

Introduction to Database Systems CSE 414

Lecture 28: Intro to Query Optimization

CSE 414 - Spring 2018

1

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

2

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!

CSE 414 - Spring 2018


3

Review: Join Algorithms

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

CSE 414 - Spring 2018

4



Hash Join

CSE 414 - Spring 2018

5

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

CSE 414 - Spring 2018

6

Hash Join Example

Patient(pid, name, address)
 Insurance(pid, provider, policy_nb)
 Patient ⋈ Insurance

Two tuples per page

Patient			Insurance		
1	'Bob'	'Seattle'	2	'Blue'	123
2	'Ela'	'Everett'	4	'Prem'	432
3	'Jill'	'Kent'	4	'Prem'	343
4	'Joe'	'Seattle'	3	'GrpH'	554

7

Hash Join Example

Patient ⋈ Insurance

Memory M = 21 pages

Showing pid only

Some large-enough #

This is one page with two tuples

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

8

Hash Join Example

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

Memory M = 21 pages

Hash h: pid % 5

Input buffer

buckets

bucket	pid
=0	5
=1	1, 6, 2
=2	
=3	3, 8, 4, 9
=4	

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

9

Hash Join Example

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

Memory M = 21 pages

Hash h: pid % 5

Input buffer

Output buffer

Write to disk or pass to next operator

bucket	pid
=0	5
=1	1, 6, 2
=2	
=3	3, 8, 4, 9
=4	

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

10

Hash Join Example

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

Memory M = 21 pages

Hash h: pid % 5

Input buffer

Output buffer

bucket	pid
=0	5
=1	1, 6, 2
=2	
=3	3, 8, 4, 9
=4	

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

11

Hash Join Example

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

Memory M = 21 pages

Hash h: pid % 5

Input buffer

Output buffer

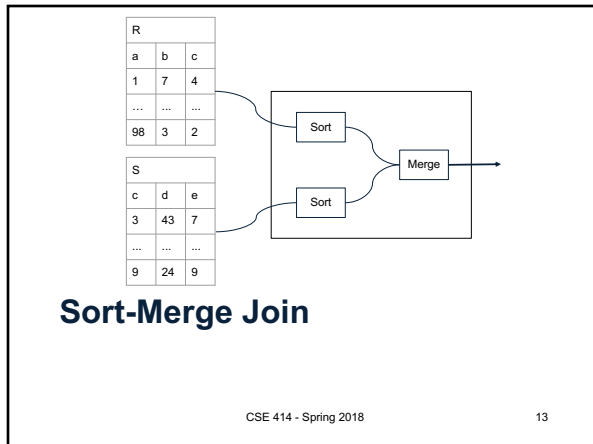
Keep going until read all of Insurance

Cost: B(R) + B(S)

bucket	pid
=0	5
=1	1, 6, 2
=2	
=3	3, 8, 4, 9
=4	

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

12



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

CSE 414 - Spring 2018 14

Sort-Merge Join Example

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

Memory M = 21 pages

Disk

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

15

Sort-Merge Join Example

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

Memory M = 21 pages *sorted*

Disk

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

16

Sort-Merge Join Example

Step 3: **Merge** Patient and Insurance

Memory M = 21 pages

Disk

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

Output buffer: 1 1

17

Sort-Merge Join Example

Step 3: **Merge** Patient and Insurance

Memory M = 21 pages

Disk

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

Output buffer: 2 2

Keep going until end of first relation

18

Index Joins

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

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

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

19

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)
    
```

CSE 414 - Spring 2018 20

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

CSE 414 - Spring 2018 21

Review: Logical vs Physical Plans

CSE 414 - Spring 2018 22

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

Review: Physical Query Plan 1

(On the fly) π_{sname}

(On the fly) $\sigma_{scity='Seattle' \text{ and } sstate='WA' \text{ and } pno=2}$

(Nested loop)

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'
    
```

23

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

Review: Physical Query Plan 2

(On the fly) π_{sname}

(On the fly) $\sigma_{scity='Seattle' \text{ and } sstate='WA' \text{ and } pno=2}$

(Hash join)

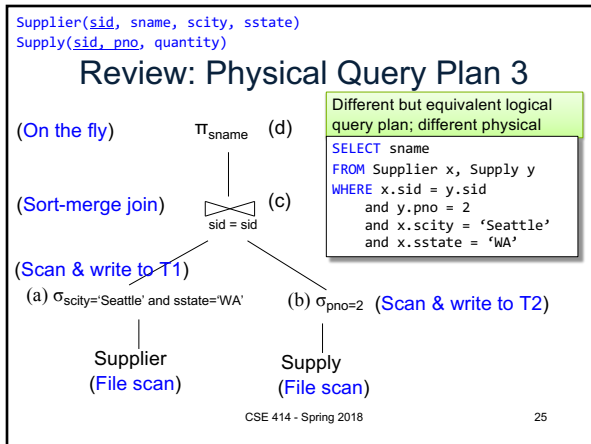
Supplier (File scan) Supply (File scan)

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'
    
```

24



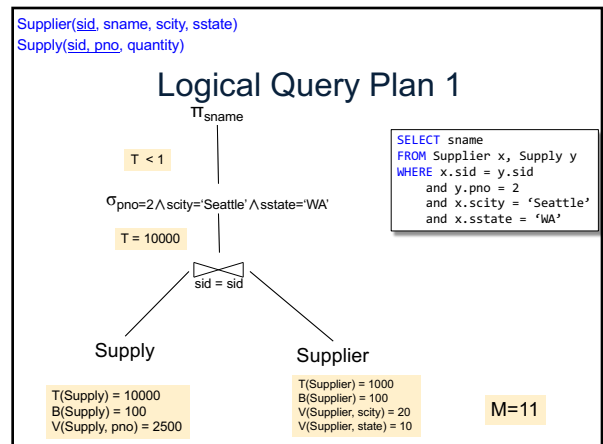
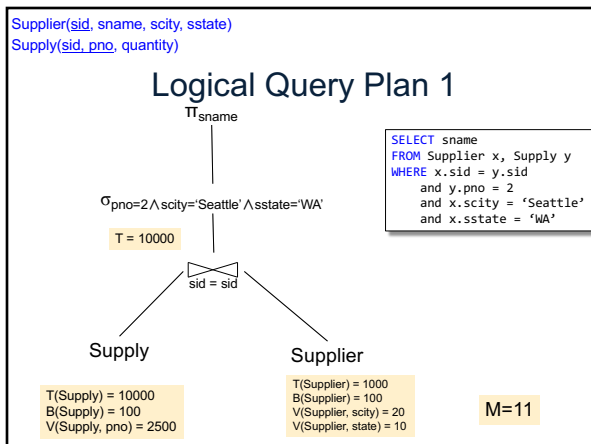
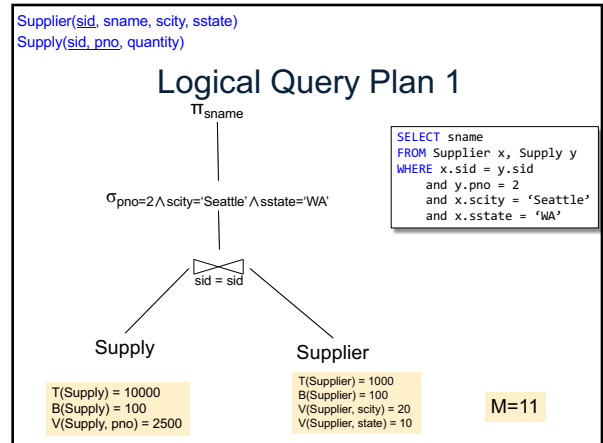
Query Optimization: Overview

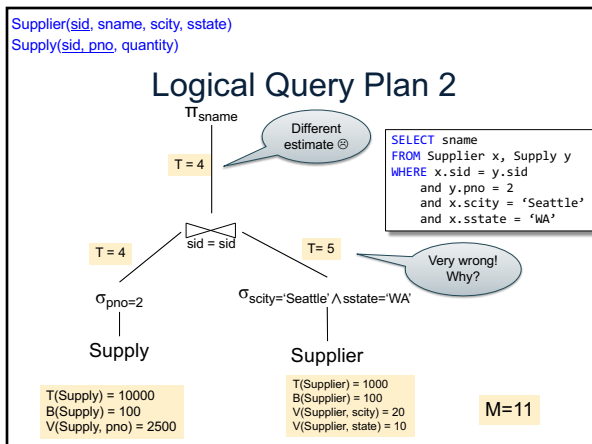
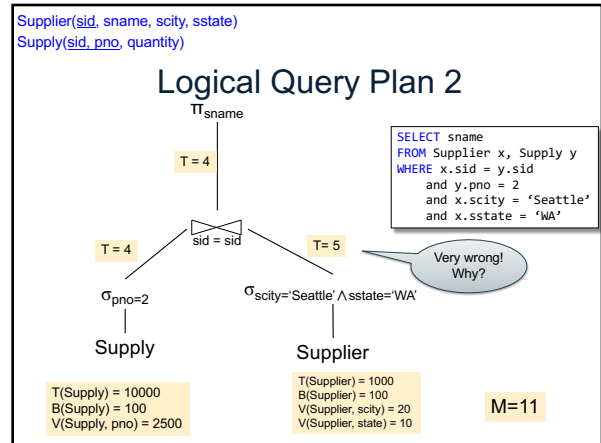
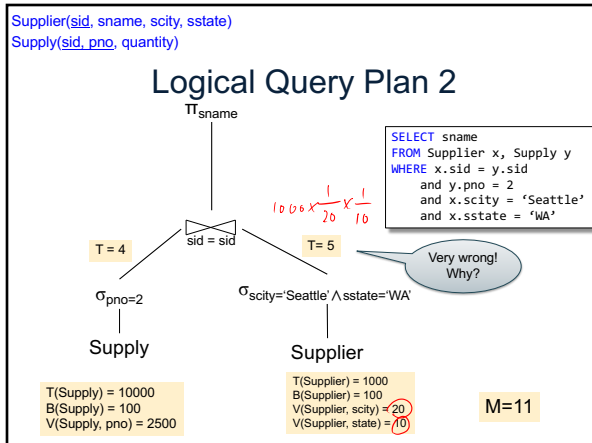
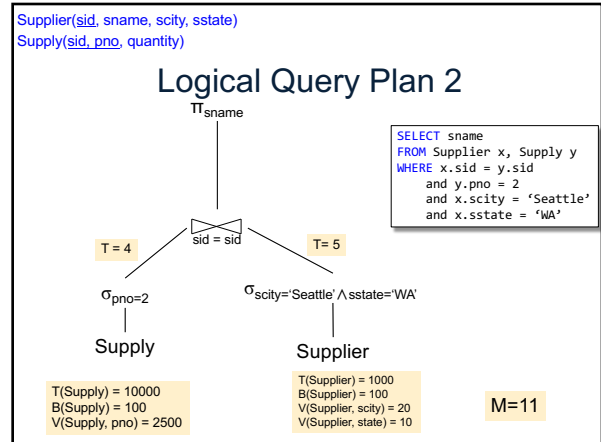
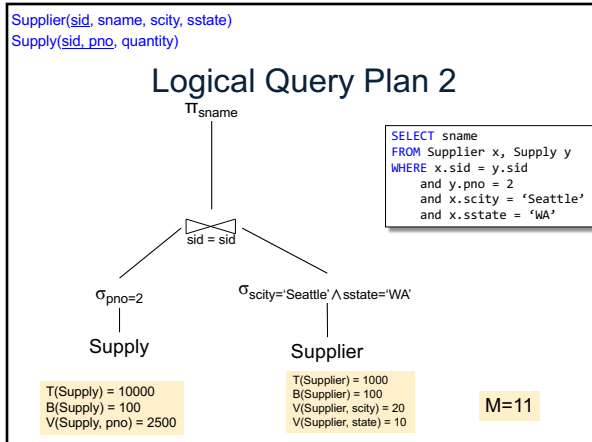
- Compute cost of each operator, which depends on:
 - Table statistics (# of tuples produced)
 - Algorithm used to implement each operator
- Cost of a physical plan = sum(each operator cost)
- Cost each plan and choose the one with lowest cost

CSE 414 - Spring 2018 26

Estimating Table Statistics

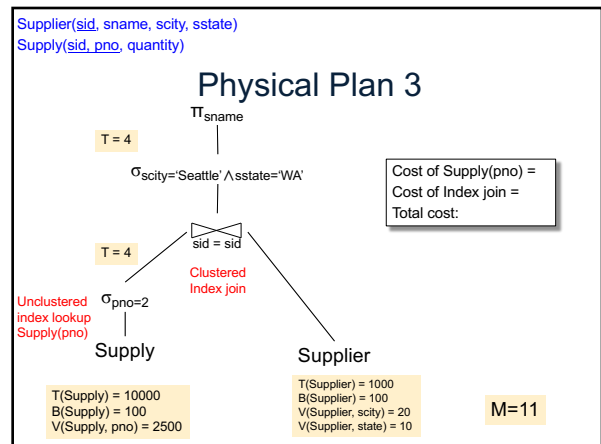
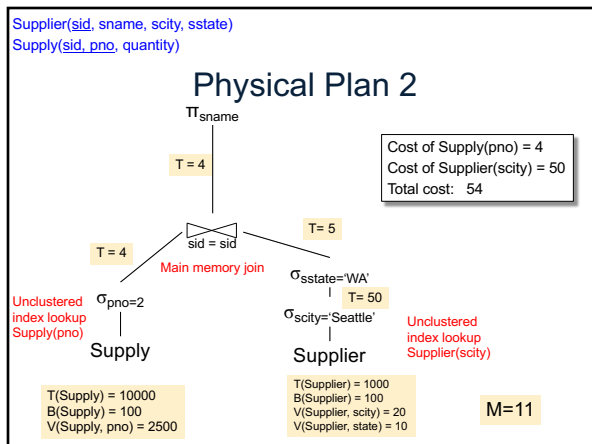
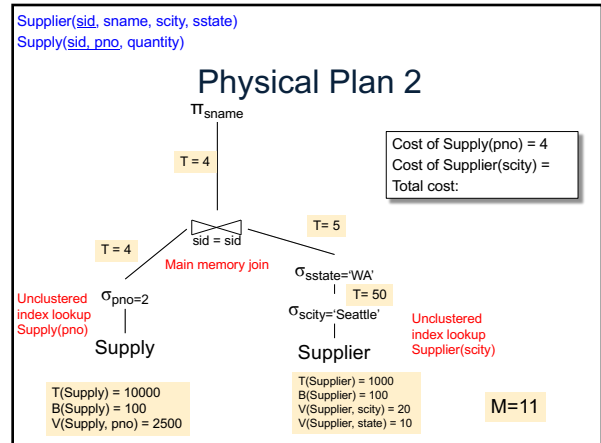
