

*Department of Electrical Engineering and Computer Science*

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

## **6.5830/31 Database Systems: Fall 2023 Quiz II**

There are 21 questions and 9 pages in this quiz booklet. To receive credit for a question, answer it according to the instructions given. *You can receive partial credit on questions.* You have **80 minutes** to answer the questions.

**Write your name on this cover sheet AND at the bottom of each page of this booklet.**

Some questions may be harder than others. Attack them in the order that allows you to make the most progress. If you find a question ambiguous, be sure to write down any assumptions you make. Be neat. If we can't understand your answer, we can't give you credit!

**THIS IS AN OPEN BOOK, OPEN NOTES QUIZ.  
LAPTOPS MAY BE USED FOR NOTES AND SLIDES; NO PHONES, INTERNET, NOR  
ON LAPTOP LLMS, SQL SHELLS, OR OTHER ASSISTANTS.**

**Name:**

## I Transactions

Consider the following schedule (here the transactions are numbered by their commit / validate points):

T1	T2	T3
Start TX		
WA		
WC		
		Start TX
		RA
	Start TX	
	RD	
RC		
		RB
(Validate) & Commit		
	WD	
		WA
	WA	
	(Validate) & Commit	
		RB
		(Validate) & Commit

1. [4 points]: Draw the conflict/precedence graph. Does it contain a cycle?  
(Draw your solution in the space below and circle one.)

Cycle    No Cycle

2. [6 points]: Is this schedule:  
(Circle 'T' or 'F' for each choice.)

T   F   View serializable

T   F   Conflict serializable

3. [4 points]: Would this schedule be allowed under (non-strict, non-rigorous) two phase locking?  
(Circle Yes or No.)

Yes    No

Short explanation:

Name:

4. [4 points]: Would this schedule be allowed under strict two phase locking?

(Circle Yes or No.)

Yes No

Short explanation:

Assume the transactions would be executed under snapshot isolation. That is, a write at time  $t$  only modifies a local copy and the write only becomes visible to other transactions after validation and commit. Also assume the transaction reads a consistent snapshot when the transaction starts.

5. [4 points]: Would this schedule be allowed under Snapshot Isolation?

(Circle Yes or No.)

Yes No

Short explanation:

Assume the transactions would be executed under optimistic concurrency control (OCC). That is a write at time  $t$  only modifies a local copy and the write only becomes visible to other transactions after validation and commit. Assume the use of the OCC serial validation protocol.

6. [4 points]: Would transaction T1 commit or abort under OCC?

(Circle Commit or Abort.)

Commit Abort

Short explanation:

7. [4 points]: Would transaction T2 commit or abort under OCC?

(Circle Commit or Abort.)

Commit Abort

Short explanation:

8. [4 points]: Would transaction T3 commit or abort under OCC?

(Circle Commit or Abort.)

Commit Abort

Short explanation:

Name:

## II ARIES

Consider the following log from a system running ARIES:

LSN	Type	Tid	PrevLSN	Data (Page)
1	SOT	1		
2	UP	1	1	A
3	SOT	2		
4	UP	2	3	B
5	UP	2	4	C
Flush				
6	SOT	3		
7	EOT	2	5	
8	UP	3	6	C
CRASH				

9. [8 points]: Assuming there are no log records or checkpoint before LSN 1, and that the Flush operation flushes all dirty pages to disk. What is the state of the xactionTable and dirtyPgTable after the analysis phase?

(Fill in the tables below.)

xactionTable

lastLSN	TID

dirtyPgTable

pgNo	recLSN

Name:

**10. [4 points]:** After the analysis phase, which transactions would have to be rolled back?  
**(Write your answer below.)**

**11. [4 points]:** ARIES does not redo all operations. Which of the following Update LSNs would be redone (e.g., applied to the page) during the redo phase?  
**(Circle 'T' if the operations has to be redone, 'F' otherwise.)**

**T F** LSN 2 (page A)

**T F** LSN 4 (page B)


**T F** LSN 5 (page C)

**T F** LSN 8 (page C)

**Name:**

### III Cardinality Estimation

Consider the following histogram showing the distribution of two attributes A and B in two different tables in T1 and T2.



Both attributes are integers distributed between 1 and 40. Here the histograms represent the number of records in each range of the attribute shown on the X axis.

**12. [4 points]:** Using a uniformity assumption, estimate the selectivity of the predicate  $T1.A > 10$   
(Choose the best answer.)

