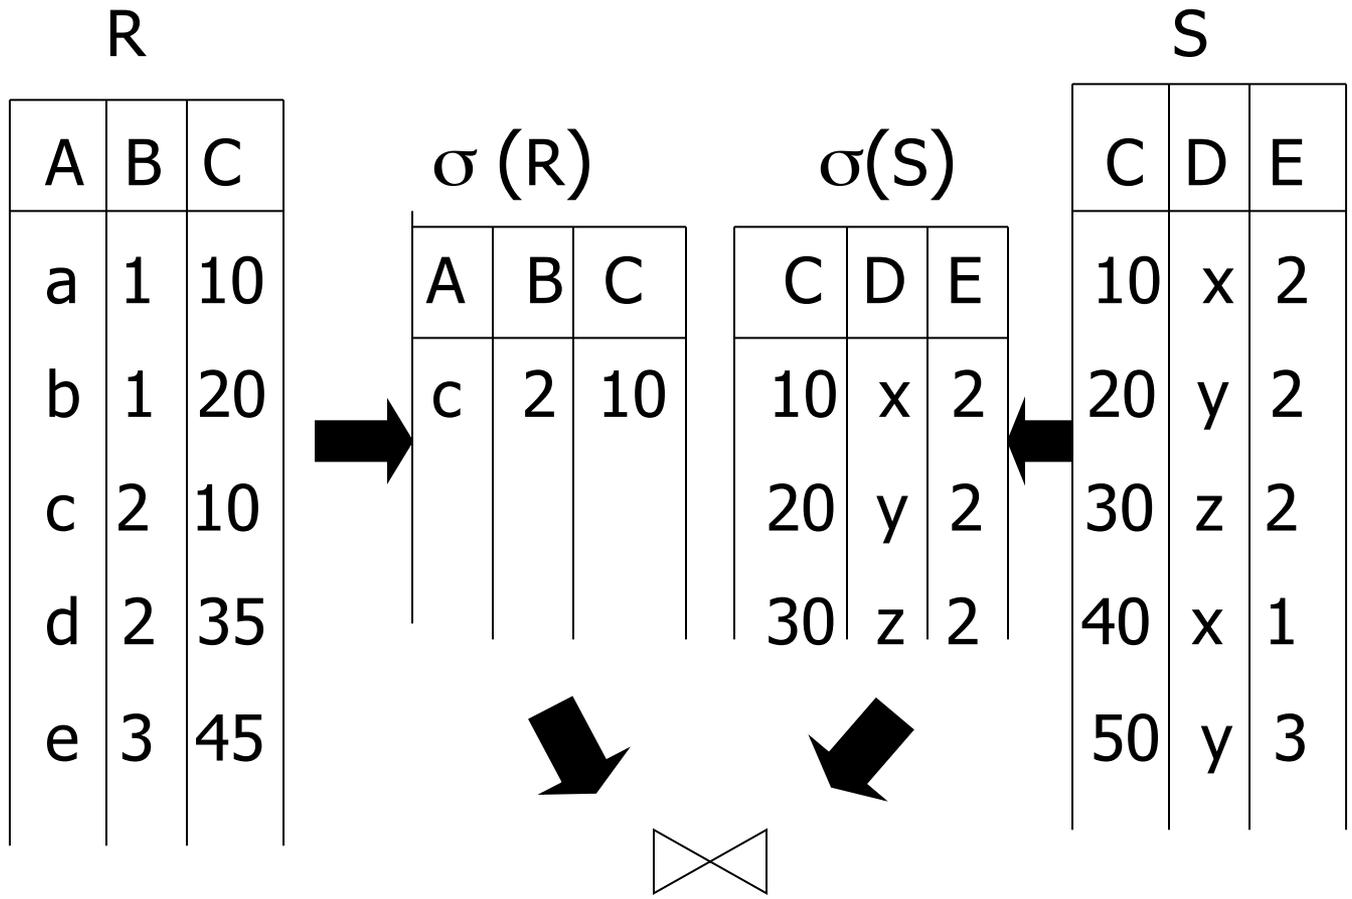
natural join

Select B,D

From R,S

Where $R.A = \text{"c"} \wedge$

$S.E = 2 \wedge R.C = S.C$



Select B,D

From R,S

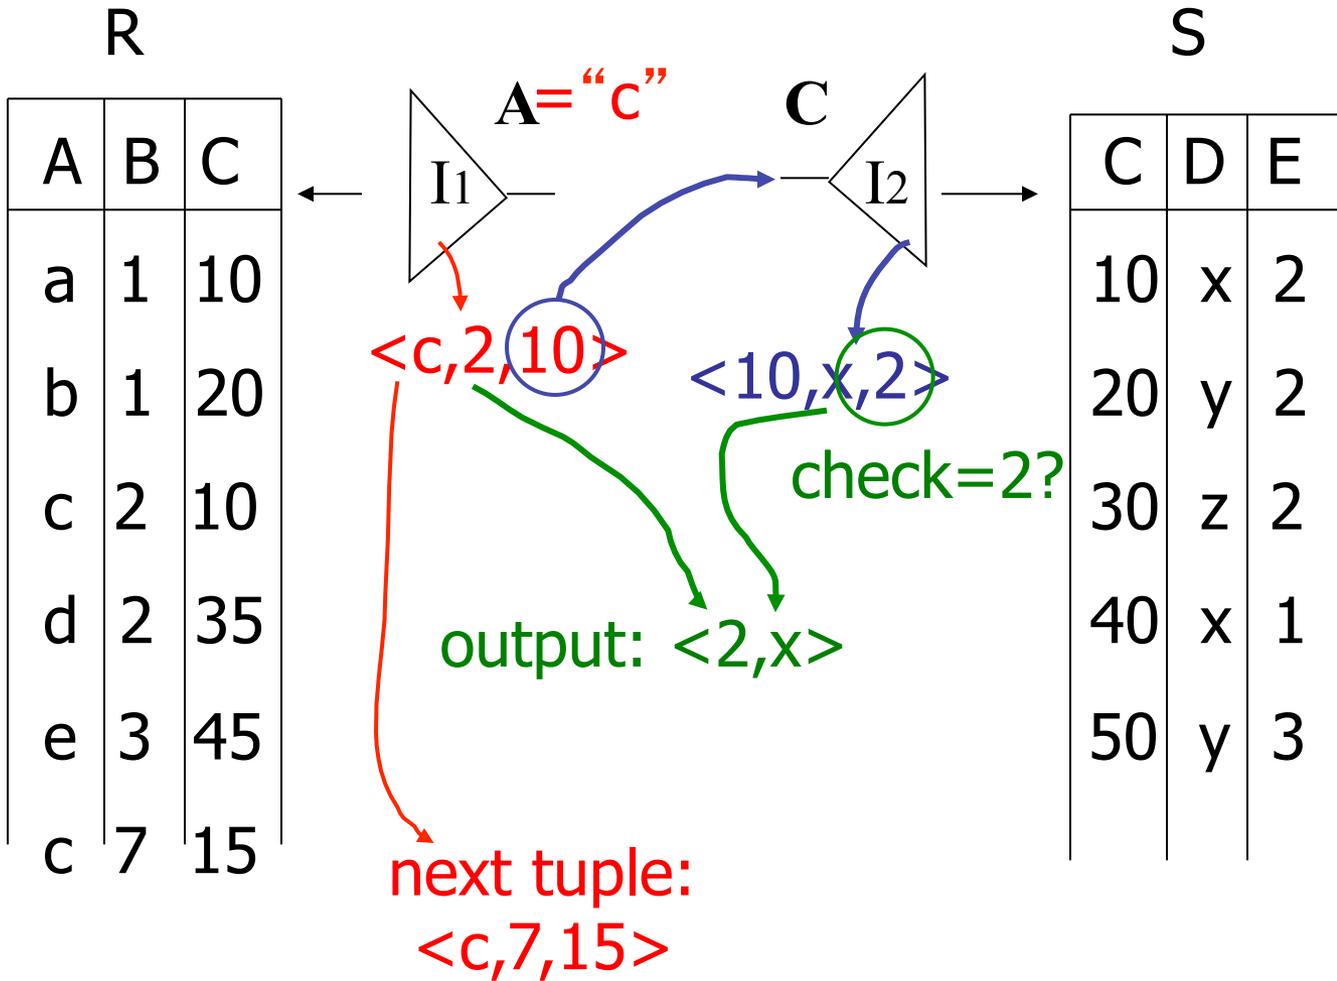
Where R.A = "c" \wedge

S.E = 2 \wedge R.C=S.C

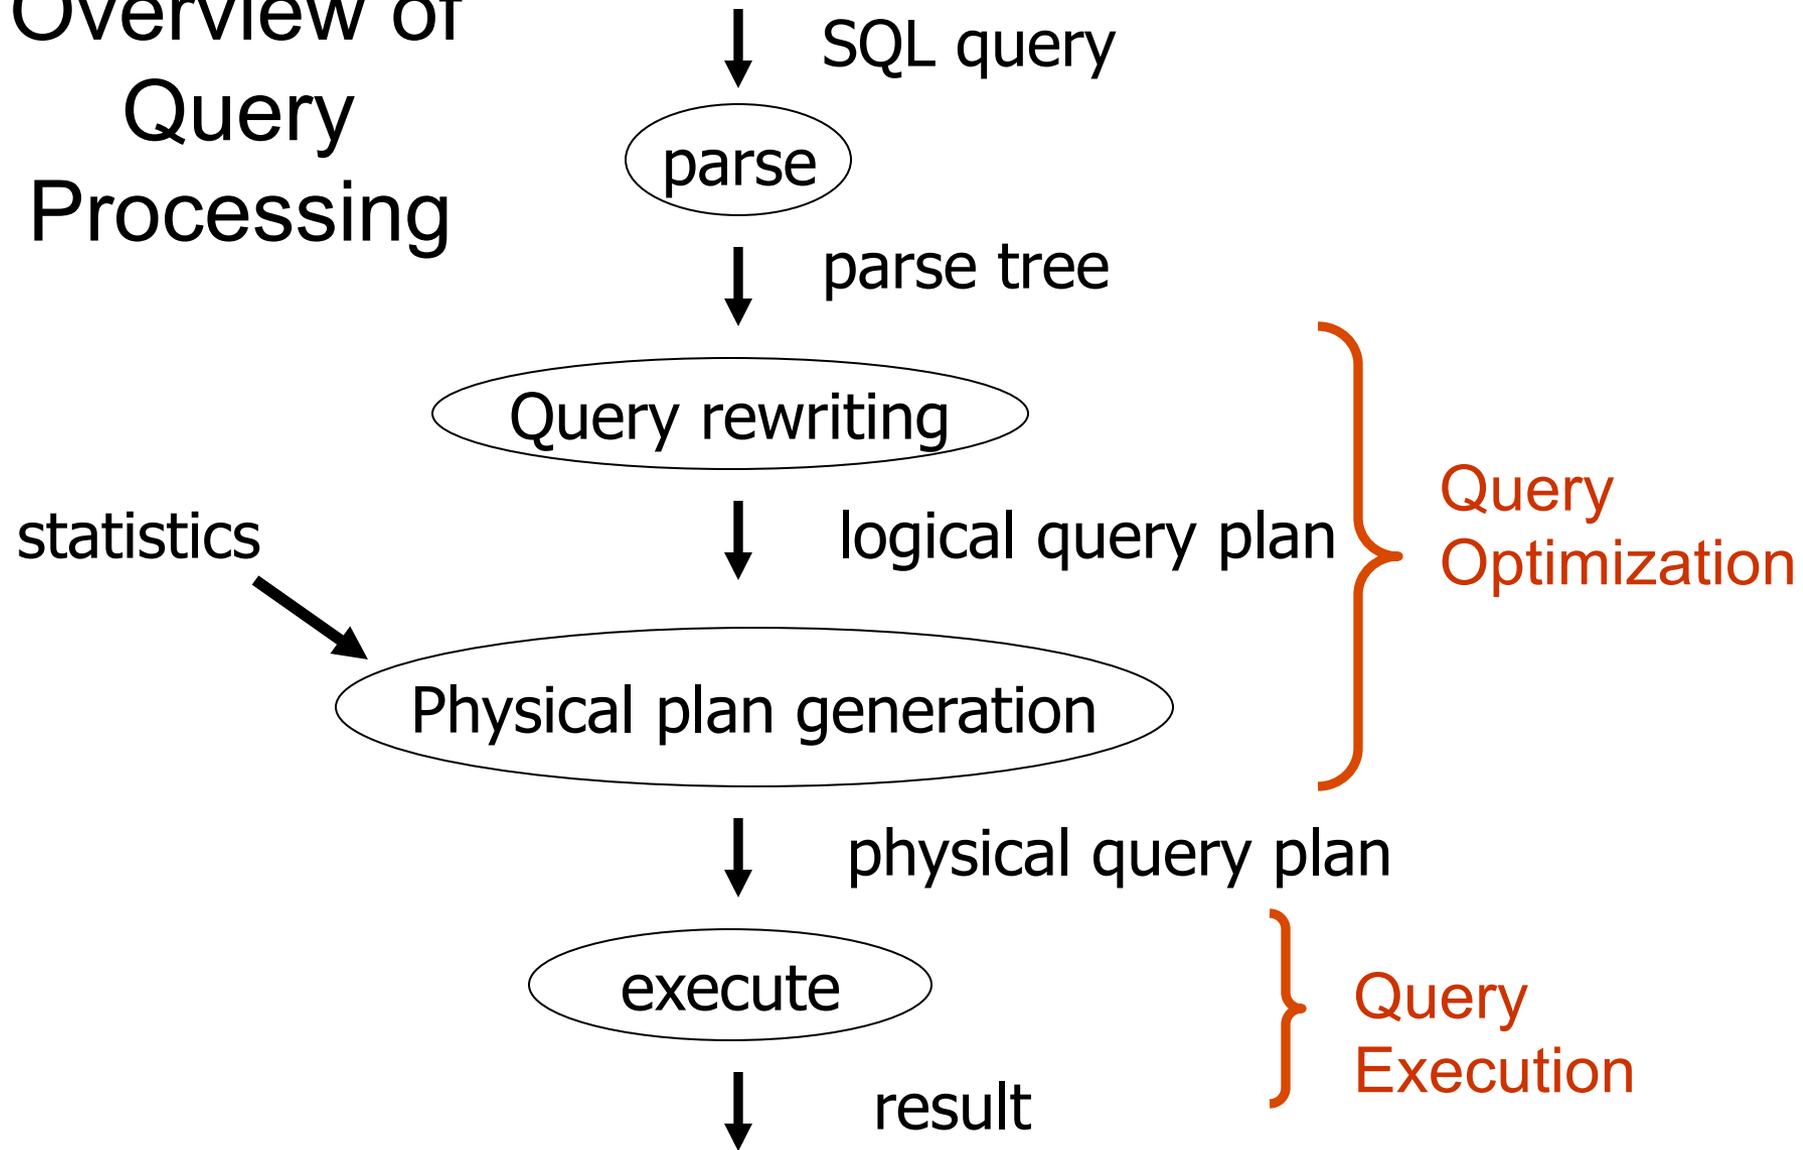
Plan III

Use R.A and S.C **Indexes**

- (1) Use R.A index to select R tuples with R.A = "c"
- (2) For each R.C value found, use S.C index to find matching tuples
- (3) Eliminate S tuples S.E \neq 2
- (4) Join matching R,S tuples, project B,D attributes, and place in result



