Chapter 20

Transaction Management

© Pearson Education Limited 1995, 2005

Chapter 20 - Objectives

- ◆ Function and importance of transactions.
- ♦ Properties of transactions.
- Concurrency Control
 - Deadlock and how it can be resolved.
 - Granularity of locking.

Chapter 20 - Objectives

- ♦ Recovery Control
 - Some causes of database failure.
 - Purpose of transaction log file.
 - Purpose of checkpointing.
 - How to recover following database failure.

© Pearson Education Limited 1995, 2005

3

Transaction Support

Transaction

Action, or series of actions, carried out by user or application, which reads or updates contents of database.

- ◆ Logical unit of work on the database.
- Application program is series of transactions with nondatabase processing in between.
- Transforms database from one consistent state to another, although consistency may be violated during transaction.

	delete(staffNo = x)
	for all PropertyForRent records, pno
read(staffNo = x, salary)	begin
salary = salary * 1.1	read(propertyNo = pno, staffNo)
<pre>write(staffNo = x, new_salary)</pre>	if $(staffNo = x)$ then
	begin
	staffNo = newStaffNo
	write(propertyNo = pno, staffNo)
	end
	end
(a)	(b)

© Pearson Education Limited 1995, 2005

5

6

Transaction Support

- Can have one of two outcomes:
 - Success transaction *commits* and database reaches a new consistent state.
 - Failure transaction *aborts*, and database must be restored to consistent state before it started.
 - Such a transaction is *rolled back* or *undone*.
- Committed transaction cannot be aborted.
- Aborted transaction that is rolled back can be restarted later.

State Transition Diagram for Transaction



Properties of Transactions

♦Four basic (*ACID*) properties of a transaction are:

<u>Atomicity</u> 'All or nothing' property.

<u>Consistency</u> Must transform database from one consistent state to another.

7

Isolation Partial effects of incomplete transactions should not be visible to other transactions.

<u>Durability</u> Effects of a committed transaction are permanent and must not be lost because of later failure.

DBMS Transaction Subsystem



Concurrency Control

Process of managing simultaneous operations on the database without having them interfere with one another.

- Prevents interference when two or more users are accessing database simultaneously and at least one is updating data.
- Although two transactions may be correct in themselves, interleaving of operations may produce an incorrect result.

Need for Concurrency Control

- Three examples of potential problems caused by concurrency:
 - Lost update problem.
 - Uncommitted dependency problem.
 - Inconsistent analysis problem.

© Pearson Education Limited 1995, 2005

Lost Update Problem

- Successfully completed update is overridden by another user.
- T_1 withdrawing £10 from an account with bal_x , initially £100.
- ◆ T₂ depositing £100 into same account.
- ♦ Serially, final balance would be £190.

Lost Update Problem

Time	T ₁	T ₂	bal _x
t ₁		begin_transaction	100
t ₂	begin_transaction	$read(bal_x)$	100
t ₃	$read(bal_x)$	$bal_{X} = bal_{X} + 100$	100
t ₄	$bal_{X} = bal_{X} - 10$	write(bal _x)	200
t ₅	write(bal _x)	commit	90
t ₆	commit		90

Loss of T₂'s update avoided by preventing T₁ from reading bal_x until after update.

© Pearson Education Limited 1995, 2005

Uncommitted Dependency Problem (dirty read)

- Occurs when one transaction can see intermediate results of another transaction before it has committed.
- T_4 updates bal_x to £200 but it aborts, so bal_x should be back at original value of £100.
- ♦ T₃ has read new value of bal_x (£200) and uses value as basis of £10 reduction, giving a new balance of £190, instead of £90.

