

CARNEGIE MELLON UNIVERSITY
COMPUTER SCIENCE DEPARTMENT
15-445/645 – DATABASE SYSTEMS (FALL 2025)
PROF. ANDY PAVLO

Homework #5 (by Will) – Solutions
Due: **Sunday Nov 16, 2025 @ 11:59pm**

IMPORTANT:

- Enter all of your answers into **Gradescope by 11:59pm on Sunday Nov 16, 2025**.
- **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)
- Each part is all or nothing. There is no partial credit.

Revision : 2025/11/26 10:29

Question	Points	Score
Serializability and 2PL	30	
Hierarchical Locking	30	
Optimistic Concurrency Control	25	
Multi-Version Concurrency Control	15	
Total:	100	

Question 1: Serializability and 2PL.....[30 points]

(a) True/False Questions:

i. **[2 points]** 2PL is an optimistic concurrency control protocol.

True False

Solution: False. 2PL is considered a pessimistic concurrency control protocol because it uses locks to prevent conflicts before they can occur. By ensuring transactions follow a strict locking protocol, 2PL aims to prevent any possible conflicts, in contrast to optimistic protocols which allow conflicts but then resolve them.

ii. **[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

Solution: False. While strict 2PL prevents cascading aborts by holding all the locks until a transaction reaches its commit point, ensuring that other transactions do not see the intermediate, uncommitted data, it does not inherently resolve deadlocks. Deadlocks, which are situations where two or more transactions prevent each other from progressing by holding locks that the other needs, still require specific detection or prevention mechanisms in strict 2PL.

iii. **[2 points]** Using regular (i.e., not strong strict) 2PL guarantees a conflict serializable schedule.

True False

Solution: True. Using regular 2PL guarantees a conflict serializable schedule because it generates schedules whose precedence graph is acyclic. This is because this ordering of acquiring resources (via the locks) ensures that there is a order of acquisition precedence.

iv. **[2 points]** Every view-serializable schedule is always conflict-equivalent to at least one conflict-serializable schedule.

True False

Solution: False. View-Serializable schedules permits schedules with “blind” writes that may not be conflict-serializable.

v. **[2 points]** Conflict-serializable schedules prevent unrepeatable reads and dirty reads.

True False

Solution: False. Conflict-serializability ensures that the schedule is conflict equivalent to a serial schedule, thus protecting against unrepeatable reads. However, dirty reads can still occur in conflict-serializable schedules, especially when considering cascading aborts. If a transaction reads data written by another transaction which later gets aborted, then the first transaction has effectively read “dirty” data.

(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			$W(B)$	$W(A)$						
T_2	$R(C)$					$R(A)$				
T_3		$W(B)$			$W(C)$				$R(D)$	
T_4							$R(D)$	$W(B)$		$W(A)$

Table 1: A schedule with 4 transactions

i. [3 points] Is this schedule serial?

Yes No

Solution: This schedule isn’t serial because this schedule interleaves the actions of different transactions.

ii. [8 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.

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

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

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

Solution: The answer is:

For A, there are W-R conflict from T_1 to T_2 , R-W conflict from T_2 to T_4 , and W-W conflict from T_1 to T_4 ;

For B, there are W-W conflict from T_3 to T_1 , T_1 to T_4 , and T_3 to T_4 ;

For C, there is R-W conflict from T_2 to T_3 ;

For D, there is no conflict.

iii. [3 points] Is this schedule conflict serializable?

Yes No

Solution: This schedule isn’t conflict serializable because there is a cycle in its data dependency graph. Moreover, this schedule is not conflict equivalent (every pair of conflicting operations is ordered in the same way) to any serial schedule of transaction execution.

iv. [3 points] Is this schedule view serializable?

Yes No

Solution: To determine if the schedule is view serializable, we need to check it against three criteria for view equivalence for original schedule S_1 and serial schedule S_2 :

1. If T_i reads the initial value of A in S_1 , then T_i must also read the initial value of A in S_2 .
2. If T_i reads a value of A written by T_j in S_1 , then T_i must also read that value of A written by T_j in S_2 .
3. If T_i writes the final value of A in S_1 , then T_i must also write the final value of A in S_2 .

From our analysis, we found a possible serial schedule that is view-equivalent to the given schedule: $T_1 \rightarrow T_2 \rightarrow T_3 \rightarrow T_4$. This serial schedule fulfills the criteria for view equivalence:

- T_1 writes B and A .
- T_2 reads the initial value of C and the value of A written by T_1 .
- T_3 writes B , writes the final version of C , and reads the initial value of D .
- T_4 reads the initial value of D , writes the final version of B , and writes the final version of A .

Therefore, this serial schedule is view-equivalent to the original one, so it is proved to be view-serializable. Importantly, compared to conflict serializability, view serializability allows all conflict serializable schedules along with “blind writes”, ensuring database consistency after all transactions are committed.

v. [3 points] Is this schedule possible under regular 2PL?

Yes
 No

Solution: This schedule is not possible under 2PL because it is not conflict serializable, and 2PL is guaranteed to produce conflict serializable schedules.

Question 2: Hierarchical Locking [30 points]

Consider a database D consisting of two tables C (which stores information about companies) and P (which stores information about products). Specifically:

- P(pid, name, company_id, category, status, market, release_year, units_sold)
- C(cid, name, type, headquarters, ceo_gender, founded_year)

Table P spans 2500 pages, which we denote P1 to P2500. Table C spans 150 pages, which we denote C1 to C150. Each page contains 160 records. We use the notation P3.20 to denote the twentieth record on the third page of table P. 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 (P and C), page-level (e.g., P10), and record-level (e.g., P10.42). We use the notation IS(D) to mean a shared database-level intention lock, and X(C2.20-C3.80) 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 C.

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) [5 points] Find the product record that has the largest units_sold where release_year = 1996.

- IX(D), X(P)
- IS(D), S(P)
- S(D)
- IS(D), IS(P)

Solution: The correct choice is IS(D), S(P). The query needs to scan all records in P to find the product with the largest units sold. So, it requires a shared lock on P.

- S(D) is incorrect because it doesn't acquire the intended parent locks necessary to obtain the S(P) lock for reading the table P.
- IS(D), IS(P) is incorrect because it only states the intention to read-access the table P but doesn't actually acquire the lock to perform the read.
- IX(D), X(P) is incorrect because it accesses an intended exclusive lock, which is not needed for a read-only query.

(b) [4 points] Update the type of all companies in table C with the name = 'MegaCorp' to 'Conglomerate'

- X(D)
- SIX(D), X(C)
- IX(D), IX(C)
- IX(D), X(C)

Solution: The correct answer choice is IX(D), X(C). This is because we potentially need to modify records in table C where the company's name is 'MegaCorp', and this choice accesses the intended exclusive parent lock to get the exclusive lock on those specific rows in table C.

- X(D) is incorrect because it gains an exclusive lock on the entire database D when it only needs to modify specific rows in table C.
- SIX(D), X(C) is incorrect because it gains a shared intention lock on the database D when it has no need to read the contents of D.
- IX(D), IX(C) is incorrect because it does not gain the exclusive lock on the specific rows in table C that match the condition, which it needs to modify the type of those companies.

(c) **[5 points]** Increment the units_sold for the 5th record on P2450.

- IX(D), IX(P), IX(P2450), IX(P2450.5)
- SIX(D), SIX(P), SIX(P2450), X(P2450.5)
- IX(D), IX(P), IX(P2450), X(P2450.5)
- IS(D), IS(P), IS(P2450), X(P2450.5)

Solution: The correct choice is IX(D), IX(P), IX(P2450), X(P2450.5)). This choice is correct because it accesses all intended exclusive locks for all parent levels necessary, and then accesses the exclusive lock for the particular record.

