


CSE232: Database System Principles

Concurrency Control

1

Concurrency Control



2

Example:

T1: Read(A)	T2: Read(A)
A ← A+100	A ← A×2
Write(A)	Write(A)
Read(B)	Read(B)
B ← B+100	B ← B×2
Write(B)	Write(B)

Constraint: A=B

3

Serial Schedule A ("good" by definition)

T1	T2	A	B
		25	25
Read(A); A ← A+100			
Write(A);		125	
Read(B); B ← B+100;			
Write(B);			125
	Read(A); A ← A×2;		
	Write(A);	250	
	Read(B); B ← B×2;		
	Write(B);		250
		250	250

4

Serial Schedule B (equally "good")

T1	T2	A	B
		25	25
	Read(A); A ← A×2;		
	Write(A);	50	
	Read(B); B ← B×2;		
	Write(B);		50
Read(A); A ← A+100			
Write(A);		150	
Read(B); B ← B+100;			
Write(B);			150
		150	150

5

Interleaved Schedule C (good because it is equivalent to A)

T1	T2	A	B
		25	25
Read(A); A ← A+100			
Write(A);		125	
	Read(A); A ← A×2;		
	Write(A);	250	
Read(B); B ← B+100;			
Write(B);			125
	Read(B); B ← B×2;		
	Write(B);		250
		250	250

6

Scoping "equivalence" is tricky; for now think that A and C are equivalent because if they start from same initial values they end up with same results

Interleaved Schedule D (bad!)

T1	T2	A	B
		25	25
Read(A); A ← A+100			
Write(A);		125	
	Read(A); A ← A×2;		
	Write(A);	250	
	Read(B); B ← B×2;		
	Write(B);		50
Read(B); B ← B+100;			
Write(B);			150
		250	150

7

Same as Schedule D
but with new T2'

Schedule E (good by "accident")

T1	T2'	A	B
		25	25
Read(A); A ← A+100			
Write(A);		125	
	Read(A); A ← A×1;		
	Write(A);	125	
	Read(B); B ← B×1;		
	Write(B);		25
Read(B); B ← B+100;			
Write(B);			125
		125	125

The accident being the particular semantics of T2' 8

- Want schedules that are "good", I.e., equivalent to serial regardless of
 - initial state and
 - transaction semantics
- Only look at order of read and writes

Example:

SC=r₁(A)w₁(A)r₂(A)w₂(A)r₁(B)w₁(B)r₂(B)w₂(B)

9

Example:

$$SC = r_1(A)w_1(A)r_2(A)w_2(A)r_1(B)w_1(B)r_2(B)w_2(B)$$

$$SC' = r_1(A)w_1(A) \underbrace{r_1(B)w_1(B)r_2(A)w_2(A)}_{T_1} \underbrace{r_2(B)w_2(B)}_{T_2}$$

10

However, for Schedule D:

$$SD = r_1(A)w_1(A)r_2(A)w_2(A) \underbrace{r_2(B)w_2(B)r_1(B)w_1(B)}_{\text{red X}}$$

- as a matter of fact, T_2 must precede T_1 in any equivalent schedule, i.e., $T_2 \rightarrow T_1$
- And vice versa

11


- $T_2 \rightarrow T_1$
- Also, $T_1 \rightarrow T_2$



- ⇒ SD cannot be rearranged into a serial schedule
- ⇒ SD is not "equivalent" to any serial schedule
- ⇒ SD is "bad"

12

Returning to Sc



- no cycles \Rightarrow SC is "equivalent" to a serial schedule (in this case T_1, T_2)

13

Concepts

Transaction: sequence of $r_i(x), w_i(x)$ actions

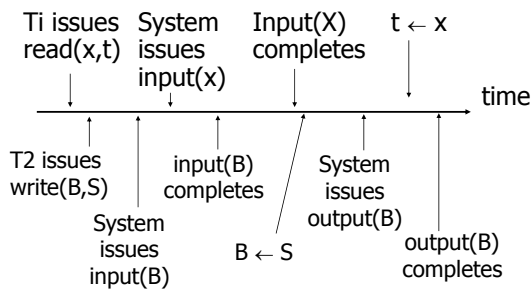
Conflicting actions: $r_1(A) \left\{ \begin{array}{l} w_2(A) \\ w_1(A) \end{array} \right.$
 $w_2(A) \left\{ \begin{array}{l} r_1(A) \\ w_2(A) \end{array} \right.$

Schedule: represents chronological order in which actions are executed

Serial schedule: no interleaving of actions or transactions

14

What about concurrent actions?




