Question 1

SP Bumi Transport has decided to build a database to help it manage Its operations. It is Interested in keeping track of It buses, the terminals to which they stop, passengers and the trips they have taken, its employees, and data about the maintenance of its buses.

a) Using this introductory description of (and hints about) the SP Burni Transport, make a list of the SEVEN(7) ' things' in the company's environment about which you think the company would want to maintain data.

Ans:- a) List of the seven 'things' in SP Bumi Transport's environment about which the company would want to maintain data:

1. Buses: Information about each bus in the fleet, including make, model, age, maintenance history, seating capacity, and license details.

2. Terminals: Data related to the terminals where buses stop, including location, facilities, operating hours, and routes served.

3. Passengers: Records of passengers, including their names, contact information, booking history, and preferences.

4. Trips: Details of each trip taken by the buses, including departure and arrival times, dates, routes, fares, and occupancy rates.

5. Employees: Information about company employees, such as drivers, maintenance staff, and administrative personnel, including their names, contact details, roles, and schedules.

6. Maintenance: Data on bus maintenance schedules, repairs, and inspections, including dates, maintenance logs, costs, and the status of each vehicle.

7. Operations: Information about the day-to-day operations, including schedules, routes, ticket sales, and financial data like revenue and expenses.

b) Suggest THREE (3) reports that SP Bum' Transport could use to Improve its operations and gain a competitive advantage?

Ans:- Three reports that SP Bumi Transport could use to improve its operations and gain a competitive advantage:

1. Occupancy Analysis Report: This report can provide insights into the occupancy rates of different routes and trips, helping the company optimize its schedules, allocate resources efficiently, and adjust pricing strategies to maximize revenue.

2. Maintenance Forecast Report: By analyzing maintenance data, this report can predict when each bus is likely to require maintenance or repairs. Proactive maintenance can minimize downtime, improve safety, and reduce operational costs.

3. Customer Satisfaction and Feedback Report: Gathering and analyzing passenger feedback can help the company identify areas for improvement in service quality, amenities, and customer satisfaction. Addressing passenger concerns can enhance the company's reputation and customer loyalty.

c)Describe the THREE (3) main characteristics of the database approach and contrast It with the file-based approach.

Ans:- Ans:- Three main characteristics of the database approach and contrast with the file-based approach:

Database Approach

1. Data Integration: Databases integrate data from various sources into a centralized system, ensuring data consistency and reducing redundancy. In contrast, file-based systems may have multiple, separate files with redundant data.

2. Data Independence: In a database, data is independent of the application programs using it. Changes in data structure (schema) can be made without affecting the applications. In file-based systems, changes in data often require modifying the application code.

3. Concurrency Control: Databases support concurrent access by multiple users, ensuring data integrity and consistency through mechanisms like locks and transactions. File-based systems often lack such control, leading to data anomalies in multi-user environments.

File-Based Approach

1. Data Redundancy: File-based systems often result in data redundancy since the same data may be stored in multiple files. This redundancy can lead to inconsistencies and increased storage requirements.

2. Limited Data Sharing: Sharing data among different applications can be challenging in file-based systems. Each application typically manages its data independently, making it difficult to access and update shared information.

3. Data Dependence: File-based systems tightly couple data with application programs. Any changes to data structures require corresponding changes to all applications that use that data, making maintenance more complex.

d) Explain the differences of 'Data Definition Language (DDL)" and "Data Manipulation Language (DML)".

Ans:- Differences between "Data Definition Language (DDL)" and "Data Manipulation Language (DML)":

Data Definition Language (DDL)

1. Purpose: DDL is used to define and manage the structure of a database. It includes creating, altering, and deleting database objects like tables, indexes, and constraints.

2. Operations: Common DDL commands include CREATE (to create database objects), ALTER (to modify object structures), and DROP (to delete objects).

3. Examples: CREATE TABLE, ALTER TABLE, CREATE INDEX, DROP TABLE, etc.

Data Manipulation Language (DML)

1. Purpose: DML is used to manipulate the data stored in the database. It involves querying, inserting, updating, and deleting data in the database.

2. Operations: Common DML commands include SELECT (to retrieve data), INSERT (to add new data), UPDATE (to modify existing data), and DELETE (to remove data).

3. Examples: SELECT * FROM table_name, INSERT INTO table_name, UPDATE table_name SET column_name = value, DELETE FROM table_name, etc.

Question 2

a) Read the following description carefully. State appropriate assumptions and devise the corresponding EAR model:

The Selangor Stage Coach Co. provides services to the greater Selangor municipal, Including various towns around the state capital.

It owns a substantial number of buses. Each bus Is allocated to a particular route although some routes may have several buses. Each bus has a unique bus number. It Is Important to store information about the seating capacity and the make/type of all buses.

Each route, distinguish by a route number, passes through a number of towns. Several routes may serve the same town. Information Is available on the average number of passengers carried per day for each route.

Due to long travelling time, one or more drivers are assigned to each stage of a route, which corresponds to a Journey through a town on a route. Drivers have an employee number, name, address and sometimes a telephone number.

Ans:- Assumptions for the EAR model:

1. Each bus has a single unique bus number.

2. Each bus has one seating capacity and one make/type.

3. Each route has a unique route number.

4. Each route passes through multiple towns.

5. Multiple buses can be assigned to the same route.

6. Each route has information about the average number of passengers carried per day.

7. Each stage of a route corresponds to a journey through a town.

8. Each driver has a unique employee number.

9. Each driver has a name, address, and optionally a telephone number.

EAR Model:

Entities:

- Bus (BusNumber, SeatingCapacity, Make/Type)

- Route (RouteNumber)

- Town (TownName)

- Driver (EmployeeNumber, Name, Address, TelephoneNumber)

- RouteStage (StageID, RouteNumber, TownName)

- RoutePassengerStats (RouteNumber, AveragePassengersPerDay)

Relationships

1. Bus is associated with Route (1 to many): A bus can be assigned to one or more routes.

2. Route is associated with Town (many to many): A route passes through multiple towns, and multiple routes can serve the same town.

3. RouteStage is associated with Route and Town (1 to many): Each stage of a route corresponds to a journey through a town on a route.

4. Driver is associated with RouteStage (many to many): Multiple drivers can be assigned to different stages of a route.

5. Route has RoutePassengerStats (1 to 1): Information on average passengers carried per day is associated with a route.

b) Devise a relational database schema to demonstrate how you will map your EAR model in (a) onto tables in a relational database. Identify the primary keys / foreign keys in each entity.

Ans:- Relational Database Schema:

Table: Bus

- Primary Key: BusNumber

- Attributes: BusNumber (PK), SeatingCapacity, Make/Type

Table: Route

- Primary Key: RouteNumber

- Attributes: RouteNumber (PK)

Table: Town

- Primary Key: TownName

- Attributes: TownName (PK)

Table: Driver

- Primary Key: EmployeeNumber
- Attributes: EmployeeNumber (PK), Name, Address, TelephoneNumber

Table: RouteStage

- Primary Key: StageID
- Foreign Keys: RouteNumber (references Route), TownName (references Town)
- Attributes: StageID (PK), RouteNumber (FK), TownName (FK)

Table: RoutePassengerStats

- Primary Key: RouteNumber
- Foreign Key: RouteNumber (references Route)
- Attributes: RouteNumber (PK, FK), AveragePassengersPerDay

Question 3

Table 1 shows list of dentist / patient appointment data. A patient Is given an appointment at a specific time and date with a dentist located at a particular surgery. On each day of patient's appointments, a dentist is allocated to a specific surgery for that day.

Table 1 — Dentist Patient Uata Stall

Staff_No	Dentist_Name	Pat_No	Pat_Name	Appointment date time	Surgery_No
S1011	Zara	P100	Roiana	12-Sep-17 10:00	S15
S1011	Zara	P105	Tiagu	12-Sep-17 12:00	S15
S1024	Aruna	P108	Andrew	12-Sep-17 10:00	S10
S1024	Aruna	P108	Andrew	14-Sep-17 12:00	S10
S1032	Robin	P105	Wong	14-Sep-17 10:00	S15

a) **The shown Table is susceptible to update anomalies. Provide ONE (1) example of insertion, deletion and update anomalies.**

Ans:-The shown Table 1 is susceptible to update anomalies. Here are examples of insertion, deletion, and update anomalies:

Insertion Anomaly: If you want to insert a new patient record but that patient doesn't have an appointment yet, you cannot do it in this table because it is structured to store both patient and appointment information together. This is an insertion anomaly.

Deletion Anomaly: If you delete a row for a patient's appointment, you would also delete information about the patient's dentist, which might be relevant for other appointments of the same patient. This is a deletion anomaly.

Update Anomaly: If a dentist changes their name, you would need to update multiple rows with the same dentist name. If you miss updating one of them, it can lead to inconsistency. This is an update anomaly.

b) Illustrate the process of normalizing Table 1 to 3NF relations. Identify the primary key, alternate and foreign keys in your 3NF relations.

Ans:- To normalize Table 1 to 3NF (Third Normal Form), we need to follow these steps:

Step 1: Identify the Candidate Keys

- First, let's identify the candidate keys in the table. A candidate key is a minimal set of attributes that can uniquely identify each tuple (row) in the table.

- In this table, 'Staff_No' and 'Appointment date time' together can uniquely identify each appointment because a dentist can have multiple appointments on different dates and times. Therefore, (Staff_No, Appointment date time) is a candidate key.

Step 2: Eliminate Partial Dependencies

- Partial dependencies occur when non-prime attributes (attributes not part of the candidate key) depend on only part of the candidate key. To eliminate partial dependencies, we'll create two new tables:

1. Dentist_Info (Dentist_Name, Staff_No, Surgery_No)
2. Patient_Appointments (Pat_No, Pat_Name, Appointment date time)

Step 3: Eliminate Transitive Dependencies

- Transitive dependencies occur when non-prime attributes depend on other non-prime attributes. To eliminate transitive dependencies, we'll create a new table:

3. Surgery_Info (Surgery_No, Staff_No)

Here are the three normalized tables in 3NF:

Table 1: Dentist_Info

- Primary Key: Staff_No

- Alternate Key: Dentist_Name

- Foreign Key: Surgery_No (references Surgery_Info)

Staff_No	Dentist_Name	Surgery_No
S1011	Zara	S15
S1024	Aruna	S10
S1032	Robin	S15

Table 2: Patient_Appointments

- Primary Key: Appointment date time

- Alternate Key: Pat_No

- Foreign Key: Staff_No (references Dentist_Info)

Pat_No	Pat_Name	Appointment date time
P100	Roiana	12-Sep-17 10:00
P105	Tiagu	12-Sep-17 12:00
P108	Andrew	12-Sep-17 10:00
P108	Andrew	14-Sep-17 12:00
P105	Wong	14-Sep-17 10:00

Table 3: Surgery_Info

- Primary Key: Surgery_No

- Foreign Key: Staff_No (references Dentist_Info)

Surgery_No	Staff_No
S15	S1011
S10	S1024
S15	S1032

c.Define Boyce Codd Normal Form (BCNF) and identify the relation you normalized in (b) is in BCNF already or not .

Ans:- Boyce-Codd Normal Form (BCNF) is a higher level of normalization that ensures that for every non-trivial functional dependency X → Y in a relation, X is a superkey. In other words, it ensures that there are no partial dependencies.

In the normalized relations from part (b), Table 4 (Appointments) is in BCNF because there are no partial dependencies, and for any functional dependency, the left side (X) is a superkey.

d) What Is the purpose of normalization?

Ans:- The purpose of normalization in database design is to minimize data redundancy and eliminate data anomalies (insertion, deletion, and update anomalies) while preserving data integrity. Normalization achieves this by organizing data into multiple related tables, each with a specific purpose, and by ensuring that data is stored efficiently without unnecessary duplication. This leads to better data consistency, accuracy, and maintainability in the database, ultimately improving data quality and reducing the likelihood of errors.

Question 4

Consider the following relational database for Nepal Airlines.

Nepal Airlines has to keep track of its flight and airplane history. A flight is uniquely identified by the combination of a flight number and a date. Every passenger who has flown on Nepal Airlines has a unique passenger number. For a particular passenger who has taken a particular flight, the company wants to keep track of the fare that she paid for it and the date that she made the reservation for it. Clearly, a passenger may have taken many flights (he must have taken at least one to be in the database) and every flight has had many passengers on it.

