

# Introduction to Data Management

## CSE 344

### Lecture 12: Cost Estimation Relational Calculus

# Announcements

- HW3 due tonight
- WQ4 and HW4 out
  - Due on Thursday 2/9

# Midterm!

- Monday, February 13<sup>th</sup> in class
- Contents
  - Lectures and sections through February 8th
  - Homework 1 through 4
  - Webquiz 1 through 4
- Closed book. No computers, phones, watches, etc.!
- Can bring one letter-sized piece of paper with notes
  - Can write on both sides
  - You might want to save it for the final

# Today's Outline

- Finish cost estimation
- Relational calculus

# Review

- Estimate cost of physical query plans
  - Based on # of I/O operations
  - Estimate cost for each operator
  - Cost of entire plan =  $\Sigma$  operator cost
- Cost for selection operator
- Cost for join operator

# Review: Cost Parameters

- ~~Cost = I/O + CPU + Network BW~~
  - We will focus on I/O in this class
- **Parameters:**
  - **$B(R)$**  = # of blocks (i.e., pages) for relation R
  - **$T(R)$**  = # of tuples in relation R
  - **$V(R, a)$**  = # of distinct values of attribute a
    - When  **$a$**  is a key,  **$V(R, a) = T(R)$**
    - When  **$a$**  is not a key,  **$V(R, a)$**  can be anything  $\leq T(R)$
- Where do these values come from?
  - DBMS collects **statistics** about data on disk

# Index Based Selection

- Example:

$$\begin{aligned} B(R) &= 2000 \\ T(R) &= 100,000 \\ V(R, a) &= 20 \end{aligned}$$

$$\text{cost of } \sigma_{a=v}(R) = ?$$

- Table scan:
- Index based selection:

# Index Based Selection

- Example:

$$\begin{aligned} B(R) &= 2000 \\ T(R) &= 100,000 \\ V(R, a) &= 20 \end{aligned}$$

$$\text{cost of } \sigma_{a=v}(R) = ?$$

- Table scan:  $B(R) = 2,000$  I/Os
- Index based selection:



# Index Based Selection

- Example: 

$B(R) = 2000$
$T(R) = 100,000$
$V(R, a) = 20$

cost of $\sigma_{a=v}(R) = ?$
-------------------------------
- Table scan:  $B(R) = 2,000$  I/Os
- Index based selection:
  - If index is clustered:
  - If index is unclustered:

# Index Based Selection

- Example: 

$B(R) = 2000$
$T(R) = 100,000$
$V(R, a) = 20$

cost of $\sigma_{a=v}(R) = ?$
-------------------------------
- Table scan:  $B(R) = 2,000$  I/Os
- Index based selection:
  - If index is clustered:  $B(R) * \underline{1/V(R,a)} = 100$  I/Os
  - If index is unclustered:

# Index Based Selection

- Example: 

$B(R) = 2000$
$T(R) = 100,000$
$V(R, a) = 20$

cost of $\sigma_{a=v}(R) = ?$
-------------------------------
- Table scan:  $B(R) = 2,000$  I/Os
- Index based selection:
  - If index is clustered:  $B(R) * 1/V(R,a) = 100$  I/Os
  - If index is unclustered:  $T(R) * 1/V(R,a) = 5,000$  I/Os

# Index Based Selection

- Example:

$$\begin{aligned} B(R) &= 2000 \\ T(R) &= 100,000 \\ V(R, a) &= 20 \end{aligned}$$

$$\text{cost of } \sigma_{a=v}(R) = ?$$

- ① • Table scan:  $B(R) = 2,000$  I/Os
- Index based selection:
  - ② – If index is clustered:  $B(R) * 1/V(R,a) = 100$  I/Os
  - If index is unclustered:  $T(R) * 1/V(R,a) = 5,000$  I/Os

Lesson: Don't build unclustered indexes when  $V(R,a)$  is small !

# Cost of Executing Operators (Focus on Joins)

# Outline

- **Join operator algorithms**
  - One-pass algorithms (Sec. 15.2 and 15.3)
  - Index-based algorithms (Sec 15.6)
- Note about readings:
  - In class, we discuss only algorithms for joins
  - Other operators are easier: read the book

# Join Algorithms

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

# Nested Loop Joins

- Tuple-based nested loop  $R \bowtie S$
- R is the outer relation, S is the inner relation

```
for each tuple  $t_1$  in R do  
  for each tuple  $t_2$  in S do  
    if  $t_1$  and  $t_2$  join then output  $(t_1, t_2)$ 
```

What is the **Cost**?



