

**CS 245**  
**Final Exam – Winter 2017**

This exam is open book and notes. You can use a calculator and your laptop to access course notes and videos (but not to communicate with other people or to browse the Internet). You have 180 minutes to complete the exam.

Best of luck!!!

Print your name: \_\_\_\_\_

\_\_\_\_\_

The Honor Code is an undertaking of the students, individually and collectively:

1. that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
2. that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.

The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.

While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

I acknowledge and accept the Honor Code.

Signed: \_\_\_\_\_

Problem	Points	Maximum
1		10
2		10
3		10
4		10
5		10
6		10
7		20
Total		80

*Please note: Problem 7 (true/false) is worth 20 points.*

## Problem 1: Recovery (10 Points)

Consider a database with two elements, X and Y. The initial values of X and Y are both 0. We consider three transactions U, V and W that modify these elements concurrently:

- $T_1$ :  $X := 42$
- $T_2$ :  $Y := 20, X := 10$
- $T_3$ :  $X := 100, Z := 101$

While the transactions execute, the database system crashes. We consider three recovery mechanisms in turn below; your task is to first complete the missing parts of the log and then answer the following questions:

(a) Using **pure undo logging**:

1.  $\langle START\ T_2 \rangle$
2.  $\langle START\ T_3 \rangle$
3.  $\langle T_2, X, 0 \rangle$
4.  $\langle T_2, Y, 0 \rangle$
5.  $\langle COMMIT\ T_2 \rangle$
6.  $\langle START\ CKPT(T_3) \rangle$
7.  $\langle T_3, X, ??? \rangle$
8.  $\langle START\ T_1 \rangle$
9.  $\langle T_3, Z, 0 \rangle$
10.  $\langle T_1, X, ??? \rangle$

(i) Please complete the undo log by providing in the appropriate values for the log below:

Line 6: \_\_\_\_\_

Line 10: \_\_\_\_\_

(ii) If the database crashes immediately after writing the above log entries, then subsequently performs recovery, which of the following statements is *true*:

- A. After a successful recovery, the state of Y is 0.
- B. At the time of crash, the state of X (on disk) must be 10.
- C. At the time of crash, the state of Z (on disk) must be 101.
- D. After a successful recovery, the state of X is 10.

True statement (one of A, B, C, or D): \_\_\_\_\_

(b) Using **pure redo logging**:

1.  $\langle \text{START } T_2 \rangle$
2.  $\langle \text{START } T_3 \rangle$
3.  $\langle T_2, X, 10 \rangle$
4.  $\langle T_2, Y, 20 \rangle$
5.  $\langle \text{COMMIT } T_2 \rangle$
6.  $\langle \text{START CKPT}(T_3) \rangle$
7.  $\langle T_3, X, ??? \rangle$
8.  $\langle \text{START } T_1 \rangle$
9.  $\langle T_3, Z, ??? \rangle$
10.  $\langle T_1, X, 42 \rangle$
11.  $\langle \text{END CKPT} \rangle$
12.  $\langle \text{COMMIT } T_3 \rangle$
13.  $\langle \text{START CKPT}(T_1) \rangle$
14.  $\langle \text{COMMIT } T_1 \rangle$

(i) Line 7: \_\_\_\_\_

Line 9: \_\_\_\_\_

(ii) If the database crashes immediately after writing the above log entries, then subsequently performs recovery, which of the following statements is *true*:

- A. After a successful recovery, the state of X is 100.
- B. At the time of crash, the state of X (on disk) must be 42.
- C. At the time of crash, the state of Y (on disk) must be 20.
- D. After a successful recovery, the state of Z is 0.

True statement (one of A, B, C, or D): \_\_\_\_\_

(c) Using **undo-redo logging**, with log format (*txn*, *item*, before-val, after-val):

1.  $\langle START\ T_1 \rangle$
2.  $\langle START\ T_2 \rangle$
3.  $\langle START\ T_3 \rangle$
4.  $\langle T_3, X, 0, 100 \rangle$
5.  $\langle T_3, Z, 0, 101 \rangle$
6.  $\langle COMMIT\ T_3 \rangle$
7.  $\langle T_2, X, 100, 10 \rangle$
8.  $\langle START\ CKPT(T_1, T_2) \rangle$
9.  $\langle T_2, Y, ???, ??? \rangle$
10.  $\langle END\ CKPT \rangle$
11.  $\langle START\ CKPT(T_1, T_2) \rangle$
12.  $\langle COMMIT\ T_2 \rangle$
13.  $\langle T_1, X, ???, ??? \rangle$

(i) Please complete the undo-redo log by providing in the appropriate values for the log below:

Line 9: \_\_\_\_\_

Line 13: \_\_\_\_\_

(ii) If the database crashes immediately after writing the above log entries, which of the following statements is *true*:

- A. During recovery, the log entry in line 4 will be redone.
- B. During recovery, the log entry in line 7 will be ignored.
- C. During recovery, the log entry in line 9 will be ignored.
- D. During recovery, the log entry in line 13 will be undone.