Overview of Query Processing



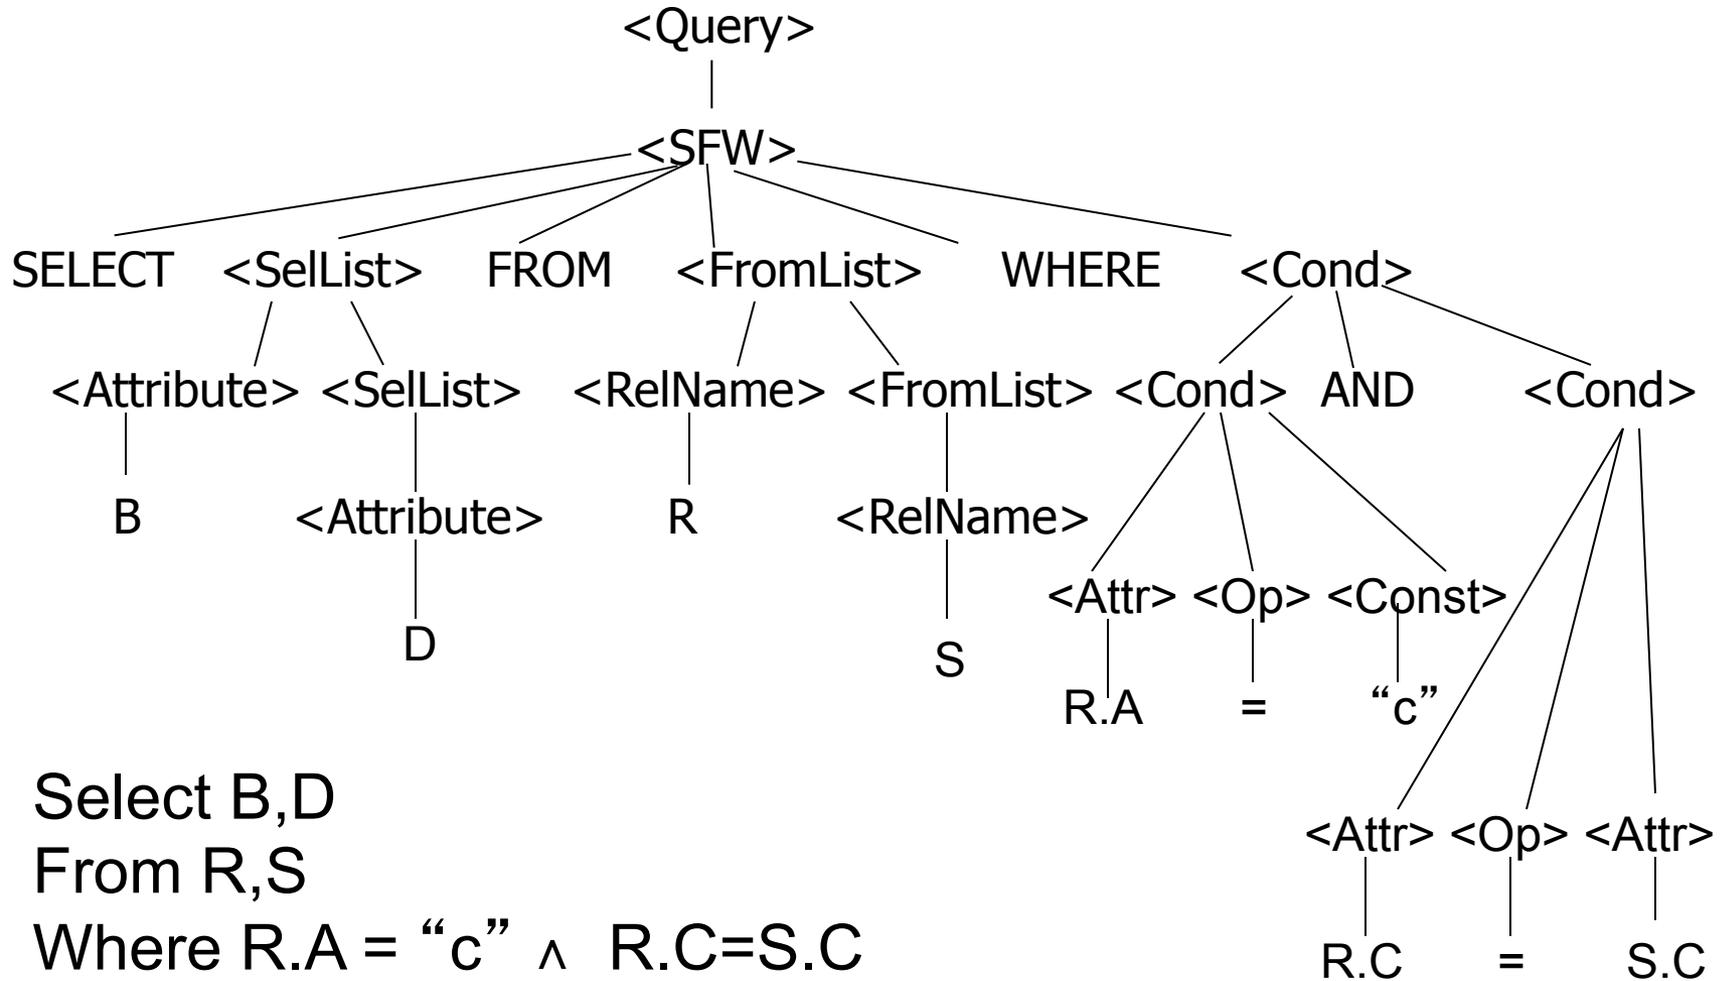
Example Query

Select B,D

From R,S

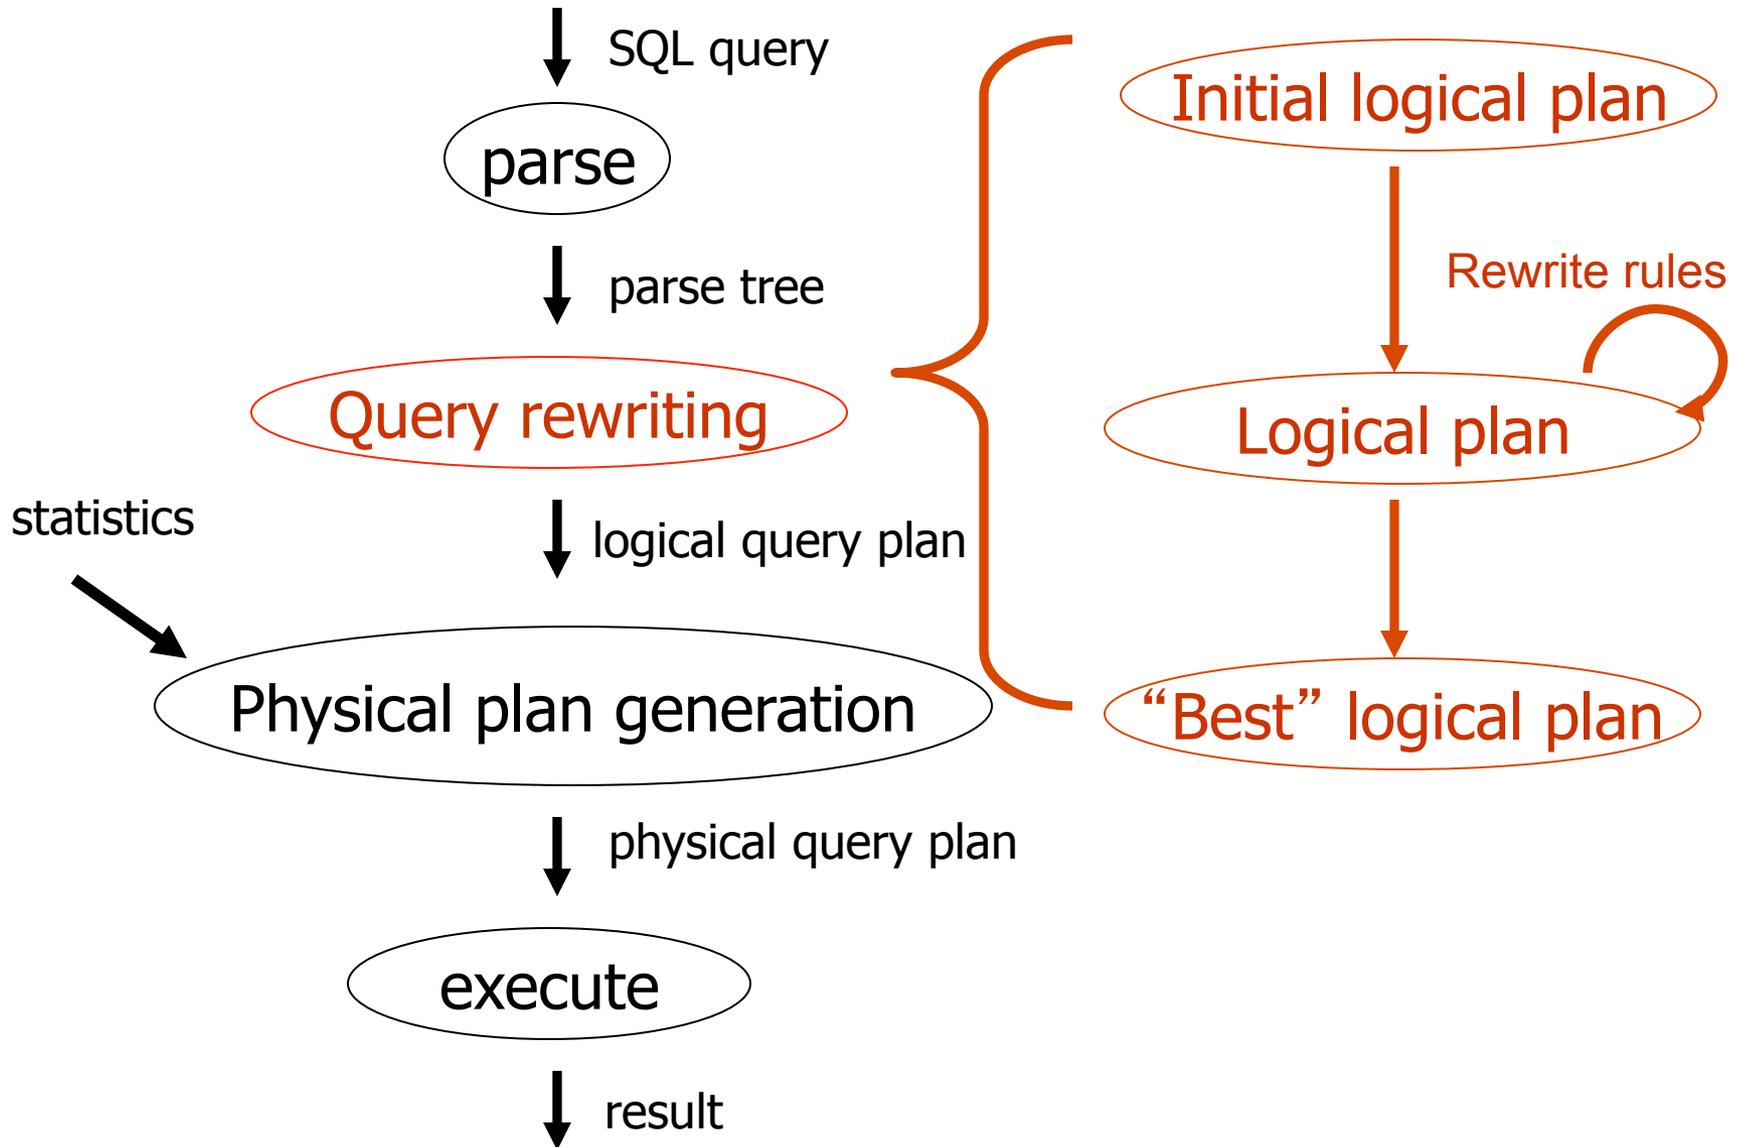
Where $R.A = \text{"c"} \wedge R.C=S.C$

Example: Parse Tree



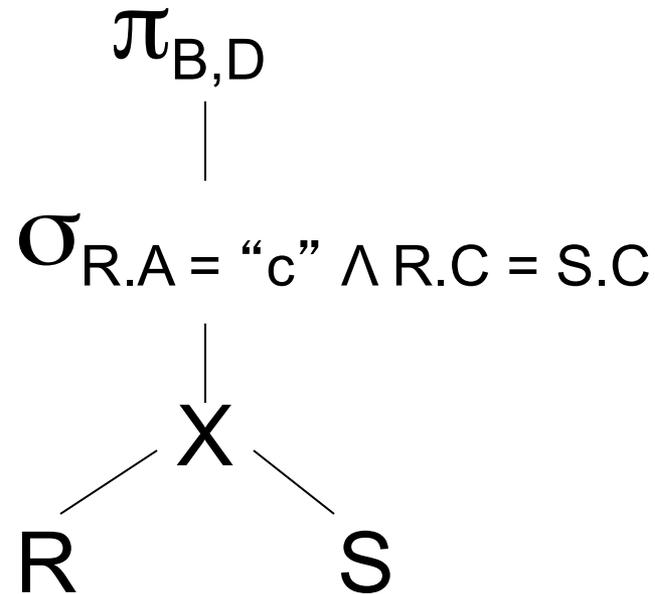
Along with Parsing ...

