

Lecture 25:

Monday, December 27, 2002

1

Administrative

- Homework 5 is due on Monday, 12/9
- Project demos will be on Tuesday 12/10
 - 10am – 12pm: 6teams
 - 2pm – 4pm: 6teams
 - 4pm – 6pm: 6teams
- Please send me email for an appointment
 - First come first served...

2

Outline

- Cost-based optimization: 16.5, 16.6
- Completing the physical query plan: 16.7
- Cost estimation: 16.4 (will continue next time)

3

Cost-based Optimizations

- Main idea: apply algebraic laws, until estimated cost is minimal
- Practically: start from partial plans, introduce operators one by one
 - Will see in a few slides
- Problem: there are too many ways to apply the laws, hence too many (partial) plans

4

Cost-based Optimizations

Approaches:

- **Top-down**: the partial plan is a top fragment of the logical plan
- **Bottom up**: the partial plan is a bottom fragment of the logical plan

5

Search Strategies

- **Branch-and-bound**:
 - Remember the cheapest complete plan P seen so far and its cost C
 - Stop generating partial plans whose cost is $> C$
 - If a cheaper complete plan is found, replace P, C
- **Hill climbing**:
 - Remember only the cheapest partial plan seen so far
- **Dynamic programming**:
 - Remember the all cheapest partial plans

6

Dynamic Programming

Unit of Optimization: select-project-join


- Push selections down, pull projections up

7

Join Trees

• $R1 \bowtie R2 \bowtie \dots \bowtie Rn$

- Join tree:




- A plan = a join tree
- A partial plan = a subtree of a join tree

8

Types of Join Trees


- Left deep:



9

Types of Join Trees


- Bushy:



10

Types of Join Trees

- Right deep:



11

Problem

- Given: a query $R1 \bowtie R2 \bowtie \dots \bowtie Rn$
- Assume we have a function $cost()$ that gives us the cost of every join tree
- Find the best join tree for the query

12

Dynamic Programming

- Idea: for each subset of $\{R_1, \dots, R_n\}$, compute the best plan for that subset
- In increasing order of set cardinality:
 - Step 1: for $\{R_1\}, \{R_2\}, \dots, \{R_n\}$
 - Step 2: for $\{R_1, R_2\}, \{R_1, R_3\}, \dots, \{R_{n-1}, R_n\}$
 - ...
 - Step n: for $\{R_1, \dots, R_n\}$
- It is a bottom-up strategy
- A subset of $\{R_1, \dots, R_n\}$ is also called a *subquery*

13

Dynamic Programming

- For each subquery $Q \subseteq \{R_1, \dots, R_n\}$ compute the following:
 - $\text{Size}(Q)$
 - A best plan for Q : $\text{Plan}(Q)$
 - The cost of that plan: $\text{Cost}(Q)$

14

Dynamic Programming

- **Step 1:** For each $\{R_i\}$ do:
 - $\text{Size}(\{R_i\}) = B(R_i)$
 - $\text{Plan}(\{R_i\}) = R_i$
 - $\text{Cost}(\{R_i\}) = (\text{cost of scanning } R_i)$

15

Dynamic Programming

- **Step i:** For each $Q \subseteq \{R_1, \dots, R_n\}$ of cardinality i do:
 - Compute $\text{Size}(Q)$ (later...)
 - For every pair of subqueries Q', Q'' s.t. $Q = Q' \cup Q''$ compute $\text{cost}(\text{Plan}(Q') \bowtie \text{Plan}(Q''))$
 - $\text{Cost}(Q) =$ the smallest such cost
 - $\text{Plan}(Q) =$ the corresponding plan

16

Dynamic Programming

- Return $\text{Plan}(\{R_1, \dots, R_n\})$

17

Dynamic Programming

To illustrate, we will make the following simplifications:

- $\text{Cost}(P_1 \bowtie P_2) = \text{Cost}(P_1) + \text{Cost}(P_2) + \text{size}(\text{intermediate result}(s))$
- Intermediate results:
 - If P_1 is a join, then the size of the intermediate result is $\text{size}(P_1)$, otherwise the size is 0
 - Similarly for P_2
- Cost of a scan = 0

18

Dynamic Programming

- Example:
- $\text{Cost}(R5 \bowtie R7) = 0$ (no intermediate results)
- $\text{Cost}((R2 \bowtie R1) \bowtie R7)$
 $= \text{Cost}(R2 \bowtie R1) + \text{Cost}(R7) + \text{size}(R2 \bowtie R1)$
 $= \text{size}(R2 \bowtie R1)$

19

Dynamic Programming

- Relations: R, S, T, U
- Number of tuples: 2000, 5000, 3000, 1000
- Size estimation: $T(A \bowtie B) = 0.01 * T(A) * T(B)$

