Databases – Transactions

Jörg Endrullis

VU University Amsterdam

Transactions :: Concurrency Anomalies

Example Transaction



A withdrawal of 100 euro causes the ATM to perform a **transaction in the bank's database**.

ATM Transaction

 $balance \leftarrow read(account)$ $balance \leftarrow balance - 100$ write(account, balance)

Example Transaction



A withdrawal of 100 euro causes the ATM to perform a **transaction in the bank's database**.

ATM Transaction

 $balance \leftarrow read(account) \\ balance \leftarrow balance - 100 \\ write(account, balance) \\$

The account is properly updated to reflect the new balance.

Money Transfer from Checking to Saving

Subtract money from source (checking) account:

- 1. *checking_balance* \leftarrow read(*checking_account*)
- 2. checking_balance \leftarrow checking_balance 500
- 3. write(*checking_account*, *checking_balance*)

Add money to the target (saving) account:

- 4. $saving_balance \leftarrow read(saving_account)$
- 5. $saving_balance \leftarrow saving_balance + 500$
- 6. write(*saving_account*, *saving_balance*)

Money Transfer from Checking to Saving

Subtract money from source (checking) account:

- 1. *checking_balance* \leftarrow read(*checking_account*)
- 2. checking_balance \leftarrow checking_balance 500
- 3. write(*checking_account*, *checking_balance*)

Add money to the target (saving) account:

- 4. $saving_balance \leftarrow read(saving_account)$
- 5. $saving_balance \leftarrow saving_balance + 500$ System crash!
- 6. write(*saving_account*, *saving_balance*)

Before the transaction gets to step 6, the system crashes. (power outage, disk failure or software bug)

Money Transfer from Checking to Saving

Subtract money from source (checking) account:

- 1. *checking_balance* \leftarrow read(*checking_account*)
- 2. checking_balance \leftarrow checking_balance 500
- 3. write(*checking_account*, *checking_balance*)

Add money to the target (saving) account:

- 4. $saving_balance \leftarrow read(saving_account)$
- 5. $saving_balance \leftarrow saving_balance + 500$ System crash!
- 6. write(*saving_account*, *saving_balance*)

Before the transaction gets to step 6, the system crashes. (power outage, disk failure or software bug)

My money is lost!

Money Transfer from Checking to Saving

Subtract money from source (checking) account:

- 1. *checking_balance* \leftarrow read(*checking_account*)
- 2. checking_balance \leftarrow checking_balance 500
- 3. write(*checking_account*, *checking_balance*)

Add money to the target (saving) account:

- 4. $saving_balance \leftarrow read(saving_account)$
- 5. $saving_balance \leftarrow saving_balance + 500$ System crash!
- 6. write(*saving_account*, *saving_balance*)

Before the transaction gets to step 6, the system crashes. (power outage, disk failure or software bug)

My money is lost!

Transactions should be **atomic** (executed fully or not at all).

My wife and I have credit cards for the same account.

What if we use the cards at the same time (concurrently)?

My wife and I have **credit cards for the same account**. What if we use the cards at the same time (**concurrently**)?

Concurrent	ATM Transaction
------------	-----------------

me (withdraws 100)	my wife (withdraws 200)	state
$balance \leftarrow read(account)$		1200
=1200	$balance \leftarrow read(account)$	1200
$\textit{balance} \leftarrow \textit{balance} - 100$	=1200	1200
=1100 =1200	$\underline{balance} \leftarrow \underline{balance} - 200$	1200
	=1000 =1200 write(<i>account</i> , <i>balance</i>)	1000
write(account, balance)	=1000	1100
=1100		

My wife and I have **credit cards for the same account**. What if we use the cards at the same time (**concurrently**)?

Concurrent .	ATM Transaction
--------------	-----------------

me (withdraws 100)	my wife (withdraws 200)	state
$balance \leftarrow read(account)$		1200
=1200	$balance \leftarrow read(account)$	1200
$balance \leftarrow balance - 100$	=1200	1200
=1100 =1200	$balance \leftarrow balance - 200$	1200
	=1000 =1200 write(<i>account</i> , <i>balance</i>)	1000
write(account, balance)	=1000	1100
=1100		

The update of my wife was lost during this execution. Lucky me!

My wife and I have **credit cards for the same account**. What if we use the cards at the same time (**concurrently**)?

Concurrent .	ATM Transaction
--------------	-----------------

me (withdraws 100)	my wife (withdraws 200)	state
$balance \leftarrow read(account)$		1200
=1200	$balance \leftarrow read(account)$	1200
$balance \leftarrow balance - 100$	=1200	1200
=1100 =1200	$balance \leftarrow balance - 200$	1200
	=1000 =1200 write(<i>account</i> , <i>balance</i>)	1000
write(account, balance)	=1000	1100
=1100		

The update of my wife was lost during this execution. Lucky me!

This is known as **lost update anomaly**.

Concurrent Access: Inconsistent Read

Reconsider the transfer from checking to saving account:

Transaction 1	Transaction 2
 update Accounts set balance = balance - 500 where customer = 1904 and account_type = 'Checking' 	
	 select sum(balance) from Accounts where customer = 1904
 update Accounts set balance = balance + 500 where customer = 1904 and account_type = 'Saving' 	

Concurrent Access: Inconsistent Read

Reconsider the transfer from checking to saving account:

Transaction 1	Transaction 2
 update Accounts set balance = balance - 500 where customer = 1904 and account_type = 'Checking' 	
	 select sum(balance) from Accounts where customer = 1904
<pre>5. update Accounts 6. set balance = balance + 500 7. where customer = 1904 8. and account_type = 'Saving'</pre>	

Transaction 2 sees a temporary, inconsistent database state.

Concurrent Access: Inconsistent Read

Reconsider the transfer from checking to saving account:

Transaction 1	Transaction 2
 update Accounts set balance = balance - 500 where customer = 1904 and account_type = 'Checking' 	
	 select sum(balance) from Accounts where customer = 1904
 update Accounts set balance = balance + 500 where customer = 1904 and account_type = 'Saving' 	

Transaction 2 sees a temporary, **inconsistent database state**.

This is known as inconsistent read anomaly.

Again, my wife and I are doing a transaction at the same time. This time, **my transaction gets cancelled**!

+ ATNA Tue we and the w

Again, my wife and I are doing a transaction at the same time. This time, **my transaction gets cancelled**!

me	my wife	state
$balance \leftarrow read(account)$		1200
$balance \leftarrow balance - 100$		1200
write(<i>account</i> , <i>balance</i>)		1100
	$balance \leftarrow read(account)$	1100
	$balance \leftarrow balance - 200$	1100
abort		1200
	write(account, balance)	900

Concurrent ATM Transaction

Again, my wife and I are doing a transaction at the same time. This time, **my transaction gets cancelled**!

me	my wife	state
$balance \leftarrow read(account)$		1200
$balance \leftarrow balance - 100$		1200
write(<i>account</i> , <i>balance</i>)		1100
	$balance \leftarrow read(account)$	1100
	$balance \leftarrow balance - 200$	1100
abort		1200
	write(account, balance)	900

My wife's transaction has read the modified balance before my transaction was **rolled back** (i.e., the effects are undone).

Again, my wife and I are doing a transaction at the same time. This time, **my transaction gets cancelled**!

me	my wife	state	
$balance \leftarrow read(account)$		1200	
$balance \leftarrow balance - 100$		1200	
write(account, balance)		1100	
	$balance \leftarrow read(account)$	1100	
	$balance \leftarrow balance - 200$	1100	
abort		1200	
	write(account, balance)	900	

Concurrent ATM Transaction

My wife's transaction has read the modified balance before my transaction was **rolled back** (i.e., the effects are undone).

This is known as **dirty read anomaly**.

Concurrency Anomalies

Lost Update Anomaly

The effects of one transaction are lost due to an uncontrolled overwrite performed by a second transaction.

Inconsistent Read

A transaction reads the partial result of another transaction.

Dirty Read

A transaction reads changes made by another transaction that is not yet committed (and might get aborted & rolled back).

Unrepeatable Read

A transaction reads a value which is afterwards changed by another transaction (before the former transaction is finished). So the first transaction operates on stale data.

ACID Properties

To prevent the mentioned problems...

Database management system ensures **ACID properties**

Atomicity:

transaction executes fully (commit) or not at all (abort)

Consistency:

transactions always leave the database in a consistent state where all defined integrity constraints hold

Isolation:

multiple users can modify the database at the same time but will not see each others partial actions

Durability:

once a transaction is committed successfully, the modified data is persistent, regardless of disk crashes

Transactions :: Transactions, Schedules and Serializability

Transactions

Formally, transactions are defined as:

A transaction is a list of actions.

The actions are

- reads (written R(O)) and
- writes (written W(O))
- of database objects O.

Transactions end with Commit or Abort.

These are sometimes omitted if not relevant.

Example Transaction

 T_1 : R(V), R(Y), W(V), W(C), Commit

The **scheduler** decides the execution order of concurrent database access.



Schedules

A schedule is a list of actions from a set of transactions. Intuitively, this is a plan on how to execute transactions.

The **order** in which 2 actions of a transaction T appear in a schedule must be the same order as they appear in T.

Schedules

A schedule is a list of actions from a set of transactions. Intuitively, this is a plan on how to execute transactions.

The **order** in which 2 actions of a transaction T appear in a schedule must be the same order as they appear in T.

 $T_1 : \mathsf{R}(\mathsf{V}) \mathsf{W}(\mathsf{V})$ $T_2 : \mathsf{R}(\mathsf{Y}) \mathsf{W}(\mathsf{Y})$

Which of the following is a schedule of these transactions?



A schedule is **serial** if the actions of the different transactions are not interleaved; they are executed one after another.

$$S_1: \frac{T_1}{T_2} = \frac{R(V) \quad W(V)}{R(Y)}$$

A schedule is **serial** if the actions of the different transactions are not interleaved; they are executed one after another.

$$S_1: \frac{T_1}{T_2} = \frac{R(V) \quad W(V)}{R(Y)}$$

A schedule is **serializable** if its effect on the database is the same as that of some serial schedule.

A schedule is **serial** if the actions of the different transactions are not interleaved; they are executed one after another.

$$S_1: \begin{array}{ccc} T_1 & \mathsf{R}(\mathsf{V}) & \mathsf{W}(\mathsf{V}) \\ \hline T_2 & \mathsf{R}(\mathsf{Y}) & \mathsf{W}(\mathsf{Y}) \end{array}$$

A schedule is **serializable** if its effect on the database is the same as that of some serial schedule.

Quiz Serializable Schedules

We usually only want to allow serializable schedules. Why?

Two actions in a schedule **conflict** if they:

- are from different transactions,
- involve the same data item, and
- one of the actions is a write.

$$T_1$$
R(Y)W(Y)W(X) T_2 R(Y)W(Z)

Two actions in a schedule **conflict** if they:

- are from different transactions,
- involve the same data item, and
- one of the actions is a write.

$$T_1$$
R(Y)W(Y)W(X) T_2 R(Y)W(Z)

Two actions in a schedule **conflict** if they:

- are from different transactions,
- involve the same data item, and
- one of the actions is a write.

$$T_1$$
R(Y)W(Y)W(X) T_2 R(Y)W(Z)

There are several types of conflicts:

- write read (WR)
- read write (RW)
- write write (WW)

Two actions in a schedule **conflict** if they:

- are from different transactions,
- involve the same data item, and
- one of the actions is a write.

$$T_1$$
R(Y)W(Y)W(X) T_2 R(Y)W(Z)

There are several types of conflicts:

- write read (WR)
- read write (RW)
- write write (WW)

Such conflicts may cause a schedule to be not serializable.

There is a **WR conflict** between T_1 and T_2 if there is an item Y: **T**₁ writes Y and afterwards, T_2 reads Y

If T_1 has not committed this is a **dirty read**.

There is a **WR conflict** between *T*₁ and *T*₂ if there is an item Y: ■ *T*₁ writes Y and afterwards, *T*₂ reads Y

If T_1 has not committed this is a **dirty read**.

Find all WR conflicts in the following schedule											
	<i>T</i> ₁			W(Y)							
	<i>T</i> ₂	R(V)			R(Y)	W(Z)					
	<i>T</i> ₃		W(V)								

There is a WR conflict between *T*₁ and *T*₂ if there is an item Y: *T*₁ writes Y and afterwards, *T*₂ reads Y

If T_1 has not committed this is a **dirty read**.

Find all WR conflicts in the following schedule												
	<i>T</i> ₁			W(Y)								
	<i>T</i> ₂	R(V)			R(Y)	W(Z)						
	<i>T</i> ₃		W(V)									
This read becomes unrepeatable.

There is a **RW conflict** between T_1 and T_2 if there is an item Y: T_1 reads Y and afterwards, T_2 writes Y

This read becomes unrepeatable.

Find all RW conflicts in the following schedule								
	<i>T</i> ₁			W(Y)				-
	<i>T</i> ₂	R(V)			R(Y)	W(Z)	R(V)	-
	<i>T</i> ₃		W(V)					

There is a **RW conflict** between T_1 and T_2 if there is an item Y: T_1 reads Y and afterwards, T_2 writes Y

This read becomes unrepeatable.

Find all RW conflicts in the following schedule								
	<i>T</i> ₁			W(Y)				-
	<i>T</i> ₂	R(V)			R(Y)	W(Z)	R(V)	-
	<i>T</i> ₃		W(V)					

Find	all \	WW conflic	ts in th	ne follo	wing so	chedule		
	<i>T</i> ₁	W(Y)						
	<i>T</i> ₂	W	(V)		R(Z)	W(Y)	W(Z)	
	<i>T</i> ₃		١	W(V)				

Find	all V	WW conflicts in	the follo	owing s	chedule		
	<i>T</i> ₁	W(Y)					
	<i>T</i> ₂	W(V)		R(Z)	W(Y)	W(Z)	
	<i>T</i> ₃		W(V)				

Find all WW conflicts in the following schedule								
	<i>T</i> ₁	W(Y)						
	<i>T</i> ₂		W(V)		R(Z)	W(Y)	W(Z)	
	<i>T</i> ₃			W(V)				

Swapping Actions

We can **swap actions** (of different transactions) without changing the outcome, if the actions are **non-conflicting**.

Swapping Actions

We can **swap actions** (of different transactions) without changing the outcome, if the actions are **non-conflicting**.



Swapping Actions

We can **swap actions** (of different transactions) without changing the outcome, if the actions are **non-conflicting**.



Conflict Equivalent Schedules

Two schedules are **conflict equivalent** if they can be transformed into each other by a sequence of **swaps of non-conflicting, adjacent actions** (of different transactions).

Are any of the following schedules conflict equivalent? *S*₁ : $\begin{array}{ccc} T_1 & W(V) & R(V) \\ \hline T_2 & & R(V) \end{array}$ W(V)*S*₂ : R(V) $S_3: \begin{array}{c} T_1 & W(V) & R(V) & W(V) \\ \hline T_2 & R(V) \end{array}$



Schedules S_1 and S_2 are conflict equivalent (RR swap).

A schedule is **conflict-serializable** if it is conflict equivalent to some serial schedule.

A schedule is **conflict-serializable** if it is conflict equivalent to some serial schedule.



A schedule is **conflict-serializable** if it is conflict equivalent to some serial schedule.



A schedule is **conflict-serializable** if it is conflict equivalent to some serial schedule.



A schedule is **conflict-serializable** if it is conflict equivalent to some serial schedule.



- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

T_1	W(V)		W(V)		
<i>T</i> ₂		R(V)			
<i>T</i> ₁	R(V)		W(V)		
<i>T</i> ₂		R(V)			
<i>T</i> ₁		W(Y)			
<i>T</i> ₂	R(V)			R(Y)	۷
<i>T</i> ₃			W(V)		

- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

<i>T</i> ₁	<i>T</i> ₂
R(Y) W(Z)	
-	R(Y) W(Z)

- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

<i>T</i> ₁	W(V)	W(V)			<i>τ</i> ~	¥ ₇
<i>T</i> ₂	R(V)				¹¹ K	12
<i>T</i> ₁	R(V)	W(V)				
<i>T</i> ₂	R(V)					
<i>T</i> ₁	W(Y)					
<i>T</i> ₂	R(V)		R(Y)	W(Z)		
<i>T</i> ₃		W(V)				

- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

<i>T</i> ₁	W(V)	W(V)	-		Τ.	\frown	Τ.	
<i>T</i> ₂	R(\	V)	_		/1	K	12	
<i>T</i> ₁	R(V)	W(V)			Τ.		τ.	
<i>T</i> ₂	R(\	V)	_		/1		12	
<i>T</i> ₁	W(Y)						
<i>T</i> ₂	R(V)		R(Y)	W(Z)				
<i>T</i> ₃		W(V)						

- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

<i>T</i> ₁	W(V)	W(V)			τ -	\rightarrow_{τ}	
<i>T</i> ₂	R(V)				^{/1} ≮	\sim $^{\prime 2}$	
<i>T</i> ₁	R(V)	W(V)			τ.,	τ	
<i>T</i> ₂	R(V)				<i>1</i> 1 ←		
<i>T</i> ₁	W(Y)						
<i>T</i> ₂	R(V)		R(Y)	W(Z)			
<i>T</i> ₃		W(V)					

- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

T_1 W(V) W(V)	
<i>T</i> ₂ R(V)	$^{\prime_1} \smile ^{\prime_2}$
T_1 R(V) W(V)	Т. Д. Т.
<i>T</i> ₂ R(V)	$I_1 \leftarrow I_2$
<i>T</i> ₁ W(Y)	 T T
<i>T</i> ₂ R(V)	R(Y) W(Z)
<i>T</i> ₃ W(V)	<i>T</i> ₃

- The graph has a node for each transaction.
- There is an edge from T₁ to T₂ if there is a conflicting action between T₁ and T₂ in which T₁ occurs first.