Uncommitted Dependency Problem

Time	T ₃	T ₄	bal _x
t ₁		begin_transaction	100
t ₂		read(bal _x)	100
t ₃		$bal_x = bal_x + 100$	100
t ₄	begin_transaction	write(bal _x)	200
t ₅	read(bal _x)	1	200
t ₆	$bal_x = bal_x - 10$	rollback	100
t ₇	write(bal _x)		190
t ₈	commit		190

Problem avoided by preventing T₃ from reading bal_x until after T₄ commits or aborts.

© Pearson Education Limited 1995, 2005

Inconsistent Analysis Problem

- Occurs when transaction reads several values but second transaction updates some of them during execution of first.
- ◆ Sometimes referred to as *fuzzy read* or *unrepeatable read*.
- ♦ T_6 is totaling balances of account x (£100), account y (£50), and account z (£25).
- ♦ Meantime, T₅ has transferred £10 from bal_x to bal_z, so T₆ now has wrong result (£10 too high).

Inconsistent Analysis Problem

Time	T ₅	T ₆	bal _x	bal _y	balz	sum
t ₁		begin_transaction	100	50	25	
t ₂	begin_transaction	sum = 0	100	50	25	0
t ₃	$read(bal_x)$	$read(bal_x)$	100	50	25	0
t ₄	$bal_x = bal_x - 10$	$sum = sum + bal_x$	100	50	25	100
t ₅	write(bal _x)	$read(bal_y)$	90	50	25	100
t ₆	read(bal _z)	$sum = sum + bal_y$	90	50	25	150
t ₇	$bal_{z} = bal_{z} + 10$		90	50	25	150
t ₈	write(bal _z)		90	50	35	150
t9	commit	read(bal _z)	90	50	35	150
t ₁₀		$sum = sum + bal_z$	90	50	35	185
t ₁₁		commit	90	50	35	185

Problem avoided by preventing T₆ from reading bal_x and bal_z until after T₅ completed updates.

© Pearson Education Limited 1995, 2005

Concurrency Control Techniques

- ◆ Two basic concurrency control techniques:
 - Locking,
 - Timestamping.
- Both are conservative approaches: delay transactions in case they conflict with other transactions.
- Optimistic methods assume conflict is rare and only check for conflicts at commit.

Locking

Transaction uses locks to deny access to other transactions and so prevent incorrect updates.

- Most widely used approach to ensure serializability.
- Generally, a transaction must claim a shared (read) or exclusive (write) lock on a data item before read or write.
- Lock prevents another transaction from modifying item or even reading it, in the case of a write lock.

© Pearson Education Limited 1995, 2005

Locking - Basic Rules

- ♦ If transaction has shared lock on item, can read but not update item.
- If transaction has exclusive lock on item, can both read and update item.
- Reads cannot conflict, so more than one transaction can hold shared locks simultaneously on same item.
- Exclusive lock gives transaction exclusive access to that item.

Locking - Basic Rules

Some systems allow transaction to upgrade read lock to an exclusive lock, or downgrade exclusive lock to a shared lock.

© Pearson Education Limited 1995, 2005

Example - Incorrect Locking Schedule

Time	T ₉	T ₁₀
t ₁	begin_transaction	
t ₂	$read(bal_x)$	
t ₃	$bal_x = bal_x + 100$	
t ₄	$write(bal_x)$	begin_transaction
t ₅		read(bal _x)
t ₆		$bal_x = bal_x * 1.1$
t ₇		write(bal _x)
t ₈		$read(bal_y)$
t9		$bal_y = bal_y *1.1$
t ₁₀		write(baly)
t ₁₁	read(bal _y)	commit
t ₁₂	$bal_y = bal_y - 100$	
t ₁₃	write(baly)	
t ₁₄	commit	

