

Chapter 4 Parallel Transactions and Execution Control

Chapter 4.1 Concept of Transaction (see chapter 1.1)

4.1.1 Syntax: **begintrans** (BOT)


a1;

a2;

...

ak

endtrans (EOT)



Transaction $t_j = (a_{j1}; a_{j2}; \dots; a_{jk})$

is a sequential program

$a_{j|}$ are actions:

$r(x)$: read Objekt x

$w(x)$: write Objekt x

Concept of Object: single tuple, attribute in tuple, page, relation, Index-page, Index-entry, SQL-Statement (Debit-Credit transaction, see Chapt. 1.1)

- *Atomicity* (Atomarität): all changes of a transaction are performed completely or not at all.
- *Consistency* (Konsistenz): a transaction transforms a consistent database again into a consistent (may be the same) database, i. e., all integrity constraints of the databases are respected by the program.
- *Isolation*: changes of running transactions are not visible for other transactions. A transaction behaves as if there were no other.
- *Durability* (Permanenz): after the successful completion (Commit) of a transaction all updates are permanent, i. e., they survive later errors and failures.

Exchange of information between DB and AS:

1. $r(x) \rightarrow X$ (main store var.)
2. $X \rightarrow w(x)$

4.1.2 Interleaving and Anomalies

- 4 Anomalies:
- Lost Update
 - Dirty Read
 - Unrepeatable Read
 - Phantom

Example Lost Update

T ₁	T ₂
BOT()	
$r(x) \rightarrow X$	
$X := X + 1000$	
	BOT()
	$r(x) \rightarrow Y$
	$Y := X + 5$
$X \rightarrow w(x)$	
Commit()	
	$Y \rightarrow w(x)$
	Commit()

Example Dirty Read:

T ₁	T ₂
BOT()	
$r(x) \rightarrow X$	
$X := X + 1000$	
$X \rightarrow w(x)$	
	BOT()
Abort()	$r(x) \rightarrow Y$
	$Y := Y + 5$

Example Unrepeatable Read (inconsistent retrieval)

T ₁	T ₂
BOT()	BOT()
	$r(y) \rightarrow Y$
$r(y) \rightarrow Y$	
$Y := Y - 1000$	
$Y \rightarrow w(y)$	
$r(x) \rightarrow X$	
$X := X + 1000$	
$X \rightarrow w(x)$	
Commit()	
	$r(x) \rightarrow X$
	balance = X - Y

Example Phantom

T_1	T_2
	BOT()
BOT()	
$r(y) \rightarrow Y$	
$r(x) \rightarrow X$	
	$r(y) \rightarrow Y$
	$r(\text{sum}) \rightarrow \text{Sum}$
	$\text{Sum} := \text{Sum} - Y$
	delete(y)
	$\text{Sum} \rightarrow w(\text{sum})$
$r(\text{sum}) \rightarrow Z$	
Commit()	

Degrees of Isolation (consistency levels)

- Is. Dg. 3: Repeatable Read (no anomalies)
long read and write locks
- Is. Dg. 2: Consistent Read (not repeatable)
short read locks, long write locks
- Is. Dg. 1: Browse: no Lost Update
but dirty reads
no readlocks, but long write locks
- Is. Dg. 0: Chaos: Lost Updates, but "physically" correct
data structures
only short write locks

Chapter 4.2 Serializability

Def.: - **Schedule for** $t_1 = (a_{11}; a_{12}, \dots, a_{1k_1})$
(History) \dots
 $t_e = (a_{e1}; \dots; a_{ek_e})$

H = Sequence of a_j , so that sequence of actions within t_i is preserved (large number of histories)

- **complete Schedule**
for every t_i abort $_i$ (ab_i)
or commit $_i$ (c_i)
is contained in H
- t_i, t_j are **interleaved** in H :
 a_{i1} **before** a_{jm} **before** a_{in}
we write $a_{i1} <_H a_{jm} <_H a_{in}$
i.e. $\dots a_{i1} \dots a_{jm} \dots a_{in} \dots = H$
- H is **serial** if no transactions in H are interleaved

Note: for I Transactions there are $!!$ serial schedules,
many more schedules

Example: $t_1 = r_1(x) w_1(x) c_1$
 $t_2 = r_2(x) w_2(x) c_2$

$H = r_1 r_2 w_1 w_2 c_1 c_2$

Problem: serial schedules not useful for practical purposes

- slow I/O
- slow process communication
- slow human Interactions

⇒ which interleaved schedules are acceptable?
 Equivalence to serial schedule!

