

Carnegie Mellon Univ.
Dept. of Computer Science
15-415/615 - DB Applications

C. Faloutsos – A. Pavlo

Lecture#21: Concurrency Control – Part 1
(R&G ch. 17)

Last Class

- Introduction to Transactions
- ACID
- Concurrency Control
- Crash Recovery

Last Class

- For **Isolation** property, serial execution of transactions is safe but slow
 - We want to find schedules equivalent to serial execution but allow interleaving.
- The way the DBMS does this is with its *concurrency control* protocol.


Today's Class

- Serializability
- Two-Phase Locking
- Deadlocks
- Lock Granularities

Formal Properties of Schedules

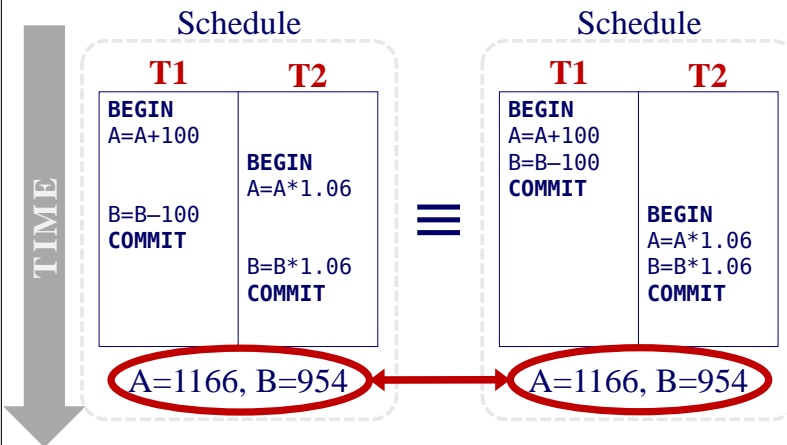
- **Serializable Schedule:** A schedule that is *equivalent* to some serial execution of the transactions.
- **Note:** If each transaction preserves consistency, every serializable schedule preserves consistency.

Example

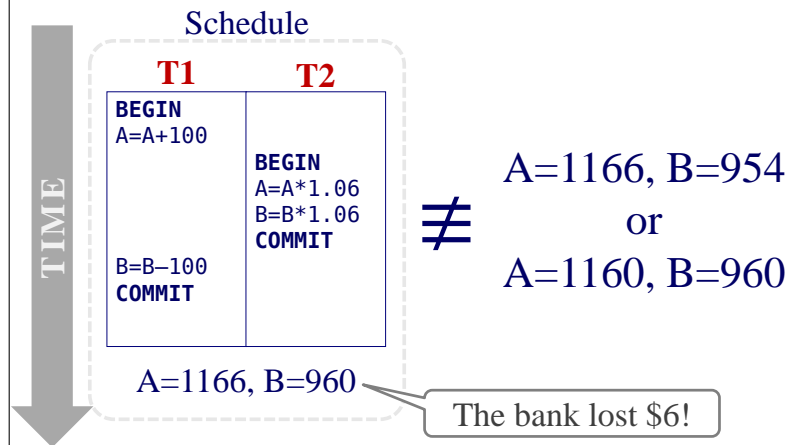


- Consider two txns:
 - T1 transfers \$100 from B's account to A's
 - T2 credits both accounts with 6% interest.
- Assume at first A and B each have \$1000.

Interleaving Example (Good)



Interleaving Example (Bad)



Formal Properties of Schedules

- There are different levels of serializability:
 - **Conflict Serializability** All DBMSs support this.
 - **View Serializability**

This is harder but allows for more concurrency.

Nobody does this.

Conflicting Operations

- We need a formal notion of equivalence that can be implemented efficiently...
 - Base it on the notion of “conflicting” operations
- Definition: Two operations conflict if:
 - They are by different transactions,
 - They are on the same object and at least one of them is a write.

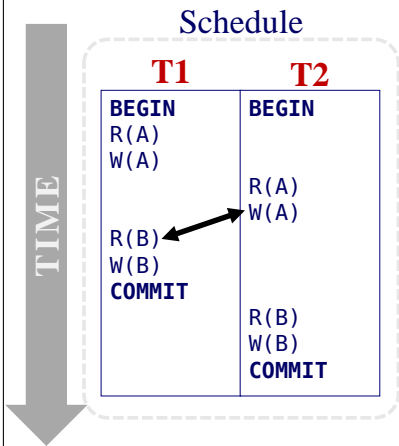
Conflict Serializable Schedules

- Two schedules are *conflict equivalent* iff:
 - They involve the same actions of the same transactions, and
 - Every pair of conflicting actions is ordered the same way.
- Schedule *S* is *conflict serializable* if:
 - *S* is conflict equivalent to some serial schedule.

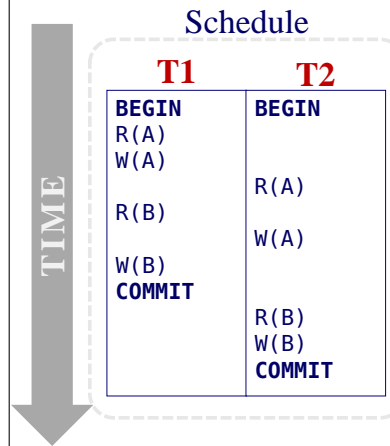
Conflict Serializability Intuition

- Schedule *S* is *conflict serializable* if:
 - You are able to transform *S* into a serial schedule by swapping consecutive non-conflicting operations of different transactions.

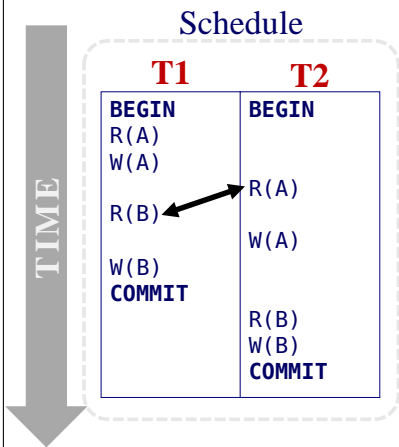
Conflict Serializability Intuition



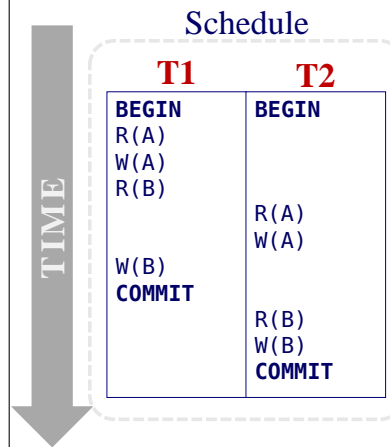
Conflict Serializability Intuition



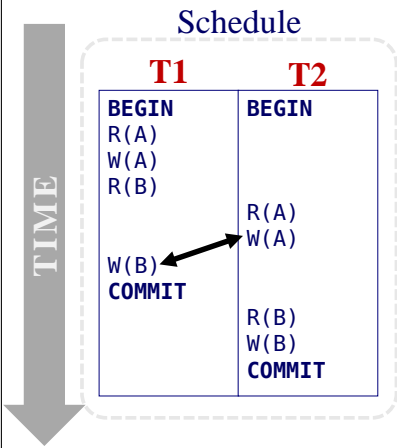
Conflict Serializability Intuition



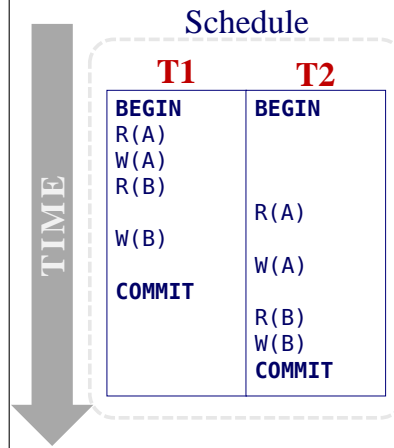
Conflict Serializability Intuition



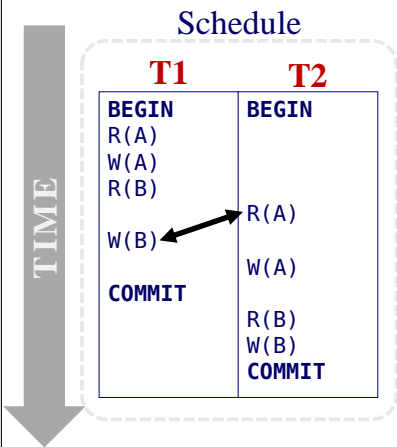
Conflict Serializability Intuition



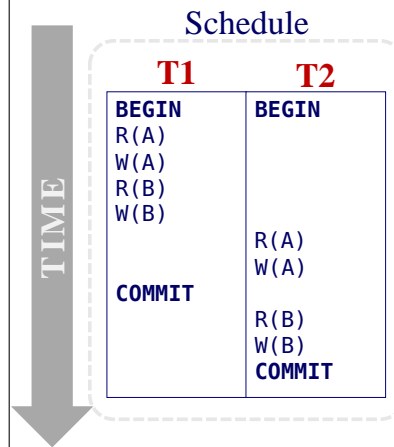
Conflict Serializability Intuition



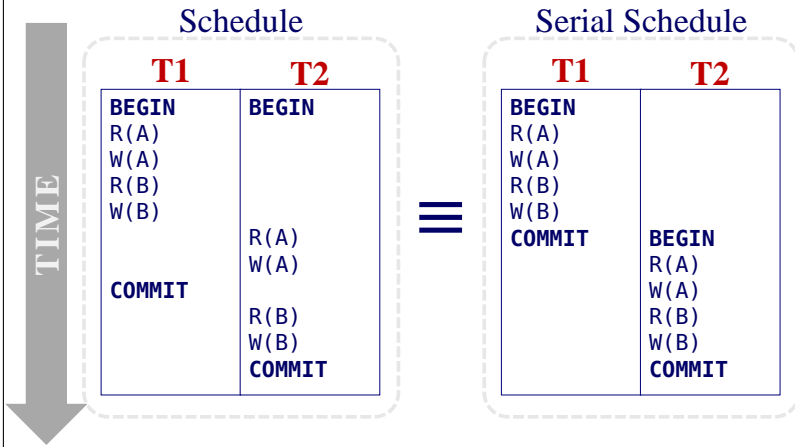
Conflict Serializability Intuition



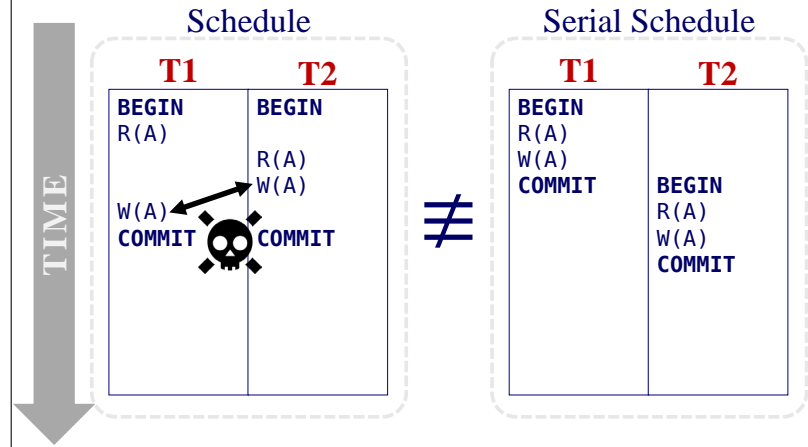
Conflict Serializability Intuition



Conflict Serializability Intuition



Conflict Serializability Intuition




Serializability

- **Q:** Are there any faster algorithms to figure this out other than transposing operations?

Dependency Graphs

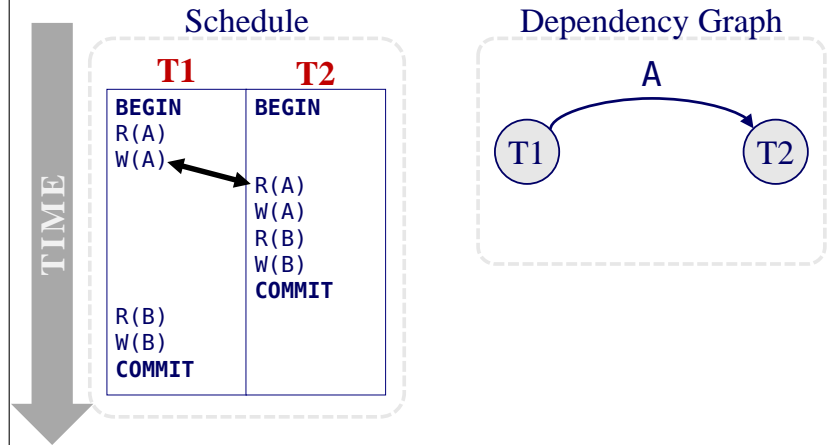
- One node per txn.
- Edge from T_i to T_j if:
 - An operation O_i of T_i conflicts with an operation O_j of T_j and
 - O_i appears earlier in the schedule than O_j .
- Also known as a “precedence graph”



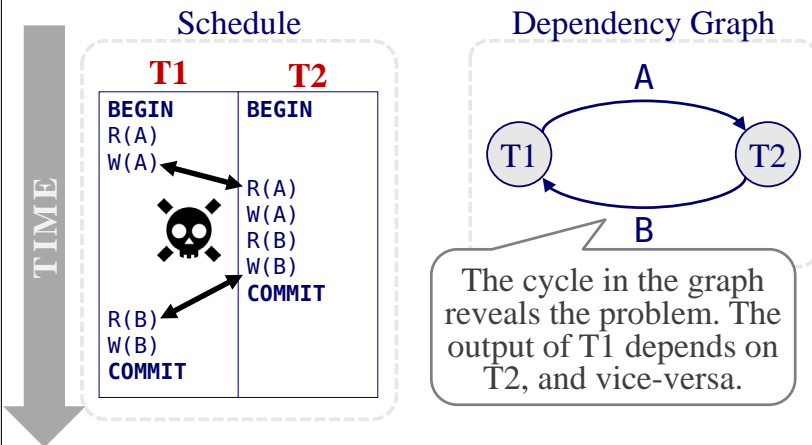
Dependency Graphs

- Theorem:** A schedule is *conflict serializable* if and only if its dependency graph is acyclic.