# Nested Loop Joins

- Tuple-based nested loop  $R \bowtie S$
- $R$  is the outer relation,  $S$  is the inner relation

```
for each tuple  $t_1$  in  $R$  do  
  for each tuple  $t_2$  in  $S$  do  
    if  $t_1$  and  $t_2$  join then output  $(t_1, t_2)$ 
```

- **Cost:**  $B(R) + T(R) B(S)$
- Multiple-pass since  $S$  is read many times

What is the **Cost**?

# Page-at-a-time Refinement

```
for each page of tuples r in R do  
  for each page of tuples s in S do  
    for all pairs of tuples t1 in r, t2 in s  
      if t1 and t2 join then output (t1,t2)
```

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

What is the **Cost**?

# Hash Join

Hash join:  $R \bowtie S$

- Scan R, build buckets 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


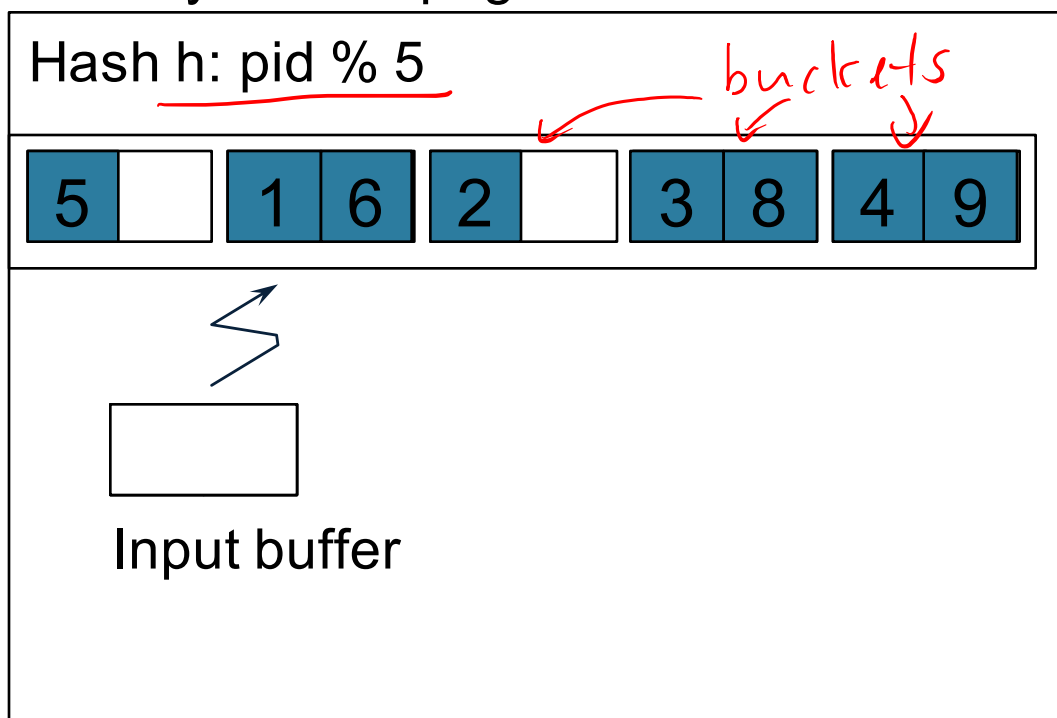


This is one page with two tuples

# Hash Join Example

Step 1: Scan Patient and **build** hash table in memory  
Can be done in method open()


