

CARNEGIE MELLON UNIVERSITY
COMPUTER SCIENCE DEPARTMENT
15-445/645 – DATABASE SYSTEMS (FALL 2023)
PROF. ANDY PAVLO AND JIGNESH PATEL

Homework #4 (by Aolei Zhou)
Due: **Sunday November 12, 2023 @ 11:59pm**

IMPORTANT:

- Enter all of your answers into **Gradescope by 11:59pm on Sunday November 12, 2023.**
- **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 : 2023/10/27 23:23

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]** For a query that involves a range condition, zone maps can help in determining if a block contains any relevant data.
 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]** For a table with a well-maintained B-Tree index, an index scan will always be faster than a sequential scan because of fewer I/O operations and faster run-time.
 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]** In the context of query optimization, where accurate statistics are crucial for plan choices, sampling requires examining every row in a table to compute these vital statistics.
 True False

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

(a) True/False Questions:

- i. **[2 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. **[2 points]** For a schedule following strong strict 2PL, the dependency graph is guaranteed to be acyclic.
 True False
- iii. **[2 points]** Using regular (i.e., not strong strict) 2PL does not guarantee a conflict serializable schedule.
 True False
- iv. **[2 points]** Conflict serializable schedules prevent unrepeatable reads and dirty reads, even in the event of cascading aborts.
 True False
- v. **[2 points]** 2PL is a pessimistic concurrency control protocol.
 True False
- vi. **[2 points]** Every conflict-serializable schedule is always conflict-equivalent to at least one view-serializable schedule.
 True False

(b) Serializability:

Consider the schedule of 4 transactions in Table 1. R(·) and W(·) 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(C)			R(A)					
T_2							W(A)	R(D)	W(B)	
T_3	W(B)					W(C)				R(D)
T_4			W(B)	W(A)						

Table 1: A schedule with 4 transactions

- i. **[2 points]** Is this schedule serial?
 Yes No
- ii. **[2 points]** Is this schedule conflict serializable?
 Yes No
- iii. **[2 points]** Is this schedule view serializable?
 Yes No
- iv. **[6 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.

- $T_1 \rightarrow T_2$
- $T_2 \rightarrow T_1$
- $T_1 \rightarrow T_3$
- $T_3 \rightarrow T_1$

- $T_2 \rightarrow T_3$
- $T_3 \rightarrow T_2$
- $T_1 \rightarrow T_4$
- $T_4 \rightarrow T_1$

- $T_2 \rightarrow T_4$
- $T_4 \rightarrow T_2$
- $T_3 \rightarrow T_4$
- $T_4 \rightarrow T_3$

v. **[2 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]** Update the type of all musical artists in table A with the name = 'Lunar Echoes' to 'Band'
- X(D)
 - IX(D), X(A)
 - SIX(D), X(A)
 - IX(D), IX(A)
- (b) **[4 points]** Find the release record whose year equals the most recent year in all releases.
- S(D)
 - IS(D), IS(R)
 - IX(D), X(R)
 - IS(D), S(R)
- (c) **[4 points]** Modify the 30th record on A12.
- IX(D), IX(A), IX(A12), IX(A12.30)
 - IX(D), IX(A), IX(A12), X(A12.30)
 - SIX(D), SIX(A), SIX(A12), X(A12.30)
 - IS(D), IS(A), IS(A12), X(A12.30)
- (d) **[4 points]** Scan all records in R and modify the 3rd record on R4
- IX(D), SIX(R), IX(R4), X(R4.3)
 - IS(D), S(R), IX(R4.3)
 - S(D), IX(R), X(R4.3)
 - IS(D), SIX(R), X(R4)
- (e) **[4 points]** Scan all records in R and modify the 23rd record on R7.
- IX(D), SIX(R), IX(R7), X(R7.23)
 - IS(D), IS(R), IS(R7), X(R7.23)

- SIX(D), SIX(R), IX(R_7), X(R7.23)
 - IS(D), SIX(R), X(R7.23)
- (f) **[4 points]** Two users are trying to access data. User A is scanning all the records in R to read, while User B is trying to modify the 23th record in A1. Which of the following sets of locks are most suitable for this scenario?
- User A: SIX(D), S(R), User B: SIX(D), IX(A), X(A1.23)
 - User A: S(D), User B: X(D)
 - User A: IS(D), S(R), User B: IX(D), IX(A), X(A1.23)
 - User A: IS(D), S(R), User B: SIX(D), IX(A), X(A1.23)

Question 4: Optimistic Concurrency Control [26 points]

Consider the following set of transactions accessing a database with object A, B, C, D . The questions below assume that the transaction manager is using **optimistic concurrency control (OCC)**. Assume that 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.

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

You can assume that the DBMS is using the serial validation protocol discussed in class where only one transaction can be in the validation phase at a time, and each transaction is doing backward validation (i.e. Each transaction, when validating, checks whether it intersects its read/write sets with any transactions that have already committed. You may assume there are no other transactions in addition to the ones shown below.)

time	T_1	T_2	T_3
1		READ(A)	
2		WRITE(A)	
3	READ(A)		
4	WRITE(A)		
5			READ(C)
6	READ(B)		
7	READ(E)		
8	VALIDATE?		
9		WRITE(B)	
10	WRITE?		
11		WRITE(C)	
12		VALIDATE?	
13			WRITE(C)
14		WRITE?	
15			READ(D)
16			WRITE(D)
17			VALIDATE?
18			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 = 3, will T_1 read A written by T_2 ?
- 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