


Lecture #08

Granularities of Locks and Degrees of Consistency

Fall 2024

» Prof. Jignesh Patel



ANNOUNCEMENTS

- Talk from Oracle on Tuesday, October 1, @ noon in 6501 GHC.
 - Unifying relational and document/JSON management.
- Exam: Oct 9th in **GHC 8102 between 1-4 pm**. Open book.
 - Start anytime. Stop 90 minutes later.
 - Let me know if you have a conflict by the end of next week 9/20.

TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```

Bank Balance : \$100



TRANSACTION MANAGEMENT

Read (A);
Check (A > \$25);
Pay (\$25);
A = A - 25;
Write (A);

You

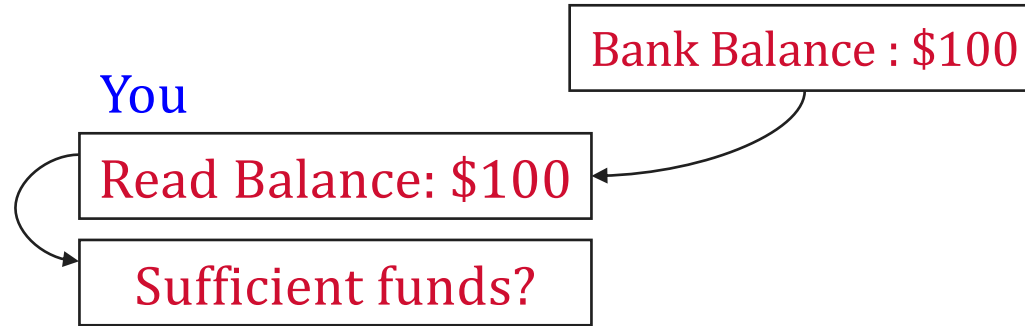
Bank Balance : \$100

Read Balance: \$100



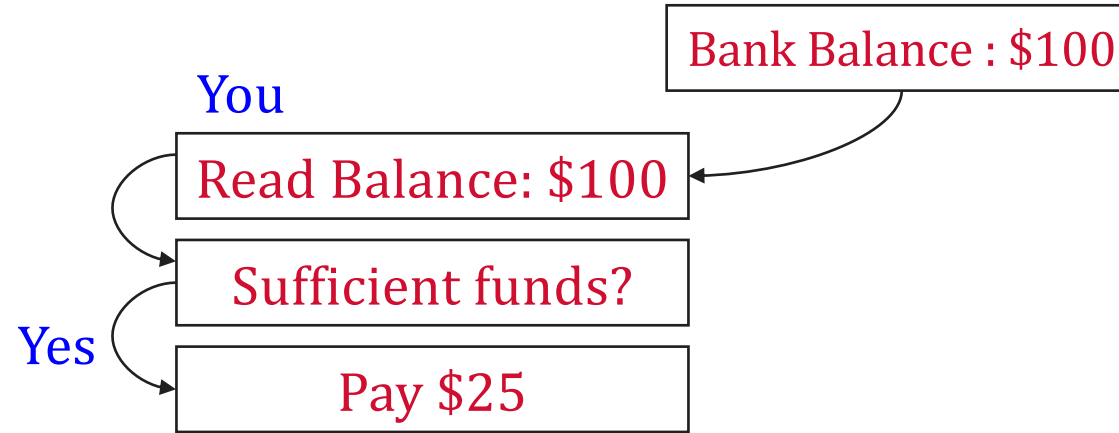
TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



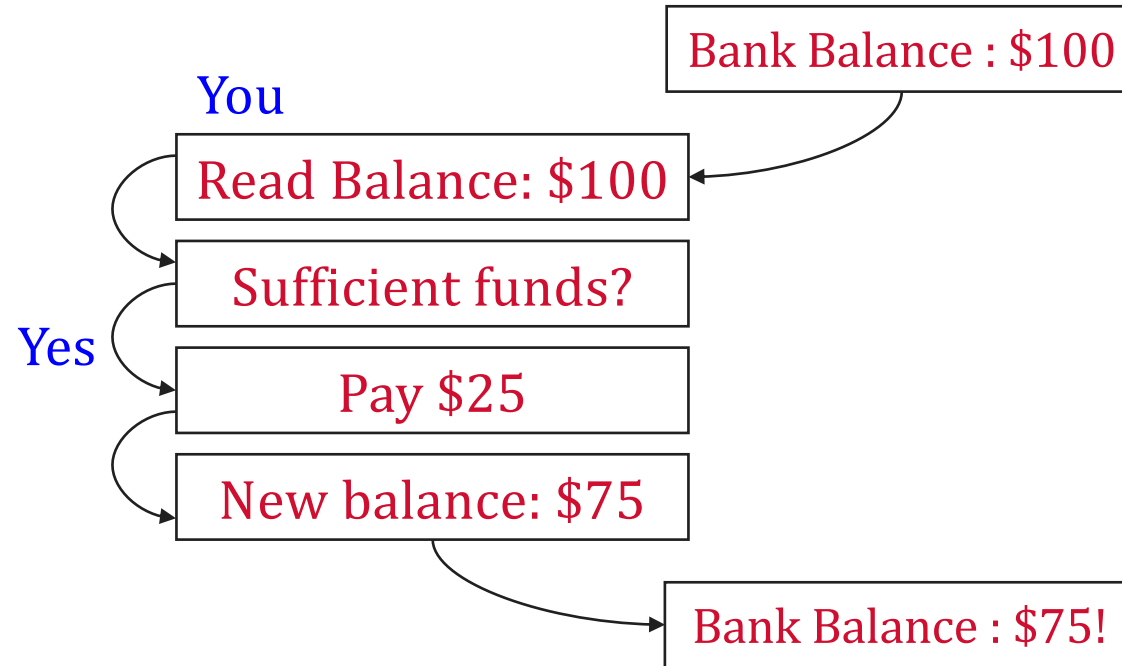
TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```




TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



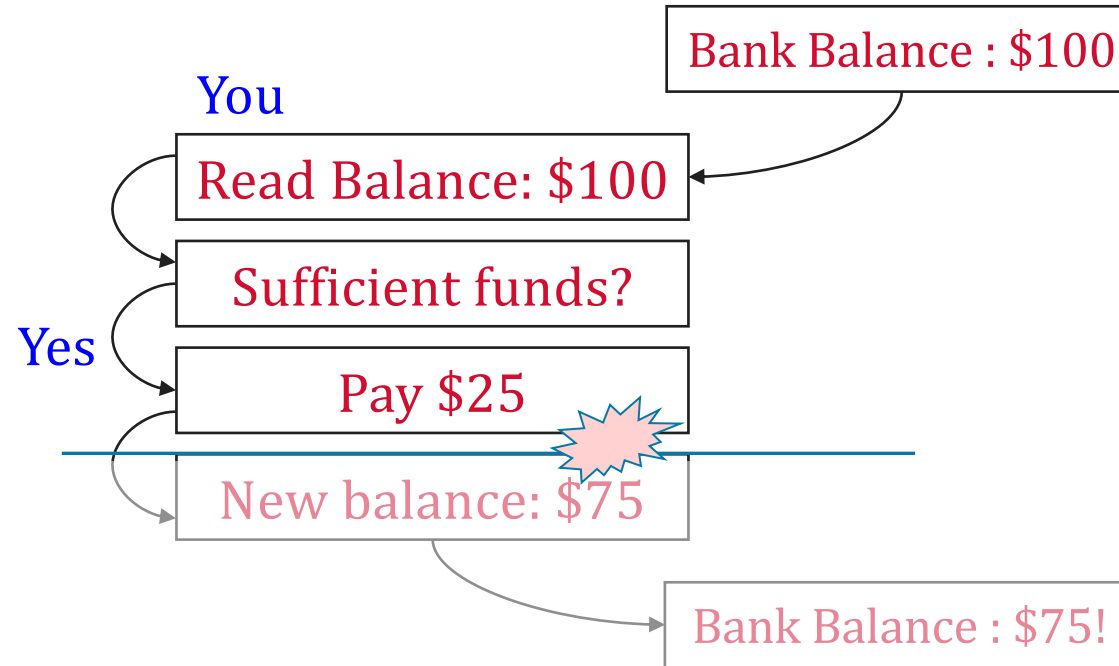

TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```

You


Bank Balance : \$100

Your Significant Other




TRANSACTION MANAGEMENT

Read (A);
Check (A > \$25);
Pay (\$25);
A = A - 25;
Write (A);



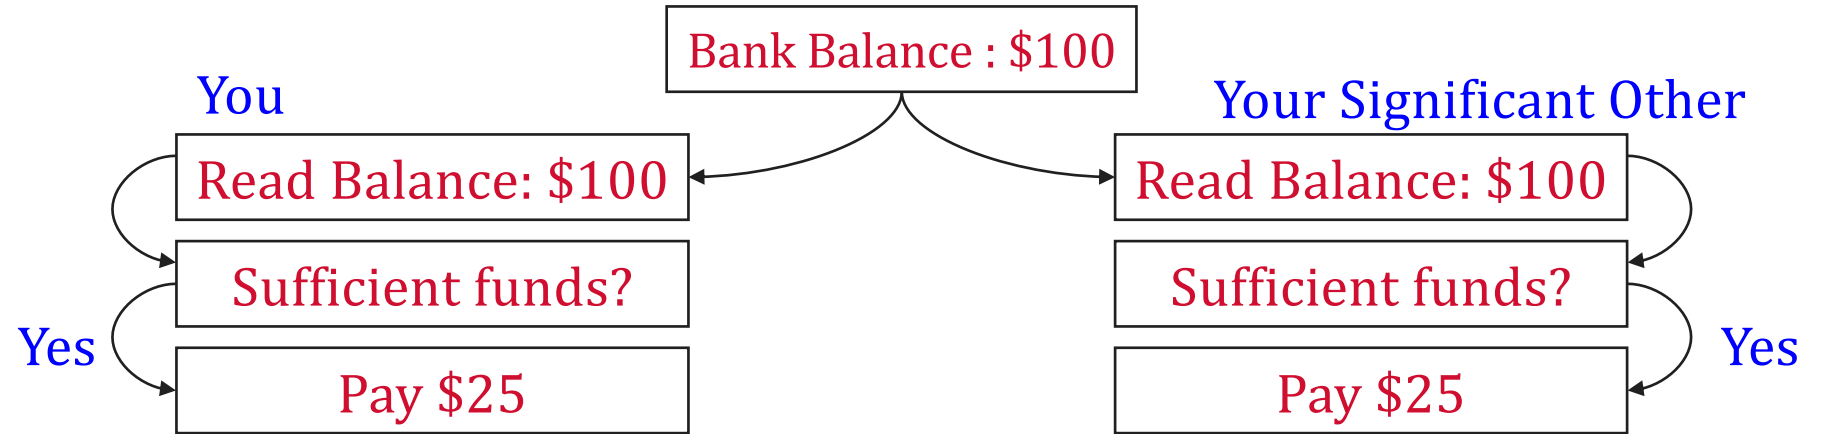
TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



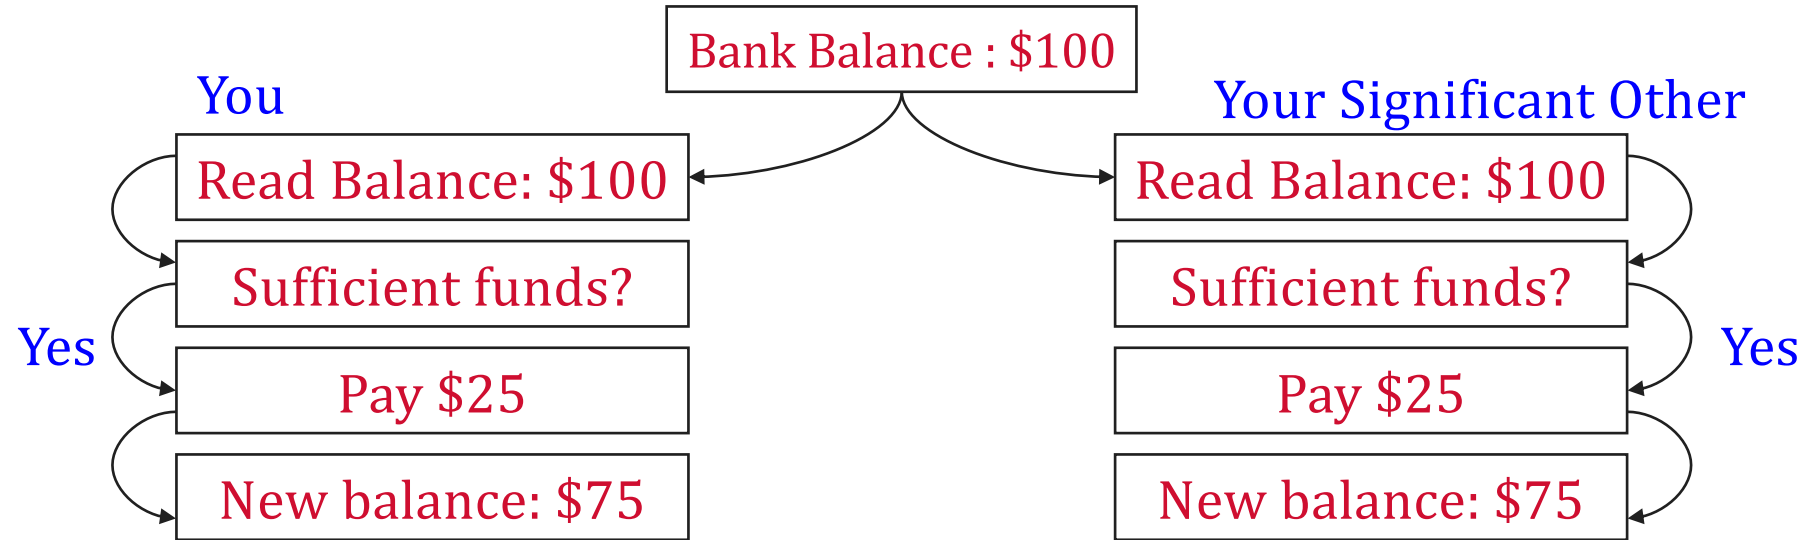
TRANSACTION MANAGEMENT

Read (A);
Check ($A > \$25$);
Pay (\$25);
 $A = A - 25$;
Write (A);



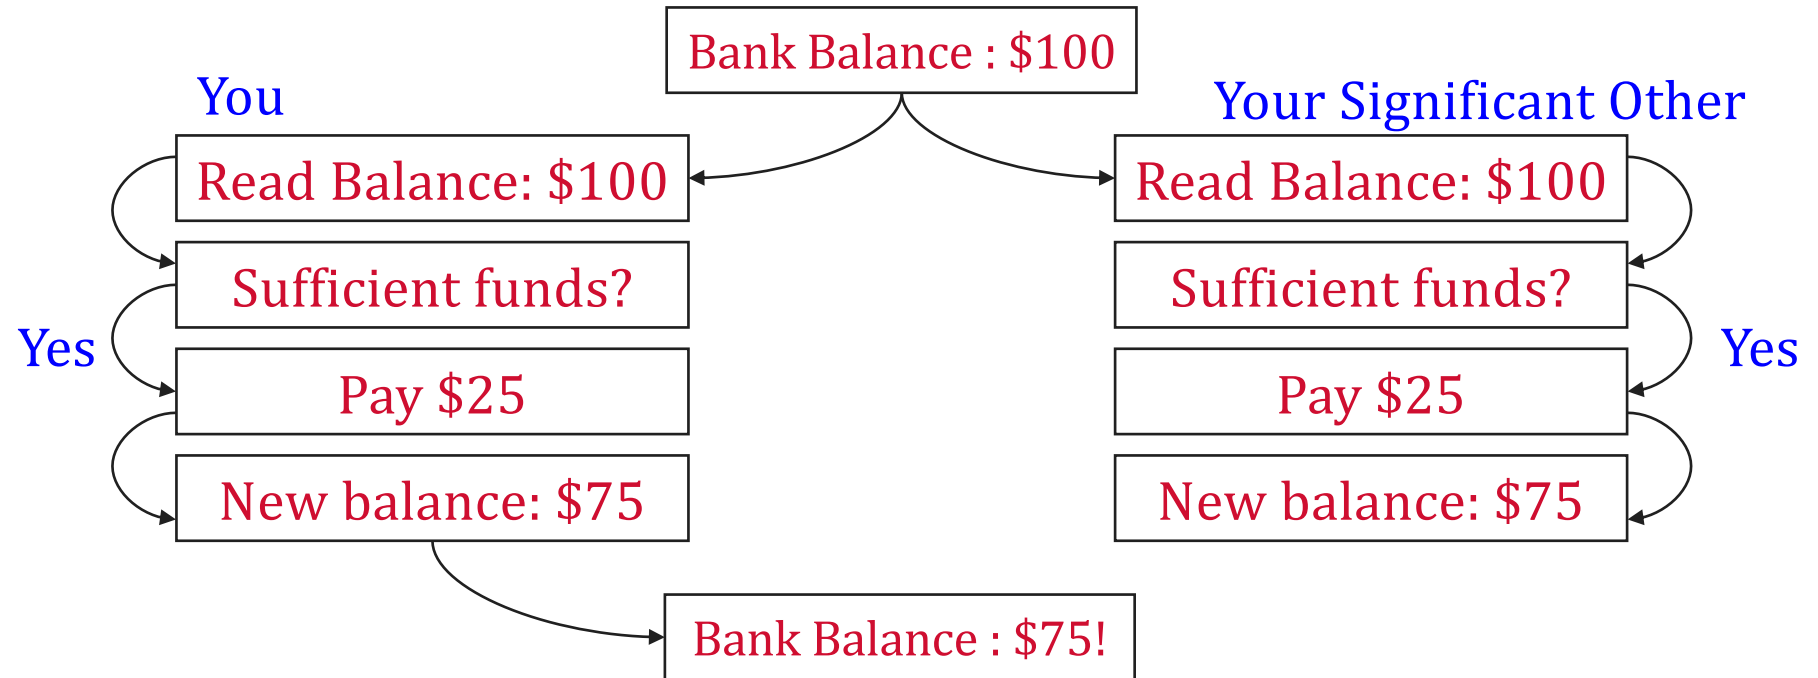
TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



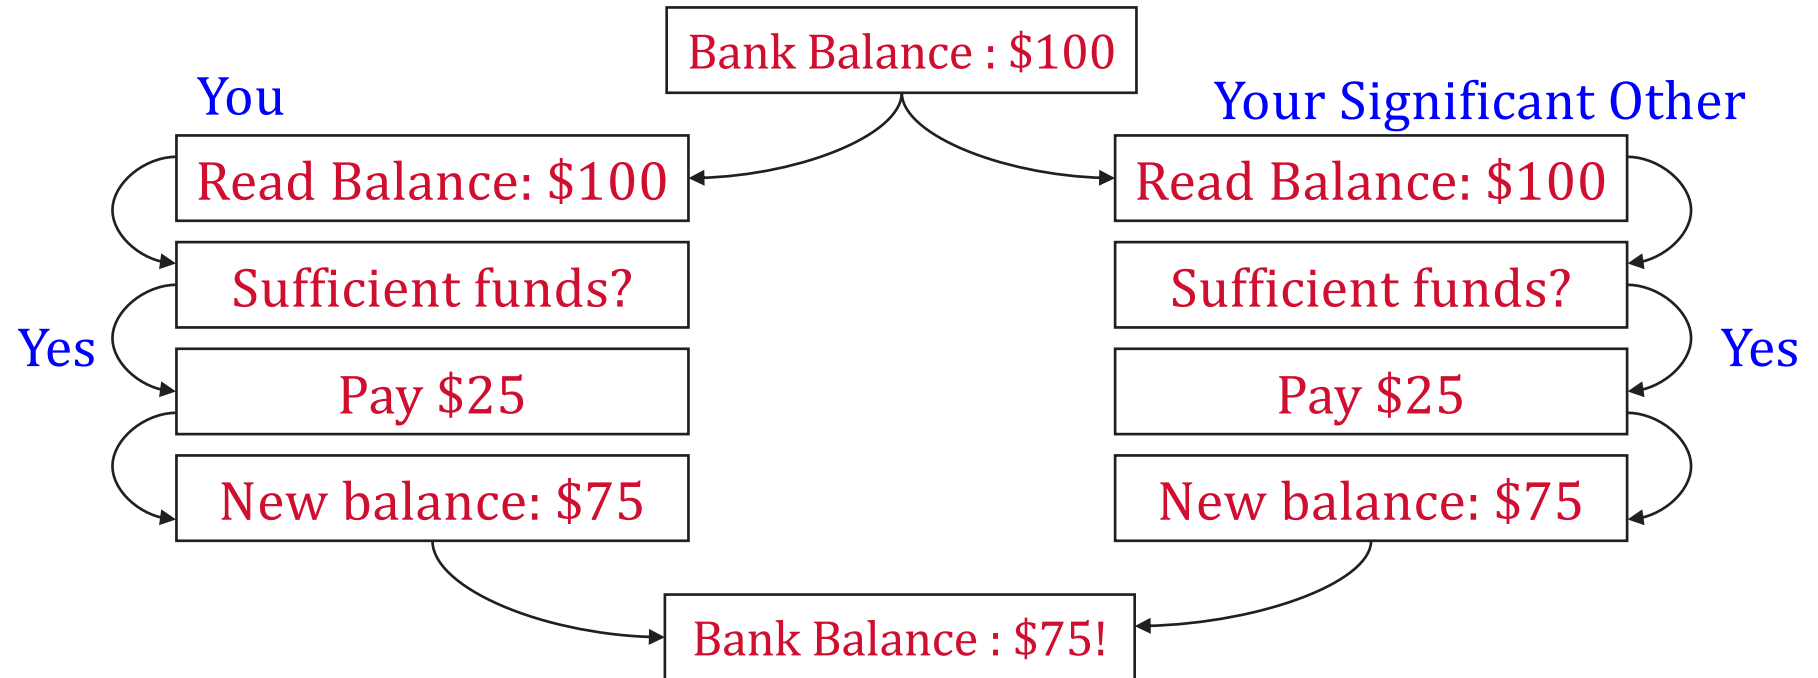
TRANSACTION MANAGEMENT

```
Read (A);  
Check (A > $25);  
Pay ($25);  
A = A - 25;  
Write (A);
```



TRANSACTION MANAGEMENT

Read (A);
Check ($A > \$25$);
Pay (\$25);
 $A = A - 25$;
Write (A);



TRANSACTION MANAGEMENT

Redo/Undo
mechanism

Atomicity

All actions in txn happen, or none happen.
“All or nothing...”