$ \frac{\overline{T_1 \ W(V) \ W(V)}}{\overline{T_2 \ R(V)}} $	$T_1 \longrightarrow T_2$
$ \overline{\begin{array}{ccc} T_1 & R(V) & W(V) \\ \overline{T_2} & R(V) \end{array} $	$T_1 \longleftarrow T_2$
$\overline{\begin{array}{ccc} \hline T_1 & W(Y) \\ \hline T_2 & B(Y) & B(Y) & W(Z) \\ \end{array}}$	$T_1 \longrightarrow T_2$
$\frac{T_2 - T_1(V)}{T_3} = W(V)$	τ_3

Checking Conflict-Serializability

A schedule is **conflict-serializable** if and only if there is **no cycle** in the precedence graph!

Checking Conflict-Serializability

A schedule is **conflict-serializable** if and only if there is **no cycle** in the precedence graph!

<i>T</i> ₁	W(V)	W(V)		$T \rightarrow T_{c}$
<i>T</i> ₂	R(V)			$^{\prime_1} \sim ^{\prime_2}$
<i>T</i> ₁	R(V)	W(V)		τ, τ
<i>T</i> ₂	R(V)			$I_1 \longleftarrow I_2$
<i>T</i> ₁		W(Y)		
<i>T</i> ₂	R(V)	F	R(Y) W(Z)	
<i>T</i> ₃	W(V)			T_3

Checking Conflict-Serializability

A schedule is **conflict-serializable** if and only if there is **no cycle** in the precedence graph!

<i>T</i> ₁	W(V) V	V(V)	$\tau \rightarrow \tau_{c}$
<i>T</i> ₂	R(V)		$^{\prime_1} \sim ^{\prime_2}$
<i>T</i> ₁	R(V) V	V(V)	т, т
<i>T</i> ₂	R(V)		$I_1 \longleftarrow I_2$
T_1	V	V(Y)	T T.
<i>T</i> ₂	R(V)	R(Y) W(Z)	$1_1 \longrightarrow 1_2$
<i>T</i> ₃	W(V)		T_3

Schedules 2 and 3 have no cycles in their precedence graph. They are conflict serializable!

If the precedence graph has no cycles, then an equivalent serial schedule is obtained by a **topological sort** of the precedence graph.

If the precedence graph has no cycles, then an equivalent serial schedule is obtained by a **topological sort** of the precedence graph.

$$\begin{array}{c|c} \hline T_1 & W(Y) \\ \hline T_2 & R(V) & R(Y) & W(Z) \\ \hline T_3 & W(V) & & T_3 \end{array} \xrightarrow{T_2} T_2$$

If the precedence graph has no cycles, then an equivalent serial schedule is obtained by a **topological sort** of the precedence graph.



• There is an edge from T_1 to T_2 thus T_1 must be before T_2 .

If the precedence graph has no cycles, then an equivalent serial schedule is obtained by a **topological sort** of the precedence graph.

<i>T</i> ₁	W(Y)		
<i>T</i> ₂	R(V)	R(Y)	W(Z)	$1_1 \longrightarrow 1_2$
<i>T</i> ₃		W(V)		T_3
Ther	e is an edg	e from T_1 to 7	5₂ thus ∃	T_1 must be before T_2 .

• There is an edge from T_2 to T_3 thus T_2 must be before T_3 .

If the precedence graph has no cycles, then an equivalent serial schedule is obtained by a **topological sort** of the precedence graph.

<i>T</i> ₁	٧	V(Y)				
<i>T</i> ₂	R(V)		R(Y)	W(Z)	$I_1 \longrightarrow I_2$	
<i>T</i> ₃		W(V)			T_3	
Ther	e is an e	edge from	T_1 to	T_2 thus	s T_1 must be before T_2	
Ther	e is an e	edge from	T ₂ to	T_3 thus	s T_2 must be before T_3	
 		C 101 11				

The sorting which fulfils these criteria is: T_1 , T_2 , T_3 .

If the precedence graph has no cycles, then an equivalent serial schedule is obtained by a **topological sort** of the precedence graph.

T_1	W(Y)			Γ.	\mathbf{x}			
<i>T</i> ₂ R	(V)	R(Y) V	V(Z)	1	\rightarrow 12			
<i>T</i> ₃	١	N(V)		T_3	∠			
There is	• There is an edge from T_1 to T_2 thus T_1 must be before T_2 .							
There is	• There is an edge from T_2 to T_3 thus T_2 must be before T_3 .							
The sorting which fulfils these criteria is: T_1 , T_2 , T_3 .								
This yields the equivalent serial schedule:								
	<i>T</i> ₁ W(Y)						
	-							

<i>T</i> ₂	R(V)	R(Y)	W(Z)	
<i>T</i> ₃				W(V)

Example

Is the following schedule conflict-serializable?

<i>T</i> ₁	R(V)		W(V)	
<i>T</i> ₂		W(V)		
<i>T</i> ₃				W(V)

Example

Is the following schedule conflict-serializable?

<i>T</i> ₁	R(V)		W(V)	
<i>T</i> ₂		W(V)		
<i>T</i> ₃				W(V)

The precedence graph is:

T₁ T₂ T₃

Example

Is the following schedule conflict-serializable?

<i>T</i> ₁	R(V)		W(V)	
<i>T</i> ₂		W(V)		
<i>T</i> ₃				W(V)

The precedence graph is:


Is the following schedule conflict-serializable?

<i>T</i> ₁	R(V)		W(V)	
<i>T</i> ₂		W(V)		
<i>T</i> ₃				W(V)

The precedence graph is:



There is a cycle, thus not conflict-serializable!

Is the following schedule conflict-serializable?

<i>T</i> ₁	R(V)		W(V)	
<i>T</i> ₂		W(V)		
<i>T</i> ₃				W(V)

The precedence graph is:



There is a cycle, thus not conflict-serializable!

However, the schedule is serializable: T_1 , T_2 , T_3 ! The writes of T_1 and T_2 are **blind writes**.

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	R(Y)		W(V)		
<i>T</i> ₃		W(V)				W(Z)

Is this following schedule conflict-serializable?

<i>T</i> ₁	R(V)			R(Z)		R(Y)	
<i>T</i> ₂	F	R(Y)			W(V)		
<i>T</i> ₃		۷	V(V)				W(Z)

Is this following schedule conflict-serializable?

The precedence graph is:

$$T_1 \qquad T_2$$

 T_3

<i>T</i> ₁	R(V)	R(Z)	R(Y)	
<i>T</i> ₂	R(Y	1	W(V)	
<i>T</i> ₃		W(V)		W(Z)

Is this following schedule conflict-serializable?



<i>T</i> ₁	R(V)	R(Z)	R(Y)	
<i>T</i> ₂	R(Y)		W(V)	
<i>T</i> ₃		W(V)		W(Z)

Is this following schedule conflict-serializable?



<i>T</i> ₁	R(V)			R(Z)		R(Y)	
<i>T</i> ₂		R(Y)			W(V)		
<i>T</i> ₃			W(V)				W(Z)

Is this following schedule conflict-serializable?



<i>T</i> ₁	R(V)	R(Z)	R(Y)	
<i>T</i> ₂	R(Y)		W(V)	
<i>T</i> ₃		W(V)		W(Z)

Is this following schedule conflict-serializable?



<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	R(Y)		W(V)		
<i>T</i> ₃		W(V)				W(Z)

Is this following schedule conflict-serializable?