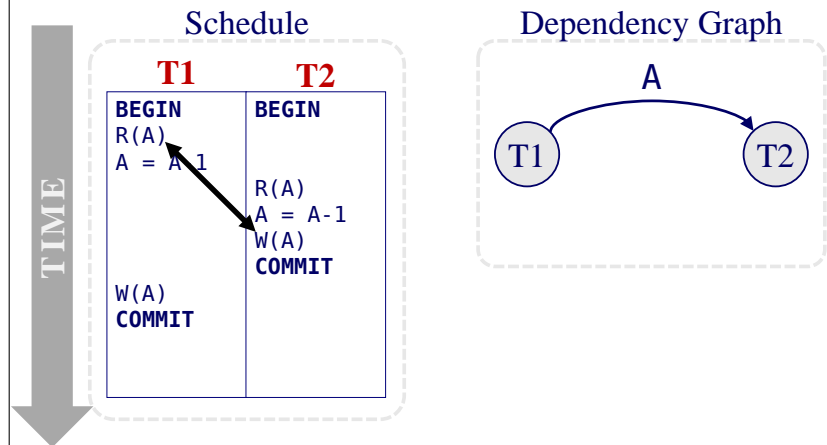
Example #1



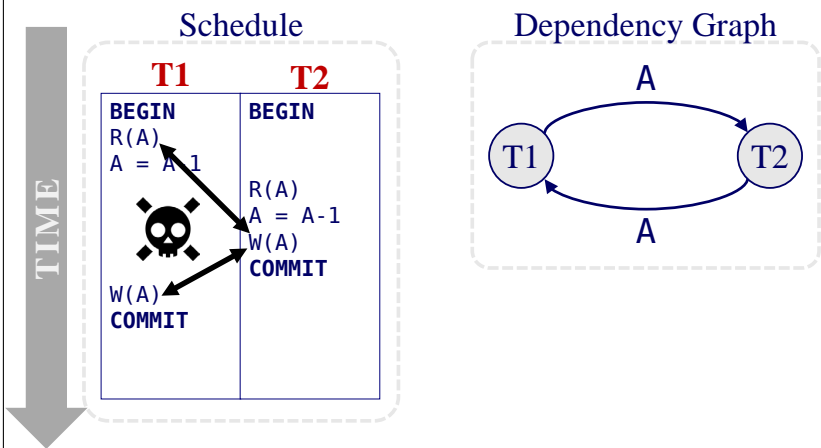
Example #1



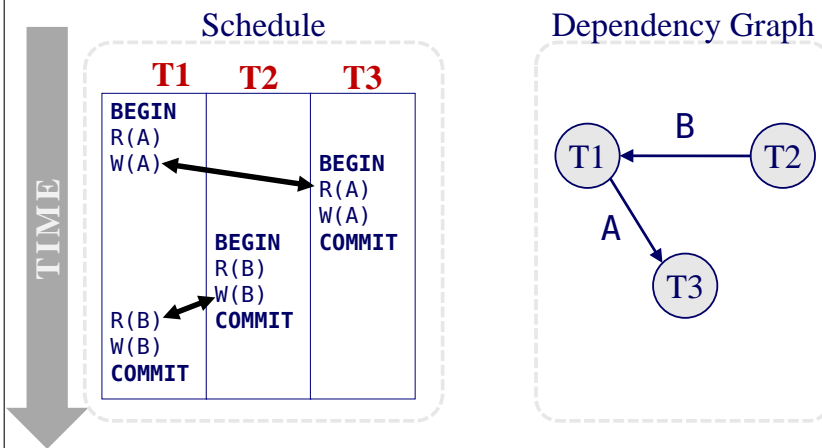
Example #2 – Lost Update



Example #2 – Lost Update



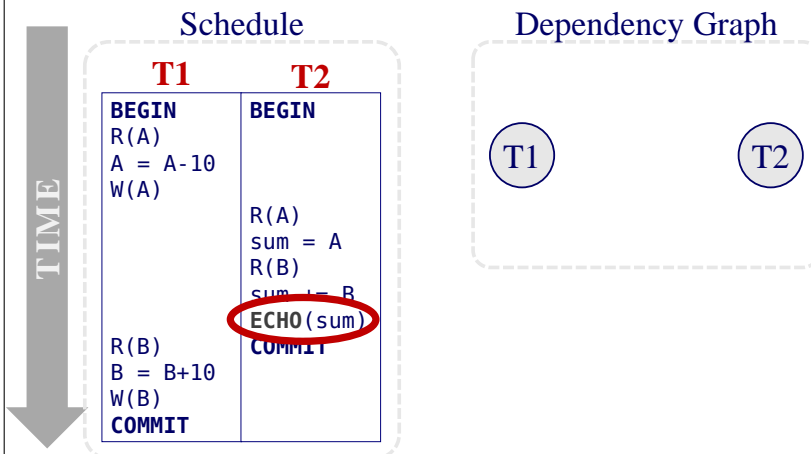
Example #3 – Threesome



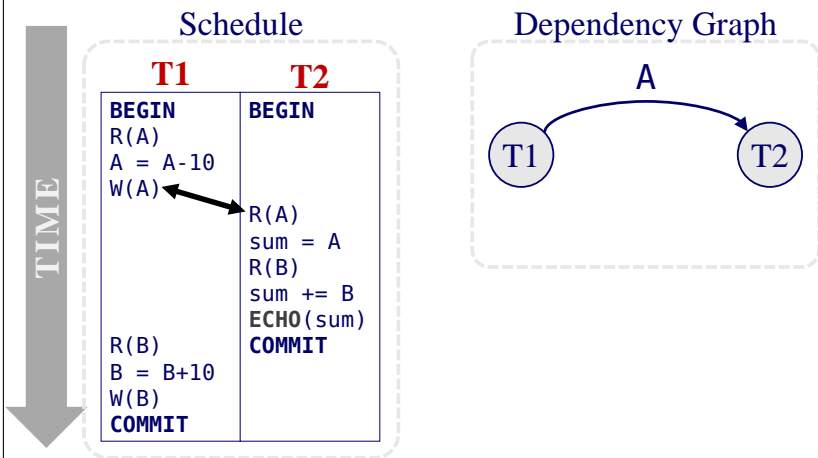
Example #3 – Threesome

- **Q:** Is this equivalent to a serial execution?
- **A:** Yes (T2, T1, T3)
 - Notice that T3 should go after T2, although it starts before it!