Integrity
Constraints

Consistency

Key constraints, CHECKS, TRIGGERS, ...
hold before and after the txn completes.

If each txn is consistent and the DB starts
consistent, then it ends up consistent.
“It looks correct to me...”

Concurrency
Control

Isolation


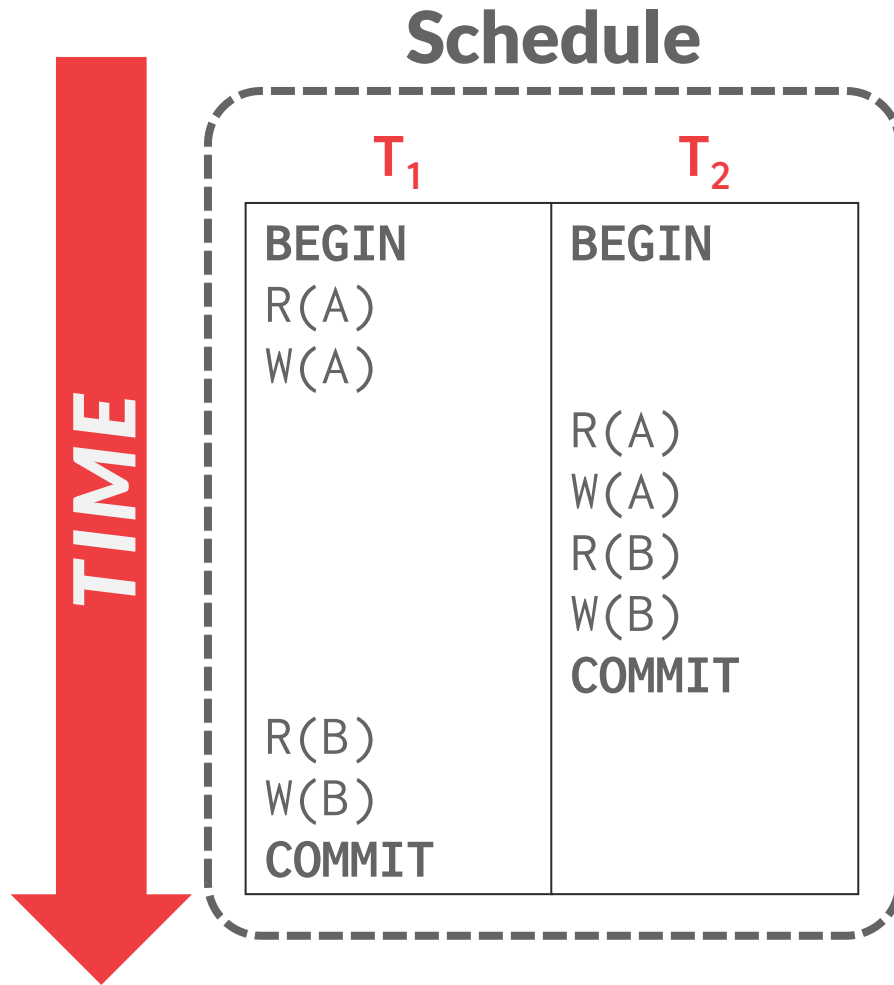
Execution of one txn is isolated from that
of other txns.
“All by myself...”

Redo/Undo
mechanism


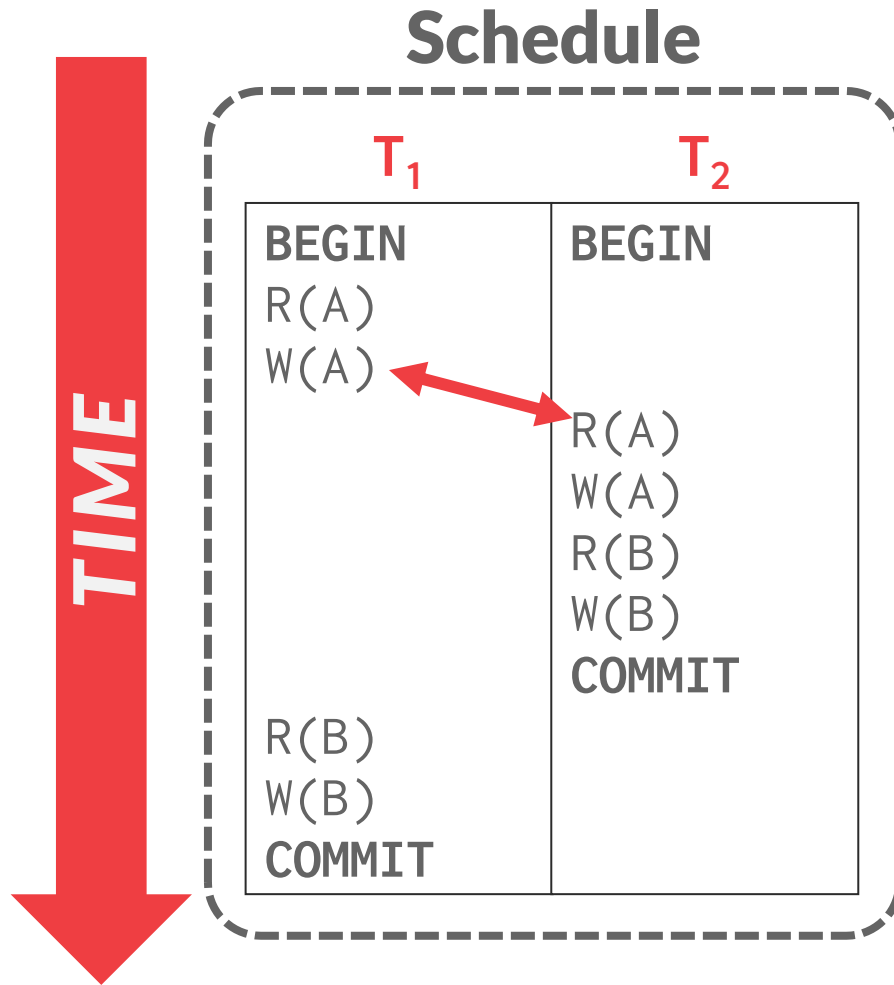
Durability

If a txn commits, its effects persist.
“I will survive...”



CONCURRENCY CONTROL AND DEPENDENCE GRAPHS




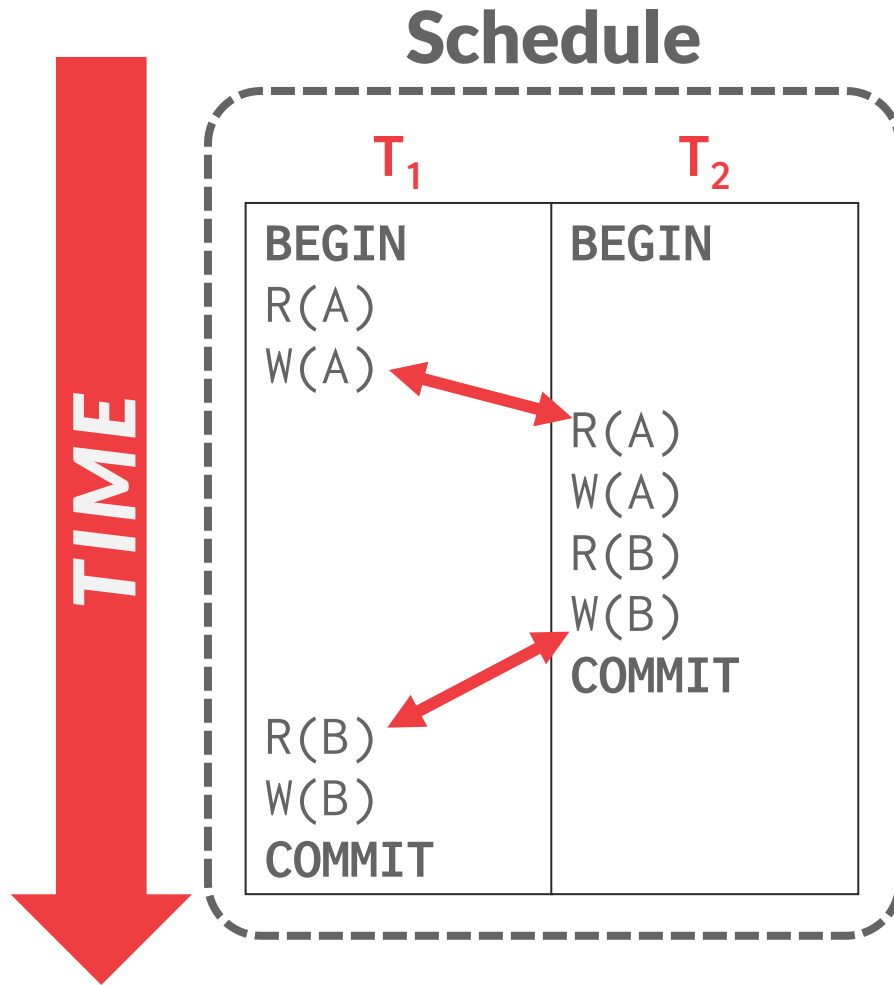
CONCURRENCY CONTROL AND DEPENDENCE GRAPHS




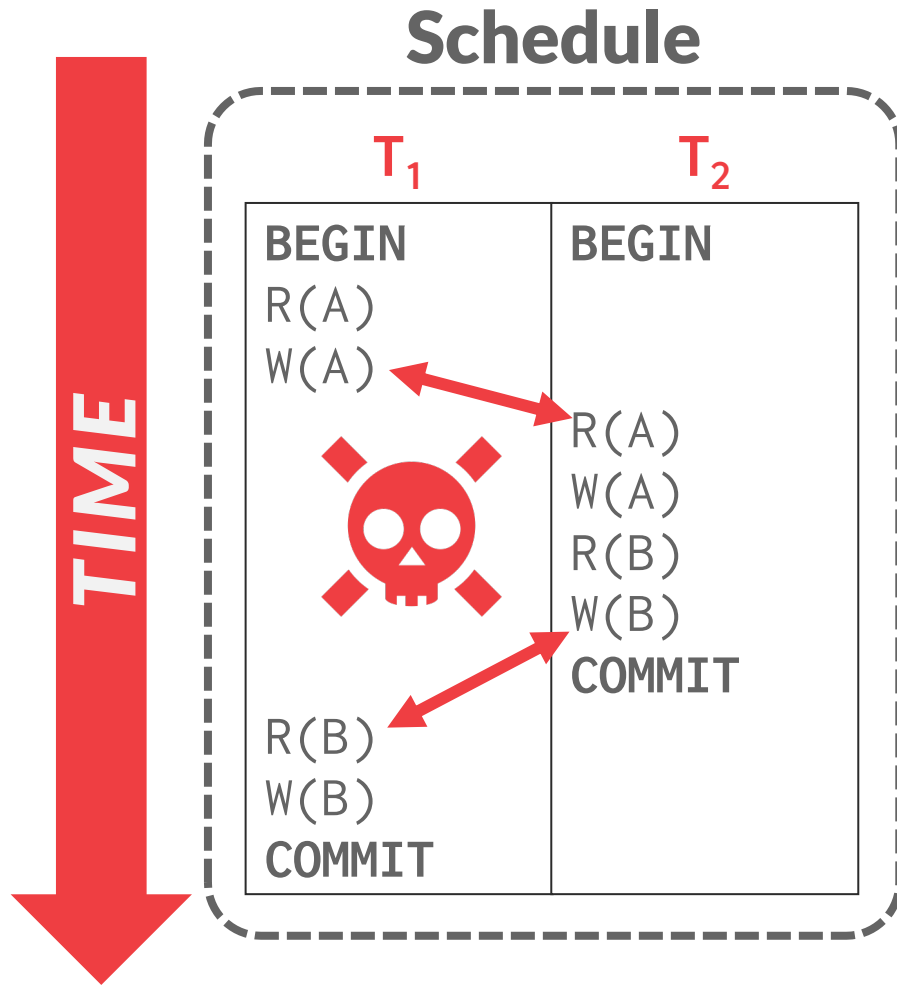
CONCURRENCY CONTROL AND DEPENDENCE GRAPHS




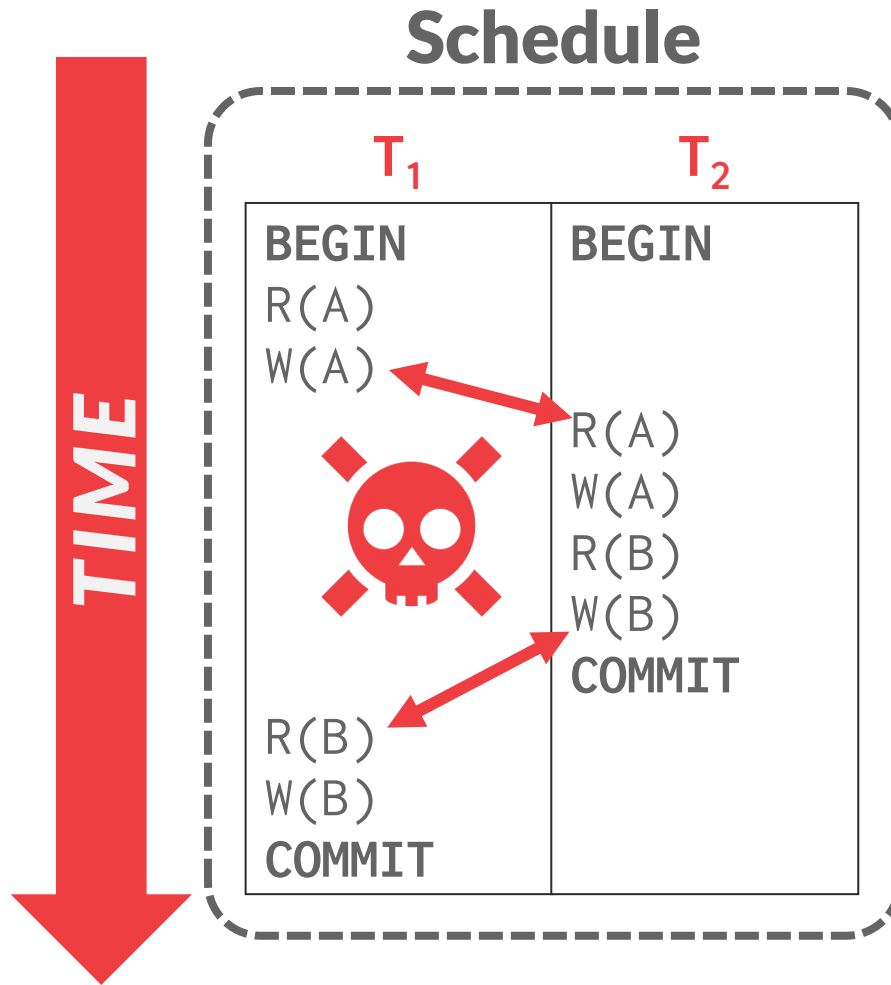
CONCURRENCY CONTROL AND DEPENDENCE GRAPHS