The precedence graph is:



There is no cycle, thus the schedule is conflict-serializable!

<i>T</i> ₁	R(V)	R(Z)	R(Y)				
<i>T</i> ₂						R(Y)	W(V)
<i>T</i> ₃				W(V)	W(Z)		

Transactions :: Strategies for Concurrency Control

Ensuring Serializability

So far, we have seen a sufficient condition that allows us to check whether a schedule is serializable.

But how to ensure serializability during runtime?

Ensuring Serializability

So far, we have seen a sufficient condition that allows us to check whether a schedule is serializable.

But how to ensure serializability during runtime?

Challenge: the system does not know in advance which transactions will run and which items they will access.

Ensuring Serializability

So far, we have seen a sufficient condition that allows us to check whether a schedule is serializable.

But how to ensure serializability during runtime?

Challenge: the system does not know in advance which transactions will run and which items they will access.

Different strategies for ensuring serializability

- 1. Pessimistic
 - Iock-based concurrency control (needs deadlock detection)
 - timestamp based concurrency control (not discussed here)

2. Optimistic

- read-set/write-set tracking
- validation before commit (transaction might abort)
- 3. Multi-version techniques
 - less concurrency control overhead for read-only queries

Transactions :: Two Phase Locking

Pessimistic: Lock-based Concurrency Control

Lock-based concurrency control

Transactions must **lock** objects before using them.

Pessimistic: Lock-based Concurrency Control

Lock-based concurrency control

Transactions must **lock** objects before using them.

Types of locks

- Shared lock (S-lock) is acquired on Y before reading Y. Many transactions can hold a shared lock on Y.
- Exclusive lock (X-lock) is acquired on Y before writing Y.

A transaction can hold an exclusive lock on Y only if no other transaction holds any lock on Y.

If a transaction has an X-lock on Y it can also read Y.

Pessimistic: Lock-based Concurrency Control

Schedule with explicit lock actions

<i>T</i> ₁		X(B) W(B)	U(B)			
<i>T</i> ₂	S(A) R(A)			X(B) W(B)	U(A)	U(B)

Here we use the following abbreviations:

- S(A) = shared lock on A
- X(A) = exclusive lock on A
- U(A) = unlock A, or if more precision is needed
 - US(A) = unlock shared lock on A
 - UX(A) = unlock exclusive lock on A

2 Phase Locking Protocol

2 PL is the concurrency control protocol used in most DBMSs:

2 Phase Locking (2 PL)

Each transaction must get,

- an S-lock on an object before reading it, and
- an X-lock on an object before writing it.

A transaction cannot get new locks once it releases any lock.

2 Phase Locking Protocol

2 PL is the concurrency control protocol used in most DBMSs:

2 Phase Locking (2 PL)

Each transaction must get,

- an S-lock on an object before reading it, and
- an X-lock on an object before writing it.

A transaction cannot get new locks once it releases any lock.



2 Phase Locking Protocol

2 PL is the concurrency control protocol used in most DBMSs:

2 Phase Locking (2 PL)

Each transaction must get,

- an S-lock on an object before reading it, and
- an X-lock on an object before writing it.

A transaction cannot get new locks once it releases any lock.



Theorem

Any schedule that confirms to 2 PL is conflict-serializable.

$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	Whic	ch of t	he fol	lowin	g con	forms	to th	e 2PL	. proto	ocol?
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	<i>T</i> ₁				X(B)	W(B)	U(B)			
$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	<i>T</i> ₂	S(A)	R(A)	U(A)				X(B)	W(B)	U(B)
T_2 S(A) X(B) R(A) W(B) U(A) U(B) T_1 X(B) W(B) U(B) U(B) U(B) U(B) U(B) T_2 S(A) R(A) U(A) X(B) W(B) U(B) T_1 X(B) W(B) U(B) U(B) U(B) U(B)	<i>T</i> ₁				X(B)	W(B)	U(B)			
T1 X(B) W(B) U(B) T2 S(A) R(A) U(A) X(B) W(B) U(B) T1 X(B) W(B) U(B) U(A) X(B) W(B) U(B)	<i>T</i> ₂	S(A)	X(B)	R(A)				W(B)	U(A)	U(B)
T2 S(A) R(A) U(A) X(B) W(B) U(B) T1 X(B) W(B) U(B) U(B) U(B) U(B)	<i>T</i> ₁			X(B)	W(B)		U(B)			
<i>T</i> ₁ X(B) W(B) U(B)	<i>T</i> ₂	S(A)	R(A)			U(A)		X(B)	W(B)	U(B)
	<i>T</i> ₁			X(B)	W(B)	U(B)				
T_2 S(A) R(A) X(B) U(A) W(B) U(B)	<i>T</i> ₂	S(A)	R(A)				X(B)	U(A)	W(B)	U(B)

Whic	h of t	he fol	lowin	g con	forms	to th	e 2PL	. proto	ocol?	
<i>T</i> ₁				X(B)	W(B)	U(B)				Nia
<i>T</i> ₂	S(A)	R(A)	U(A)				X(B)	W(B)	U(B)	INO
<i>T</i> ₁				X(B)	W(B)	U(B)				
<i>T</i> ₂	S(A)	X(B)	R(A)				W(B)	U(A)	U(B)	
<i>T</i> ₁			X(B)	W(B)		U(B)				
<i>T</i> ₂	S(A)	R(A)			U(A)		X(B)	W(B)	U(B)	
<i>T</i> ₁			X(B)	W(B)	U(B)					
<i>T</i> ₂	S(A)	R(A)				X(B)	U(A)	W(B)	U(B)	

Whic	ch of t	he fol	lowin	g con	forms	to th	e 2PL	. proto	ocol?	
<i>T</i> ₁				X(B)	W(B)	U(B)				NIa
<i>T</i> ₂	S(A)	R(A)	U(A)				X(B)	W(B)	U(B)	INO
<i>T</i> ₁				X(B)	W(B)	U(B)				NI-
<i>T</i> ₂	S(A)	X(B)	R(A)				W(B)	U(A)	U(B)	INO
<i>T</i> ₁			X(B)	W(B)		U(B)				
<i>T</i> ₂	S(A)	R(A)			U(A)		X(B)	W(B)	U(B)	
<i>T</i> ₁			X(B)	W(B)	U(B)					
<i>T</i> ₂	S(A)	R(A)				X(B)	U(A)	W(B)	U(B)	

Whic	ch of t	he fol	lowin	g con	forms	to th	e 2PL	. proto	ocol?	
<i>T</i> ₁				X(B)	W(B)	U(B)				NIa
<i>T</i> ₂	S(A)	R(A)	U(A)				X(B)	W(B)	U(B)	INO
<i>T</i> ₁				X(B)	W(B)	U(B)				
<i>T</i> ₂	S(A)	X(B)	R(A)				W(B)	U(A)	U(B)	NO
<i>T</i> ₁			X(B)	W(B)		U(B)				
<i>T</i> ₂	S(A)	R(A)			U(A)		X(B)	W(B)	U(B)	NO
T_1			X(B)	W(B)	U(B)					
<i>T</i> ₂	S(A)	R(A)	. ,	. ,		X(B)	U(A)	W(B)	U(B)	
-										

Which of the following conforms to the 2PL protocol?	
<i>T</i> ₁ X(B) W(B) U(B)	NIa
$T_2 S(A) R(A) U(A) \qquad \qquad X(B) W(B) U(B)$	NO
T ₁ X(B) W(B) U(B)	Na
$T_2 S(A) X(B) R(A) \qquad \qquad W(B) U(A) U(B)$	NO
<i>T</i> ₁ X(B) W(B) U(B)	
$T_2 S(A) R(A) \qquad \qquad U(A) \qquad X(B) W(B) U(B)$	NO
T_1 X(B) W(B) U(B)	
T_2 S(A) R(A) X(B) U(A) W(B) U(B)	res

Example: ATM Transaction

Conci	Concurrent ATM Transaction				
	Transaction 1	Transaction 2	DB state		
	slock(<i>account</i>) read(<i>account</i>) unlock(<i>account</i>)		1200		
		slock(<i>account</i>) read(<i>account</i>) unlock(<i>account</i>)			
	xlock(<i>account</i>) write(<i>account</i>) unlock(<i>account</i>)		1100		
		xlock(account) write(account) unlock(account)	1000		

⁴ Once a lock has been released, no new lock can be acquired.

To comply with the 2PL, the ATM transaction must not acquire new locks after a lock has been released.