Def.: State equivalence (final-state equivalence):
 schedule H is fs-equivalent to serial schedule H_s ,
 if both lead to the same DB-state (semantic
 equivalence) (too strong)

Def.: $a_i(x)$ and $a_j(x)$ are in conflict if
 $a_i(x) = w(x)$ or $\overline{a_j(x) = w(x)}$
 formally: $\text{conf}(a_i(x), a_j(x))$

Conflict relation conf_H for schedule H:
 $\text{conf}_H = \{(a,b) \mid a,b \text{ are in conflict} \wedge a <_H b\}$
 we also write $\text{conf}_H(a,b)$ or $a \text{ conf}_H b$.

Note: $\overline{\text{conf}}$ is symmetric
 conf_H is not symmetric

Example: $H = r_1(x) r_2(x) w_1(x) w_2(x) c_1 c_2$
 $\text{conf}_H = \{(r_1, w_2)$
 (r_2, w_1)
 $(w_1, w_2)\}$

Def.: H_1, H_2 are called **conflict equivalent**,
 - if H_1, H_2 have the same actions and
 - $\text{conf}_{H_1} = \text{conf}_{H_2}$

Def.: complete Schedule H is called **conflict serializable** if there exists H' :
 - H, H' have the same set of actions
 - $\text{conf}_H = \text{conf}_{H'}$
 - H' is serial

Lemma: conflict serializability
 ⇒ fs-equivalence (final state equivalence)

Note: conflict serializability is the synchronization
 method of almost all commercial DBS,
 achievable by synchronization protocol =
 Concurrency Control Protocol

Def.: Interleaved (parallel) execution $\{t_1, \dots, t_n\}$ of
 Trans. t_1, \dots, t_n is **correct (consistent)** if:
 $\exists i_1, \dots, i_n [\{t_1, t_2, \dots, t_n\} \equiv (ti_1; ti_2; \dots; ti_n)]$

many executions (Histories) are possible
 specific execution ~ Schedule,
 admissible are only conflict-serializable schedules.

Chapter 4.3 Synchronization

Def: Let H_s be serial schedule

$t_1 \prec t_2$ in H_s , if all actions of t_1 are before all actions of t_2 .

\prec is linear order. Denote $H_s = H_{\prec}$.

Note: For a conflict serializable schedule H only conflicting actions are relevant.

(1) as soon as $\exists a_{i,m} \exists a_{j,n} : \text{conf}(a_{i,m}, a_{j,n})$ in H we must have

(2) $\forall a_{i,m} \forall a_{j,n} : \overline{\text{conf}}(a_{i,m}, a_{j,n}) \supset a_{i,m}$ **before** $a_{j,n}$
resp. ... $\overline{\text{conf}}(\dots) \supset \text{conf}(\dots)$

How can (1) \Rightarrow (2) be achieved?

Strategies for Serialization

1. Serialization according to the sequence of transaction calls
2. Parallelism (interleaving) only between conflict free transactions, i.e.
 - all objects needed by t must be known in advance
 - check conflicts with running t'

- start t only, if there are no conflicts
 \Rightarrow all locks of t must be placed before starting t , i.e. $t \prec t'$;
no deadlock problem

3. Dynamic acquisition of lock, as soon as t wants to process object o .
 - 3.1 lock granted $\Rightarrow t$ continues
 - 3.2 lock not granted, i.e. o is locked by t' with incompatible lock.
 $\Rightarrow t$ must wait for t'

Deadlock - Problem:

Def.: Follow - Relation \rightarrow for transactions in H:

$t_j \rightarrow t_i$ if (1)

i. e. $\exists a_i m \exists a_j n : \text{conf}(a_i m, a_j n)$

and t_j follows t_i

Lemma: H is conflict serializable, if \rightarrow is acyclic

Proof: Construct topological order \leftarrow^* corresponding

$t_o \rightarrow H$

\leftarrow^* is conflict equivalent to H

Basic Idea: Object locks

Compute $\text{conf}(a_i l, a_j k)$ and simultaneously

\rightarrow incrementally with the aid of locks on objects

reason: conflicts arise because of access to objects

$r_i(o); r_j(o)$: no conflict
compatible locks

$r_i(o); r_j(o)$: t_i places R-lock on o

$w_j(o); r_j(o)$: t_i places X-lock on o

$w_i(o); w_j(o)$: t_i places X-lock on o

update \rightarrow and check for acyclicity $t_j \rightarrow t_i$

means t_j waits for t_i

Lock Compatibility

	R	X
R		
X		

2-Phase Lock Protocol

Example:

1. t_1, t_2, t_3 request R-lock on o1
2. t_4 requests X-lock on o2
3. t_5 requests X-lock on o1
4. t_6 requests R-lock on o2
5. t_7 requests R-lock on o1

Passing of transactions?

