

# Introduction to Data Management

## CSE 344

### Lecture 9: Relational Algebra and Query Evaluation

# Today

- Relational algebra
- Physical plans and query evaluation

# Relational Algebra Operators

- Union  $\cup$ , intersection  $\cap$ , difference  $-$
- Selection  $\sigma$
- Projection  $\pi$
- Cartesian product  $\times$ , join  $\bowtie$
- Rename  $\rho$

RA

- Duplicate elimination  $\delta$
- Grouping and aggregation  $\gamma$
- Sorting  $\tau$

Extended RA

All operators take in 1 or more relations as inputs and return another relation

# Join Summary

- **Theta-join:**  $R \bowtie_{\theta} S = \sigma_{\theta} (R \times S)$ 
  - Join of R and S with a join condition  $\theta$
  - Cross-product followed by selection  $\theta$
- **Equijoin:**  $R \bowtie_{\theta} S = \pi_A (\sigma_{\theta} (R \times S))$ 
  - Join condition  $\theta$  consists only of equalities
  - Projection  $\pi_A$  drops all redundant attributes
- **Natural join:**  $R \bowtie S = \pi_A (\sigma_{\theta} (R \times S))$ 
  - Equijoin
  - Equality on **all** fields with same name in R and in S
  - Projection  $\pi_A$  drops all redundant attributes

# So Which Join Is It ?

When we write  $R \bowtie S$  we usually mean an equijoin, but we often omit the equality predicate when it is clear from the context

# More Joins

- **Outer join**
  - Include tuples with no matches in the output
  - Use NULL values for missing attributes
  - Does not eliminate duplicate columns
- Variants
  - Left outer join
  - Right outer join
  - Full outer join

# Outer Join Example

AnonPatient P

age	zip	disease
54	98125	heart
20	98120	flu
33	98120	lung

AnnonJob J

job	age	zip
lawyer	54	98125
cashier	20	98120

P  $\bowtie$  J

P.age	P.zip	disease	job	J.age	J.zip
54	98125	heart	lawyer	54	98125
20	98120	flu	cashier	20	98120
33	98120	lung	null	null	null

$\bowtie$  LOJ  
 $\ltimes$  ROJ  
 $\ltimes$  FOJ

# Some Examples

Supplier(sno, sname, scity, sstate)

Part(pno, pname, psize, pcolor)

Supply(sno, pno, qty, price)

Name of supplier of parts with size greater than 10

$\pi_{\text{sname}}(\text{Supplier} \bowtie (\text{Supply} \bowtie (\sigma_{\text{psize}>10}(\text{Part}))))$

2. ← execution order  
1. ←

Name of supplier of red parts or parts with size greater than 10

$\pi_{\text{sname}}(\text{Supplier} \bowtie (\text{Supply} \bowtie (\sigma_{\text{psize}>10}(\text{Part}) \cup \sigma_{\text{pcolor}='red'}(\text{Part}))))$

