

CARNEGIE MELLON UNIVERSITY  
DEPARTMENT OF COMPUTER SCIENCE  
15-445/645 – DATABASE SYSTEMS (FALL 2017)  
PROF. ANDY PAVLO

Homework 5 (by Joy Arulraj)  
Due: **Monday Nov 13, 2017 @ 11:59pm**

**IMPORTANT:**

- **Upload this PDF** with your answers to **Gradescope by 11:59pm on Monday Nov 13, 2017.**
- **Plagiarism:** Homework may be discussed with other students, but all homework is to be completed **individually.**
- **You have to use this PDF for all of your answers.**

For your information:

- Graded out of **100** points; **4** questions total

*Revision : 2017/10/30 22:21*

Question	Points	Score
Serializability and 2PL	20	
Deadlock Detection and Prevention	30	
Hierarchical Locking - A Blogging Website	30	
B+ tree Locking	20	
Total:	100	

**Question 1: Serializability and 2PL.....[20 points]**

(a) Yes/No questions:

- i. **[2 points]** Schedules under Strict 2PL could have cascading aborts.  
 Yes    No
- ii. **[2 points]** A conflict serializable schedule need not always be view serializable.  
 Yes    No
- iii. **[2 points]** Schedules under Strict 2PL could have deadlocks.  
 Yes    No
- iv. **[2 points]** There could be schedules under 2PL that are not serializable.  
 Yes    No
- v. **[2 points]** If a precedence graph associated with a schedule has cycles, then it is not conflict serializable.  
 Yes    No

(b) Serializability:

Consider the schedule given below in Table 1. R(·) and W(·) stand for ‘Read’ and ‘Write’, respectively.

time	$t_1$	$t_2$	$t_3$	$t_4$	$t_5$	$t_6$	$t_7$	$t_8$	$t_9$	$t_{10}$
$T_1$	R(A)	W(A)	W(E)	R(B)		W(B)	R(C)	W(C)	R(D)	W(D)
$T_2$			R(A)		W(A)	R(D)	W(D)	W(E)		W(A)
$T_3$				R(C)	W(C)			R(F)	W(F)	

Table 1: A schedule with 3 transactions

- i. **[1 point]** Is this schedule serial?  
 Yes    No
- ii. **[3 points]** Give the dependency graph of this schedule. List each edge in the dependency graph like this:  $T_x \rightarrow T_y$  because of  $Z$  (i.e., Transaction  $T_y$  reads/writes  $Z$  which was last written by  $T_x$ ). Order the edges in ascending order with respect to  $x$ .
- iii. **[1 point]** Is this schedule conflict serializable?  
 Yes    No
- iv. **[3 points]** If you answer “yes” to (iii), provide the equivalent serial schedule. If you answer “no”, briefly explain why.
- v. **[2 points]** Is this schedule possible under 2PL?  
 Yes    No

**Question 2: Deadlock Detection and Prevention.....[30 points]**

(a) Deadlock Detection:

Consider the following lock requests in Table 2. And note that

- $S(\cdot)$  and  $X(\cdot)$  stand for ‘shared lock’ and ‘exclusive lock’, respectively.
- $T_1$ ,  $T_2$ , and  $T_3$  represent three transactions.
- $LM$  stands for ‘lock manager’.
- Transactions will never release a granted lock.

time	$t_1$	$t_2$	$t_3$	$t_4$	$t_5$	$t_6$	$t_7$
$T_1$	S(A)				S(B)	S(C)	
$T_2$		S(B)		X(C)			X(B)
$T_3$			X(A)				
$LM$	g						

Table 2: Lock requests of 3 transactions

- [3 points]** For the lock requests in Table 2, determine which lock will be granted or blocked by the lock manager. Please write ‘g’ in the LM row to indicate the lock is granted and ‘b’ to indicate the lock is blocked. For example, in the table, the first lock (S(D) at time  $t_1$ ) is marked as granted.
- [4 points]** Give the wait-for graph for the lock requests in Table 2. List each edge in the graph like this:  $T_x \rightarrow T_y$  because of  $Z$  (i.e.,  $T_x$  is waiting for  $T_y$  to release its lock on resource  $Z$ ). Order the edges in ascending order with respect to  $x$ .
- [3 points]** Determine whether there exists a deadlock in the lock requests in Table 2, and briefly explain why.