True statement (one of A, B, C, or D): \_\_\_\_\_

## Problem 2: Serializability and Precedence Graphs (10 Points)

For each of the following schedules, draw the precedence graph and determine whether the schedule is conflict serializable. If the schedule is conflict serializable, write the conflict equivalent serial schedule of *transactions* (e.g.,  $T_5; T_7; T_2$ ). If the schedule is not conflict serializable, identify a set of conflicting actions that prevents it from being conflict serializable (e.g.,  $W_5(C)$ ,  $W_7(D)$ ,  $W_2(D)$ ).

To reiterate: for each, (a) draw the precedence graph, and, (b) if conflict serializable, give the conflict equivalent serial ordering of transactions or, if not, identify the conflicting operations.

(a)  $R_1(A)W_2(B)R_2(A)W_1(B)W_3(B)W_1(C)R_3(B)W_1(A)$

(b)  $W_2(C)R_1(A)W_3(B)R_1(C)R_3(B)R_3(A)W_1(B)$

(c)  $R_3(B)R_2(A)R_1(A)R_2(C)W_3(A)W_2(B)W_1(C)$

(d)  $W_1(A)W_2(B)W_1(B)W_2(A)$

### Problem 3: 2PL and Validation (10 points)

Consider the following schedules and decide which schedules are valid under a legal, two-phase-locking scheduler. (Assume that only exclusive locks are allowed.) If yes, augment the schedule with lock and unlock actions to show it can be generated by 2PL. For example,  $l_1(A)$  and  $u_1(A)$  for locking and unlocking object A by transaction 1. If not, circle the **first** action in the schedule that is problematic.

(a)  $r_1(A); w_1(A); w_2(B); w_3(C); r_1(D); w_2(D)$

(b)  $r_3(A); w_4(B); r_1(C); r_3(D); w_3(B); w_2(D); r_3(A); w_1(D); r_3(B); r_2(C); r_1(A)$

The following 5 transactions execute under a validation concurrency control protocol:

- $T_1$ : Read Set:  $\{B, E\}$ , Write Set:  $\{C, E, G\}$
- $T_2$ : Read Set:  $\{G\}$ , Write Set:  $\{A, B\}$
- $T_3$ : Read Set:  $\{A\}$ , Write Set:  $\{C\}$
- $T_4$ : Read Set:  $\{A, C\}$ , Write Set:  $\{B, D\}$
- $T_5$ : Read Set:  $\{G\}$ , Write Set:  $\{C, D\}$

Consider the following sequence of events:

1.  $T_1$  starts
2.  $T_2$  starts
3.  $T_2$  begins validation
4.  $T_1$  begins validation
5.  $T_3$  starts
6.  $T_2$  finishes validation
7.  $T_4$  starts
8.  $T_5$  starts
9.  $T_4$  begins validation
10.  $T_1$  finishes validation
11.  $T_5$  begins validation
12.  $T_3$  begins validation
13.  $T_4$  finishes validation
14.  $T_3$  finishes validation
15.  $T_5$  finishes validation



State which transactions validate successfully and which do not. If a transaction successfully validates, write “YES.” If not, write ‘NO’ and state the conflict by stating the transactions that conflict, and which item(s) they conflict on (e.g., “ $T_1$  and  $T_2$  have a read-write conflict on F”) .

(a) Does  $T_1$  validate successfully? If not, state the conflict as instructed above.

(b) Does  $T_2$  validate successfully? If not, state the conflict as instructed above.

(c) Does  $T_3$  validate successfully? If not, state the conflict as instructed above.

(d) Does  $T_4$  validate successfully? If not, state the conflict as instructed above.

(e) Does  $T_5$  validate successfully? If not, state the conflict as instructed above.

## Problem 4: Multi-Granularity Locking (10 points)

Consider a multi-granularity locking system, with lock modes S, X, IS, IX, and SIX as in lecture, with a hierarchy from Relation, Table, Record, to Field. In this question, we consider a database with the following contents:

- Relation R has 2 underlying tables, A and B.
  - Table A has 2 records t1 and t2, each with 3 fields E, F, and G.
  - Table B has 2 records t3 and t4, each with 3 fields H, I, and J.
- (a) Provide the sequence of lock requests required to perform the given operation. Your answer should use the least restrictive locks possible, allowing the greatest degree of concurrency. For example, the answer  $X(R)$  avoids conflicts in all cases below, but is too restrictive. Provide your answers using the provided spaces, with one lock or unlock operation per space proceeding left to right (you need not use all the spaces provided). For example:

$X(R); SIX(t_1); IX(A); X(B)$   
 $SIX(t_4); SIX(I)$

i) Read all records in A:



ii) Modify the H field of t3:


iii) Perform the operation  $t1.E = t3.H$  (a read of t3.H and a write of t1.E are needed):


(b) Consider the following four schedules, each with two transactions, U and V, executing under the multi-granularity locking system and database as above. Circle the schedules that can be achieved under *strict* 2-phase locking. Below,  $w_U(A)$  refers to write of table A by transaction U.

