



Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.5830/31 Database Systems: Fall 2022 Quiz II

There are 13 questions and 11 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; NO PHONES OR INTERNET ALLOWED.**

This exam is worth a total of 100 points.

Your quiz booklet will be scanned for grading. Please *do not* write too close to the edge of the page.

Name:

Please print your full name clearly in the box above.

I Cardinality Estimation

Consider a table with two columns, A and B. In this problem we consider the effect of correlations in these columns on cardinality estimation. Specifically the correlation between the values can range from -1 (maximum negative correlation) to 0 (no correlation / independent), to 1 (maximum positive correlation). We consider two tables:

1. T1 has 10K points with A and B values highly correlated (correlation = 0.9).
2. T2 has 10K points with A and B values not correlated / generated independently (correlation = 0).

For both tables we consider queries that have filters on both the columns, like:

```
SELECT COUNT(*) FROM T1
WHERE A < 'a1'
AND B < 'b1'
```

1. [8 points]: We are interested in how PostgreSQL's cardinality estimation will work in these scenarios. For each of the following statements about PostgreSQL's cardinality estimation for these tables, circle true if the statement is true. Otherwise, circle false.

(Circle 'T' or 'F' for each choice.)

T F PostgreSQL estimates will be more accurate in T1 (data with corr = 0.9) than in T2 (data with corr = 0.0).

T F PostgreSQL will usually under-estimate (as compared to the true values) the size of queries on T1.

Answer: F, T

Name:

II Serializability

Consider the following three transaction schedules. Here, “C” means “Commit”.

Schedule 1			Schedule 2			Schedule 3		
T1	T2	T3	T1	T2	T3	T1	T2	T3
RA					RA			RA
WA			RA			RA		
C				RB				WA
		WA	WA					C
		C	C				RB	
	RA			WB		WB		
	WA			C		C		
	C				RC		WA	
					C		C	

2. [12 points]: For each of the schedules above, indicate whether it has a serial equivalent. If yes, give one such serial schedule as well (e.g., T1, T2, T3).

(Indicate yes or no and give a serial schedule if applicable.)

- i. **Schedule 1:** Has a serial equivalent? (yes/no): _____ Serial schedule (if applicable): _____
- ii. **Schedule 2:** Has a serial equivalent? (yes/no): _____ Serial schedule (if applicable): _____
- iii. **Schedule 3:** Has a serial equivalent? (yes/no): _____ Serial schedule (if applicable): _____

Answer:

i. Yes: T1, T3, T2

ii. Yes: T3, T1, T2 or T3, T2, T1

iii. No

Name:

III Two Phase Locking

Consider the following three transactions which are run concurrently in some unknown interleaving.

T1	T2	T3
RA*	RB	RA
RB*	WB	RC
WA	RD	WC
	WD _x	

* Runs without blocking

x Runs last

Suppose you know that the database executes T1 RA and T1 RB without blocking for another transaction and that the last operation to complete is T2 WD.

3. [12 points]: Which serial equivalent orders are possible under non-strict, non-rigorous two-phase locking:

(Circle all that apply.)

- A. T1, T2, T3
- B. T1, T3, T2
- C. T2, T1, T3
- D. T2, T3, T1
- E. T3, T1, T2
- F. T3, T2, T1

Answer: A, B, E

Name:

IV OCC

Suppose you have a database system running OCC with serial validation. Suppose three transactions with the following read/write sets all finish their read phase at the same time. None of these transactions have been assigned a TID yet, and they can enter the validation phase in any order.

T_x Read set: $\{A, B\}$
Write set: $\{B\}$

T_y Read set: $\{B, C\}$
Write set: $\{C\}$

T_z Read set: $\{D\}$
Write set: $\{A\}$

4. [12 points]: Assuming the system does not re-execute aborted transactions, which of the following commit orders could result? If a transaction does not appear in the listed order, this indicates it had to abort.

(Circle all that apply.)

A. T_x, T_y, T_z

B. T_y, T_x, T_z

C. T_z, T_y, T_x

D. T_x, T_y

E. T_x, T_z

F. T_y, T_z

Answer: B, E, F

Name:

V Logging, Recovery, and ARIES

Consider the following three transactions, executed with non-strict/non-rigorous two-phase locking (these transactions run concurrently but this depiction is not meant to show an interleaving).

T1	T2	T3
RA	RC	RB
WA	WC	WB
RB	RA	
WB		

Suppose you have the following log on disk, in a system running ARIES that crashes after the last log write:

1. BEGIN T1
2. BEGIN T2
3. T1 WA
4. BEGIN T3
5. T3 WB
6. CP
7. COMMIT T3
8. T2 WC

5. [4 points]: Given the log fragment above, what are the possible equivalent serial orders that the system could have resulted for these transactions if the system had not crashed?