(b) Deadlock Prevention:

Consider the following lock requests in Table 3. Again,

- $S(\cdot)$  and  $X(\cdot)$  stand for ‘shared lock’ and ‘exclusive lock’, respectively.
  - $T_1$ ,  $T_2$ ,  $T_3$ , and  $T_4$  represent four transactions.
  - $LM$  represents a ‘lock manager’.
  - Transactions will never release a granted lock.
- [3 points]** For the lock requests in Table 3, determine which lock request will be granted, blocked or aborted by the lock manager ( $LM$ ), if it has no deadlock prevention policy. Please write ‘g’ for grant, ‘b’ for block and ‘a’ for abort. Again, example is given in the first column.
  - [4 points]** Give the wait-for graph for the lock requests in Table 3. List each edge in the graph like this:  $T_x \rightarrow T_y$  because of  $Z$  (i.e.,  $T_x$  is waiting for  $T_y$  to release its lock on resource  $Z$ ). Order the edges in ascending order with respect to  $x$ .

time	$t_1$	$t_2$	$t_3$	$t_4$	$t_5$	$t_6$	$t_7$	$t_8$
$T_1$	S(A)		S(B)					
$T_2$		X(B)						X(D)
$T_3$				S(C)	X(D)		X(A)	
$T_4$						X(C)		
$LM$	g							

Table 3: Lock requests of 4 transactions

- iii. **[3 points]** Determine whether there exists a deadlock in the lock requests in Table 3, and briefly explain why.
- iv. **[5 points]** To prevent deadlock, we use the lock manager ( $LM$ ) that adopts the Wait-Die policy. We assume that in terms of priority:  $T_1 > T_2 > T_3 > T_4$ . Here,  $T_1 > T_2$  because  $T_1$  is older than  $T_2$  (i.e., older transactions have higher priority). Determine which lock request will be granted ('g'), blocked ('b') or aborted ('a'). Follow the same format as the previous question.
- v. **[5 points]** Now we use the lock manager ( $LM$ ) that adopts the Wound-Wait policy. We assume that in terms of priority:  $T_1 > T_2 > T_3 > T_4$ . Here,  $T_1 > T_2$  because  $T_1$  is older than  $T_2$  (i.e., older transactions have higher priority). Determine which lock request will be granted ('g'), blocked ('b') or aborted ('a'). Follow the same format as the previous question.

### Question 3: Hierarchical Locking - A Blogging Website . . . . . [30 points]

Consider a Database (D) consisting of two tables, Users (U) and Posts (P). Specifically,

- Users(uid, first\_name, last\_name), spans 300 pages, namely  $U_1$  to  $U_{300}$
- Posts(pid, uid, title, body), spans 600 pages, namely  $P_1$  to  $P_{600}$

Further, **each page contains 100 records**, and we use the notation  $U_3 : 20$  to represent the 20<sup>th</sup> record on the third page of the Users table. Similarly,  $P_5 : 10$  represents the 10<sup>th</sup> record on the fifth page of the Posts table.

We use Multiple-granularity locking, with **S, X, IS, IX** and **SIX** locks, and **four levels of granularity**: (1) *database-level (D)*, (2) *table-level (U, P)*, (3) *page-level ( $U_1 - U_{300}$ ,  $P_1 - P_{600}$ )*, (4) *record-level ( $U_1 : 1 - U_{300} : 100$ ,  $P_1 : 1 - P_{600} : 100$ )*.

For each of the following operations on the database, please determine the sequence of lock requests that should be generated by a transaction that wants to efficiently carry out these operations by maximizing concurrency.

Please follow the format of the examples listed below:

- write **“IS(D)”** for a request of **database-level IS lock**
  - write **“X( $P_2 : 30$ )”** for a request of **record-level X lock for the 30<sup>th</sup> record on the second page of the Posts table**
  - write **“S( $P_2 : 30 - P_3 : 100$ )”** for a request of **record-level S lock from the 30<sup>th</sup> record on the second page of the Posts table to the 100<sup>th</sup> record on the third page of the Posts table.**
- (a) [6 points] Fetch the 10<sup>th</sup> record on page  $P_{100}$ .
  - (b) [6 points] In remembrance of our TA (RIP Christopher “Inf” Wallace), set the last\_name of all the users to be “Inf”.
  - (c) [6 points] Scan all the records on pages  $P_1$  through  $P_{20}$ , and modify the record  $P_5 : 10$ .
  - (d) [6 points] Order all the posts by their title.
  - (e) [6 points] Delete ALL the records from ALL tables.

**Question 4: B+ tree Locking ..... [20 points]**

Consider the following B+ tree:




Figure 1: B+ tree locking

To lock this B+ tree, we would like to use the **Bayer-Schkolnick** algorithm. **Important:** we use the version as presented in the lecture, which **does not** use lock upgrade.

For each of the following transactions, give the sequence of lock/unlock requests. For example, please write  $S(A)$  for a request of shared lock on node A,  $X(B)$  for a request of exclusive lock on node B and  $U(C)$  for a request of unlock node C.

**Important notes:**

- Each of the following transactions is applied on the *original tree*, i.e., ignore any tree changes described in other problems.
- For simplicity, *ignore* the changes on the pointers between leaves.

- [5 points] Search for data entry “15\*”
- [5 points] Delete data entry “28\*”
- [5 points] Insert data entry “9\*”
- [5 points] Insert data entry “45\*”