

Introduction to Database Systems

CSE 414

Lecture 28: Intro to Query Optimization

Final Exam

- Thursday 6/7, 2:30-4:20pm
- Location: here
- Comprehensive exam
 - Covers all lectures, sections, web quizzes, HWs, and readings
- Can bring 2 letter-size sheets of notes
 - Handwritten or printed
- More info on course website
- Review session:
 - Sunday 6/3, 2:30-5pm, SMI 102

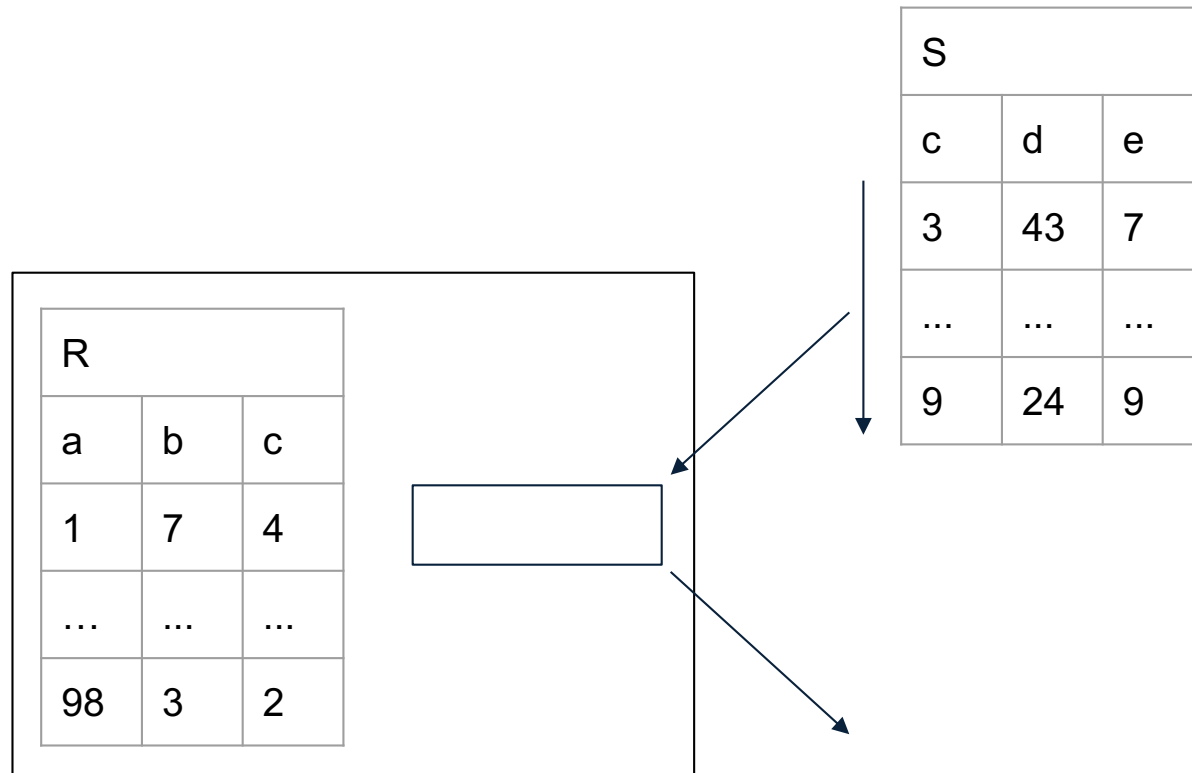
Big Picture

- How to choose the “best” query plan to run? (aka query optimization)
- To answer this question we need to understand:
 - Data organization on the disk
 - Index structures and how they are used in queries
 - A way to model query “costs”
 - Compute cost for each query operator
 - Compute cost for each physical plan

Last topics
this quarter!

Review: Join Algorithms

- Nested loop join
- Hash join
- Sort-merge join



Hash Join

Hash Join

Hash join: $R \bowtie S$

- Scan R, build hash table in main memory
 - Then scan S and join
 - Cost: $B(R) + B(S)$
 - Which relation to build the hash table on?
-
- One-pass algorithm when $B(R) \leq M$
 - M = number of memory pages available

Hash Join Example

Patient(pid, name, address)

Insurance(pid, provider, policy_nb)

Patient ⋈ Insurance

Patient

1	'Bob'	'Seattle'
2	'Ela'	'Everett'

3	'Jill'	'Kent'
4	'Joe'	'Seattle'

Insurance

2	'Blue'	123
4	'Prem'	432

4	'Prem'	343
3	'GrpH'	554

Two tuples
per page

Hash Join Example


Patient \bowtie Insurance

Some large-enough #

Memory M = 21 pages

Showing
pid only

Patient		Insurance	
1	2	2	4
3	4	4	3
9	6	2	8
8	5	8	9




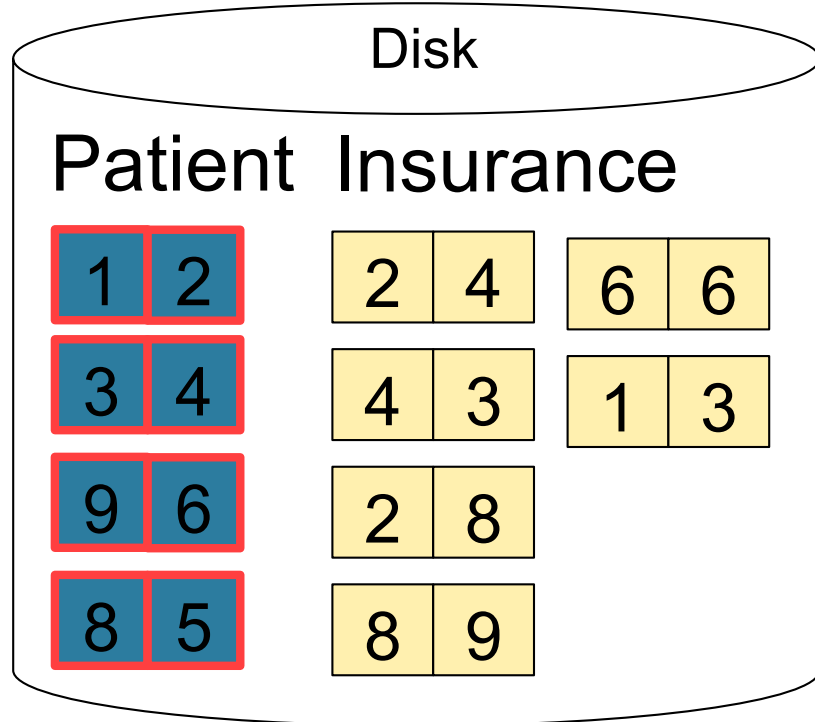

This is one page
with two tuples

Hash Join Example

Step 1: Scan Patient and **build** hash table in memory

Memory M = 21 pages

Hash h: pid % 5



Hash Join Example

Step 2: Scan Insurance and **probe** into hash table

Memory M = 21 pages

Hash h: pid % 5

= 0	= 1	= 2	= 3	= 4
5	1 6	2	3 8	4 9

Patient		Insurance	
1	2	2	4
3	4	4	3
9	6	2	8
8	5	8	9

2	4
---	---

Input buffer

2	2
---	---

Output buffer

Write to disk or
pass to next
operator

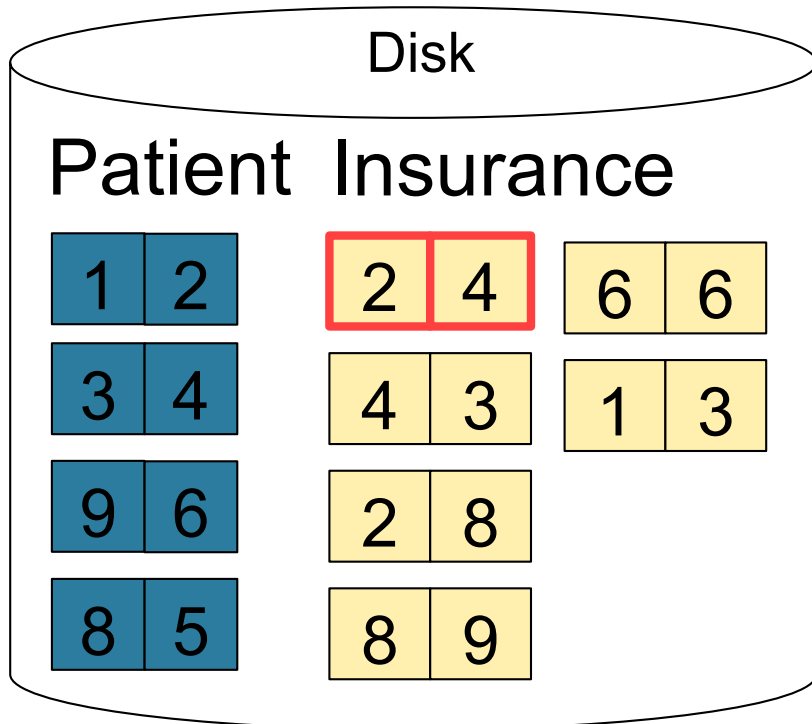
Hash Join Example

Step 2: Scan Insurance and **probe** into hash table


Memory M = 21 pages

Hash h: pid % 5

= 0	= 1	= 2	= 3	= 4
5	1 6	2	3 8	4 9



Input buffer



Output buffer

Hash Join Example

Step 2: Scan Insurance and **probe** into hash table

Memory M = 21 pages

Hash h: pid % 5

= 0	= 1	= 2	= 3	= 4
5	1 6	2	3 8	4 9

Disk			
Patient		Insurance	
1	2	2	4
3	4	4	3
9	6	2	8
8	5	8	9

4	3
---	---

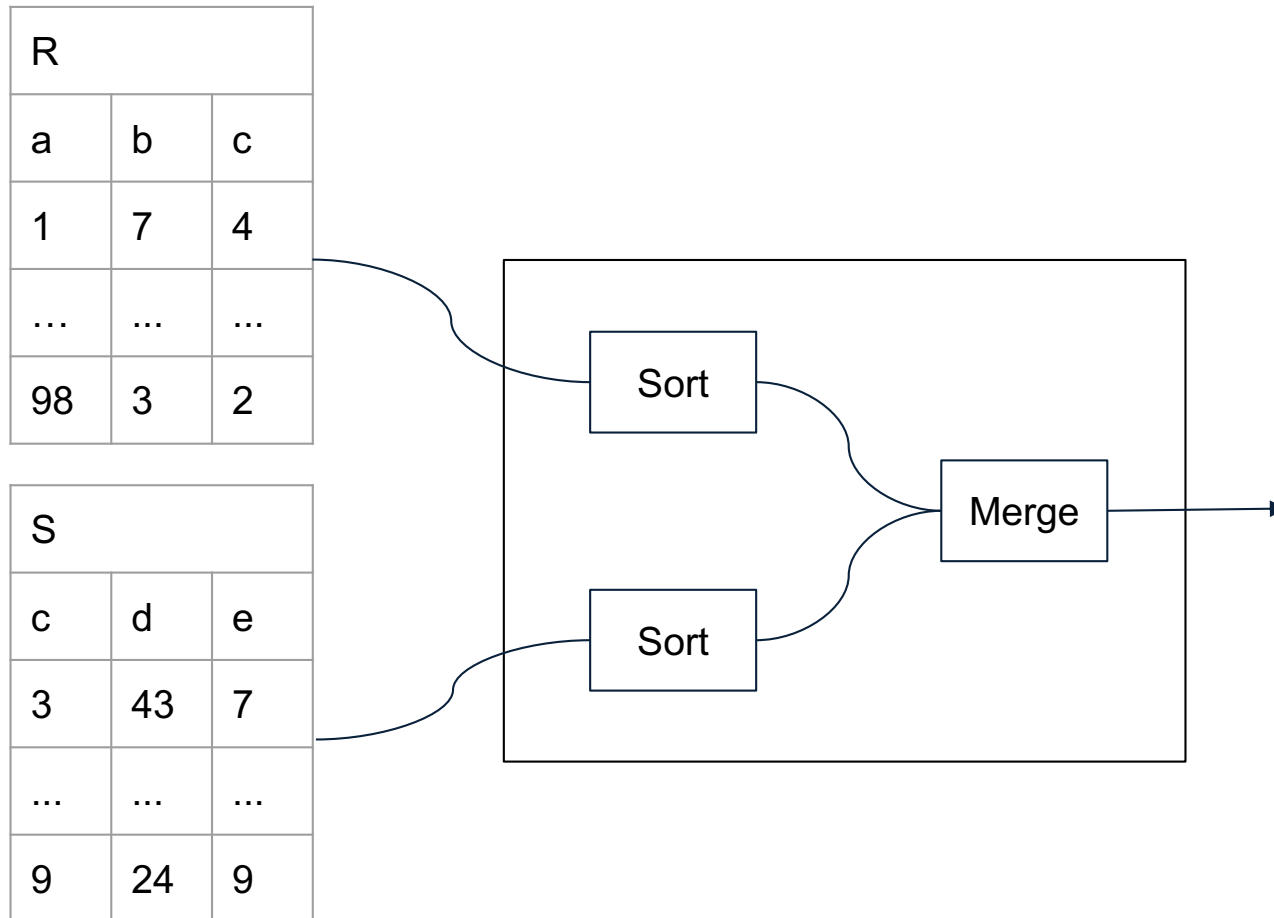
Input buffer

4	4
---	---

Output buffer

Keep going until read all of Insurance

Cost: $B(R) + B(S)$



Sort-Merge Join

Sort-Merge Join

Sort-merge join: $R \bowtie S$

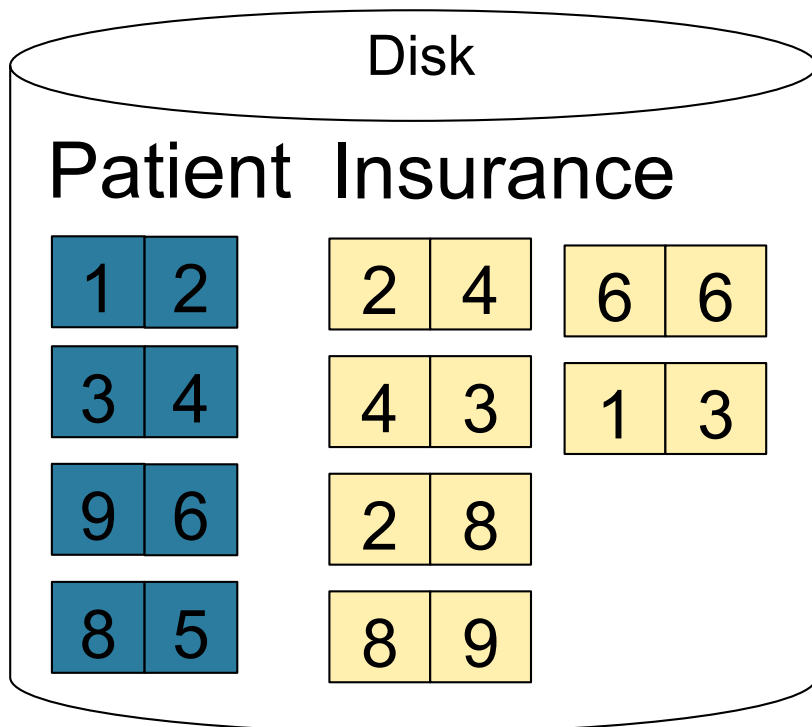
- Scan R and sort in main memory
 - Scan S and sort in main memory
 - Merge R and S
-
- Cost: $B(R) + B(S)$
 - One pass algorithm when $B(S) + B(R) \leq M$
 - Typically, this is NOT a one pass algorithm

Sort-Merge Join Example

Step 1: Scan Patient and **sort** in memory

Memory M = 21 pages

1	2	3	4	5	6	8	9
---	---	---	---	---	---	---	---





Sort-Merge Join Example

Step 2: Scan Insurance and **sort** in memory

Memory M = 21 pages



sorted



Sort-Merge Join Example

Step 3: **Merge** Patient and Insurance



Memory M = 21 pages




Sort-Merge Join Example

Step 3: Merge Patient and Insurance

Memory M = 21 pages




R		
a	b	c
1	7	4
...
98	3	2



S		
c	d	e
c	d	e
3	43	7
...
9	24	9

■ ■ ■



S		
c	d	e
c	d	e
3	43	7
...
9	24	9

Index Joins

Index Nested Loop Join

$R \bowtie S$

- Assume S has an index on the join attribute
- Iterate over R, for each tuple fetch corresponding tuple(s) from S

```
for r in R
    // use index to lookup
    for s' in S that should be joined with r
        s = fetch S tuple pointed to by s' from disk
    output (r,s)
```


Index Nested Loop Join

$R \bowtie S$

```
for r in R
    // use index to lookup
    for s' in S that should be joined with r
        s = fetch S tuple pointed to by s' from disk
        output (r,s)
```

- **Cost:**

- If index on S is clustered: $B(R) + T(R) * (B(S) * 1/V(S,a))$
- If index on S is unclustered: $B(R) + T(R) * (T(S) * 1/V(S,a))$

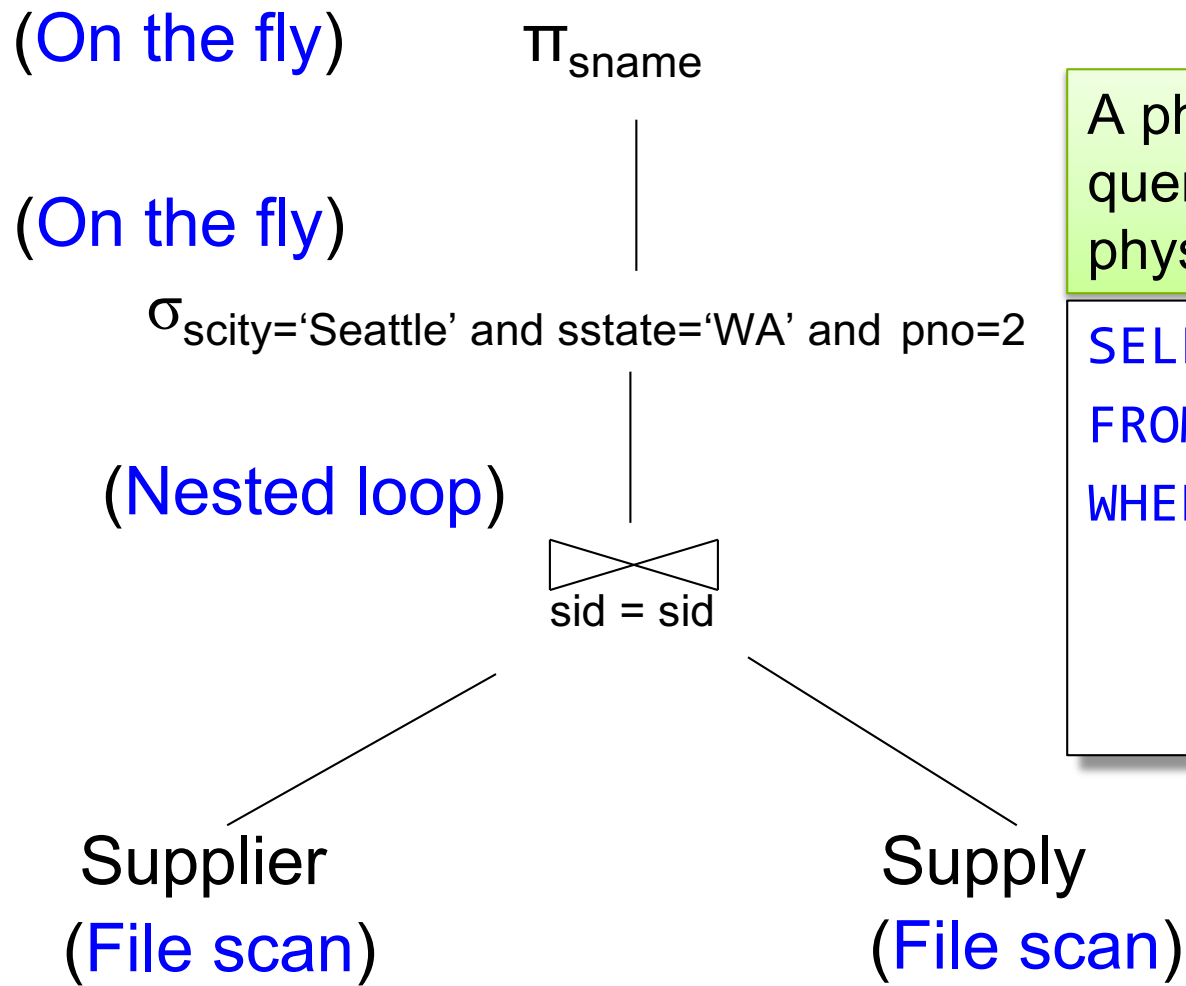


Review:

Logical vs Physical Plans

Supplier(sid, sname, scity, sstate)
Supply(sid, pno, quantity)

Review: Physical Query Plan 1



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)

Review: Physical Query Plan 2

(On the fly)

π_{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)

Review: Physical Query Plan 3

(On the fly)

π_{sname} (d)

(Sort-merge join)

(c)
sid = sid

(Scan & write to T1)

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

Supplier
(File scan)

(b) $\sigma_{\text{pno}=2}$ (Scan & write to T2)

Supply
(File scan)

Different but equivalent logical query plan; different physical

```
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: Overview


- Compute cost of each operator, which depends on:
 - Table statistics (# of tuples produced)
 - Algorithm used to implement each operator
- Cost of a physical plan =
 $\text{sum}(\text{each operator cost})$
- Cost each plan and choose the one with lowest cost

Estimating Table Statistics

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 1



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

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500


T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 1



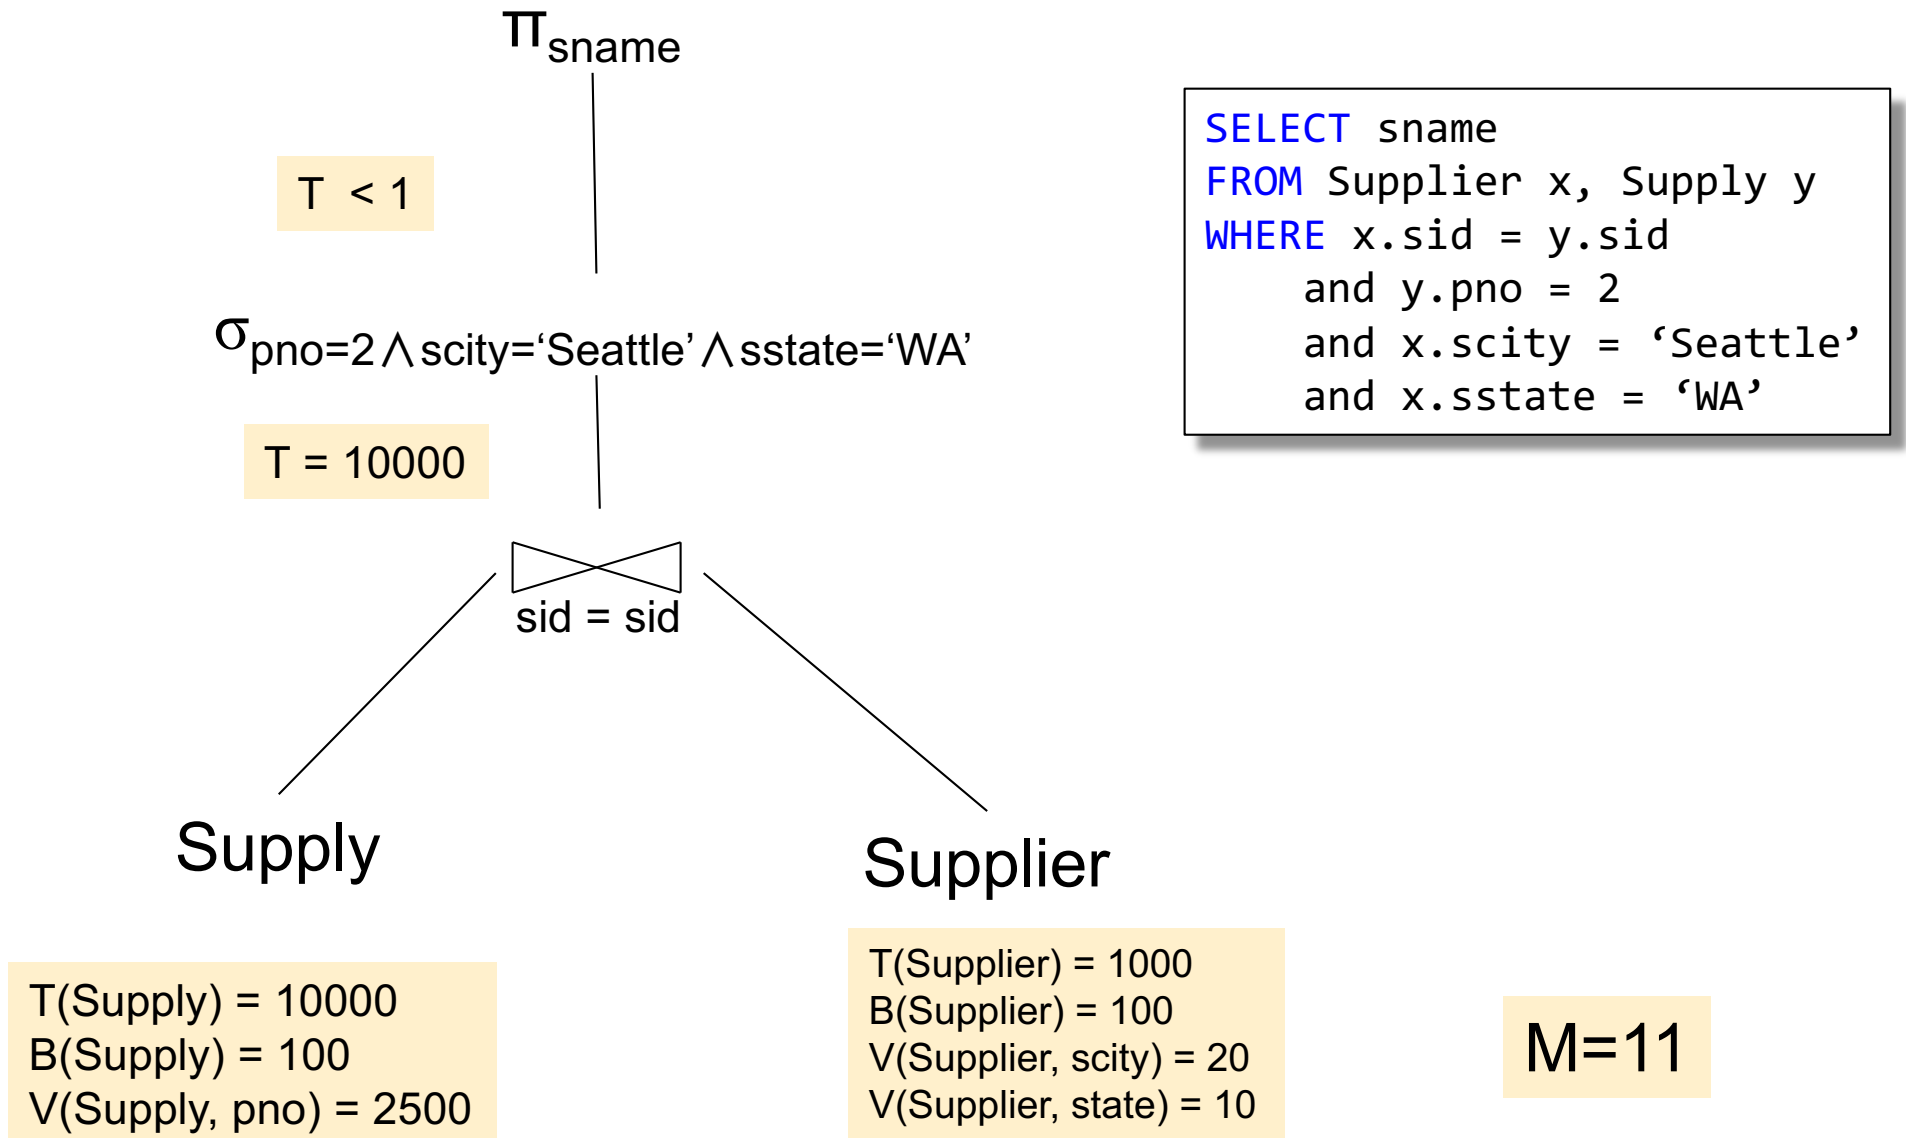
```
SELECT sname
FROM Supplier x, Supply y
WHERE x.sid = y.sid
      and y.pno = 2
      and x.scity = 'Seattle'
      and x.sstate = 'WA'
```

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)


Logical Query Plan 1



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 2



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

T(Supply) = 10000
B(Supply) = 100
V(Supply, pno) = 2500

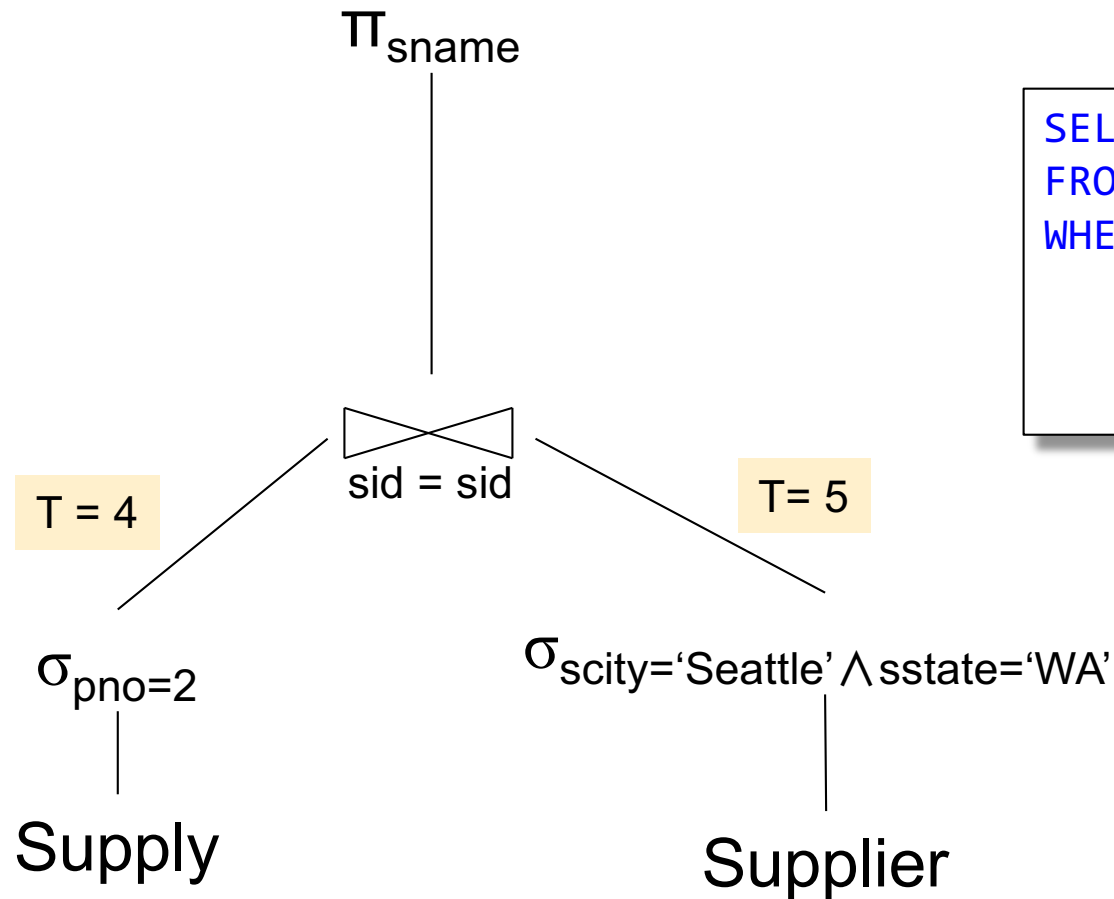
T(Supplier) = 1000
B(Supplier) = 100
V(Supplier, scity) = 20
V(Supplier, state) = 10

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 2



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

$T(\text{Supply}) = 10000$
 $B(\text{Supply}) = 100$
 $V(\text{Supply}, \text{pno}) = 2500$

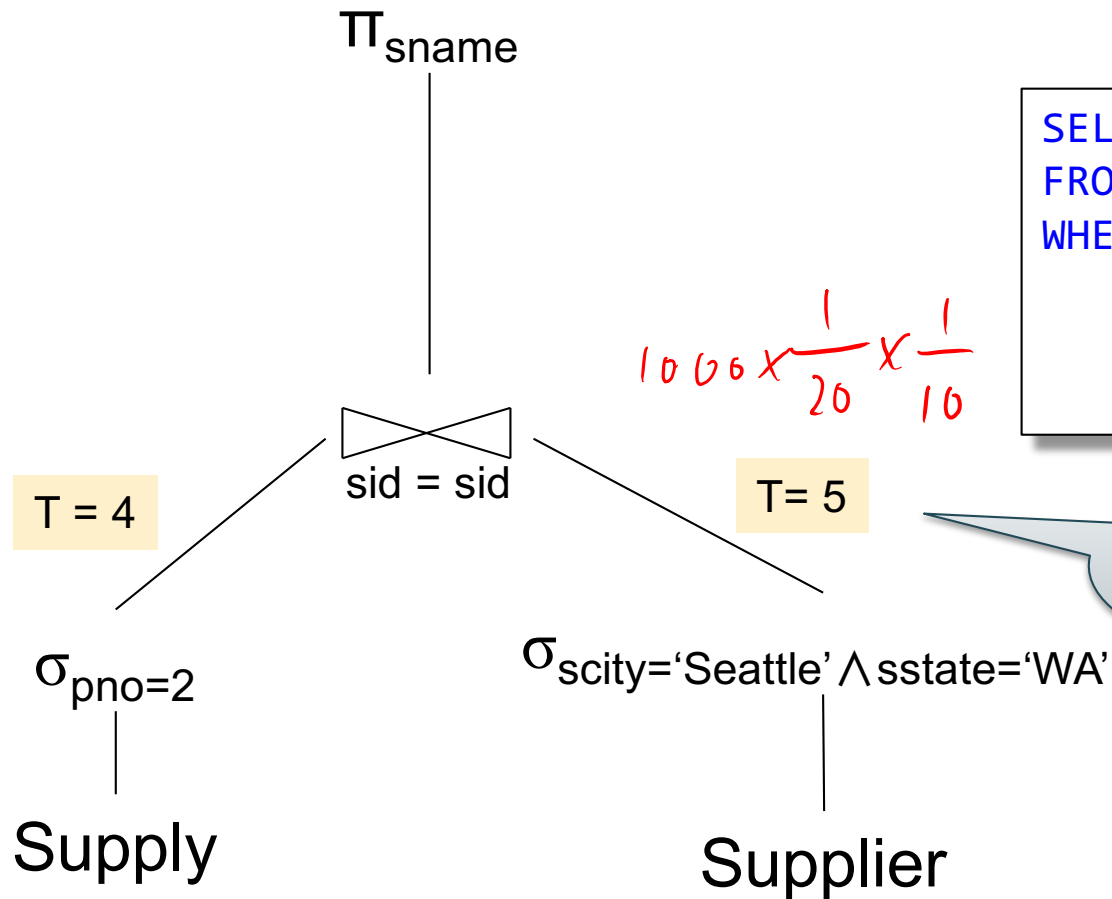
$T(\text{Supplier}) = 1000$
 $B(\text{Supplier}) = 100$
 $V(\text{Supplier}, \text{scity}) = 20$
 $V(\text{Supplier}, \text{state}) = 10$

M=11

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 2



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

Very wrong!
Why?

$T(\text{Supply}) = 10000$
 $B(\text{Supply}) = 100$
 $V(\text{Supply}, pno) = 2500$

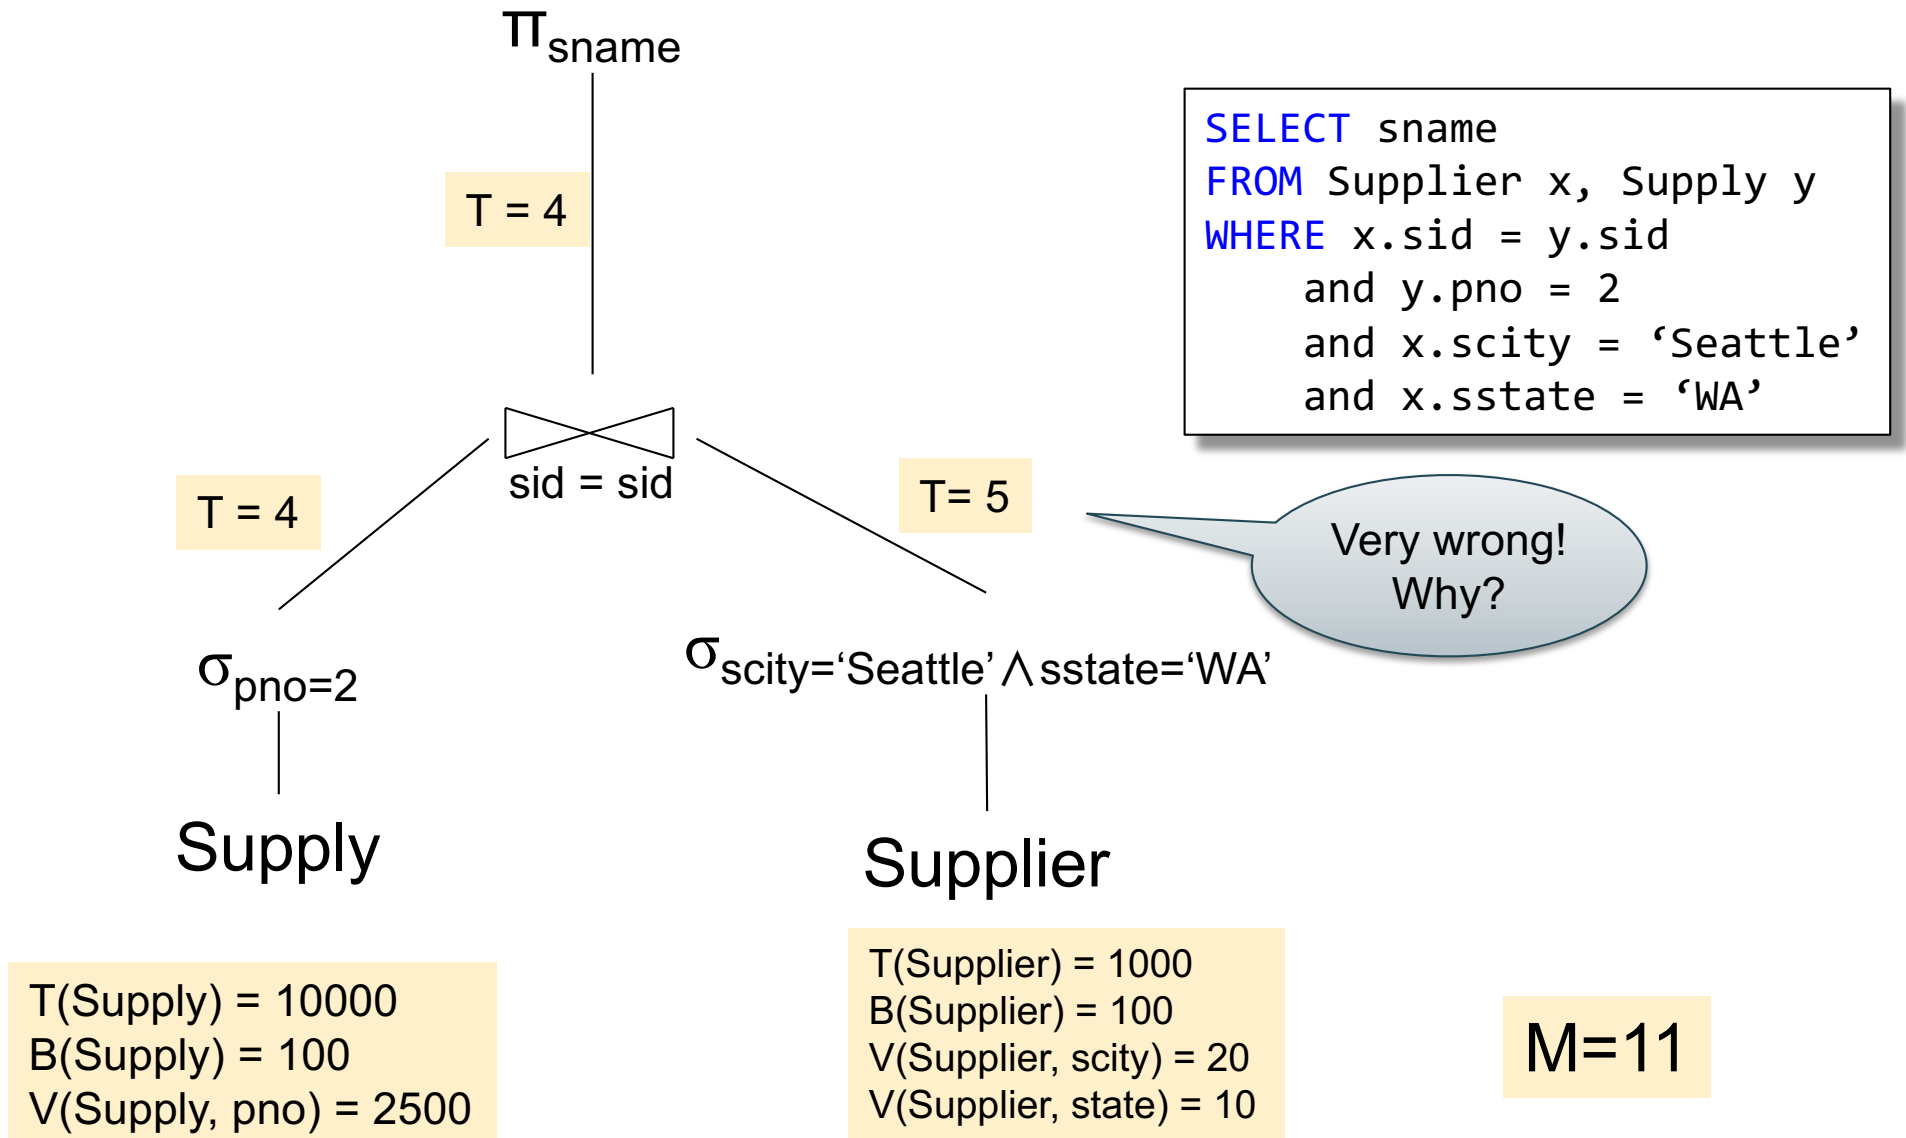
$T(\text{Supplier}) = 1000$
 $B(\text{Supplier}) = 100$
 $V(\text{Supplier}, scity) = 20$
 $V(\text{Supplier}, state) = 10$

$M=11$

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

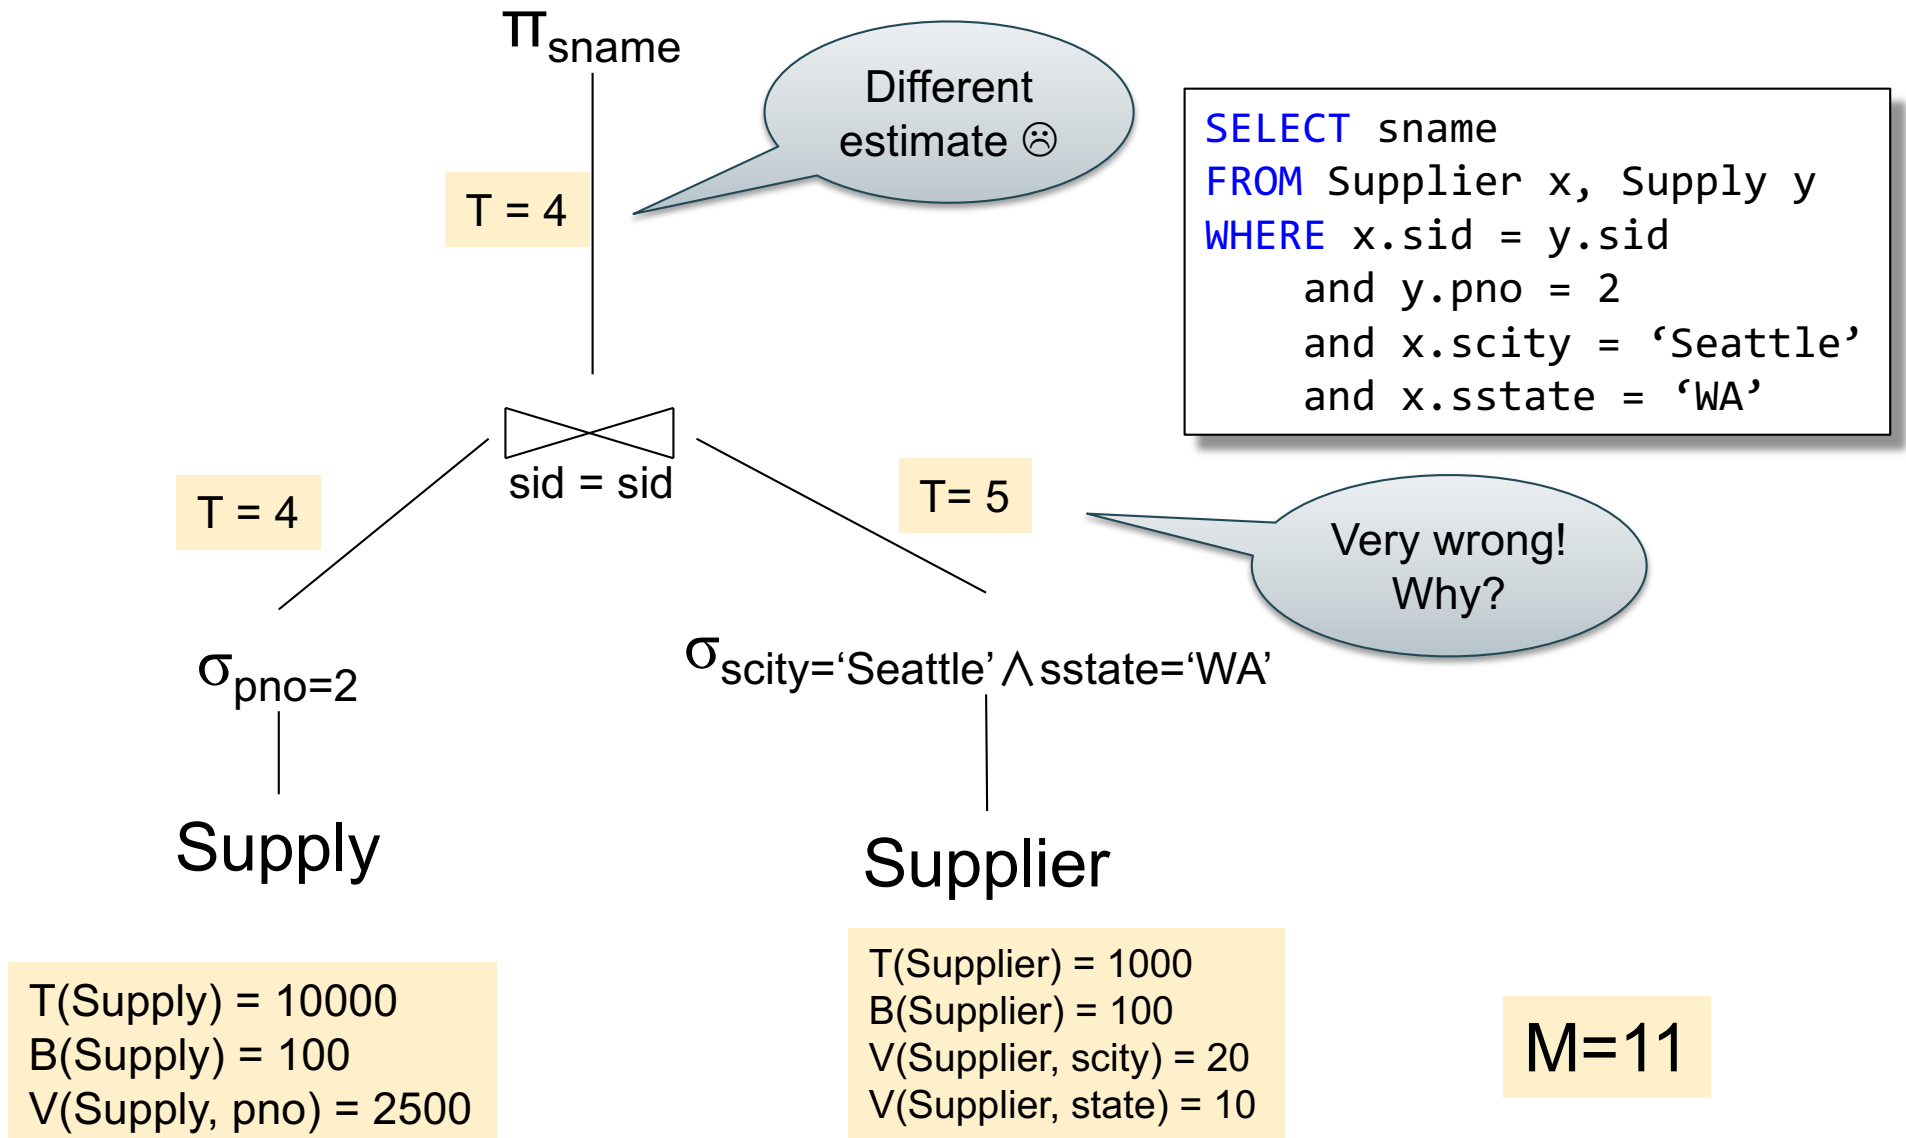
Logical Query Plan 2



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Logical Query Plan 2

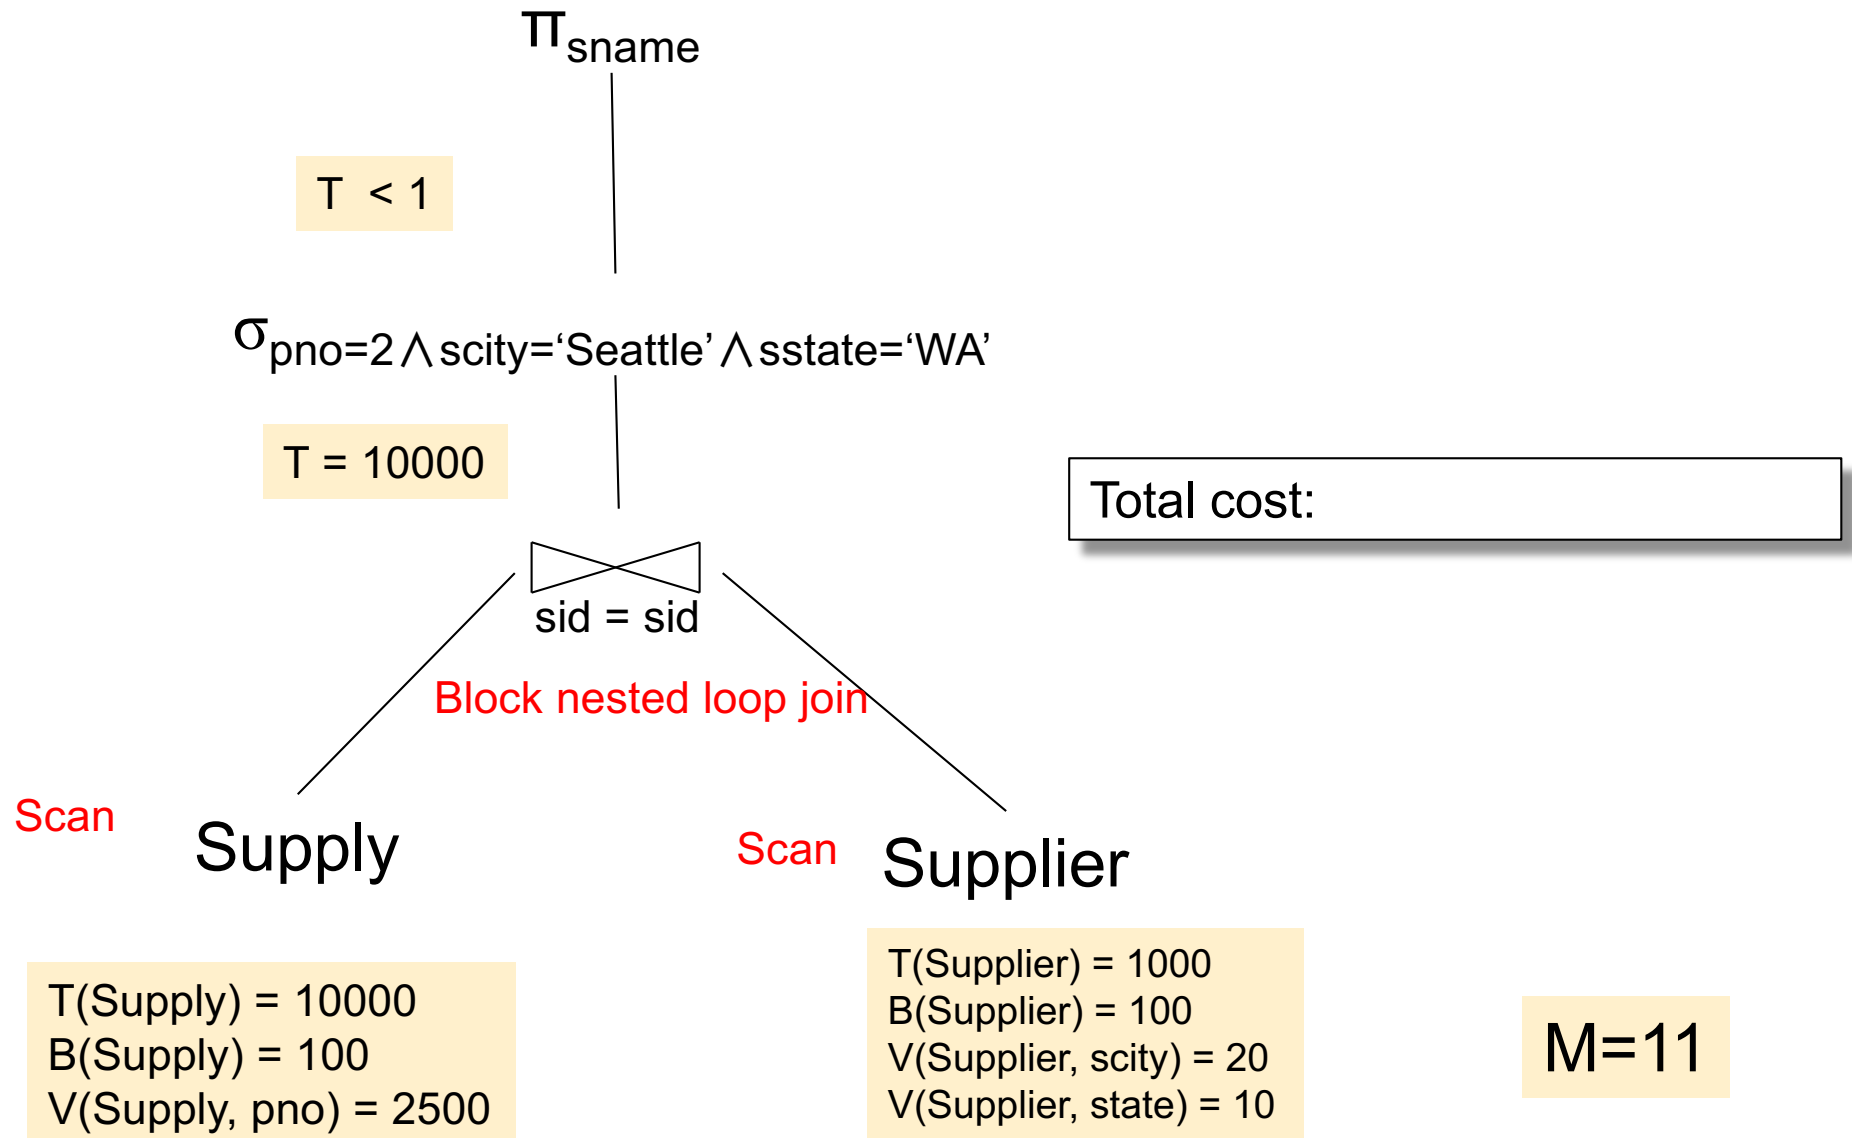


Computing Plan Costs

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

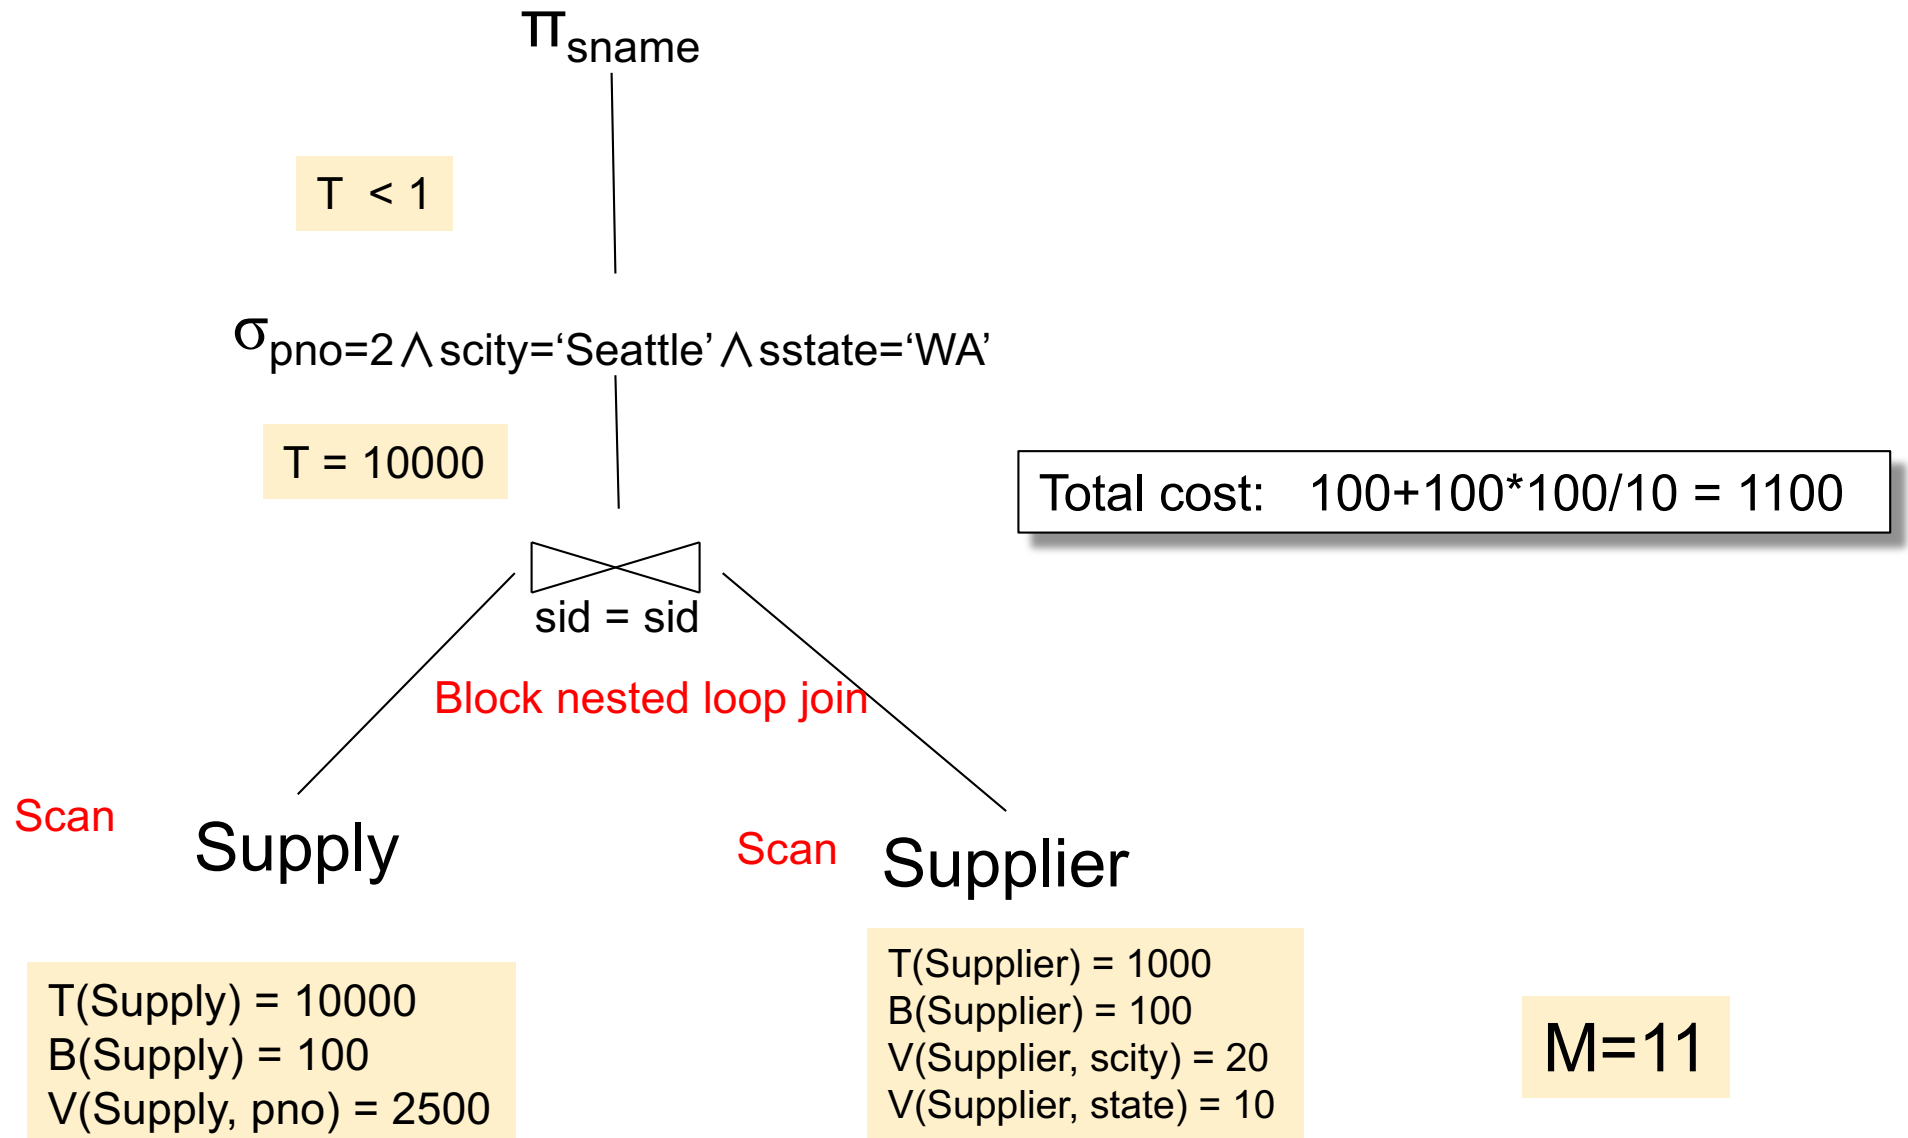