15

So net effect is either

- $S = \dots r_1(x) \dots w_2(B) \dots$ or
- $S = \dots w_2(B) \dots r_1(x) \dots$

16

What about conflicting, concurrent actions on same object?



- Assume equivalent to either $r_1(A) w_2(A)$ or $w_2(A) r_1(A)$
- \Rightarrow low level synchronization mechanism
- Assumption called "atomic actions"

17

Definition

S_1, S_2 are conflict equivalent schedules if S_1 can be transformed into S_2 by a series of swaps on non-conflicting actions.

18

Definition

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

19

Precedence graph $P(S)$ (S is schedule)

Nodes: transactions in S

Arcs: $T_i \rightarrow T_j$ whenever

- $p_i(A), q_j(A)$ are actions in S
- $p_i(A) <_S q_j(A)$
- at least one of p_i, q_j is a write

20

Exercise:

- What is $P(S)$ for
 $S = w_3(A) w_2(C) r_1(A) w_1(B) r_1(C) w_2(A) r_4(A) w_4(D)$

- Is S serializable?

21

Lemma

S_1, S_2 conflict equivalent $\Rightarrow P(S_1)=P(S_2)$

Proof:

Assume $P(S_1) \neq P(S_2)$

$\Rightarrow \exists T_i, T_j: T_i \rightarrow T_j$ in S_1 and not in S_2

$\Rightarrow S_1 = \dots p_i(A) \dots q_j(A) \dots$	} p_i, q_j conflict
$S_2 = \dots q_j(A) \dots p_i(A) \dots$	

$\Rightarrow S_1, S_2$ not conflict equivalent

22

Note: $P(S_1)=P(S_2) \not\Rightarrow S_1, S_2$ conflict equivalent

Counter example:

$S_1 = w_1(A) \ r_2(A) \quad w_2(B) \ r_1(B)$

$S_2 = r_2(A) \ w_1(A) \quad r_1(B) \ w_2(B)$

23

Theorem

$P(S_1)$ acyclic $\Leftrightarrow S_1$ conflict serializable

(\Leftarrow) Assume S_1 is conflict serializable

$\Rightarrow \exists S_s: S_s, S_1$ conflict equivalent

$\Rightarrow P(S_s) = P(S_1)$

$\Rightarrow P(S_1)$ acyclic since $P(S_s)$ is acyclic

24

Theorem

$P(S_1)$ acyclic $\iff S_1$ conflict serializable

(\implies) Assume $P(S_1)$ is acyclic

Transform S_1 as follows:


(1) Take T_1 to be transaction with no incident arcs

(2) Move all T_1 actions to the front

$S_1 = \dots q_j(A) \dots p_1(A) \dots$

(3) we now have $S_1 = \langle T_1 \text{ actions} \rangle \langle \dots \text{rest} \dots \rangle$

(4) repeat above steps to serialize rest!



25


How to enforce serializable schedules?

Option 1: run system, recording $P(S)$;
check for $P(S)$ cycles and
declare if execution was good;
or abort transactions as soon
as they generate a cycle

26

How to enforce serializable schedules?

Option 2: prevent $P(S)$ cycles from
occurring




27

A locking protocol

Two new actions:

lock (exclusive): $li(A)$

unlock: $ui(A)$



28

Rule #1: Well-formed transactions

T_i : ... $li(A)$... $pi(A)$... $ui(A)$...

29

Rule #2 Legal scheduler

$S = \dots li(A) \dots ui(A) \dots$
 $\leftarrow \rightarrow$
no $lj(A)$

30

Exercise:

- What schedules are legal?
What transactions are well-formed?

S1 = $l_1(A)l_1(B)r_1(A)w_1(B)l_2(B)u_1(A)u_1(B)$
 $r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

S2 = $l_1(A)r_1(A)w_1(B)u_1(A)u_1(B)$
 $l_2(B)r_2(B)w_2(B)l_3(B)r_3(B)u_3(B)$

S3 = $l_1(A)r_1(A)u_1(A)l_1(B)w_1(B)u_1(B)$
 $l_2(B)r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

31

Exercise:

- What schedules are legal?
What transactions are well-formed?