- Semantic checks
 - Do the projected attributes exist in the relations in the From clause?
 - Ambiguous attributes?
 - Type checking, ex: $R.A > 17.5$
- Expand **views**



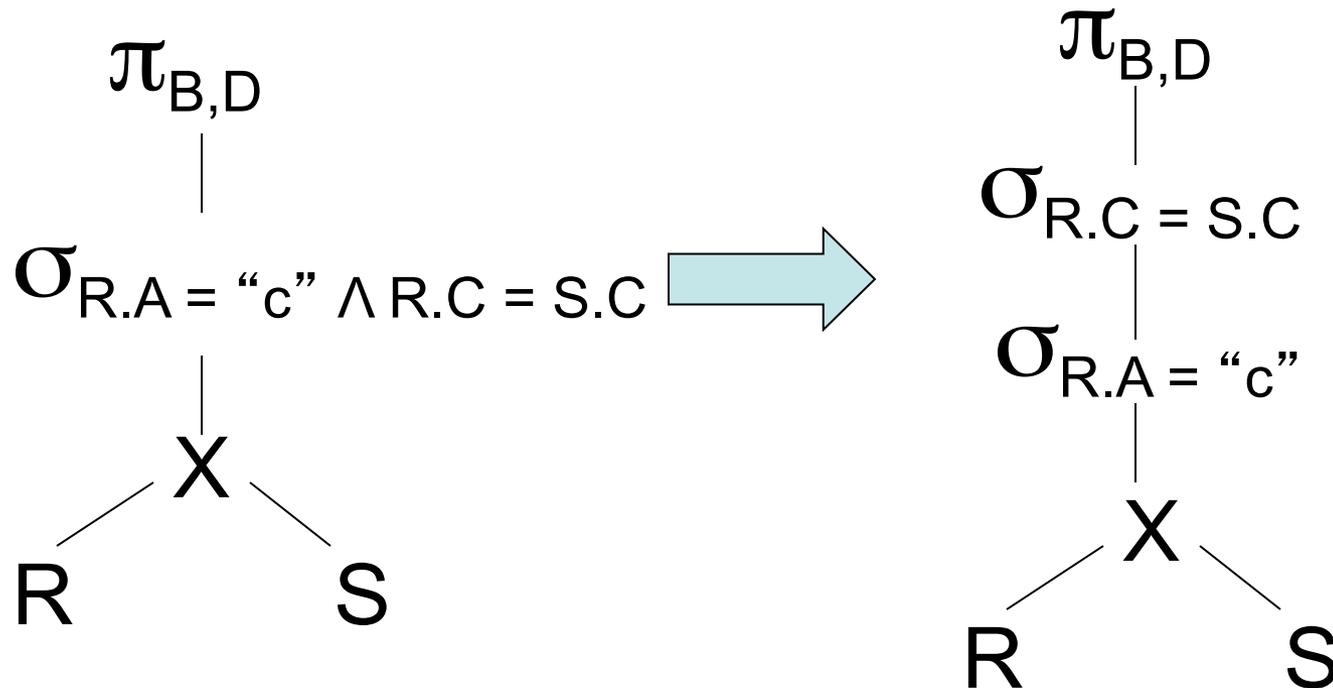
Initial Logical Plan

Select B,D
From R,S
Where R.A = "c" \wedge
R.C=S.C



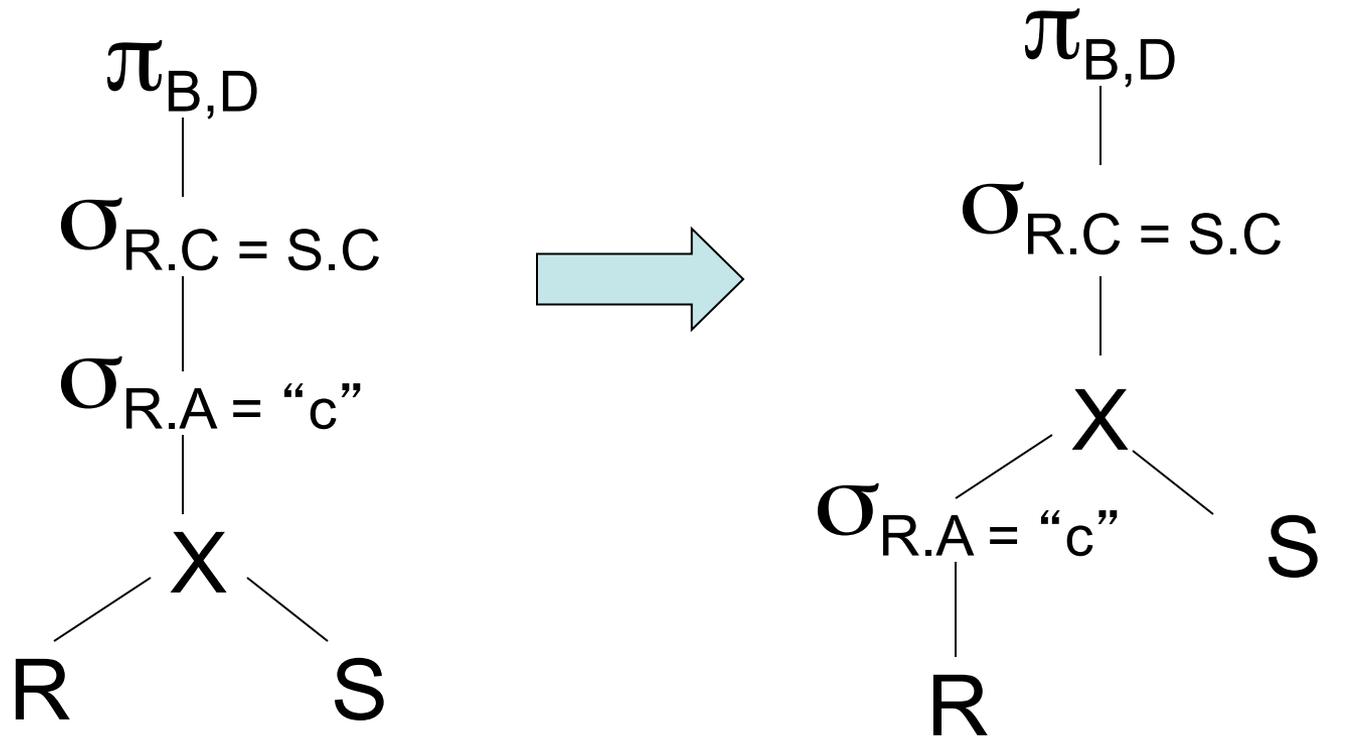
Relational Algebra: $\Pi_{B,D} [\sigma_{R.A="c" \wedge R.C = S.C} (RXS)]$

Apply Rewrite Rule (1)



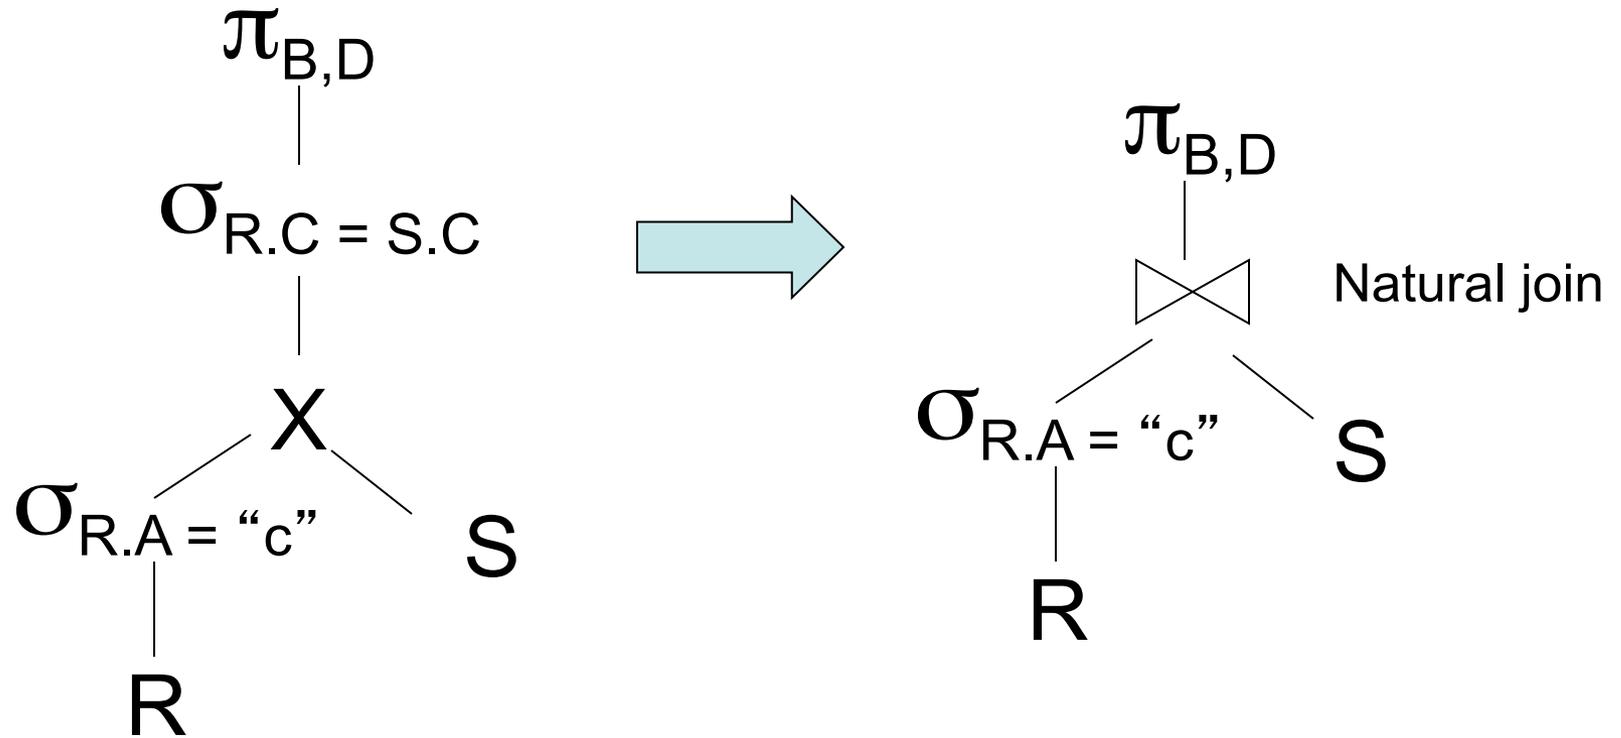
$$\Pi_{B,D} [\sigma_{R.C=S.C} [\sigma_{R.A=\text{"c"}}(R X S)]]$$

