

CS 5300 module6

Name

Student ID

Problem #1 (10 Points)

a) Consider the three transactions T_1 , T_2 , and T_3 , and the schedules S_1 and S_2 .

T_1 : $r_1(X); r_1(Z); w_1(X);$

T_2 : $r_2(Y); r_2(Z); w_2(Y);$

T_3 : $w_3(X); r_3(Y); w_3(Z);$

S_1 : $w_3(X); r_1(X); r_3(Y); r_2(Y); w_3(Z); r_2(Z); r_1(Z); w_2(Y); w_1(X);$

S_2 : $w_3(X); r_3(Y); w_3(Z); r_2(Y); r_2(Z); w_2(Y); r_1(X); r_1(Z); w_1(X);$

Fill out the **best response** and **justify** your answer:

1) S_2 is a serial schedule (why)?

Actions of different transactions are not interleaved

2) S_1 and S_2 are equivalent schedules (why)?

1) Every read operation reads from the same write operation in both schedule, and

2) Both schedules have the same final writes (results).

3) S_1 is a serializable schedule (why)?

Because it is equivalent to S_2

or

Its graph does not have a cycle (it is acyclic).

b) Improper scheduling of transactions may result in unpleasant course of events such as, “Dirty read”, “unrepeatable read”, ... For the following transactions and schedules, name the nature of the anomalies (justify your answer):

<p>T1: Read (x); $x := x - N$; Write (x); Read (y); $y := y + N$; Write (y);</p>	<p>T2 Read (x); $x := x + M$; Write (x);</p>
--	--

a) Schedule1:

<p>T1 Read (x); $x := x - N$; Write (x); Read (y); $y := y + N$; Write (y);</p>	<p>T2 Read (x); $x := x + M$; Write (x);</p>
---	--

Lost update; X has an incorrect value since its update by T1 is lost. Assume initially X is 80 and N is 5 and M is 4,

b) Schedule2:

<p>T1 Read (x); $x := x - N$; Write (x); Read(y); • • •</p>	<p>T2 Read (x); $x := x + M$; Write (x);</p>
--	--

This is a dirty read. Assume T1 fails after Read(y), so its effect has to be undone, however, T2 has already read the temporary incorrect value of X,

Problem #2 (8 Points)

- 1) Define the term transaction within the scope of databases.

A query does not change the data in base sources (e.g., relations), however, a transaction may do so.

It is a set of operations that transfer databases from one consistent state to another consistent state.

- 2) Relative to query processing, name issues that complicate transaction processing.

Dependence between transactions, and system failure during the execution of a request.

- 3) ACID property stands for:

Atomicity, consistency, isolation, and durability

- 4) Define the term concurrency control:

The concurrency control is a protocol to generate a serializable schedule for execution of a sequence of transactions

Problem #3 (10 Points)

a) Define the following:

a. A schedule is **Recoverable** if:

A recoverable schedule is a schedule that for each pair of transactions T_i and T_j such that T_j reads a data item previously written by T_i , T_i commits before T_j commits.

T_i writes x , T_j reads x and T_i commits before T_j .

b. A schedule is **Cascadeless** if:

A schedule is a cascadeless schedule where each pair of transactions T_i and T_j such that T_j reads a data item previously written by T_i , the commit operation of T_i appears before read operation of T_j .

T_i writes x and T_i commits, before T_j reads x .

c. A schedule is **Strict** if:

A schedule is a strict schedule in which transactions cannot read or write an item X until the last transaction that wrote X is committed (or aborted).

b) Consider the following schedules. Determine whether each schedule is recoverable, cascadeless, or strict (**make sure to justify and support your answers**).

S₁: r₁(x); r₂(z); r₁(z); r₃(x); r₃(y); w₁(x); c₁; w₃(y); c₃; r₂(y); w₂(z); w₂(y); c₂;

S₂: r₁(x); r₂(z); r₁(z); r₃(x); r₃(y); w₁(x); w₃(y); r₂(y); w₂(z); w₂(y); c₁; c₂; c₃;

S₃: r₁(x); r₂(z); r₃(x); r₁(z); r₂(y); r₃(y); w₁(x); c₁; w₂(z); w₃(y); w₂(y); c₃; c₂;

Recoverable schedule: A schedule is recoverable if the following condition is satisfied:

T_j commits after T_i if T_j has read any data item written by T_i.

