

Introduction to Database Systems CSE 414

Lecture 27: More Operator Costs

CSE 414 - Spring 2018

1

Announcements

- HW8 and WQ7 both due tonight!
- Please fill out course evals online!
- Last lecture on Friday

CSE 414 - Spring 2018

2

Final Exam

- Thursday 6/7, 2:30-4:20pm
- Location: here
- 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

CSE 414 - Spring 2018

3

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

4

Big Picture

Why do we care about all these internal details?

5

Cost Parameters

- Cost = I/O + CPU + Network BW
 - We will focus on I/O in this class
- Parameters (a.k.a. statistics):
 - $B(R)$ = # 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 $\leq T(R)$

- DBMS collects **statistics** about base tables
must infer them for intermediate results

6

Join Algorithms

- Nested loop join (short review)
- Hash join
- Sort-merge join

CSE 414 - Spring 2018 7

Nested Loop Joins (review)

CSE 414 - Spring 2018 8

Nested Loop Joins

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

```

for each tuple t1 in R do
  for each tuple t2 in S do
    if t1 and t2 join then output (t1,t2)
        
```

What is the Cost?

CSE 414 - Spring 2018 9

Nested Loop Joins

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

```

for each tuple t1 in R do
  for each tuple t2 in S do
    if t1 and t2 join then output (t1,t2)
        
```

What is the Cost?

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

CSE 414 - Spring 2018 10

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)
        
```

What is the Cost?

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

CSE 414 - Spring 2018 11

Page-at-a-time Refinement

CSE 414 - Spring 2018 12

Page-at-a-time Refinement

Do any pairs of these join?

1	2
4	3

Input buffer for Patient

Input buffer for Insurance

Output buffer

Disk

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

13

Page-at-a-time Refinement

Do any pairs of these join?

1	2
2	8

Input buffer for Patient

Input buffer for Insurance

Output buffer

Disk

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

14

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

Hash Join

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

→

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

15

CSE 414 - Spring 2018

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

16

CSE 414 - Spring 2018

Hash Join Example

Patient(pid, name, address)
 Insurance(pid, provider, policy_nb)
 Patient \bowtie 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

17

Hash Join Example

Patient \bowtie Insurance

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

Some large-enough #

This is one page with two tuples

18

Hash Join Example

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

Memory M = 21 pages

Hash h: pid % 5

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

Input buffer

Disk

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

19

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: 2 4

Output buffer: 2 2

Write to disk or pass to next operator

Disk

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

20

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: 2 4

Output buffer: 4 4

Disk

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

21

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: 4 3

Output buffer: 4 4

Keep going until read all of Insurance

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

Disk

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

22

Sort-Merge Join

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

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

Sort

Merge

23

CSE 414 - Spring 2018

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

24

CSE 414 - Spring 2018

Sort-Merge Join Example

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

Memory M = 21 pages

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

Disk

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

25

Sort-Merge Join Example

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

Memory M = 21 pages

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

Disk

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

26

Sort-Merge Join Example

Step 3: **Merge** Patient and Insurance

Memory M = 21 pages

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

Output buffer

Disk

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

27

Sort-Merge Join Example

Step 3: **Merge** Patient and Insurance

Memory M = 21 pages

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

Output buffer

Keep going until end of first relation

Disk

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

28

Index Joins

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

...

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

29


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

Cost:

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



CSE 414 - Spring 2018 30

Index Nested Loop Join

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

Still have to scan in R

Why is the multiplier term T(R)?

What does $1/V(S,a)$ represent?

T(R) must be used because we cannot assume that a whole block of R (B(R)) will have the same attribute to join on, and thus use the same index access on S for.

$1/V(S,a)$ represents the nature of the B+ Tree index. We are only scanning as much as we need. Note that the performance of the index join will decrease as V decreases.

31

Index Nested Loop Join

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

Why did this change from B(R) to T(R)?

Remember that tuples are stored on contiguous blocks. In a clustered index from before we know we can scan a single chunk of the disk to get the entire desired range. In an unclustered index we no longer can assume contiguous access. Thus we estimate that every tuple needs its own I/O operation.

32

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

Generating Query Plans (review)

CSE 414 - Spring 2018 33

Review: Logical vs Physical Plans

- Logical plans:
 - Created by the parser from the input SQL text
 - Expressed as a relational algebra tree
 - Each SQL query has many possible logical plans
- Physical plans:
 - Goal is to choose an efficient implementation for each operator in the RA tree
 - Each logical plan has many possible physical plans

CSE 414 - Spring 2018 34

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

Review: 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"

CSE 414 - Spring 2018 35

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

Review: Physical Query Plan 1

(On the fly)

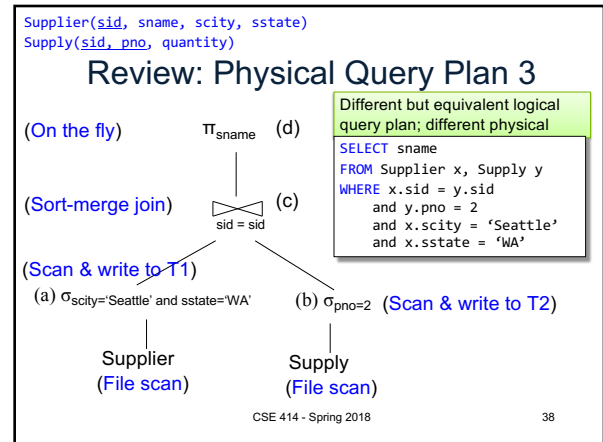
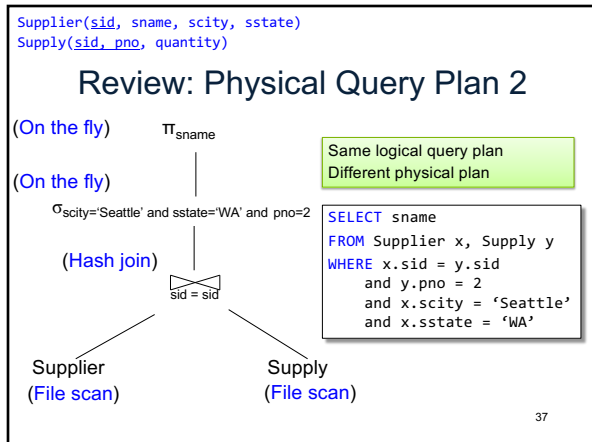
(On the fly)

(Nested loop)

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'
```

CSE 414 - Spring 2018 36



- ### Query Optimization: Overview
- Compute cost of each operator
 - This depends on:
 - Table statistics (# of tuples etc)
 - Algorithm used
 - Cost of a physical plan = sum(each operator cost)
 - Cost each plan and choose the one with lowest cost
- CSE 414 - Spring 2018 39