- For two transactions above, a valid schedule using these rules is:
- $S = \{ write_lock(T_0, bal_x), read(T_0, bal_x), \}$ write(T_0 , bal_x), unlock(T_0 , bal_x), write_lock(T_{10} , bal_x), read(T_{10} , bal_x), write(T_{10} , bal_x), unlock(T_{10} , bal_x), write_lock(T_{10} , bal_v), read(\tilde{T}_{10} , bal_v), write(T_{10} , bal_v), unlock(T_{10} , bal_v), $\operatorname{commit}(\mathbf{T}_{10}),$ write $lock(T_0,$ bal,), $read(T_0, bal_v),$ write(T_o, bal_v), unlock(T_0 , bal_v), commit(T_0) }

Example - Incorrect Locking Schedule

Time	T ₉	T ₁₀
t ₁	begin_transaction	
t ₂	$read(bal_x)$	
t ₃	$bal_x = bal_x + 100$	
t ₄	write(bal _x)	begin_transaction
t ₅		$read(\mathbf{bal}_{\mathbf{x}})$
t ₆		$bal_x = bal_x * 1.1$
t ₇		write(bal _x)
t ₈		$read(bal_y)$
t9		$bal_y = bal_y * 1.1$
t ₁₀		write(bal _y)
t ₁₁	read(bal _y)	commit
t ₁₂	$bal_y = bal_y - 100$	•
t ₁₃	write(baly)	•
t ₁₄	commit	

- If at start, bal_x = 100, bal_y = 400, result should be:
 - $bal_x = 220$, $bal_y = 330$, if T_9 executes before T_{10} , or
 - $bal_x = 210$, $bal_y = 340$, if T_{10} executes before T_9 .
- However, result gives $bal_x = 220$ and $bal_y = 340$.

◆ S is not a serializable schedule.

© Pearson Education Limited 1995, 2005

Example - Incorrect Locking Schedule

- Problem is that transactions release locks too soon, resulting in loss of total isolation and atomicity.
- ♦ To guarantee serializability, need an additional protocol concerning the positioning of lock and unlock operations in every transaction.

Two-Phase Locking (2PL)

Transaction follows 2PL protocol if all locking operations precede first unlock operation in the transaction.

- Two phases for transaction:
 - Growing phase acquires all locks but cannot release any locks.
 - Shrinking phase releases locks but cannot acquire any new locks.

© Pearson Education Limited 1995, 2005

Preventing Lost Update Problem using 2PL

Time	T_1	T ₂	bal _x
t ₁		begin_transaction	100
t ₂	begin_transaction	write_lock(bal _x)	100
t ₃	write_lock(bal _x)	read(bal _x)	100
t ₄	WAIT	$bal_x = bal_x + 100$	100
t ₅	WAIT	write(bal _x)	200
t ₆	WAIT	commit/unlock(bal _x)	200
t ₇	read(bal _x)		200
t ₈	$bal_x = bal_x - 10$		200
t ₉	write(bal _x)		190
t ₁₀	commit/unlock(bal _x)		190

Preventing Uncommitted Dependency Problem using 2PL

Time	T ₃	T ₄	bal _x
t ₁		begin_transaction	100
t ₂		write_lock(bal _x)	100
t3		read(bal _x)	100
t ₄	begin_transaction	$bal_x = bal_x + 100$	100
t ₅	write_lock(bal _x)	write(bal _x)	200
t ₆	WAIT	rollback/unlock(bal _x)	100
t ₇	read(bal _x)		100
t ₈	$bal_x = bal_x - 10$		100
t9	write(bal _x)		90
t ₁₀	$commit/unlock(\mathbf{bal}_{\mathbf{X}})$		90