ATM withdrawal with 2 Phase Locking

- 1. xlock(account)
- 2. $balance \leftarrow read(account)$
- 3. *balance* \leftarrow *balance* -100
- 4. write(*account*, *balance*)
- 5. unlock(*account*)

Concur	Concurrent ATM Transaction				
	Transaction 1	Transaction 2	DB state		
	xlock(<i>account</i>) read(<i>account</i>)		1200		
	write(<i>account</i>) unlock(<i>account</i>)	xlock(<i>account</i>)	1100		
		xlock(account) read(account) write(account) unlock(account)	900		

Transaction 2 blocked until transaction 1 releases the lock.

Note: now both transactions are correctly executed!

Transactions :: Two Phase Locking - Deadlock Handling

Like many lock-based protocols, 2PL has the risk of deadlocks.

Like many lock-based protocols, 2PL has the risk of deadlocks.

Transaction 1	Transaction 2

Like many lock-based protocols, 2PL has the risk of deadlocks.

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
:	
do something	
:	

Like many lock-based protocols, 2PL has the risk of deadlocks.

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
:	$\operatorname{xlock}(B)$
do something	:
:	do something
	÷

Like many lock-based protocols, 2PL has the risk of deadlocks.

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
:	$\operatorname{xlock}(B)$
do something	:
:	do something
lock(B)	:

Like many lock-based protocols, 2PL has the risk of deadlocks.

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
:	xlock(B)
do something	÷
:	do something
lock(B)	:
(waiting for T_2 to unlock B)	

Like many lock-based protocols, 2PL has the risk of deadlocks.

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
:	$\operatorname{xlock}(B)$
do something	:
:	do something
lock(B)	:
(waiting for T_2 to unlock B)	lock(A)
Deadlocks

Like many lock-based protocols, 2PL has the risk of deadlocks.

A Deadlock Situation

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
:	$\operatorname{xlock}(B)$
do something	:
:	do something
lock(B)	:
(waiting for T_2 to unlock B)	lock(A)
	(waiting for T_1 to unlock A)

Deadlocks

Like many lock-based protocols, 2PL has the risk of deadlocks.

A Deadlock Situation

Transaction 1	Transaction 2
$\operatorname{xlock}(A)$	
÷	$\operatorname{xlock}(B)$
do something	:
:	do something
lock(B)	:
(waiting for T_2 to unlock B)	lock(A)
	(waiting for T_1 to unlock A)

Both transactions would wait for each other **indefinitely**. We need to detect deadlocks!

Deadlock Handling via Wait-for-Graphs

Deadlock Detection via Wait-for-Graphs

The system maintains a waits-for-graph:

- Nodes of the graph are transactions.
- Edge $T_1 \rightarrow T_2$ means T_1 is blocked by a lock held by T_2 . Hence T_1 waits for T_2 to release the lock.

Deadlock Handling via Wait-for-Graphs

Deadlock Detection via Wait-for-Graphs

The system maintains a waits-for-graph:

- Nodes of the graph are transactions.
- Edge $T_1 \rightarrow T_2$ means T_1 is blocked by a lock held by T_2 . Hence T_1 waits for T_2 to release the lock.

The system checks periodically for cycles in the graph:

If a cycle is detected, then the deadlock is resolved by aborting one or more transactions.

Deadlock Handling via Wait-for-Graphs

Deadlock Detection via Wait-for-Graphs

The system maintains a waits-for-graph:

- Nodes of the graph are transactions.
- Edge $T_1 \rightarrow T_2$ means T_1 is blocked by a lock held by T_2 . Hence T_1 waits for T_2 to release the lock.

The system checks periodically for cycles in the graph:

If a cycle is detected, then the deadlock is resolved by aborting one or more transactions.

Selecting the victim is a challenge

- Aborting young transactions might lead to starvation. The same transaction may be cancelled again and again.
- Aborting old transactions may cause a lot of computational investment to be thrown away.

Deadlock Detection via Timeout

Let transactions block on a lock request only for a limited time.

After timeout, assume a deadlock has occurred and abort T.

Transactions :: Two Phase Locking - Cascading Rollbacks



Assume that:

- T₁ is aborted (due to a conflict with another transaction)
- T₂ tries to commit

What is the problem here?



Assume that:

- T₁ is aborted (due to a conflict with another transaction)
- T₂ tries to commit

What is the problem here?

• T_2 has read a value written by T_1 .



Assume that:

- T₁ is aborted (due to a conflict with another transaction)
- T₂ tries to commit
- What is the problem here?
 - T_2 has read a value written by T_1 .
 - Thus if T₁ is aborted, then T₂ needs to be aborted too. The commit will result in an abort.



What happens here?



What happens here?

Note: T_2 and T_3 cannot commit until the fate of T_1 is known.



What happens here?

Note: T_2 and T_3 **cannot commit** until the fate of T_1 is known. When T_1 aborts:

- **T**₂ and T_3 have already read data written by T_1 (dirty read)
- T₂ and T₃ need to be rolled back too (cascading roll back)



What happens here?

Note: T_2 and T_3 cannot commit until the fate of T_1 is known. When T_1 aborts:

*T*₂ and *T*₃ have already read data written by *T*₁ (dirty read)
*T*₂ and *T*₃ need to be rolled back too (cascading roll back)
Since *T*₃ is aborted, *T*₄ needs to be aborted as well.

Definition: Recoverable Schedule Delay commits:

If T₂ reads a value written by T₁, the commit of T₂ must be delayed until after the commit of T₁.

Definition: Recoverable Schedule Delay commits:

If T₂ reads a value written by T₁, the commit of T₂ must be delayed until after the commit of T₁.

Schedules should always be recoverable!

Definition: Recoverable Schedule Delay commits:

If T₂ reads a value written by T₁, the commit of T₂ must be delayed until after the commit of T₁.

Schedules should always be recoverable!

Definition: Cascadeless Schedule

Delay reads: only read values produced by already committed transactions.

If T₂ reads a value written by T₁, then the read is delayed until after the commit of T₁.

Definition: Recoverable Schedule Delay commits:

If T₂ reads a value written by T₁, the commit of T₂ must be delayed until after the commit of T₁.

Schedules should always be recoverable!

Definition: Cascadeless Schedule

Delay reads: only read values produced by already committed transactions.

If T₂ reads a value written by T₁, then the read is delayed until after the commit of T₁.

No dirty reads, thus abort (rollback) does not cascade!

Definition: Recoverable Schedule Delay commits:

If T₂ reads a value written by T₁, the commit of T₂ must be delayed until after the commit of T₁.

Schedules should always be recoverable!

Definition: Cascadeless Schedule

Delay reads: only read values produced by already committed transactions.

If T₂ reads a value written by T₁, then the read is delayed until after the commit of T₁.

No dirty reads, thus abort (rollback) does not cascade!

All cascadeless schedules are recoverable.

<i>T</i> ₁	X(A) S(B) W(A) U(A)	R(B) U(B) Commit				
<i>T</i> ₂	S(A)R(A) X(A)W(A	A) U(A) Commit				
Is this schedule cascadeless?						
ls tl	nis schedule recoverable?					

<i>T</i> ₁	X(A) S(B)W(A) U(A)	R(B) U(B) Commit				
<i>T</i> ₂	S(A)R(A) X(A)W(A) U(A) Commit					
ls th ∎	his schedule cascadeless? No If the commit of T_1 fails, then T_2 need	ls to be rolled back.				
ls th	nis schedule recoverable?					

<i>T</i> ₁	X(A) S(B) W(A) U(A)	R(B) U(B) Commit
<i>T</i> ₂	S(A) R(A) X(A)	W(A) U(A) Commit
الم ما		
is th	IIS SCHEQUIE CASCADEIESS? INO	
	If the commit of T_1 fails, then T_2	needs to be rolled back.
ls th	is schedule recoverable? No	
_	The commit of T is not delayed	$ \dots $

• The commit of T_2 is not delayed until after commit of T_1 .

<i>T</i> ₁	X(A) S(B) W (A) U(A) R	(B) U(B) Commit
<i>T</i> ₂	S(A)R(A) X(A)W(A) U(A) (Commit
ls th ∎	his schedule cascadeless? No If the commit of <i>T</i> ₁ fails, then <i>T</i> ₂ needs to b	e rolled back.
ls th	his schedule recoverable? No	
	The commit of T ₂ is not delayed until after	commit of T_1 .

Not	Cascadeless, But Recoverable	
<i>T</i> ₁	X(A) S(B) W(A) U(A)	R(B) U(B) Commit
<i>T</i> ₂	S(A)R(A) X(A)W(A) U	(A) Commit

Transactions :: Strict & Preclaiming Two Phase Locking

Strict 2 Phase Locking Protocol

Strict 2 Phase Locking

Like in 2 PL, each transaction must get,

- an S-lock on an object before reading it, and
- an X-lock on an object before writing it.

But moreover:

A transaction releases all locks only when the transaction is completed (i.e. when performing commit/rollback).



Strict 2 Phase Locking Protocol

Strict 2 Phase Locking

Like in 2 PL, each transaction must get,

- an S-lock on an object before reading it, and
- an X-lock on an object before writing it.

But moreover:

A transaction releases all locks only when the transaction is completed (i.e. when performing commit/rollback).



This protocol is **cascadeless**, avoids cascading aborts.

But there still are deadlocks!

Preclaiming 2 Phase Locking Protocol



Preclaiming 2 Phase Locking Protocol



All needed locks are declared at the beginning of the transaction.

Advantage: **No deadlocks!** But rollbacks are cascading.

Preclaiming 2 Phase Locking Protocol



All needed locks are declared at the beginning of the transaction.

Advantage: **No deadlocks!** But rollbacks are cascading.

Disadvantage

Not applicable in multi-query transactions. 쉵

(Queries might depend on the results of the previous queries)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂		R(Y)		W(V)		
<i>T</i> ₃		V	/(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	R(V)			R(Z)		R(Y)	
<i>T</i> ₂		R(Y)			W(V)		
<i>T</i> ₃			W(V)				W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

T_1	S(V)R(V)	S(Z)R(Z)	S(Y)R(Y)	
<i>T</i> ₂	S(Y)R(Y)	X(V)W(V)	
<i>T</i> 3	X(V)W(V)	2	X(Z)W(Z)

<i>T</i> ₁	R(V)			R(Z)		R(Y)	
<i>T</i> ₂		R(Y)			W(V)		
<i>T</i> ₃			W(V)				W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(V)R(V)	S(Z)R(Z)	S(Y)R(Y)	
<i>T</i> ₂	S(Y)R(Y)	X(V	W(V)	
<i>T</i> ₃	X(V)W	(V)	X(Z)V	V(Z)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> 2	R	(Y)		W(V)		
<i>T</i> ₃		W(V))			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

T ₁ S(V)R(V)		S(Z)R(Z)	S(Y)R(Y)	
<i>T</i> ₂	S(Y)R(Y)	X(V)W(V)	
<i>T</i> 3	X(V)W(V)	X(Z	Z)W(Z)

<i>T</i> ₁	R(V)		R	(Z)	R(Y)	
<i>T</i> ₂	F	R(Y)		W(V)	
<i>T</i> ₃		١	N(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z)	R(Y)	
<i>T</i> ₂	S(Y)R(Y)		X(V)W(V)	
<i>T</i> ₃		X(V)W(V)		X(Z)W(Z)

<i>T</i> ₁	R(V)		R	(Z)	R(Y)	
<i>T</i> ₂	F	R(Y)		W(V)	
<i>T</i> ₃		١	N(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z)	R(Y)	
<i>T</i> ₂	S(Y)R(Y)		<mark>X(V)</mark> W(V)	
<i>T</i> 3		X(V)W(V)		X(Z)W(Z)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	R()	')		W(V)		
<i>T</i> ₃		W(V)				W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z)	I	R(Y)		
<i>T</i> ₂	S(Y)R(Y))	X(V)W(V)			
<i>T</i> 3		X(V)W(V)	U(V)	X(Z)W(Z)		
<i>T</i> ₁	R(V)		R(Z)		R(Y)	
-----------------------	------	------------	------	------	------	------
<i>T</i> ₂	R(`	')		W(V)		
<i>T</i> ₃		W(V)				W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z)	F	R(Y)
<i>T</i> ₂	S(Y)R(Y))	X(V)W(V)	
<i>T</i> ₃		X(V)W(V)	U(V)	X(Z)W(Z)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	R(`	')		W(V)		
<i>T</i> ₃		W(V)				W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R	(Z)	R(Y)
<i>T</i> ₂	S(Y)R(Y	.)	X(V))W(V)
<i>T</i> ₃		X(V)W(V)	X(Z) U(V)	W(Z)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	F	R(Y)		W(V)		
<i>T</i> ₃		W	V(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R	(Z)	R(Y)
<i>T</i> ₂	S(Y)R(Y	.)	X(V)	W(V)
<i>T</i> ₃		X(V)W(V)	X(Z) U(V)	W(Z)

<i>T</i> ₁	R(V)	R(Z)	R(Y)	
<i>T</i> ₂	R(Y)		W(V)	
<i>T</i> ₃		W(V)		W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z) U(Z)	R(Y)
<i>T</i> ₂	S(Y)R(Y)		X(V)W(V)	
<i>T</i> ₃		X(V)W(V)	X(Z) U(V)	W(Z)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	R(`	')		W(V)		
<i>T</i> ₃		W(V)				W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z) U(Z)	R(Y) U(Y)	
<i>T</i> ₂	S(Y)R(Y)		X(V)W(V) U(V)	Y)
<i>T</i> 3	Х	X(V)W(V) $X(Z)U(V)$	W(Z) U(Z	Z)

<i>T</i> ₁	R(V)		R(Z)		R(Y)	
<i>T</i> ₂	F	R(Y)		W(V)		
<i>T</i> ₃		W	V(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z) U(Z)	R(Y) U(Y)
<i>T</i> ₂	S(Y)R(Y)		X(V)W(V) U(VY)
<i>T</i> ₃	>	X(V)W(V) X(Z) U(V	/) W(Z) U(Z)

Impossible with Strict 2 PL: T_1 must hold lock on V until commit. This schedule is 2 PL !

<i>T</i> ₁	R(V)		R	(Z)	R(Y)	
<i>T</i> ₂	F	R(Y)		W(V)	
<i>T</i> ₃		١	N(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z) U(Z)	F	R(Y)	U(Y)
T ₂	S(Y)R(Y)		X(V)W(V)		U(VY)
<i>T</i> ₃	>	X(V)W(V)	X(Z) U(V)	W(Z)	U(Z)

Impossible with Strict 2 PL: T_1 must hold lock on V until commit.

This schedule is 2 PL !

Can it be achieved using Preclaiming 2 PL?

<i>T</i> ₁	R(V)		R	(Z)	R(Y)	
<i>T</i> ₂	F	R(Y)		W(V)	
<i>T</i> ₃		١	N(V)			W(Z)

- Can it be achieved using 2 PL?
- Can it be achieved using Strict 2 PL?

Add the corresponding lock and unlock statements.

<i>T</i> ₁	S(VZY) R(V) U(V)	R(Z) U(Z)	F	R(Y)	U(Y)
<i>T</i> ₂	S(Y)R(Y)		X(V)W(V)		U(VY)
<i>T</i> ₃	>	X(V)W(V)	X(Z) U(V)	W(Z)	U(Z)

Impossible with Strict 2 PL: T_1 must hold lock on V until commit.

This schedule is 2 PL !

Can it be achieved using Preclaiming 2 PL? No

Transactions :: Granularity of Locking



Idea: multi-granularity locking...

Decide the granularity of locks held **for each transaction**. *Depending on the characteristics of the queries.*

Decide the granularity of locks held **for each transaction**. *Depending on the characteristics of the queries.*

For example, acquire a row lock for:

```
Q<sub>1</sub>: row-selecting query (id is a key)
```

select *
from Customers
where id = 42

Decide the granularity of locks held **for each transaction**. *Depending on the characteristics of the queries.*

For example, acquire a row lock for:

```
Q<sub>1</sub>: row-selecting query (id is a key)
```

```
select *
from Customers
where id = 42
```

For example, acquire a table lock for:

*Q*₂: table scan query

select *
from Customers

Decide the granularity of locks held **for each transaction**. *Depending on the characteristics of the queries.*

For example, acquire a row lock for:

```
Q<sub>1</sub>: row-selecting query (id is a key)
```

select *
from Customers
where id = 42

For example, acquire a table lock for:

```
Q<sub>2</sub>: table scan query
```

select *
from Customers

How do such transactions know of each others locks? Note that the locks are on different granularity levels!

Databases use an additional type of locks: intention locks.

- Lock mode intention share (IS)
- Lock mode intention exclusive (IX)

Before introducing S (or X) lock, first IS (or IX) locks on all coarser levels of granularity.

Databases use an additional type of locks: intention locks.

- Lock mode intention share (IS)
- Lock mode intention exclusive (IX)

Before introducing S (or X) lock, first IS (or IX) locks on all coarser levels of granularity.

Extended lock confl	ict m	natr	ix			
		S	Х	IS	IX	
-	S		Х		Х	
	Χ	Х	Х	Х	Х	
	IS		Х			
	IX	Х	Х			

Intention Locks

Multi-granularity Locking Protocol

- Before a granule g can be locked in S (or X) mode, the transaction has to obtain an IS (or IX) lock on **all coarser** granularities that contain g.
- After all intention locks are granted, the transaction can lock g in the announced mode.

Intention Locks

Multi-granularity Locking Protocol

- Before a granule g can be locked in S (or X) mode, the transaction has to obtain an IS (or IX) lock on all coarser granularities that contain g.
- After all intention locks are granted, the transaction can lock g in the announced mode.

The query Q_1 would for example:

- obtain an IS lock on the database
- obtain an IS lock on the table Customers

Afterwards obtain an **S lock** on the **row** with id = 42.

Intention Locks

Multi-granularity Locking Protocol

- Before a granule g can be locked in S (or X) mode, the transaction has to obtain an IS (or IX) lock on all coarser granularities that contain g.
- After all intention locks are granted, the transaction can lock g in the announced mode.

The query Q_1 would for example:

- obtain an IS lock on the database
- obtain an IS lock on the table Customers

Afterwards obtain an **S lock** on the **row** with id = 42.

The query Q_2 would for example:

obtain an IS lock on the database

Afterwards obtain an **S lock** on the **table** Customers.

Now suppose an updating query comes in:

Q₃: update request

update Customers set name = 'Pete' where id = 17

Now suppose an updating query comes in:

Q₃: update request

update Customers set name = 'Pete' where id = 17

The query Q_3 will try to:

- obtain an IX lock on the database
- obtain an IX lock on the table Customers

Afterwards obtain an X lock on the row with id = 17.

Now suppose an updating query comes in:

Q₃: update request

update Customers set name = 'Pete' where id = 17

The query Q_3 will try to:

- obtain an IX lock on the database
- obtain an IX lock on the table Customers

Afterwards obtain an X lock on the row with id = 17.

 compatible with Q₁ (no conflict between IS of Q₁ and IX lock of Q₃ on table)

Now suppose an updating query comes in:

Q₃: update request

update Customers set name = 'Pete' where id = 17

The query Q_3 will try to:

- obtain an IX lock on the database
- obtain an IX lock on the table Customers

Afterwards obtain an X lock on the row with id = 17.

- compatible with Q₁
 (no conflict between IS of Q₁ and IX lock of Q₃ on table)
- incompatible with Q₂
 (conflict between S lock of Q₂ and IX lock of Q₃ on table)

Transactions :: Isolation Levels

Isolation Levels

Some degree of **inconsistency** may be acceptable for specific applications to gain **increased concurrency & performance**.

E.g. accept inconsistent read anomaly and be rewarded with improved concurrency. Relaxed consistency guarantees can lead to increased throughput!

Isolation Levels

Some degree of **inconsistency** may be acceptable for specific applications to gain **increased concurrency & performance**.

E.g. accept inconsistent read anomaly and be rewarded with improved concurrency. Relaxed consistency guarantees can lead to increased throughput!

SQL-92 Isolation Levels & Consistency Guarantees						
isolation level	dirty read	non-repeat. read	phantom rows			
read uncommitted	possible	possible	possible			
read committed	not possible	possible	possible			
repeatable read	not possible	not possible	possible			
serializable	not possible	not possible	not possible			

Isolation Levels

Some degree of **inconsistency** may be acceptable for specific applications to gain **increased concurrency & performance**.

E.g. accept inconsistent read anomaly and be rewarded with improved concurrency. Relaxed consistency guarantees can lead to increased throughput!

SQL-92 Isolation Levels & Consistency Guarantees						
isolation level	dirty read	non-repeat. read	phantom rows			
read uncommitted	possible	possible	possible			
read committed	not possible	possible	possible			
repeatable read	not possible	not possible	possible			
serializable	not possible	not possible	not possible			

Different DBMS support different levels of isolation.

Phantom Row Problem

Transaction 1	Transaction 2	Effect
scan relation R	insert new row into R commit	T_1 locks all rows T_2 locks new row T_2 's lock released reads new row too!

Phantom Row Problem

Transaction 1	Transaction 2	Effect
scan relation R	insert new row into R commit	T_1 locks all rows T_2 locks new row T_2 's lock released reads new row too!

Both transactions properly follow the 2 PL protocol!

Nevertheless, T_1 observed an effect caused by T_2 .

- Isolation violated!
- Cause of the problem: *T*₁ can only lock existing rows.

Phantom Row Problem

Transaction 1	Transaction 2	Effect
scan relation R	insert new row into R commit	T_1 locks all rows T_2 locks new row T_2 's lock released reads new row too!

Both transactions properly follow the 2 PL protocol!

Nevertheless, T_1 observed an effect caused by T_2 .

- Isolation violated!
- Cause of the problem: *T*₁ can only lock existing rows.

Solutions

- 1. multi-granularity locking (locking the table)
- 2. declarative locking: key-range or predicate locking

Basic idea: use variations of strict 2 PL.

SQL-92 Isolation Levels

read uncommitted (also dirty read or browse)

Only write locks are acquired. Any row read may be concurrently changed by other transactions.

Basic idea: use variations of strict 2 PL.

SQL-92 Isolation Levels

read uncommitted (also dirty read or browse)
 Only write locks are acquired. Any row read may be concurrently changed by other transactions.

read committed (also cursor stability)

Read locks are held for as long as the application cursor sits on a particular, **current row**. Write locks as usual. Rows may be changed between repeated reads.

Basic idea: use variations of strict 2 PL.

SQL-92 Isolation Levels

read uncommitted (also dirty read or browse)
 Only write locks are acquired. Any row read may be concurrently changed by other transactions.

read committed (also cursor stability)

Read locks are held for as long as the application cursor sits on a particular, **current row**. Write locks as usual. Rows may be changed between repeated reads.

repeatable read

Strict 2 PL locking. Nevertheless, a transaction may read **phantom rows** if it executes an aggregation query twice.

Basic idea: use variations of strict 2 PL.

SQL-92 Isolation Levels

read uncommitted (also dirty read or browse)
 Only write locks are acquired. Any row read may be concurrently changed by other transactions.

read committed (also cursor stability)

Read locks are held for as long as the application cursor sits on a particular, **current row**. Write locks as usual. Rows may be changed between repeated reads.

repeatable read

Strict 2 PL locking. Nevertheless, a transaction may read **phantom rows** if it executes an aggregation query twice.

serializable

Strict 2 PL + multi-granularity locking. No phantom rows.

Transactions :: SQL Transaction Control

SQL Transaction Control

SQL Transaction Control

- set autocommit on/off
 - on: each SQL query is one transaction
- start transaction
- commit
- rollback
- set transaction isolation level ...
SQL Transaction Control

SQL Transaction Control

- set autocommit on/off
 - on: each SQL query is one transaction
- start transaction
- commit
- rollback
- set transaction isolation level ...

Many applications do not need full serializability Selecting a weaker, yet acceptable isolation level is important part of **database tuning**.

Transactions :: Optimistic Concurrency Control

Up to now we have seen **pessimistic** concurrency control:

- assume that transactions will conflict
- protect the database integrity by locks and lock protocols

Up to now we have seen **pessimistic** concurrency control:

- assume that transactions will conflict
- protect the database integrity by locks and lock protocols

Optimistic concurrency control

- hope for the best
- Iet transactions freely proceed with read/write operations
- only at commit, check that no conflicts have happened

Up to now we have seen **pessimistic** concurrency control:

- assume that transactions will conflict
- protect the database integrity by locks and lock protocols

Optimistic concurrency control

- hope for the best
- Iet transactions freely proceed with read/write operations
- only at commit, check that no conflicts have happened

Rationale:

- non-serializable conflicts are not that frequent
- save the locking overhead
- only invest effort if really required

Under **optimistic concurrency control**, transactions proceed in **three phases**:

1. Read phase:

Execute transaction, but do **not** write data back to disk. Collect updates in the transaction's **private workspace**.

2. Validation phase:

When the transaction wants to **commit**, the DBMS test whether its execution was correct (only acceptable conflicts happened). If not, **abort** the transaction.

3. Write phase:

Transfer data from private workspace into database.

Note: phases 2 and 3 are performed in a uninterruptible critical section (also called val-write phase).

Optimistic Concurrency Control: Validation

Validation is typically implemented by maintaining:

- a read set $RS(T_k)$ (attributes read by T_k), and
- a write set WS(T_k) (attributes written by T_k)

for every transaction T_k .

Optimistic Concurrency Control: Validation

Validation is typically implemented by maintaining:

- a read set $RS(T_k)$ (attributes read by T_k), and
- a write set WS(T_k) (attributes written by T_k)

for every transaction T_k .

Backward-oriented optimistic concurrency control (BOCC)

On commit, compare T_k against all **committed** transactions T_i . Check **succeeds** if

 T_i committed before T_k started **or** $RS(T_k) \cap WS(T_i) = \emptyset$ If the check fails, abort T_k .

Optimistic Concurrency Control: Validation

Validation is typically implemented by maintaining:

- a read set $RS(T_k)$ (attributes read by T_k), and
- a write set WS(T_k) (attributes written by T_k)

for every transaction T_k .

Backward-oriented optimistic concurrency control (BOCC)

On commit, compare T_k against all **committed** transactions T_i . Check **succeeds** if

 T_i committed before T_k started **or** $RS(T_k) \cap WS(T_i) = \emptyset$ If the check fails, abort T_k .

Forward-oriented optimistic concurrency control (FOCC)

On commit, compare T_k against all **running** transactions T_i . Check **succeeds** if

 $WS(T_k) \cap RS(T_i) = \emptyset$

If the check fails, abort T_k or T_i .

Transactions :: Multiversion Concurrency Control

Is this schedule serializable?										
Т	- 1	R(X)	W(X)			R(Y)	W(Z)			
T	- 2			R(X)	W(Y)					

Is this schedule serializable?								
	<i>T</i> ₁	R(X)	W(X)			R(Y)	W(Z)	
	<i>T</i> ₂			R(X)	W(Y)			
No								

Is this schedule serializable?									
	<i>T</i> ₁	R(X)	W(X)			R(Y)	W(Z)	-	
	<i>T</i> ₂			R(X)	W(Y)			_	
No									

But what if we had a copy of the old values available?

Is this schedule serializable?									
	<i>T</i> ₁	R(X)	W(X)			R(Y)	W(Z)		
	<i>T</i> ₂			R(X)	W(Y)				
No									

But what if we had a copy of the old values available?

Then we could do:

Mult	i-ver	sion						
	<i>T</i> ₁	R(X)	W(X)			R(Y-old)	W(Z)	-
	<i>T</i> ₂			R(X)	W(Y)			-

Is this schedule serializable?								
	<i>T</i> ₁	R(X)	W(X)			R(Y)	W(Z)	-
	<i>T</i> ₂			R(X)	W(Y)			
No								

But what if we had a copy of the old values available?

Then we could do:

111011										
Multi-version										
	<i>T</i> ₁	R(X)	W(X)			R(Y-old)	W(Z)			
	<i>T</i> ₂			R(X)	W(Y)					
This is can be serialised to:										
	<i>T</i> ₁	R(X)	W(X)	R(Y)	W(Z)					
	T_2					R(X)	W(Y)			

Multiversion Concurrency Control

Multiple versions of the data are stored. The isolation level determines what version a transaction sees.

Multiversion Concurrency Control

Multiple versions of the data are stored. The isolation level determines what version a transaction sees.

Common Isolation Level: Snapshot Isolation

Each transaction sees a consistent snapshot of the database that corresponds to the state at the moment it started.

Multiversion Concurrency Control

Multiple versions of the data are stored. The isolation level determines what version a transaction sees.

Common Isolation Level: Snapshot Isolation

Each transaction sees a consistent snapshot of the database that corresponds to the state at the moment it started.

Read-only transactions never need to be blocked!

- might see outdated, but consistent version of the data
- as if the entire query happened at the moment it started

Multiversion Concurrency Control

Multiple versions of the data are stored. The isolation level determines what version a transaction sees.

Common Isolation Level: Snapshot Isolation

Each transaction sees a consistent snapshot of the database that corresponds to the state at the moment it started.

Read-only transactions never need to be blocked!

- might see outdated, but consistent version of the data
- as if the entire query happened at the moment it started

With snapshot isolation:

- Read-only transactions do not have to lock anything!
- Transactions conflict if they write the same object:
 - Pessimistic concurrency control: only writes are locked
 - Optimistic concurrency control: only write-sets interesting

Disadvantages:

- versioning requires space and management overhead
- update transactions still need concurrency control!

Snapshot isolation does not guarantee serializability! But

- the anomalies
 - dirty read,
 - unrepeatable read,
 - phantom rows

do not occur.

however, write skew occurs!

Disadvantages:

- versioning requires space and management overhead
- update transactions still need concurrency control!

Snapshot isolation does not guarantee serializability! But

- the anomalies
 - dirty read,
 - unrepeatable read,
 - phantom rows

do not occur.

however, write skew occurs!

Snapshot isolation is used in the Oracle SQL Server. (Oracle has no real serializability).

- Constraint: X + Y < 2
- Initially: X = 0 and Y = 0

- Constraint: *X* + *Y* < 2
- Initially: X = 0 and Y = 0
- $T_1: X = X + 1$; it sees X = 1 and Y = 0 and commits

- Constraint: *X* + *Y* < 2
- Initially: X = 0 and Y = 0
- $T_1: X = X + 1$; it sees X = 1 and Y = 0 and commits
- T_2 : Y = Y + 1; it sees X = 0 and Y = 1 and commits

- Constraint: *X* + *Y* < 2
- Initially: X = 0 and Y = 0
- $T_1: X = X + 1$; it sees X = 1 and Y = 0 and commits
- T_2 : Y = Y + 1; it sees X = 0 and Y = 1 and commits
- T_1 and T_2 have an empty write intersection (**no conflict**).

- Constraint: *X* + *Y* < 2
- Initially: X = 0 and Y = 0
- $T_1: X = X + 1$; it sees X = 1 and Y = 0 and commits
- T_2 : Y = Y + 1; it sees X = 0 and Y = 1 and commits
- *T*₁ and *T*₂ have an empty write intersection (**no conflict**).
- End result: X = 1 and Y = 1

- Constraint: *X* + *Y* < 2
- Initially: X = 0 and Y = 0
- $T_1: X = X + 1$; it sees X = 1 and Y = 0 and commits
- T_2 : Y = Y + 1; it sees X = 0 and Y = 1 and commits
- *T*₁ and *T*₂ have an empty write intersection (**no conflict**).
- End result: X = 1 and $Y = 1 \oint X + Y \ge 2$

Write Skew Anomaly

- Constraint: *X* + *Y* < 2
- Initially: X = 0 and Y = 0
- $T_1: X = X + 1$; it sees X = 1 and Y = 0 and commits
- T_2 : Y = Y + 1; it sees X = 0 and Y = 1 and commits
- *T*₁ and *T*₂ have an empty write intersection (**no conflict**).
- End result: X = 1 and $Y = 1 \oint X + Y \ge 2$

This problem does not occur if this is a check constraint comparing values of the same row since the finest locking granularity are rows.

Therefore write skew anomalies occur with complex assertions that involve multiple tuples.

Transactions :: Optimising Performance

Optimising Performance

Suppose you have a typical log of queries for your database.

Optimising Performance

Suppose you have a typical log of queries for your database.

For each query in the log:

- Analyse average time and variance for this type of query.
 - Long delays or frequent aborts may indicate contention.
- Is it is a read-only or updating query?
 - Compute the read-sets and write-sets.
 - Will it require row or table locks? Shared or exclusive?

How do read- and write-sets of the different queries intersect?

What is the chance of conflicts? (delays/rollbacks)

Optimising Performance

Suppose you have a typical log of queries for your database.

For each query in the log:

- Analyse average time and variance for this type of query.
 - Long delays or frequent aborts may indicate contention.
- Is it is a read-only or updating query?
 - Compute the read-sets and write-sets.
 - Will it require row or table locks? Shared or exclusive?

How do read- and write-sets of the different queries intersect?

What is the chance of conflicts? (delays/rollbacks)

Once you understand your query workload, you might improve performance by:

- Rewriting queries to have smaller read- and write-sets.
- Change scheduling of queries to reduce contention.
 E.g. rewrite applications to do large aggregation queries at night.
- Use a different isolation level for the queries.