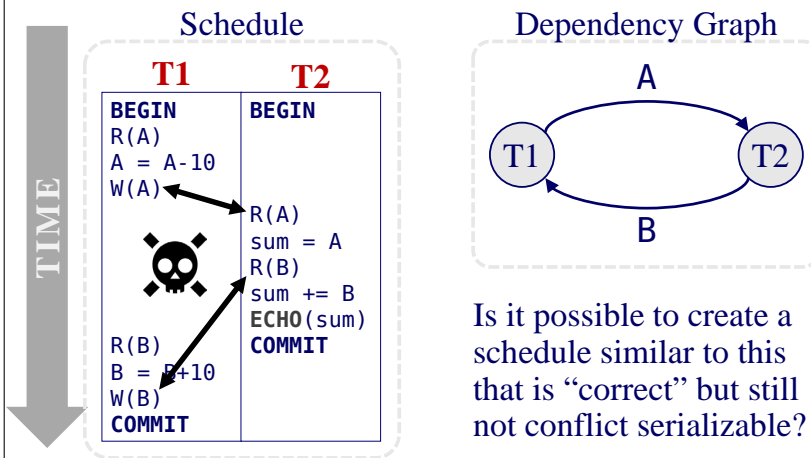
Example #4 – Inconsistent Analysis



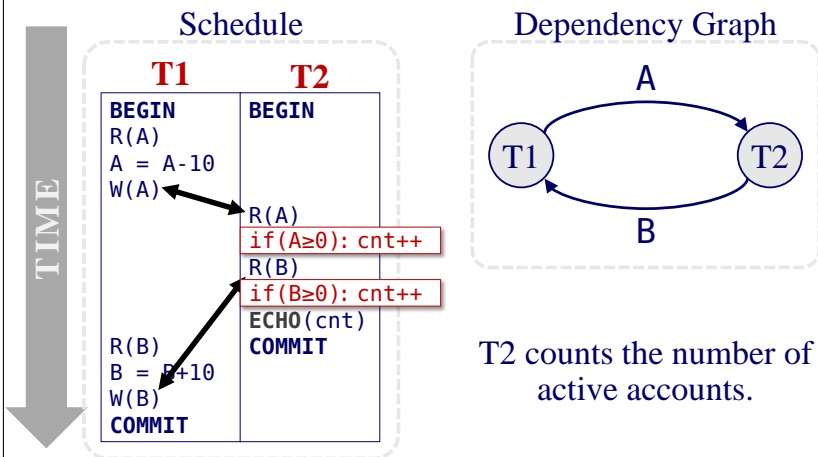
Example #4 – Inconsistent Analysis



Example #4 – Inconsistent Analysis



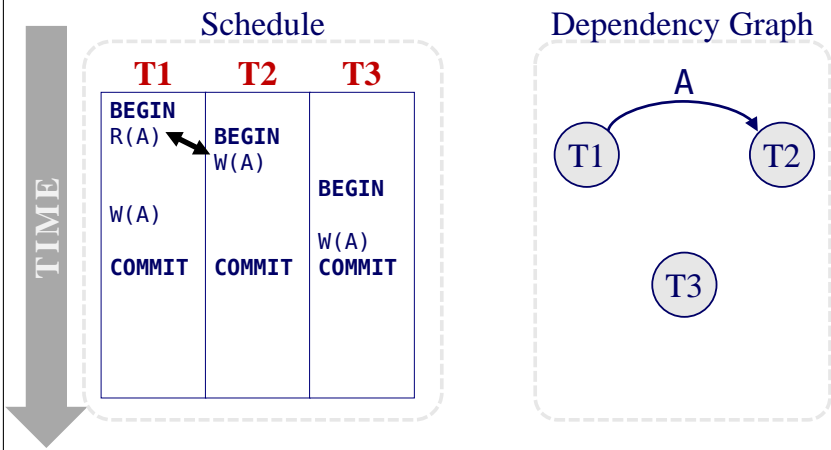
Example #4 – Inconsistent Analysis



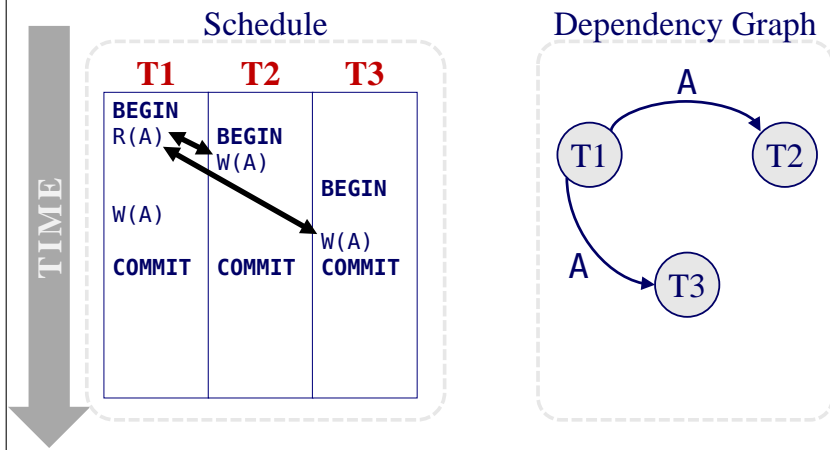
View Serializability

- Alternative (weaker) notion of serializability.
- Schedules S1 and S2 are *view equivalent* if:
 - If T1 reads initial value of A in S1, then T1 also reads initial value of A in S2.
 - If T1 reads value of A written by T2 in S1, then T1 also reads value of A written by T2 in S2.
 - If T1 writes final value of A in S1, then T1 also writes final value of A in S2.

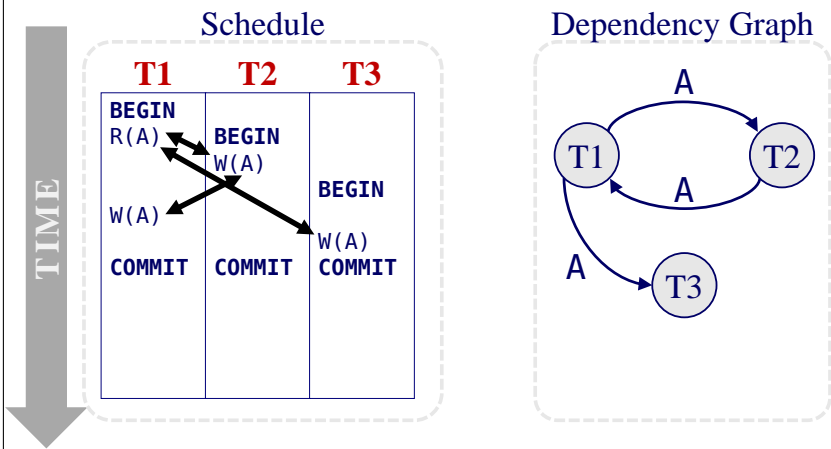
View Serializability



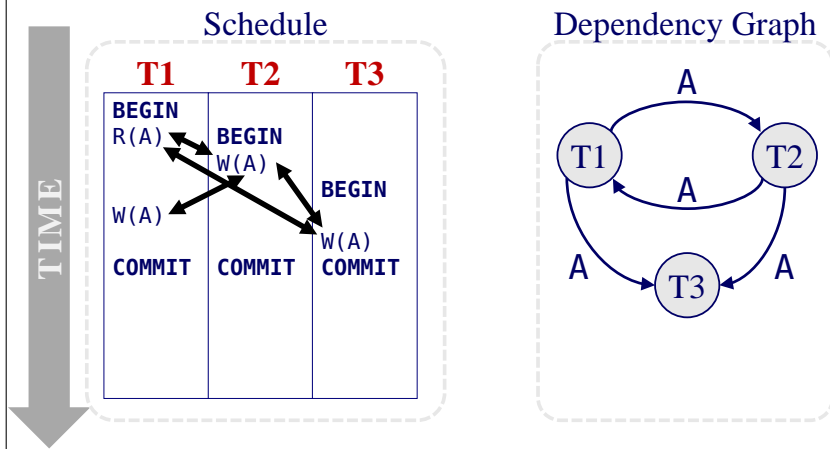
View Serializability



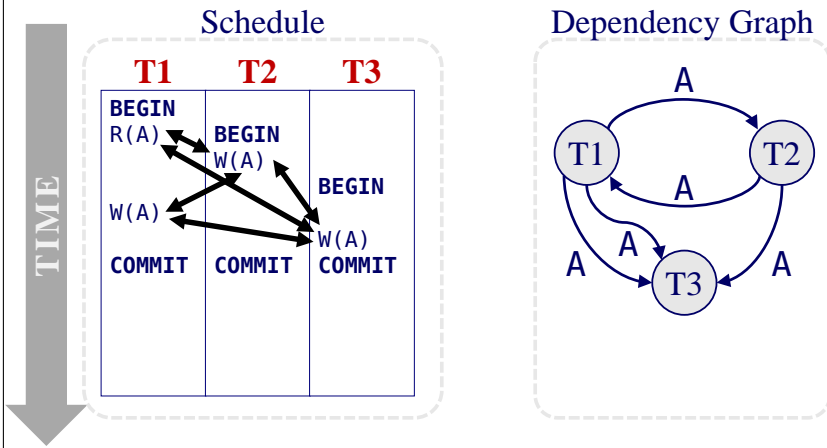
View Serializability



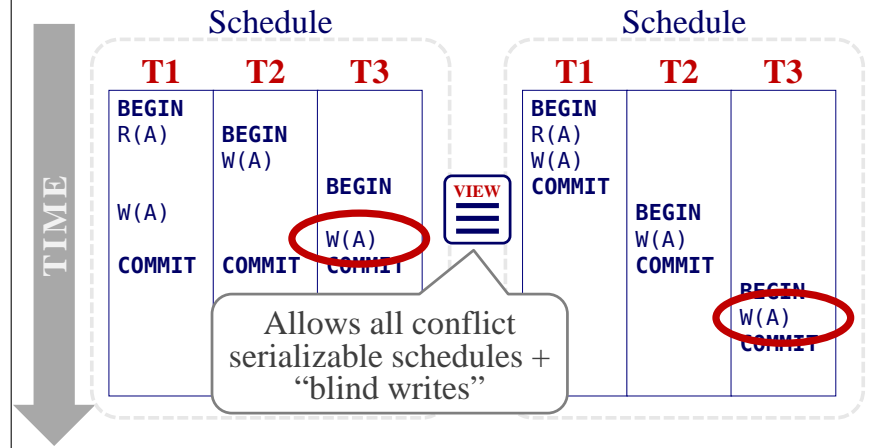
View Serializability



View Serializability



View Serializability



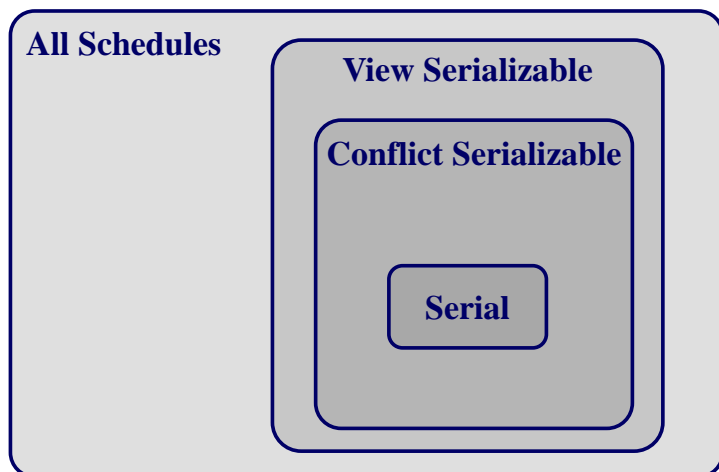
Serializability

- **View Serializability** allows (slightly) more schedules than **Conflict Serializability** does.
 - But is difficult to enforce efficiently.
- Neither definition allows all schedules that you would consider "serializable".
 - This is because they don't understand the meanings of the operations or the data (recall example #4)