© Pearson Education Limited 1995, 2005

Preventing Inconsistent Analysis Problem using 2PL

Time	T ₅	T ₆	bal _x	bal _y	balz	sum
t ₁		begin_transaction	100	50	25	
t ₂	begin_transaction	sum = 0	100	50	25	0
t ₃	write_lock(bal _x)		100	50	25	0
t ₄	read(bal _x)	read_lock(bal _x)	100	50	25	0
t ₅	$bal_x = bal_x - 10$	WAIT	100	50	25	0
t ₆	write(bal _x)	WAIT	90	50	25	0
t ₇	write_lock(bal _z)	WAIT	90	50	25	0
t ₈	read(bal _z)	WAIT	90	50	25	0
t9	$bal_z = bal_z + 10$	WAIT	90	50	25	0
t ₁₀	write(bal _z)	WAIT	90	50	35	0
t ₁₁	commit/unlock(bal_x, bal_z)	WAIT	90	50	35	0
t ₁₂		read(bal _x)	90	50	35	0
t ₁₃		$sum = sum + bal_x$	90	50	35	90
t ₁₄		read_lock(bal y)	90	50	35	90
t ₁₅		read(bal _y)	90	50	35	90
t ₁₆		$sum = sum + bal_y$	90	50	35	140
t ₁₇		read_lock(bal _z)	90	50	35	140
t ₁₈		read(bal _z)	90	50	35	140
t19		$sum = sum + bal_z$	90	50	35	175
t20		commit/unlock(bal _x , bal _y , bal _z)	90	50	35	175

Deadlock

An impasse that may result when two (or more) transactions are each waiting for locks held by the other to be released.

Time	T ₁₇	T ₁₈
t ₁	begin_transaction	
t ₂	write_lock(bal _x)	begin_transaction
t ₃	read(bal _x)	write_lock(bal y)
t ₄	$bal_x = bal_x - 10$	read(bal y)
t ₅	write(bal _x)	$bal_y = bal_y + 100$
t ₆	write_lock(bal y)	write(bal _y)
t ₇	WAIT	write_lock(bal _x)
t ₈	WAIT	WAIT
t9	WAIT	WAIT
t ₁₀	:	WAIT
t ₁₁	÷	:

© Pearson Education Limited 1995, 2005

Deadlock

- Only one way to break deadlock: abort one or more of the transactions.
- Deadlock should be transparent to user, so DBMS should restart transaction(s).
- Three general techniques for handling deadlock:
 - Timeouts.
 - Deadlock prevention.
 - Deadlock detection and recovery.

Timeouts

- Transaction that requests lock will only wait for a system-defined period of time.
- If lock has not been granted within this period, lock request times out.
- ♦ In this case, DBMS assumes transaction may be deadlocked, even though it may not be, and it aborts and automatically restarts the transaction.

© Pearson Education Limited 1995, 2005

Deadlock Prevention

- DBMS looks ahead to see if transaction would cause deadlock and never allows deadlock to occur.
- Could order transactions using transaction timestamps:
 - <u>Wait-Die</u> only an older transaction can wait for younger one, otherwise transaction is aborted (*dies*) and restarted with same timestamp.

 <u>Wound-Wait</u> - only a younger transaction can wait for an older one. If older transaction requests lock held by younger one, younger one is aborted (*wounded*).

© Pearson Education Limited 1995, 2005

Deadlock Detection and Recovery

- DBMS allows deadlock to occur but recognizes it and breaks it.
- Usually handled by construction of wait-for graph (WFG) showing transaction dependencies:
 - Create a node for each transaction.
 - Create edge $T_i \rightarrow T_j$, if T_i waiting to lock item locked by T_j .
- Deadlock exists if and only if WFG contains cycle.
- ◆ WFG is created at regular intervals.





© Pearson Education Limited 1995, 2005

Recovery from Deadlock Detection

- ♦ Several issues:
 - choice of deadlock victim;
 - how far to roll a transaction back;
 - avoiding starvation.

Process of restoring database to a correct state in the event of a failure.

- ♦ Need for Recovery Control
 - Two types of storage: volatile (main memory) and nonvolatile.
 - Volatile storage does not survive system crashes.
 - Stable storage represents information that has been replicated in several nonvolatile storage media with independent failure modes.

© Pearson Education Limited 1995, 2005

Types of Failures

- System crashes, resulting in loss of main memory.
- Media failures, resulting in loss of parts of secondary storage.
- ♦ Application software errors.
- ♦ Natural physical disasters.
- Carelessness or unintentional destruction of data or facilities.
- ◆ Sabotage.