Apply Rewrite Rule (2)



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

Apply Rewrite Rule (3)



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

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

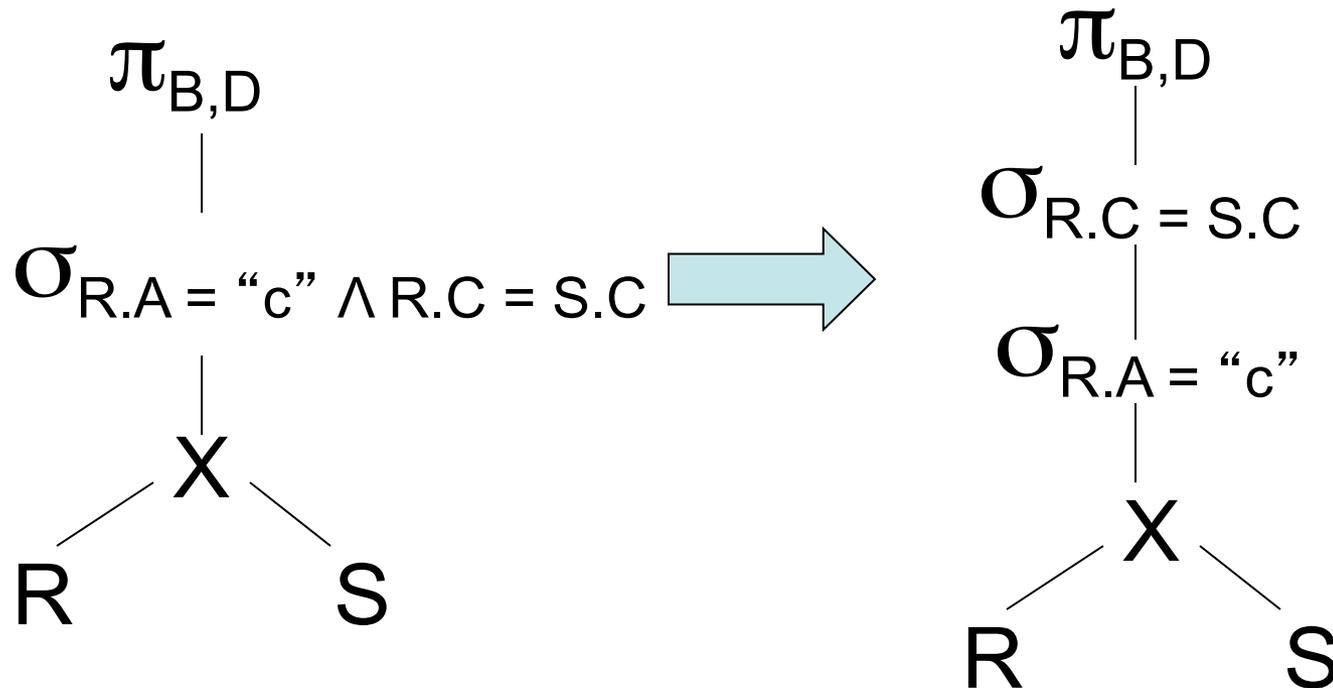
$$R \times S = S \times R$$

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

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

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

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

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

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

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?