Starvation?

Note: Locks alone do not guarantee serialization
(consistency)

$t_1 : A := 1; \quad t_2 : B := 1;$
 $B := 0 \quad A := 0$

Special interleaving (parallelism) with sequence w.r. to time:

$\{t_1, t_2\}^* \equiv$	$A_1 := 1;$	(t_1, A, X)
	$B_2 := 1;$	$A := 1$
	$B_1 := 0;$	release (t_1, A, X)
	$A_2 := 0$.
		.
		.


$\Rightarrow A = 0 \wedge B = 0$

but : $(t_1; t_2) \Rightarrow B = 1 \wedge A = 0$
 $(t_2; t_1) \Rightarrow A = 1 \wedge B = 0$


$(t_1; t_2) \not\equiv \{t_1, t_2\}^* \not\equiv (t_2; t_1)$

i.e. $\{t_1, t_2\}^*$ is not conflict-serializable (no consistent parallel execution)

2 Phases w.r. to Locks:




Strict 2-Ph. Lock-Protocol:



Theorem: 2-phase lock protocol results in conflict-serializable schedules
(consistent execution control)

Proof:



$\tau_s(t)$: time, when last lock is granted

$\tau_s(t)$: t has locked all needed objects, may continue without inhibition, no further waiting!

$\tau_s(t_1) < \tau_s(t_2)$: t_1 in serialization before t_2 : $t_2 \rightarrow t_1$
 \rightarrow

Assumption: t_1 wants to lock $\{o_1, \dots, o_n\} =: O$
 t_2 wants to lock $\{p_1, \dots, p_m\} =: P$

Case 1: $O \cap P = \emptyset$
 $\{t_1, t_2\} \equiv (t_1; t_2) \equiv (t_2; t_1)$
 no conflicts!

Case 2: $O \cap P \neq \emptyset$
 let $O_j = P_k \in O \cap P$ locked simultaneously by t_1, t_2 ,
 only possible for R-locks $\Rightarrow o_j$ only read by t_1, t_2 ,
 $(t_1; t_2) \equiv (t_2; t_1)$


Case 3: $O \cap P \neq \emptyset$
 let $o_j = p_k$ with conflicting locks: "who locks first, finishes first"
 e.g. $t_2 \rightarrow t_1$ w.r. to o_j

Let $O_K \subseteq O \cap P$; O_K set of objects requested by t_1, t_2 with conflicting locks:


a) t_1 locks all of O_K first :
 $t_2 \rightarrow t_1$ and $t_1 \ll t_2$

b) t_2 locks all of O_K first:
 $t_1 \rightarrow t_2$ and $t_2 \ll t_1$

c) t_1 locks $O_1 \subset O_K$ and $O_1 \neq \emptyset$
 \neq
 t_2 locks $O_2 \subset O_K$ and $O_2 \neq \emptyset$
 \neq
 then $t_1 \rightarrow t_2$ to free O_2
 $t_2 \rightarrow t_1$ to free O_1



d) general case: $O_{K_i,j}$ = set of objects needed by t_i and t_j with conflicts.




Note: 2-Phase protocol determines sequence if conflicts arise

locally: $t_1 \rightarrow t_2$ or $t_2 \rightarrow t_1$

globally: wait graph and detection of cycles
= deadlock

Note: Sink of the wait graph and isolated nodes
= transactions in process state "ready"
(rechenwillig)

Example:



1. t_1 updates o to V_n^1 , frees o
2. t_2, t_3 read V_n^1
3. t_4 reads V_n^1 , changes V_n^1 to V_n^2
4. t_1 fails, is aborted,
e.g. because of diskfailure, integrity violation
 $\Rightarrow t_2, t_3, t_4$ must be aborted, cascading!!

Theorem: Strict 2-Phase locking protocol avoids
cascading when transactions are aborted.

Proof: reason for cascading:
versions of objects were read, which were not yet
committed, since locks (and objects) were freed
before commit.

Example with strict protocol?