









SELECT S.sname Example FROM Reserves R, Sailors S WHERE R.sid=S.sid AND R.bid=100 AND S.rating>5 ∏ sname ∏ sname ∏ sname $\sigma_{\rm rating > 5}$ $^{
m O}_{
m bid=100}$ $^{
m /}$ rating > 5 \bowtie sid=sid \bowtie sid=sid ਾating > 5 ∽_{bid=100} \searrow [♂]bid=100 Sailors sid=sid Reserves Sailors Reserves Reserves Sailors CS5208: Query Optimization





























































































Review

- Consider the join R JOIN(R.a=S.b) S, given the following information about the relations to be joined. The cost metric is the number of page I/Os, and the cost of writing out the result should be ignored.
 - R contains 10,000 tuples and has 10 tuples per page.
 - S contains 20,000 tuples and has 10 tuples per page.
 - S.b is the primary key for S.
 - Both relations are stored as simple heap files.
 - 102 buffer pages are available (inclusive of input/output buffers).
- What is the cost of joining R and S using a block nested-loops join algorithm? What is the minimum number of buffer pages required for this cost to remain unchanged?
- What is the cost of joining R and S using a sort-merge join algorithm? What is the minimum number of buffer pages required for this cost to remain unchanged?

CS5208: Query Optimization

Review

- Block Nested Loops Join.
 - Using R as the outer relation, and 1 page for input and output buffer.
 - $\cos t = 10,000/10 + (10,000/10)/100 \times 20,000/10 = 21,000$
 - minimum number of buffer page s= 102 (no change)
- Sort-merge Join
 - Each relation needs 2 passes to sort
 - Cost to sort R = 2*2*10,000/10; cost to sort S = 2*2*20,000/10
 - Cost = 4000+8000+1000+2000 = 15,000
 - min buffer required is the same as that required to sort the larger relation, which is S. So, min buffer = 46

CS5208: Query Optimization

54



































Partial Result: U =	R1 🖂 R2	
$T(U) = 1000 \times 2000$	V(U,A) = 50	
200	V(U,B) = 100	
	V(U,C) = 300	
Z = U ⊳ R3		
$T(Z) = 1000 \times 2000 \times 3000$	V(Z,A) = 50	
200×300	V(Z,B) = 100	
	V(Z,C) = 90	
	V(Z,D) = 500	
CS5208: Query Optimization		72

















More on Randomized Techniques

- Two states are neighboring states if one move suffices to go from one state to the other
- A local minimum in the state space is a state such that its cost is lower than that of all neighboring states
- A global minimum is a state which has the lowest cost among all local minima
 - · at most one global minimum
- A move that takes one state to another state with a lower cost is called a downward move; otherwise it is an upward move
 - · in a local/global minimum, all moves are upward moves

CS5208: Query Optimization





Issues on Local Optimization	
 How is the start state obtained? The state in which we start a run. The start state of the first run is the initial state. All start states should be different. Should be obtained quickly random greedy heuristics making a number of moves from the local minimum, except that this time each move is accepted irrespective of whether it increases or decreases the cost 	
How is the local minimum detected?	
How is the stopping criterion detected?	
CS5208: Query Optimization	84







Comparison between Exhaustive, Greedy and Randomized Algorithms

- Plan quality
- Optimization overhead

Dynamic Programming (Left-Deep Trees)

- The algorithm proceeds by considering increasingly larger subsets of the set of all relations.
- Plans for a set of cardinality *i* are constructed as extensions of the best plan for a set of cardinality *i*-1
- Search space can be pruned based on the principal of optimality
 - if two plans differ only in a subplan, then the plan with the better subplan is also the better plan

CS5208: Query Optimization



Dynamic Programming (Left-Deep Trees)

- accessPlan(R) produces the best plan for relation R
- joinPlan(p1,R) extends the join plan p1 into another plan p2 in which the result of p1 is joined with R in the best possible way
- Optimal plans for subsets are stored in optplan() array and are reused rather than recomputed

CS5208: Query Optimization



Dynamic Programming Example Consider the join of 4 relations, R, S, T and U

Each table has 1000 tuples

Assume intermediate result size (tuples) as cost metrics

R(a,b)	S(b,c)	T(c,d)	U(d,a)
V(R,a)=100			V(U,a)=50
V(R,b)=200	V(S,b)=100		
	V(S,c)=500	V(T,c)=20	
		V(T,d)=50	V(U,d)=1000

CS5208: Query Optimization

		E.	xample ((Cont)
	{R}	{S}	{T}	{U}
Size	1,000	1,000	1,000	1,000
Cost	0	0	0	0
BestPlan	R	S	Т	U

		Example (Cont)				
	$\{R,S\}$	$\{R,T\}$	{ R , U }	$\{S,T\}$	$\{S,U\}$	{ T , U }
Size	5,000	1 M	10,000	2,000	1 M	1,000
Cost	0	0	0	0	0	0
BestPlan	R⋈S	RMT	R⊠U	SMT	S 🖂 U	ΤᢂU
What about S \bowtie R since its cost is also 0??						
CS5208: Query Opti	imization					95

		Example (Cont)		
	$\{R,S,T\}$	{ R , S , U }	{ R , T , U }	{ S , T , U }
Size	10,000	50,000	10,000	2,000
Cost	2,000	5,000	1,000	1,000
BestPlan	$(S \bowtie T) \bowtie R$		$(T \bowtie U) \bowtie R$	
		$(\mathbf{R} \bowtie \mathbf{S}) \bowtie \mathbf{U}$		$(T \bowtie U) \bowtie S$
CS5208: Query Op	timization			96

	Example (Cont)
Grouping	Cost
$((S \bowtie T) \bowtie R) \bowtie U)$	12,000
$((\mathbf{R} \bowtie \mathbf{S}) \bowtie \mathbf{U}) \bowtie \mathbf{T})$	55,000
$((T \bowtie U) \bowtie R) \bowtie S)$	11,000
$((T \bowtie U) \bowtie S) \bowtie R)$	3,000
$(T \bowtie U) \bowtie (R \bowtie S)$	6,000
$(\mathbf{R} \bowtie \mathbf{T}) \bowtie (\mathbf{S} \bowtie \mathbf{U})$	2M
$(S \bowtie T) \bowtie (R \bowtie U)$	12,000
CS5208: Query Optimization	9

г

	Example (Cont)
Grouping	Cost
$((S \bowtie T) \bowtie R) \bowtie U)$	12,000
$((\mathbf{R} \bowtie \mathbf{S}) \bowtie \mathbf{U}) \bowtie \mathbf{T})$	55,000
$((T \bowtie U) \bowtie R) \bowtie S)$	11,000
$((\mathbf{T} \bowtie \mathbf{U}) \bowtie \mathbf{S}) \bowtie \mathbf{R})$	3,000
$(T \bowtie U) \bowtie (R \bowtie S)$	6,000
$(\mathbf{R} \bowtie \mathbf{T}) \bowtie (\mathbf{S} \bowtie \mathbf{U})$	2M
$(S \bowtie T) \bowtie (R \bowtie U)$	12,000
	i
CS5208: Query Optimization	98

.