CONCURRENCY CONTROL AND DEPENDENCE GRAPHS



CONCURRENCY CONTROL AND DEPENDENCE GRAPHS

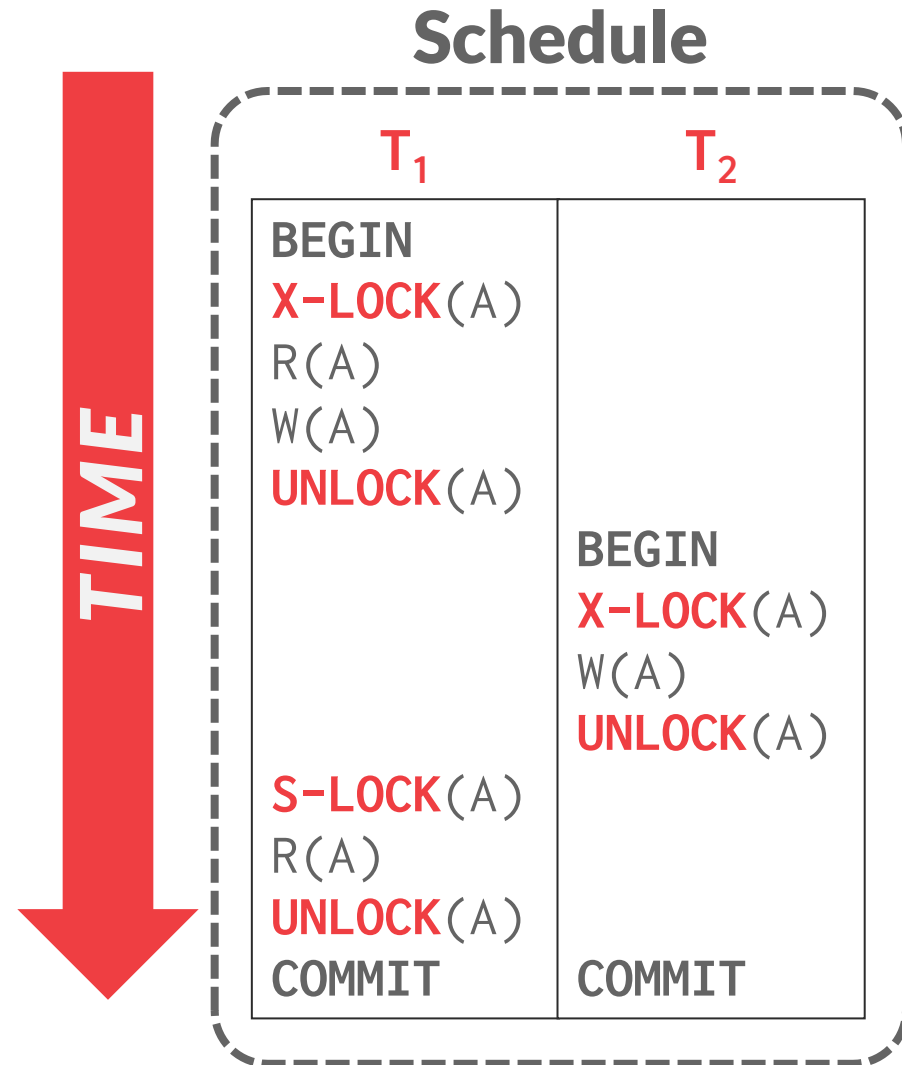


SERIALIZABLE SCHEDULE

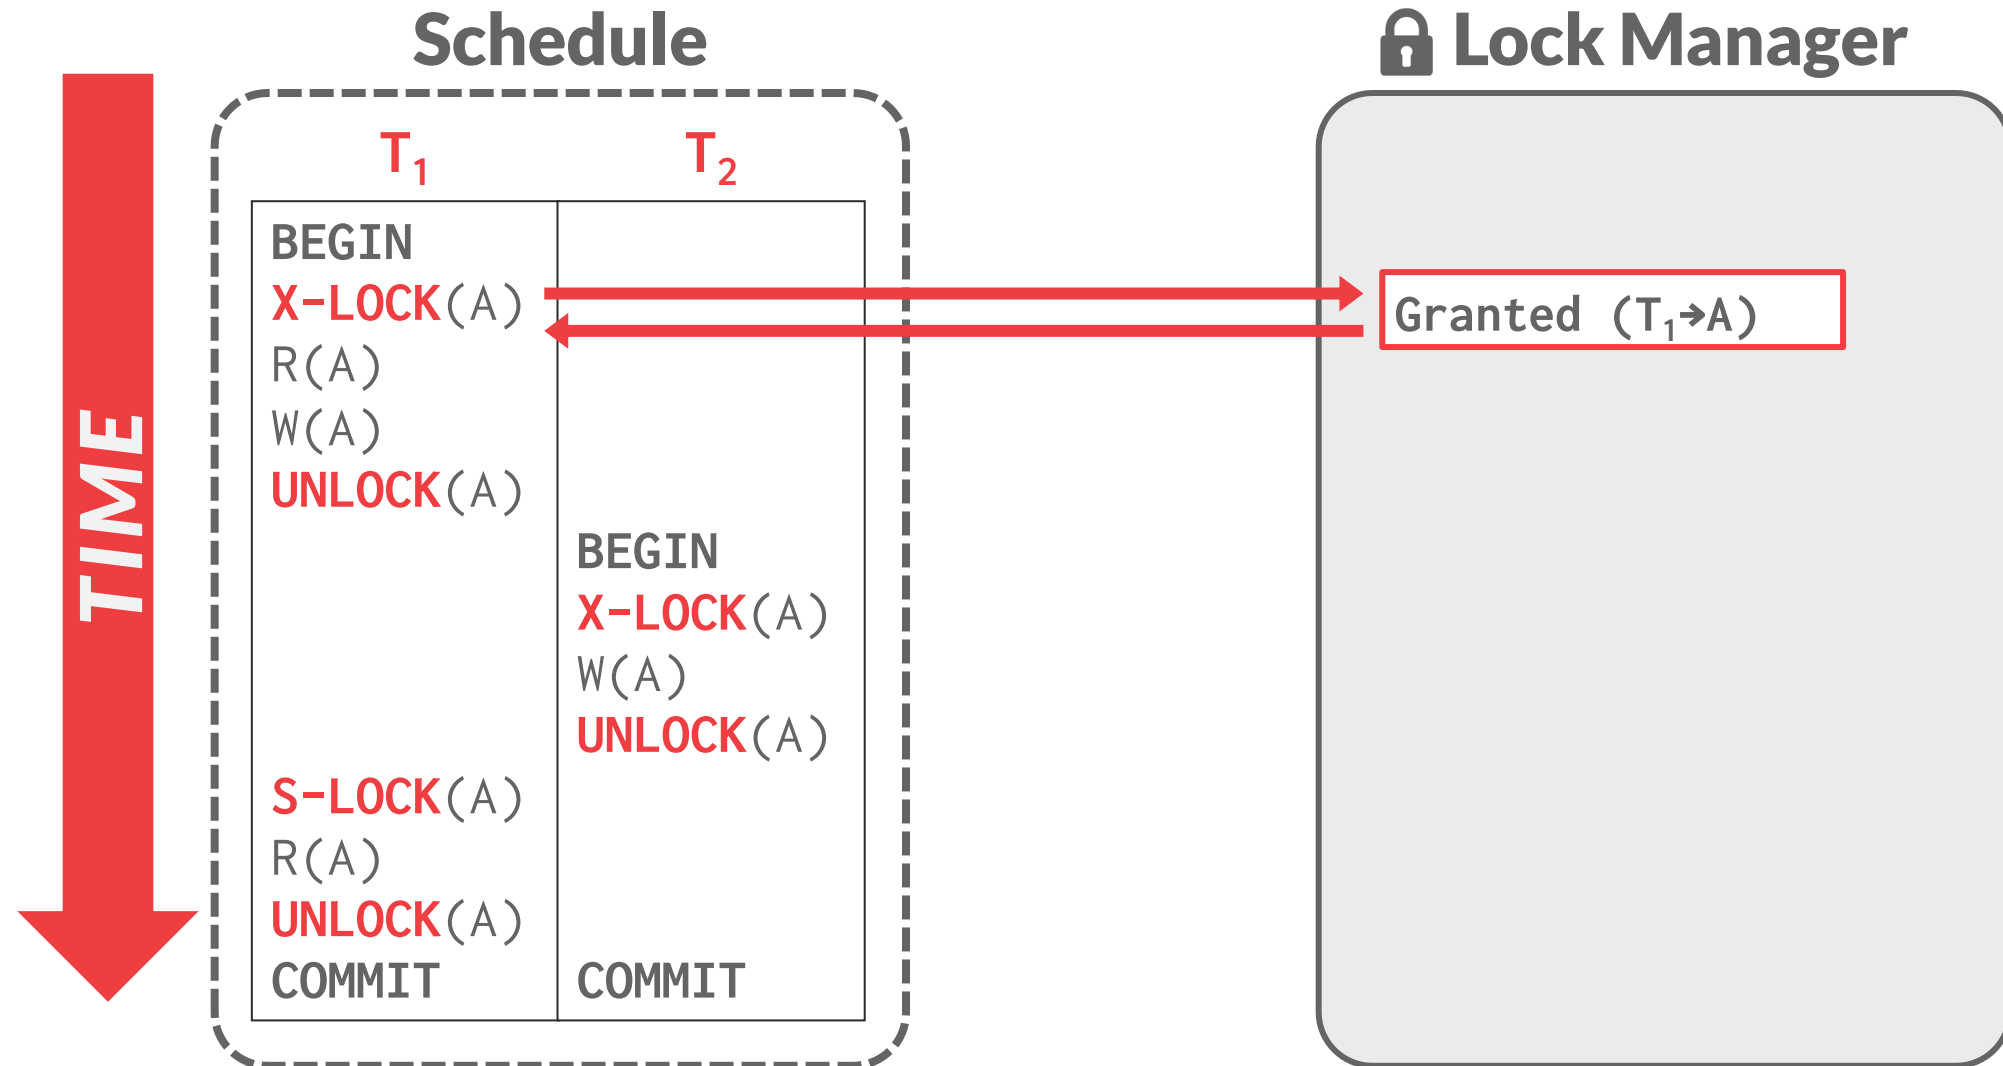


- A schedule that is equivalent to some serial execution of the transactions.
- Need to reason about conflicting operations.
- Two operations **conflict** if:
 - They are by different transactions,
 - They are on the same object and one of them is a write.
- Interleaved Execution **Anomalies**
 - Read-Write Conflicts (**R-W**). Also called **Unrepeatable Read**.
 - Write-Read Conflicts (**W-R**). Also called **Dirty Read**.
 - Write-Write Conflicts (**W-W**). Also called **Lost Update**.

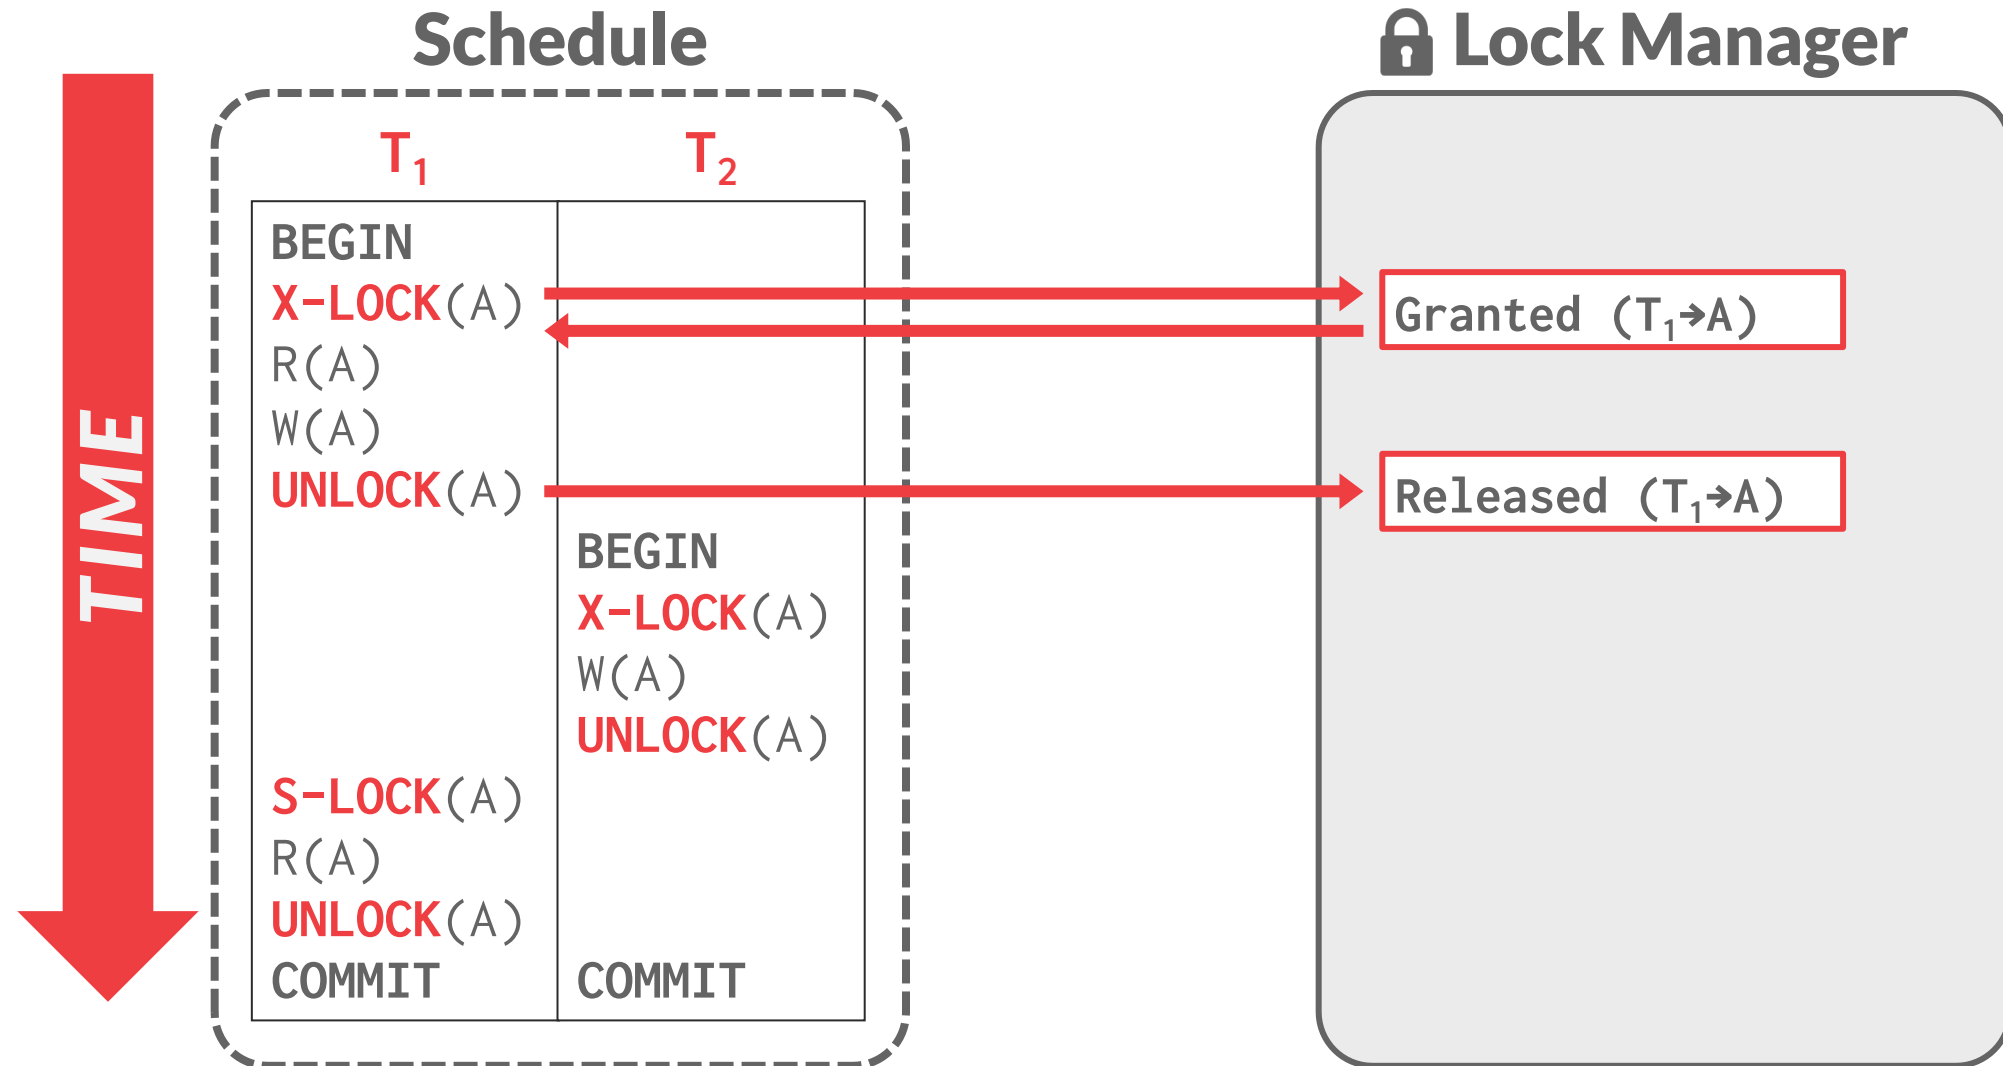
LOCKING WITH S AND X LOCKS AT THE ROW-LEVEL



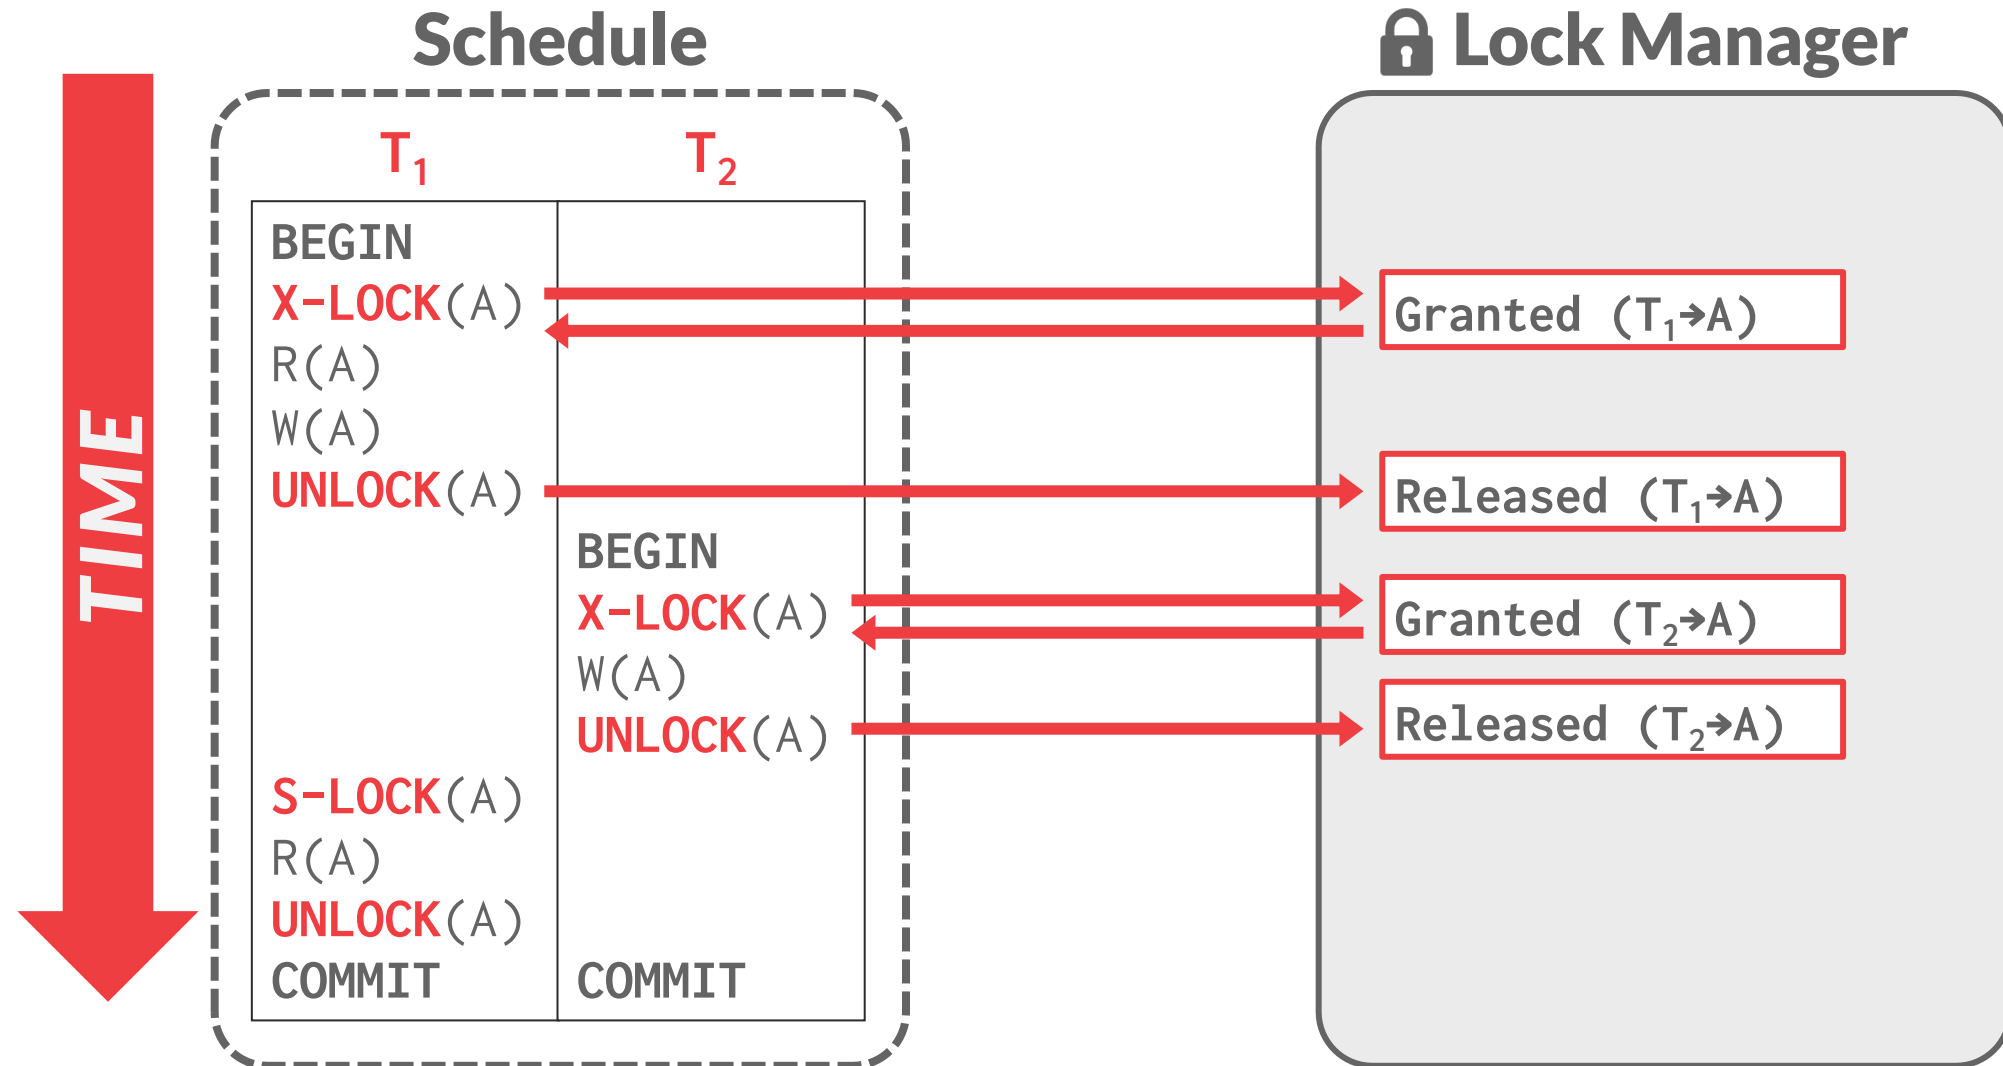
LOCKING WITH S AND X LOCKS AT THE ROW-LEVEL



