

CARNEGIE MELLON UNIVERSITY  
COMPUTER SCIENCE DEPARTMENT  
15-445/645 – DATABASE SYSTEMS (SPRING 2024)  
PROF. JIGNESH PATEL

Homework #4 (by Amy Cheng, Ritu Pathak, Shivang Dalal)  
Due: **Saturday April 6, 2024 @ 11:59pm**

**IMPORTANT:**

- Enter all of your answers into **Gradescope by 11:59pm on Saturday April 6, 2024.**
- **Plagiarism:** Homework may be discussed with other students, but all homework is to be completed **individually**.

For your information:

- Graded out of **100** points; **4** questions total
- Rough time estimate:  $\approx 2 - 4$  hours (0.5 - 1 hours for each question)

*Revision : 2024/04/02 09:42*

Question	Points	Score
Query Execution, Planning, and Optimization	24	
Serializability and 2PL	26	
Hierarchical Locking	24	
Optimistic Concurrency Control	26	
Total:	100	

**Question 1: Query Execution, Planning, and Optimization . . . . [24 points]**

- (a) **[3 points]** The zone map optimization is more effective in speeding up OLAP queries as opposed to OLTP queries.  
 True  False
- (b) **[3 points]** For OLAP queries, which often involve complex operations on vast datasets, intra-query parallelism is typically not preferred to optimize performance.  
 True  False
- (c) **[3 points]** The process per DBMS worker approach provides better fault isolation than the thread per DBMS worker approach.  
 True  False
- (d) **[3 points]** In OLAP workload, the vectorized model's performance improvements come mainly from the reduction in the number of disk I/O operations.  
 True  False
- (e) **[3 points]** An index scan is always better (fewer I/O operations, faster run-time) than a sequential scan if the query contains an ORDER BY clause matching the index key.  
 True  False
- (f) **[3 points]** The query optimizer in a database management system always guarantees the generation of an optimal execution plan by exhaustively evaluating all possible plans to ensure the lowest cost for query execution.  
 True  False
- (g) **[3 points]** Predicate pushdown involves moving filter conditions closer to the node where the data is stored, while projection pushdown involves transferring only the necessary columns of the data.  
 True  False
- (h) **[3 points]** The vectorized query processing model (which uses SIMD instructions to parallelize operations) is an example of intra-query parallelism.  
 True  False

**Question 2: Serializability and 2PL.....[26 points]**

(a) True/False Questions:

- i. **[3 points]** Strong strict Two-Phase Locking (2PL) prevents the occurrence of cascading aborts and inherently avoids deadlocks without the need for additional prevention or detection techniques.  
 True    False
- ii. **[3 points]** For a schedule following strong strict 2PL, the dependency graph is guaranteed to be acyclic.  
 True    False
- iii. **[2 points]** Using 2PL guarantees a conflict-serializable schedule.  
 True    False
- iv. **[2 points]** Conflict-serializable schedules prevent unrepeatable reads and dirty reads.  
 True    False
- v. **[2 points]** A schedule that is view-serializable is also conflict-serializable.  
 True    False

(b) Serializability:

Consider the schedule of 4 transactions in Table 1.  $R(\cdot)$  and  $W(\cdot)$  stand for ‘Read’ and ‘Write’, respectively, and time increases from left to right. (This is in contrast to the diagrams in class, where time proceeded downward.)

	$t_1$	$t_2$	$t_3$	$t_4$	$t_5$	$t_6$	$t_7$	$t_8$	$t_9$	$t_{10}$
$T_1$		R(B)		W(C)	W(A)					
$T_2$	W(E)			R(E)			W(C)	R(D)	W(A)	
$T_3$						R(A)	W(B)			R(B)
$T_4$			R(D)	R(B)		W(D)				

Table 1: A schedule with 4 transactions

- i. **[3 points]** Is this schedule serial?  
 Yes    No
- ii. **[3 points]** Is this schedule conflict-serializable?  
 Yes    No
- iii. **[5 points]** Compute the conflict dependency graph for the schedule in Table 1, selecting all edges (and the object that caused the dependency) that appear in the graph.
 

<input type="checkbox"/> $T_1 \rightarrow T_2$	<input type="checkbox"/> $T_2 \rightarrow T_3$	<input type="checkbox"/> $T_2 \rightarrow T_4$
<input type="checkbox"/> $T_2 \rightarrow T_1$	<input type="checkbox"/> $T_3 \rightarrow T_2$	<input type="checkbox"/> $T_4 \rightarrow T_2$
<input type="checkbox"/> $T_1 \rightarrow T_3$	<input type="checkbox"/> $T_1 \rightarrow T_4$	<input type="checkbox"/> $T_3 \rightarrow T_4$
<input type="checkbox"/> $T_3 \rightarrow T_1$	<input type="checkbox"/> $T_4 \rightarrow T_1$	<input type="checkbox"/> $T_4 \rightarrow T_3$
- iv. **[3 points]** Is this schedule possible under regular 2PL?  
 Yes  
 No

**Question 3: Hierarchical Locking ..... [24 points]**

Consider a database D consisting of two tables A (which stores information about musical artists) and R (which stores information about the artists' releases). Specifically:

- R(rid, name, artist\_credit, language, status, genre, year, number\_sold)
- A(id, name, type, area, gender, begin\_date\_year)

Table R spans 1000 pages, which we denote R1 to R1000. Table A spans 50 pages, which we denote A1 to A50. Each page contains 100 records. We use the notation R3.20 to denote the twentieth record on the third page of table R. There are no indexes on these tables.

Suppose the database supports shared and exclusive hierarchical intention locks (S, X, IS, IX and SIX) at four levels of granularity: database-level (D), table-level (R and A), page-level (e.g., R10), and record-level (e.g., R10.42). We use the notation IS(D) to mean a shared database-level intention lock, and X(A2.20-A3.80) to mean a set of exclusive locks on the records from the 20th record on the second page to the 80th record on the third page of table A.

For each of the following operations below, what sequence of lock requests should be generated to **maximize the potential for concurrency** while guaranteeing correctness?

- (a) **[4 points]** Fetch the records of all musical artists in A with type = 'Orchestra'.
- SIX(D), S(A)
  - IX(D), S(A)
  - IS(D), S(A)
  - S(D)
- (b) **[4 points]** Update the genre for all release records with language = 'English' to 'Musical theatre'.
- IX(D), X(R)
  - SIX(D), X(R)
  - IX(D), IX(R)
  - IX(D), SIX(R)
- (c) **[4 points]** Modify the 7<sup>th</sup> record on R80.
- IS(D), IS(R), IS(R80), X(R80.7)
  - IX(D), IX(R), IX(R80), IX(R80.7)
  - SIX(D), IX(R), IX(R80), X(R80.7)
  - IX(D), IX(R), IX(R80), X(R80.7)
- (d) **[4 points]** Increment the number\_sold for the 6<sup>th</sup> record on R999.
- IX(D), IX(R), SIX(R999), X(R999.6)
  - IS(D), IS(R), IS(R999), S(R999.6)
  - IX(D), IX(R), S(R999), X(R999.6)
  - IX(D), IX(R), IX(R999), X(R999.6)
- (e) **[4 points]** Scan all records on pages A10 to A50 and modify the 20<sup>th</sup> record on A14.
- IX(D), S(A), X(A14)
  - SIX(D), IX(A), IX(A14), X(A14.20)

- IX(D), IX(A), IX(A10-A50), IX(A14), X(A14.20)
  - IX(D), SIX(A), IX(A14), X(A14.20)
- (f) **[4 points]** Delete records in A if type='Band'.
- SIX(D), SIX(A)
  - IX(D), X(A)
  - IX(D), IX(A)
  - SIX(D), X(A)

**Question 4: Optimistic Concurrency Control . . . . . [26 points]**

Consider the following set of transactions accessing a database with object  $A$ ,  $B$ ,  $C$ ,  $D$ . You should make the following assumptions:

- The transaction manager is using **optimistic concurrency control** (OCC).
- A transaction begins its read phase with its first operation and switches from the READ phase immediately into the VALIDATION phase after its last operation executes.
- The DBMS is using the serial validation protocol discussed in class where only one transaction can be in the validation phase at a time.
- Each transaction is doing **forward validation** (i.e. Each transaction, when validating, checks whether it intersects its read/write sets with any active transactions that have not committed yet).
- There are no other transactions in addition to the ones shown below.

Note: VALIDATION may or may not succeed for each transaction. If validation fails, the transaction will get immediately aborted.

time	$T_1$	$T_2$	$T_3$
1	READ(A)		
2			READ(B)
3			WRITE(B)
4		READ(A)	
5		READ(D)	
6	READ(B)		WRITE(C)
7			READ(D)
8			VALIDATE?
9	READ(C)		
10			WRITE?
11	VALIDATE?		
12	WRITE?		
13		WRITE(D)	
14		WRITE(C)	
15		VALIDATE?	
16		WRITE?	

Figure 1: An execution schedule

- (a) [4 points] When is each transaction's timestamp assigned in the transaction process?
- The start of the write phase.
  - Timestamps are not necessary for OCC.
  - The start of the validation phase.
  - The start of the read phase.
- (b) [4 points] When time = 9, will  $T_1$  read  $C$  written by  $T_3$ ?
- Yes    No

- (c) **[4 points]** Will T1 abort?  
 Yes  
 No
- (d) **[4 points]** Will T2 abort?  
 Yes  
 No
- (e) **[4 points]** Will T3 abort?  
 Yes  
 No
- (f) **[2 points]** OCC works best when concurrent transactions access the same subset of data in a database.  
 True    False
- (g) **[2 points]** Transactions can suffer from *phantom reads* in OCC.  
 True    False
- (h) **[2 points]** Aborts due to OCC are wasteful because they happen after a transaction has already finished executing.  
 True    False