

CSC 742

Database Management Systems

Topic #12: Transaction Concepts

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 1

What is a transaction?

- Informally, a transaction is an execution of a program that satisfies certain properties.
 - ◆ Atomicity
 - ◆ Consistency
 - ◆ Isolation
 - ◆ Durability

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 2

Database Model for Transaction Processing

- A database is modeled as a collection of named items.
- Operations to database items are simplified to:
 - ◆ Read(x)
 - ◆ Write(x)
- The read-set (write-set) of a transaction is the set of items that the transaction reads (writes).

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 3

An Example Transaction

T1:

Read(X)
X := X - 10
Write(X)
Read(Y)
Y := Y + 10
Write(Y)

Read-set (T1) = ?
Write-set (T1) = ?

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 4

What are transactions for?

- Concurrency control
- Recovery

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 5

Currency Control

- Currently running processes are actually interleaved.
- They may access same database items in these interleaved operations.
- The purpose of transactions (in the context of concurrency control) is to avoid the problems that may arise from interleaving

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 6

Problem 1: Lost Update

- The effect of a transaction on the DB is accidentally overwritten by another transaction
- $T1 = r1(a); \rightarrow a++; \quad w1(a)$
- $T2 = r2(a); \rightarrow a++; \quad w2(a)$


Spring 2002

CSC 742: DBMS by Dr. Peng Ning

7

Problem 2: Dirty Read

- A transaction prematurely reads a value from the DB that is later invalidated by another transaction



Spring 2002

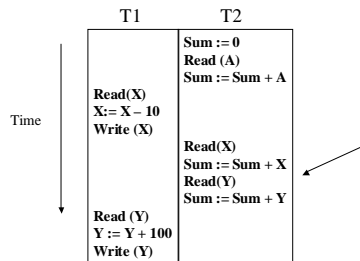
CSC 742: DBMS by Dr. Peng Ning

8

Problem 3: Incorrect Summary

Also termed *inconsistent analysis*

- A transaction observes an incorrect state of the DB
 - ◆ typically with some sort of aggregation



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

9

Recovery

- The DBMS needs to make sure that
 - ◆ Either all the operations of a transaction are completed and the effect is recorded *permanently* in the database,
 - ◆ Or the transaction has no effect at all.
- Recovery is necessary to make sure this is true *in case of failures*.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

10

Desirable Transaction Properties

- These define traditional transactions
 - ◆ Atomicity or Failure Atomicity
 - ◆ All-or-nothing
 - ◆ If failed then no changes or messages
 - ◆ Consistency
 - ◆ Don't violate DB integrity constraints
 - ◆ Isolation or Isolation Atomicity
 - ◆ As if the transaction is being executed alone.
 - ◆ Durability
 - ◆ Effects (of transactions that "happened" or committed) are forever.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

11

Atomicity or Failure Atomicity

- All or nothing
- If failed then
 - ◆ no changes
 - ◆ no messages

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

12

Consistency

- Don't violate DB integrity constraints
 - ◆ These constraints include
 - ◆ both the integrity constraints defined on database schemas and
 - ◆ other constraints that should hold on the database.
 - Rarely evaluated formally
 - Mostly just enforced by the programmer

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

13

Isolation or Isolation Atomicity

- Concurrent transactions are executed as if they are all executed alone.
- Partial results are hidden from users and other transactions

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

14

Durability

- Effects (of transactions that "happened" or committed) are forever.
- Should not be affected by failures.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

15

Programming with Transactions

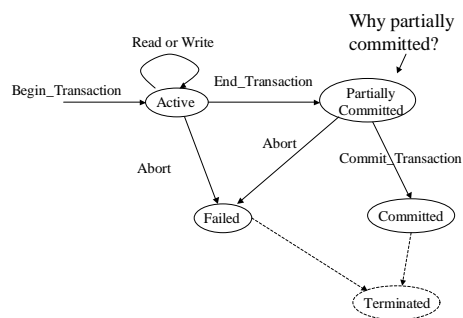
- Transactions simplify programming
 - ◆ the programmer guarantees correctness of individual transactions
 - ◆ the system eventually recovers
 - ◆ then the system guarantees the ACID properties for the entire set of transactions that execute concurrently