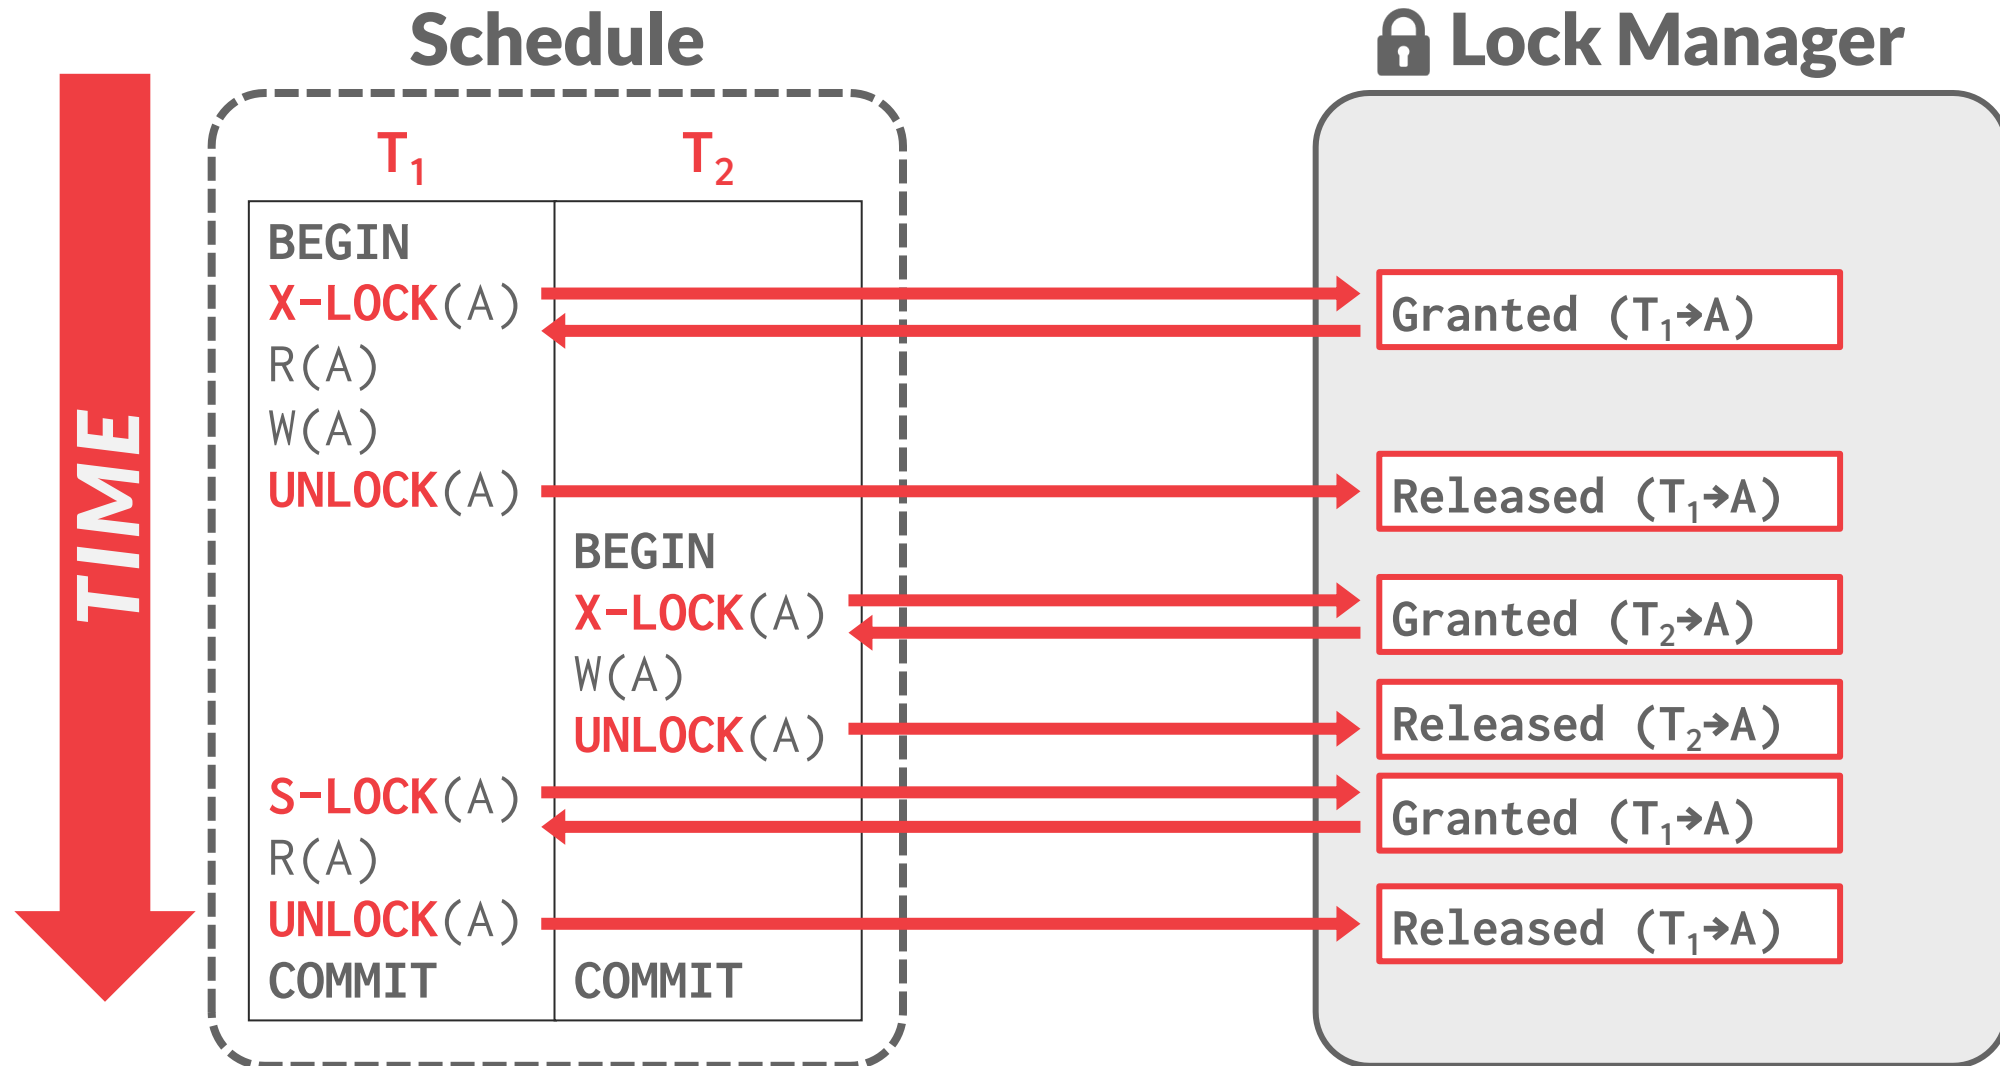
LOCKING WITH S AND X LOCKS AT THE ROW-LEVEL



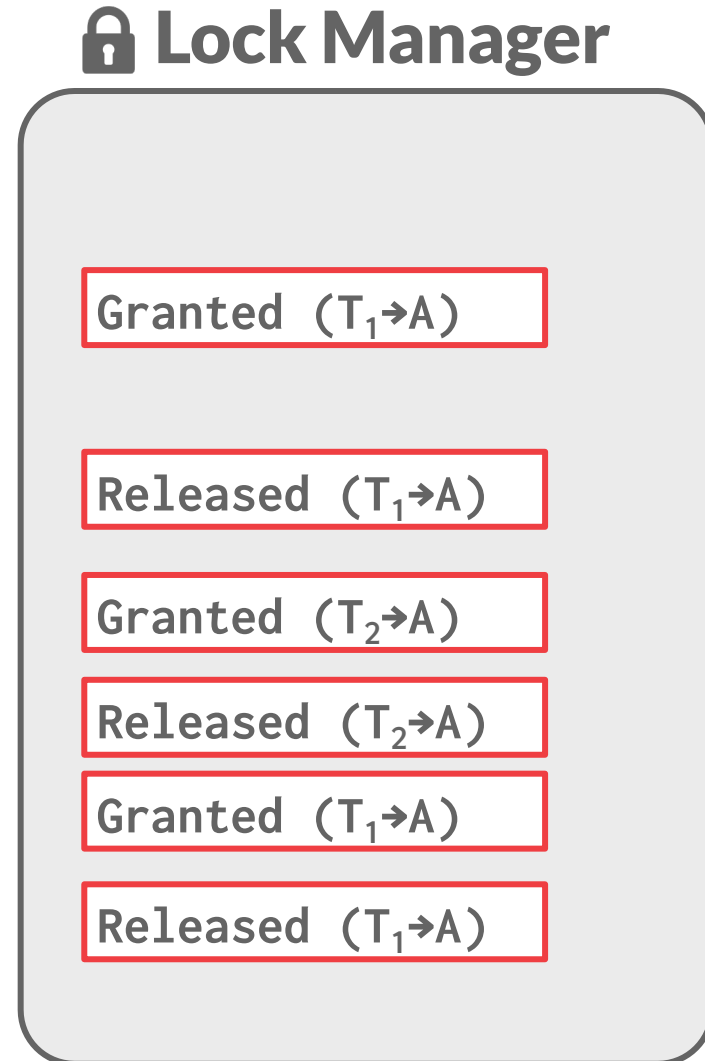

LOCKING WITH S AND X LOCKS AT THE ROW-LEVEL



LOCKING WITH S AND X LOCKS AT THE ROW-LEVEL




LOCKING WITH S AND X LOCKS AT THE ROW-LEVEL



GRANULARITIES PAPER: KEY CHALLENGE


A Resource Hierarchy



GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.


A Resource Hierarchy



GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?


A Resource Hierarchy



GRANULARITIES PAPER: KEY CHALLENGE

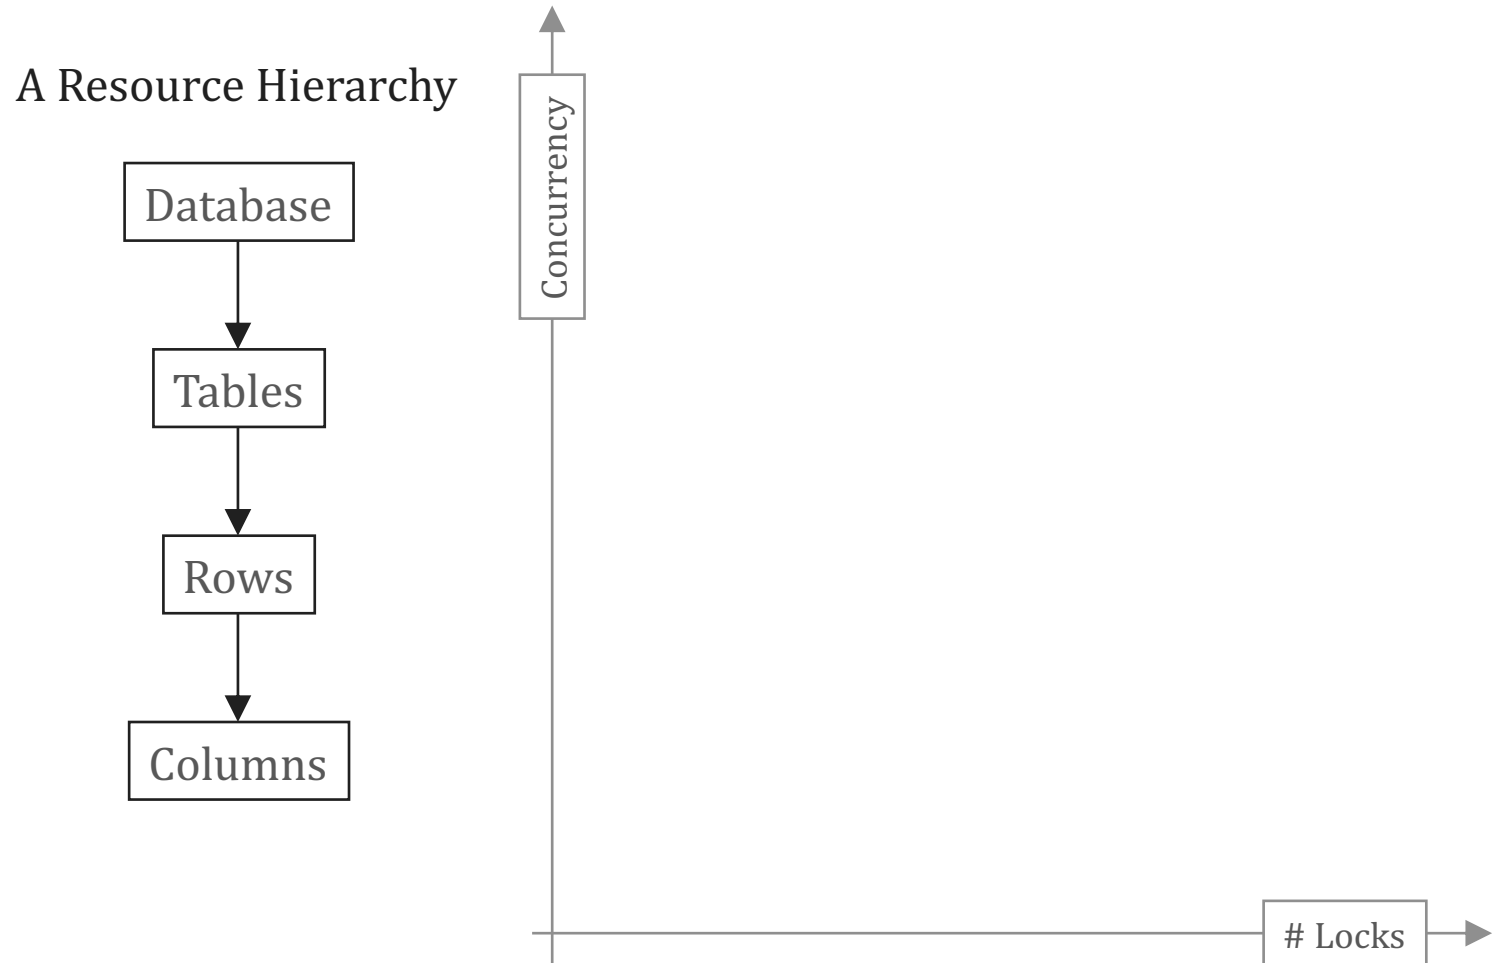
- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.

A Resource Hierarchy



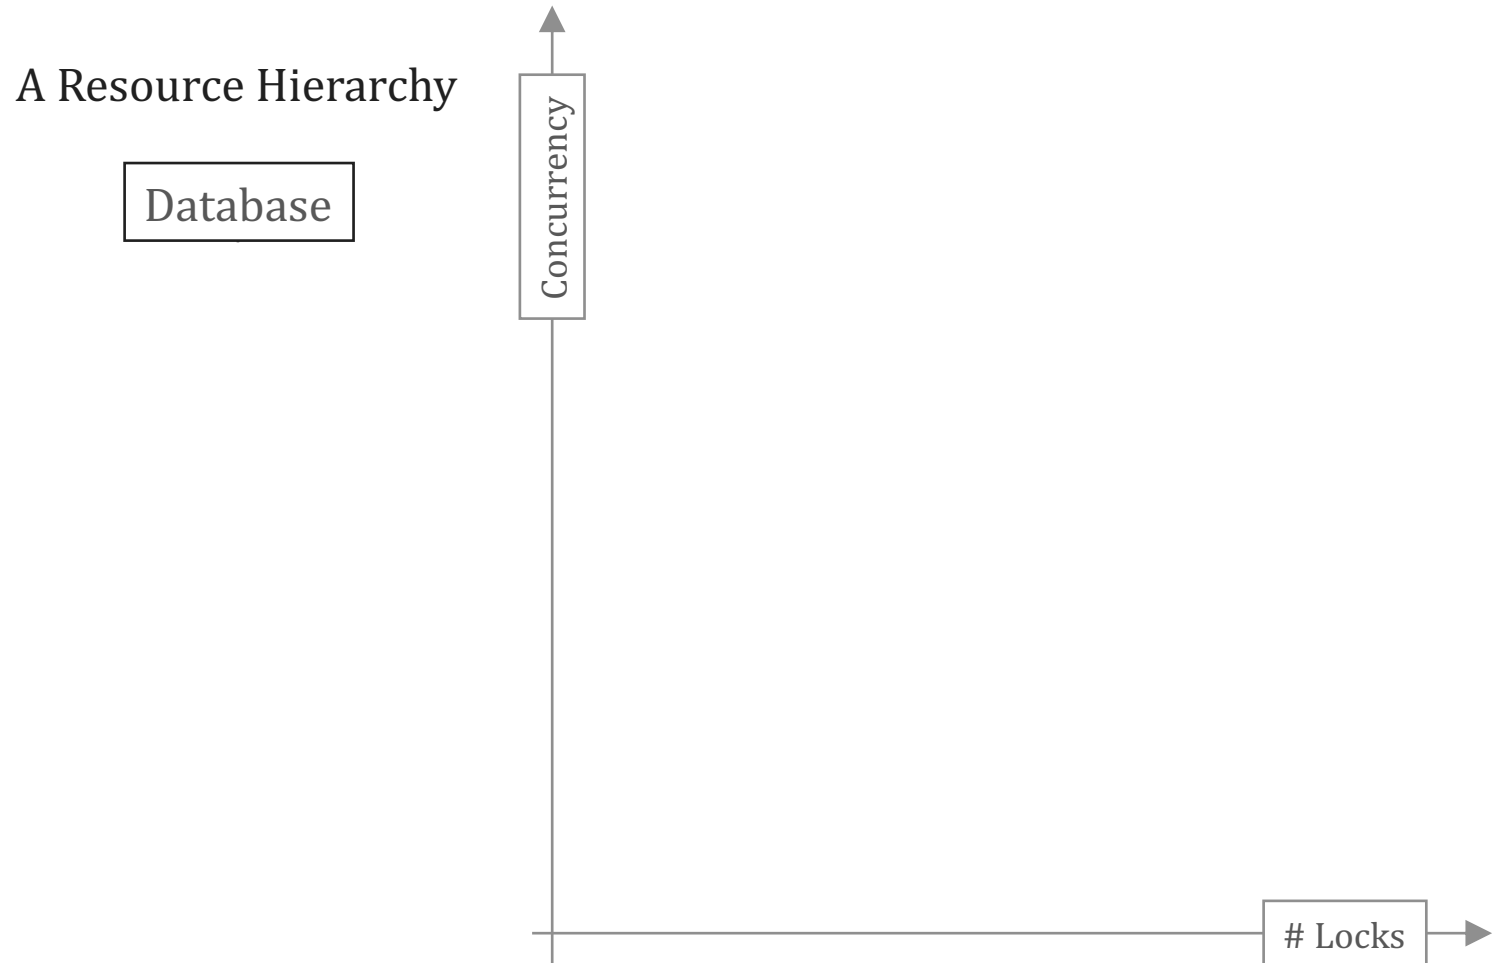
GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.



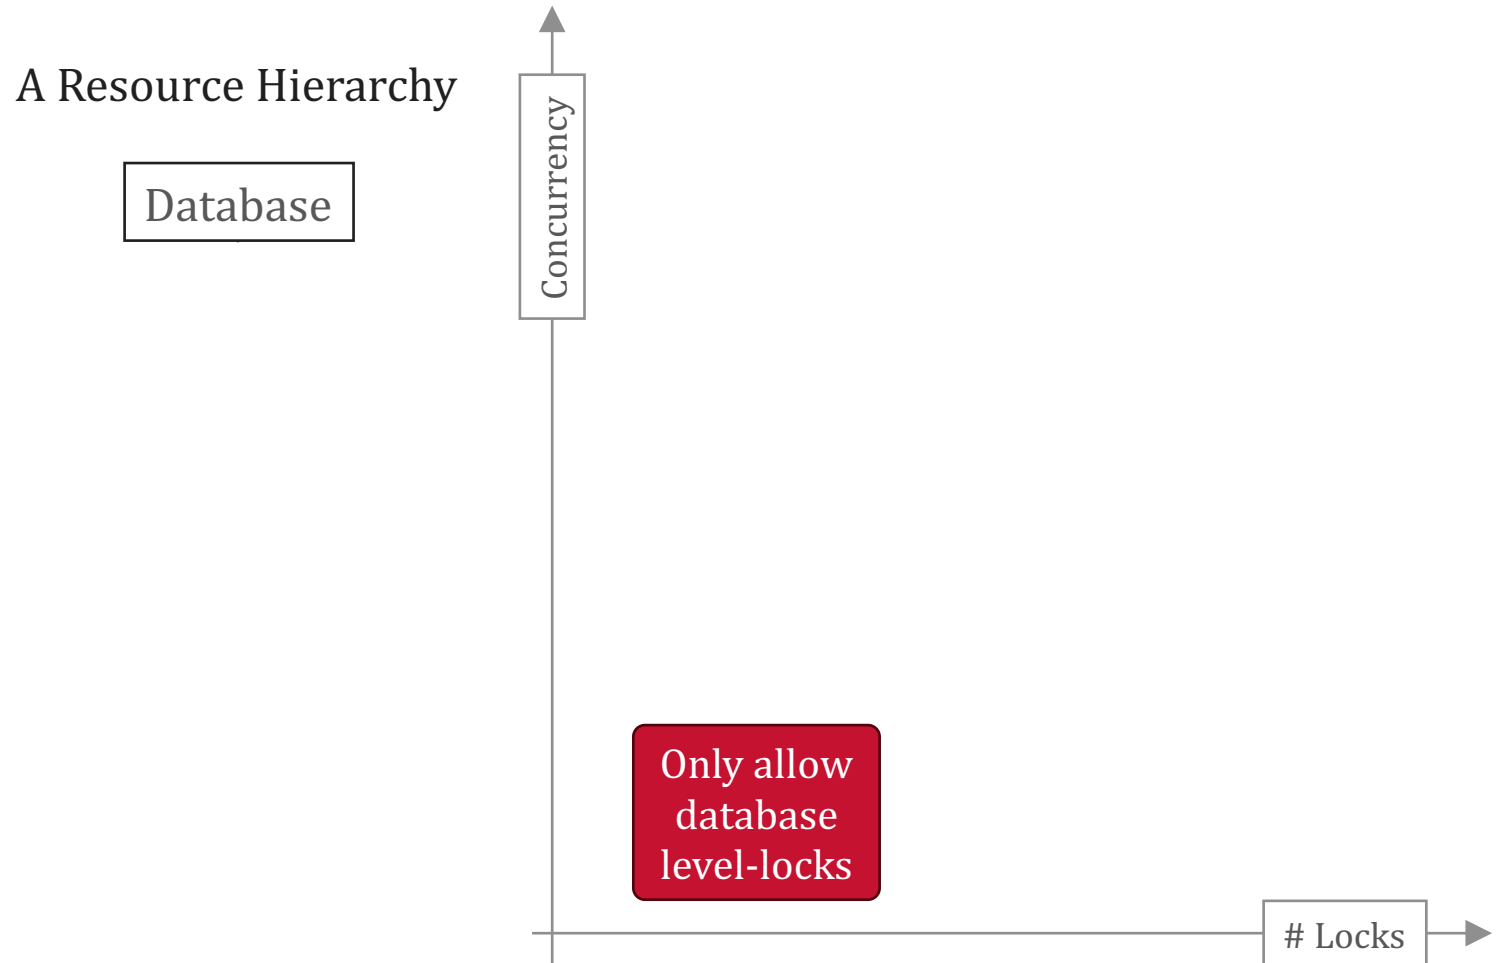
GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.



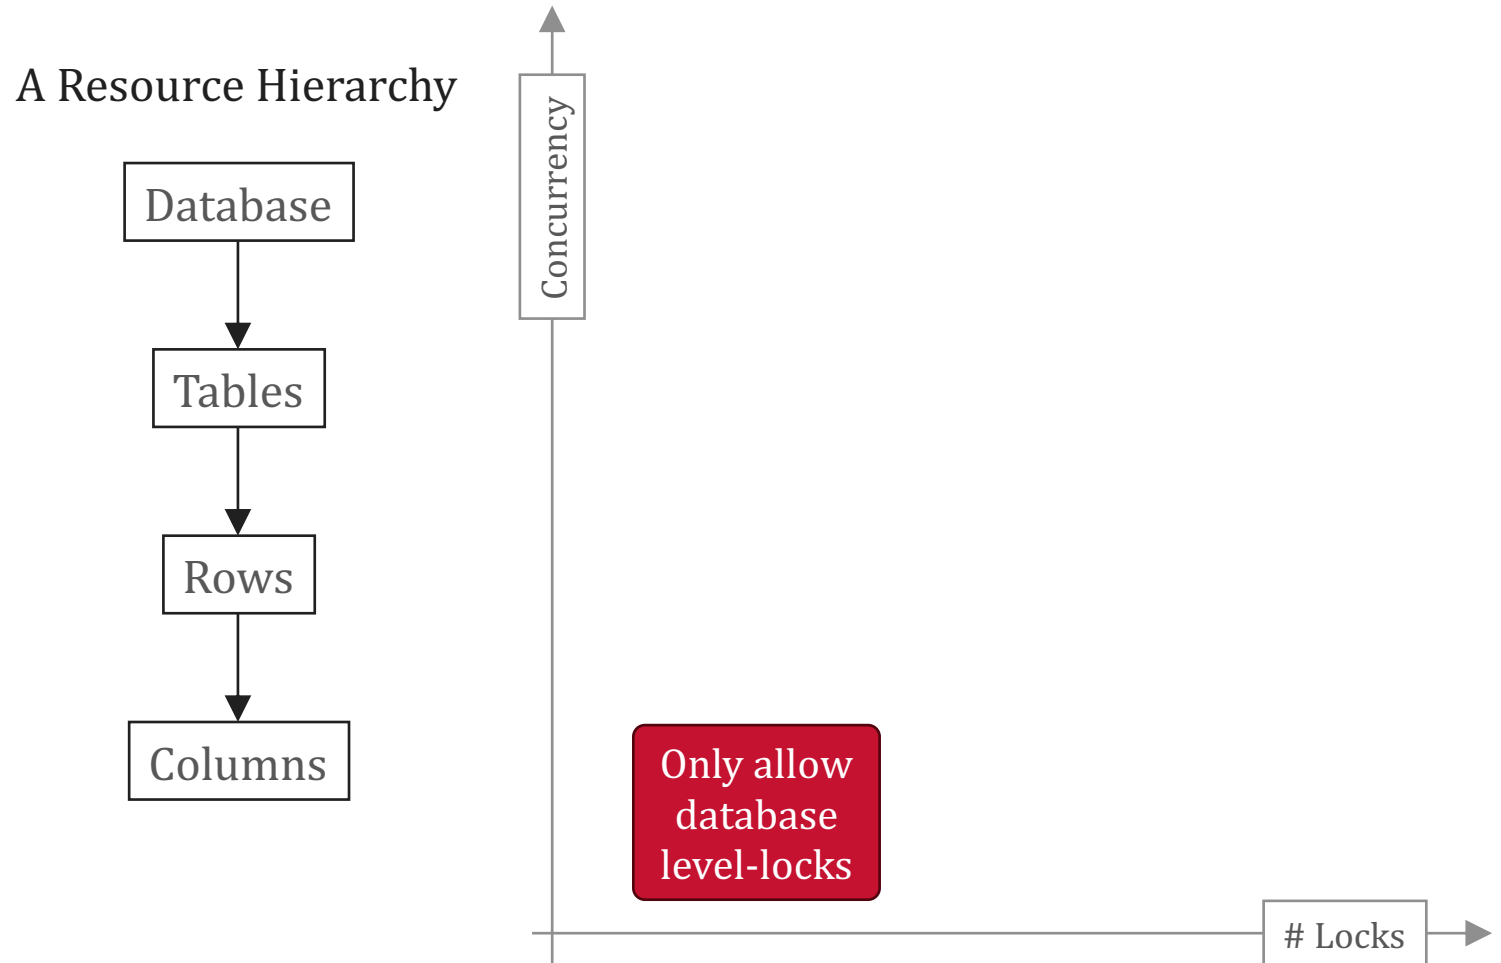
GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.



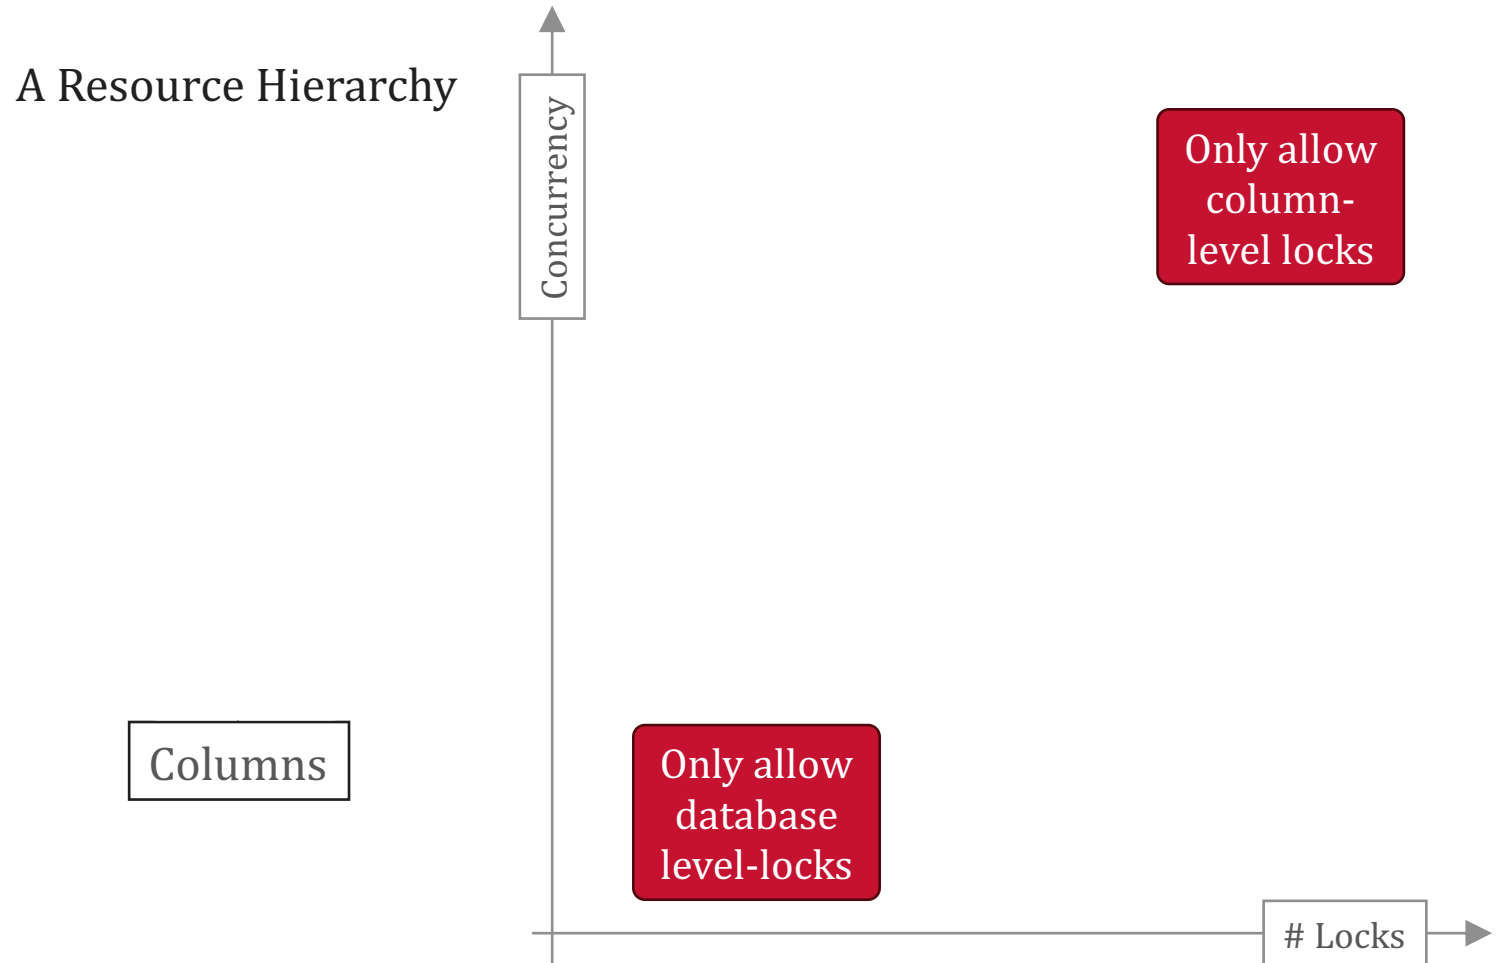
GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.



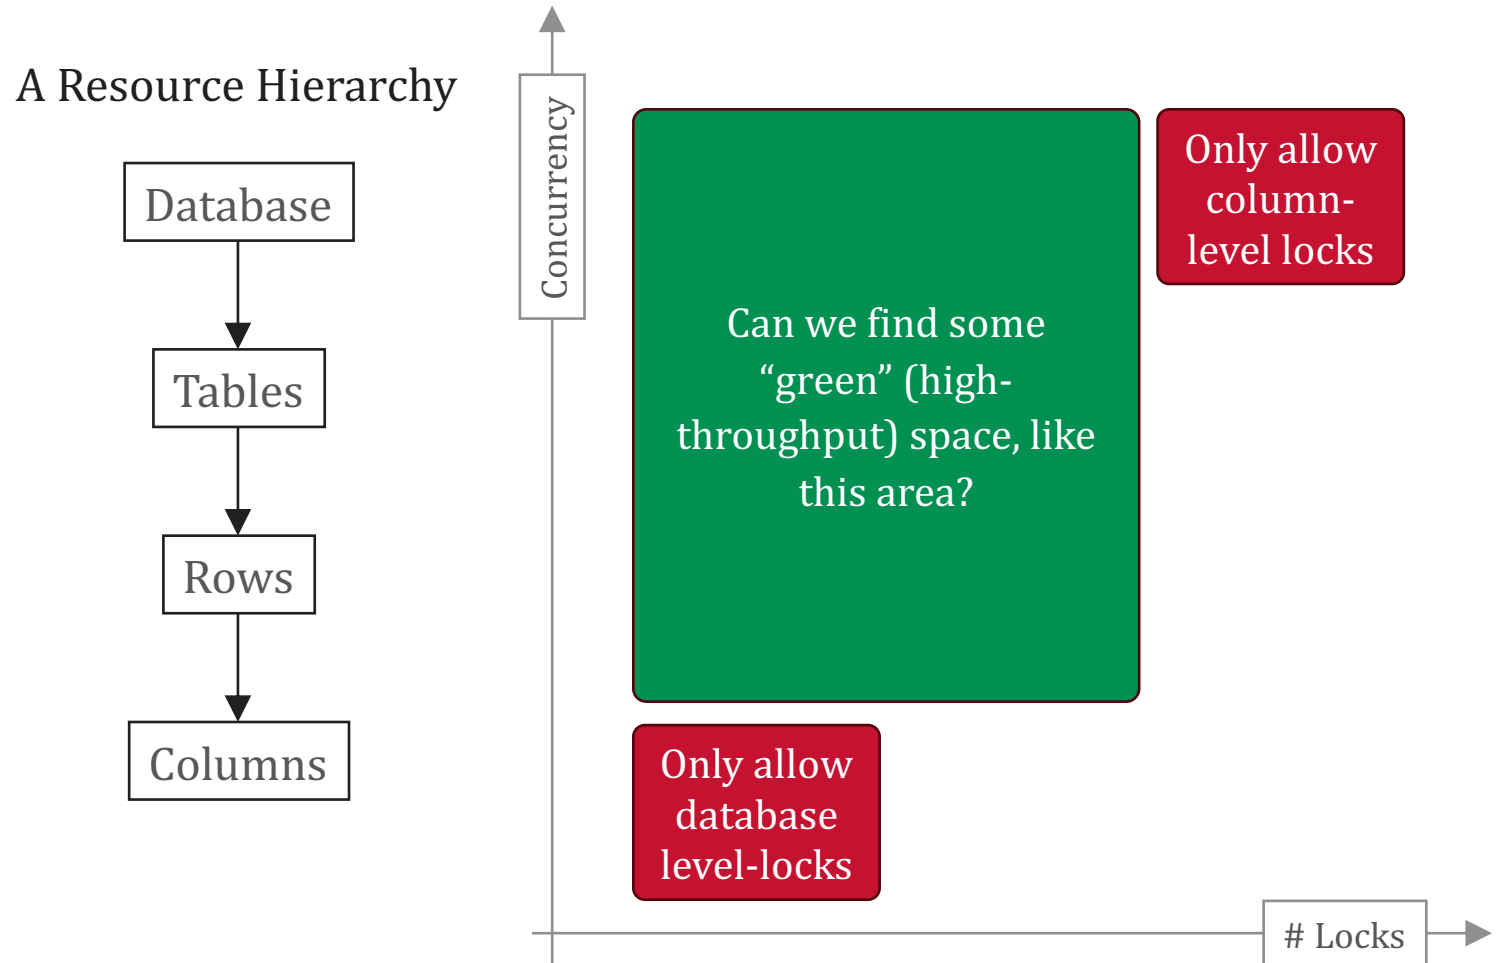
GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.



GRANULARITIES PAPER: KEY CHALLENGE

- Correct (provably so) locking protocol that balances.
 - # locks that are acquired.
 - Amount of concurrency that is allowed by the protocol.
- Key metric is the transaction (txn) throughput: how many txns can we commit per second?
- Want **High** throughput, but simple methods result in **Low** throughput.



THE SOLUTION (DECEPTIVELY SIMPLE, AND HENCE BRILLIANT)

- Work with complex/**DAG resource graphs**.
- Use a well-formed protocol to acquire locks in top-down manner.
- Introduce the notion of **“intention” locks** to allow a txn to indicate that they will grab a “regular” lock (S or X) on a resource below in the hierarchy.
- Develop a novel lock-compatibility matrix that allows **balancing # locks with the “granularity” of locking**.
- Can **offer different degrees of consistency** (trading performance for lower consistency), allowing **concurrent transaction** to operate at different consistency levels.




1998 ACM Turing Award

INTENTION LOCKS

- An intention lock allows a higher-level node to be locked in **shared** or **exclusive** mode without having to check all descendent nodes.
- If a node is locked in an intention mode, then some txn is doing explicit locking at a lower level in the tree.

A Resource Hierarchy



INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



An X lock will be acquired somewhere in the gray region.
An intention lock discloses the intent to do additional locking below.

INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



Now, this entire gray region is considered locked in X mode. No need to acquire X locks here. The lock at the top level “covers” this whole region.

INTENTION LOCKS

- **Intention-Shared (IS)**

- Indicates explicit locking at lower level with **S** locks.
- Intent to get **S** lock(s) at finer granularity.


- **Intention-Exclusive (IX)**

- Indicates explicit locking at lower level with **X** locks.
- Intent to get **X** lock(s) at finer granularity.

- **Shared+Intention-Exclusive (SIX)**

- The subtree rooted by that node is locked explicitly in **S** mode and explicit locking is being done at a lower level with **X** locks.

A Resource Hierarchy



COMPATIBILITY MATRIX


		T_2 Wants				
		IS	IX	S	SIX	X
T_1 Holds	IS	✓	✓	✓	✓	✗
	IX	✓	✓	✗	✗	✗
	S	✓	✗	✓	✗	✗
	SIX	✓	✗	✗	✗	✗
	X	✗	✗	✗	✗	✗

LOCKING PROTOCOL: WELL FORMED


- Each txn obtains an appropriate lock at the highest level of the hierarchy.
- To get **S** or **IS** lock on a node, the txn must hold at least **IS** on parent node.
- To get **X**, **IX**, or **SIX** on a node, must hold at least **IX** on parent node.
- All lock are acquired top-down, so if a txn has an intention lock, every other txn will see that before they acquire lock at a lower level in the resource hierarchy.
- Locks released leaf to root, or all at once at the end of the txn.
- Need non-intention locks somewhere in the resource hierarchy (so can't have txns that only do intention locks).

The partial ordering of the lock modes.

Higher is more restrictive.




A Resource Hierarchy



EXAMPLE


- T_1 – Get the balance of Alice's account.
- T_2 – Increase Bob's account by 1%.
- What locks should these txns obtain?
 - Exclusive + Shared for leaf nodes of lock tree.
 - Special Intention locks for higher levels.



EXAMPLE

- T_1 – Get the balance of Alice's account.
- T_2 – Increase Bob's account by 1%.
- What locks should these txns obtain?
 - Exclusive + Shared for leaf nodes of lock tree.
 - Special Intention locks for higher levels.


Read Alice's record in **R**.



EXAMPLE

- T_1 – Get the balance of Alice's account.
- T_2 – Increase Bob's account by 1%.
- What locks should these txns obtain?
 - Exclusive + Shared for leaf nodes of lock tree.
 - Special Intention locks for higher levels.


Read Alice's record in **R**.



EXAMPLE


- T_1 – Get the balance of Alice's account.
- T_2 – Increase Bob's account by 1%.
- What locks should these txns obtain?
 - Exclusive + Shared for leaf nodes of lock tree.
 - Special Intention locks for higher levels.

Read Alice's record in **R**.




EXAMPLE

- T_1 – Get the balance of Alice's account.
- T_2 – Increase Bob's account by 1%.
- What locks should these txns obtain?
 - Exclusive + Shared for leaf nodes of lock tree.
 - Special Intention locks for higher levels.




EXAMPLE

- T_1 – Get the balance of Alice's account.
- T_2 – Increase Bob's account by 1%.
- What locks should these txns obtain?
 - Exclusive + Shared for leaf nodes of lock tree.
 - Special Intention locks for higher levels.





DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.




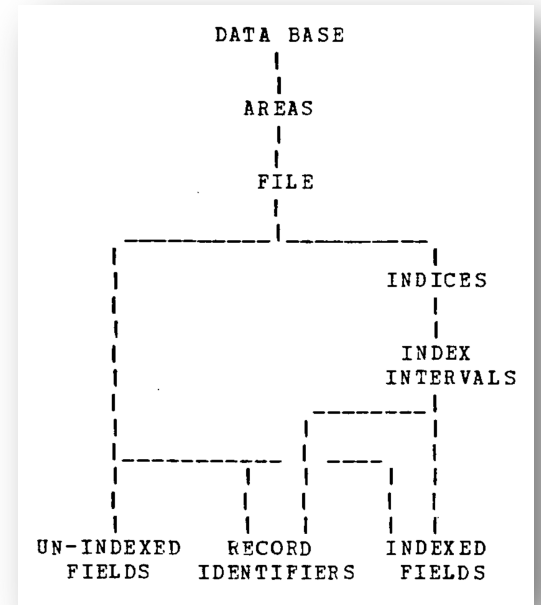
DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.





DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.




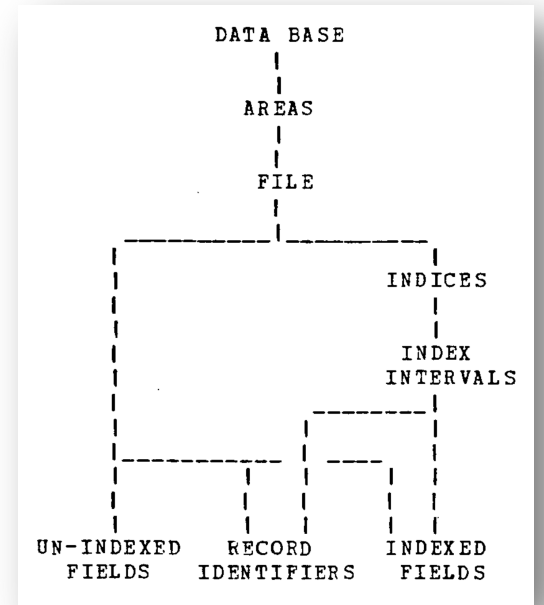
DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.





DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.




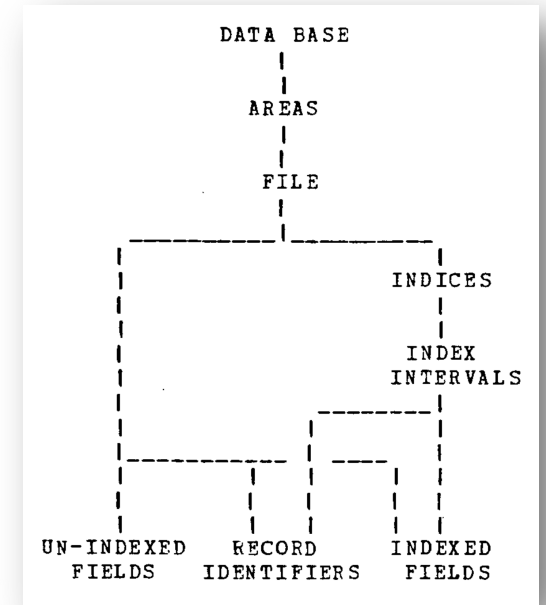
DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.
- Now in the new area, we may not have the appropriate locks on the ancestors.




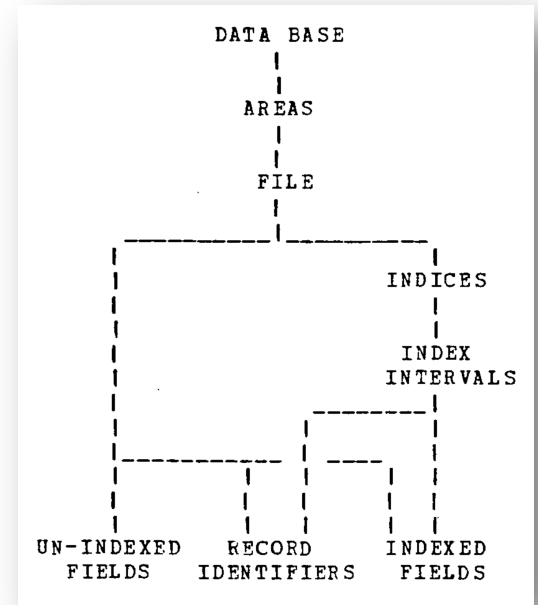
DYNAMIC LOCK GRAPH

- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.
- Now in the new area, we may not have the appropriate locks on the ancestors.



DYNAMIC LOCK GRAPH


- Data can move around (in the resource graph) in the same transaction; e.g., an update that moves data from one part of an index to another.
- Now in the new area, we may not have the appropriate locks on the ancestors.
- Solution: Before moving data, **both the old and new locations** must have an X lock, and the **well-formed protocol must be preserved** so that top-down traversal by another txn does not grab a conflicting lock.



LOCK SCHEDULES AND UPGRADES


- A single resource may have multiple locks. Group mode is the highest level of lock on that resource.

e.g., if a group has an mix of IS and IX locks, the group mode is IX.




LOCK SCHEDULES AND UPGRADES

- A single resource may have multiple locks. Group mode is the highest level of lock on that resource.
e.g., if a group has an mix of IS and IX locks, the group mode is IX.




Current Mode	New Mode				
	IS	IX	S	SIX	X
IS	IS	IX	S	SIX	X
IX	IX	IX	SIX	SIX	X
S	S	SIX	S	SIX	X
SIX	SIX	SIX	SIX	SIX	X
X	X	X	X	X	X

The partial ordering of the lock modes.



LOCK SCHEDULES AND UPGRADES


- A single resource may have multiple locks. Group mode is the highest level of lock on that resource.
e.g., if a group has an mix of IS and IX locks, the group mode is IX.



- Some notion of fairness (e.g., FIFO) is needed so some txn does not wait forever on a lock request.
- Lock upgrade request: Give priority to a txn in the pending queue if it is already part of the granted group.
This txn is already holding a resource. Try to get this txn to finish quickly, and free up this resource.


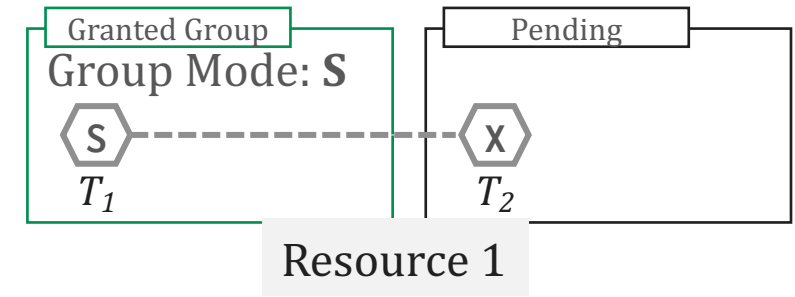
Current Mode	New Mode				
	IS	IX	S	SIX	X
IS	IS	IX	S	SIX	X
IX	IX	IX	SIX	SIX	X
S	S	SIX	S	SIX	X
SIX	SIX	SIX	SIX	SIX	X
X	X	X	X	X	X

The partial ordering of the lock modes.




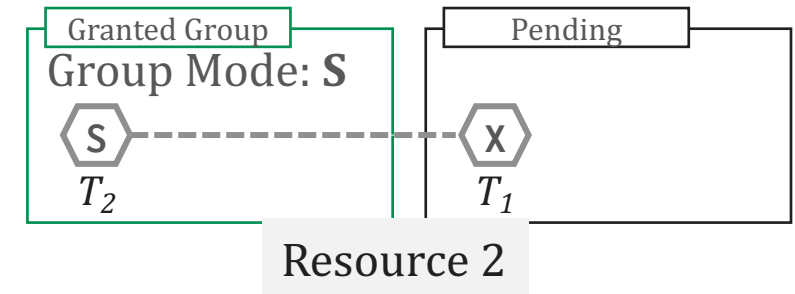

DEADLOCKS

- We can now have deadlocks.





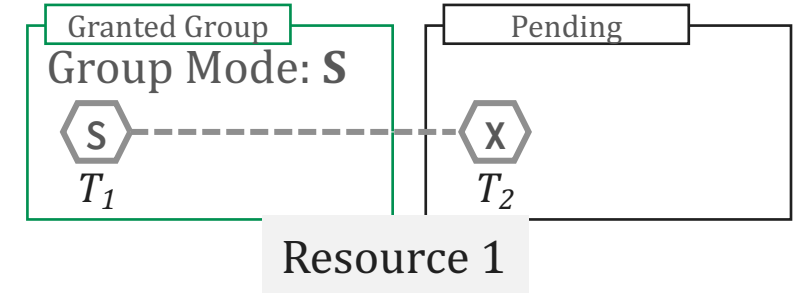
DEADLOCKS

- We can now have deadlocks.



DEADLOCKS


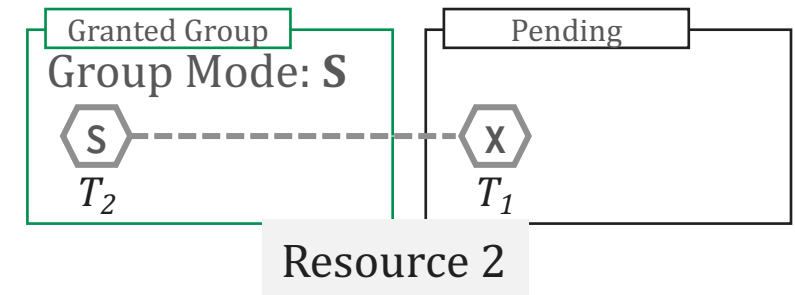
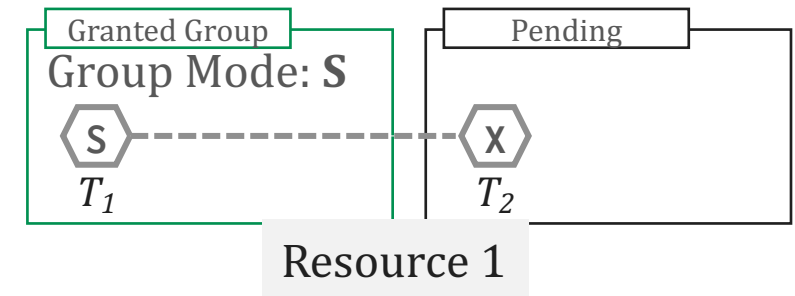
- We can now have deadlocks.
- Need a mechanism to either detect deadlocks, or prevent deadlocks.



DEADLOCKS

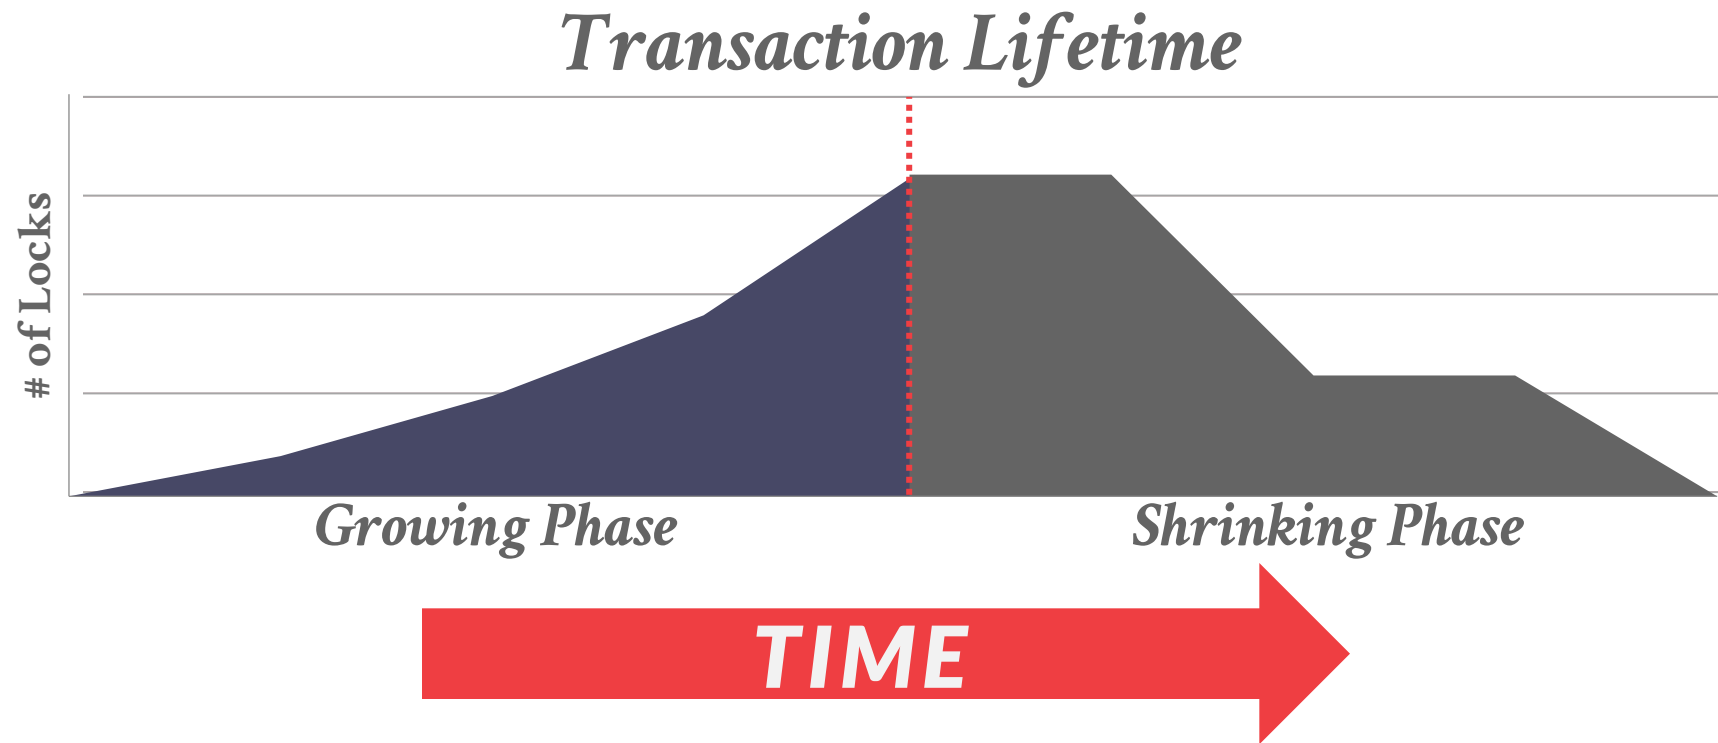
- We can now have deadlocks.
- Need a mechanism to either detect deadlocks, or prevent deadlocks.
- Deadlock detection: Construct and periodically examine the wait-for-graph. Pick a victim (oldest txn or newest txn) to break the deadlock. Abort the victim txn and restart it (perhaps after some sleep/delay).
- Deadlock prevention: Abort a txn as soon as it waits for another txn.

Wound-wait or wait-die (see the intro DB class for details).



