CIS 556 – FINAL December 14 2022 Total Points: 100

Instructions

This midterm exam covers chapters 12, 16, 17, 18, 19, 20, 21, 22, 23, 24, and 30. There are various types of questions. Attempt all questions and provide the best possible answers. You have 24 hours to complete this test in Canvas. After 24 hours the exam time will be closed automatically and you will not be able to attend it again.

You can use books and Internet Browser:

You cannot share or discuss your answers with any other student in the class.

You cannot copy any materials from the Internet for your answers.

No discussion board conversation allowed during the exam.

How to submit:

You must submit ONLY ONE PDF for all the answers else it will not be accepted, and you will fail the exam

You can draw diagrams by hand but the disgrams must be very clear

For all questions including mutiple choice questions, just include question number and the answers in a single PDF document. Don't include question description.

Upload ONLY ONE PDF containing all your answers to the Exam (similar to what you did for Midterm 2).

There will be no extensions to the deadline.

You can write me an email if you have any questions or concerns pertaining to questions and submission.

Any violations will be strictly dealt with as per course and university policies.

Q1. (10 points)

A disk unit has the following Disk specifications for creating a BOOK database for a Book Library System:

seek time s=20 msec

rotational delay rd=10 msec

transfer rate=1500 bytes/msec

block size B=3000 bytes

interblock gap size G=400 bytes

Blocking factor (bfr) = 28

Number of disk blocks in data file = 1786

The BOOK file has records of fixed-length format, and unspanned blocking. Answer the following questions for the above EMPLOYEE file:

(a) Calculate the record size in bytes (2 points)
(b) Calculate number of records in data file (2 points)
(c) Calculate the wasted space in each disk block because of the unspanned organization (2 points)
(d) Calculate the block transfer time (2 points)
(e) Calculate the average number of block accesses needed to search for an specific record in the ordered file, using linear search. (2 points)

Answer:

Q2. (20 points)

Consider the following SQL query:

```
SELECT r.roomNo, r.type, r.price
FROM Room r, Booking b, Hotel h
WHERE r.roomNo = b.roomNo AND b.hotelNo = h.hotelNo AND
      h.hotelName = 'Grosvenor' AND r.price > 100;
```

(a) Write its relational algebra expression (5 points)

(b) Draw at least two query trees that can represent this query (5 points)

(c) Draw the initial query tree for this query, and then show how the query tree is optimized by the algorithm outlined in Section 18.7. (10 points)

Answer:

Q3. (20 points)

A file of 32768 blocks is to be sorted with an available buffer space of 256 blocks. (5 points)

(a) How many passes will be needed in the merge phase of the external sort-merge algorithm? (10 points)
(b) Calculate the number of disk block reads and writes before sorting of the whole file is completed (5 points)
(c) Calculate the worst-case number of disk block reads and writes with minimum number of main memory buffers of 3 before sorting of the whole file is completed (5 points)

Answer:

Q4. (25 points)

For each of the following 5 schedules, state whether the schedule is serializable, conflict-serializable, recoverable, and whether it avoids cascading aborts. Also draw a precedence graph for each of the schedules. (5 points for each part)

(a) $r_1(X); r_2(X); w_1(X); w_2(X); c_1; c_2$

(b) $r_1(X); r_2(Y); w_3(X); r_2(X); r_1(Y); c_1; c_2; c_3$

(c) $r_1(X); w_2(X); w_1(X); a_2; c_1$

(d) $w_1(X); r_2(X); w_1(X); c_2; a_1$

(e) $r_1(X); w_2(X); w_1(X); r_3(X); c_1; c_2; c_3$

Answer:

*Q*5. (15 points).

Given the following 3 schedules for transactions T7 and T8:

Time	T_7	T_8
t_1	begin_transaction	
t_2	read(bal_x)	
t_3	write(bal_x)	
t_4		begin_transaction
t_5		read(bal_x)
t_6		write(bal_x)
t_7	read(bal_y)	
t_8	write(bal_y)	
t_9	commit	
t_{10}		read(bal_y)
t_{11}		write(bal_y)
t_{12}		commit

(a) Schedule S_1

T_7	T_8
begin_transaction	
read(bal_x)	
write(bal_x)	
	begin_transaction
	read(bal_x)
read(bal_y)	
	write(bal_x)
write(bal_y)	
commit	
	read(bal_y)
	write(bal_y)
	commit

(b) Schedule S_2

T_7	T_8
begin_transaction	
read(bal_x)	
write(bal_x)	
read(bal_y)	
write(bal_y)	
commit	
	begin_transaction
	read(bal_x)
	write(bal_x)
	read(bal_y)
	write(bal_y)
	commit

(c) Schedule S_3

(a) Explain which schedules are serial and which are nonserial (5 points)

(b) Which schedules are equivalent and why? (5 points)

(c) Which schedules are conflict serializable and show its precedence graph? (5 points)

Answer:

Q6. (5 points)

Describe the two methods that gaurantee serializability

Answer:

Q7. (5 points)

Describe 5 deadlock prevention protocols along with their strength and weaknesses.

Answer: