


Question 1

Correct

Marked out of 4.00



The inventory system of a Switch game card rental store tracks the information of customers, cards, and rental transactions between customers and cards. Customers are identified with their names, telephone numbers, and unique IDs, while each game card is associated with a short introduction and one unique ID. To better track the availability of game cards, the store needs to record the borrow time, agreed return time, and final return time of each transaction.

Below is an ER model to describe the above scenario.

Fill in the gaps with the following words. Words can be reused multiple times or not be used at all.

Note: If multiple words can be filled in gaps without distinction, please place them in alphabetical order.

Entity E1 : ✓

Key K1 : ✓

Attribute A1 : ✓

Attribute A2 : ✓

Entity E2 : ✓

Key K2 : ✓

Attribute A3 : ✓


Relationship R1: ✓

Attribute A4 : ✓

Attribute A5 : ✓

Attribute A6 : ✓

The correct answer is:



The inventory system of a Switch game card rental store tracks the information of customers, cards, and rental transactions between customers and cards. Customers are identified with their names, telephone numbers, and unique IDs, while each game card is associated with a short introduction and one unique ID. To better track the availability of game cards, the store needs to record the borrow time, agreed return time, and final return time of each transaction.

Below is an ER model to describe the above scenario.

Fill in the gaps with the following words. Words can be reused multiple times or not be used at all.

Note: If multiple words can be filled in gaps without distinction, please place them in alphabetical order.

Entity E1 : [Customers] Key K1 : [Customer ID]

Attribute A1 : [Customer name] Attribute A2 : [Customer telephone number]

Entity E2 : [Game cards] Key K2 : [Card ID]

Attribute A3 : [Introduction] Relationship R1: [Borrow]

Attribute A4 : [Agreed return time] Attribute A5 : [Borrow time]

Attribute A6 : [Return time]

Question 2

Correct

Marked out of
2.00

Which of the following questions can be answered from the ER diagram in Question 1?

- a. Who will be the next customers to borrow one specific card?
- b. Who are the customers that borrowed the most? ✓
- c. Which game cards are the most popular (borrowed the most)? ✓
- d. Who are the customers that did not return cards on time? ✓

The correct answers are: Who are the customers that borrowed the most?,
Who are the customers that did not return cards on time?,
Which game cards are the most popular (borrowed the most)?

Question 3

Correct

Marked out of
2.00

What is the type of relationship R1 between the entity E1 and E2 in Question 1?

- a. Many-to-Many ✓
- b. One-to-Many
- c. One-to-One
- d. Many-to-one

The correct answer is:
Many-to-Many

Question 4

Correct

Marked out of
2.00

Consider the ER model introduced in Question 1, and fill in the gaps in the sentences.

Words can be reused multiple times or not be used at all.

The participation of the E1 entity in the R1 relationship is ✓

The participation of the E2 entity in the R1 relationship is ✓

The correct answer is:

Consider the ER model introduced in Question 1, and fill in the gaps in the sentences.

Words can be reused multiple times or not be used at all.

The participation of the E1 entity in the R1 relationship is [Partial]

The participation of the E2 entity in the R1 relationship is [Partial]

Question 5

Correct

Marked out of
2.00

Are E1 and E2 in Question 1 weak entities?

- a. E2 is a weak entity but E1 is not a weak entity
- b. None of them is a weak entity ✓
- c. E1 is a weak entity but E2 is not a weak entity
- d. Both are weak entities

The correct answer is:

None of them is a weak entity

Question 6

Correct

Marked out of
2.00

Consider a weak entity set W and its identifying entity set O . The primary key of W is composed of

- a. All attributes of W and the primary key of O .
- b. The partial key of W and the primary key of O . ✓
- c. The partial key of W and all attributes of O .
- d. All attributes of W and all attributes of O .


The correct answer is:

The partial key of W and the primary key of O .

Information

Suppose a fast food company hands out membership cards to their users, which they can choose to enter when ordering food. It is however not required to be a user, nor to have a card, in order to make an order.

Consider the following translation from an ER model to the relational model



And the following relational model:

- User(userid, name)
- Card(cardid, code)
- UserHolds(userid, cardid, since)
- Order(orderid, item, amount)
- OrderUses(userid, cardid, orderid)

Where fields underlined are the primary keys of the relation. All fields are not null.

Question 7

Correct

Marked out of
2.00

Choose the **fewest** number of choices of key constraints to be added, so the relational model completely represents all key constraints of the ER model.

- a. Foreign keys: OrderUses(orderid) REFERENCES Order(orderid) ✓
- b. Unique: OrderUses(orderid) ✓
- c. Foreign keys: UserHolds(userid) REFERENCES User(userid), UserHolds(cardid) REFERENCES Card(cardid) ✓
- d. Unique: UserHolds(cardid) ✓
- e. Foreign keys: OrderUses(userid, cardid) REFERENCES UserHolds(userid, cardid) ✓
- f. Unique: UserHolds(userid)
- g. Unique: OrderUses(userid, cardid)
- h. Foreign keys: OrderUses(userid) REFERENCES User(userid), OrderUses(cardid) REFERENCES Card(cardid)

The correct answers are:

Foreign keys: UserHolds(userid) REFERENCES User(userid), UserHolds(cardid) REFERENCES Card(cardid),

Unique: UserHolds(cardid),

Foreign keys: OrderUses(orderid) REFERENCES Order(orderid),

Foreign keys: OrderUses(userid, cardid) REFERENCES UserHolds(userid, cardid),

Unique: OrderUses(orderid)

Question 8

Correct

Marked out of
1.00

Is it possible to merge the tables Order and OrderUses?

- a. No ✓
- b. Yes

The correct answers are:

No,

Yes

Question 9

Correct

Marked out of
1.00

Is it possible to merge UserHolds with another table?

- a. Yes, UserHolds can be merged with OrderUses
- b. No, UserHolds can not be merged at all
- c. Yes, UserHolds can be merged with User
- d. Yes, UserHolds can be merged with Card ✓

The correct answer is:

Yes, UserHolds can be merged with Card

Question 10

Correct

Marked out of
1.00

Suppose we want to enforce that each card can only be used on one order. Which constraints have to be added into OrderUses?

- a. UNIQUE (orderid)
- b. UNIQUE (cardid, orderid)
- c. UNIQUE (cardid) ✓
- d. No constraints needs to be added, it is already enforced with the current constraints

The correct answer is:


UNIQUE (cardid)

Information

Suppose a book publishing company wants to digitally lend its books to customers and transfer parts of its lending price to the books' authors. The company wants to store the following:

- For each author, their id and IBAN account number.
- For each book, its title, year, one primary author, zero or more secondary authors and its lending price.
- For each customer, their email address, which should be unique to the user.
- A record of all current lendings, each consisting of a customer lending a single book with an expiration date, and the date the transaction was made.
- Finally, as a measure against mass lending, each customer can only lend **one book per day**.

The company came up with the following ER model:



Question 11

Correct

Marked out of
10.00

Finish the following SQL statements to create the relational model corresponding to the requirements, guided by the ER model.

Assume that all string fields are at most 100 characters long, and we only use CHAR(..) (where .. is filled with a number), INTEGER, DATE and FLOAT for data types.

If multiple columns are needed, list them in the order they are defined in the table, with one comma and one space in between (for example: "first_col, second_col").

```
CREATE TABLE Author (  
  id INTEGER NOT NULL,  
  iban CHAR(25) NOT NULL UNIQUE,  
  PRIMARY KEY (id)  
);  
CREATE TABLE Book (  
  id INTEGER NOT NULL,  
  title CHAR(100) ✓ NOT NULL,  
  year INTEGER ✓ NOT NULL,  
  price FLOAT NOT NULL,  
  primary_author INTEGER ✓ NOT NULL,  
  PRIMARY KEY (id),  
  FOREIGN KEY (primary_author) REFERENCES Author ✓ ( id ✓ )
```

```
);  
CREATE TABLE SecondaryAuthor (  
  author_id INTEGER ✓ NOT NULL,  
  book_id INTEGER ✓ NOT NULL,  
  PRIMARY KEY ( author_id, book_id ✓ ),  
  FOREIGN KEY (author_id) REFERENCES Author ✓ ( id ✓ ),  
  FOREIGN KEY (book_id) REFERENCES Book ✓ ( id ✓ )  
);
```

```
CREATE TABLE Customer (  
  id CHAR(100) NOT NULL,  
  email CHAR(100) ✓ NOT NULL,  
  UNIQUE( email ✓ ),  
  PRIMARY KEY (id)
```

```
);  
CREATE TABLE Lend(  
  customer_id CHAR(100) ✓ NOT NULL,  
  book_id INTEGER ✓ NOT NULL,  
  on_date DATE NOT NULL,  
  expiry_date DATE NOT NULL,  
  PRIMARY KEY ( customer_id, book_id ✓ ),  
  UNIQUE ( customer_id, on_date ✓ ),  
  FOREIGN KEY (customer_id) REFERENCES Customer(id),  
  FOREIGN KEY (book_id) REFERENCES Book ✓ ( id ✓ )
```

```
);
```

Question 12

Correct

Marked out of
2.00

The following is the description of the database. The questions come afterwards.

Pubs Database Schema

author(author_id, first name, last name)

author_pub(author_id, pub_id, author position)

book(book_id, book_title, month, year, editor)

pub(pub_id, title, book_id)

- author_id in author_pub is a foreign key referencing author
- pub_id in author_pub is a foreign key referencing pub
- book_id in pub is a foreign key referencing book
- editor in book is a foreign key referencing author(author_id)
- Primary keys are underlined

Pubs Database State:

$r(author)$

author_id	first_name	last_name
1	John	McCarthy
2	Dennis	Ritchie
3	Ken	Thompson
4	Claude	Shannon
5	Alan	Turing
6	Alonzo	Church
7	Perry	White
8	Moshe	Vardi
9	Roy	Batty

$r(book)$

book_id	book_title	month	year	editor
1	CACM	April	1960	8
2	CACM	July	1974	8
3	BST	July	1948	2
4	LMS	November	1936	7
5	Mind	October	1950	NULL
6	AMS	Month	1941	NULL
7	AAAI	July	2012	9
8	NIPS	July	2012	9

$r(author_pub)$

author_id	pub_id	author_position
1	1	1
2	2	1
3	2	2
4	3	1
5	4	1
5	5	1
6	6	1

$r(pub)$

pub_id	title	book_id
1	LISP	1
2	Unix	2
3	Info Theory	3
4	Turing Machines	4
5	Turing Test	5
6	Lambda Calculus	6

How many tuples will be returned by the following relational algebra query?

$\pi_{book_title}(book)$

- a. 9
- b. 6
- c. 8
- d. 7 ✓

The correct answer is:

7

Question 13

Correct

Marked out of
2.00

How many tuples are returned by the following relational algebra expression?

$|\pi_{author_id}(author) - \pi_{editor}(book)|$

- a. 1
- b. 5 ✓
- c. 10
- d. 0
- e. 9
- f. 8

The correct answer is:

5

Question 14

Correct

Marked out of
3.00

Which of the following relational algebra expressions returns the names of all authors who are book editors?

- a. $\pi_{first_name, last_name}(author \bowtie_{author_id=pub_id} author_pub)$
- b. $\pi_{first_name, last_name}(author \bowtie_{author_id=editor} book)$ ✓
- c. $\pi_{first_name, last_name}(author \bowtie_{author_id=author_id} author_pub)$
- d. $\pi_{first_name, last_name}(author \bowtie_{author_id=book_id} pub)$

The correct answer is:

$$\pi_{first_name, last_name}(author \bowtie_{author_id=editor} book)$$

Question 15

Correct

Marked out of
3.00

Which of the followings is a relational algebra expression that returns the names of all authors who are not book editors?

- a. $\pi_{first_name, last_name}(author) \cup \pi_{first_name, last_name}(author \bowtie_{author_id=editor} book)$
- b. $\pi_{first_name, last_name}(author) - \pi_{first_name, last_name}(author \bowtie_{author_id=editor} book)$ ✓
- c. $\pi_{first_name, last_name}(author) \cap \pi_{first_name, last_name}(author \bowtie_{author_id=editor} book)$
- d. $\pi_{first_name, last_name}((\pi_{author_id}(author) - \pi_{editor}(book)) * author)$ ✓

The correct answers are:

$$\pi_{first_name, last_name}(author) - \pi_{first_name, last_name}(author \bowtie_{author_id=editor} book)$$

$$\pi_{first_name, last_name}((\pi_{author_id}(author) - \pi_{editor}(book)) * author)$$

Question 16

Correct

Marked out of
3.00

What question does the following relational algebra expression answer?

$$author * (author_pub * (\sigma_{month='July'}(book) * pub))$$

- a. Which authors authored a pub that was published in July? ✓
- b. Which authors edited books that were published in July?
- c. Which authors were born in July?

