

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 a company in the organization of parties celebrating a person. Of each *party* we are interested in the code (unique within the year in which the party takes place), the date on which it takes place, the venue in which it takes place, the number of expected guests, the expected cost, and the person who commissioned it. There are two types of parties, depending on how many persons are being celebrated: single parties (one celebrant) and multiple parties (at least two celebrants). Of each *single party*, we are interested to know who is the celebrant, and for what occasion the party is held (birthday, promotion, etc.). Of each *multiple party* we are interested to know who are the celebrants, and what percentage of the expenses each of them will cover. Of each *person* we are interested in the social security number, gender, date of birth, and city of birth. In addition, of each celebrant we are interested in the code (identifier assigned by the company) and the city in which the person resides. Of each person commissioning a party we are interested in the profession (which, however, is significant only if the person is employed). Of each *venue* we are interested in the code (unique within the city), the address (consisting of city, street, and house number), and the capacity. Of each *city* we are interested in the name (unique with the region), the region, and the number of inhabitants. Moreover, of each venue we are also interested in all the maintenance operations that it has undergone, where each *maintenance operation* targets one or more venues, takes place on a certain date, and is supervised by an operator, with the rule that no operator can supervise more than one maintenance operation per day. Notice that an *operator* is a person employed by the company whose year of employment is of interest. There are no restrictions on who can be operator, celebrant, or commission a party.

**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) Celebrant's data are accessed using the code assigned by the company. (ii) Null values in the database should be avoided. (iii) When accessing the data on a maintenance operation, we always want to know the operator who supervised it.

In your design you should follow the methodology adopted in the course, and 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).

You should motivate explicitly how the above indications affect your design.

**Problem 3** [18%] Consider a database that includes the relations *Person* and *WorksIn*.

The relation *Person*(*ssn*, *city*, *yearob*, *gender*) stores for each person the social security number, the city and year of birth, and the gender. The relation *WorksIn*(*ssn*, *city*, *year*) stores information about which persons (represented by the *ssn*) have worked in which cities and in which years (for example, a person may have worked in Bolzano in the years 2015, 2016, 2020, and 2021).

Express the following queries in SQL:

1. Compute the women who have worked at least once in the city in which they were born.
2. For each person who has worked at least once, compute the age at which they first worked.
3. Compute the cities in which at least one male worked and in which no female born in that city has ever worked.

**Problem 4** [10%] Consider a relational schema containing two relations whose schemas are  $R(A, B, C)$  and  $Q(D, E, F)$ , with all attributes of type INTEGER. Show how the following constraints can be expressed within the SQL CREATE TABLE statement related to relation R:

- (i) No pair of values appears simultaneously in  $PROJ_{A,B}(R)$  and  $PROJ_{E,F}(Q)$ .
- (ii) Every value that appears in  $PROJ_C(R)$  also appears in  $PROJ_F(SEL_{D=1}(Q))$ .