



Session ?? Examination

ITC114 / ITC423 Database Management Systems

SAMPLE EXAM SOLUTIONS

Question Paper MAY NOT be retained by the Candidate

WRITING DURING READING TIME IS PERMITTED ON ALL EXAMINATION MATERIALS

LECTURER: Peter White

WRITING TIME: 2 hours 10 minutes

MATERIALS SUPPLIED BY UNIVERSITY:

1 x 12 page answer book

1 x General Purpose Answer Sheet

MATERIALS PERMITTED IN EXAMINATION:

2B Pencil/Eraser, pens, ruler

NUMBER OF QUESTIONS: Five (5)

VALUE: 100 marks

INSTRUCTIONS TO CANDIDATES:

1. Enter your name and student number and sign the space provided at the bottom of the page.
2. This is a closed book examination, therefore no written material, reference books or notes will be permitted in the examination room.
3. The examination paper, together with all answer booklets, must be returned at the completion of the examination.
4. Question 1 consists of 20 multiple choice questions, each worth 1 mark. The answers to these must be marked on the GPAS supplied with a 2B pencil.
5. Questions 2 - 5 consist of written questions. The number of marks allocated to each question is indicated. Write your answers in the answer booklet provided.
6. You should answer ALL questions.

STUDENT NAME: STUDENT NUMBER:
STUDENT SIGNATURE:
Calculator model (if used):.....

Question 1: MULTIPLE CHOICE (20 marks)

Answer on the General Purpose Answer Sheet.

For each question select the single best response from the options given.

Attempt all questions.

All questions have equal value (1 mark each).

1. Data ____ is defined as "the condition in which all of the data in the database are consistent with the real-world events and conditions."
 - a. redundancy
 - b. inconsistency
 - c. **integrity**
 - d. verification
2. Data ____ exists when it is possible to make changes in the data storage characteristics without affecting the application program's ability to access the data.
 - a. integrity
 - b. **independence**
 - c. inconsistency
 - d. mining
3. ____ independence exists when it is possible to make changes in the file structure without affecting the application program's ability to access the data.
 - a. logical
 - b. **physical**
 - c. fragmentation
 - d. structural
4. The ____ provides a detailed description of all tables found within the user/designer-created database.
 - a. index
 - b. primary key
 - c. **data dictionary**
 - d. data set
5. Which of the following is true of the entity integrity rule?
 - a. **all primary key entries must be unique**
 - b. a part of the key may be null
 - c. foreign key values cannot reference primary key values
 - d. duplicate object values are allowed
6. A primary key that consists of more than one attribute is called a ____ key.
 - a. foreign
 - b. secondary
 - c. group

- d. **composite**
7. The referential integrity rule requires that ____.
- every null foreign key value must reference an existing primary key value
 - it makes it possible for an attribute to have a corresponding value
 - every non-null foreign key value must reference an existing primary key value**
 - it makes it possible to delete a row in one table whose primary key does not have a matching foreign key value in another table
8. The one-to-many (1:M) relationship is easily implemented in the relational model by putting the primary key of the "1" side in the table of the "many" side as a ____ key.
- composite
 - natural
 - join
 - foreign**
9. A(n) ____ is an orderly arrangement used to logically access rows in a table.
- entity
 - index**
 - join
 - key
10. A(n) ____ links tables on the basis of an equality condition that compares specified columns of each table.
- equijoin**
 - natural join
 - left outer join
 - right outer join
11. In BCNF, every ____ in a table is a candidate key.
- determinant**
 - entity
 - primary key
 - atomic attribute
12. A ____ key is an artificial PK introduced by the designer with the purpose of simplifying the assignment of primary keys to tables.
- surrogate**
 - composite
 - candidate
 - foreign
13. A table in ____ contains no transitive dependencies.
- 1NF
 - 2NF
 - 3NF**
 - none of the above