- (i)  $r_U(t1.E), w_V(A), commit(U), commit(V)$
- (ii)  $r_V(A), w_U(t3), r_U(t1), commit(U), commit(V)$
- (iii)  $w_V(t1), r_U(A), r_V(A), commit(U), commit(V)$
- (iv)  $w_U(t1), r_U(A), w_V(t3), commit(U), commit(V)$

(c) Consider the following index consisting of nodes N1 through N10:



What is the sequence of lock and unlock operations required to lock node 10? As above, provide your answer in the spaces below, with one lock or unlock operation per space proceeding left to right (you need not use all the spaces provided). Only depict lock or unlock operations (no intention locks). For example:

$L(N_5)U(N_5)L(N_1)U(N_1)$   
 $L(N_4)L(N_1)L(N_{10})$

*hint:* the above example is not correct.


## Problem 5: Deadlock Prevention (10 points)

Consider three transactions,  $T_1$ ,  $T_2$  and  $T_3$  that exclusive lock (L) and unlock (U) database items  $A$ ,  $B$  and  $C$  in the following order:

- $T_1 : L(B)L(C)U(C)U(B)$
- $T_2 : L(A)L(C)U(A)U(C)$
- $T_3 : L(C)L(B)U(C)U(B)$

For this problem we want to simulate how these transactions execute under either the wait-die or wound-wait deadlock prevention strategies.

For the simulation, assume that these transactions execute in a round-robin fashion ( $T_1$  then  $T_2$  then  $T_3$ ). When it is a transaction's turn, it executes its next lock or unlock step if it can, and otherwise dies or waits or wounds the holder of the lock, as appropriate. If a transaction is waiting when it is its turn, it skips the step and continues waiting. When a waiting transaction is restarted (due to the action of some other transaction) it becomes eligible to run at its very next turn.

Assume that the transactions have timestamps in the order of  $T_1 < T_2 < T_3$  ( $T_1$  is the oldest).

Fill in the following tables to show the sequence of steps that the three transactions make. Use  $L(X)$  to denote locking an object  $X$ ,  $U(X)$  to denote unlocking  $X$ , "Die" to denote the transaction dying or being wounded, and "Wait" to denote the transaction waiting. If a transaction A is wounded by some other transaction B, transaction B's next action is 'Die'.

In the grid, the cell for  $[i, j]$  (row  $i$  column  $T_j$ ) represents the  $i^{th}$  action of  $T_j$ . We have filled out a few cells to illustrate how you should fill out the rest of the table. Simulate actions until all three transactions complete (not all cells in the grid may be necessary). If a transaction is already complete, its turn is skipped.

(a) Simulation when **wait-die** deadlock prevention is used:

Step	$T_1$	$T_2$	$T_3$
1	L(B)	L(A)	L(C)
2	Wait	Wait	Die
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			

(b) Simulation when **wound-wait** deadlock prevention is used:

Step	$T_1$	$T_2$	$T_3$
1	L(B)	L(A)	L(C)
2	L(C)	Wait	Die
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
14			

- (c) Circle each of the following statements that are true:
- (i) Wound-wait always outperforms wait-die.
  - (ii) If we know all records that each transaction will access, we can prevent deadlocks by imposing a total order on lock acquisitions.
  - (ii) Wait-die always outperforms wound-wait.
  - (iv) In general, locking uses fewer resources than optimistic methods that rely on validation.

## Problem 6: Pre-Midterm Potpourri (10 Points)

- (a) Query Processing I/Os: Consider performing a natural join operation on two relations  $R(A, B)$  and  $S(B, C)$ . Assume that  $R$  contains 36000 records stored contiguously on 1800 disk blocks (i.e.,  $B(R) = 1800$ ; 20 tuples per block), and that  $S$  contains 1000 records stored non-contiguously and that each block of memory and disk can store 20 tuples of  $S$ .
- (i) Suppose we have just two blocks of memory. Ignoring output I/Os, what is the minimum number of I/Os required to perform the natural join via a iteration join if use our two blocks of memory to store one block each of  $R$  and  $S$  in memory at a time?:
- (ii) Suppose we have just enough memory to allow a two-pass hash join (*not* hybrid hash join!). Ignoring output I/Os, Using a two-pass hash join (as described in class notes 7, slides 61-64), what is the total number of I/Os required to perform the natural join via a hash join?:
- (b) Cardinality Estimation: Consider relations  $M(A, B)$ ,  $N(C, D)$ , each containing 500 records.
- (i) What is the expected number of records in  $[M \bowtie N]$ ?
- (ii) Given a selection predicate  $B = 2$  with expected selectivity 0.01 (i.e., 1 in 100 records will match the predicate). What is the expected number of records in  $[\sigma_{B=2}M \bowtie N]$ ?