Serializability

- In practice, **Conflict Serializability** is what gets used, because it can be enforced efficiently.
- To allow more concurrency, some special cases get handled separately at the application level.

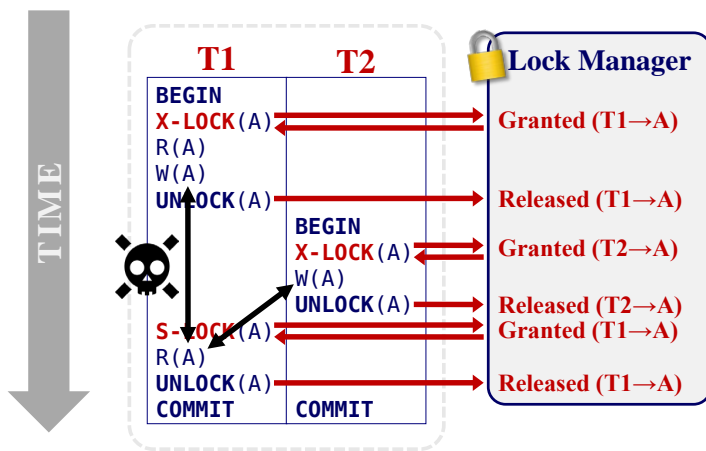
Schedules



Today's Class

- Serializability
- Two-Phase Locking
- Deadlocks
- Lock Granularities

Executing with Locks

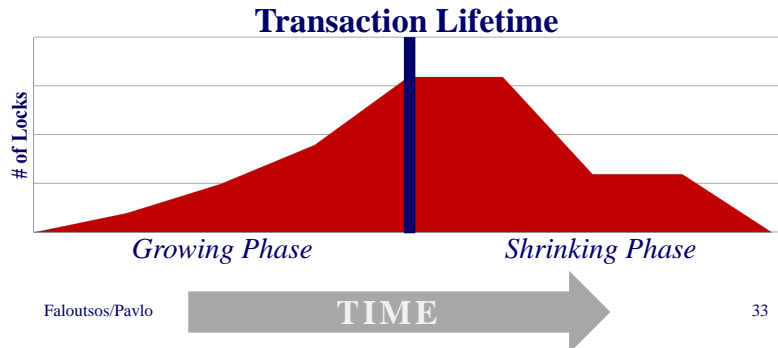


Two-Phase Locking

- **Phase 1: Growing**
 - Each txn requests the locks that it needs from the DBMS's lock manager.
 - The lock manager grants/denies lock requests.
- **Phase 2: Shrinking**
 - The txn is allowed to only release locks that it previously acquired. It cannot acquire new locks.


Two-Phase Locking

- The txn is not allowed to acquire/upgrade locks after the growing phase finishes.

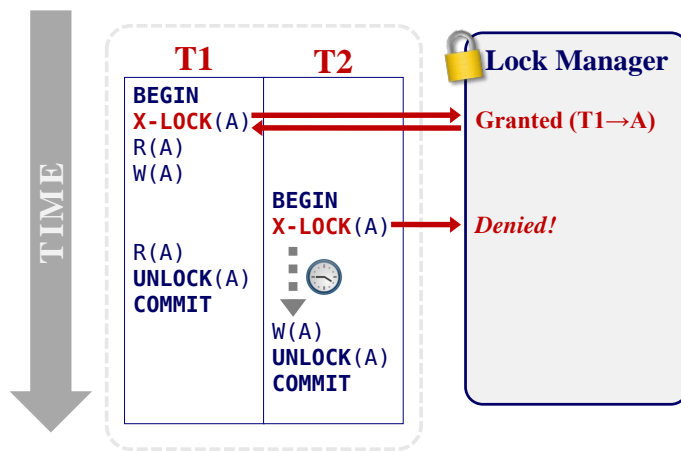


Two-Phase Locking

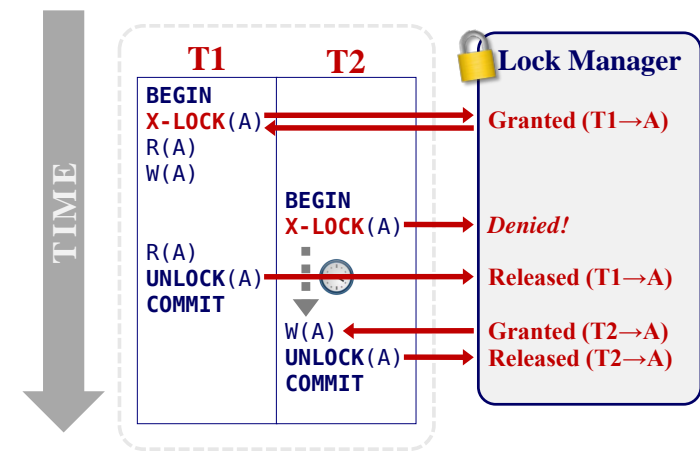
- The txn is not allowed to acquire/upgrade locks after the growing phase finishes.



Executing with 2PL



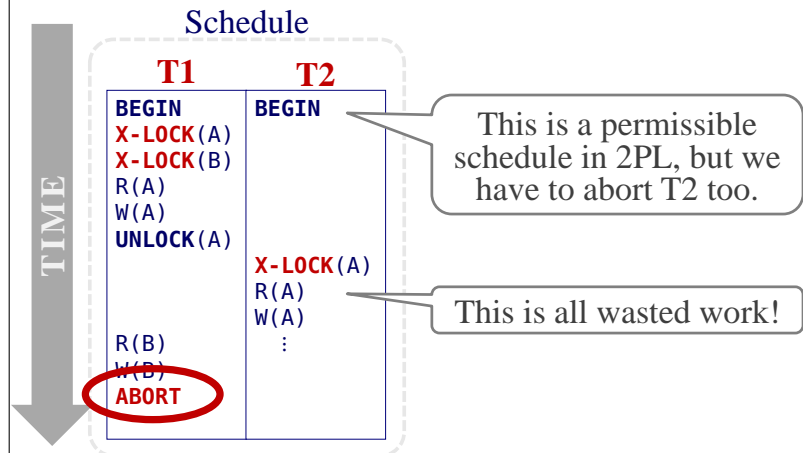
Executing with 2PL



Two-Phase Locking

- 2PL on its own is sufficient to guarantee conflict serializability (i.e., schedules whose precedence graph is acyclic), but, it is subject to *cascading aborts*.

2PL – Cascading Aborts




2PL Observations

- There are schedules that are serializable but would not be allowed by 2PL.
- Locking limits concurrency.
- May lead to deadlocks.
- May still have “dirty reads”
 - Solution: **Strict 2PL**

Strict Two-Phase Locking

- The txn is not allowed to acquire/upgrade locks after the growing phase finishes.
- Allows only conflict serializable schedules, but it is actually stronger than needed.