20

Subquery	Size	Cost	Plan
RS			
RT			
RU			
ST			
SU			
TU			
RST			
RSU			
RTU			
STU			
RSTU			

21

Subquery	Size	Cost	Plan
RS	100k	0	RS
RT	60k	0	RT
RU	20k	0	RU
ST	150k	0	ST
SU	50k	0	SU
TU	30k	0	TU
RST	3M	60k	(RT)S
RSU	1M	20k	(RU)S
RTU	0.6M	20k	(RU)T
STU	1.5M	30k	(TU)S
RSTU	30M	60k+50k=110k	(RT)(SU)

22

Dynamic Programming

- Summary: computes optimal plans for subqueries:
 - Step 1: {R1}, {R2}, ..., {Rn}
 - Step 2: {R1, R2}, {R1, R3}, ..., {Rn-1, Rn}
 - ...
 - Step n: {R1, ..., Rn}
- We used naïve size/cost estimations
- In practice:
 - more realistic size/cost estimations (next time)
 - heuristics for Reducing the Search Space
 - Restrict to left linear trees
 - Restrict to trees “without cartesian product”
 - need more than just one plan for each subquery:
 - “interesting orders”

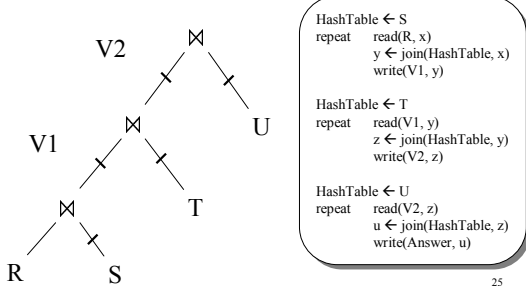
23

Completing the Physical Query Plan

- Choose algorithm to implement each operator
 - Need to account for more than cost:
 - How much memory do we have ?
 - Are the input operand(s) sorted ?
- Decide for each intermediate result:
 - To materialize
 - To pipeline

24

Materialize Intermediate Results Between Operators



25

Materialize Intermediate Results Between Operators

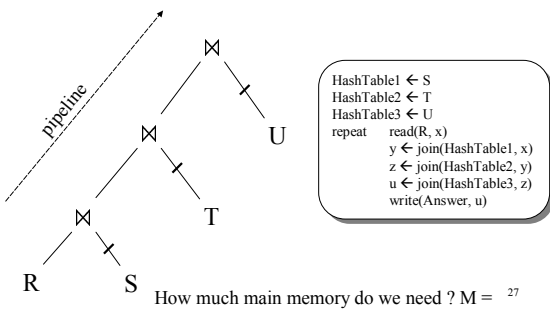
Question in class

Given $B(R)$, $B(S)$, $B(T)$, $B(U)$

- What is the total cost of the plan ?
 - Cost =
- How much main memory do we need ?
 - M =

26

Pipeline Between Operators



How much main memory do we need ? M = 27

27

Pipeline Between Operators

Question in class

Given $B(R)$, $B(S)$, $B(T)$, $B(U)$

- What is the total cost of the plan ?
 - Cost =
- How much main memory do we need ?
 - M =

28


Completing the Physical Query Plan

- Choose algorithm to implement each operator
 - Need to account for more than cost:
 - How much memory do we have ?
 - Are the input operand(s) sorted ?
- Decide for each intermediate result:
 - To materialize
 - To pipeline

29

Example 16.36


- Logical plan is:



- Main memory $M = 101$ buffers

30

Example 16.36




Naïve evaluation:

- 2 partitioned hash-joins
- Cost $3B(R) + 3B(S) + 4k + 3B(U) = 75000 + 4k$

31

Example 16.36




Smarter:

- Step 1: hash R on x into 100 buckets, each of 50 blocks; to disk
- Step 2: hash S on x into 100 buckets; to disk
- Step 3: read each R_i in memory (50 buffer) join with S_i (1 buffer); hash result on y into 50 buckets (50 buffers) -- here we pipeline
- Cost so far: $3B(R) + 3B(S)$

32

Example 16.36




Continuing:

- How large are the 50 buckets on y? Answer: $k/50$.
- If $k \leq 50$ then keep all 50 buckets in Step 3 in memory, then:
- Step 4: read U from disk, hash on y and join with memory
- Total cost: $3B(R) + 3B(S) + B(U) = 55,000$

33

Example 16.36




Continuing:

- If $50 < k \leq 5000$ then send the 50 buckets in Step 3 to disk
 - Each bucket has size $k/50 \leq 100$
- Step 4: partition U into 50 buckets
- Step 5: read each partition and join in memory
- Total cost: $3B(R) + 3B(S) + 2k + 3B(U) = 75,000 + 2k$

34

Example 16.36



Continuing:

- If $k > 5000$ then materialize instead of pipeline
- 2 partitioned hash-joins
- Cost $3B(R) + 3B(S) + 4k + 3B(U) = 75000 + 4k$

35

Example 16.36

Summary:

- If $k \leq 50$, cost = 55,000
- If $50 < k \leq 5000$, cost = 75,000 + 2k
- If $k > 5000$, cost = 75,000 + 4k

36

Estimating Sizes

- Need size in order to estimate cost
- Example:
 - Cost of partitioned hash-join $E1 \bowtie E2$ is $3B(E1) + 3B(E2)$
 - $B(E1) = T(E1)/\text{block size}$
 - $B(E2) = T(E2)/\text{block size}$
 - So, we need to estimate $T(E1), T(E2)$

37

Estimating Sizes

Estimating the size of a projection

- Easy: $T(\Pi_L(R)) = T(R)$
- This is because a projection doesn't eliminate duplicates

38

Estimating Sizes

Estimating the size of a selection

- $S = \sigma_{A=c}(R)$
 - $T(S)$ can be anything from 0 to $T(R) - V(R,A) + 1$
 - Mean value: $T(S) = T(R)/V(R,A)$
- $S = \sigma_{A < c}(R)$
 - $T(S)$ can be anything from 0 to $T(R)$
 - Heuristics: $T(S) = T(R)/3$

39

Estimating Sizes

Estimating the size of a natural join, $R \bowtie_A S$

- When the set of A values are disjoint, then $T(R \bowtie_A S) = 0$
- When A is a key in S and a foreign key in R, then $T(R \bowtie_A S) = T(R)$
- When A has a unique value, the same in R and S, then $T(R \bowtie_A S) = T(R) T(S)$

40

Estimating Sizes

Assumptions:

- Containment of values: if $V(R,A) \leq V(S,A)$, then the set of A values of R is included in the set of A values of S
 - Note: this indeed holds when A is a foreign key in R, and a key in S
- Preservation of values: for any other attribute B, $V(R \bowtie_A S, B) = V(R, B)$ (or $V(S, B)$)

41

Estimating Sizes

Assume $V(R,A) \leq V(S,A)$

- Then each tuple t in R joins *some* tuple(s) in S
 - How many?
 - On average $S/V(S,A)$
 - t will contribute $S/V(S,A)$ tuples in $R \bowtie_A S$
- Hence $T(R \bowtie_A S) = T(R) T(S) / V(S,A)$

In general: $T(R \bowtie_A S) = T(R) T(S) / \max(V(R,A), V(S,A))$

42

Estimating Sizes

Example:

- $T(R) = 10000$, $T(S) = 20000$
- $V(R,A) = 100$, $V(S,A) = 200$
- How large is $R \bowtie_A S$?

Answer: $T(R \bowtie_A S) = 10000 \cdot 20000 / 200 = 1M$

43

Estimating Sizes

Joins on more than one attribute:

- $T(R \bowtie_{A,B} S) =$

$$T(R) T(S) / \max(V(R,A), V(S,A)) \max(V(R,B), V(S,B))$$

44

Histograms

- Statistics on data maintained by the RDBMS
- Makes size estimation much more accurate (hence, cost estimations are more accurate)

45

Histograms

Employee(ssn, name, salary, phone)

- Maintain a histogram on salary:

Salary:	0..20k	20k..40k	40k..60k	60k..80k	80k..100k	> 100k
Tuples	200	800	5000	12000	6500	500

- $T(\text{Employee}) = 25000$, but now we know the distribution

46

Histograms

Ranks(rankName, salary)

- Estimate the size of $\text{Employee} \bowtie_{\text{Salary}} \text{Ranks}$

Employee	0..20k	20k..40k	40k..60k	60k..80k	80k..100k	> 100k
	200	800	5000	12000	6500	500

Ranks	0..20k	20k..40k	40k..60k	60k..80k	80k..100k	> 100k
	8	20	40	80	100	2

47

Histograms

- Assume:

- $V(\text{Employee}, \text{Salary}) = 200$
- $V(\text{Ranks}, \text{Salary}) = 250$

- Then $T(\text{Employee} \bowtie_{\text{Salary}} \text{Ranks}) =$
 $= \sum_{i=1,6} T_i T_i' / 250$
 $= (200 \times 8 + 800 \times 20 + 5000 \times 40 +$
 $12000 \times 80 + 6500 \times 100 + 500 \times 2) / 250$
 $= \dots$

48