- (iii) Given a selection predicate  $B = 2$  with expected selectivity 0.01. What is the expected number of records in  $\sigma_{B=2}[M \bowtie N]$ ?
- (iv) Given a selection predicate  $B = 2$  with expected selectivity 0.01 and a selection predicate  $C = 4$  with expected selectivity 0.1. What is the expected number of records in  $[\sigma_{B=2}M \bowtie \sigma_{C=4}N]$ ?
- (c) Storage Costs: Consider relation  $D(A, B, C)$  such that column  $A$  contains a 8 byte string, column  $B$  contains a 4-byte floating point number, and column  $C$  contains a binary blob of up to 2048 bytes represented by a 2-byte header followed by the data. If we store  $D$  on a disk with block size 512:
- (i) What is the minimum number of blocks required to store 1000 records of  $D$ ?
- (i) Ignoring the cost of additional headers required to mark records that span multiple blocks, what is the maximum number of blocks required to store 1000 records of  $D$ ?

## Problem 7: True/False, Advanced Topics (20 points)

State if the following statements are true or false. Please write TRUE or FALSE in the space provided.

1. To maintain data durability in the presence of up to two fail-stop server failures, we must ensure modifications to data are persisted on three or more servers.

ANSWER: \_\_\_\_\_

2. Many NoSQL databases such as MongoDB provide flexible schemas, making it easy to add new columns and even nested columns to data.

ANSWER: \_\_\_\_\_

3. Serializability is always required to maintain application invariants in the presence of concurrent operations.

ANSWER: \_\_\_\_\_

4. MapReduce was the first data processing system to allow users to write user-defined functions.

ANSWER: \_\_\_\_\_

5. All transactional databases such as Oracle and SAP HANA guarantee serializability.

ANSWER: \_\_\_\_\_

6. In two-phase commit, it is always impossible for the coordinator to unblock the protocol if a single participant fails.

ANSWER: \_\_\_\_\_

7. In distributed model training, sequential execution always results in faster convergence time.

ANSWER: \_\_\_\_\_

8. For read queries that access contiguous sets of keys (e.g., IDs 1000-2000), a distributed database that uses hash-based partitioning will have to contact fewer partitions than than a distributed database that uses range-based partitioning.

ANSWER: \_\_\_\_\_

9. In the absence of failures, Spark enables faster execution than MapReduce because it avoids writing as much intermediate state to disk.

ANSWER: \_\_\_\_\_

10. Flexible schemas can make query processing less efficient because the data is less regularly structured and database has less control over how data is represented.

ANSWER: \_\_\_\_\_

11. In the presence of unbounded message delays and fail-stop server failures, there exists at least one scenario for every atomic commitment protocol in which the protocol must block indefinitely.

ANSWER: \_\_\_\_\_

12. To decide on a primary for each partition in a database cluster, we can use a consensus service like Zookeeper.

ANSWER: \_\_\_\_\_

13. In HogWild!-style asynchronous model training, increasing asynchrony is guaranteed to decrease convergence time.

ANSWER: \_\_\_\_\_

14. Consider the following protocol: write to servers R1, R2, and read from *either* servers R2, R3 or R4. Is this protocol a valid instance of quorum replication (as defined in class)?

ANSWER: \_\_\_\_\_

15. Given that it takes 134ms for the speed of light to circle Earth's equator, two datacenters simultaneously executing conflicting serializable transactions on opposite sides of the planet can commit more than 7 transactions per second (without relying on quantum communication/entanglement).

ANSWER: \_\_\_\_\_

16. Instead of providing serializability, Amazon's Dynamo key-value store allows users to write their own conflict resolution policies in the presence of conflicting writes.

ANSWER: \_\_\_\_\_

17. In distributed computing, it's always a good idea to assume the network is reliable.

ANSWER: \_\_\_\_\_

18. In distributed computing, it's always a good idea to assume that latency is zero.

ANSWER: \_\_\_\_\_

19. It is impossible to combine relational query processing with parameter server-style model training in a single system.

ANSWER: \_\_\_\_\_

20. As a Stanford student and CS245 graduate, you have the knowledge and the awesome power to build the next generation of data management systems.

ANSWER: \_\_\_\_\_