Physical Plan 1



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

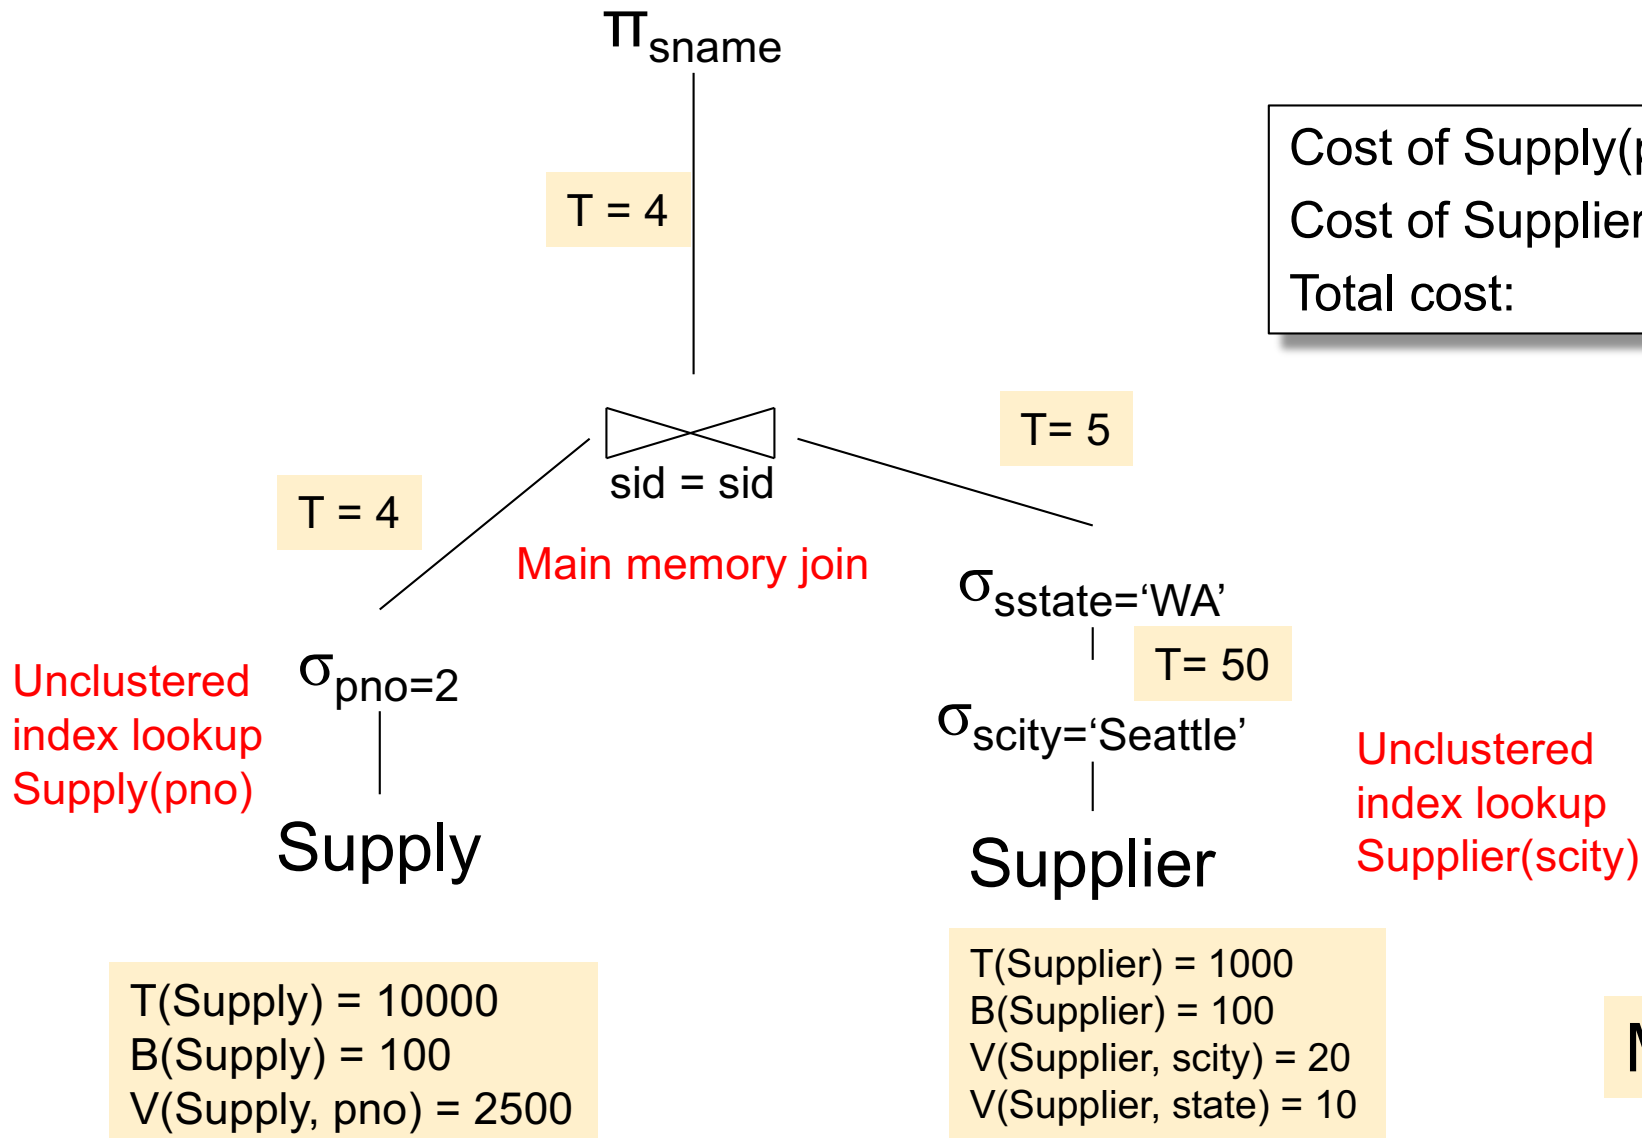
Physical Plan 1



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 2

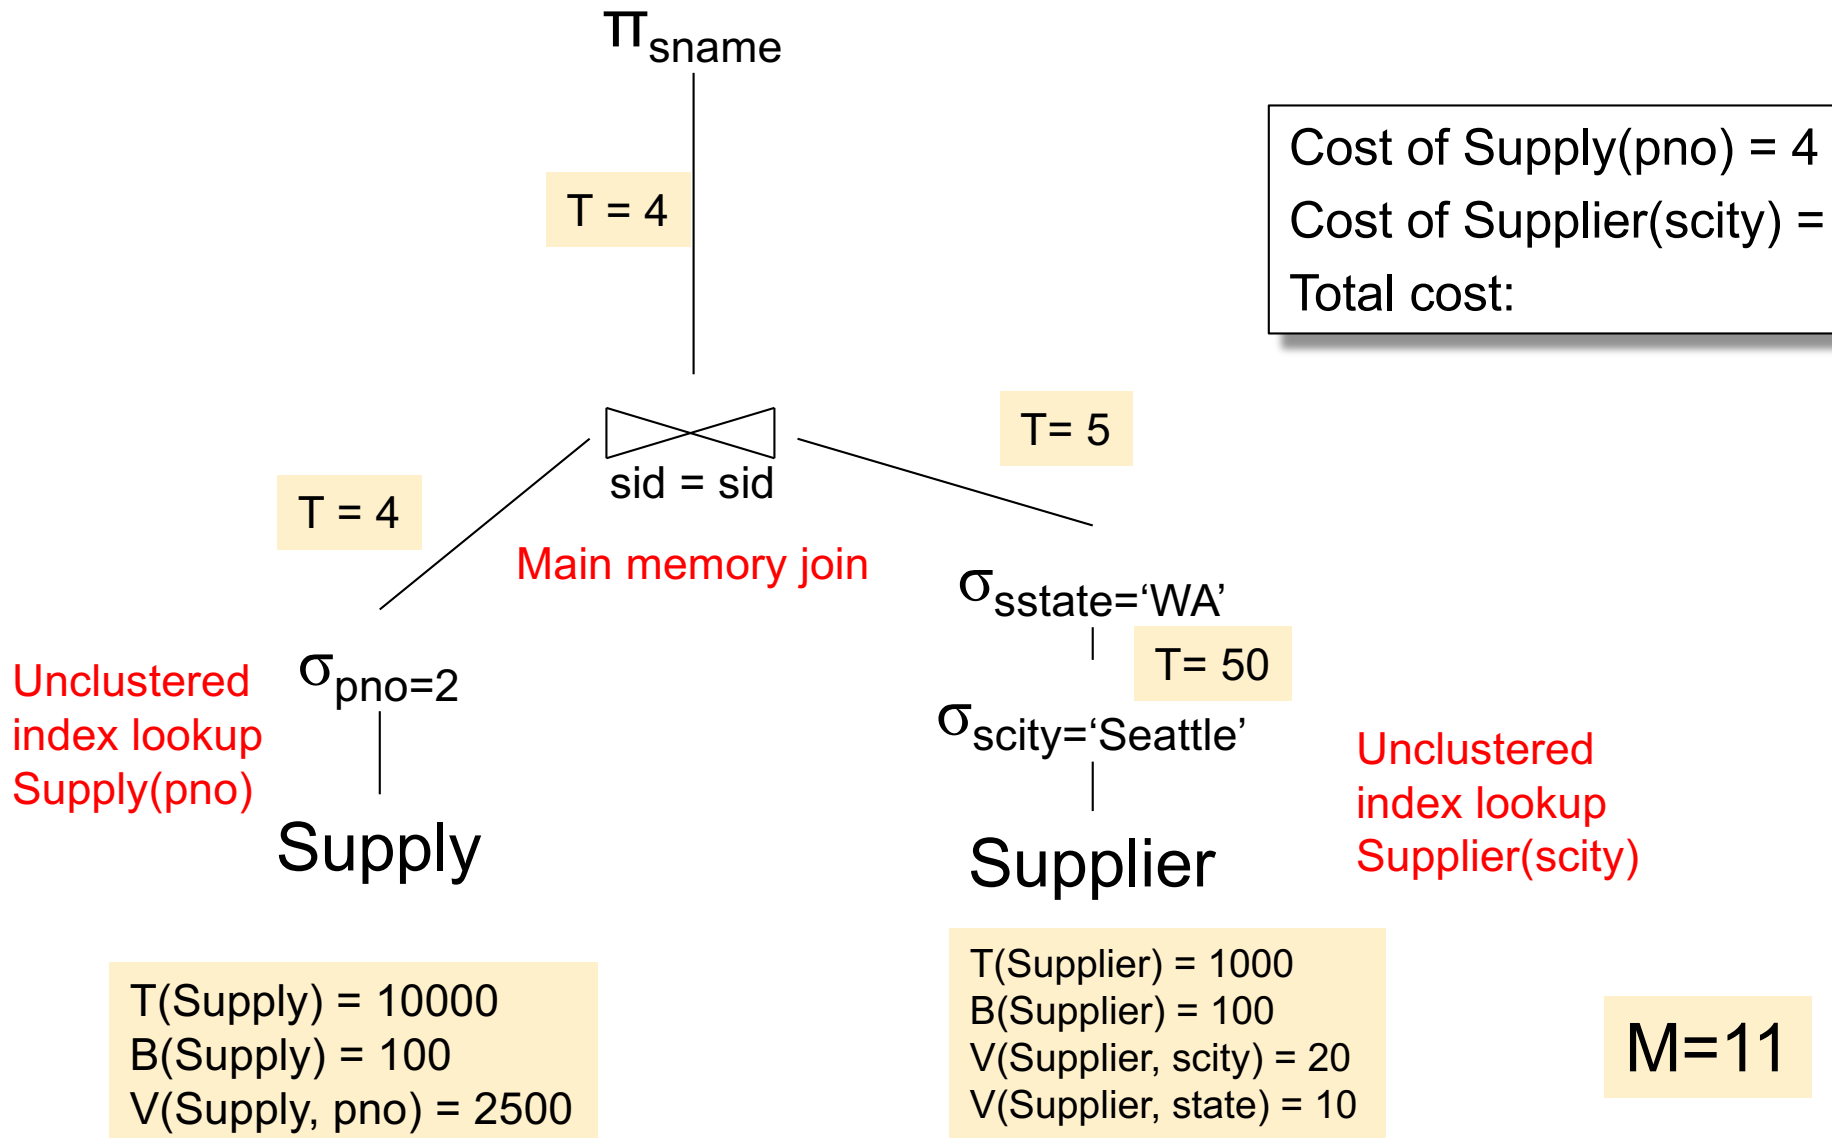


Cost of Supply(pno) =
Cost of Supplier(scity) =
Total cost:

Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

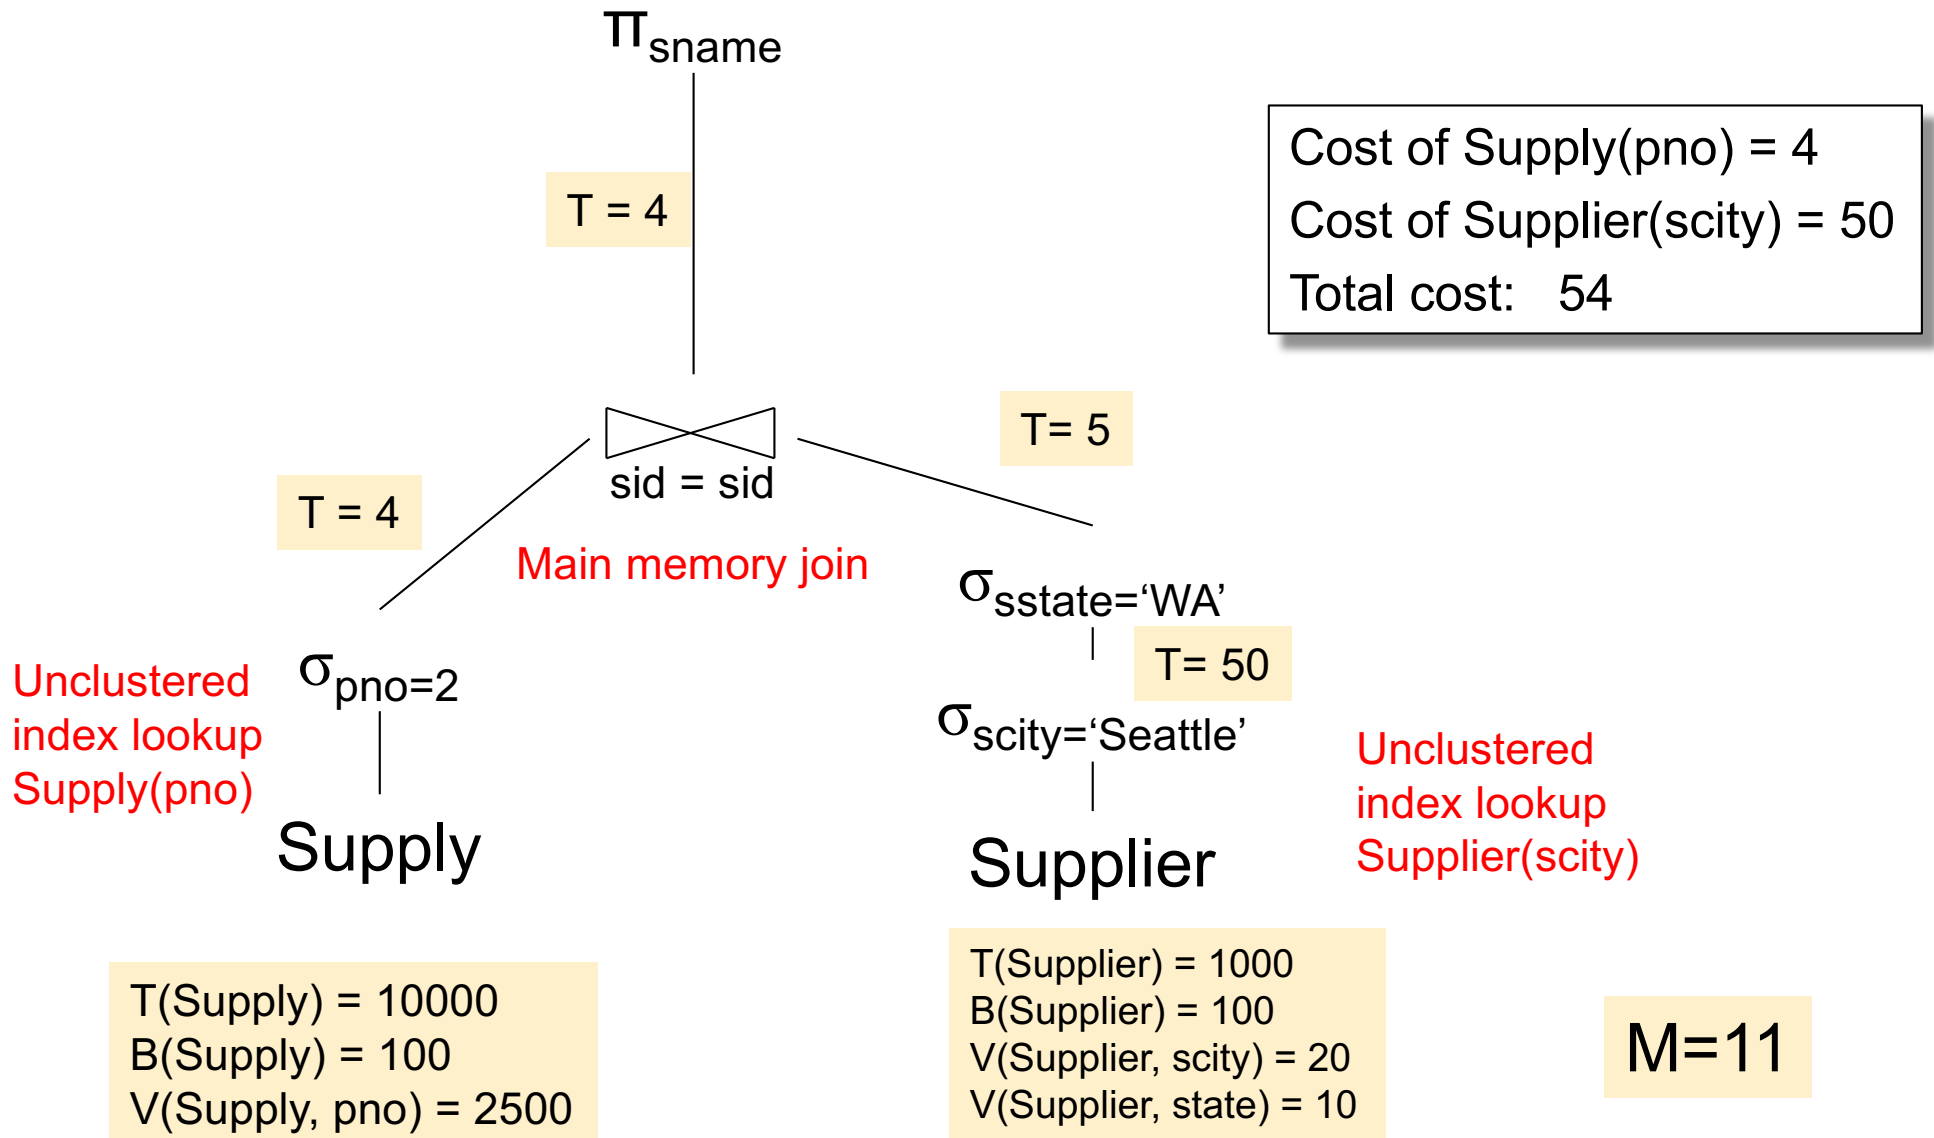
Physical Plan 2



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

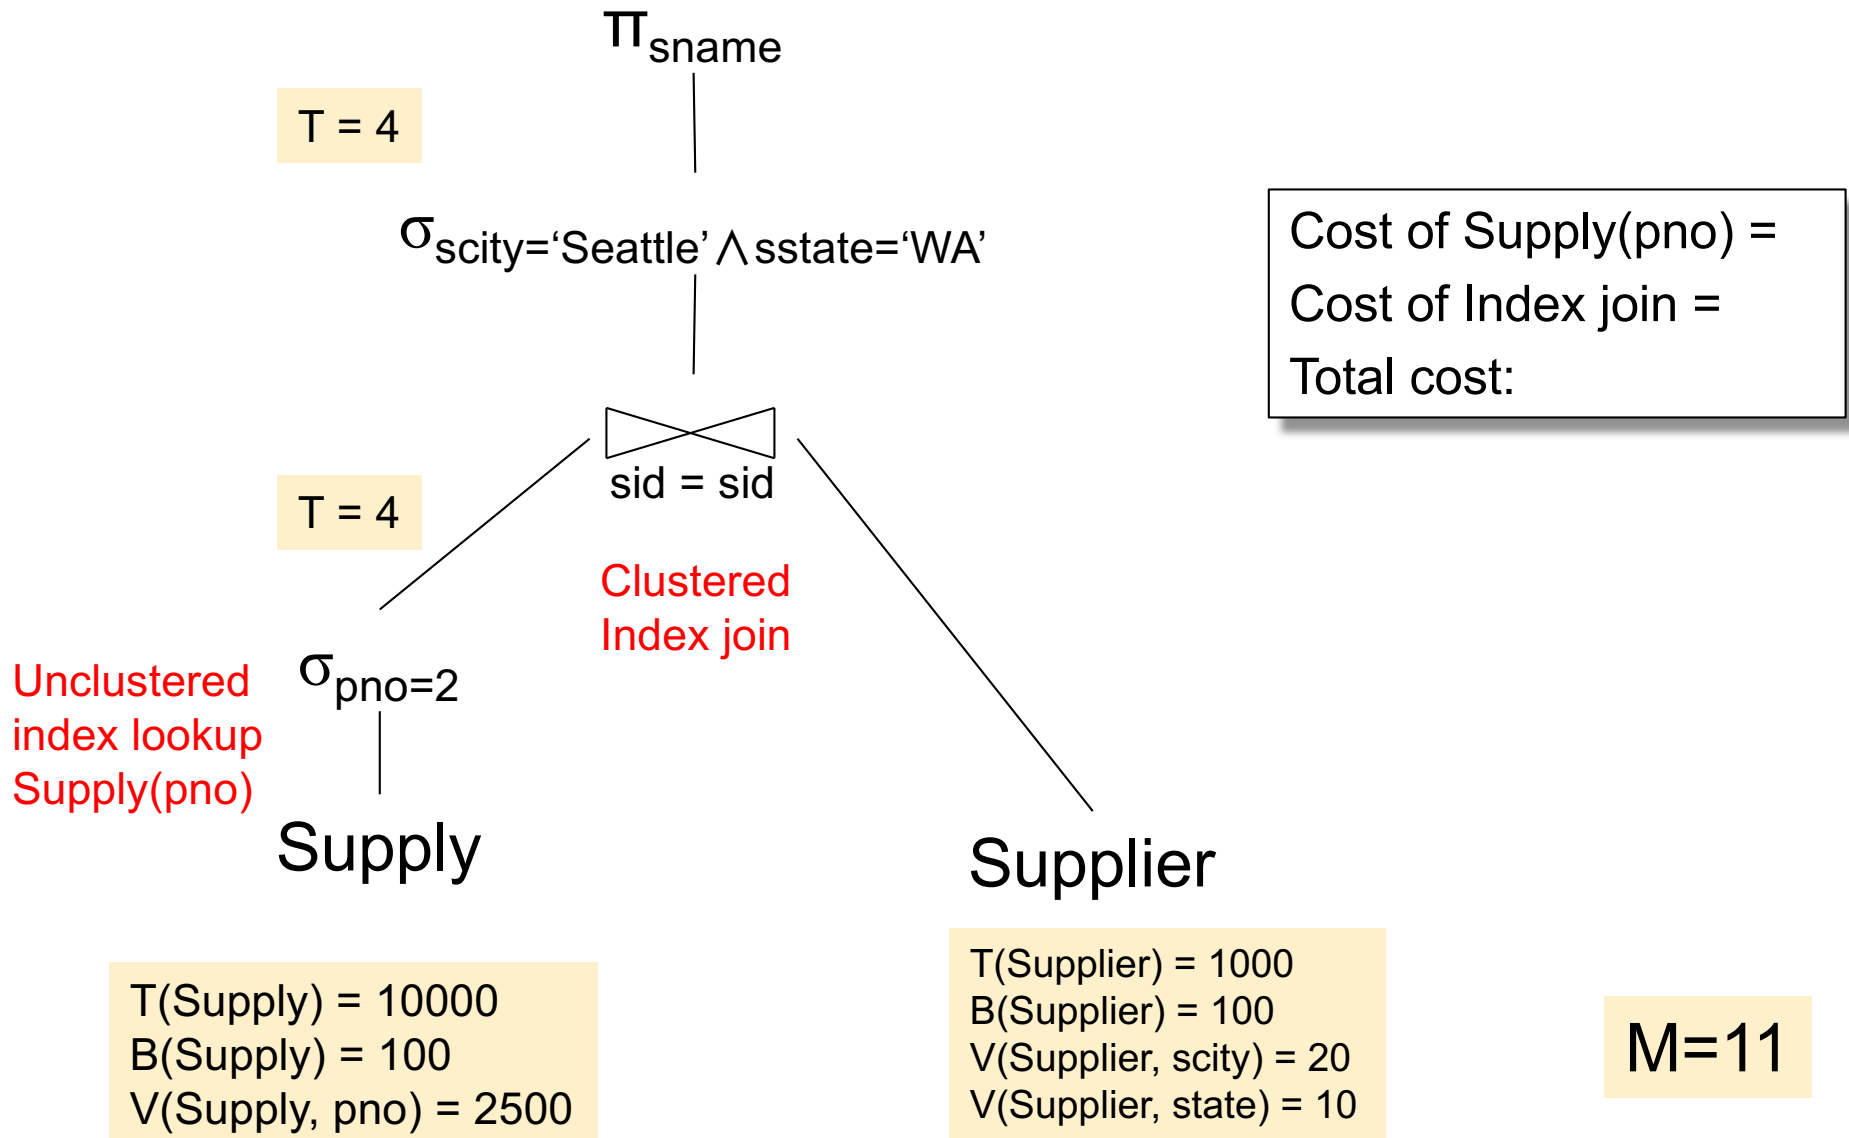
Physical Plan 2



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