- IX(D), IX(P), IX(P2450), IX(P2450.5) is incorrect because it only gets the intention exclusive lock for the record.
- SIX(D), SIX(P), SIX(P2450), X(P2450.5) is incorrect because the DBMS only intends to write to a tuple, not read any data. Therefore it should not grab the shared-exclusive intention locks.
- IS(D), IS(P), IS(P2450), X(P2450.5) is incorrect because the DBMS intends to write to a tuple, so it should not grab the shared intention parent locks.

(d) **[5 points]** Delete records in C if type = 'Shell'.

- SIX(D), SIX(C)
- IX(D), IX(C)
- SIX(D), X(C)
- IX(D), X(C)

Solution: The correct choice is IX(D), X(C). This choice is correct because it accesses all intended locks and the exclusive lock on C, since we potentially need to modify all records in C.

- SIX(D), SIX(C) is incorrect because it gains a shared intention parent lock for

both D and C, when it does not need to read any contents, and it does not gain the exclusive lock for the records of C it needs to delete..

- IX(D), IX(C) is incorrect because it does not gain an exclusive lock for the records of C it needs to delete.
- SIX(D), X(C) is incorrect because it gains a shared lock for the database, when it does not need to read any contents.

(e) **[5 points]** Scan all records between P30 and P260 and modify the 5th record on P75

- IX(D), SIX(P), IX(P75), X(P75.5)
- IS(D), S(P), IX(P75.5)
- IX(D), IX(P), IX(P30-P260), IX(P75), X(P75.5)
- IS(D), SIX(P), X(P75)

Solution: The correct choice is IX(D), SIX(P), IX(P75), X(P75.5). This choice is correct because it accesses all intended locks and the exclusive lock X(P75.5). It also gains a shared intention lock on P, so it can read and modify records in P.

- IS(D), S(P), IX(P75.5) is incorrect because it accesses an intended shared parent lock for both D and P, when we plan to modify a record in P.
- IX(D), IX(P), IX(P30-P260), IX(P75), X(P75.5) is incorrect because we do plan on reading and modifying in relation P, but we are only gaining a exclusive lock on the whole P table, which is too much.
- IS(D), SIX(P), X(P75) is incorrect because while it gains a shared intention lock for P and the database, it tries to exclusively lock the whole page P4 instead of just the specific record.

(f) **[6 points]** Two users are trying to access data. User C is scanning all the records in P to read, while User B is trying to modify the 7th record in C16. Which of the following sets of locks are most suitable for this scenario?

- User C: IS(D), S(P), User B: IX(D), IX(C), IX(C16), X(C16.7)
- User C: SIX(D), S(P), User B: SIX(D), IX(C), IX(C16), X(C16.7)
- User C: S(D), User B: X(D)
- User C: IS(D), S(P), User B: SIX(D), IX(C), IX(C16), X(C16.7)

Solution: The correct choice is User C: IS(D), S(P), User B: IX(D), IX(C), IX(C16), X(C16.7). This choice is correct because:

- User C only intends to read all records in P. Thus, he acquires an intention shared lock on the database D and a shared lock on the table P.
- User B intends to modify a specific record in table C. He acquires an intention exclusive lock on the database D, an intention exclusive lock on the table C and C16, and then an exclusive lock on the specific record C16.7.

Other choices are incorrect because:

- User C: SIX(D), S(P) and User B: SIX(D), IX(C), IX(C16), X(C16.7) is incorrect because User C doesn't need to acquire shared intention exclusive locks for a mere read operation, and neither does User B for a specific record modification.
- User C: S(D) and User B: X(D) is problematic as it locks the entire database either in shared mode or exclusive mode, preventing concurrent operations.
- User C: IS(D), S(P) and User B: SIX(D), IX(C), IX(C16), X(C16.7) is incorrect because while User C's locks are appropriate for his read operation, User B does not need a shared intention exclusive lock on the database for modifying a specific record.