14. All relational tables satisfy the ____ requirements.
- a. **1NF**
 - b. 2NF
 - c. BCNF
 - d. 3NF
15. The objective of ____ control is to ensure the serializability of transactions in a multiuser database environment.
- a. integrity
 - b. transaction
 - c. **isolation**
 - d. concurrency
16. Two-phase locking guarantees serializability, but it does not prevent ____.
- a. **deadlocks**
 - b. inconsistency
 - c. lost updates
 - d. uncommitted data
17. ____ means that the data used during the execution of a transaction cannot be used by a second transaction until the first one is completed.
- a. Atomicity
 - b. Consistency
 - c. Durability
 - d. **Isolation**
18. Deadlocks are possible only when one of the transactions wants to obtain a(n) ____ lock on a data item.
- a. binary
 - b. **exclusive**
 - c. shared
 - d. complete
19. The ____ administrator controls and supervises the DBMS.
- a. data
 - b. **database**
 - c. program
 - d. system
20. A(n) ____ data dictionary is automatically updated by the DBMS with every database access.
- a. re-entrant
 - b. **active**
 - c. passive
 - d. static

Question 2: 25 Marks

Answer all questions

1. Explain how a database designer uses business rules to identify entities and relationships. (5 marks)

- Describe how entities are identified using business rules (use of nouns)
- Describe how relationships between those entities are identified (use of verbs)
- Describe how the entities are re-evaluated to find further entities or to even discard them as attributes

2. Describe the difference between a Wait/ Die and a Wound/Die lock system in an RDBMS. (5 marks)

- **Wait/die:**
 - If the transaction requesting the lock is the older of the two transactions, it will wait until the other transaction is completed and it releases its locks
 - If the transaction requesting the lock is the younger transaction, it will die (roll back) and is rescheduled using the same time stamp
 -
- **Wound/wait :**
 - If the transaction requesting the lock is the older of the two transactions, it will pre-empt (wound) the younger by rolling it back. The pre-empted transaction is rescheduled using the same time stamp
 - If the transaction requesting the lock is the younger transaction, it will wait until the other transaction is completed and it releases its locks

3. Describe the three data anomalies that can occur with lower-normal-form tables. Provide examples of each anomaly using the Repayments table. (5 marks)

Insertion anomaly: To enter a new entity customer you need to enter a receipt.

Deletion anomaly: When you delete the last receipt for the customer you also delete the customer details.

Update anomaly: To update the customer name you need to update each row that contains the customer.

4. Answer the following questions using the Repayments table.

- a) Identify all the functional dependencies. (5 marks)
- b) What is the primary key of the table? (1 mark)
- c) What is the normal form of the table? (2 marks)
- d) Describe why the table has the chosen normal form. (2 marks)

Repayments

customer id	receipt no	amount received	customer name	customer address	receipt date
Z23	35	\$201.30	Joe	Bathurst	3 Aug 2009
Y72	39	\$173.12	Vicki	Wagga	15 July 2008
Z23	27	\$8.60	Joe	Bathurst	12 Sep 2009
X11	32	\$93.92	Joe	Albury	22 Feb 2007
Y72	41	\$11.91	Vicki	Wagga	7 Apr 2008

Each customer has one address and there is one customer for each account. The receipt is made out to the customer when a payment is received. Each receipt is for one customer only. Each customer has one address.

a. Functional dependencies

customer_id → customer_name, customer_address

receipt_no → amount_received, receipt_date, customer_id

b. What is the primary key of the table.

receipt_no

c. What is the normal form

2NF

d. Why

2NF as it is in first normal form(no repeating groups and identified PK), and all non-prime attributes are fully functionally dependent on receipt_no.

Not in 3NF because it contains a transitive dependency between non-prime attributes (customer_id → customer_name, customer_address)

Question 3 (20 marks)

The Thirlmere Thoroughbreds Database contains the following tables.

HORSE (Horse_id, Name, Colour, Sire, Dam, Born_Date, Died_Date, Gender)
PRIMARY KEY Horse_Id
FOREIGN KEY (Sire) REFERENCES TABLE Horse
FOREIGN KEY (Dam) REFERENCES TABLE Horse

SHOW (Show_Id, Name, Show_Address)
PRIMARY KEY Show_Id

JUDGE (Judge_Code, Name, Judge_Address)
PRIMARY KEY Judge_Code

EVENT(Event_id, Show_Id, Name, Judge_code)
PRIMARY KEY Event_Id
FOREIGN KEY (Show_Id) REFERENCES TABLE Show
FOREIGN KEY (Judge_Code) REFERENCES TABLE Judge

ENTRY (Event_Id, Horse_id, Place)
PRIMARY KEY (Event_Id, Horse_Id)
FOREIGN KEY (Event_Id) REFERENCES TABLE Event
FOREIGN KEY (Horse_Id) REFERENCES TABLE Horse

PRIZE (Event_id, Place, Money)
PRIMARY KEY (Event_Id, Place)
FOREIGN KEY (Event_Id) REFERENCES TABLE Event

Notes on the data:

- The mother of a horse is called its Dam.
- The father of a horse is called its Sire.
- The gender of a horse is recorded as “M” for mare (ie female) or “S” for stallion (ie male).
- Live horses have no entry in the Died_date attribute
- Place is numeric 1, 2, 3 etc. Place is null for events that have not been held.

Notes on the requirements:

- Remove duplicate rows from your results where appropriate.
- Do not use VIEWS or the INCLUDE statement.

Write SQL SELECT statements to retrieve the following information from the database.

1. List all data from the HORSE table about horses who are still alive. Display the youngest horse first. (4 marks)
2. For every event, list the name and id of the event and the ID of the horse that was first. (4 marks)

3. For each event with at least two prizes, list, in descending frequency (with the event having the most prizes first) the event_id, show_id, number of prizes, and total amount of prize money available. (4 marks)
4. For every entry in each event, list the name of the show, name of the event, name of the horse, name of the judge. (4 marks)
5. List the name of each mare, together with the names of all her “children” in the following format: (4 marks)

Daffodil is the mother of Tulip
 Sally is the mother of Blaze
 Sally is the mother of Flash

1.

```
SELECT *
FROM horse
WHERE died_date IS NOT NULL
ORDER BY born_date ASC
```
2.

```
SELECT event.event_id, event.name, entry.horse_id
FROM event, entry
WHERE event.event_id = entry.event_id
AND entry.place = 1
```
3.

```
SELECT event.event_id, event.show_id, count(*) as “number of prizes”, sum(money)
FROM event, prize
WHERE event.event_id = prize.event_id
GROUP BY event.event_id, event.show_id
HAVING 'number of prizes' > 2
```
4.

```
SELECT show.name, event.name, horse.name, judge.name
FROM show, event, entry, horse, judge
WHERE show.show_id = event.show_id
AND judge.judge_code = event.judge_code
AND event.event_id = entry.event_id
AND horse.horse_id = entry.horse_id
```
5.

```
SELECT mother.name, “is the mother of “, child.name
FROM horse mother, horse child
WHERE mother.horse_id = child.dam
```


QUESTION 4 (15 marks)

Case Study:

You are required to:

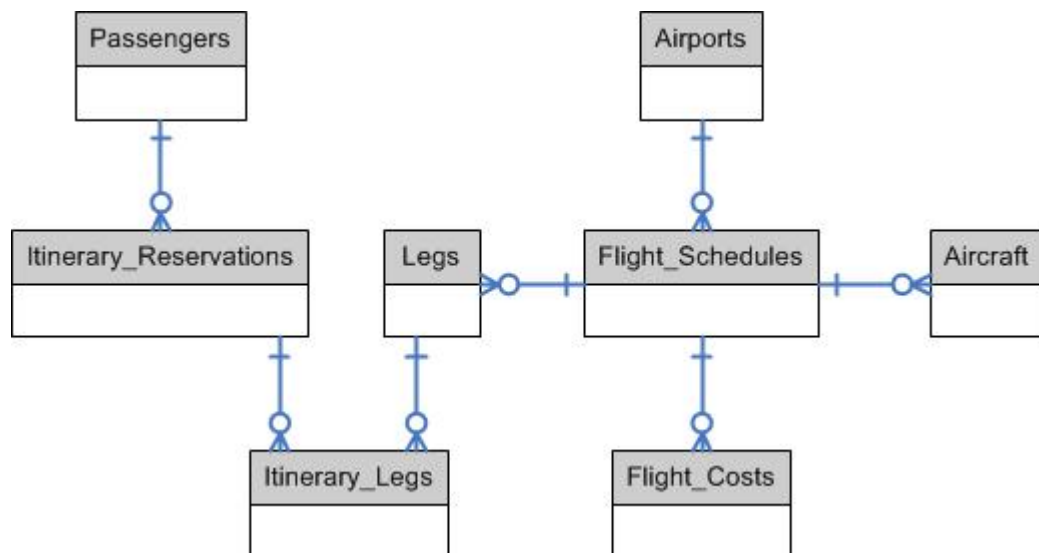
1. Draw an Entity Relationship Diagram (ERD) that represents this case study.

Central West Airways (CWA) is headquartered in Bathurst NSW and flies to many regional and capital city destinations in Central and Eastern Australia. CWA has a fleet of 5 types of aircraft that range in size from 10 seat planes to 90 seat aircraft. Aircraft of either 50 or 90 seat capacity have both Business and Economy class seats. Flights are scheduled to most destinations on a daily basis, but each flight may consist of zero or many legs to intermediate stops before reaching the destination.

CWA makes all reservations on the CWA website and does not use booking agents. Each passenger must make a reservation for a flight to a destination. Each passenger must hold a valid ticket for flight to their destination on the specified date. Passengers must select their date of flight, Flight number, class of travel (Business or Economy), and request a seat allocation at the time of reservation. When a passenger has completed their reservation, an itinerary of their flight(s) must be emailed to them.

Notes:

- The third party payment site is out of scope for this question and can be ignored.
- The ERD must show names for both entities and relationships.
- The ERD must indicate the cardinality (ie. one-to-one, one-to-many or many-to-many) of all relationships.
- The ERD does not need to show attributes.
- List any assumptions that you make.



Question 5 (25 marks)

Draw an Extended Entity Relationship Diagram (EERD) for Hare and Tortoise case study described below. Note:

- Your Entity Relationship Diagram must use Crows Foot notation.
- Give meaningful names to all entities and relationships in your EERD
- Show all attributes in the EERD
- Explicitly indicate the minimum and maximum cardinalities of all relationships (the 'one' or 'many' symbols) in your EERD. Do not omit the symbol at the 'one' end, thinking that it will be assumed.
- List any assumptions that you have made.

Hare and Tortoise Case Study

The Hare and Tortoise is a chain of noodle and dumpling restaurants in London serving Asian influenced meals. The following information has been gathered about the restaurant.

There are several restaurants, for each you need to record the restaurant location (Bloomsbury, Putney, Ealing, Kensington and Blackfrairs) and the address and telephone number.

The chain has a standard range of dishes that are produced at each restaurant. Each dish is described by the dish code, dish name, description (eg R2, 'Malaysian Chicken Curry and Rice', 'chicken on the bone stewed ...'). This information is displayed on the menu available at each restaurant.

Each restaurant is free to construct its own menu choosing from the standard dishes and to set the price of the dish at the restaurant. A sample menu with the menu dishes available at the Bloomsbury locations is shown below.

Hare and Tortoise: Bloomsbury		
Noodle and Dumpling Bar		
Menu		
R2	Malaysian Chicken Curry and Rice chicken on the bone stewed with traditional spices, curry leaves, lemon grass.	£4.50
T3X	Char Siu Pork marinated fillet of pork, oven roasted, and braised with cane syrup.	£4.30
S2C	Kani Salad snow crab leg and green leaves with home	£5.50

	made Japanese dressing.	
M3	Futomaki	£5.00
	big seaweed roll filled with omelette, avocado, crab stick, cucumber, pickle and sakura denbu - 8pcs	
C1	Yaki Udon	\$6.00
	wok fried thick white noodles with prawns, squid, pork, crab sticks, Chinese mushrooms, onions and beansprouts.	

Figure 1 Extract from the Menu

Orders are taken by a waiter recording the dish(es), the table number and takes the order to the cashier. The cashier enters the order information for each table, recording the table number, the dishes ordered and the waiter that serves the table. The system will record the date and the shift. The order is sent to the kitchen for preparation.

A sample order (figure 2) for five people is shown below. A dish ordered twice is repeated on the order.

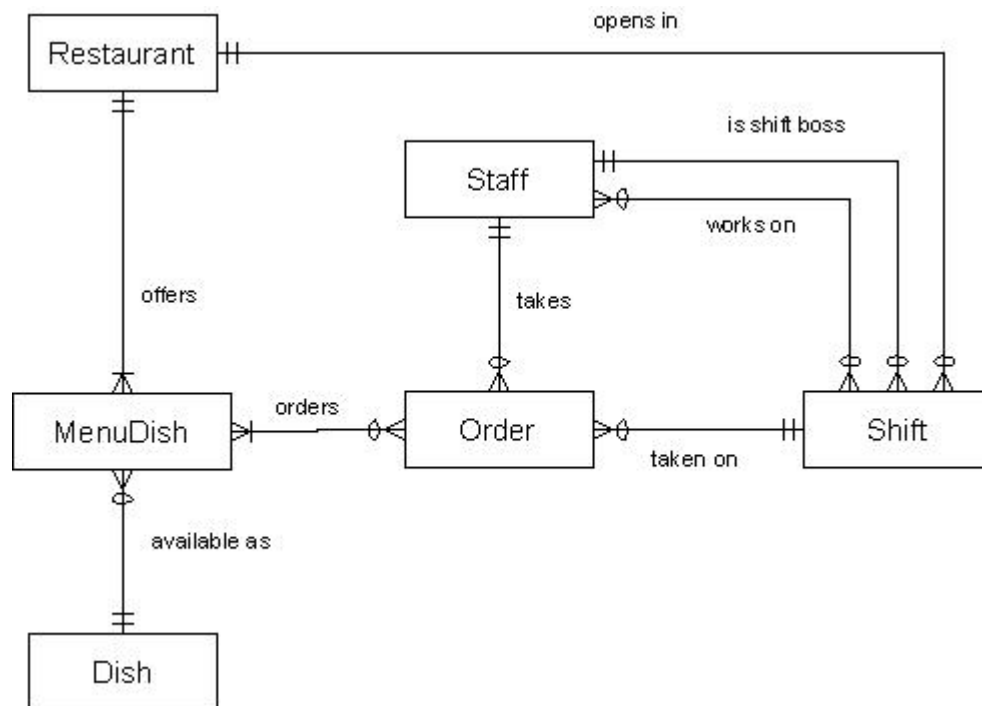
Hare and Tortoise : Bloomsbury		
Noodle and Dumpling Bar		
15 / 17 Brunswick Shopping Centre		
London WC1 Tel 020 7278 4945		
Date: 11-09-2009 Shift: Lunch		
Table : 12 Waiter : ES		
Code	Description	Price
R2	Malaysian Chicken Curry and Rice	£4.50
R2	Malaysian Chicken Curry and Rice	£4.50
T3X	Char Siu Pork	£4.30
SC2	Kani Salid	£5.50
C1	Yaki Udon	£6.00
Total:		£24.80

Figure 2 A sample order

Staff are employed on a casual basis and may work at any of the restaurants in the chain. Each staff member has an id, a unique set of initials (ES in the above figure) and a name.

Organizing the staff into the two shifts offered each day at each restaurant is a difficult task. Responsibility for this task falls to one staff member, the shift boss on each shift.

The system needs to record for each shift offered by each restaurant the staff that work on the shift and the shift boss. The two shifts are lunch: 11:30 – 3:00 and evening: 5:00 – 10:00. It is sufficient to record only Lunch and Evening.



Restaurant (restaurant_id, location, address, phone_number)
Primary Key restaurant_id

Dish(dish_code, dish_name, dish_description)
Primary Key dish_code

MenuDish(restaurant_id, dish_code, dish_price)
Primary Key (restaurant_id, dish_code)
A surrogate key could be introduced – but then the natural key must be listed as an alternate key.
Foreign Key restaurant_id References Restaurant
Foreign Key dish_code References Dish

Staff(staff_id, staff_initials, staff_name)
Primary Key (staff_id)
Alternate Key staff_initials

Shift(shift_id, shift_date, shift_code, restaurant_id, shift_boss_staff_id)
Primary Key shift_id

Foreign Key restaurant_id References Restaurant
Foreign Key shift_boss_staff_id References Staff
Alternate Key (restaurant_id, shift_date, shift_code)
shift_code is either Lunch or Evening

StaffShift(staff_id, shift_id)

Primary Key staff_id, shift_id
Foreign Key c staff_id References Staff
Foreign Key shift_id References Shift

Order(order_id, shift_id, waiter_staff_id, table_number)

Primary Key order_id
Foreign Key shift_id References Shift
Foreign Key waiter_staff_id References Staff

OrderDish(order_dish_id, order_id, restaurant_id, dish_code)

Primary Key order_dish_id
Foreign Key order_id References Order
Foreign Key (restaurant_id, dish_code) References MenuDish