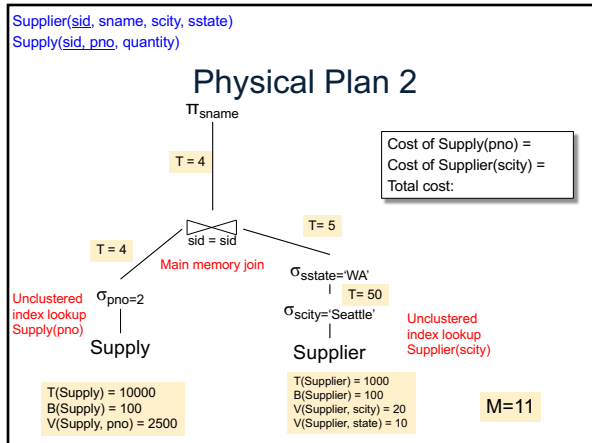
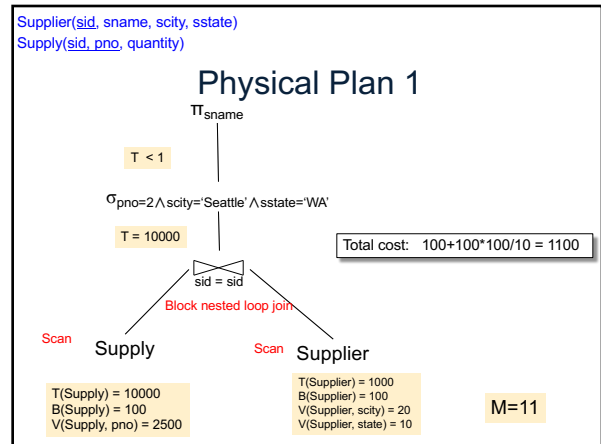
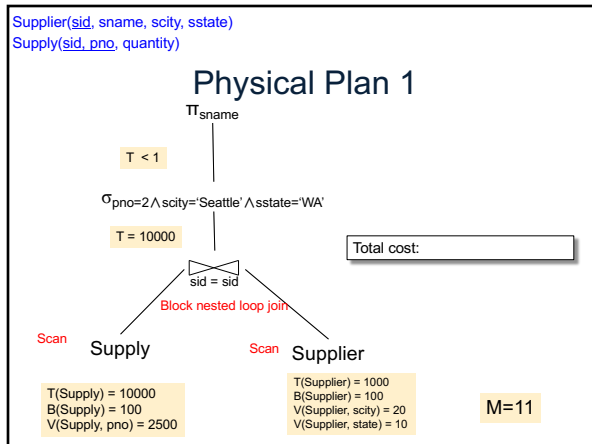
CSE 414 - Spring 2018 27

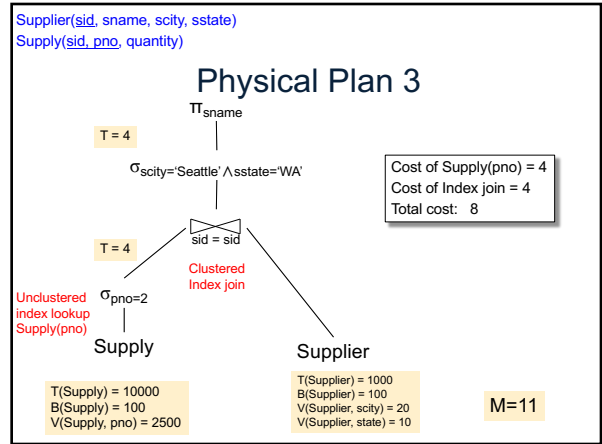
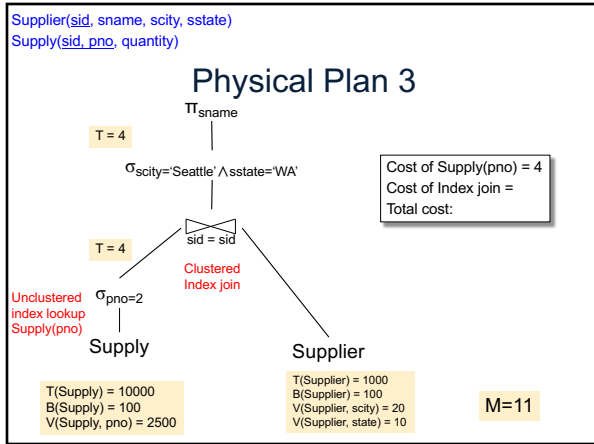




Computing Plan Costs

CSE 414 - Spring 2018 36





- ### 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
- CSE 414 - Spring 2018 45