

! This quiz has been regraded; your score was affected.

Midterm Exam

Due Oct 19 at 3:50pm

Points 100

Questions 15

Available Oct 19 at 2:30pm - Oct 19 at 3:50pm about 1 hour

Time Limit 80 Minutes

This quiz was locked Oct 19 at 3:50pm.

Attempt History

	Attempt	Time	Score	Regraded
LATEST	Attempt 1	80 minutes	36 out of 100	82 out of 100

! Correct answers are hidden.

Score for this quiz: **82** out of 100

Submitted Oct 19 at 3:50pm

This attempt took 80 minutes.

Consider the following database schema about flights and airports:

```
Airport (aircode, name, city, state)
```

```
Flights (number, airline, departure-airport, arrival-airport, departure-time, arrival-time)
```

The following constraints hold for the above schema:

aircode is a primary key for Airport

number is a primary key for Flights

departure-airport, arrival-airport are both foreign keys referring to Airport.aircode

Question 1**10 / 10 pts**

Express the following query in SQL: *what are the names of the airports that you can reach with a direct flight from Madison?* You can use the fact that Madison has aircode = 'MSN'.

Your Answer:

```
SELECT DISTINCT a.name
FROM Flights f, Airport a
WHERE f.departure-airport == "MSN"
AND f.arrival-airport == a.aircode;
```

Question 2**10 / 10 pts**

Express the following query in SQL: *how many flight itineraries exist from Madison (aircode = 'MSN') to Seattle (aircode = 'SEA') that have exactly one stop (or equivalently, two legs)?*

Your Answer:

```
SELECT COUNT(*)
FROM Flights f1, Flights f2
WHERE f1.departure-airport == "MSN"
AND f1.arrival-airport <> f1.departure-airport AND f2.arrival-airport <>
f2.departure-airport
AND f1.arrival-airport == f2.departure-airport
AND f2.arrival-airport == "SEA"
AND f1.arrival-time < f2.departure-time;
```

Question 3**8 / 8 pts**

Suppose we define the following view:

```
CREATE VIEW DeltaFlights AS
SELECT DISTINCT departure-airport AS A, arrival-airport AS B
FROM Flights
WHERE airline = 'Delta' ;
```

A round-trip is an itinerary with two legs: from airport 1 to airport 2, and from airport 2 back to airport 1. Write a SQL query that counts the number of round-trips such that both legs are operated by Delta. **Your query must use the above view!**

Your Answer:

```
SELECT COUNT(*)
FROM DeltaFlights d1, DeltaFlights d2
WHERE d1.A == d2.B AND d1.B == d2.A
AND d1.A <> d1.B AND d2.A <> d2.B
AND d1.departure-time < d2.arrival-time;
```

Question 4**7 / 10 pts**

Write a Relational Algebra expression equivalent to the following SQL query. You can use the "Insert Math Equation" feature in Canvas to write the RA in Latex, or embed an image where you have drawn a RA plan.

```
SELECT DISTINCT airline
FROM Flights f, Airport a
```

```
WHERE f. departure-airport = a.aircode AND a.state = 'WI'
```

```
UNION
```

```
SELECT DISTINCT airline
```

```
FROM Flights f, Airport a
```

```
WHERE f. arrival-airport = a.aircode AND a.state = 'WI' ;
```

Your Answer:

Objective: Select all distinct airlines that operates flights in or of Wisconsin airports.

Solution:

$$WIAirports = \sigma_{state='WI'}(Airport)$$

$$Flights = \Pi_{num, airline, d-air, a-air}(\Pi_{number, airline, departure-airport, arrival-airport}(Flights))$$

$$Overall = Flights \bowtie P_{d-air, d-name, d-city, d-state}(WIAirports)$$

$$\bowtie P_{a-air, a-name, a-city, a-state}(WIAirports)$$

$$Ans = \Pi_{airline}(Overall)$$

overall should do a union, and not a join

Incorrect

Question 5

0 / 5 pts

The following two SQL queries will always give the same result (assume no NULLs in the database):

```
SELECT city FROM Airport GROUP BY city ;
```

```
SELECT city FROM Airport ;
```

True

False

Incorrect

Question 6

0 / 6 pts

The following two RA expressions are equivalent:

$\pi_{number}(Flights \bowtie_{aircode=departure-airport} \sigma_{state='WI'}(Airport))$

$\pi_{number}(\sigma_{state='WI'}(Flights \bowtie_{aircode=departure-airport} \pi_{state,aircode}(Airport)))$

True

False

Question 7

5 / 5 pts

The following SQL query is valid:

```
SELECT MIN(departure-time), MAX(arrival-time)
FROM Flights
GROUP BY departure-airport ;
```

True

False

Partial

Question 8

6 / 8 pts

Consider the relation $R(A, B, C, D)$. For which of the following sets of FDs is $\{A, B\}$ a **key** in R ?

$A \rightarrow C, B \rightarrow D$

$A B \rightarrow C$

$A \rightarrow B, B \rightarrow C D$

$B \rightarrow C D$

Question 9

5 / 5 pts

Consider the relation $R(A, B, C, D)$ with the following set of FDs:

$A \rightarrow B$

$B \rightarrow A$

$C \rightarrow D$

$D \rightarrow C$

Recall that the BCNF decomposition algorithm can lead to different decompositions, depending on the FD that we use to split at each step. How many BCNF decompositions can be produced for the above schema?

