

CS 564 Midterm Exam

Fall 2016

Answers

A: ER AND RELATIONAL MODELS [12%]

For the following questions, just answer (by circling) either True or False.

1. [3%] The ER model can not model relationships between more than two entity sets.
FALSE
2. [3%] A relationship in the ER model is always translated to its own separate relation in the Relational model.
FALSE
3. [3%] There is no way to express foreign key constraints in the ER model.
TRUE
4. [3%] In an ISA hierarchy, a subclass entity set always inherits the key from its superclass entity set.
TRUE

B: NORMALIZATION AND DEPENDENCY THEORY [30%]

Consider the following relation that contains information about airline reservations:

Reservations (rid, pid, p_fname, p_lname, agent_id, agent_name, rdate).

with the following set of functional dependencies:

rid \rightarrow pid, agent_id, rdate
pid \rightarrow p_fname, p_lname
agent_id \rightarrow agent_name

Answer the following questions.

1. [20%] Choose by circling the right option(s) and explain the reasoning behind your answer. If your explanation is invalid or incomplete, you will be awarded only half the points.

(a) [10%] The following attribute sets are *superkeys* in **Reservations**:

- (i) agent_id, pid (ii) rid, rdate (iii) rid (iv) agent_id, pid, rdate

ANSWER: (ii) and (iii)

EXPLANATION:

The attribute {rid} is a key for the relation, since it functionally determines all the attributes in the relations. Since {rid, rdate} is a superset of {rid}, it is a superkey as well. The remaining choices do not functionally determine all attributes, so they are not superkeys.

(b) [10%] The following FDs are *logically implied* by the set of FDs that hold for the attributes in **Reservations**:

- (i) agent_id, rdate \rightarrow rid (ii) rid \rightarrow p_fname, agent_name

- (iii) pid \rightarrow p_fname, rdate (iv) pid, agent_id \rightarrow agent_name

ANSWER: (ii) and (iv)

EXPLANATION:

Since rid is not implied by any FD, (i) cannot be correct. (ii) is correct because rid is a key, so it functionally determines all attributes. (iii) is wrong since pid cannot infer rdate. Finally (iv) is correct because the FD agent_id \rightarrow agent_name holds.

2. [10%] Perform BCNF decomposition of the relation **Reservations** using the FDs provided. Name your decomposed relations appropriately, and specify the keys for each decomposed relation. Write and explain all steps in your decomposition.

ANSWER: a possible BCNF decomposition of the relations is as follows:

ReservationsNew(rid, pid, agent_id, rdate)

Passenger(pid, p_fname, p_lname)

Agent(agent_id, agent_name)

C: SQL & RELATIONAL ALGEBRA [38%]

We will be using the following database schema for part C:

Wine (wid, wname, price, color, makerid)

Wine.makerid is a foreign key referring to Winemaker.makerid

Winemaker (makerid, name, countrycode)

Winemaker.countrycode is a foreign key referring to Country.countrycode

Country (countrycode, countryname, population)

1. [10%] Express the following query in Relational Algebra: *what are the names of the winemakers that produce **no** white wines?*

$$\pi_{name}(\pi_{makerid,name}(Winemaker) - \pi_{makerid,name}(\sigma_{color=white}(Wine) \bowtie Winemaker))$$

2. [19%] Write one SQL query for each of the following questions:

- (a) [9%] Output the **number** of (distinct) winemakers that produce wines with price more than \$100.

```
SELECT    COUNT (DISTINCT makerid)
FROM      Winemaker
WHERE     price > 100;
```

- (b) [10%] For each country, output the **name** of the country and the **number** of *red* wines it produces.

```
SELECT    c.countryname, COUNT(w.wid)
FROM      Wine w, Winemaker m, Country c
WHERE     w.makerid = m.makerid AND m.countrycode = c.countrycode
AND      w.color = 'red'
GROUP BY  c.countrycode, c.countryname ;
```

The answer that groups by c.countryname only will also receive full points.

3. [9%] For the following questions, just answer (by circling) either True or False.

- (a) [3%] Creating a view increases the size of the data stored in the database.

FALSE

- (b) [3%] Suppose that we define the schema of **Winemaker** in SQL as follows:

```
CREATE TABLE Winemaker (
    makerid      INTEGER,
    name         CHAR(20),
    countrycode  INTEGER,
    FOREIGN KEY (countrycode) REFERENCES Country(countrycode)
                ON DELETE CASCADE);
```

When we update the countrycode of a tuple in **Country**, the database will update the corresponding countrycode in the **Winemaker** table.

FALSE

- (c) [3%] The following query will return at least one tuple if the relation **Wine** is not empty:

```
SELECT    wname
FROM      Wine
WHERE     color = 'red' OR color <> 'red' ;
```

FALSE

D: DISKS AND BUFFER MANAGEMENT [20%]

1. [8%] You are given a hard disk (HDD) with the following properties:

RPM: 10,000

maximum seek time: 15 ms

transfer rate: 100 MB/s

What is the maximum time that the hard disk needs to read a single block of 8KB? Explain your computation in detail.

ANSWER: 21.08 ms.

EXPLANATION:

The time to read the block consists of: (a) the maximum seek time (15ms), (b) the maximum rotational delay, which can be one full rotation (6 ms), (c) the transfer time, so time it takes to read the 8KB block ($8\text{KB}/100\text{ MB/s} = 0.08\text{ms}$). Summing up, we obtain 21.08 ms.

2. [12%] In this question, you have to count the number of page I/Os (reads and writes) for a given page access sequence for a buffer manager that uses the LRU replacement policy.

There are 2 query processes. A page "Request(*i*)" means the query process $i = 1, 2$ wants to read that page's values for computations. A page "Modify(*i*)" means the query process $i = 1, 2$ is modifying the values on that page. A page "Release(*i*)" means the query process $i = 1, 2$ has finished using that page and notifies the buffer manager accordingly.

The number of page frames in the buffer pool is 3. Initially, all buffer frames are free and clean, and none of the pages to be accessed are in RAM.

Access Sequence: Request(1) A, Request(2) A, Release(1) A, Request(1) B, Request(2) C, Modify(1) B, Release(1) B, Request(2) D, Modify(2) A, Release(2) A.

What is the total number of page I/Os (number of pages transferred from the disk to the RAM, and from RAM to the disk)? Explain each page I/O in your answer.

ANSWER: 5 page I/Os.

EXPLANATION:

Buffer frames F1, F2, and F3. Initially, all of their pincount and dirtybit are 0.

Request(1) A: I/O to read A into F1. Set F1.pincount = 1.

Request(2) A: A is already into F1. Set F1.pincount = 2.

Release(1) A. Set F1.pincount = 1.

Request(1) B: I/O to read B into F2. Set F2.pincount = 1.

Request(2) C: I/O to read C into F3. Set F3.pincount = 1.

Modify(1) B: Set F2.dirty = 1.

Release(1) B: Set F2.pincount = 0.

Request(2) D: Need to find a frame with pincount 0 to replace. F2 is the only available. But F2.dirty = 1. Thus, I/O to write (flush) B in F2 to disk. Set F2.dirty = 0. I/O to read D into F2. Set F2.pincount = 1.

Modify(2) A: Set F1.dirty = 1.

Release(2) A. Set F1.pincount = 0.