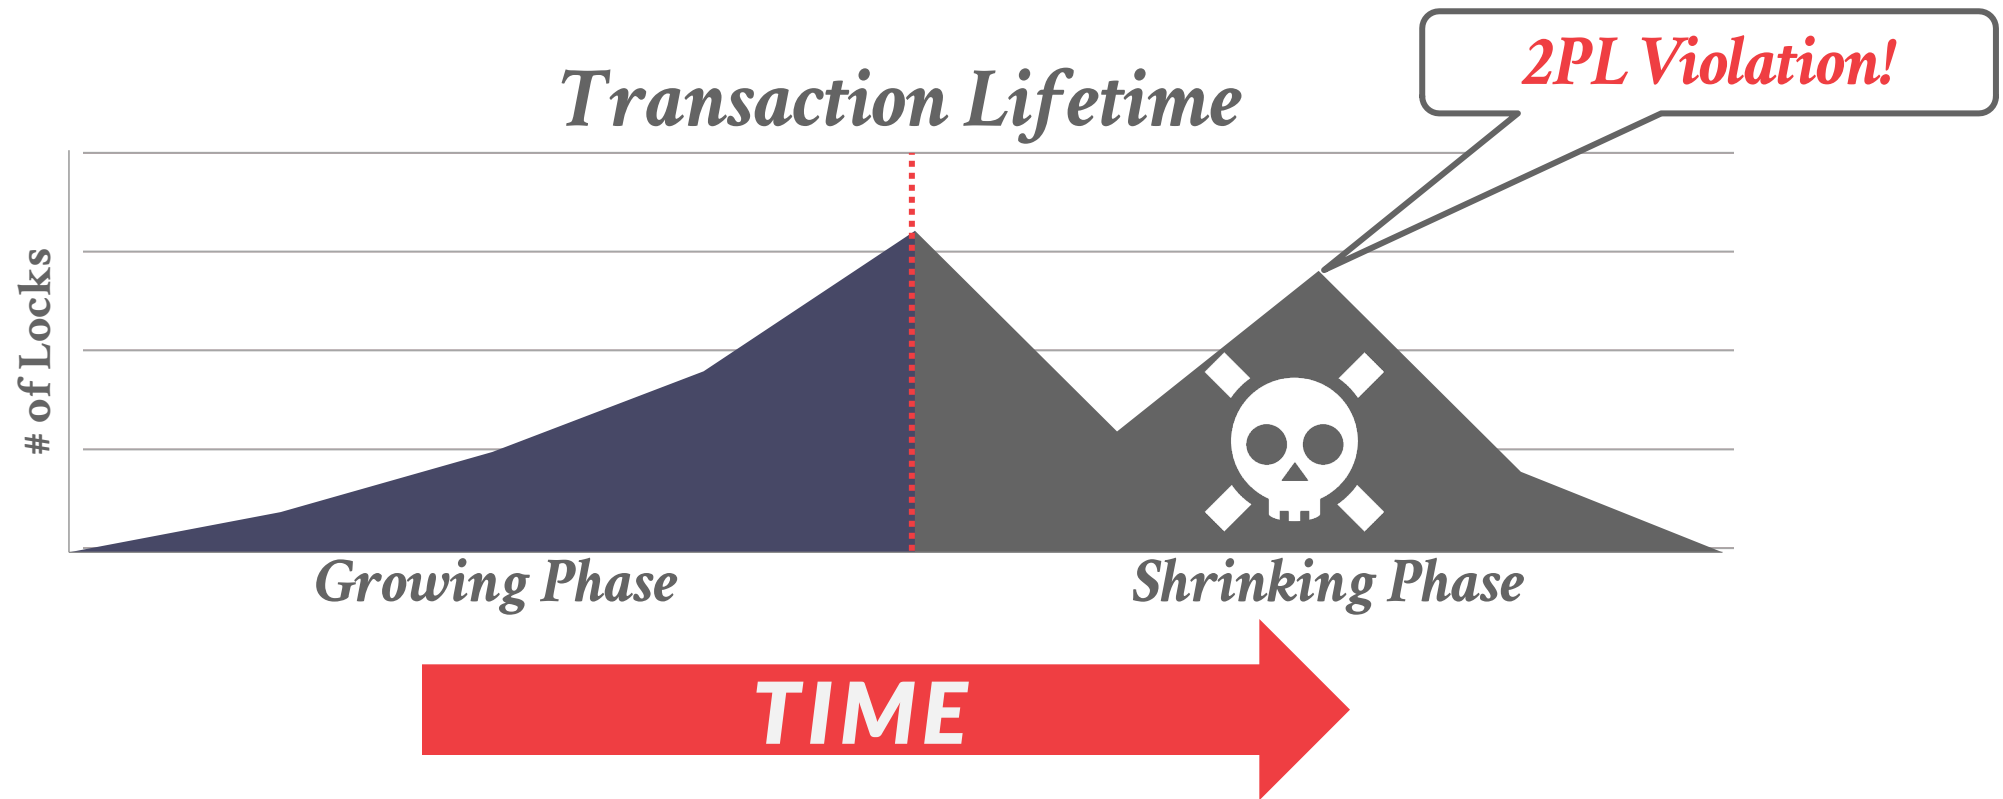
TWO PHASE LOCKING (2PL)

- Txn has 2 phases: a growing (acquire lock phase) and a subsequent drop lock phase.

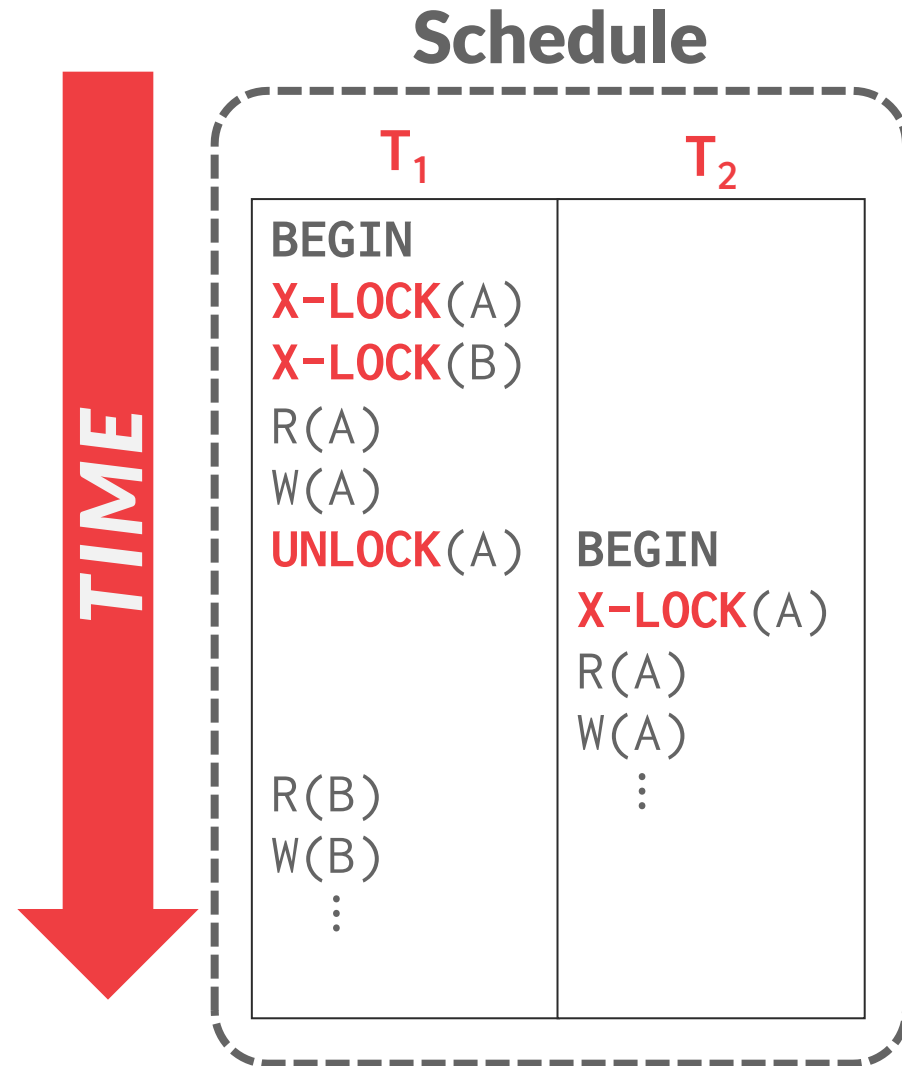


TWO PHASE LOCKING (2PL)

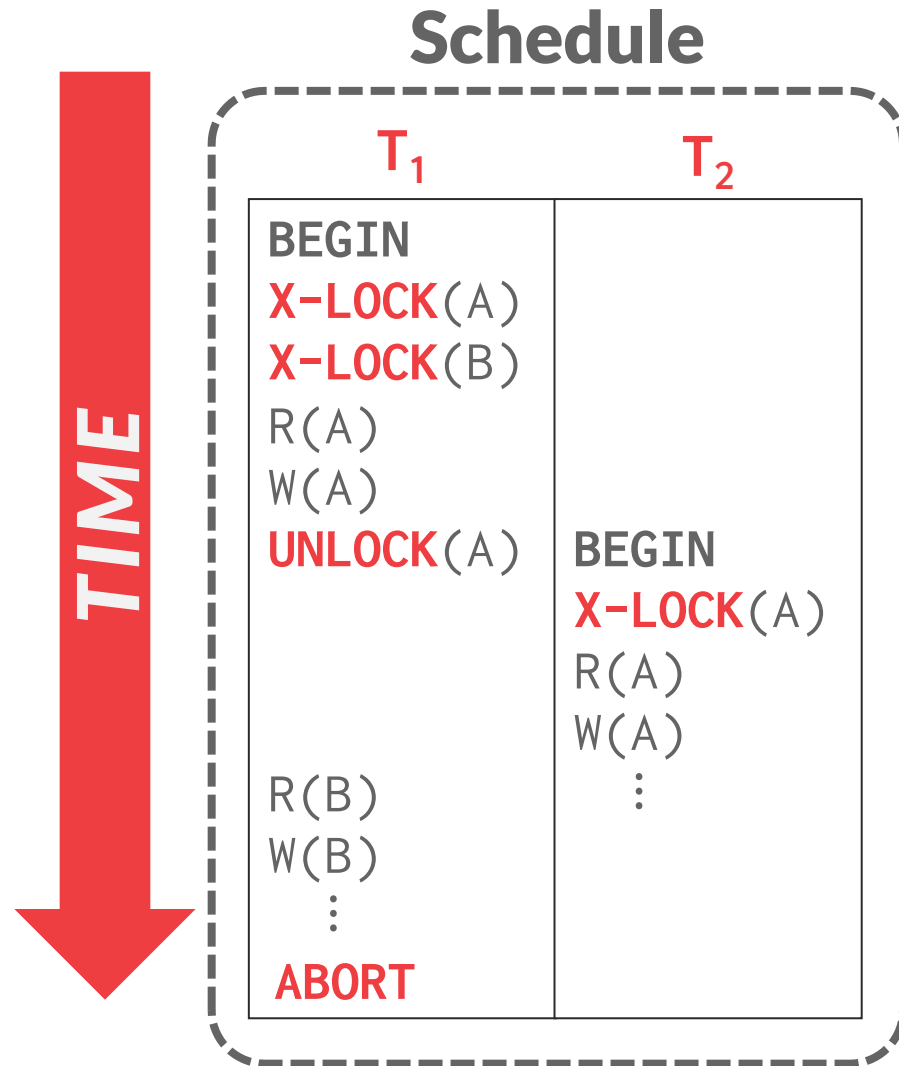
- Txn has 2 phases: a growing (acquire lock phase) and a subsequent drop lock phase.
- Can't acquire a lock after the first lock is released.




2PL PERMITS CASCADING ABORTS



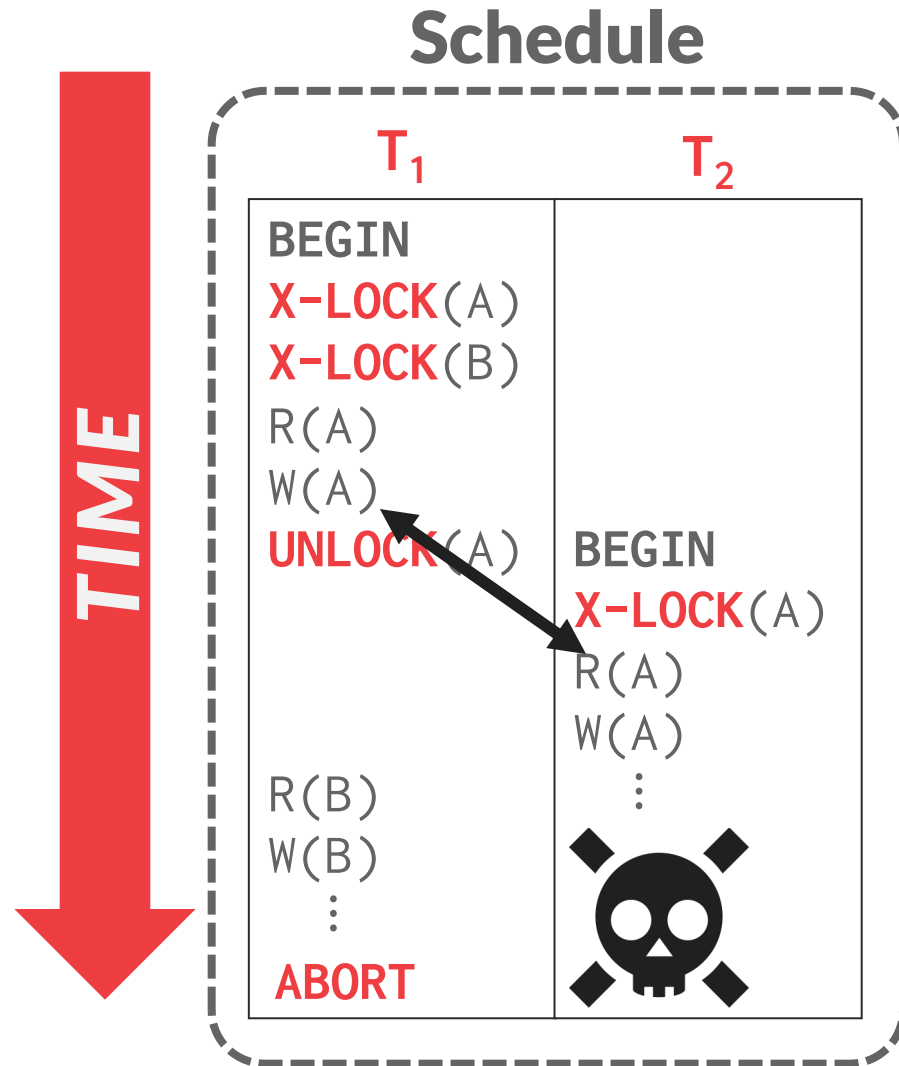
2PL PERMITS CASCADING ABORTS



2PL PERMITS CASCADING ABORTS

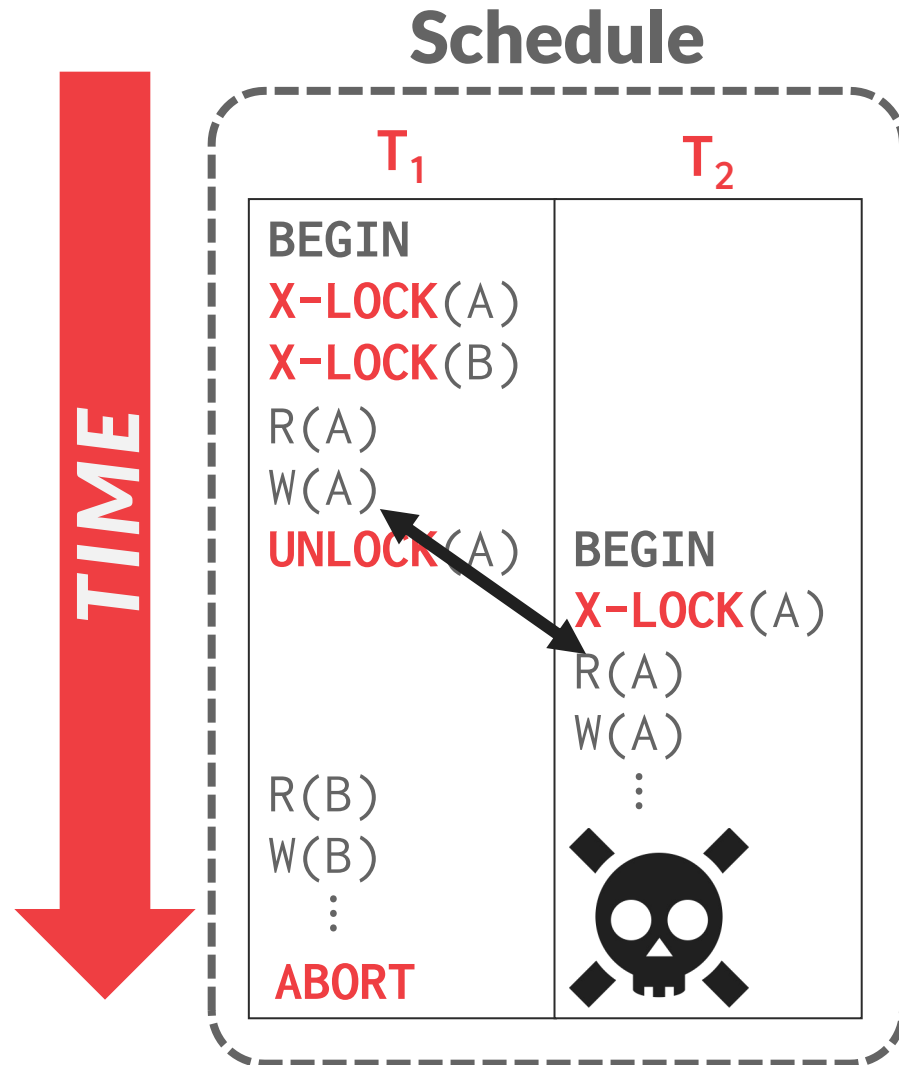


2PL PERMITS CASCADING ABORTS



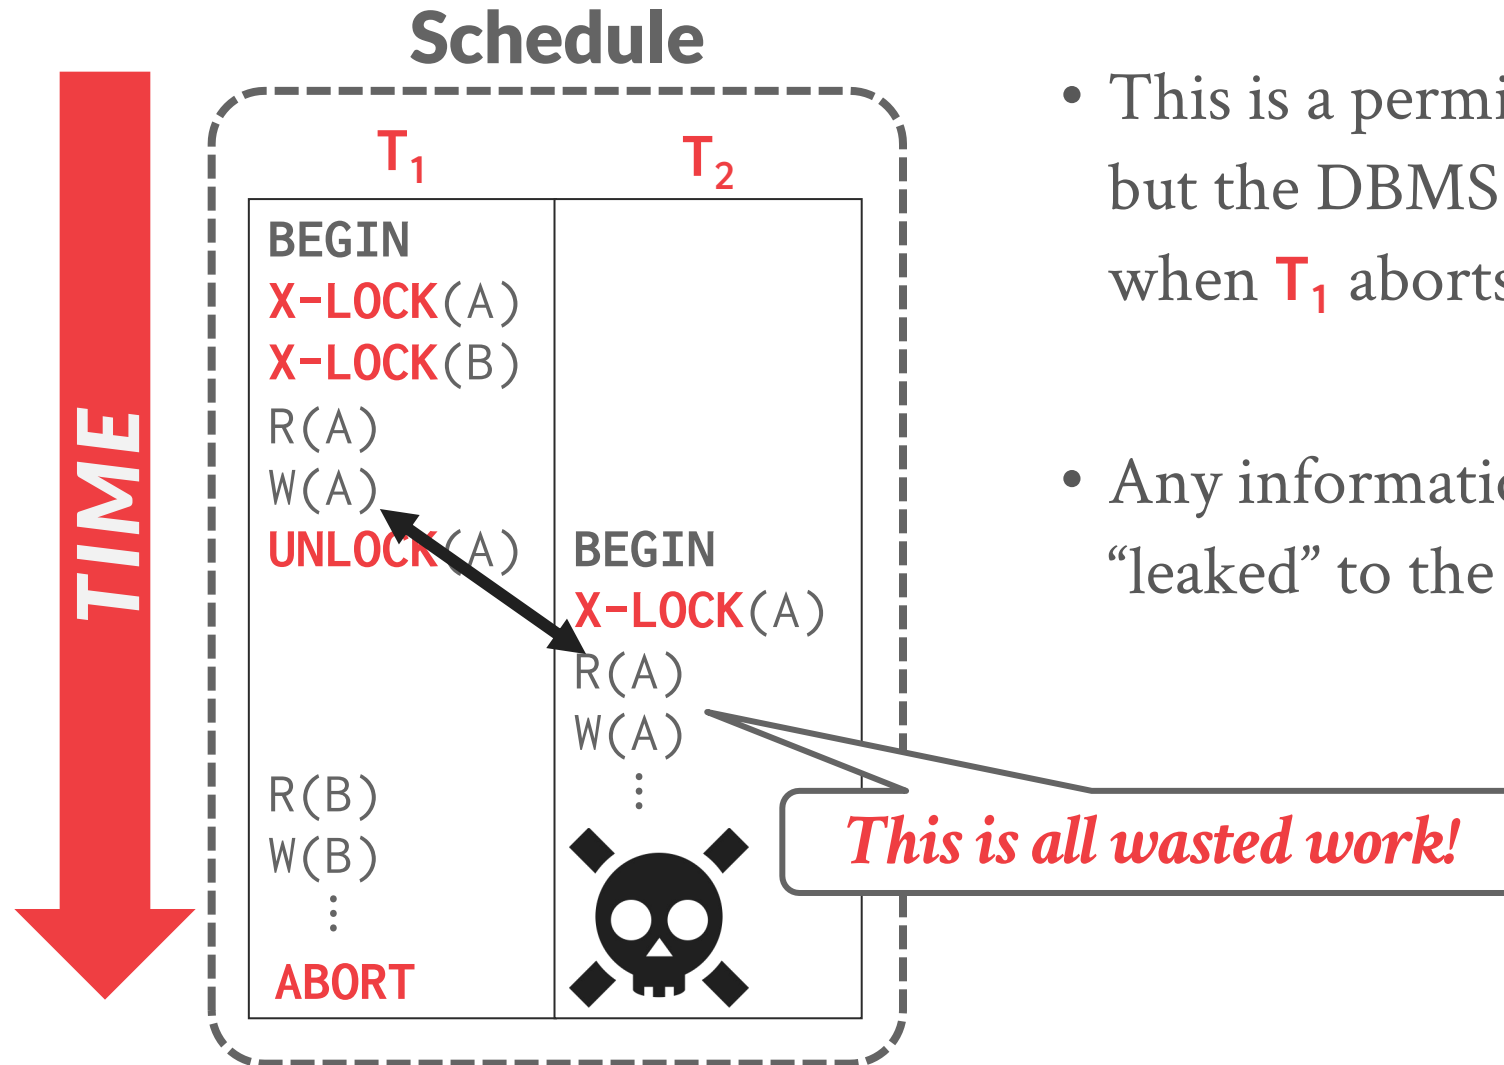
- This is a permissible schedule in 2PL, but the DBMS has to also abort T_2 when T_1 aborts.

2PL PERMITS CASCADING ABORTS



- This is a permissible schedule in 2PL, but the DBMS has to also abort T_2 when T_1 aborts.
- Any information about T_1 cannot be “leaked” to the outside world.

2PL PERMITS CASCADING ABORTS

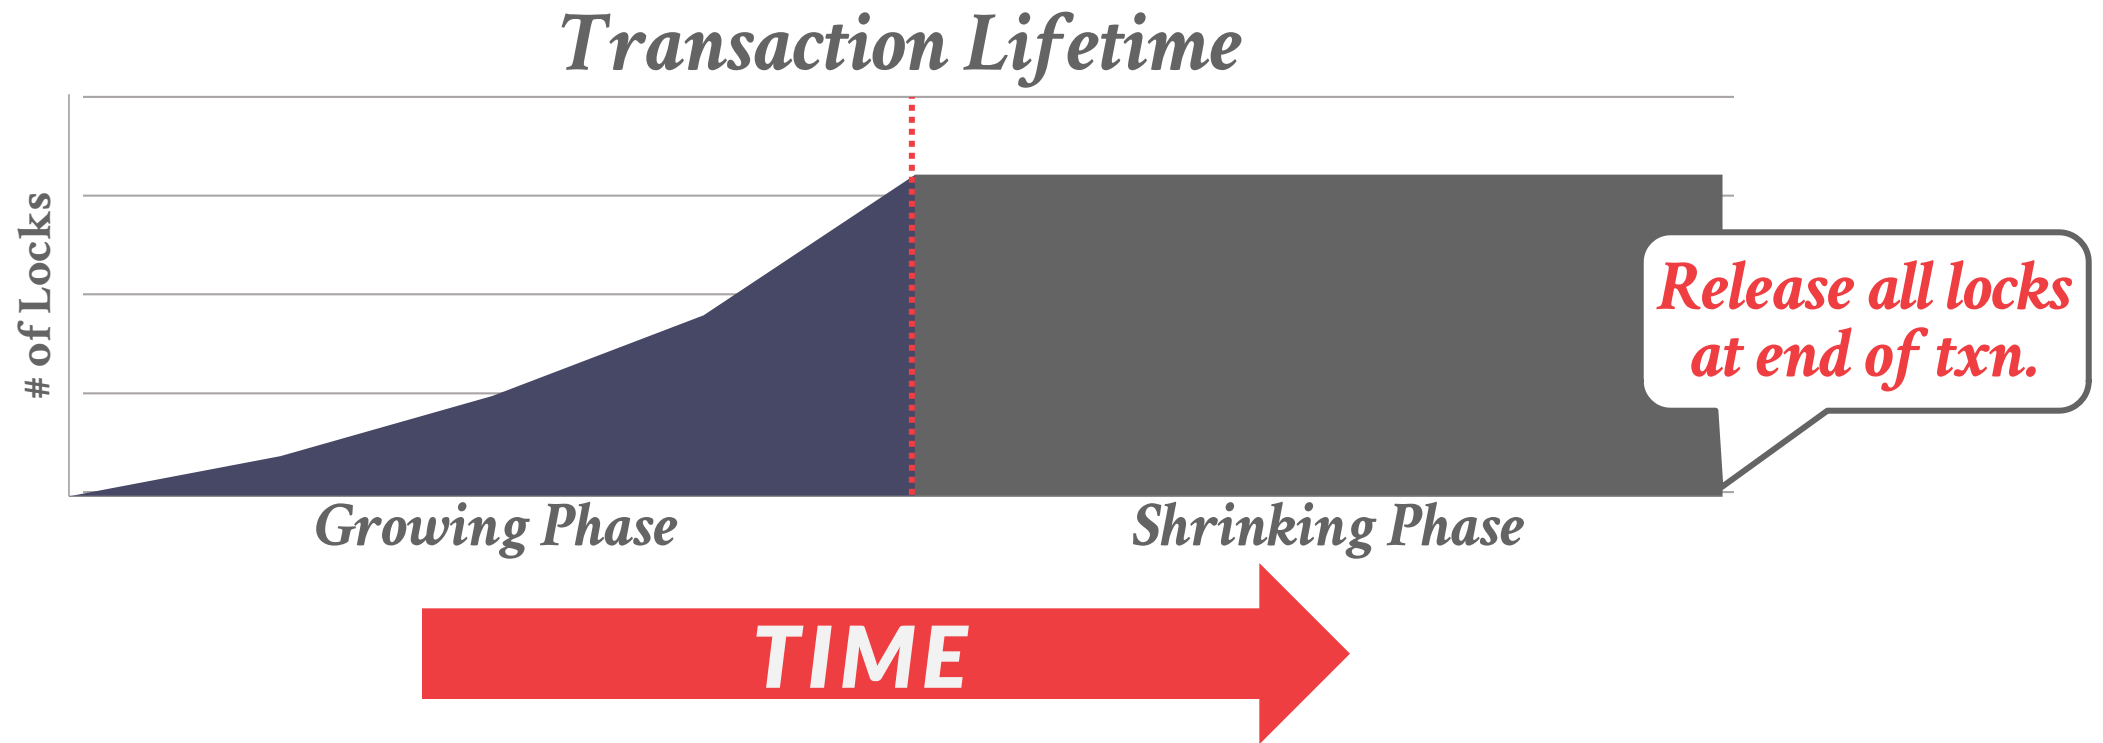


- This is a permissible schedule in 2PL, but the DBMS has to also abort **T₂** when **T₁** aborts.
- Any information about **T₁** cannot be “leaked” to the outside world.

STRONG STRICT TWO PHASE LOCKING

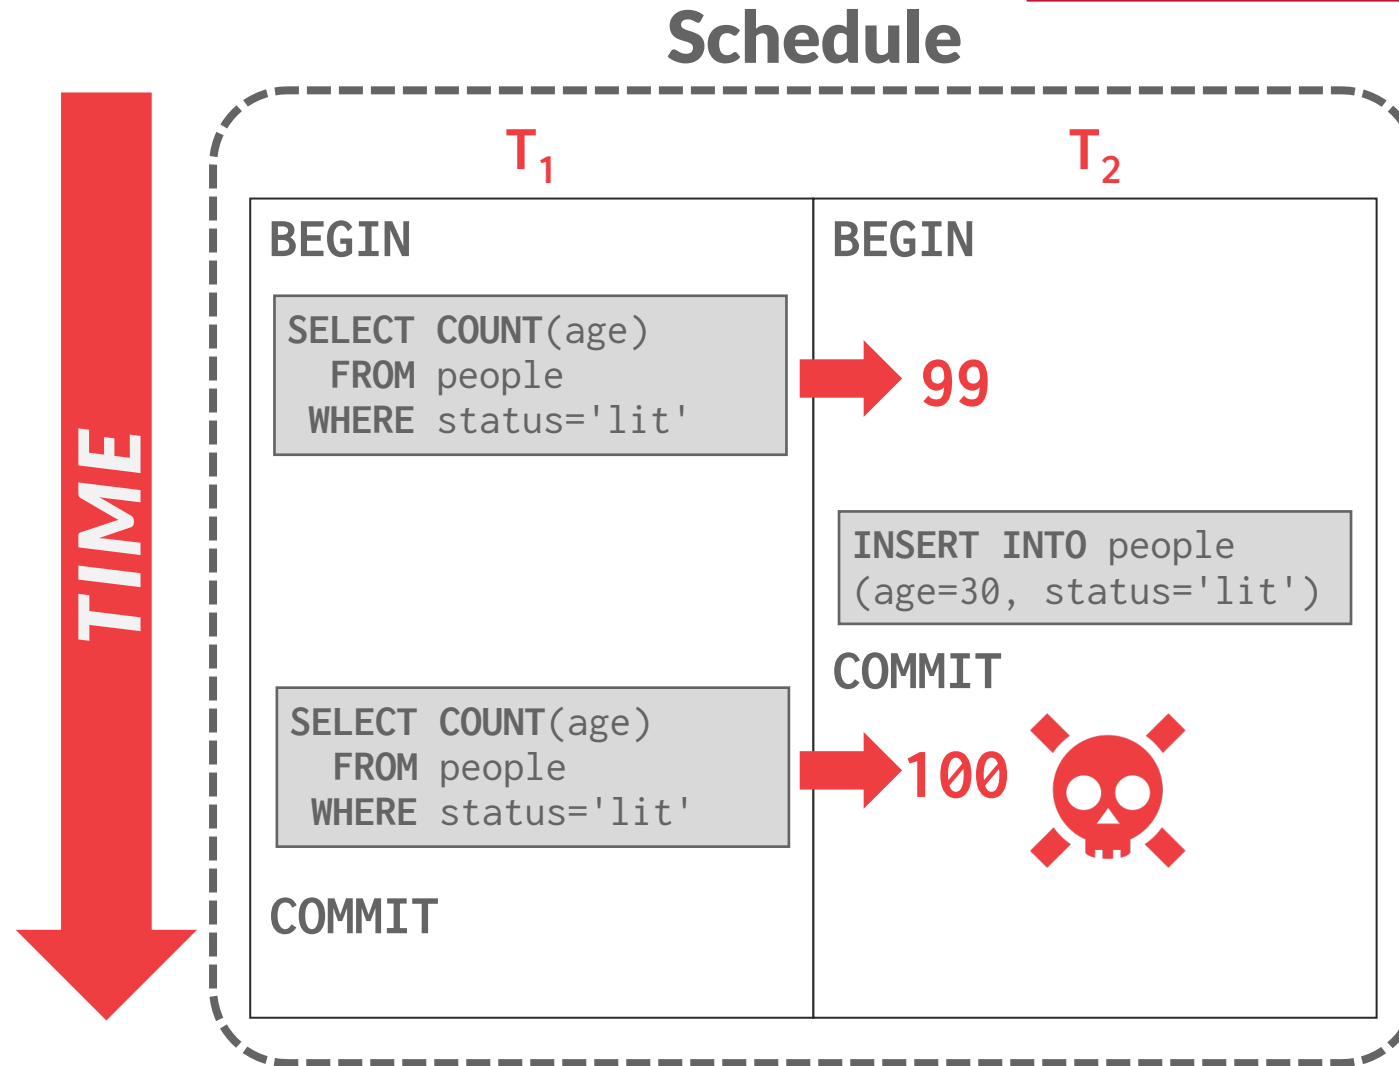
Allows only “conflict serializable” schedules.
ACA – Avoids Cascading Aborts.

- Txn has 2 phases: a growing (acquire lock phase) and a subsequent drop lock phase.
- Can't acquire a lock after the first lock is released.
- Txn holds all locks till the end (abort or commit) and drops them then.



PHANTOMS

Lock key ranges in the B-tree to prevent phantoms, aka. predicate locking.



```
CREATE TABLE people (
  id SERIAL,
  name VARCHAR,
  age INT,
  status VARCHAR
);
```

WEAKER LEVELS OF ISOLATION

- Want to allow various “degrees of consistency” in the same system and concurrent txns.
- Some txns may be ok with lower levels of consistency; e.g., statistics update query.

	<i>Dirty Read</i>	<i>Unrepeatable Read</i>	<i>Phantom</i>
SERIALIZABLE	No	No	No
REPEATABLE READ	No	No	Maybe
READ COMMITTED	No	Maybe	Maybe
READ UNCOMMITTED	Maybe	Maybe	Maybe

ISOLATION LEVELS

- **SERIALIZABLE**: Obtain all locks first; plus index locks, plus strong strict 2PL.
- **REPEATABLE READS**: Same as above, but no index locks.
- **READ COMMITTED**: Same as above, but **S** locks are released immediately.
- **READ UNCOMMITTED**: Same as above but allows dirty reads (no **S** locks).

Part of SQL, and you can explicitly set the isolation levels.

```
SET TRANSACTION ISOLATION LEVEL  
<isolation-level>;
```

```
BEGIN TRANSACTION ISOLATION LEVEL  
<isolation-level>;
```

ISOLATION LEVELS AND ANOMALIES

Table 4. Isolation Types Characterized by Possible Anomalies Allowed.								
Isolation level	P0 Dirty Write	P1 Dirty Read	P4C Cursor Lost Update	P4 Lost Update	P2 Fuzzy Read	P3 Phantom	A5A Read Skew	A5B Write Skew
READ UNCOMMITTED == Degree 1	Not Possible	Possible	Possible	Possible	Possible	Possible	Possible	Possible
READ COMMITTED == Degree 2	Not Possible	Not Possible	Possible	Possible	Possible	Possible	Possible	Possible
Cursor Stability	Not Possible	Not Possible	Not Possible	Sometimes Possible	Sometimes Possible	Possible	Possible	Sometimes Possible
REPEATABLE READ	Not Possible	Not Possible	Not Possible	Not Possible	Not Possible	Possible	Not Possible	Not Possible
Snapshot	Not Possible	Not Possible	Not Possible	Not Possible	Not Possible	Sometimes Possible	Not Possible	Possible
ANSI SQL SERIALIZABLE == Degree 3 == Repeatable Read Date, IBM, Tandem, ...	Not Possible	Not Possible	Not Possible	Not Possible	Not Possible	Not Possible	Not Possible	Not Possible


Hal Berenson, Philip A. Bernstein, Jim Gray, Jim Melton, Elizabeth J. O'Neil, Patrick E. O'Neil:
A Critique of ANSI SQL Isolation Levels. SIGMOD 1995

ISOLATION LEVELS

	<i>Default</i>	<i>Maximum</i>
Action Ingres	SERIALIZABLE	SERIALIZABLE
IBM DB2	CURSOR STABILITY	SERIALIZABLE
CockroachDB	SERIALIZABLE	SERIALIZABLE
Google Spanner	STRICT SERIALIZABLE	STRICT SERIALIZABLE
MSFT SQL Server	READ COMMITTED	SERIALIZABLE
MySQL	REPEATABLE READS	SERIALIZABLE
Oracle	READ COMMITTED	SNAPSHOT ISOLATION
PostgreSQL	READ COMMITTED	SERIALIZABLE
SAP HANA	READ COMMITTED	SERIALIZABLE
VoltDB	SERIALIZABLE	SERIALIZABLE
YugaByte	SNAPSHOT ISOLATION	SERIALIZABLE

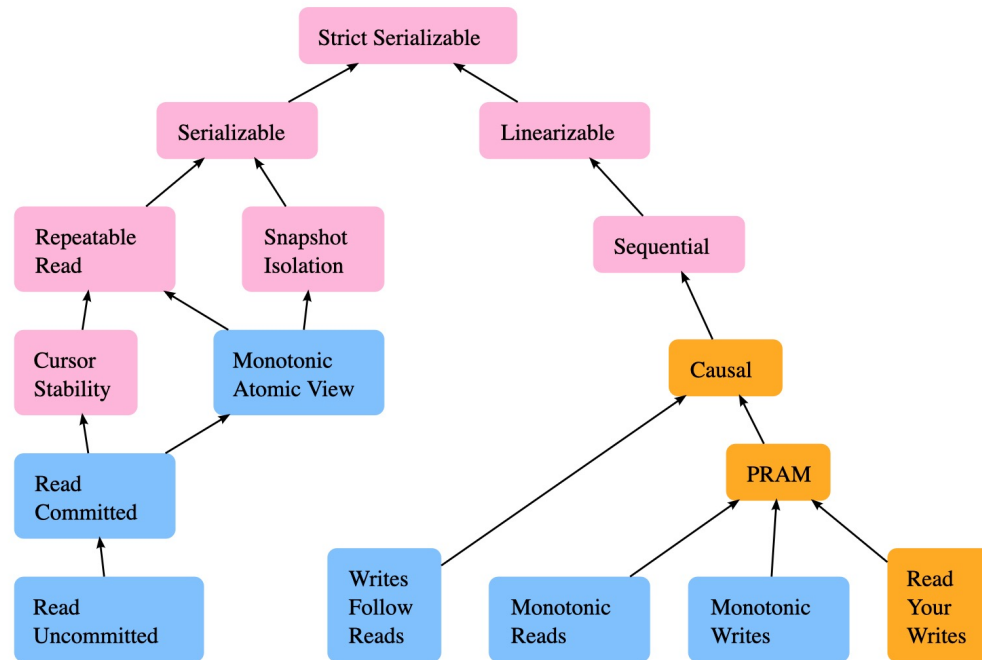
SUMMARY AND OUTLOOK

- This paper directionally set the way concurrency control is implemented in a data platforms (granularity of locking and degrees of consistency), and influenced the SQL standard.
- For the longest time, many database platforms only used pure locking-based protocols for concurrency control.
- But, there were other approaches, including OCC.
- Also, MVCC influences that concurrency protocol. MVCC is about creating versions of data on an update rather than update-in-place and can be used with Locking (or OCC).
 - Revise MVCC from your intro to DB class if you have forgotten it.
- Lot of different way to do concurrency control today with various tradeoffs in the “degree of consistency” and performance.



Consistency Models

This clickable map (adapted from *Bailis, Davidson, Fekete et al* and *Viotti & Vukolic*) shows the relationships between common consistency models for concurrent systems. Arrows show the relationship between consistency models. For instance, strict serializable implies both serializability and linearizability, linearizability implies sequential consistency, and so on. Colors show how available each model is, for a distributed system on an asynchronous network.



Legend

Unavailable

Not available during some types of network failures. Some or all nodes must pause operations in order to ensure safety.

Sticky Available

Available on every non-faulty node, so long as clients only talk to the same servers, instead of switching to new ones.

Total Available

Available on every non-faulty node, even when the network is completely down.

<https://jepsen.io/consistency>