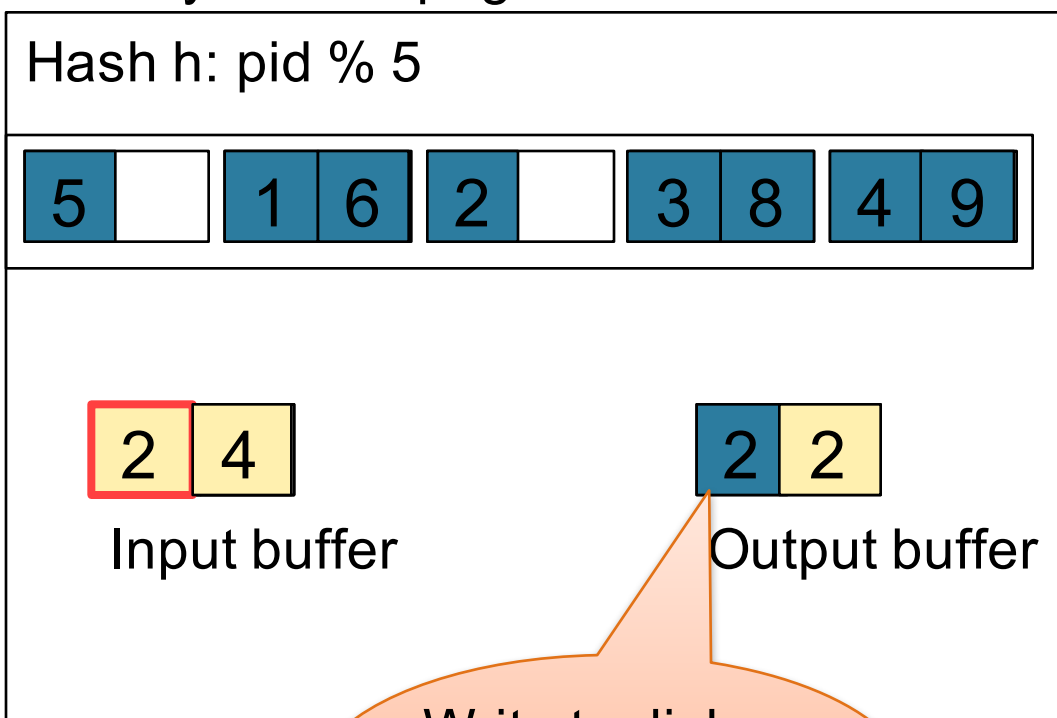
Memory M = 21 pages



# Hash Join Example

Step 2: Scan Insurance and **probe** into hash table  
 Done during  
 calls to next()

Memory M = 21 pages

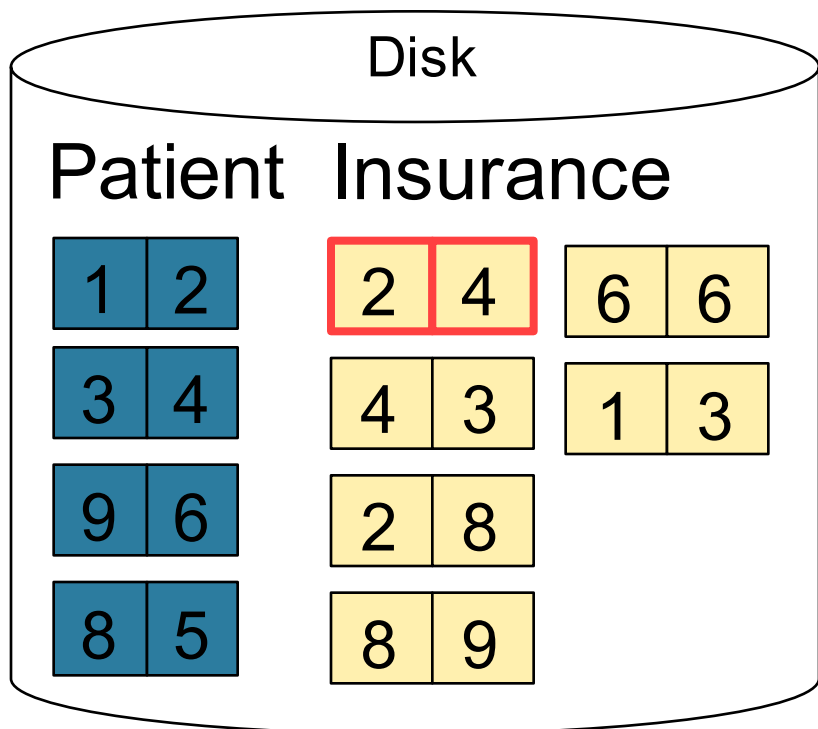



Write to disk or  
 pass to next  
 operator

# Hash Join Example

Step 2: Scan Insurance and **probe** into hash table  
 Done during calls to next()

Memory M = 21 pages

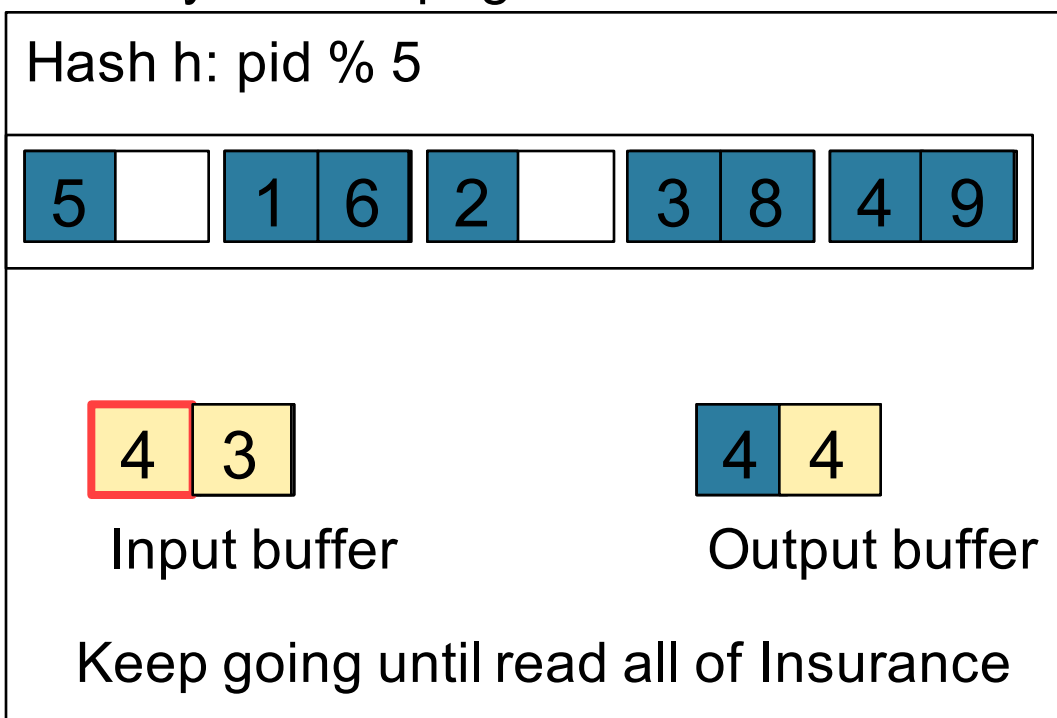





# Hash Join Example

Step 2: Scan Insurance and **probe** into hash table  
 Done during calls to next()

Memory M = 21 pages



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



# 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

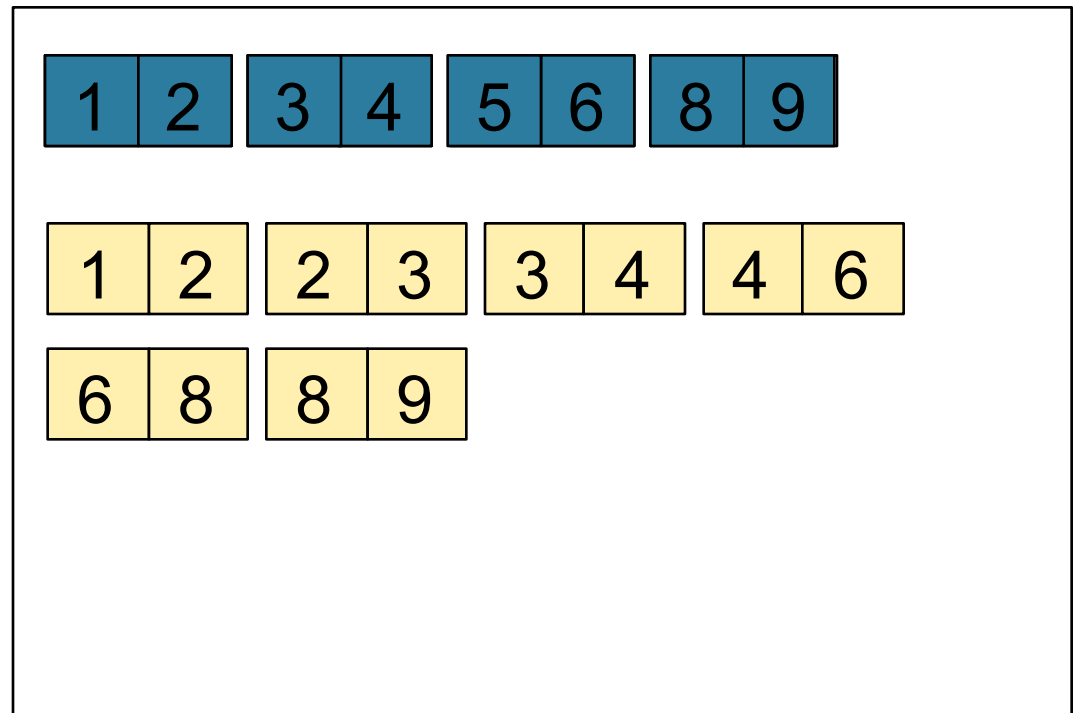

Memory M = 21 pages



# Sort-Merge Join Example

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


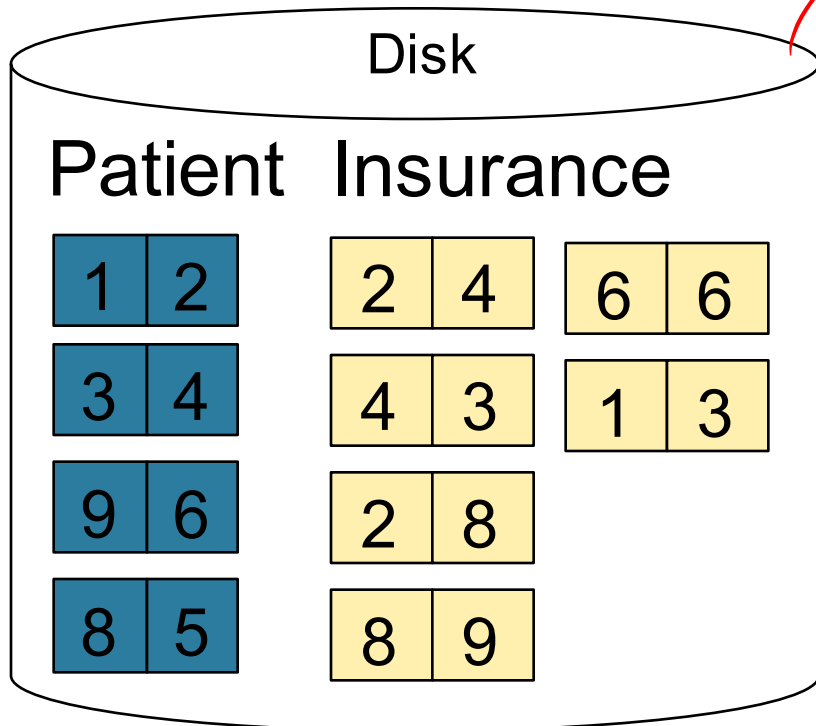
Memory M = 21 pages



# Sort-Merge Join Example

Step 3: Merge Patient and Insurance

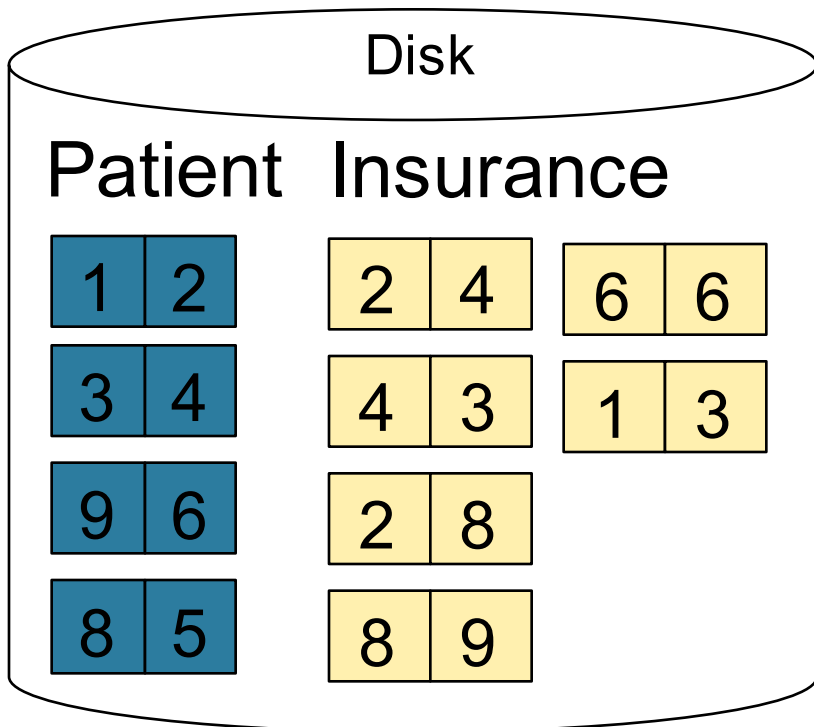
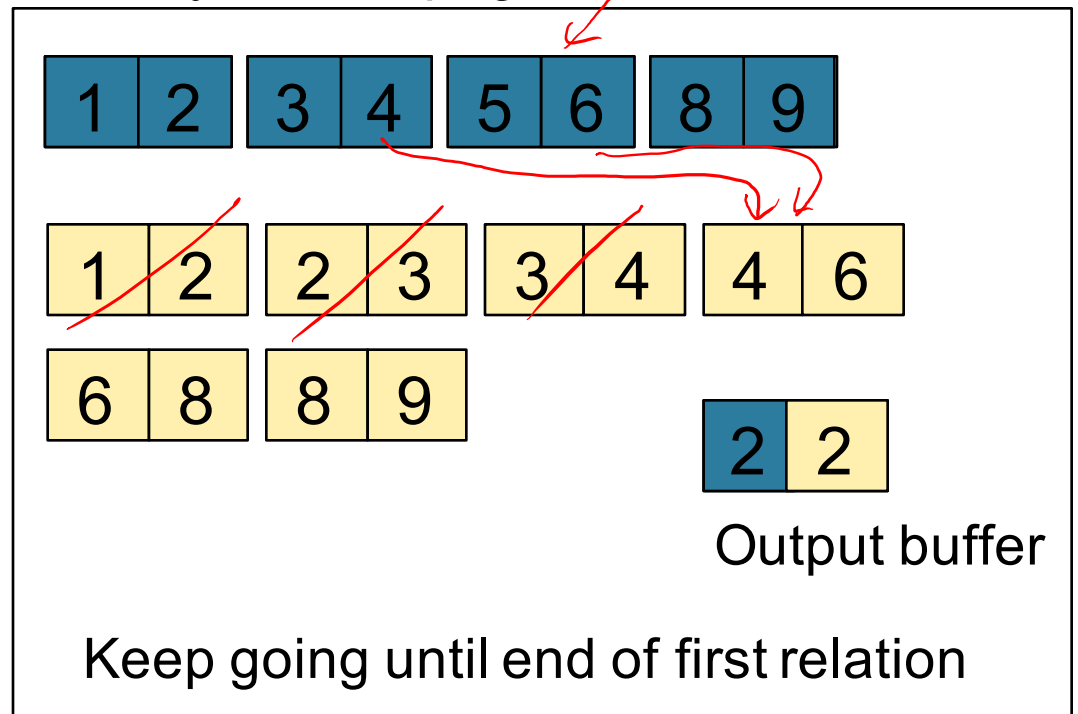
Memory M = 21 pages



# Sort-Merge Join Example

Step 3: **Merge** Patient and Insurance

Memory M = 21 pages



# Cost of Query Plans

T(Supplier) = 1000  
T(Supply) = 10,000

B(Supplier) = 100  
B(Supply) = 100

V(Supplier,scity) = 20  
V(Supplier,state) = 10  
V(Supply,pno) = 2,500

M = 11

# Physical Query Plan 1

(On the fly)

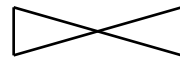
$\Pi_{\text{sname}}$

Selection and project on-the-fly  
→ No additional cost.