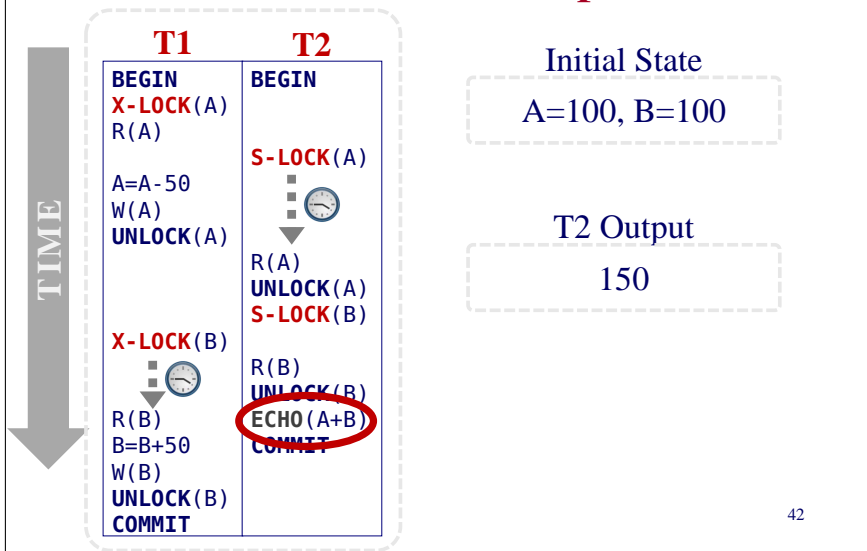
Strict Two-Phase Locking

- A schedule is *strict* if a value written by a txn is not read or overwritten by other txns until that txn finishes.
- Advantages:
 - Does not incur cascading aborts.
 - Aborted txns can be undone by just restoring original values of modified tuples.

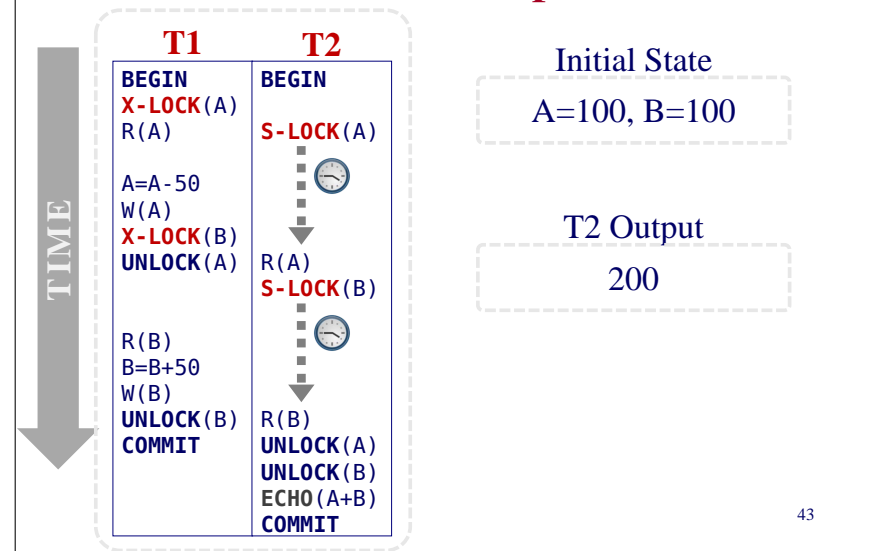
Examples

- **T1:** Move \$50 from Andy's account to his bookie's account.
- **T2:** Compute the total amount in all accounts and return it to the application.
- Legend:
 - **A** → Andy's account.
 - **B** → The bookie's account.

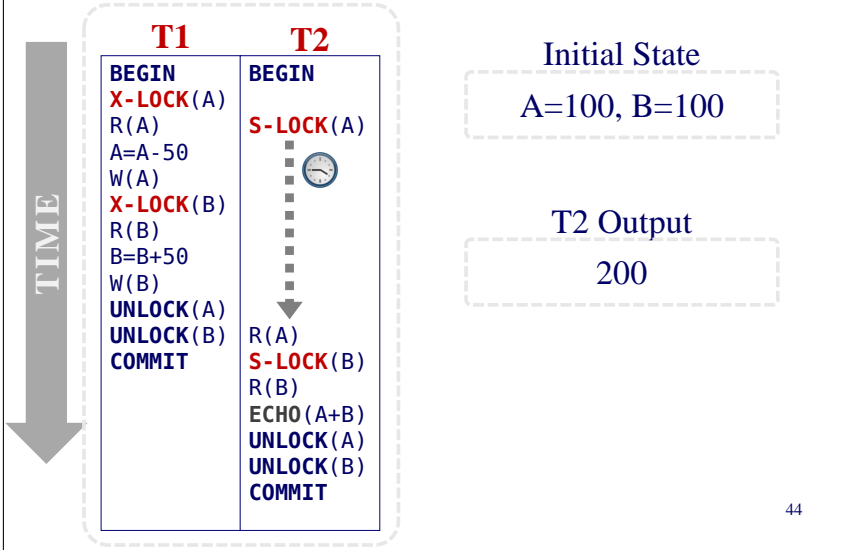
Non-2PL Example



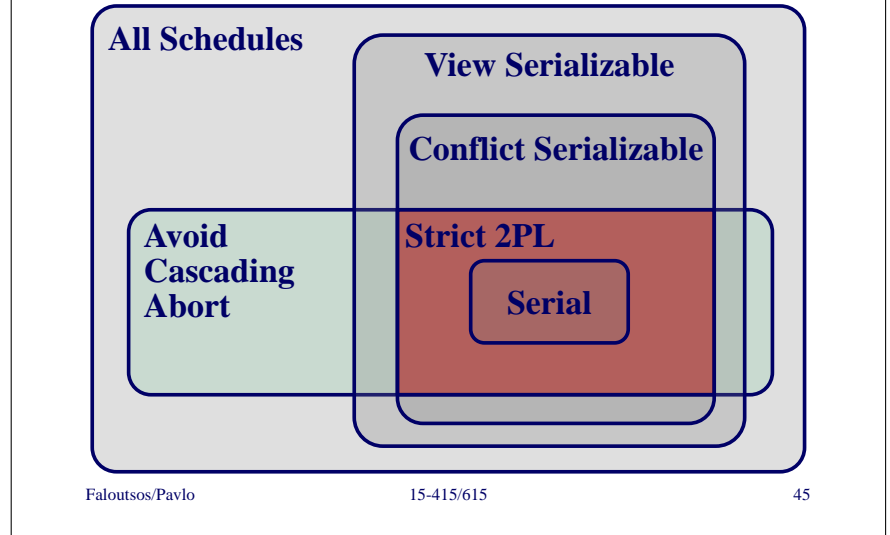
2PL Example



Strict 2PL Example



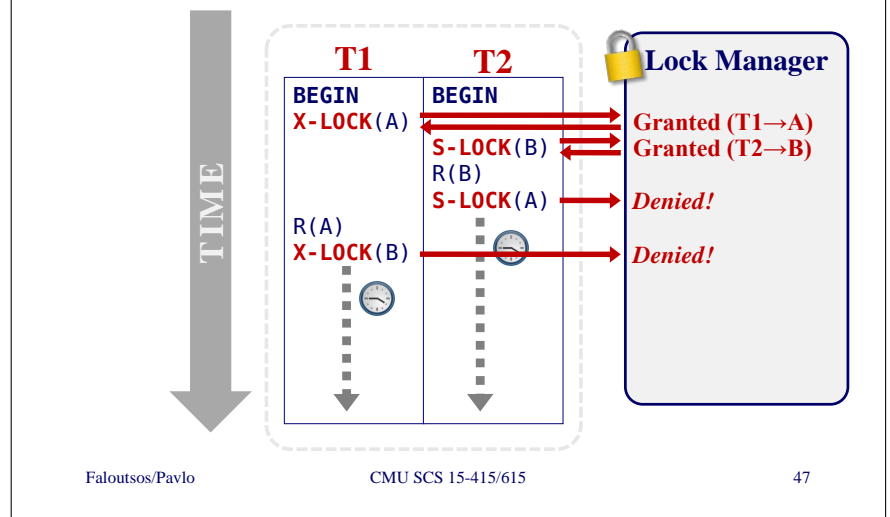
Schedules



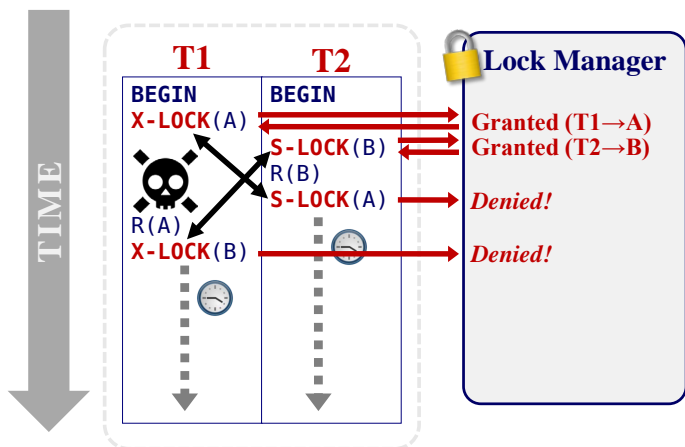
Today's Class

- Serializability
- Two-Phase Locking
- Deadlocks
- Lock Granularities

It Just Got Real, Son



It Just Got Real, Son



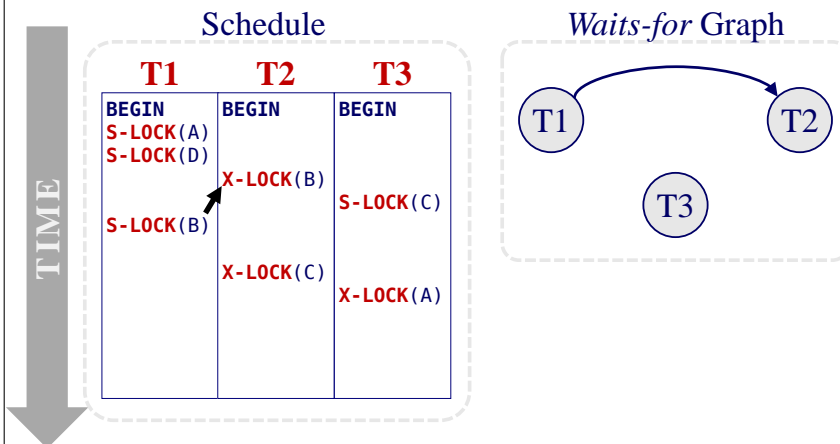
Deadlocks

- **Deadlock:** Cycle of transactions waiting for locks to be released by each other.
- Two ways of dealing with deadlocks:
 - Deadlock **detection**
 - Deadlock **prevention**


Deadlock Detection

- The DBMS creates a *waits-for* graph:
 - Nodes are transactions
 - Edge from T_i to T_j if T_i is waiting for T_j to release a lock
- The system periodically check for cycles in *waits-for* graph.

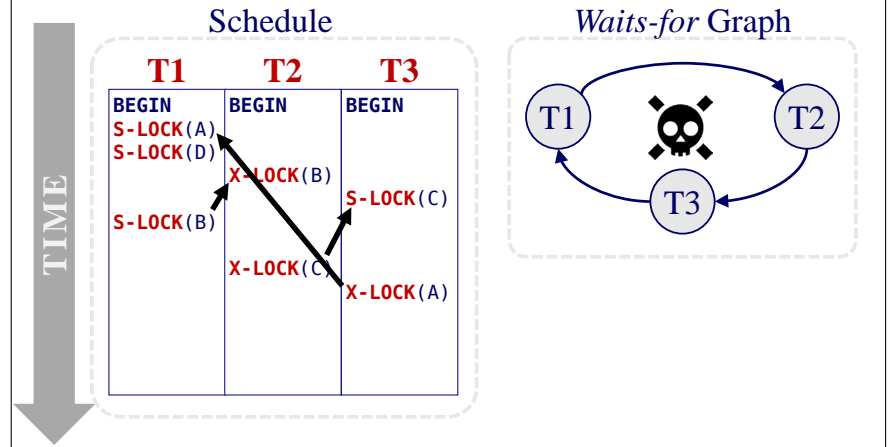
Deadlock Detection



Deadlock Detection



Deadlock Detection




Deadlock Detection

- How often should we run the algorithm?
- How many txns are typically involved?
- What do we do when we find a deadlock?


Deadlock Handling

- **Q:** What do we do?
- **A:** Select a “victim” and rollback it back to break the deadlock.




Deadlock Handling

- **Q:** Which one do we choose?
- **A:** It depends...
 - By age (lowest timestamp)
 - By progress (least/most queries executed)
 - By the # of items already locked
 - By the # of txns that we have to rollback with it
- We also should consider the # of times a txn has been restarted in the past.



Deadlock Handling

- **Q:** How far do we rollback?
- **A:** It depends...
 - Completely
 - Minimally (i.e., just enough to release locks)

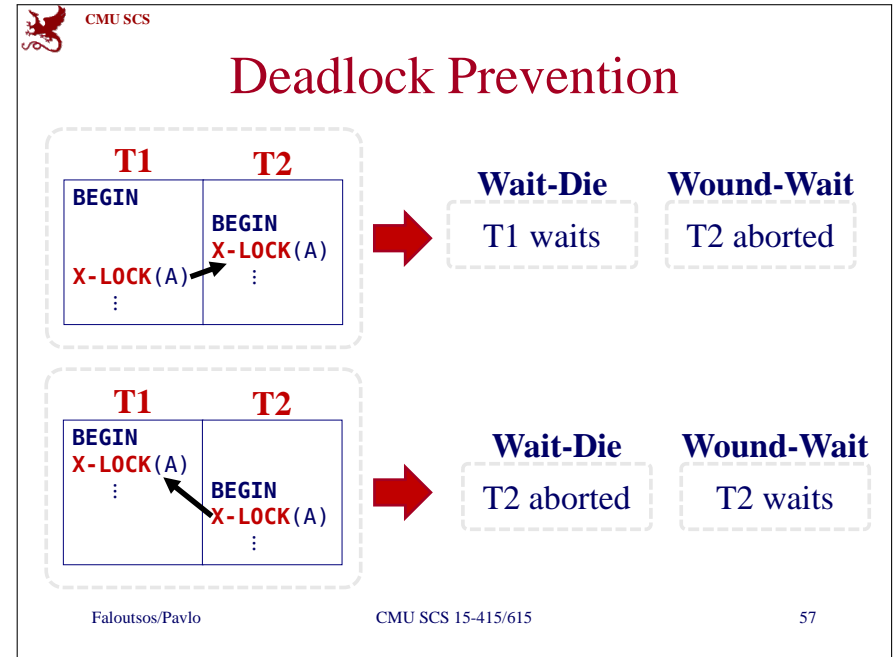
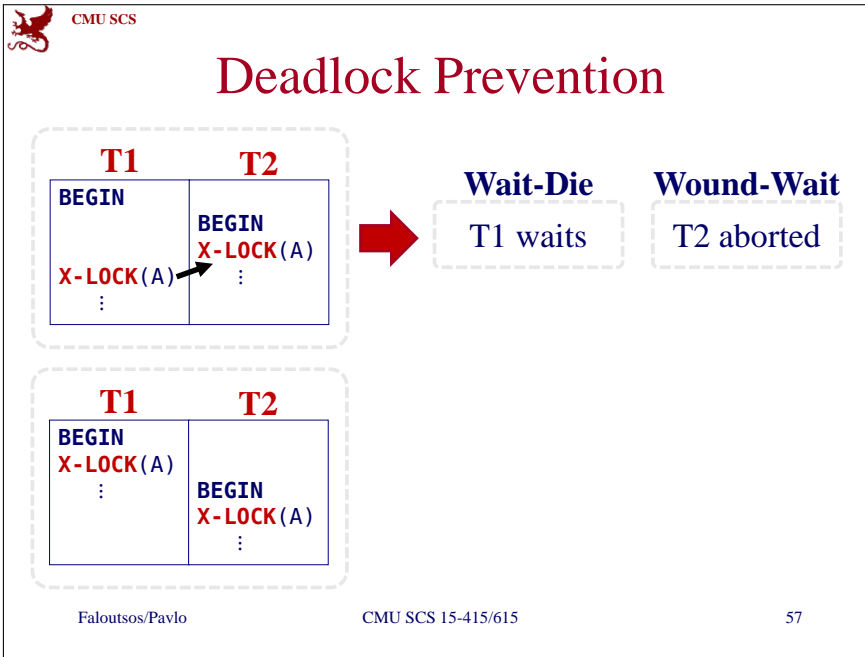


Deadlock Prevention

- When a txn tries to acquire a lock that is held by another txn, kill one of them to prevent a deadlock.
- No *waits-for* graph or detection algorithm.

Deadlock Prevention

- Assign priorities based on timestamps:
 - Older → higher priority (e.g., $T1 > T2$)
- Two different prevention policies:
 - **Wait-Die:** If T1 has higher priority, T1 waits for T2; otherwise T1 aborts (“old wait for young”)
 - **Wound-Wait:** If T1 has higher priority, T2 aborts; otherwise T1 waits (“young wait for old”)




- CMU SCS
- ## Deadlock Prevention
- **Q:** Why do these schemes guarantee no deadlocks?
 - **A:** Only one “type” of direction allowed.
 - **Q:** When a transaction restarts, what is its (new) priority?
 - **A:** Its original timestamp. Why?
- Faloutsos/Pavlo CMU SCS 15-415/615 58

- CMU SCS
- ## Today’s Class
- Serializability
 - Two-Phase Locking
 - Deadlocks
 - Lock Granularities
- Faloutsos/Pavlo CMU SCS 15-415/615 59

Lock Granularities

- When we say that a txn acquires a “lock”, what does that actually mean?
 - On an Attribute? Tuple? Page? Table?
- Ideally, each txn should obtain fewest number of locks that is needed...

Database Lock Hierarchy



Example



- **T1:** Get the balance of Andy’s shady off-shore bank account.
- **T2:** Increase Christos’ bank account balance by 1%.
- **Q:** What locks should they obtain?

Example



- **Q:** What locks should they obtain?
- **A:** Multiple
 - **Exclusive + Shared** for leafs of lock tree.
 - Special **Intention** locks for higher levels

Intention Locks



- Intention locks allow a higher level node to be locked in **S** or **X** mode without having to check all descendent nodes.
- If a node is in an intention mode, then explicit locking is being done at a lower level in the tree.

Intention Locks



- **Intention-Shared (IS)**: Indicates explicit locking at a lower level with shared locks.
- **Intention-Exclusive (IX)**: Indicates locking at lower level with exclusive or shared locks.

Intention Locks



- **Shared+Intention-Exclusive (SIX)**: The subtree rooted by that node is locked explicitly in shared mode and explicit locking is being done at a lower level with exclusive-mode locks.

Compatibility Matrix




		T2 Wants				
		IS	IX	S	SIX	X
T1 Holds	IS	✓	✓	✓	✓	✗
	IX	✓	✓	✗	✗	✗
	S	✓	✗	✓	✗	✗
	SIX	✓	✗	✗	✗	✗
	X	✗	✗	✗	✗	✗

Example – Two-level Hierarchy

Read Andy's record in R.


T1



Example – Two-level Hierarchy

Read Andy's record in R.

T1




Example – Two-level Hierarchy

Update Christos' record in R.

T1

T2




Example – Two-level Hierarchy

Update Christos' record in R.

T1

T2



CMU SCS

Example – Two-level Hierarchy

in R.

		T2 Wants				
		IS	IX	S	SIX	X
T1 Holds	IS	✓	✓	✓	✓	✗
	IX	✓	✓	✗	✗	✗
	S	✓	✗	✓	✗	✗
	SIX	✓	✗	✗	✗	✗
	X	✗	✗	✗	✗	✗

Faloutsos/Pavlo CMU SCS 15-415/615 68

CMU SCS

Example – Two-level Hierarchy

Update Christos' record in R.

Faloutsos/Pavlo CMU SCS 15-415/615 68

CMU SCS

Example – Threesome

- Assume three txns execute at same time:
 - T1: Scan R and update a few tuples.
 - T2: Read a single tuple in R.
 - T3: Scan all tuples in R.

Faloutsos/Pavlo CMU SCS 15-415/615 69

CMU SCS

Example – Threesome


Scan R and update a few tuples.

T1

Faloutsos/Pavlo CMU SCS 15-415/615 70

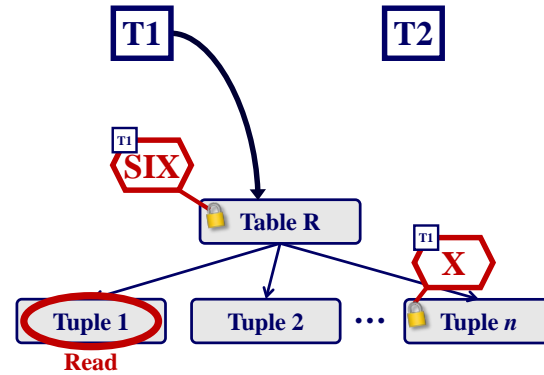
Example – Threesome

Scan **R** and update a few tuples.




Example – Threesome

Read a single tuple in **R**.




Example – Threesome

Read a single tuple in **R**.

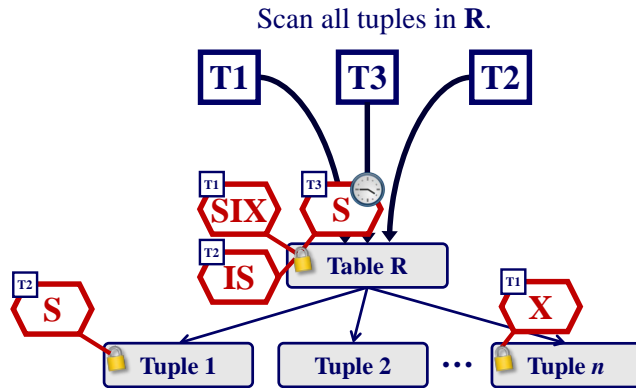


Example – Threesome

Scan all tuples in **R**.



Example – Threesome



Multiple Lock Granularities

- Useful in practice as each txn only needs a few locks.
- Intention locks help improve concurrency:
 - **Intention-Shared (IS)**: Intent to get **S** lock(s) at finer granularity.
 - **Intention-Exclusive (IX)**: Intent to get **X** lock(s) at finer granularity.
 - **Shared+Intention-Exclusive (SIX)**: Like **S** and **IX** at the same time.

Locking in Practice

- You typically don't set locks manually.
- Sometimes you will need to provide the DBMS with hints to help it to improve concurrency.
- Also useful for doing major changes.

LOCK TABLE

Postgres

```
LOCK TABLE <table> IN <mode> MODE;
```

MySQL

```
LOCK TABLE <table> <mode>;
```

- Explicitly locks a table.
- Not part of the SQL standard.
 - Postgres Modes: **SHARE, EXCLUSIVE**
 - MySQL Modes: **READ, WRITE**

SELECT...FOR UPDATE

```
SELECT * FROM <table>  
WHERE <qualification> FOR UPDATE;
```

- Perform a select and then sets an exclusive lock on the matching tuples.
- Can also set shared locks:
 - Postgres: **FOR SHARE**
 - MySQL: **LOCK IN SHARE MODE**

Concurrency Control Summary

- Conflict Serializability ↔ Correctness
- Automatically correct interleavings:
 - Locks + protocol (2PL, S2PL ...)
 - Deadlock detection + handling
 - Deadlock prevention