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

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

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

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

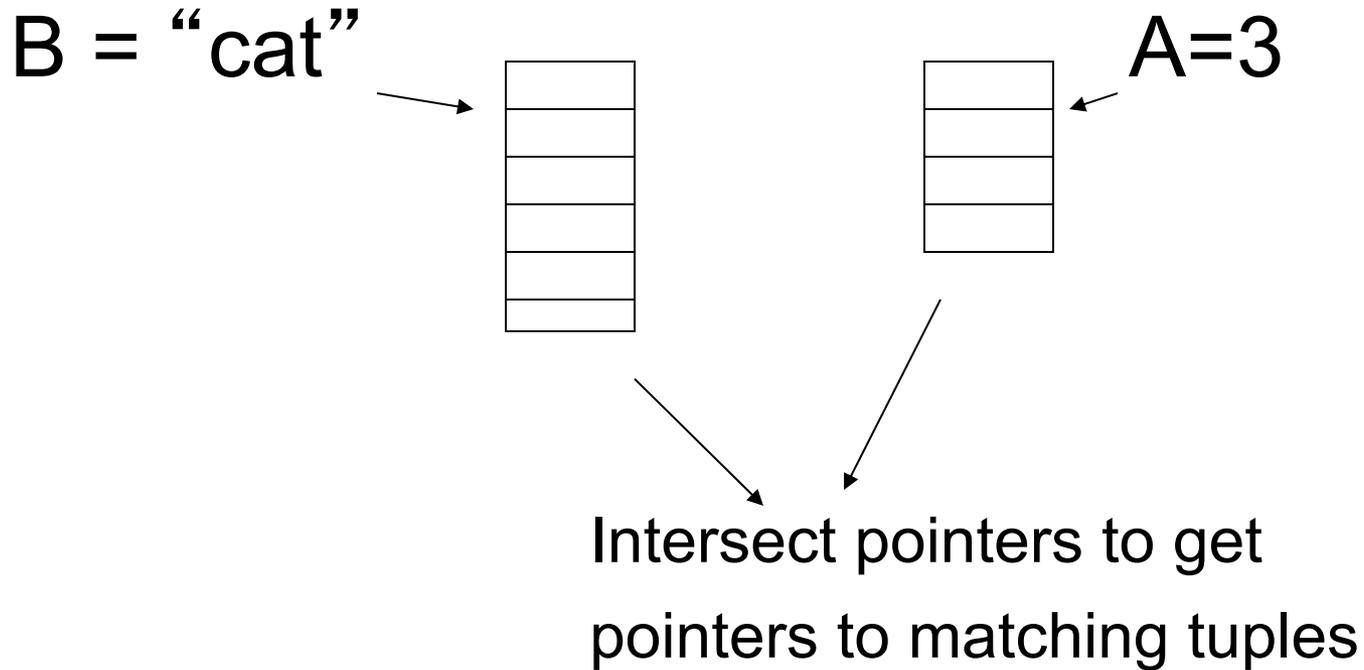
Conventional wisdom: do projects early

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

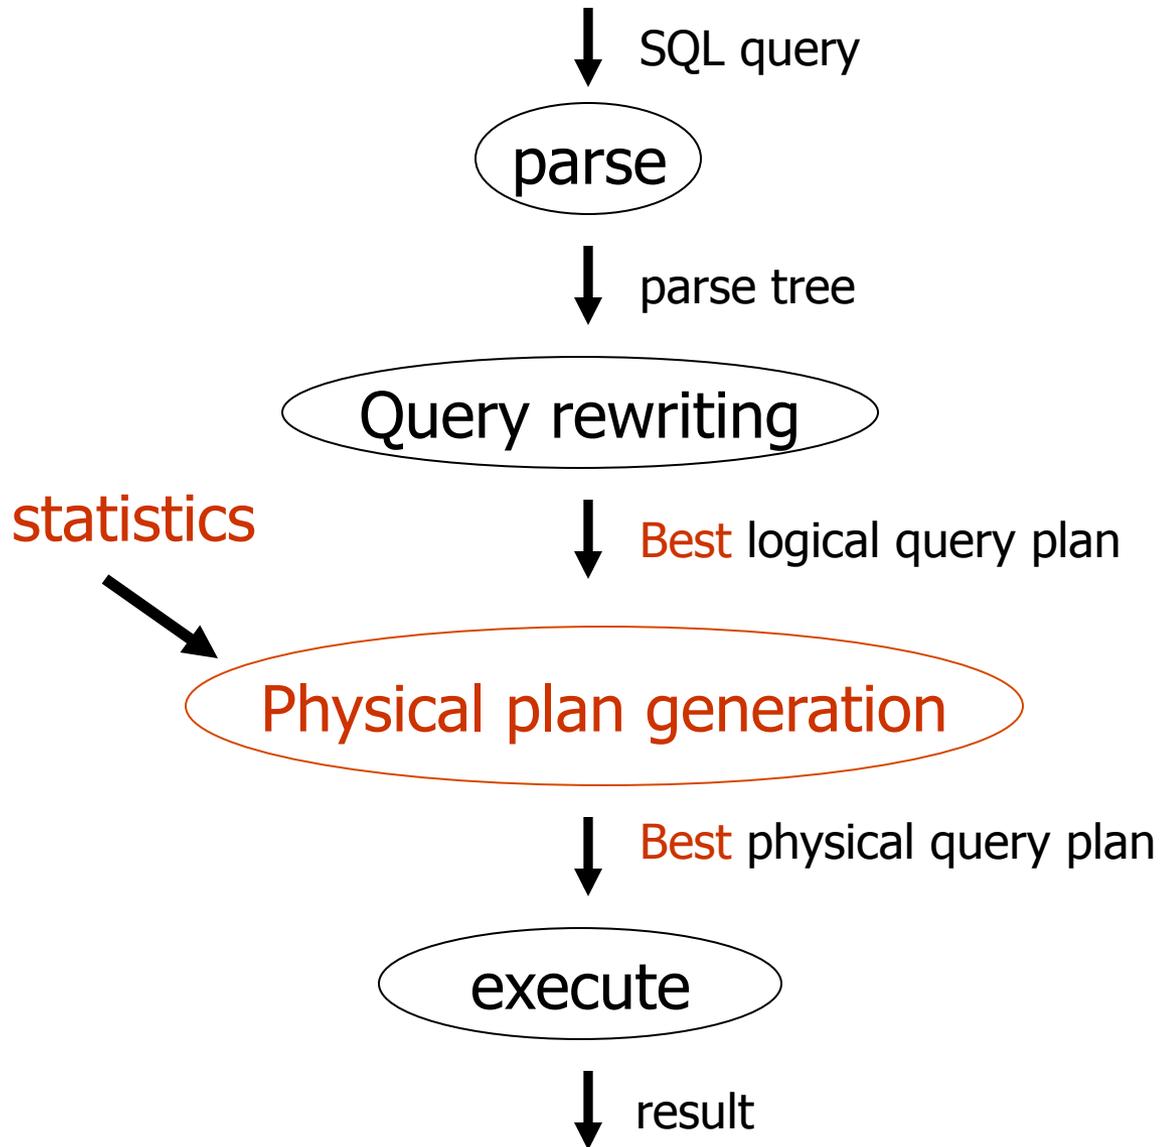
Avoid Cross Products (if possible)

Select B,D
From R,S,T,U
Where $R.A = S.B \wedge$
 $R.C = T.C \wedge R.D = U.D$

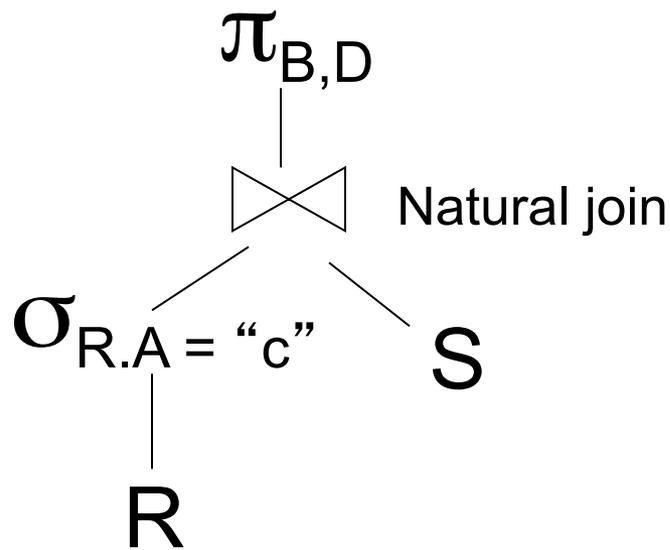
- Which join trees avoid cross-products?
- If you can't avoid cross products, perform them as late as possible

More Query Rewrite Rules

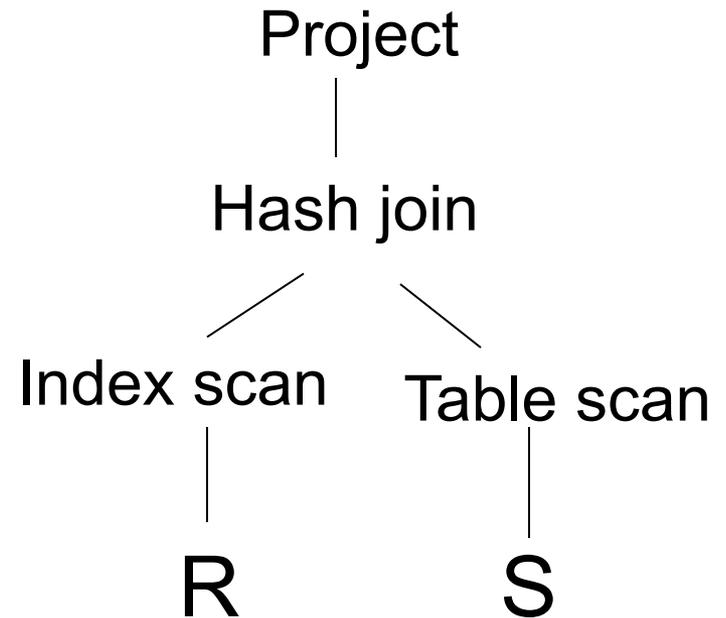
- Transform one **logical plan** into another
 - Do not use statistics
- Equivalences in relational algebra
- Push-down predicates
- Do projects early
- Avoid cross-products if possible
- **Use left-deep trees**
- **Subqueries → Joins**
- **Use of constraints, e.g., uniqueness**

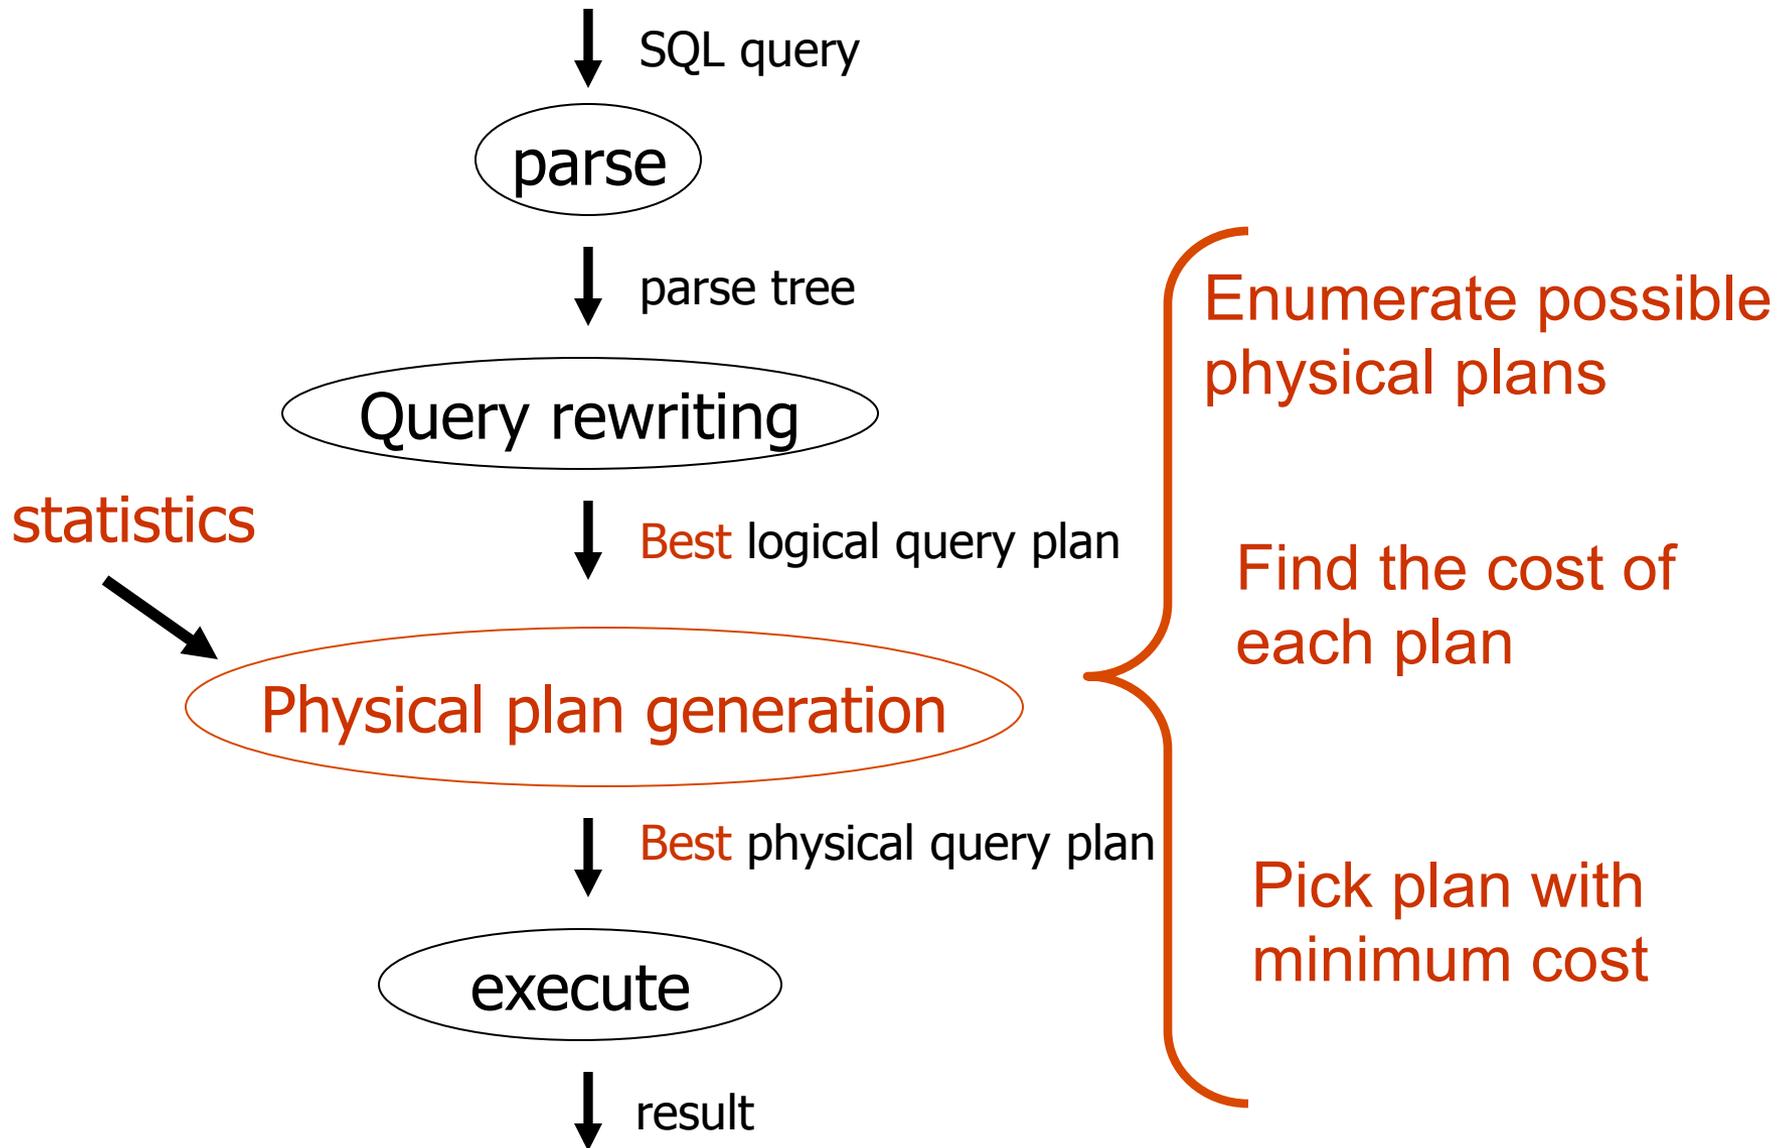


Physical Plan Generation

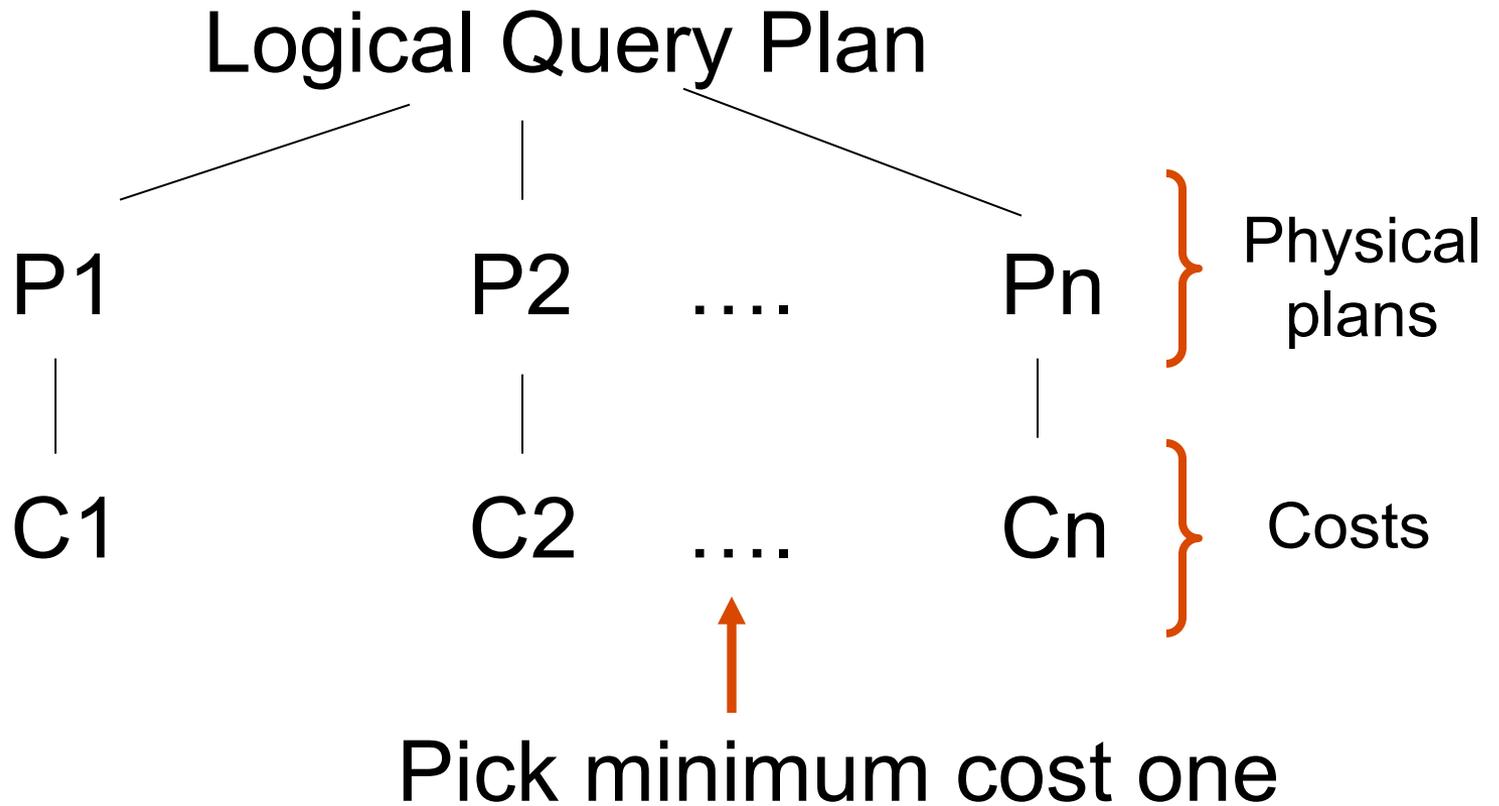


Best logical plan

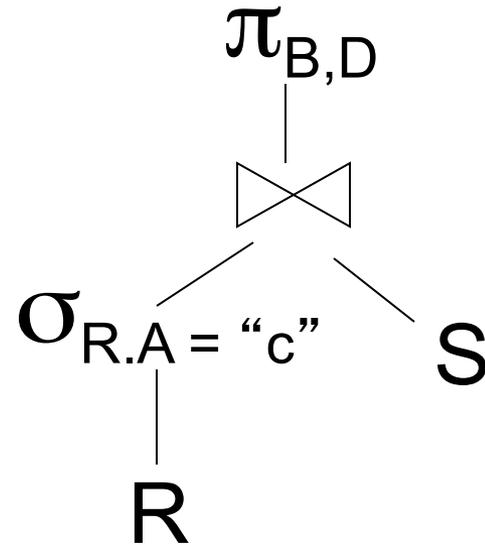




Physical Plan Generation

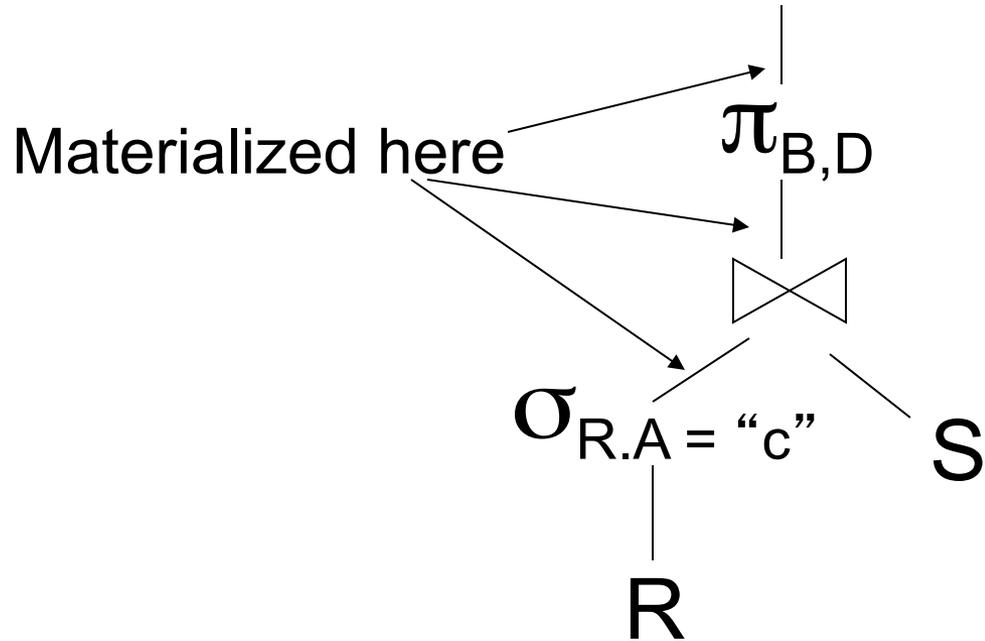


Operator Plumbing

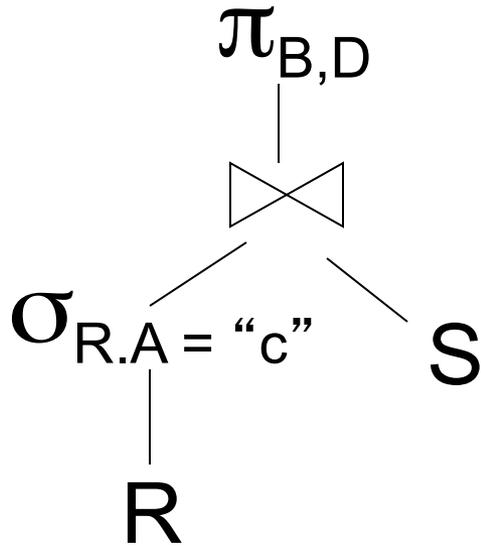


- **Materialization:** output of one operator written to disk, next operator reads from the disk
- **Pipelining:** output of one operator directly fed to next operator