4

Question 10**5 / 5 pts**

Explain your answer to the above question.

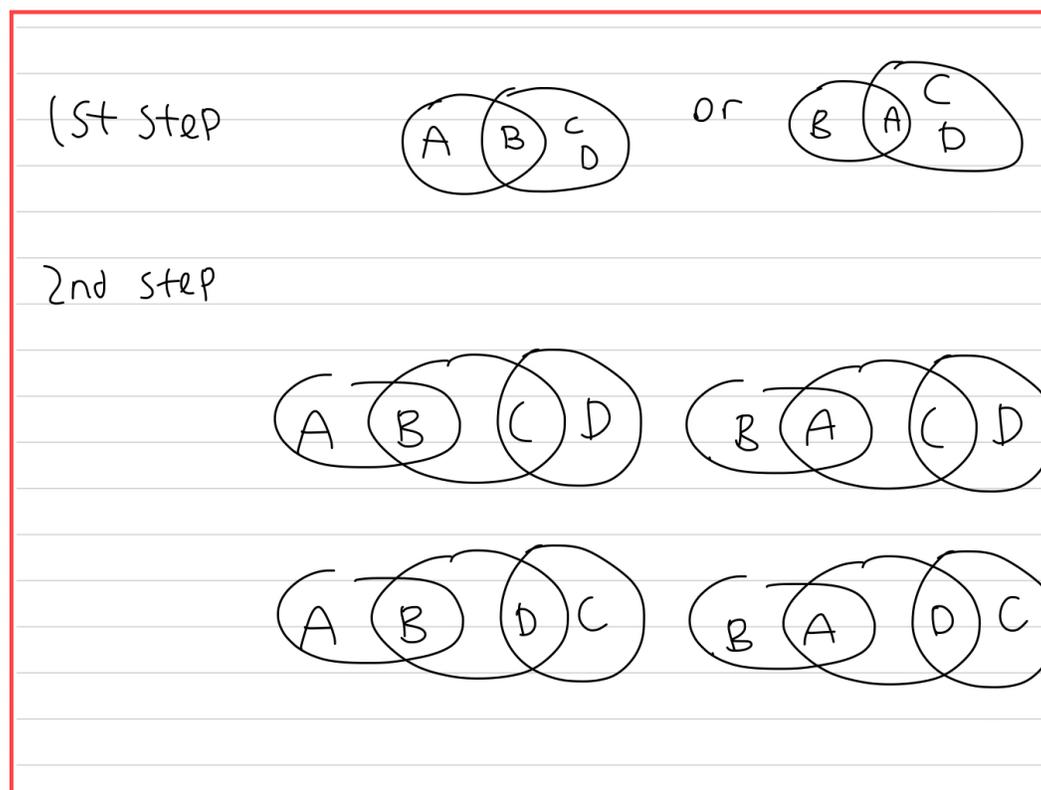
Your Answer:

Q8:

- (1) A, B can infer all other attributes and neither A nor B individually can.
- (2) Can't infer D, not superkey
- (3) Is superkey, but A alone infers all other attributes, so not key.
- (4) B doesn't infer A, not superkey

Q9:

For the original (A, B, C, D), all FDs are violated. Let's look at A & B. When doing decomposition, one of A and B is decomposed into a new relation and the other is left. So right now we have two possibilities. When looking at C & D, two other possibilities are available. Therefore, altogether there are $2 \times 2 = 4$ possibilities. See pic below



Question 11 Original Score: 0 / 5 pts **Regraded Score: 5 / 5 pts**

⚠ This question has been regraded.

Consider the relation $R(A, B, C, D, E)$ with the following set F of FDs:

$A B \rightarrow C$

$A \rightarrow C$

$D \rightarrow E$

The above set of FDs forms a minimal basis for F .

True

False

Question 12**3 / 5 pts**

Explain your answer to the above question.

Your Answer:

$A B \rightarrow A \rightarrow C$. Therefore, $A \rightarrow C$ is redundant as if we remove it then the closure doesn't change.

$A, B \rightarrow C$ does not imply $A \rightarrow C$.

Consider a buffer pool with N frames and a file that consists of 10 pages. Consider the following sequence of page requests:

1, 2, 3, ..., 10, 1, 2, 3, ..., 10.

Each page request is followed directly by a page release (before the next page request). Initially, all buffer frames are free, and none of the pages to be accessed are in RAM.

We say that a page request is a **page hit** if the page is already in the buffer pool. The **buffer hit ratio** is the percentage of page requests that are page hits.

Question 13

6 / 6 pts

Suppose that the buffer manager uses the **LRU replacement policy**. What is the buffer hit ratio when the size of the buffer pool has $N = 10$ frames?

Question 14

6 / 6 pts

Suppose that the buffer manager uses the **LRU replacement policy**. What is the buffer hit ratio when the size of the buffer pool reduces to $N = 9$ frames?

Question 15**6 / 6 pts**

Explain your answer to the above question.

Your Answer:

N = 10: First 10 accesses are all misses. The last 10 access are all hits.
 $10/20 = 0.5$.

N = 9: First 9 accesses are all misses. When accessing page 10, page 1 was least recently used, so it's ejected. Then, when reading page 1, page 2 is ejected, and so on. Finally, when accessing page 9, page 10 is ejected. When accessing page 10, page 1 is ejected. Thus, all requests lead to buffer misses, and the hit ratio is 0.

Quiz Score: **82** out of 100