Question 3: Optimistic Concurrency Control [25 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		WRITE(B)	
2	READ(A)		READ(C)
3		WRITE(C)	
4	WRITE(A)		
5			
6		READ(A)	
7	READ(B)		
8	VALIDATE?		WRITE(C)
9	WRITE?		
10			
11		WRITE(A)	READ(D)
12		VALIDATE?	
13		WRITE?	
14			READ(A)
15			
16			WRITE(D)
17			VALIDATE?
18			WRITE?

Figure 1: An execution schedule

(a) [3 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.

Solution: Each transaction's timestamp is assigned at the beginning of the validation phase.

(b) [3 points] Will T1 abort?

- Yes
- No

Solution: T_1 will not abort because T_1 's read set does not intersect with any other transactions read-write set before T_1 commits in backward validation.

(c) [3 points] Will T2 abort?

- Yes
- No

Solution: T_2 will abort because T_2 should have read the value of A written by T_1 .

(d) [3 points] Will T3 abort?

- Yes
- No

Solution: T_3 won't abort because the write sets between T_3 and T_1 do not overlap.

(e) [3 points] When time = 14, will T_3 read A written by T_1 ?

- Yes
- No

Solution: Yes. In OCC, when a transaction commits then the write set is written to the global database and when a transaction wants to read and bring the value of a object in its private workspace it picks the value up from the global database at the time of reading. So, T_3 will read A at 14 which was written by T_1 .

(f) [2 points] OCC works best when concurrent transactions access disjoint subsets of data in a database.

- True
- False

Solution: OCC is good to use when the number of conflicts is low.

(g) [2 points] Aborts due to OCC are wasteful because they happen after a transaction has already finished executing.

- True
- False

(h) [2 points] Transactions can suffer from *unrepeatable reads* in OCC.

- True
- False

(i) [2 points] Transactions can suffer from *phantom reads* in OCC.

- True
- False

(j) [2 points] All Objects in the read-set of a transaction are read as a whole when the transaction starts.

- True
- False

Question 4: Multi-Version Concurrency Control [15 points]

Mark either **True** or **False** for each of the following statements.

Assume that the DBMS in each scenario supports multi-versioning of tuples.

(a) Consider a MVCC DBMS using **append-only** version storage.

i. [3 points] Newest-to-oldest (N2O) version chain ordering is faster than oldest-to-newest (O2N) for queries that need to read the latest version using an index scan on the primary key.

True **False**

Solution: True. When the DBMS does a look-up on the primary key index, the index will point to the head of the version chain for a tuple which will be the newest version.

ii. [3 points] The DBMS can use *cooperative garbage collection* with oldest-to-newest (O2N) version chain ordering only if it also uses logical pointers for secondary indexes.

True **False**

Solution: False. The secondary index scheme does not preclude the DBMS from using cooperative garbage collection, since workers will be able to remove older versions during the version chain scan. It doesn't matter how the worker got to the version chain.

(b) Consider a MVCC DBMS using **delta** version storage.

i. [3 points] *Background vacuuming* will only remove delta records for old versions and never tuples from the main table.

True **False**

Solution: False. The vacuum can also remove deleted tuples from the main table.

ii. [3 points] If the DBMS uses *newest-to-oldest* (N2O) version chain ordering, all reclaimable versions will be found in the delta storage area.

True **False**

Solution: False. A tuple with only one version will only be found in the main table storage.

iii. [3 points] When a transaction modifies an existing tuple's primary key attributes, the DBMS must delete the tuple and then re-insert it (with its new primary key) regardless of the version chain ordering.

True **False**

Solution: True. Any update to a tuple's primary key has to be treated as a DELETE/INSERT because the DBMS has to maintain two pointers in the primary key index (one to the original tuple, one to the new tuple).