



SD202 Exam: 2h

02 July 2021

Grades : I (10 points), II (10 points)

Recommandation : prefer simple solutions and brief answers.

You can answer in *English* or *French*.

Course materials (on papers **ONLY**) are authorized.

Part I : Functional Dependencies and Normalization

The following questions are independent

Exercise I.1 :

1. Consider the sets of attributes A, B, C, and D.

Do the following statements hold? If yes, prove your answer using Armstrong's axioms. Otherwise, give a relevant counter-example to support your answer, e.g., using a set of tuples from the instance of a relation.

- a) If $AB \rightarrow C$ and $B \rightarrow D$, then $AD \rightarrow C$
- b) If $A \rightarrow B$, $A \rightarrow C$, and $BC \rightarrow D$, then $A \rightarrow D$

2. Consider a relation **R1(A,B,C,D,E)** with the following FDs:

$A \rightarrow B$; $BC \rightarrow E$; $ED \rightarrow A$

- a) List all candidate keys for **R1**.
- b) Is **R1** in 3NF? Justify
- c) Is **R1** in BCNF? Justify

Exercise I.2 :

Consider a relation **R2** with schema **R2(A, B, C, D, E)** and FDs:

$AB \rightarrow CD$

$D \rightarrow E$

$A \rightarrow C$

$B \rightarrow D$

1. Decompose the relation **R2** using the BCNF decomposition algorithm. Give a short justification for each new relation while precisizing the keys and the FDs.

Part II : SQL

Exercise II.1 :

Suppose a database that stores information about airlines, flights, and pilots. The database consists of the following tables (the primary keys of the relations are underlined):

Airlines (*airlineCode*, *name*)

Pilot (*pilotId*, *name*, *flight_hours*, *airlineCode*)

Flight (*airlineCode*, *number*, *date*, *from_city*, *to_city*, *depart_time*, *arrive_time*)

Flown (*pilotId*, *airlineCode*, *number*)

Notes:

- *Flight_hours* is the total number of hours the pilot has flown during his/her career. Every pilot currently works for some airline, i.e. *airlineCode* in **Pilot** is not NULL.
- **Flight** gives information about each unique flight in the database. *depart_time* and *arrive_time* are integers giving 24-hour clock time (e.g., 2015 means 20:15pm, 1135 means 11:35am). *date* is in format MM/DD/YYYY.
- **Flown** records which pilots have flown which flights. A pilot may appear many times in this table, once for each (*airlineCode*, *number*) pair, and flight may have many different pilots.
- A pilot is currently employed by one airline but has worked for different airlines in the past, i.e., **Flown** table may contain tuples for a pilot with different *airlineCode* other than the one the pilot currently works for.
- *AirlineCode* in **Pilot** and **Flight** is a foreign key (fk) that references the table Airlines. **pilotId** in **Flown** is a fk that references **Pilot**. *number* and *AirlineCode* together in **Flown** refer to *number* and *AirlineCode* in **Flight**.

Write the corresponding queries in SQL:

1. Give the number of different pilots that have more than 1000 flight hours.
2. Give the average flight hours.
3. List the total number of pilots who flew at least one flight before the year 2010.
4. List the airline names, flight numbers, and departure times for all flights from 'Paris' to 'Tunis' on or after 10:00am on 10/10/2021.
5. List the pilot names who have no scheduled flight in January 2020.
6. List the airlineCode of the airlines that have more than one assigned flight.
7. A senior pilot is a pilot who has 2000 or more total hours of flight time. List the names of all airlines that currently employ more than 10 senior pilots.

