

CS145 Fall 2020 Homework 3

October 27, 2020

This homework is to help familiarize you with transactions. This homework, as with the rest of the homeworks in the course, is completely optional.

Question 1 [8 points] - Disk vs. Main Memory

You are working as the administrator for a database, and you are responsible for performance tuning. Assume your database receives transactions that modify 4KB of data. Your database server has a disk which can perform sequential I/O at 170MB/s, but when writing random 4KB chunks it only operates at 1MB/s.

For this problem, we will be using the convention that 1000KB = 1MB.

1.1 (2 points)

How many transactions can you process per second if you don't use a transaction log? (Assume transactions modify data in random locations on disk.)

1.2 (2 points)

How many transactions can you process per second if you set your database to only write to the sequential log without ever updating the tables disk?

1.3 (2 points)

A typical memory access takes approximately 100ns, with read/write speeds of 25GB/s. How many transactions can you process per second operating only in memory? (Assume each transaction requires 1 memory access + the time to write 4KB to memory.)

1.4 (2 points)

Imagine you can perfectly scale your database workload across servers (so 100 servers will process transactions 100x faster – in reality they would be somewhat slower). If you had a

cluster of servers durably processing transactions as in part 1.2, how many would it take to match the performance of a single server processing transactions in non-durable memory as in part 1.3?

Question 2 [4 points] - The Problem with Crashes/Aborts

Examine the following transaction:

1. Read AccountA
2. Read AccountB
3. $\text{AccountB} = \text{AccountB} + \text{AccountA}$
4. $\text{AccountA} = 0$

For each step, explain how the output of the program would change if the computer crashes right before that step is executed.

Question 3 [8 points] - ACID

For 3.1 to 3.3, it is sufficient to provide a letter choice without an explanation.

3.1 Atomicity (1 point)

What is **NOT** a possible result from an atomic transaction?

- (A) All changes are made
- (B) Some changes are made
- (C) No changes are made

3.2 Consistency (1 point)

Which of the following examples is related to consistency in transactions?

- (A) A teller looking up a balance at one time should not be allowed to see a concurrent transaction involving a deposit or withdraw from the same account. Only when the withdraw transaction commits successfully and the teller looks at the balance again will the new balance be reported.

- (B) When processing medical tests, there's a whole set of rules that medical professionals have to follow. Doctors and their staff have to fill in a requisition properly. The specimen collection center has to verify that information, take samples and pass everything on to a lab. The lab that performs the tests must ensure that everything is valid before performing the test.
- (C) To transfer funds from one account to another involves making a withdraw operation from the first account and a deposit operation on the second. If the deposit operation failed, you don't want the withdraw operation to happen either.
- (D) A data warehouse can still keep the contents of its database even facing a system crash or other failure.

3.3 Isolation (1 point)

Which of the following statements about isolation is correct?

- (A) To maintain the isolation of transaction, the transactions must be executed serially.
- (B) Isolation guarantees that once a transaction has committed, its effects remain in the database.
- (C) Isolation can be achieved by wisely scheduling the transactions as if they are executed serially.
- (D) Isolation of the transactions can be achieved by Write-Ahead Logging

3.4 Durability (2 points)

Given the following transaction T:

T	R(A)	R(B)	A := A+B	B := A+B	COMMIT
----------	------	------	----------	----------	--------

Let $A = 1$ and $B = 2$ at the start of the execution. We run this transaction on a database that does not guarantee Durability.

If our computer crashes after committing T, what are the possible values of A and B after it restarts? Please list all possible pairs of values for A and B.

3.5 ACID and Scheduling (3 points)

How can transaction scheduling affect the ACID properties of a DBMS? Note all properties that may be affected. Explain by providing an example.

Question 4 [8 points] - Write-Ahead Logging (WAL)

Transaction T performs the following operations:

1. R(A)
2. R(B)
3. $B := B+1$, e.g. W(B)
4. $A := A+1$, e.g. W(A)

4.1 (4 points)

Initially, $A = 0$, $B = 1$, and each write operation increases the record's value by 1. Assume that the main memory and disk both start out empty. Fill out the following table to describe the step-by-step changes to both the data and the log on main memory and disk during each step of T and for each of the 2 steps in which WAL writes the log and data to disk (6 steps total). Note that not all cells may need to be filled out.

Step	MM Data	MM Log	Disk Data	Disk Log
1				
2				
3				
4				
5 (Write to Disk Log)				
6 (Write to Disk Data)				

4.2 (2 points)

Imagine we alter the WAL protocol so that we first commit T, then write the data and log records to disk. Describe which ACID property we lose and give an example why.

4.3 (2 points)

Now, say we alter our WAL protocol so instead it writes the data to disk before the corresponding log record. Describe what ACID property we lose, and give an example why.

Question 5 [6 points] - Conflict Serializability

Consider the following schedules, each involving two transactions T1 and T2. For each schedule, specify whether or not the schedule is conflict serializable and why. If the schedule is conflict serializable, give the equivalent serial schedule.

5.1 Schedule 1 (2 points)

T1		R(B)	W(C)			R(C)
T2	W(B)			R(C)	W(A)	
Timestamp	0	1	2	3	4	5

5.2 Schedule 2 (2 points)

T1	R(B)	W(C)				R(C)
T2			W(B)	R(C)	W(A)	
Timestamp	0	1	2	3	4	5

5.3 Schedule 3 (2 points)

T1	R(A)		R(B)			W(A)		
T2		R(A)		R(C)	W(C)		R(B)	W(B)
Timestamp	0	1	2	3	4	5	6	7

Question 6 [3 points] - Conflict Types

Let's say you are buying a \$2.00 cup of coffee at Coupa Cafe. But just after you've paid and your debit card transaction is being processed, Stanford decides to levy a draconian tax of 2% of both your account and Coupa Cafe's account. Consider the following simplified interleaving of transactions:

T1			Your acct -= \$2.00	Coupa's acct += \$2.00				
T2	Coupa's acct* = .98					Your acct* = .98		
Timestamp	0	1	2	3	4	5	6	7

You can assume that each transaction above consists of a read and a write, in that order, with each read and each write taking up a unit of time. For example, the read in "Coupa's acct* = .98" occurs at timestamp 0, and the write occurs at timestamp 1.

What are all the pairs of actions that conflict and what type of conflict do they form? Identify them by their timestamp.

Question 7 [9 points] - Two-Phase Locking

For each of the following schedules, specify whether the schedule is conflict serializable and whether the schedule is possible / can be produced under strict 2PL. Make sure to explain your answer. Remember that the 2PL protocol grabs an X (exclusive) lock for writing and an S (shared) lock for reading.

7.1 Schedule 1 (3 points)

T1	R(A)	W(A)		R(B)	
T2			W(A)		W(B)
timestamp	0	1	2	3	4

7.2 Schedule 2 (3 points)

T1	R(A)		R(B)	W(B)	
T2		R(A)			R(A)
timestamp	0	1	2	3	4

7.3 Schedule 3 (3 points)

T1	R(A)		R(A)		
T2		W(A)		W(B)	R(B)
timestamp	0	1	2	3	4

Question 8 [4 points] - Deadlocks

In lecture, we saw that 2PL is susceptible to deadlocks. Provide a series of lock requests made by some transactions T1, T2, and T3 over shared resources A, B and C such that a deadlock occurs. Draw the corresponding waits-for graph for your solution.