The correct answer is:

Which authors authored a pub that was published in July?

Question 17

Correct

Marked out of
2.00

Which SQL statement is used to retrieve unique values from a column?

- a. `SELECT DIFFERENT column_name FROM table_name;`
- b. `SELECT DISTINCT column_name FROM table_name;` ✓
- c. `SELECT UNIQUE column_name FROM table_name;`
- d. `SELECT column_name FROM table_name;`

The correct answer is:

`SELECT DISTINCT column_name FROM table_name;`

Question 18

Correct

Marked out of
2.00

What does the SQL function COUNT(*) do?

- a. Counts all rows in a table, including those with null values ✓
- b. Counts all non-null values in a column.
- c. Counts all columns in a table.
- d. Counts all distinct values in a column.

The correct answer is:

Counts all rows in a table, including those with null values

Question 19

Correct

Marked out of
2.00

How do you select all columns from a table named 'employees'?

- a. SELECT all FROM employees;
- b. SELECT columns FROM employees;
- c. SELECT * FROM employees; ✓
- d. SELECT employees.*;

The correct answer is:

SELECT * FROM employees;

Question 20

Correct

Marked out of
1.00

Which SQL statement is used to remove a table from a database?

- a. REMOVE TABLE table_name;
- b. DELETE TABLE table_name;
- c. TRUNCATE TABLE table_name;
- d. DROP TABLE table_name; ✓

The correct answer is:

DROP TABLE table_name;

Question 21

Correct

Marked out of
2.00

What is the RAID level that provides exact mirroring without handling corruption?

- a. RAID 5
- b. RAID 1 ✓
- c. RAID 3
- d. RAID 0

The correct answer is:

RAID 1

Question 22

Correct

Marked out of
2.00

Which of these are the responsibilities of the buffer manager.

- a. Keeps active pages in memory ✓
- b. Reduces IO by buffering/caching ✓
- c. Discards/write back pages when needed ✓
- d. Uses logging that can be used for recovery ✓

The correct answers are:

Reduces IO by buffering/caching,

Keeps active pages in memory,

Discards/write back pages when needed

,

Uses logging that can be used for recovery

Information

You have designed a buffer pool capable of storing up to 4 frames at a given time.

Assuming you have three corresponding access traces (one for each policy) where each letter represents a requested page ID. The trace starts from the left and goes right in the discrete time.

While answering the questions, assume you start from an empty buffer pool for each policy, and you don't need to consider the pin-count for pages, and no page is written on.

Information

(A) For below two questions (A.1, A.2) use the following trace:

1 2 5 4 1 6 3 1 3 5 1 7 2

Question 23

Correct

Marked out of 2.00

(A.1) If you use **FIFO** algorithm for page replacement, how many I/O does the buffer manager need to issue?

Answer:



The correct answer is: 10

Question 24

Correct

Marked out of
2.00

(A.2) If you use **FIFO** algorithm for page replacement, which pages will remain in the buffer pool at the end of trace?

- a. [2] [3] [5] [7]
- b. [2] [3] [5] [6]
- c. [1] [2] [6] [7]
- d. [1] [2] [5] [7] ✓
- e. [1] [3] [5] [6]

The correct answer is:

[1] [2] [5] [7]

Information

(B) For below two questions (B.1, B.2) use the following trace:

2 1 4 5 1 6 3 1 3 5 1 7 2

Question 25

Correct

Marked out of
2.00

(B.1) If you use **LRU** algorithm for page replacement, how many I/O does the buffer manager need to issue?

Answer: ✓

The correct answer is: 8

Question 26

Correct

Marked out of
2.00

(B.2) If you use **LRU** algorithm for page replacement, which pages will remain in the buffer pool at the end of trace?

- a. [1] [2] [5] [7] ✓
- b. [1] [2] [3] [7]
- c. [1] [4] [6] [7]
- d. [1] [3] [5] [7]
- e. [1] [2] [3] [5]

The correct answer is:

[1] [2] [5] [7]

Information

(C) For below two questions (C.1, C.2) use the following trace:

5 6 2 4 1 3 6 1 3 5 1 2 7

Question 27

Correct

Marked out of
2.00

(C.1) If you use **CLOCK** algorithm for page replacement, how many I/O does the buffer manager need to issue?

Answer: ✓

The correct answer is: 10

Question 28

Correct

Marked out of
2.00

(C.2) If you use **CLOCK** algorithm for page replacement, which pages will remain in the buffer pool at the end of trace?

- a. [1] [3] [5] [7]
- b. [1] [2] [6] [7]
- c. [1] [2] [3] [7]
- d. [1] [2] [5] [7]
- e. [2] [3] [5] [7]
- f. [2] [5] [6] [7] ✓

The correct answer is:

[2] [5] [6] [7]

Question 29

Complete

Not graded

(Optional) You can briefly explain your answers to the questions above. You can use free text or formulas to explain.

Question 30

Correct

Marked out of
2.00

DataFlex Inc. stores its customer records in a database with the following schema: CUSTOMER(custID, firstName, lastName, address, city, country, joinDate, lastOrder)

The size of each data record is 256B with the following attribute sizes: custID is 8B, firstName is 32B, lastName is 32B, address is 64B, city is 32B, country is 16B, joinDate is 8B, and lastOrder is 64B (total of 256B per record).

The disk page size is 4096B. Each page can store only a whole number of records, and every page contains a 32B header and 32B footer of metadata.

How many records fit on a single page?

Answer: ✓

The correct answer is: 15

Question 31

Correct

Marked out of
1.00

Using the same CUSTOMER table and page information from the previous question, what is the size of the CUSTOMER file in number of disk pages if it contains 50,000 customer records?

Answer: ✓

The correct answer is: 3334

Question 32

Correct

Marked out of
2.00

[Record Storage] Answer the following questions (marked with [Record Storage]) assuming the following scenario.

Consider the TravelAgency database that uses the PAX (Partition Attributes Across) storage model for its BOOKING table:
BOOKING(bookingID, customerName, destination, departureDate, returnDate, price)

The size breakdown per attribute is as follows: bookingID is 8B, customerName is 48B, destination is 32B, departureDate is 8B, returnDate is 8B, and price is 8B (total of 112B per record).

The disk page size is 2048B: each page contains a 32B header and no footer. Each minipage within a PAX page also has a 4B header.

[Record Storage] How many records fit in a single PAX page?

Answer:



The correct answer is: 17

Question 33

Correct

Marked out of
3.00

[Record Storage] Calculate the offset (in bytes) from the beginning of the page to the start of the price minipage **records** in the PAX layout.

- a. 1768
- b. 1824 ✓
- c. 1800
- d. 1960

The correct answer is:
1824

Question 34

Correct

Marked out of
1.00

In NSM (N-ary Storage Model), if a query needs to access only one attribute from a table, it needs to read all attributes of all records.

- True ✓
- False

The correct answer is 'True'.

Question 35

Correct

Marked out of
1.00

DSM (Decomposition Storage Model) typically has better compression rates than NSM because values of the same attribute are stored together.

- True ✓
- False

The correct answer is 'True'.

Question 36

Correct

Marked out of
1.00

PAX (Partition Attributes Across) storage model is optimized for analytic processing workloads but performs poorly for transaction processing workloads compared to NSM.

- True ✓
- False

The correct answer is 'True'.

Question 37

Correct

Marked out of
1.00

[I/O Cost] Answer the following questions (marked with [I/O Cost]) assuming the following scenario.

InvestBank maintains a TRANSACTION table to record all financial transactions: TRANSACTION(transID, accountID, timestamp, amount, type, description)

The file contains 1 million records and is stored as a heap file with 20 records per page (50,000 pages total). The database uses NSM storage model, and the file is not sorted on any attribute.