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

16

Transaction States and Operations



Spring 2002

CSC 742: DBMS by Dr. Peng Ning

17

System Log

- System log:
 - ◆ To keep track of all transaction operations.
 - ◆ Kept on disk.
 - ◆ Necessary for recovery from failures.
- Log records: Entries in the system log
 - ◆ [Start_Transaction, T]
 - ◆ [Write, T, old_value, new_value]
 - ◆ [Read, T, X] (May not be required, depending on the protocol.)
 - ◆ [Commit, T]
 - ◆ [Abort, T]
 - ◆ [checkpoint]

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

18

Commit Point of a Transaction

- A transaction T reaches its *commit point* when
 - ◆ All its operations that access the database have been executed, and
 - ◆ The effect of all transaction operations have been recorded in the log.
- A transaction is *committed* beyond its commit point.
 - ◆ [Commit, T] in the log.
- *Force write* the log:
 - ◆ Before T 's commit point, all of the log must be written to the disk.
 - ◆ Why?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

19

Checkpoint

- [Checkpoint] is appended to the log when writes of all previously committed transactions are effected on the database.
- Purpose:
 - ◆ Reduce the work required for recovery.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

20


Schedules

- Schedules are histories of computations showing all the events of interest
 - ◆ Schedule or history of $T_1 \dots T_n$ has operations of $T_1 \dots T_n$
 - ◆ Operations may be interleaved, but must be in the same order as within each T_i
 - ◆ Events of interest:
 - ◆ Read, Write, Commit, Abort in T_i
 - ◆ Abbreviations: r_i , w_i , c_i , a_i .

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

21




Schedule:

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

22



Schedule:


Spring 2002

CSC 742: DBMS by Dr. Peng Ning

23

Conflict

- Two operations in a schedule *conflict* if
 - ◆ They are in different transactions
 - ◆ They involve the same data item
 - ◆ At least one of them is write.



What are the conflict operations?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

24

Complete Schedule

- A schedule S of n transactions T_1, \dots, T_n is a *complete schedule* if
 - ◆ The operations in S are exactly those in T_1, \dots, T_n , including a commit or abort as the last operation for each T_i ;
 - ◆ For any pair of operations from the same T_i , their order is the same as in T_i ;
 - ◆ For any two conflicting operations, one must occur before the other.
- A *partial order* in which
 - ◆ a commit or abort event for each T_i is included
 - ◆ conflicting ops are ordered

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

25

Examples

Transactions: $T_1: r_1(X); w_1(X); r_1(Y); w_1(Y); c_1;$
 $T_2: r_2(Y); w_2(Y); a_2;$
 $T_3: r_3(X); w_3(X);$

$S_1: r_1(X); r_2(Y); w_1(X); w_2(Y); r_1(Y); a_2; w_1(Y); c_1; r_3(X); w_3(X);$

$S_2: r_1(X); r_2(Y); w_1(X); w_2(Y); r_1(Y); a_2; w_1(Y); c_1;$

$S_3: r_1(X); w_2(Y); w_1(X); r_2(Y); r_1(Y); a_2; w_1(Y); c_1;$

Is S_i a complete schedule? For what?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

26

Committed Projection

- A *committed projection*: $C(S)$ of a schedule S include the operations in S only from the committed transactions.

Transactions: $T_1: r_1(X); w_1(X); r_1(Y); w_1(Y); c_1;$
 $T_2: r_2(Y); w_2(Y); a_2;$
 $T_3: r_3(X); w_3(X);$

$S_1: r_1(X); r_2(Y); w_1(X); w_2(Y); r_1(Y); a_2; w_1(Y); c_1; r_3(X); w_3(X);$

$C(S_1) = ?$

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

27

Recoverable Schedules

- Goal:
 - ◆ Once a transaction T is committed, it should never be necessary to roll back T .
 - ◆ Recoverable schedules: schedules that meet this criterion.

$S_1: r_1(X); w_1(X); r_2(X); w_2(X); c_2; a_1;$

Is this recoverable?

In terms of the ACID properties, what is the risk in allowing a non-recoverable schedule?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

28

Recoverable Schedules (Cont'd)

- A schedule S is *recoverable* if each transaction commits after all transactions it *read from* have committed
 - ◆ T reads from T_2 if the schedule contains a subsequence $w_2(x) \dots r(x)$, where w_2 is the first write on x going backwards from $r(x)$.

