Review (Past year question)

- Consider the relations R1(A,B,C), R2(C,D,E) and R3(E,F), with primary keys A, C, E respectively. Assume that R1 has 10000 tuples, R2 has 15000 tuples and R3 has 7500 tuples. For simplicity, assume that all tuples (including the query result) have the same size, and that each page can contain 10 tuples of R. Consider the query: R1 JOIN R2 JOIN R3. Assume that all attributes are of the same size, and any join output will include all attributes of all relations. Further, assume records do not span pages. Assuming the data are uniformly distributed, estimate the result size of the query.
- List all possible plans assuming only left-deep search space is considered (assuming only one join method). You may assume that cross product are to be avoided.
- Compute the cost for each of the above plans you listed to determine the optimal plan. For simplicity, you may assume that only the nestedblock join is supported, the buffer size is 100 pages, and all intermediate results are to be stored in secondary storage.



Review (Past year question)

- Assume Left Deep Tree plans and one join method. In total, there are 6 possible plans, but since cross products are not permitted, we end up with 4 plans
 - (R1 JOIN R2) JOIN R3
 - (R2 JOIN R1) JOIN R3
 - (R2 JOIN R3) JOIN R1
 - (R3 JOIN R2) JOIN R1

R1

R2













Atomicity and Durability A.C.I.D. • A transaction ends in one of two ways: • *commit* after completing all its actions • "commit" is a contract with the caller of the DB • *abort* (or be aborted by the DBMS) after executing some actions. • Or system crash while the xact is in progress; treat as abort. • Two important properties for a transaction: • Atomicity : Either execute all its actions, or none of them • *Durability* : The effects of a committed xact must survive failures. • DBMS ensures the above by *logging* all actions (Recovery): • Undo the actions of aborted/failed transactions. • *Redo* actions of committed transactions not yet propagated to disk when system crashes. CS5208 - Concurrency Control













Schedule A: S	Serial Schedule		
		А	В
T1	T2	25	25
Read(A); $A \leftarrow A+100$ Write(A); Read(B); $B \leftarrow B+100$;		125	
Write(B);	Read(A);A \leftarrow A×2; Write(A); Read(B);B \leftarrow B×2;	250	125
	Write(B);		250
		250	250
CS5208 – Concurrency Control			16

Schedule B				
		А	В	
T1	T2	25	25	
Read(A); $A \leftarrow A+100$				
Write(A);		125		
	$Read(A); A \leftarrow A \times 2;$			
	Write(A);	250		
Read(B); $B \leftarrow B+100$;				
Write(B);			125	
	Read(B); $B \leftarrow B \times 2$;			
	Write(B);		250	
		250	250	
CS5208 – Concurrency Control				17

Schedule C			
		А	В
T1	T2	25	25
Read(A); $A \leftarrow A+100$			
Write(A);		125	
	Read(A); $A \leftarrow A \times 2$;		
	Write(A);	250	
	Read(B); $B \leftarrow B \times 2$;		
	Write(B);		50
Read(B); $B \leftarrow B+100$;			
Write(B);			150
		250	150
CS5208 – Concurrency Control			18

Schedule D	Same as Schedule C but with new T2'		
		А	В
T1	T2'	25	25
Read(A); $A \leftarrow A+100$			
Write(A);		125	
	$Read(A); A \leftarrow A \times 1;$		
	Write(A);	125	
	Read(B); $B \leftarrow B \times 1$;		
	Write(B);		25
Read(B); $B \leftarrow B+100$;			
Write(B);			125
		125	125
CS5208 – Concurrency Control			19

















Conflict-Serializability is NOT necessary for Serializability

- S1: w1(Y); w1(X); w2(Y); w2(X); w3(X)
 - Serial schedule
- S2: w1(Y); w2(Y); w2(X); w1(X); w3(X)
 - Serializable schedule since effect is same as S1
- S1, S2 not conflict equivalent

CS5208 - Concurrency Control









Theorem

33

 $P(S_1)$ acyclic $\iff S_1$ conflict serializable

S₁, S₂ conflict equivalent \Rightarrow P(S₁)=P(S₂) ??? P(S₁)=P(S₂) \neq S₁, S₂ conflict equivalent ???

CS5208 – Concurrency Control









37

Rule #1: Well-formed transactions Ti: ... $li(A) \dots pi(A) \dots ui(A) \dots$

Rule #2: Legal scheduler $S = \dots li(A) \xrightarrow{ui(A)} ui(A) \dots no lj(A)$

CS5208 - Concurrency Control



















CS5208 - Concurrency Control









Rule # 3 2PL transactions

No change except for upgrades:

(I) If upgrade gets more locks $(e.g., S \rightarrow \{S, X\})$ then no change! (II) If upgrade releases read (shared) lock (e.g., $S \rightarrow X$) - can be allowed in growing phase

CS5208 - Concurrency Control



52

















<u>Note:</u> object A may be locked in different modes at the same time...

S1=...l-S1(A)...l-S2(A)...
$$\begin{cases} 1-X_3(A)...?\\ 1-S_3(A)...?\\ 1-U_3(A)...? \end{cases}$$

To grant a lock in mode t, mode t must be compatible with all currently held locks on object

CS5208 - Concurrency Control

































Parent locked in	Child can be locked in	P
IS	IS, S IS, S, IX, X, SIX	
$\frac{1X}{S}$	[S, IS] not necessary	(C)
SIX	X, IX, [SIX]	
Х	none	

<u>Rules</u>

- (1) Follow multiple granularity comp function
- (2) Lock root of tree first, any mode
- (3) Node Q can be locked by Ti in S or IS only if parent(Q) locked by Ti in IX or IS
- (4) Node Q can be locked by Ti in X,SIX,IX only if parent(Q) locked by Ti in IX,SIX
- (5) Ti is two-phase
- (6) Ti can unlock node Q only if none of Q's children are locked by Ti

CS5208 – Concurrency Control









































































- Have studied lock-based CC mechanisms
 - 2 PL
 - Multiple granularity
 - Deadlock
- Did not cover non-locking based CC (timestamp/validation-based) schemes

CS5208 – Concurrency Control