[I/O Cost] The I/O cost (number of page accesses) for finding all transactions for accountID = 54321 without any index is:

Answer: 

The correct answer is: 50000

Question 38

Correct

Marked out of
1.00

[I/O Cost] The TRANSACTION file now has an unclustered B+ tree index on the accountID attribute. Each account has exactly 100 transactions on average. The I/O cost (number of page accesses) to find all transactions for accountID = 54321 is: (Assume the B+ tree has a height of 3 and the index entries are much smaller than data records)

Answer: 

The correct answer is: 103

Question 39

Correct

Marked out of
2.00

For queries that need to retrieve records in a specific range of values (e.g., all transactions with amount between \$1,000 and \$5,000), the most efficient file organization is:

- a. B+ tree index ✓
- b. Heap file
- c. Sorted file
- d. Hash index

The correct answer is:

B+ tree index

Question 40

Correct

Marked out of
1.00

You are designing a database that requires fast insertions and point lookups but does not need range queries. Which file organization should you choose?

- a. B+ tree index
- b. Hash index ✓
- c. Heap file
- d. Sorted file

The correct answer is:

Hash index

Question 41

Correct

Marked out of
3.00**Given the employee dataset:****EMPLOYEES(jobid, name, age, department, ranking)****The size of each record is 120B: jobid (primary key) is 10B, name is 50B, age is 4B, department is 44B, ranking is 12B.****The records are organized in 8KB pages (each page stores an integer number of records, and you can ignore metadata).****You can assume that the distribution of rankings follows a Gaussian curve from 1-10, the age distribution is typical of a mid-sized tech company workforce, and the most popular department is "Engineering" while the least popular department is "Legal".****Which index to use for that query, depending on the following query patterns and dataset sizes? If it would be better to use no index than the possible options, leave each dropdown as No Index.****(1)****Dataset size of 10.5 million records with infrequent modifications and frequent queries of the format `SELECT * FROM employees WHERE age < 21`****Data structure is:**

b-tree



The correct answer is:

Given the employee dataset:**EMPLOYEES(jobid, name, age, department, ranking)**

The size of each record is 120B: jobid (primary key) is 10B, name is 50B, age is 4B, department is 44B, ranking is 12B. The records are organized in 8KB pages (each page stores an integer number of records, and you can ignore metadata). You can assume that the distribution of rankings follows a Gaussian curve from 1-10, the age distribution is typical of a mid-sized tech company workforce, and the most popular department is "Engineering" while the least popular department is "Legal".

Which index to use for that query, depending on the following query patterns and dataset sizes? If it would be better to use no index than the possible options, leave each dropdown as No Index.

(1)

Dataset size of 10.5 million records with infrequent modifications and frequent queries of the format `SELECT * FROM employees WHERE age < 21`

Data structure is: [b-tree]

Question 42

Correct

Marked out of
3.00

Format is ✓

The correct answer is:
Format is [clustered]

Question 43

Correct

Marked out of
5.00

On what key:

age



The correct answer is:

On what key: [age]

Question 44

Correct

Marked out of
3.00

Density is

sparse



The correct answer is:

Density is [sparse]

Question 45

Correct

Marked out of
3.00

Dataset size of 500,000 records with frequent updates and queries of the format `SELECT jobid FROM employees WHERE department="Legal"`

Data structure is: ✓

The correct answer is:

Dataset size of 500,000 records with frequent updates and queries of the format `SELECT jobid FROM employees WHERE department="Legal"`

Data structure is: [hash]

Question 46

Correct

Marked out of
3.00

Format is: ✓

The correct answer is:

Format is: [unclustered]

Question 47

Correct

Marked out of
5.00

On what key:



The correct answer is:

On what key: [department]

Question 48

Correct

Marked out of
3.00

Density is:



The correct answer is:

Density is: [Dense]

Question 49

Correct

Marked out of
3.00

Dataset size of 40 records with frequent updates and queries of the format `SELECT AVG(ranking) FROM employees WHERE age > 23 AND department="Engineering"`

Data structure is: ✓

The correct answer is:

Dataset size of 40 records with frequent updates and queries of the format `SELECT AVG(ranking) FROM employees WHERE age > 23 AND department="Engineering"`

Data structure is: [no index]

Question 50

Correct

Marked out of
3.00

Format is: ✓

The correct answer is:

Format is: [no index]

Question 51

Correct

Marked out of
6.00On what key: 

The correct answer is:

On what key: [no index]

Question 52

Correct

Marked out of
3.00Density is: 

The correct answer is:

Density is: [no index]

[◀ MIDTERM Attendance](#)[Midterm 2025 - Questions ▶](#)**EPFL**

Contact EPFL CH-1015 Lausanne +41 21 693 11 11

Follow the pulses of EPFL on social networks