S1 = $l_1(A)l_1(B)r_1(A)w_1(B)l_2(B)u_1(A)u_1(B)$
 $r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$

S2 = $l_1(A)r_1(A)w_1(B)u_1(A)u_1(B)$
 $l_2(B)r_2(B)w_2(B)l_3(B)r_3(B)u_3(B)$

S3 = $l_1(A)r_1(A)u_1(A)l_1(B)w_1(B)u_1(B)$
 $l_2(B)r_2(B)w_2(B)u_2(B)l_3(B)r_3(B)u_3(B)$


32

Schedule F

T1	T2
$l_1(A); Read(A)$	
$A \leftarrow A + 100; Write(A); u_1(A)$	
	$l_2(A); Read(A)$
	$A \leftarrow A \times 2; Write(A); u_2(A)$
	$l_2(B); Read(B)$
	$B \leftarrow B \times 2; Write(B); u_2(B)$
$l_1(B); Read(B)$	
$B \leftarrow B + 100; Write(B); u_1(B)$	


33

Schedule G



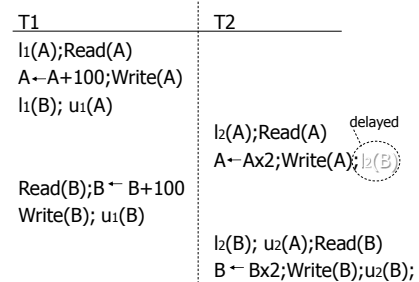
37

Schedule G




38

Schedule G



39

Schedule H (T₂ reversed)



40

- Assume deadlocked transactions are rolled back
 - They have no effect
 - They do not appear in schedule

E.g., Schedule H =
 This space intentionally left blank!

41

Next step:

Show that rules #1,2,3 ⇒ conflict-serializable schedules

42

Conflict rules for $l_i(A), u_i(A)$:

- $l_i(A), l_j(A)$ conflict
- $l_i(A), u_j(A)$ conflict

Note: no conflict $\langle u_i(A), u_j(A) \rangle, \langle l_i(A), r_j(A) \rangle, \dots$

43

Theorem Rules #1,2,3 \Rightarrow conflict
(2PL) serializable
schedule

To help in proof:

Definition $\text{Shrink}(T_i) = \text{SH}(T_i) =$
first unlock action of T_i

44

Lemma

$T_i \rightarrow T_j$ in $S \Rightarrow \text{SH}(T_i) <_S \text{SH}(T_j)$

Proof of lemma:

$T_i \rightarrow T_j$ means that

$S = \dots p_i(A) \dots q_j(A) \dots;$ p, q conflict

By rules 1,2:

$S = \dots p_i(A) \dots u_i(A) \dots l_j(A) \dots q_j(A) \dots$

By rule 3: $\text{SH}(T_i) \quad \text{SH}(T_j)$

So, $\text{SH}(T_i) <_S \text{SH}(T_j)$

45

Theorem Rules #1,2,3 \Rightarrow conflict
(2PL) serializable
schedule

Proof:

(1) Assume P(S) has cycle

$$T_1 \rightarrow T_2 \rightarrow \dots \rightarrow T_n \rightarrow T_1$$

(2) By lemma: $SH(T_1) < SH(T_2) < \dots < SH(T_n)$

(3) Impossible, so P(S) acyclic

(4) \Rightarrow S is conflict serializable

46

- Beyond this simple 2PL protocol, it is all a matter of improving performance and allowing more concurrency....
 - Shared locks
 - Multiple granularity
 - Inserts, deletes and phantoms
 - Other types of C.C. mechanisms

47

Shared locks

So far:

S = ...l₁(A) r₁(A) u₁(A) ... l₂(A) r₂(A) u₂(A) ...

Do not conflict

Instead:

S = ... ls₁(A) r₁(A) ls₂(A) r₂(A) us₁(A) us₂(A)

48

Lock actions

$l-t(A)$: lock A in t mode (t is S or X)

$u-t(A)$: unlock t mode (t is S or X)

Shorthand:

$u_i(A)$: unlock whatever modes

T_i has locked A

49

Rule #1 Well formed transactions

$T_i = \dots l-S_i(A) \dots r_i(A) \dots u_i(A) \dots$

$T_i = \dots l-X_i(A) \dots w_i(A) \dots u_i(A) \dots$

50

- What about transactions that read and write same object?

Option 1: Request exclusive lock

$T_i = \dots l-X_i(A) \dots r_i(A) \dots w_i(A) \dots u_i(A) \dots$

51

- What about transactions that read and write same object?

Option 2: Upgrade

(E.g., need to read, but don't know if will write...)