(Write your answer in the space below.)

Answer: Assuming there are no aborts, because T3's WB precedes T1's WB, T3 must come before T1 in the equivalent ordering. Because T2's WC comes after T1's WA T2 must come after T1, so the only possible equivalent ordering is T3 - T1 - T2.

Name:

Now suppose the dirty page table in the checkpoint has the following contents:

<u>pgNo</u>	<u>recLSN</u>
A	3

- 6. [8 points]:** Based on the information given above (log and dirty page table), which writes may need to be written to disk during the REDO Phase (i.e., which writes may not have been flushed to disk)?
(Write your answer in the space below.)

Answer: Since WB with Rec LSN 5 is not in the dirty page table, this write must have been flushed prior to the checkpoint. Therefore, LSN 3 and LSN 8 are the only LSNs that could possibly need to be flushed.

- 7. [4 points]:** Based on the information given above (log and dirty page table), which writes need to be removed from disk during the UNDO Phase?
(Write your answer in the space below.)

Answer: The loser transactions are T1 and T2, so their writes need to be UNDOne. These are LSNs 8 and 3 (again!).

Name:

VI Distributed Query Processing

Suppose there are two tables, A and B, which are partitioned across five nodes. We want to join A.a to B.b. A is partitioned on a, but B is not partitioned on b. The total size of A is 250 GB, while the total size of B is 10 GB.

8. [8 points]: How many gigabytes does each node transmit if we repartition B during execution of the join?

(Write your answer in the space below.)

*Answer: $((|T|/n)/n) * (n - 1) = ((10/5)/5) * (4) = 1.6 \text{ GB}$*

9. [8 points]: How many gigabytes does each node transmit if we replicate B during execution of the join?

(Write your answer in the space below.)

*Answer: $(|T|/n) * (n - 1) = ((10/5)) * (4) = 8 \text{ GB}$*

Name:

VII Distributed Transactions and Two Phase Commit

Imagine that during a distributed transaction that uses Two-Phase Commit, the following steps take place:

- The coordinator asks every worker to PREPARE a transaction,
- Every worker force-writes the PREPARE message,
- And the coordinator crashes before taking any further action.

10. [6 points]: Which of the following statements is true?
(Circle one.)

- A. Modifications during the write are now visible to other transactions
- B. The transaction will be committed when the coordinator recovers
- C. The transaction will be aborted when the coordinator recovers
- D. The transaction will be committed if more than half of the workers report the transaction as successful

Answer: C

11. [6 points]: If there are W workers in the distributed database that uses 2PC, what is the smallest possible number of network messages that must be both sent and received before the outcome of the transaction will be "COMMIT", no matter which nodes fail?
(Circle one.)

- A. W
- B. $2W$
- C. $3W$
- D. $4W$
- E. $3W+1$

Answer: $2W$ messages. After the coordinator receives YES (post-FW of PREPARE) from all workers, it can logically commit.

Name:

VIII MapReduce

Consider the following MapReduce program, which computes final grades for all students in a large class.

Note that one library method has been written previously: `parseExamQuestion()` takes as input the text of an answered exam question, and it returns a tuple that contains:

- the name of the student that took the exam (a string)
- the number of points for the question (an integer)
- whether the student's answer was correct (a Boolean)

```
def map(key, val):
    studentName, numPoints, isCorrect = parseExamQuestion(val)
    earnedPoints = 0
    if isCorrect:
        earnedPoints = numPoints
    emit(key=studentName, value=(earnedPoints, numPoints))
```

```
def reduce(key, vals[]):
    for i in range(0, len(vals)):
        ep, np = vals[i]
        totalEp += ep
        totalNp += np
    emit(key, (totalEp/float(totalNp)))
```

Imagine this program is run on input comprising pairs of the form `(studentName, examQuestion)`. (Both are strings.) The cluster has a set of worker machines, plus a coordinator node.

12. [6 points]: How many times will `map()` be invoked?
(Circle one.)

- A. The number of unique `studentNames`
- B. The number of worker nodes
- C. The number of unique questions on the exam
- D. None of the above

Answer: D

Name:

13. [6 points]: How many times will `reduce()` be invoked?
(Circle one.)

- A. The number of unique `studentNames`
- B. The number of worker nodes
- C. The number of unique questions on the exam
- D. None of the above

Answer: A

End of Quiz II!



a black and white drawing of a pair of professors holding a ham

Name: