

# ARIES (& Logging)

*Database Systems: The Complete Book*

Ch 17

# Transaction Correctness

- Reliability in database transactions guaranteed by ACID
- A - Atomicity (“Do or Do Not, there is nothing like try”) - usually ensured by logs
- C - Consistency (“Within the framework of law”) - usually ensured by integrity constraints, validations, etc.
- I - Isolation (“Execute in parallel or serially, the result should be same”) - usually ensured by locks
- D - Durability (“once committed, remain committed”) - usually ensured at hardware level

What does it mean for a transaction to be committed?

commit  
returns  
successfully

=

the xact's  
effects  
are visible  
forever

commit  
returns  
successfully

=

the xact's  
effects  
are visible  
forever

commit  
called but  
doesn't  
return


=

the xact's  
effects  
may be  
visible

# Motivation

Committed Transactions.

These should be present when the DB restarts.



Uncommitted Transactions.  
These should leave no trace

- How do we guarantee durability under failures?
- How do aborted transactions get rolled back?
- How do we guarantee atomicity under failures?

**Problem 1:** Providing durability under failures.



## Simplified Model

When a write succeeds, the data is completely written

# Problems

- A crash occurs part-way through the write.
- A crash occurs before buffered data is written.

# Write-Ahead Logging

Before writing to the database,  
first write what you plan to write  
to a log file...

**Log**  
**W ( A : 10 )**



<b>A</b>	8
<b>B</b>	12
<b>C</b>	5
<b>D</b>	18
<b>E</b>	16

# Write-Ahead Logging

Once the log is safely on disk  
you can write the database

**Log**  
W ( A : 10 )



A	<del>8</del> 10
B	12
C	5
D	18
E	16

# Write-Ahead Logging

Log is append-only,  
so writes are always  
efficient

## Log

W(A:10)  
W(C:8)  
W(E:9)



A	<del>8</del> 10
B	12
C	5
D	18
E	16

# Write-Ahead Logging



...allowing random writes  
to be safely batched

## Log

W(A:10)

W(C:8)

W(E:9)