$T_i = \dots I-S_i(A) \dots r_1(A) \dots I-X_i(A) \dots w_1(A) \dots u(A) \dots$

Think of
 - Get 2nd lock on A, or
 - Drop S, get X lock

52

Rule #2 Legal scheduler

$S = \dots I-S_i(A) \dots \dots u_i(A) \dots$
 ←→
 no $I-X_j(A)$

$S = \dots I-X_i(A) \dots \dots u_i(A) \dots$
 ←→
 no $I-X_j(A)$
 no $I-S_j(A)$

53

A way to summarize Rule #2

Compatibility matrix

Comp		S	X
S		true	false
X		false	false

54

Rule # 3 2PL transactions

No change except for upgrades:

- (I) If upgrade gets more locks
(e.g., $S \rightarrow \{S, X\}$) then no change!
- (II) If upgrade releases read (shared) lock (e.g., $S \rightarrow X$)
 - can be allowed in growing phase

55

Theorem Rules 1,2,3 \Rightarrow Conf.serializable
for S/X locks schedules

Proof: similar to X locks case

56

Lock types beyond S/X


Examples:

- (1) increment lock
- (2) update lock

57

Example (1): increment lock

- Atomic increment action: $IN_i(A)$
 $\{\text{Read}(A); A \leftarrow A+k; \text{Write}(A)\}$
- $IN_i(A), IN_j(A)$ do not conflict!



58

Comp

	S	X	I
S			
X			
I			

59


Comp

	S	X	I
S	T	F	F
X	F	F	F
I	F	F	T

60

Update locks

A common deadlock problem with upgrades:




61

Solution

If T_i wants to read A and knows it may later want to write A, it requests update lock (not shared)

62



-> symmetric table?

63

		New request			
		S	X	U	
Comp	{	S	T	F	T
		X	F	F	F
		U	F	F	F

Lock
already
held in

64

Note: object A may be locked in different modes at the same time...

$S_1 = \dots I-S_1(A) \dots I-S_2(A) \dots I-U_3(A) \dots \left\{ \begin{array}{l} I-S_4(A) \dots? \\ I- \end{array} \right.$

$U_4(A) \dots?$

- To grant a lock in mode t, mode t must be compatible with all currently held locks on object

65

How does locking work in practice?


- Every system is different

But here is one (simplified) way ...


66

Sample Locking System:

- (1) Don't trust transactions to request/release locks
- (2) Hold all locks until transaction commits




67




68

Lock table Conceptually



69


But use hash table:



If object not found in hash table, it is unlocked

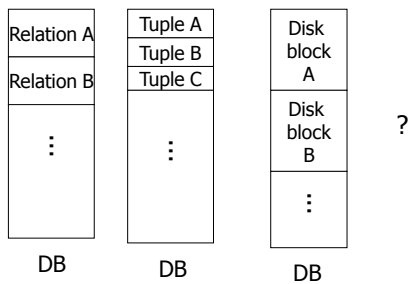
70

Lock info for A - example



71

What are the objects we lock?




72

- Locking works in any case, but should we choose small or large objects?
- If we lock large objects (e.g., Relations)
 - Need few locks
 - Low concurrency
- If we lock small objects (e.g., tuples, fields)
 - Need more locks
 - More concurrency

73


We can have it both ways!!

Ask your janitor to give you the solution...




74

Example



75

Example



76

Multiple granularity

		Requestor				
		IS	IX	S	SIX	X
Holder	IS					
	IX					
	S					
	SIX					
	X					


77

Multiple granularity

		Requestor				
		IS	IX	S	SIX	X
Holder	IS	T	T	T	T	F
	IX	T	T	F	F	F
	S	T	F	T	F	F
	SIX	T	F	F	F	F
	X	F	F	F	F	F


78

Parent locked in	Child can be locked in
IS	
IX	
S	
SIX	
X	



79

Parent locked in	Child can be locked in
IS	IS, S
IX	IS, S, IX, X, SIX
S	[S, IS] not necessary
SIX	X, IX, [SIX]
X	none



80

Rules

- (1) Follow multiple granularity comp function
- (2) Lock root of tree first, any mode
- (3) Node Q can be locked by Ti in S or IS only if parent(Q) locked by Ti in IX or IS
- (4) Node Q can be locked by Ti in X,SIX,IX only if parent(Q) locked by Ti in IX,SIX
- (5) Ti is two-phase
- (6) Ti can unlock node Q only if none of Q's children are locked by Ti

81

Insert + delete operations

A
⋮
Z
α

← Insert

82

Modifications to locking rules:

- (1) Get exclusive lock on A before deleting A
- (2) At insert A operation by T_i , T_i is given exclusive lock on A

83

Still have a problem: **Phantoms**

Example: relation R (E#,name,...)
constraint: E# is key
use tuple locking

R

	E#	Name
o1	55	Smith	
o2	75	Jones	

84


T1: Insert <99,Gore,...> into R
 T2: Insert <99,Bush,...> into R

T1	T2
S1(01)	S2(01)
S1(02)	S2(02)
Check Constraint	Check Constraint
⋮	⋮
Insert o3[99,Gore,..]	Insert o4[99,Bush,..]