- A. 0.72
- B. 0.28
- C. 260
- D. 0.86

**13. [4 points]:** Using a uniformity assumption, estimate the selectivity of the predicate  $T2.B \leq 5$   
(Choose the best answer.)

- A. 25
- B. 0.056
- C. 0.11
- D. 0.5

**14. [6 points]:** Using a uniformity assumption, estimate the **cardinality** of the expression  $T1.A = T2.B$  and  $T2.B \leq 10$ . Suppose this is not a key-foreign key join. Show your work.  
(Write your answer in the space below.)

Name:

## IV Distributed Query Execution

Dana Bass is running a key-foreign key equality join between two tables, A and B, using parallel hashing. Table A is the primary key table, and the amount of data involved in the join is 10 GB (1 GB =  $10^9$  bytes), spread across 100 M rows. Table B is 1 GB, spread across 10 M rows, and each row joins with one row of A selected uniformly at random. Each row is the same size. The join is run on machines with disks that read at 1 GB / sec, and a network that transmits at 1 GB/sec between any pair of nodes. Estimate the runtime of the following joins, ignoring setup/teardown costs, but including the network cost to transmit the join result to a single output node.

You can further assume the following:


- The single output node can receive at 1GB / sec from each worker and can send and receive data at the same time (from all workers).
- All attributes from both tables are output (including both copies of the join attribute).
- No data is sent until the join has completed.
- The hash table and other intermediate state for the joins can fit into memory.
- CPU time is negligible.
- If tables are not partitioned on the join attribute, they are repartitioned (rather than being replicated).
- At the start of execution no data is in cache.
- Networking cannot be done concurrently with disk I/O.

**15. [6 points]:** Join on 2 nodes where A and B are both partitioned on the join attribute.  
**(Give your best estimate of the runtime. Show your work.)**

**16. [6 points]:** Join on 4 nodes where A is partitioned on the join attribute and B is not.  
**(Give your best estimate of the runtime. Show your work.)**

**Name:**

## V Eventual Consistency / Dynamo



Consider the above diagram depicting the Amazon Dynamo ring.

**17. [2 points]:** Consider the diagram on the left, without any network partitions and no other failures. If  $N=3$ ,  $W=2$ ,  $R=2$ , which nodes will receive a write of key  $k$  with value = 1 coordinated by node E?  
**(Write your answer in the space below.)**

**18. [4 points]:** Now suppose the use of sloppy quorums and partitioning shown on the right, and no other failures. If  $N=3$ ,  $W=2$ ,  $R=2$ , which nodes will receive a write of key  $k$  with value = 2 coordinated by node D? If the write cannot succeed, indicate this with "DNF".  
**(Write your answer in the space below.)**

**19. [4 points]:** Now suppose the use of sloppy quorums and partitioning shown on the right, and no other failures. If  $N=3$ ,  $W=2$ ,  $R=2$ , which nodes will receive a write of key  $k$  with value = 3 coordinated by node A? If the write cannot succeed, indicate this with "DNF".  
**(Write your answer in the space below.)**

Name:



Consider the sequence of writes to  $k$  from the previous 3 questions (e.g., a write of 1 coordinated by E, without partitioning, a write of 2 coordinated by D while the network is partitioned, and a write of 3 coordinated by A while the network is partitioned.) Suppose that D and A perform a sloppy quorum read of  $k$  before writing it.

**20. [4 points]:** Which of the following represents the state of the vector clock stored at node C after this sequence of writes? Here  $[A, 1][B, 1]$  indicates that the write incorporates version 1 written by A and version 1 written by B.

**(Choose the best option.)**

- A.  $[A, 1]$
- B.  $[E, 1]$
- C.  $[E, 1][D, 1]$
- D.  $[E, 1][D, 2]$
- E.  $[D, 1]$
- F.  $[E, 1][D, 1][A, 1]$
- G. Cannot tell

## VI H-Store

**21. [10 points]:** Consider the H-Store system. Which of the following statements about it are true:  
**(Circle 'T' or 'F' for each choice.)**

- T F** It does not require disk-based REDO because it replicates writes
- T F** It relies heavily on parallel shuffle joins for its efficiency
- T F** It relies on two-phase commit for coordination of distributed transactions
- T F** It uses two-phase locking for concurrency control
- T F** It does not have a buffer pool

**End of Quiz II!**

**Name:**