Database Management (COIY028H6) – 2021

Model Answers

1. (a) i. It is assumed that Number values are not globally unique, in other words two cars could have the same RepairJob Number assigned to them.

(2 marks)

ii. A. The attributes would be License, Number and Description, with License and Number comprising the primary key.

(4 marks)

B. The attributes would be NI#, License, and Number, with all of them comprising the primary key.

(4 marks)

(b) i. (FlightNumber, FlightDate, PassportNumber) is the only key; With the closure algorithm, we can add DepAirport and ArrAirport from the first FD, SeatNo from the second, and Price from the third.

(2 marks)

ii. The prime attributes are FlightNumber, FlightDate and PassportNumber. The first and third FDs do not satisfy 3NF because the left sides are not superkeys and the right sides are not prime. The second FD satisfies 3NF because the left side is a superkey.

(3 marks)

iii. A 3NF decomposition would comprise the relation schemas (FlightNumber, DepAirport, ArrAirport), (FlightNumber, FlightDate, PassportNumber, SeatNo) and (FlightNumber, SeatNo, Price).

(3 marks)

iv. Yes, each is in BCNF because the left side of each applicable FD is a superkey in its schema.

(2 marks)

2. (a) i. If the foreign key references table R, it means that when a row in R is deleted, any values in column C which match the primary key value of the row are set to null.

(3 marks)

ii. If C is part of the primary key in table T, then its value could not be set to null because of the "entity integrity" constraint of the relational model.

(2 marks)

(b) i. Two from Round $\rightarrow$ any, Player_1 $\rightarrow$ Player_2, Player_1 $\rightarrow$ Score, Score $\rightarrow$ Player_1, and Score $\rightarrow$ Player_2.

(2 marks)

ii. Two from Player_2 $\rightarrow$ any, Score $\rightarrow$ Round, and Round Player_1 $\rightarrow$ Player_2 Score.

(2 marks)

iii. Player_2 $\rightarrow$ any would hold because a loser can only appear once in column Player_2. Score $\rightarrow$ Round would not hold because the same score could occur in matches in different rounds. Round Player_1 $\rightarrow$ Player_2 Score would hold because in each round, the winner of a match plays only a single opponent resulting in a single score.

(4 marks)

(c) i. Serial execution means that one transaction is executed in its entirety after the other. A serialisable execution is one in which the outcome is the same as some serial execution.

(3 marks)

ii. The two transactions could both be trying to book an unoccupied seat on a flight. Assume that they both find the same seat unoccupied, the first transaction books the seat but it is then overwritten by the second transaction. Then the execution of the first is not isolated from that of the second because if it were, the first would have booked the seat.

(4 marks)

3. (a) i. unknown

(1 mark)

ii. unknown

(1 mark)

(b) i. For a view to be updatable, the associated query must not contain expressions in the select clause.

(2 marks)

ii. The system could simply divide the value given for Euros by 1.15, and insert ('OnIT', 'Widget', 100) into the Sells table.

(2 marks)

(c) Nested user-defined types such as a name comprising a first name and last name, and data types such as arrays and multisets which allow multiple values to be stored in a single column.

(3 marks)

(d) i. With the `query` method, malicious user input could simply be appended to a partial query in the PHP program. With the `prepare` method, placeholders are used where user input is expected. If user input is not of the expected form, the system will return an error.

(5 marks)

ii. The system would first translate the SQL query into a canonical tree in relational algebra. This tree would be transformed using logical optimisation into one that would be expected to be more efficient to evaluate, by performing steps such as moving selections down the tree. After this, physical optimisation would take place in order to decide which algorithm should be used to evaluate each operator in the tree.

(6 marks)

4. (a) i. Use of `except` requires that the two subqueries are defined over the same attributes.

(2 marks)

ii. The query asks for people who have not taken any course.

(2 marks)

iii. Two alternative queries are

```
select id from person
where id not in (select person from taken)
```

and

```
select id from person left join taken on id=person
where person is null
```

(4 marks)

iv. The system could either scan the `person` table or the index on `id`, performing a lookup on the `person` index to remove any `id` that matches.

(2 marks)

v. For A, the answer would be m. This is because no row in `taken` can have a null value for `person` because it is part of the primary key. Because of the foreign key, every row in `taken` must join with exactly one row from `person`.
For B, the answer would be m plus the number of rows in `person` which do not match a row in `taken`.

(4 marks)

(b) For the closure of AC, we start with AC, can add B from the first FD, and can add D from the third FD, giving $ABCD$. For the closure of BE, we start with BE, can add A from the second FD, can add C from the fifth FD, and can add D from the third FD, giving $ABCDE$. So BE is a superkey, but neither B nor E is a superkey, so BE is a key.

(6 marks)