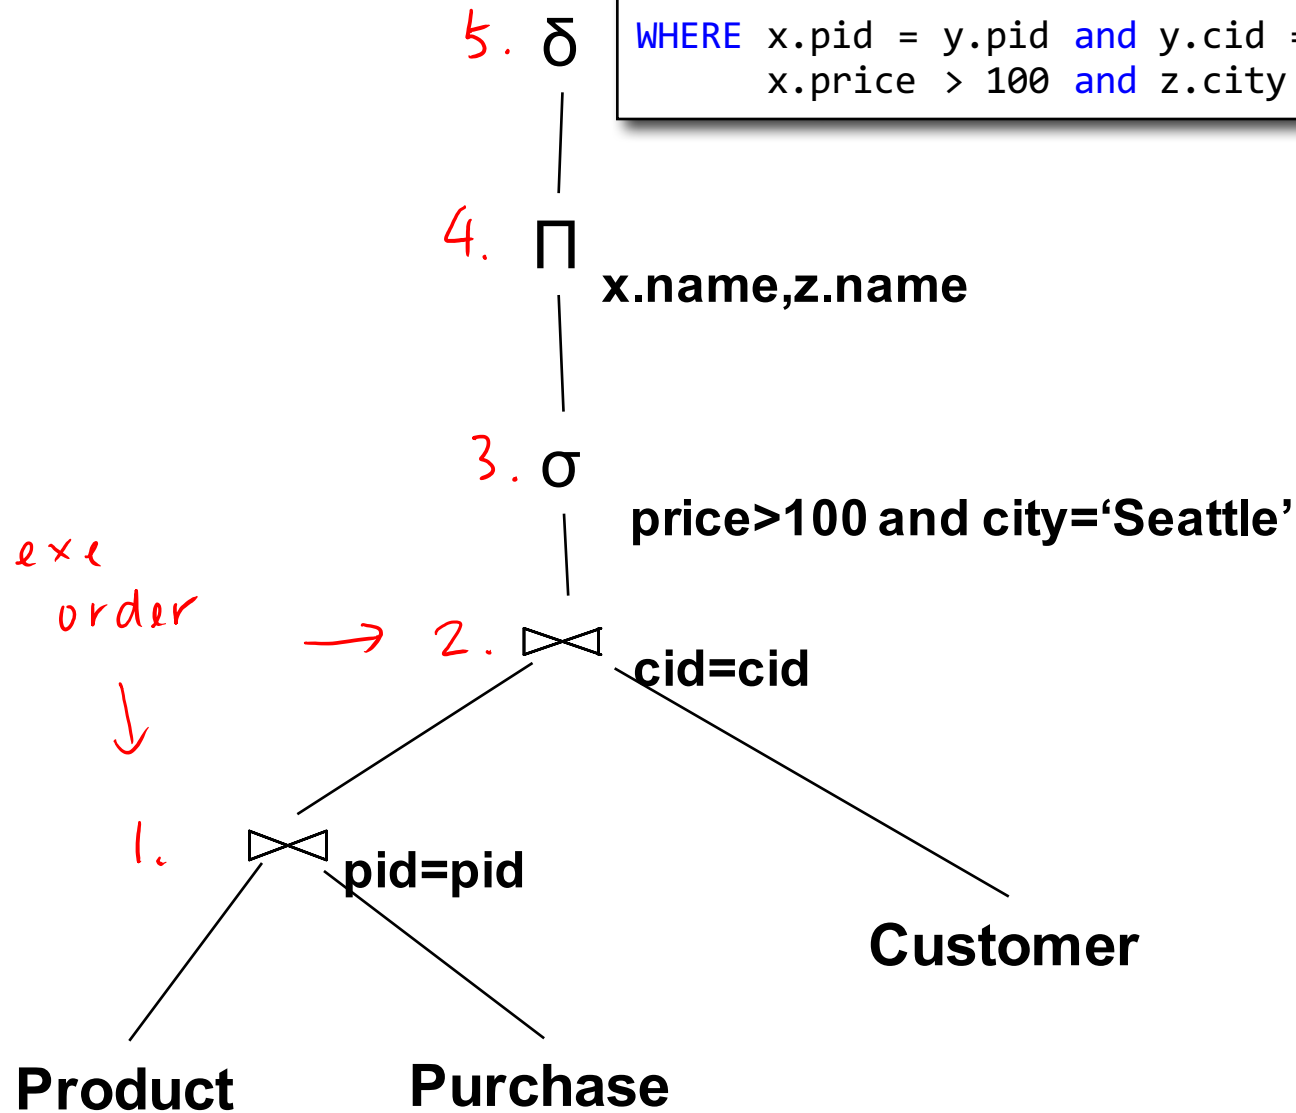
Can be represented as trees as well (as seen from lecture 7)



Product(pid, name, price)  
Purchase(pid, cid, store)  
Customer(cid, name, city)

# From SQL to RA

```
SELECT DISTINCT x.name, z.name  
FROM Product x, Purchase y, Customer z  
WHERE x.pid = y.pid and y.cid = z.cid and  
x.price > 100 and z.city = 'Seattle'
```




Product(pid, name, price)  
Purchase(pid, cid, store)  
Customer(cid, name, city)

# From SQL to RA

Can you think of a "better" plan?

```
SELECT DISTINCT x.name, z.name  
FROM Product x, Purchase y, Customer z  
WHERE x.pid = y.pid and y.cid = z.cid and  
x.price > 100 and z.city = 'Seattle'
```



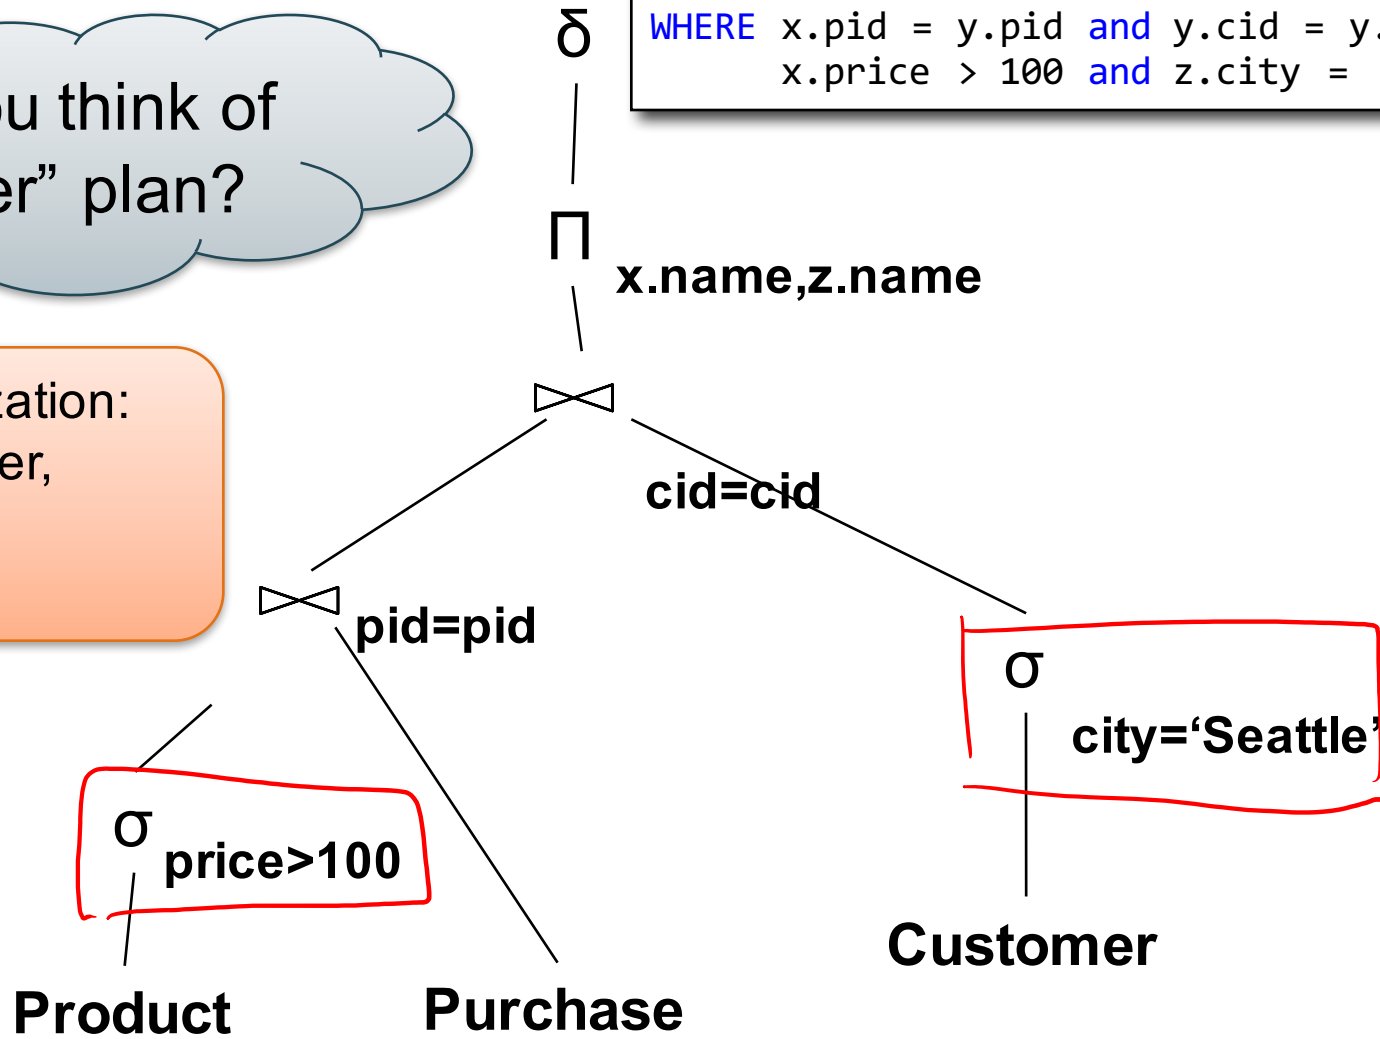
Product(pid, name, price)  
Purchase(pid, cid, store)  
Customer(cid, name, city)

# Equivalent Expression

Can you think of a "better" plan?

Query optimization:  
finding cheaper,  
equivalent  
expressions

```
SELECT DISTINCT x.name, z.name  
FROM Product x, Purchase y, Customer z  
WHERE x.pid = y.pid and y.cid = y.cid and  
x.price > 100 and z.city = 'Seattle'
```



# Extended RA: Operators on Bags

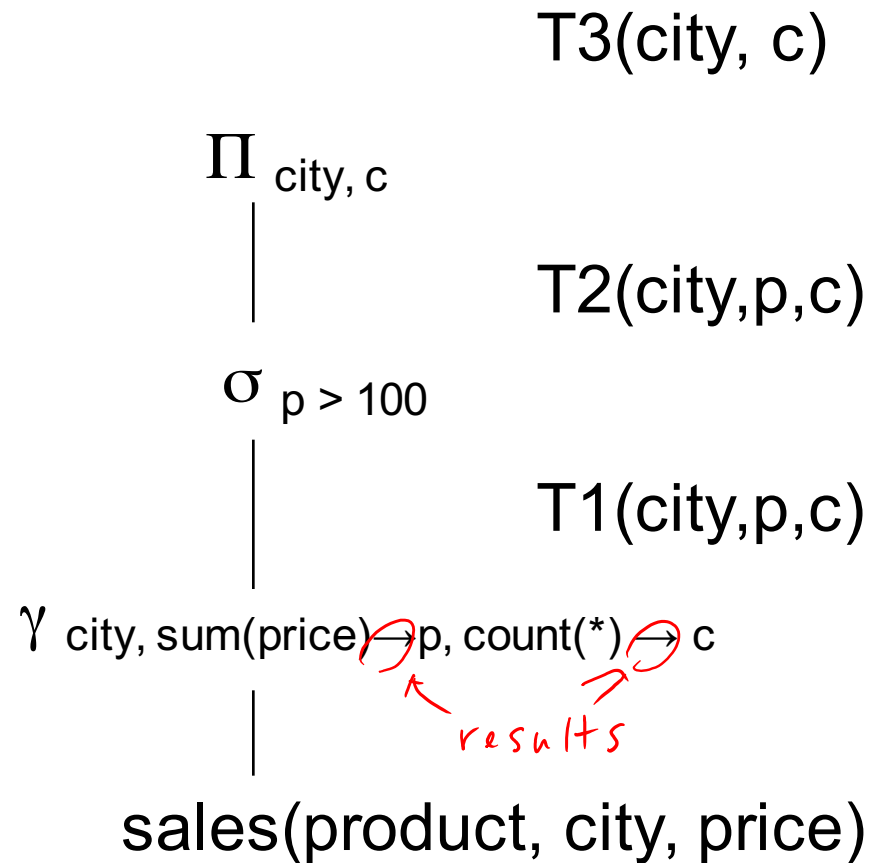
- Duplicate elimination  $\delta$
- Grouping  $\gamma$ 
  - Takes in relation and a list of grouping operations (e.g., aggregates). Returns a new relation.
- Sorting  $\tau$ 
  - Takes in a relation, a list of attributes to sort on, and an order. Returns a new relation.

# Using Extended RA Operators

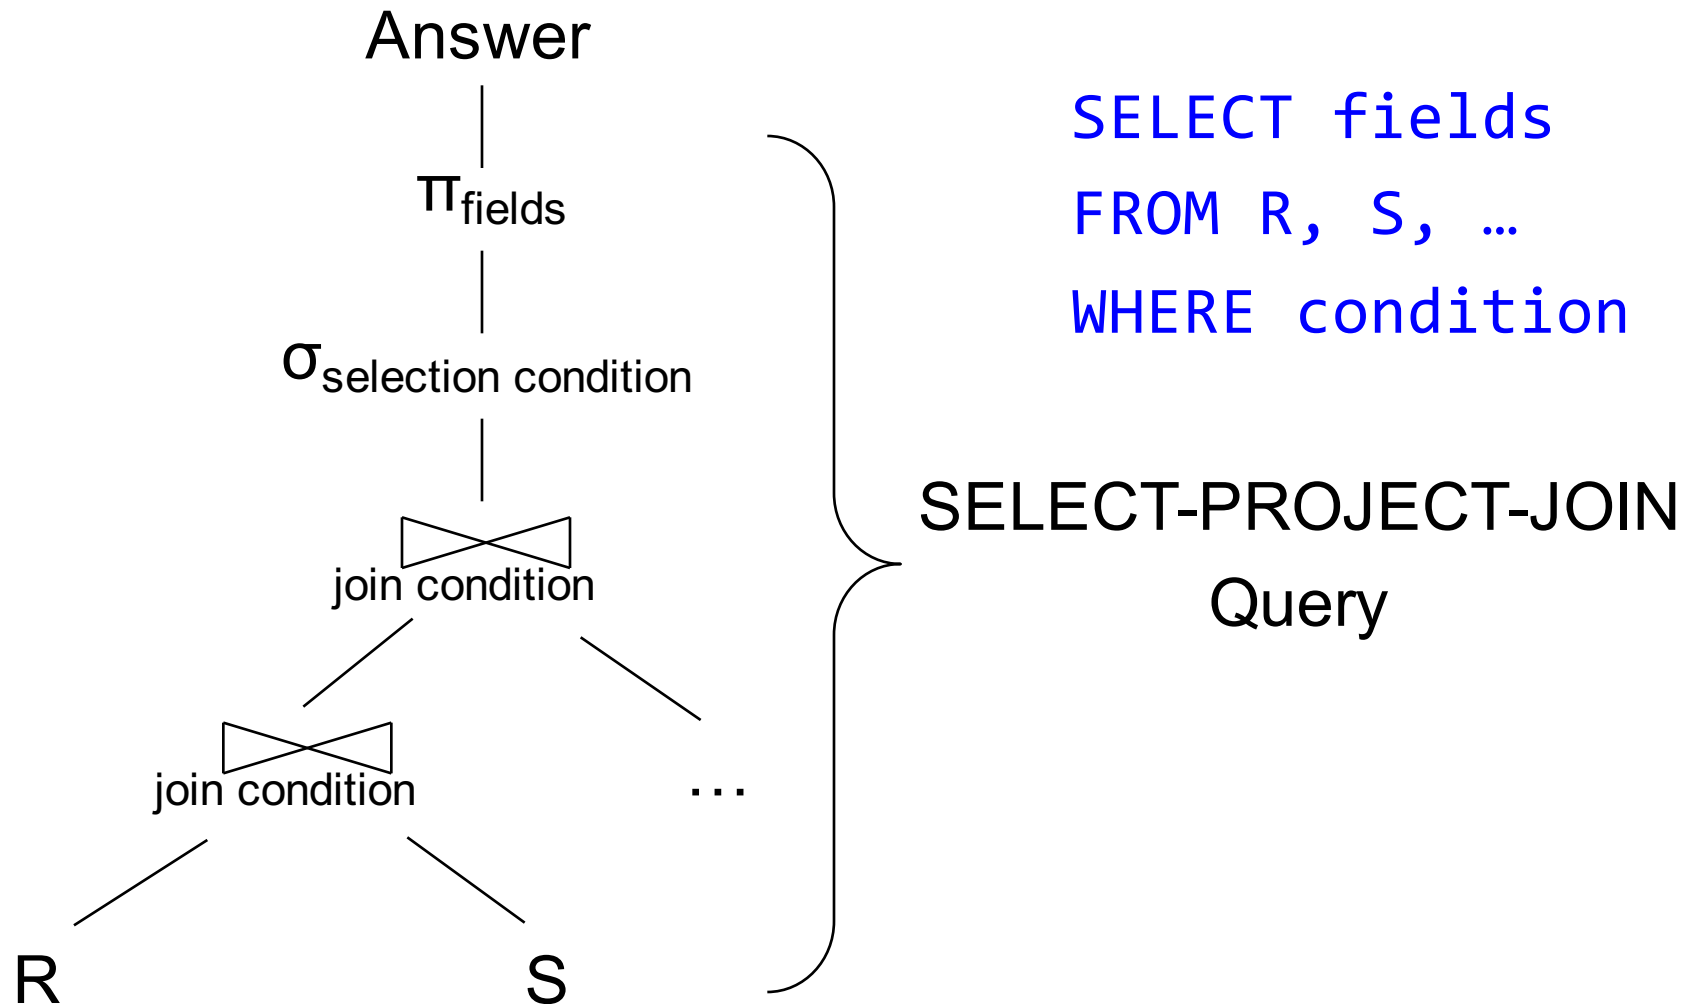
```

SELECT city, count(*)
FROM sales
GROUP BY city
HAVING sum(price) > 100
    
```

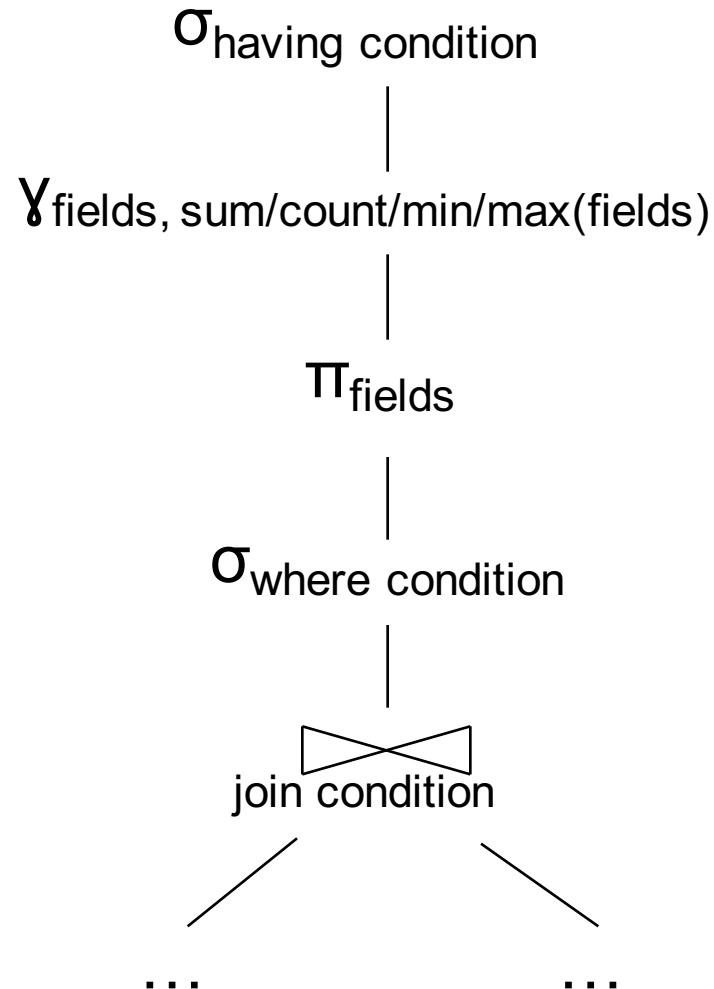
T1, T2, T3 = temporary tables



# Typical Plan for a Query (1/2)



# Typical Plan for a Query (1/2)



SELECT fields  
FROM R, S, ...  
WHERE condition  
GROUP BY fields  
HAVING condition

Supplier(sno,sname,scity,sstate)  
Part(pno,pname,psize,pcolor)  
Supply(sno,pno,price)

# How about Subqueries?

```
SELECT  Q.sno
FROM    Supplier Q
WHERE   Q.sstate = 'WA'
       and not exists
       (SELECT *
        FROM Supply P
        WHERE P.sno = Q.sno
              and P.price > 100)
```



Supplier(sno,sname,scity,sstate)  
Part(pno,pname,psize,pcolor)  
Supply(sno,pno,price)

# How about Subqueries?

```
SELECT Q.sno  
FROM Supplier Q  
WHERE Q.sstate = 'WA'  
and not exists  
(SELECT *  
FROM Supply P  
WHERE P.sno = Q.sno  
and P.price > 100)
```

Correlation !

Supplier(sno,sname,scity,sstate)  
Part(pno,pname,psize,pcolor)  
Supply(sno,pno,price)

# How about Subqueries?

```
SELECT Q.sno
FROM Supplier Q
WHERE Q.sstate = 'WA'
and not exists
(SELECT *
FROM Supply P
WHERE P.sno = Q.sno
and P.price > 100)
```

De-Correlation

```
SELECT Q.sno
FROM Supplier Q
WHERE Q.sstate = 'WA'
and Q.sno not in
(SELECT P.sno
FROM Supply P
WHERE P.price > 100)
```

Supplier(sno,sname,scity,sstate)  
Part(pno,pname,psize,pcolor)  
Supply(sno,pno,price)

# How about Subqueries?

Un-nesting

```
(SELECT Q.sno  
FROM Supplier Q  
WHERE Q.sstate = 'WA')  
EXCEPT  
(SELECT P.sno  
FROM Supply P  
WHERE P.price > 100)
```

EXCEPT = set difference

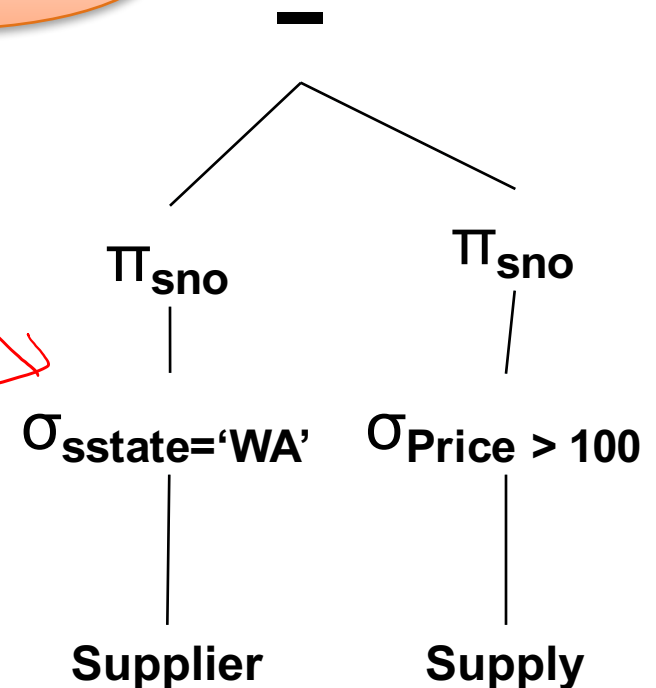
```
SELECT Q.sno  
FROM Supplier Q  
WHERE Q.sstate = 'WA'  
and Q.sno not in  
(SELECT P.sno  
FROM Supply P  
WHERE P.price > 100)
```

Supplier(sno, sname, scity, sstate)  
Part(pno, pname, psize, pcolor)  
Supply(sno, pno, price)

# How about Subqueries?


```
(SELECT Q.sno  
FROM Supplier Q  
WHERE Q.sstate = 'WA')  
EXCEPT  
(SELECT P.sno  
FROM Supply P  
WHERE P.price > 100)
```

Finally...



# From Logical RA Plans to Physical Plans

# Query Evaluation Steps Review




Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

# Relational Algebra

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```

Relational algebra expression is also called the “logical query plan”



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

# Physical Query Plan 1

(On the fly)

$\Pi_{\text{sname}}$

(On the fly)

$\sigma_{\text{scity}='Seattle' \text{ and } \text{sstate}='WA' \text{ and } \text{pno}=2}$

(Nested loop)

sid = sid

Supplier  
(File scan)

Supply  
(File scan)

A physical query plan is a logical query plan annotated with physical implementation details

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

# Physical Query Plan 2

(On the fly)

$\Pi_{\text{sname}}$

(On the fly)

$\sigma_{\text{scity}='Seattle' \text{ and } \text{sstate}='WA' \text{ and } \text{pno}=2}$

(Hash join)

sid = sid

Same logical query plan  
Different physical plan

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```

Supplier  
(File scan)

Supply  
(File scan)

Supplier(sid, sname, scity, sstate)


Supply(sid, pno, quantity)

# Physical Query Plan 3

(On the fly)

$\Pi_{\text{sname}}$  (d)

(Sort-merge join)



(Scan & write to T1)

(a)  $\sigma_{\text{scity}='Seattle' \text{ and } \text{sstate}='WA'}$

Supplier  
(File scan)

(Scan & write to T2)

(b)  $\sigma_{\text{pno}=2}$

Supply  
(File scan)

Different but equivalent logical query plan; different physical plan

```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```

# Query Optimization Problem


- For each SQL query... many logical plans
- For each logical plan... many physical plans
- How do find a fast physical plan?
  - Will discuss in a few lectures
  - First we need to understand how query operators are implemented

# Query Execution

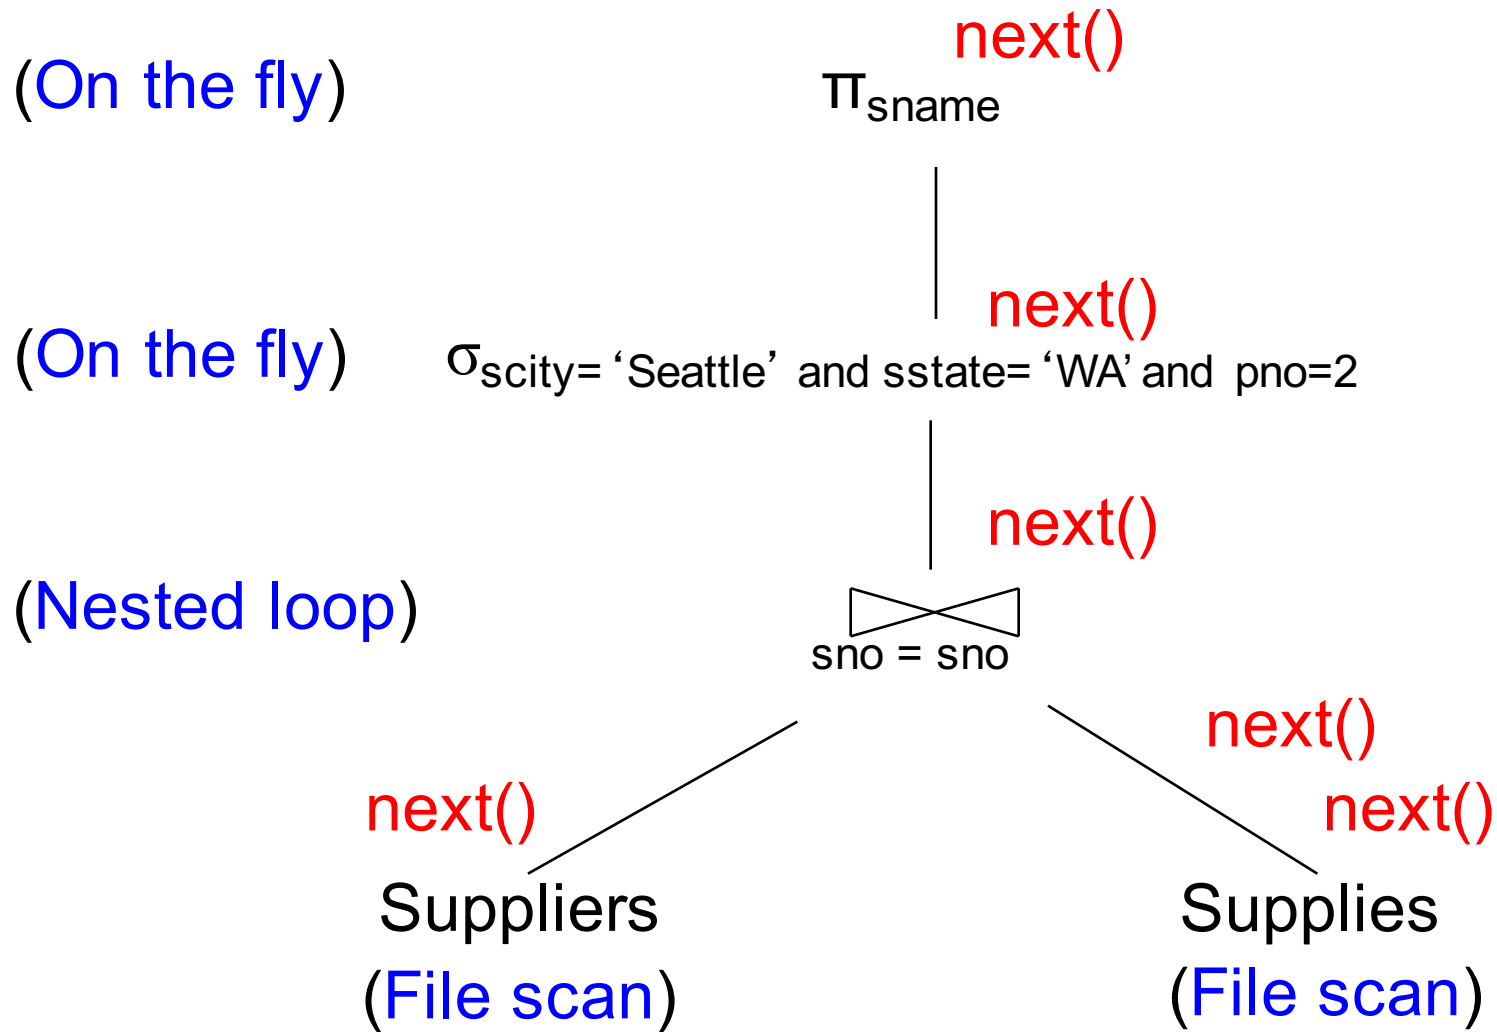
# Iterator Interface for Query Operators

- **open()**
  - Initializes operator state
  - Sets parameters such as selection condition
- **next()**
  - Operator invokes `get_next()` recursively on its inputs
  - Performs processing and produces an output tuple
- **close()**: clean-up state

# Pipelined Query Execution



# Pipelined Query Execution



# Pipelined Execution

- **Tuples generated by an operator are immediately sent to the parent**
- **Benefits:**
  - No operator synchronization issues
  - No need to buffer tuples between operators
  - Saves cost of writing intermediate data to disk
  - Saves cost of reading intermediate data from disk
- **This approach is used whenever possible**



# Query Execution Bottom Line

- SQL query transformed into **physical plan**
  - **Access path selection** for each relation
    - Scan the relation or use an index (next lecture)
  - **Implementation choice** for each operator
    - Nested loop join, hash join, etc.
  - **Scheduling decisions** for operators
    - Pipelined execution or intermediate materialization
- Pipelined execution of physical plan

# Physical Data Independence

- Applications are insulated from changes in physical storage details
- SQL and relational algebra facilitate physical data independence
  - Both languages input and output relations
  - Can choose different implementations for operators