85

Solution

- Use multiple granularity tree
- Before insert of node Q, lock parent(Q) in X mode




86

Back to example

T1: Insert<99,Gore>	T2: Insert<99,Bush>
T1	T2
X1(R)	<i>X2(R)</i> <i>delayed</i>
Check constraint	X2(R)
Insert<99,Gore>	Check constraint
U(R)	Oops! e# = 99 already in R!

87

Instead of using R, can use index on R:



88

- This approach can be generalized to multiple indexes...

89


Next:

- Tree-based concurrency control
- Validation concurrency control

90

Example


- all objects accessed through root, following pointers



← can we release A lock if we no longer need A??

91


Idea: traverse like "Monkey Bars"



92

Why does this work?

- Assume all T_i start at root; exclusive lock
- $T_i \rightarrow T_j \Rightarrow T_i$ locks root before T_j



- Actually works if we don't always start at root


93

Rules: tree protocol (exclusive locks)

- (1) First lock by T_i may be on any item
- (2) After that, item Q can be locked by T_i only if $\text{parent}(Q)$ locked by T_i
- (3) Items may be unlocked at any time
- (4) After T_i unlocks Q , it cannot relock Q

94

- Tree-like protocols are used typically for B-tree concurrency control



E.g., during insert, do not release parent lock, until you are certain child does not have to split

95

Validation-based Concurrency Control

Transactions have 3 phases:

- (1) Read
 - all DB values read
 - writes to temporary storage
 - no locking
- (2) Validate
 - check if schedule so far is serializable
- (3) Write
 - if validate ok, write to DB

96

Key idea

- Make validation atomic
- If T_1, T_2, T_3, \dots is validation order, then resulting schedule will be conflict equivalent to $S_s = T_1 T_2 T_3 \dots$


97

To implement validation, system keeps two sets:

- FIN = transactions that have finished phase 3 (and are all done)
- VAL = transactions that have successfully finished phase 2 (validation)

98

Example of what validation must prevent:




99

Example of what validation must ^{allow} prevent:

$$RS(T_2) = \{B\} \quad RS(T_3) = \{A, B\} \neq \phi$$

$$WS(T_2) = \{B, D\} \quad WS(T_3) = \{C\}$$




100

Another thing validation must prevent:

$$RS(T_2) = \{A\} \quad RS(T_3) = \{A, B\}$$

$$WS(T_2) = \{D, E\} \quad WS(T_3) = \{C, D\}$$




101

Another thing validation must ^{allow} prevent:

$$RS(T_2) = \{A\} \quad RS(T_3) = \{A, B\}$$

$$WS(T_2) = \{D, E\} \quad WS(T_3) = \{C, D\}$$



102

Validation rules for T_j:

(1) When T_j starts phase 1:

ignore(T_j) ← FIN

(2) at T_j Validation:

if check (T_j) then

[VAL ← VAL U {T_j};

do write phase;

FIN ← FIN U {T_j}]

103

Check (T_j):

For T_i ∈ VAL - IGNORE (T_j) DO

IF [WS(T_i) ∩ RS(T_j) ≠ ∅ OR

T_i ∉ FIN] THEN RETURN false;

RETURN true;

Is this check too restrictive ?

104

Improving Check(T_i)

For T_i ∈ VAL - IGNORE (T_j) DO

IF [WS(T_i) ∩ RS(T_j) ≠ ∅ OR


(T_i ∉ FIN AND WS(T_i) ∩ WS(T_j) ≠ ∅)]

THEN RETURN false;

RETURN true;

105

Exercise:



106

Validation (also called optimistic concurrency control) is useful in some cases:

- Conflicts rare
- System resources plentiful
- Have real time constraints

107

Summary

Have studied C.C. mechanisms used in practice

- 2 PL
- Multiple granularity
- Tree (index) protocols
- Validation

108