A pilot is identified by a unique pilot (or employee) number. A flight on a particular date has exactly one pilot. Each pilot has typically flown many flights but a pilot may be new to the company, is in training, and has not flown any flights, yet. Each airplane has a unique serial number. A flight on a particular date used one airplane. Each airplane has flown on many flights and dates, but a new airplane may not have been used at all, yet.

The following describes the attributes for each entity.

PILOT

PILOTNUM	PILOTNAME	BIRTHDATE	HIREDATE

FLIGHT

FLIGHTNUM	DATE	DEPTIME	ARRTIME	PILOTNUM	PLANENUM

PASSENGER

PASSNUM	PASSNAME	ADDRESS	PHONE

RESERVATION

FLIGHTNUM	DATE	PASSNUM	FARE	RESVDATE

AIRPLANE

PLANENUM	MODEL	CAPACITY	YEARBUILT	MANUF

Write SQL SELECT commands to answer the following queries.

a) Find the records for the airplanes manufactured by Boeing.

b) How many reservations are there for flight 278 on February 21,2017?

c) List the flights on March 7, 2017 that are scheduled to depart between 10 and 11AM or that are scheduled to arrive after 3PM on that date.

d) How many Boeing 737s does Grand Travel have?

e) How many of each model of Boeing aircraft does Grand Travel have?

f) List the names and dates of hire of the pilots, who flew Airbus A320 aircraft In March, 2017. (3 marks)

g) List the names, addresses, and telephone numbers of the passengers who have reservations on Flight 562 on January 15, 2017.

h) What was the total fare paid for each flight scheduled to depart between 9 and 10AM on December 23, 2003? Only Include those flights for which the total fare was at least 55,000.

I) List the smallest (in terms of passenger capacity) Boeing 737s.

Ans:-

```
a) Find the records for the airplanes manufactured by Boeing.
SELECT *
FROM AIRPLANE
WHERE MANUF = 'Boeing';

b) How many reservations are there for flight 278 on February 21, 2017?
SELECT COUNT(*)
FROM RESERVATION
WHERE FLIGHTNUM = 278 AND DATE = '2017-02-21';

c) List the flights on March 7, 2017, that are scheduled to depart between
10 and 11 AM or that are scheduled to arrive after 3 PM on that date.
SELECT *
FROM FLIGHT
WHERE DATE = '2017-03-07'
AND (DEPTIME BETWEEN '10:00:00' AND '11:00:00' OR ARRTIME > '15:00:00');

d) How many Boeing 737s does Grand Travel have?
SELECT COUNT(*)
FROM AIRPLANE
WHERE MANUF = 'Boeing' AND MODEL = '737';

e) How many of each model of Boeing aircraft does Grand Travel have?
SELECT MODEL, COUNT(*)
FROM AIRPLANE
WHERE MANUF = 'Boeing'
GROUP BY MODEL;
```

f) List the names and dates of hire of the pilots who flew Airbus A320 aircraft in March 2017.

```
SELECT PILOTNAME, HIREDATE
FROM PILOT
WHERE PILOTNUM IN (
    SELECT DISTINCT PILOTNUM
    FROM FLIGHT
    WHERE MODEL = 'Airbus A320'
    AND DATE BETWEEN '2017-03-01' AND '2017-03-31'
);
```

g) List the names, addresses, and telephone numbers of the passengers who have reservations on Flight 562 on January 15, 2017.

```
SELECT P.PASSNAME, P.ADDRESS, P.PHONE
FROM PASSENGER P
JOIN RESERVATION R ON P.PASSNUM = R.PASSNUM
WHERE R.FLIGHTNUM = 562 AND R.DATE = '2017-01-15';
```

h) What was the total fare paid for each flight scheduled to depart between 9 and 10 AM on December 23, 2003? Only include those flights for which the total fare was at least 55,000.

```
SELECT FLIGHTNUM, SUM(FARE) AS TOTAL_FARE
FROM RESERVATION
WHERE DATE = '2003-12-23' AND DEPTIME BETWEEN '09:00:00' AND '10:00:00'
GROUP BY FLIGHTNUM
HAVING TOTAL_FARE >= 55000;
```

i) List the smallest (in terms of passenger capacity) Boeing 737s.

```
SELECT PLANENUM, MODEL, CAPACITY
FROM AIRPLANE
WHERE MANUF = 'Boeing' AND MODEL = '737'
ORDER BY CAPACITY ASC;
```