$S_1: r_1(X); w_1(X); r_2(X); w_2(X); c_1; c_2;$

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

29

Avoid Cascading Aborts

- A schedule S can *avoid cascading aborts* if no transaction in S read from uncommitted transactions.
- What is the risk in doing so?

$S_1: r_1(X); w_1(X); r_2(X); w_2(X); a_1; a_2;$

$S_2: r_1(X); w_1(X); r_2(X); w_2(X); r_3(X); w_3(X); a_1; a_2; a_3;$

$S_3: r_1(X); w_1(X); a_1; r_2(X); w_2(X); c_2; r_3(X); w_3(X); c_3;$

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

30

Strict Schedules

- A schedule S is *strict* if no transaction can read or write an item X until the previous transaction to write that item has committed or aborted.

S1: r1(X); w1(X); w2(X); c1; c2;
S2: r1(X); w1(X); c1; w2(X); c2;

Avoid Cascading Abort?
Strict?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

31

Strict Schedules (Cont'd)

- Benefit:
 - ◆ Allows us to UNDO by restoring the before image

X=9
S1: w1(X, 5); w2(X, 8); a1;
X=?
Log: ...
[write, T1, X, 9, 5]
[write, T2, X, 5, 8]
[abort, T1]
...

X=9
S1: w1(X, 5); a1; w2(X, 8);
X=?
Log: ...
[write, T1, X, 9, 5]
[abort, T1]
[write, T2, X, 9, 8]
...

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

32

Note

- Strict \rightarrow Avoid Cascading Abort \rightarrow Recoverable.
- All of them are about recovery.

Spring 2002


CSC 742: DBMS by Dr. Peng Ning

33

Serial Schedules

Transactions are wholly before or after others.

- Serial schedules are correct (assuming each transaction is correct)
- Clearly, we must allow for service requests to come in slowly, one-by-one.




Spring 2002

CSC 742: DBMS by Dr. Peng Ning

34

Non-serial Schedules

X = 10, Y = 20



Are they correct?

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

35

Serializable Schedules: 1

- *Serializable schedules* are those *equivalent* to some serial schedule.
- Here equivalent can mean
 - ◆ *conflict equivalent*
 - ◆ The order of any two conflicting operations is the same in both schedules.
 - ◆ Alternatively, to be equivalent to a serial schedule, a non-serial schedule must have all pairs of conflicting operations ordered the same way.
 - ◆ *view equivalent*—each read of a transaction gets the same view in both schedules.

Spring 2002

CSC 742: DBMS by Dr. Peng Ning

36

X = 10, Y = 20

	T1	T2
Time ↓	Read(X) X := X - 10	Read(X) X := X + 100
	Write(X) Read(Y)	Write(X)
	Y := Y + 20 Write(Y)	

	T1	T2
Time ↓	Read(X) X := X - 10 Write(X)	Read(X) X := X + 100 Write(X)
	Read(Y) Y := Y + 20 Write(Y)	

Order of conflicting operations?
Are they conflict equivalent to any serial schedule?
Are they serializable w.r.t. conflict?

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 37

Serializable Schedules: 2

- Most approaches use conflict equivalence, because conflict equivalence
 - entails view equivalence
 - is relatively easy to ensure

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 38

Checking Serializability

- Build a *serialization graph* for the given concurrent mix of transactions:
 - a node for each transaction
 - a directed edge for each conflict
 - serializability \leftrightarrow the graph is acyclic

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 39

S1: r2(Z); r2(Y); w2(Y); r3(Y); r3(Z); r1(X); w1(X); w3(Y); w3(Z); r2(X); r1(Y); w1(Y); w2(X);

S2: r3(Y); r3(Z); r1(X); w1(X); w3(Y); w3(Z); r2(Z); r1(Y); w1(Y); r2(Y); w2(Y); r2(X); w2(X);

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 40

Serializable Schedules: 3

Can we find an equivalent serial schedule for a serializable schedule?

- Such a serial schedule exists by definition
- Can easily derive from serialization graph.

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 41

```

graph BT
    T3 --> T1
    T3 --> T2
  
```

Equivalent serial schedules: ?

Spring 2002 CSC 742: DBMS by Dr. Peng Ning 42

Achieving Serializability

- Optimistically: Let each transaction run, but check for serializability before committing.
- Pessimistically: Use a protocol, e.g., locking, to ensure that only serializable schedules are realized

Generally, the pessimistic approach is more common