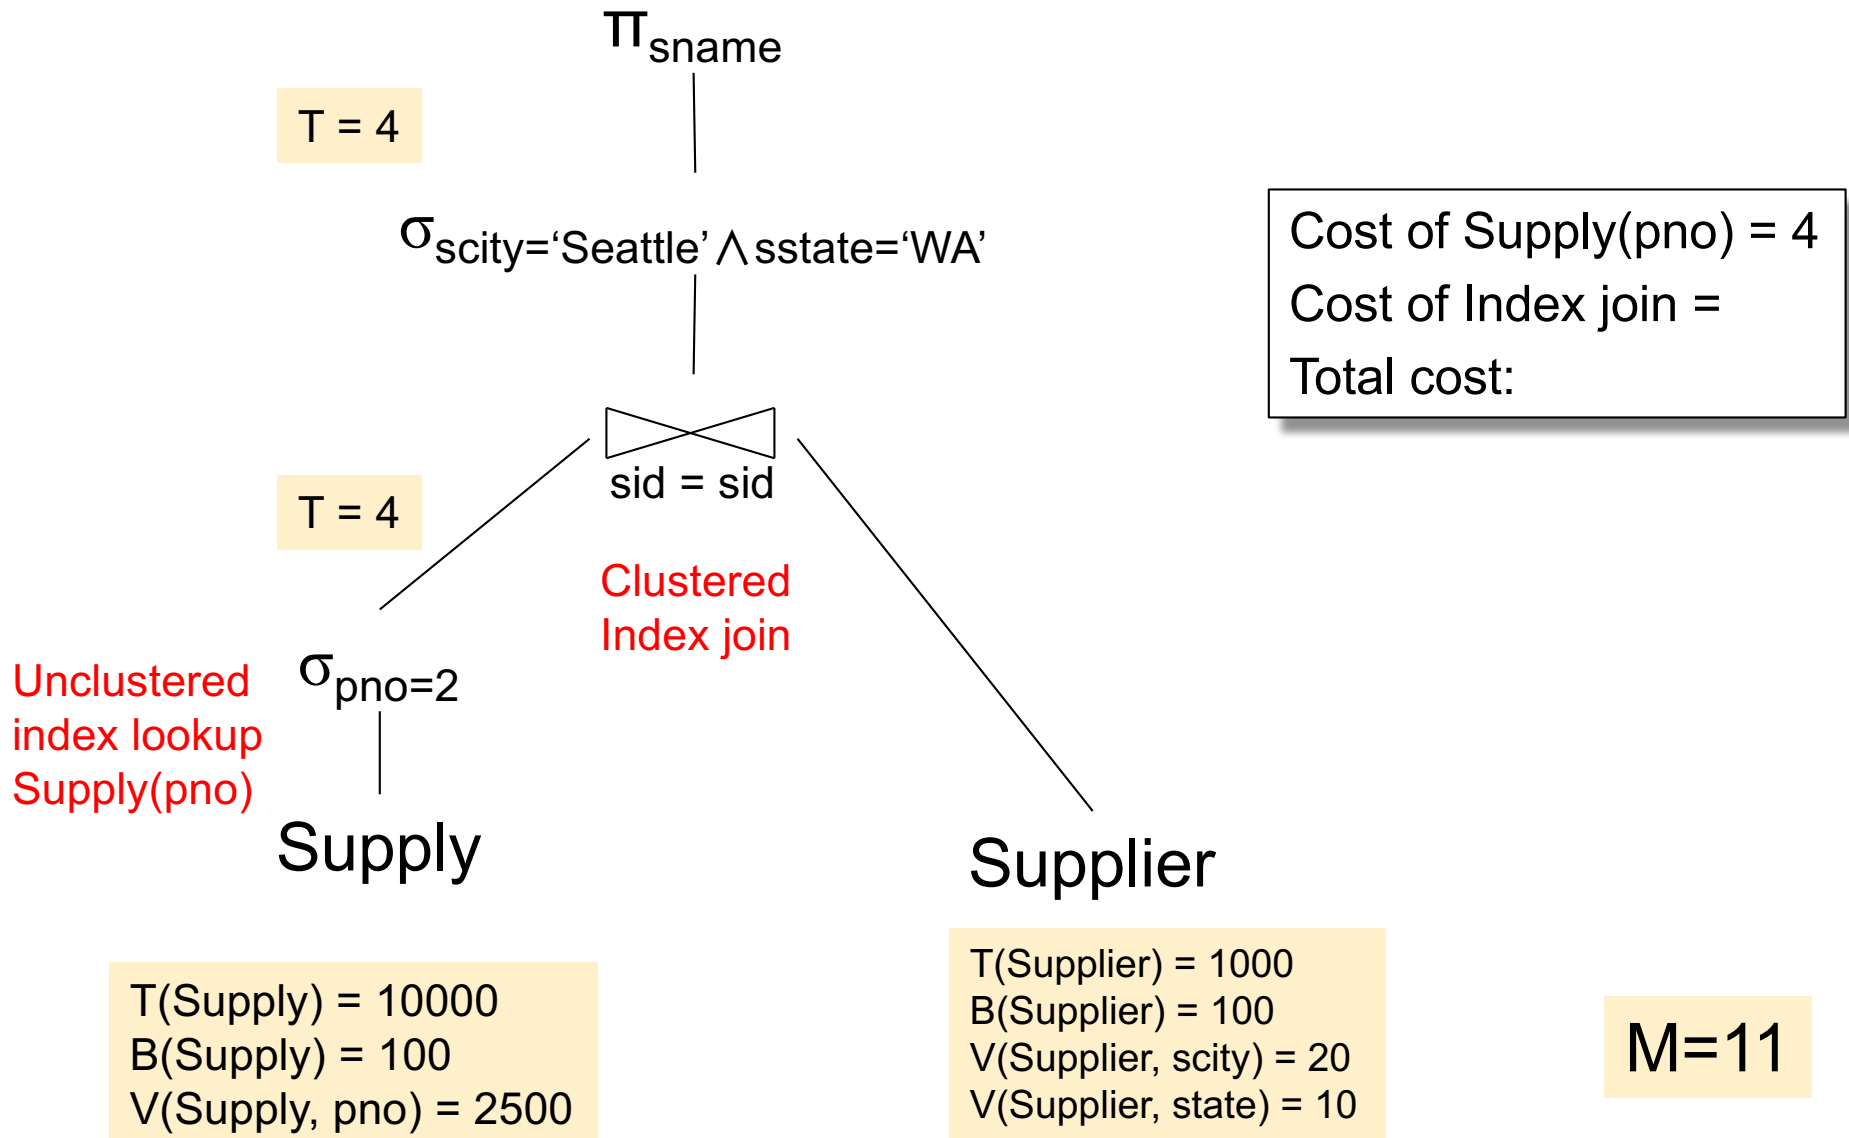
Physical Plan 3



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

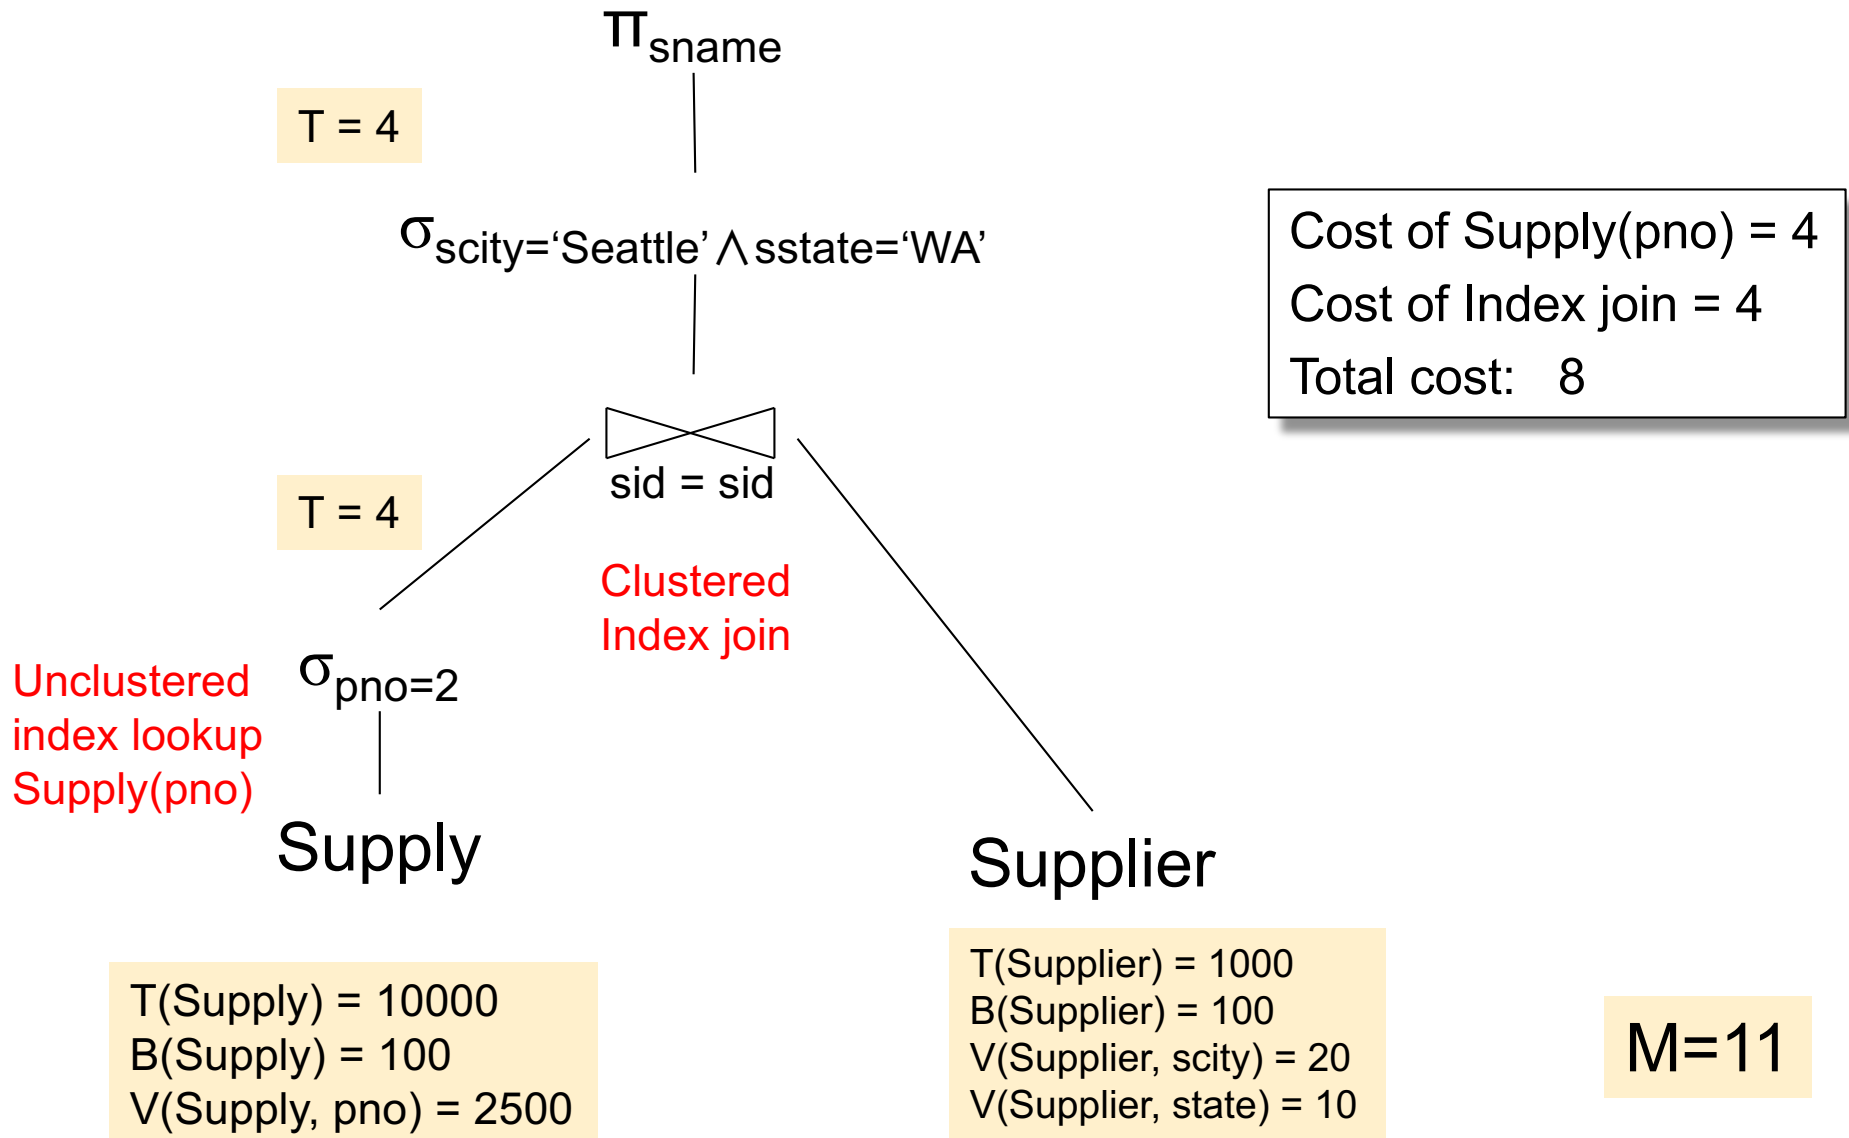
Physical Plan 3



Supplier(sid, sname, scity, sstate)

Supply(sid, pno, quantity)

Physical Plan 3



Query Optimizer Summary

- Input: A logical query plan
- Output: A good physical query plan
- Basic query optimization algorithm
 - Enumerate alternative plans (logical and physical)
 - Compute estimated cost of each plan
 - Choose plan with lowest cost
- This is called cost-based optimization
 - More in CSE 444