Materialization



Iterators: Pipelining



→ Each operator supports:

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

Iterator for Table Scan (R)

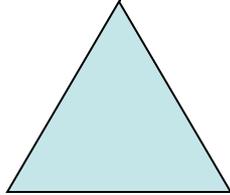
```
Open() {  
    /** initialize variables */  
    b = first block of R;  
    t = first tuple in block b;  
}
```

```
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

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

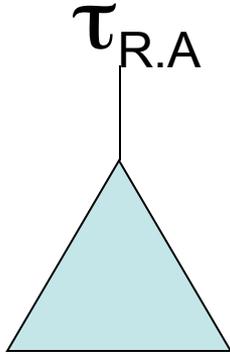


```
Open() {  
    /** initialize child */  
    Child.Open();  
}
```

```
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

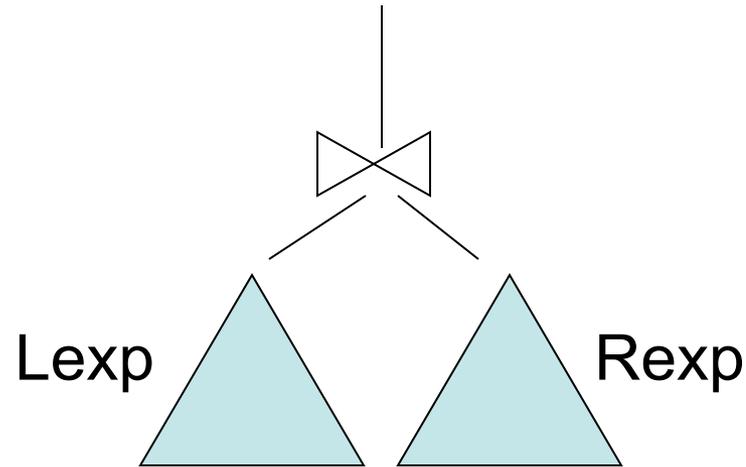


```
getNext() {  
    IF (more tuples)  
        RETURN next tuple in order;  
    ELSE RETURN EOT;  
}
```

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

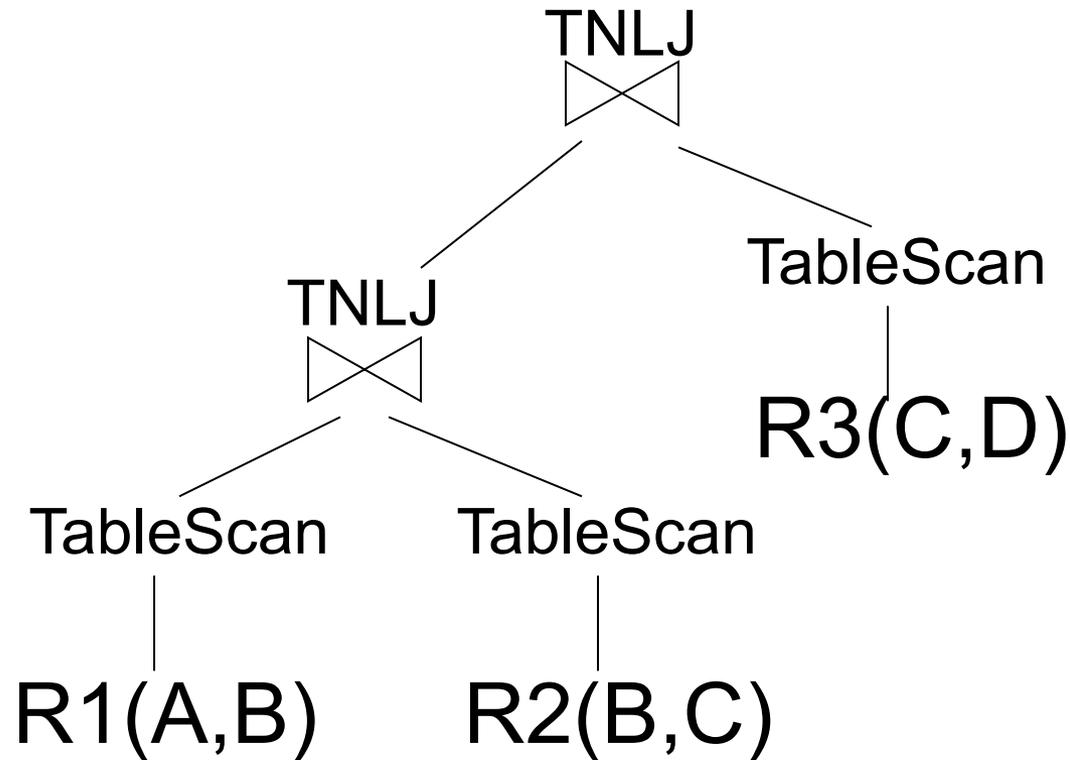
```
Close() {  
    /** inform child */  
    Child.Close();  
}
```

Iterator for Tuple Nested Loop Join



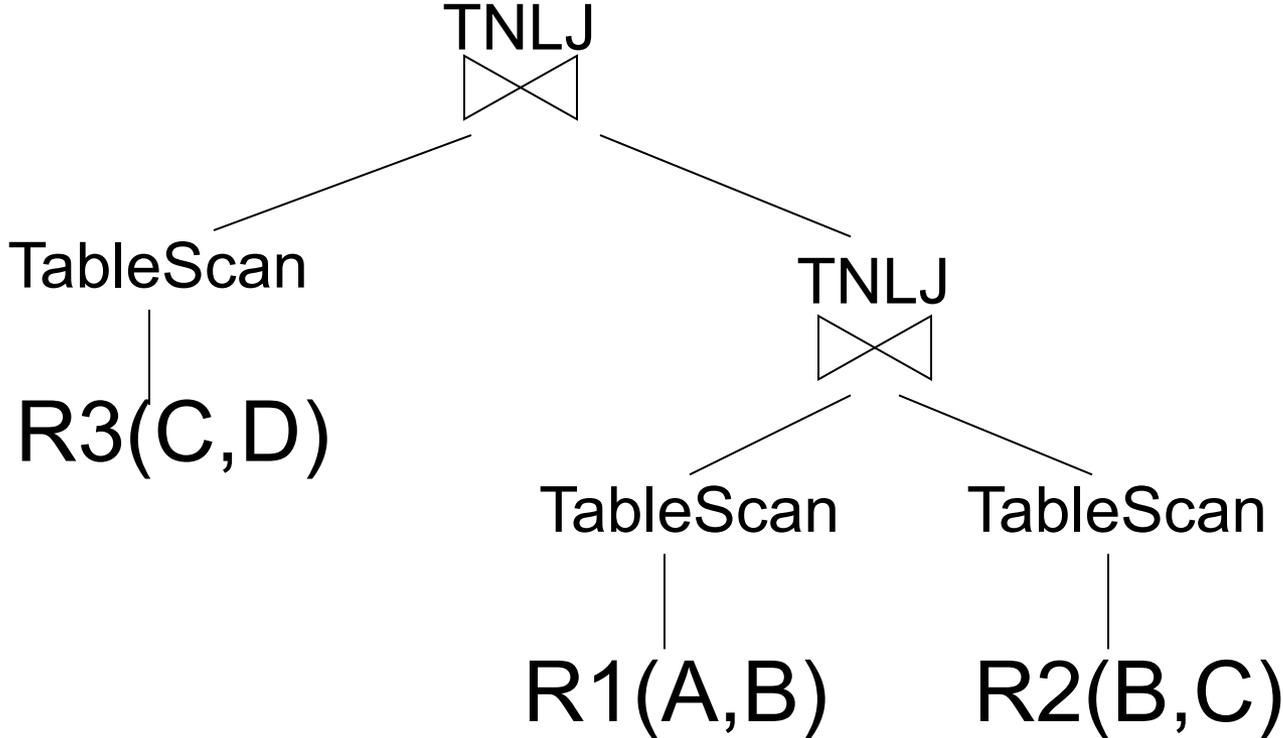
- TNLJ (conceptually)
for each $r \in L_{exp}$ do
for each $s \in R_{exp}$ do
if $L_{exp}.C = R_{exp}.C$, output r,s

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