A	<del>8</del> 10
B	12
C	<del>5</del> 8
D	18
E	<del>16</del> 9

**Problem 2:** Providing rollback.

# Single DB Model

**Txn 1**  
→ A = 20  
B = 14  
COMMIT

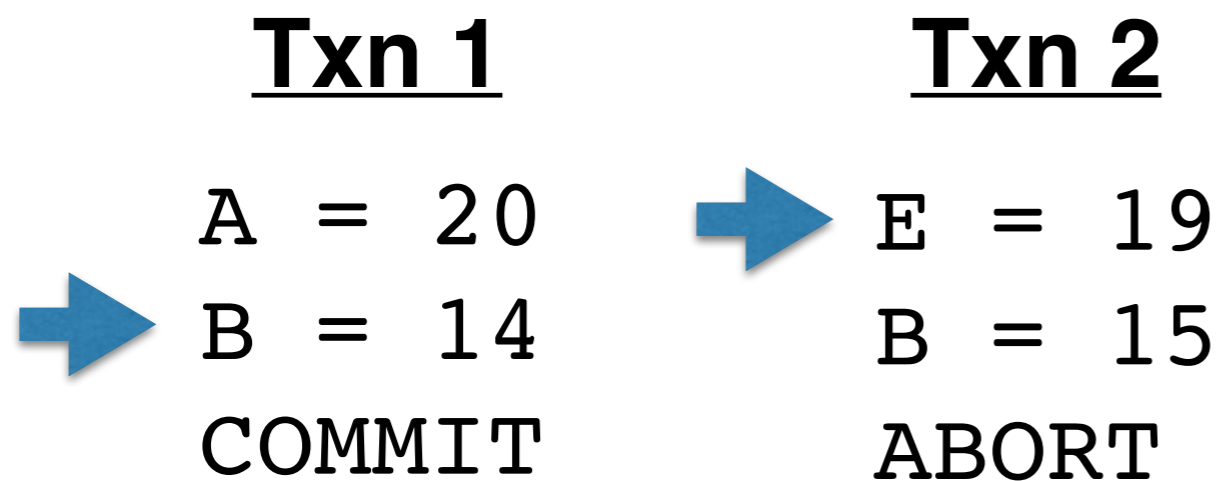
**Txn 2**  
→ E = 19  
B = 15  
ABORT



<b>A</b>	8
<b>B</b>	12
<b>C</b>	5
<b>D</b>	18
<b>E</b>	16

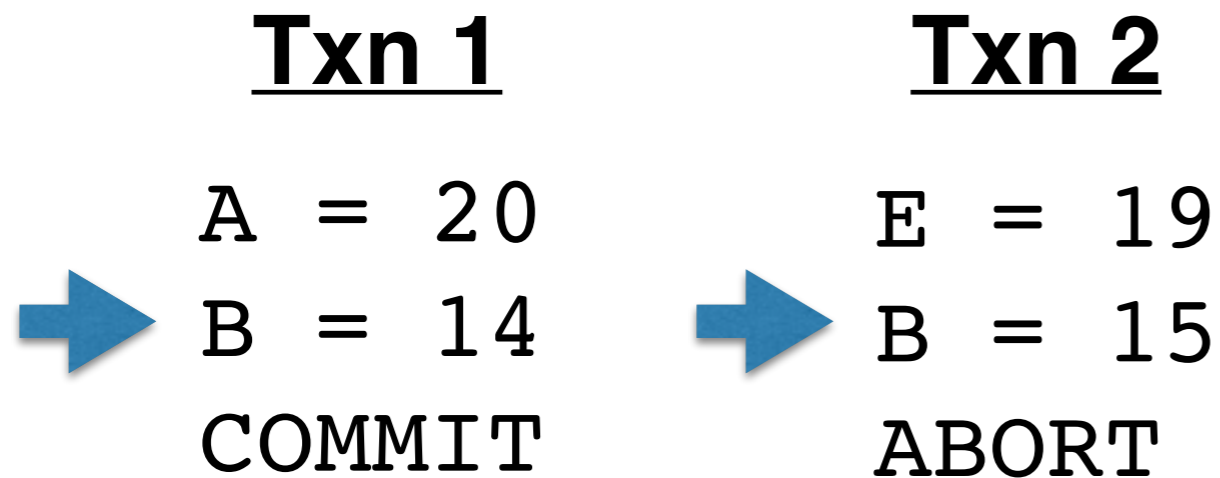


# Single DB Model



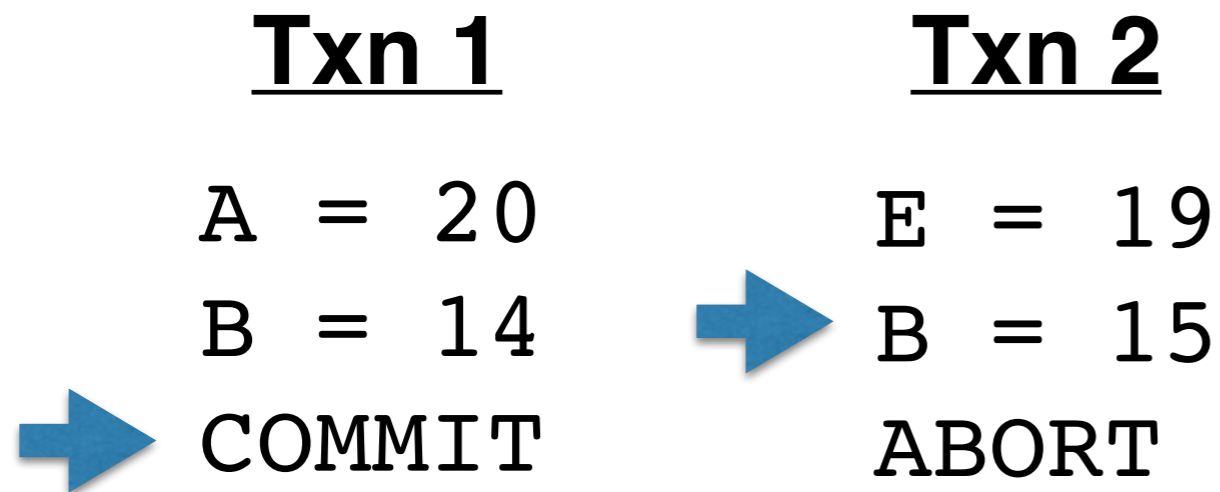
A	<del>8</del> 20
B	12
C	5
D	18
E	16

# Single DB Model



A	<del>8</del> 20
B	<del>1</del> 2
C	5
D	18
E	<del>1</del> 6 19

# Single DB Model



A	<del>8</del> 20
B	<del>12</del> 14
C	5
D	18
E	<del>16</del> 19

# Single DB Model

**Txn 1**  
A = 20  
B = 14  
→ COMMIT

**Txn 2**  
E = 19  
B = 15  
→ ABORT



<b>A</b>	<del>8</del> 20
<b>B</b>	<del>12</del> <del>14</del> 15
<b>C</b>	5
<b>D</b>	18
<b>E</b>	<del>16</del> 19

# Staged DB Model

## Txn 1

➔ A = 20  
B = 14  
COMMIT



## Txn 2

➔ E = 19  
B = 15  
ABORT



A	8
B	12
C	5
D	18
E	16

A	8
B	12
C	5
D	18
E	16

# Staged DB Model

## Txn 1

A = 20

B = 14

→ COMMIT

## Txn 2

E = 19

B = 15

→ ABORT



A	<del>8</del> 20
B	<del>12</del> 14
C	5
D	18
E	16

A	8
B	<del>12</del> 15
C	5
D	18
E	<del>16</del> 19

# Staged DB Model

## Txn 1

A = 20

B = 14

→ COMMIT



## Txn 2

E = 19

B = 15

→ ABORT

A	<del>8</del> 20
B	<del>12</del> 14
C	5
D	18
E	16

Is staging always possible?



- Staging takes up more memory.
- Merging after-the-fact can be harder.
- Merging after-the-fact introduces more latency!

for the single database model

**Problem 2:** Providing rollback.  
^

# UNDO Logging



Store both the “old” and the “new” values of the record being replaced

## Log

W ( A : 8 → 10 )

W ( C : 5 → 8 )

W ( E : 16 → 9 )

A	<del>8</del> 10
B	12
C	<del>5</del> 8
D	18
E	<del>16</del> 9

# UNDO Logging



## Active Xacts

Xact:1, Log: 45

Xact:2, Log: 32

## Log

43 : W ( A : 8 → 10 )

44 : W ( C : 5 → 8 )

45 : W ( E : 16 → 9 )

A	<del>8</del>	10
B	12	
C	<del>5</del>	8
D	18	
E	<del>16</del>	9

# UNDO Logging



A	<del>8</del>	10
B	12	
C	<del>5</del>	8
D	18	
E	<del>16</del>	9

## Active Xacts

Xact: **ABORT**, Log: 45

Xact: 2, Log: 32

## Log

43 : W ( A : 8 → 10 )

44 : W ( C : 5 → 8 )

→ 45 : W ( E : 16 → 9 )

# UNDO Logging



A	<del>8</del> 10
B	12
C	<del>5</del> 8
D	18
E	16

## Active Xacts

Xact: **ABORT**, Log: 45

Xact: 2, Log: 32

## Log

43 : W ( A : 8 → 10 )

44 : W ( C : 5 → 8 )

→ 45 : W ( E : 16 → 9 )

# UNDO Logging



A	<del>8</del> 10
B	12
C	5
D	18
E	16

## Active Xacts

Xact: **ABORT**, Log: 45  
Xact: 2, Log: 32

## Log

43 : W ( A : 8 → 10 )  
44 : W ( C : 5 → 8 )  
45 : W ( E : 16 → 9 )

# UNDO Logging



A	8
B	12
C	5
D	18
E	16

## Active Xacts

Xact: **ABORT**, Log: 45  
Xact: 2, Log: 32

## Log

→ 43 : W ( A : 8 → 10 )  
44 : W ( C : 5 → 8 )  
45 : W ( E : 16 → 9 )



# Log Sequence Number Linked Lists



Transaction Table

XID, LastLSN

LSN, Prev LSN,  
**Prev Image**, ...

LSN, Prev LSN,  
**Prev Image**, ...

ABORT  
[XID]

Log

(necessary for crash recovery)

**Problem 3:** Providing atomicity.

**Goal:** Be able to reconstruct all state at the time of the DB's crash (minus all running xacts)

What state is relevant?

# DB State

**On-Disk  
(or rebuildable)**



**In-Memory  
Only!**

**Active Xacts**

Xact:1, Log: 45

Xact:2, Log: 32

**On-Disk**

**Log**

43 : W ( A : 8 → 10 )

44 : W ( C : 5 → 8 )

45 : W ( E : 16 → 9 )

<b>A</b>	<del>8</del> 10
<b>B</b>	12
<b>C</b>	<del>5</del> 8
<b>D</b>	18
<b>E</b>	<del>16</del> 9

# Rebuilding the Xact Table

Log every COMMIT  
(replay triggers commit process)

Log every ABORT  
(replay triggers abort process)

New message: END  
(replay removes Xact from Xact Table)

What about BEGIN?  
(when does an Xact get added to the Table?)

# Transaction Commit

- Write **Commit** Record to Log
- All Log records up to the transaction's LastLSN are flushed.
- Note that Log Flushes are Sequential, Synchronous Writes to Disk
- Commit() returns.
- Write **End** record to log.

# Simple Transaction Abort (supporting crash recovery)

- Before restoring the old value of a page, write a Compensation Log Record (CLR).
  - Logging continues during UNDO processing.
  - CLR has an extra field: UndoNextLSN
    - Points to the next LSN to undo (the PrevLSN of the record currently being undone)
- CLR's are never UNDO'ed.
  - But might be REDO'ed when repeating history.
  - (Why?)



# Rebuilding the Xact Table

**Optimization:** Write the Xact Table to the log periodically.  
(checkpointing)

# ARIES Crash Recovery

- Start from checkpoint stored in master record.
- **Analysis**: Rebuild the Xact Table
- **Redo**: Replay operations from all live Xacts (even uncommitted ones).
- **Undo**: Revert operations from all uncommitted/aborted Xacts.