Transactions and Recovery

- ◆ Transactions represent basic unit of recovery.
- Recovery manager responsible for atomicity and durability.
- If failure occurs between commit and database buffers being flushed to secondary storage then, to ensure durability, recovery manager has to redo (rollforward) transaction's updates.

© Pearson Education Limited 1995, 2005

Transactions and Recovery

- ♦ If transaction had not committed at failure time, recovery manager has to *undo* (*rollback*) any effects of that transaction for atomicity.
- Partial undo only one transaction has to be undone.
- ♦ Global undo all transactions have to be undone.

Example



- DBMS starts at time t₀, but fails at time t_f. Assume data for transactions T₂ and T₃ have been written to secondary storage.
- ◆ T₁ and T₆ have to be undone. In absence of any other information, recovery manager has to redo T₂, T₃, T₄, and T₅.

© Pearson Education Limited 1995, 2005

Recovery Facilities

- DBMS should provide following facilities to assist with recovery:
 - Backup mechanism, which makes periodic backup copies of database.
 - Logging facilities, which keep track of current state of transactions and database changes.
 - Checkpoint facility, which enables updates to database in progress to be made permanent.
 - Recovery manager, which allows DBMS to restore database to consistent state following a failure.

Log File

Contains information about all updates to database:

Transaction records.

- Checkpoint records.
- Often used for other purposes (for example, auditing).

© Pearson Education Limited 1995, 2005

Log File

- Transaction records contain:
 - Transaction identifier.
 - Type of log record, (transaction start, insert, update, delete, abort, commit).
 - Identifier of data item affected by database action (insert, delete, and update operations).
 - Before-image of data item.
 - After-image of data item.
 - Log management information.

Sample Log File

Tid	Time	Operation	Object	Before image	After image	pPtr	nPtr
T1	10:12	START				0	2
T1	10:13	UPDATE	STAFF SL21	(old value)	(new value)	1	8
T2	10:14	START			Arto- 803-9	0	4
T2	10:16	INSERT	STAFF SG37		(new value)	3	5
T2	10:17	DELETE	STAFF SA9	(old value)		4	6
T2	10:17	UPDATE	PROPERTY PG16	(old value)	(new value)	5	9
T3	10:18	START			2 M 10 10 10 10 10 10 10 10 10 10 10 10 10	0	11
T1	10:18	COMMIT				2	0
	10:19	CHECKPOINT	T2, T3				
T2	10:19	COMMIT				6	0
T3	10:20	INSERT	PROPERTY PG4		(new value)	7	12
T3	10:21	COMMIT				11	0

© Pearson Education Limited 1995, 2005

Log File

- ◆ Log file may be duplexed or triplexed.
- ♦ Log file sometimes split into two separate random-access files.
- Potential bottleneck; critical in determining overall performance.

Checkpointing

Checkpoint

Point of synchronization between database and log file. All buffers are force-written to secondary storage.

- Checkpoint record is created containing identifiers of all active transactions.
- When failure occurs, redo all transactions that committed since the checkpoint and undo all transactions active at time of crash.

© Pearson Education Limited 1995, 2005

Checkpointing

♦ In previous example, with checkpoint at time t_c, changes made by T₂ and T₃ have been written to secondary storage.

♦ Thus:

- only redo T_4 and T_5 ,
- undo transactions T₁ and T6.

Recovery Techniques

- ◆ If database has been damaged:
 - Need to restore last backup copy of database and reapply updates of committed transactions using log file.
- ♦ If database is only inconsistent:
 - Need to undo changes that caused inconsistency. May also need to redo some transactions to ensure updates reach secondary storage.
 - Do not need backup, but can restore database using before- and after-images in the log file.

49

© Pearson Education Limited 1995, 2005