This is a closed book exam: the only resources allowed are blank paper, pens, and your head. Explain your reasoning. Write clearly, in the sense of logic, language, and legibility. The clarity of your explanations affects your grade. Good luck!

Write your name and student number on *all* solution sheets and here. Name: .....
At the end of the exam, hand in *all* sheets that you received, including this one. Student number: .....

Problem 1 [30%] Design the Entity-Relationship schema of an application supporting the organization of conferences. When a person requests to organize a *conference*, the conference is assigned a unique identifier and the person who issued the request is recorded, together with the expected number of participants, the start date, the duration, and the rooms (at least one and at most 5) that are considered suitable for the conference. For *finished conferences*, we are interested also in the actual number of participants, the total cost, the person who acted as conference chair, and the room in which the conference was actually held (this room must be one of those that were reported as suitable at the time of the initial request). Of each *room* we are interested in the code (identifier), the capacity, and the maintenance interventions undergone by that room. Each *maintenance intervention* is related to a room, occurs on a date, has a cost, and is carried out under the responsibility of one person, with the rule that no person can be responsible for more than one maintenance intervention per month. The rooms are classified into two categories: generic rooms and hotel rooms. Of each *generic room*, we are interested in the area and the city in which it is located. Of each *hotel room* we are interested in the hotel in which it is located, the floor, and the room number (which is unique within the hotel in which the room is located). Of each *hotel* we are interested in the city in which it is located, the name (unique within the city), the category, and the person who manages it (with the date when the direction started). Of each *city* we are interested in the name (unique within the region), the region, and the number of inhabitants. Of each *person* we are interested in the social security number (identifier), the date of birth, and the city of residence.

Problem 2 [42%] Carry out the logical design of the database, producing the complete relational schema with constraints, taking into account the following indications: *(i)* Each hotel room is always accessed by the room number and the hotel in which it is located. *(ii)* When accessing a hotel we always want to know who its manager is.

As steps in your design you should produce:

1. [7%] the restructured ER schema (possibly with external constraints),
2. [25%] the direct translation into the relational model (possibly with external constraints), and
3. [10%] the restructured relational schema (again with constraints).

Motivate explicitly how the above indications affect your design.

Problem 3 [18%] Consider a database of a library that includes the relations `Book`, `Client`, and `Borrowed`. The relation `Book(name,length)` stores for each book in the library the name and its length in pages, the relation `Client(ssn,city)` stores for each registered client of the library the ssn and the city of birth, and the relation `Borrowed(ssn,name,year)` stores which books have been borrowed by which clients in which year. Express the following queries in SQL:

1. Return the ssn of the clients born in Bolzano who since 2020 have borrowed at least one book longer than 200 pages.
2. For each client, return the average length of the books they have borrowed.
3. Return the ssn and the city of birth of those clients who have borrowed all books in the library that are at least 1000 pages long.

Problem 4 [10%] Consider a relational schema containing two relations whose schemas are `R(A,B,C,D)` and `Q(E,F)`. Provide the SQL `CREATE TABLE` statement for relation `R`, taking into account that: *(i)* all attributes are of type `INTEGER`; *(ii)* attribute `A` is the primary key of `R`; *(iii)* attribute `B` cannot take null values; *(iv)* every value that appears in `C` also appears in `E` (where `E` is the primary key of `Q`); *(v)* every value that appears in $\texttt{PROJ}_{\texttt{D}}\texttt{(R)}$ also appears in $\texttt{PROJ}_{\texttt{F}}\texttt{(SEL}_{\texttt{E=1}}\texttt{(Q))}$.