(On the fly)

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

(Nested loop)

  
sid = sid

Total cost of plan is thus cost of join:  
= B(Supplier)+B(Supplier)\*B(Supply)  
= 100 + 100 \* 100  
**= 10,100 I/Os**

Supplier  
(File scan)

Supply  
(File scan)

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



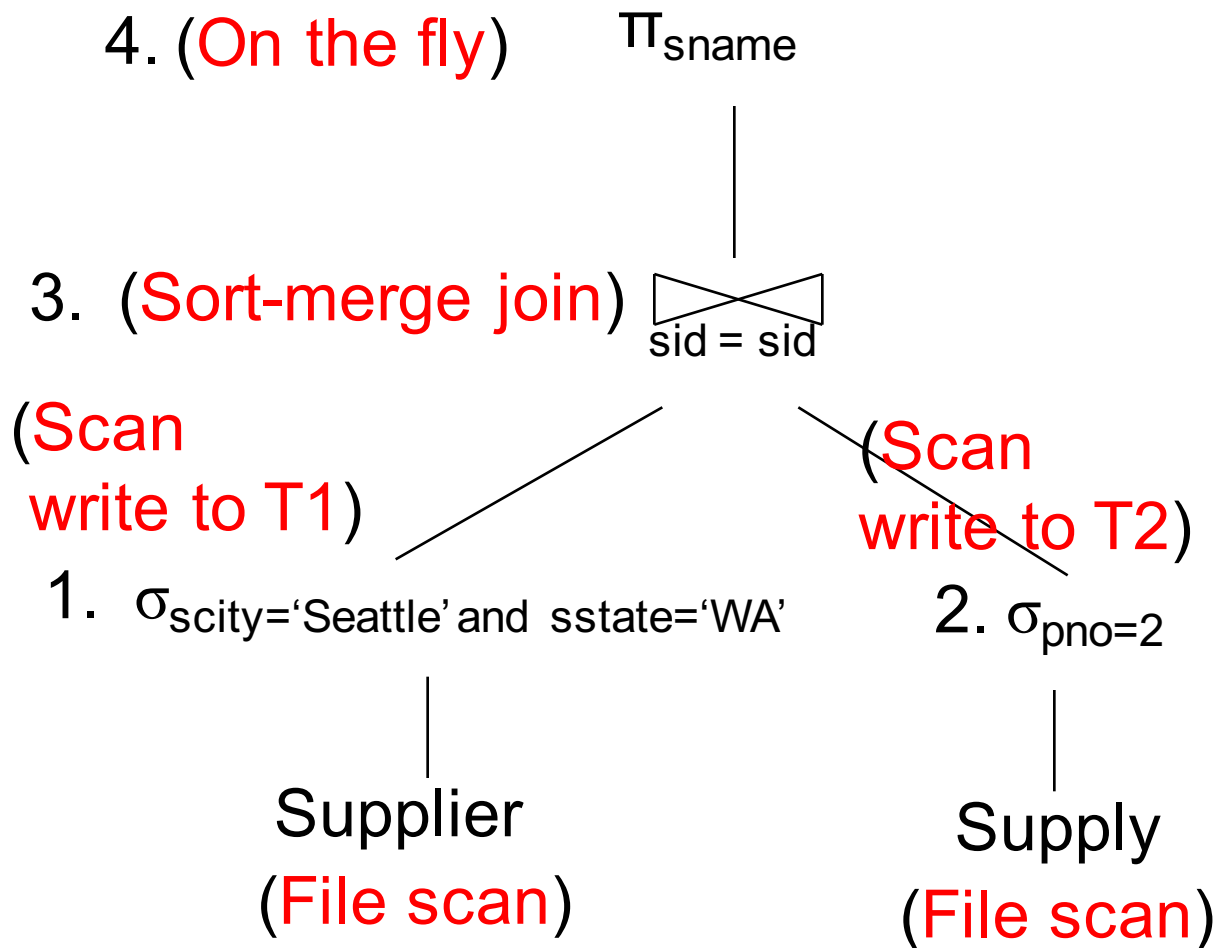
T(Supplier) = 1000  
T(Supply) = 10,000

B(Supplier) = 100  
B(Supply) = 100

V(Supplier,scity) = 20  
V(Supplier,state) = 10  
V(Supply,pno) = 2,500

M = 11

# Physical Query Plan 2



$\checkmark$

Total cost  $\checkmark \rightarrow 1$

= 100 + 100 \* 1/20 \* 1/10

(step 1)

+ 100 + 100 \* 1/2500  $\rightarrow 1$

(step 2)

+ 2

(step 3)

+ 0

(step 4)

Total cost  $\approx$  **204 I/Os**

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