- If A₁>C₃>C₂, then S₁ is recoverable because rolling back of T₁ does not affect T₂ and T₃. If C₁>A₃>C₂, S₁ is not recoverable because T₂ read the value of Y (r₂(Y)) after T₃ wrote X (w₃(Y)) and T₂ committed but T₃ rolled back. Thus, T₂ used non-existent value of Y. If C₁>C₃>A₃, then S₁ is recoverable because roll back of T₂ does not affect T₁ and T₃. Strictest condition of S₃ is C₃>C₂.
- If A₁>C₂>C₃, then S₂ is recoverable because roll back of T₁ does not affect T₂ and T₃. If C₁>A₂>C₃, then S₂ is recoverable because the roll back of T₂ will restore the value of Y that was read and written to by T₃ (w₃(Y)). It will not affect T₁. If C₁>C₂>A₃, then S₂ is not recoverable because T₃ will restore the value of Y which was not read by T₂. Strictest condition of S₄ is C₃>C₂, but it is not satisfied by S₂.
- If A₁>C₃>C₂, then S₃ is recoverable because neither T₂ nor T₃ writes to X, which is written by T₁. If C₁>A₃>C₂, then S₃ is not recoverable because T₃ will restore the value of Y, which was not read by T₂. Thus, T₂ committed with a non-existent value of Y. If C₁>C₃>A₂, then S₃ is recoverable because it will restore the value of Y to the value, which was read by T₃. Thus, T₃ committed with the right value of Y. Strictest condition of S₃ is C₃>C₂, but it is not satisfied by S₃.
- **Cascadeless schedule : A schedule is cascadeless if the following condition is satisfied:**
T_j reads X only after T_i has written to X and terminated (aborted or committed).
- Schedule S₁ is not cascadeless because T₃ reads X (r₃(X)) before T₁ commits.

- Schedule S2 is not cascadeless because T3 reads X ($r_3(X)$) before T1 commits.
- Schedule S3 is not cascadeless because T3 reads X ($r_3(X)$) before T1 commits or T2 reads Y ($r_2(Y)$) before T3 commits.
- **Strict schedule: A schedule is strict if it satisfies the following conditions:**
 - 1. Tj reads a data item X after Ti has written to X and Ti is terminated (aborted or committed)**
 - 2. Tj writes a data item X after Ti has written to X and Ti is terminated (aborted or committed)**
- Schedule S1 is not strict because T3 reads X ($r_3(X)$) before T1 has written to X ($w_1(X)$) but T3 commits after T1. In a strict schedule T3 must read X after C1.
- Schedule S2 is not strict because T3 reads X ($r_3(X)$) before T1 has written to X ($w_1(X)$) but T3 commits after T1. In a strict schedule T3 must read X after C1.
- Schedule S3 is not strict because T3 reads X ($r_3(X)$) before T1 has written to X ($w_1(X)$) but T3 commits after T1. In a strict schedule T3 must read X after C1.

Problem #4 (7 Points)

- a) Lock-based protocol has been proposed as a mechanism to allow concurrent execution of transactions, what is it (**detailed discussion, make sure to address its shortcomings**).

to ensure serializability is to require data items to be accessed in a mutual exclusive fashion — while a transaction is accessing the data item, no other transaction can access that data item, i.e., being Locked.

To access a data item, transaction T_i must first request for a lock on that data item. If the data item is already locked by another transaction in an incompatible mode, the concurrency control manager will not grant the lock until all incompatible locks held by other transactions are released.

- b) Define the term two-phase lock-based protocol, and compare and contrast it with lock-based protocol.

This protocol ensures serializability, however, it requires that each transaction issue lock and unlock requests in two phases:

- Growing Phase: A transaction may obtain locks, but may not release any lock.
- Shrinking phase: A transaction may release locks, but may not obtain any new locks.

Initially, a transaction is in the growing phase. It acquires locks as needed. Once it releases a lock, it enters the shrinking phase, and can issue no more lock requests.