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 related to the management of research projects. Of each *researcher*, we are interested in the id and level (that together identify the researcher), the name, and the projects for which they work, with the roles (at least one) they hold in each project (e.g., responsible for publicity, budget responsible, etc.). Some researchers are *seniors* and for them we are also interested in their seniority. Of each *project*, we are interested in the code (identifier), the description (a string), and the researchers (at least one) that work for it. Some projects are *individual projects* and for them we are also interested in the research area. Note that only a single researcher participates to an individual projects, and such researcher must be senior. For each project, we are also interested in the *workpackages* that belong to it. Of each workpackage, we are interested in its number, which is unique within the project to which the workpackage belongs, and in its start and end date. In the case where the workpackage is a *management workpackage*, we are also interested in its cost. Note that each project has exactly one management workpackage, and the start and end dates of all workpackages of a project must fall within the start and end dates of the management workpackage of that project.

Problem 2 [42%] Carry out the logical design of the database, producing the complete relational schema with constraints, taking into account that when we access the information about a project, we always want to know whether it is an individual project, and if it is so, we always want to know the research area and the senior researcher that works for the project.

In your design you should follow the methodology adopted in the course, and you should produce:

1. [7%] the restructured Entity-Relationship 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 containing the tables `Worker(ssn,age,department)`, where null values are allowed for the attributes `age` and `department`, and `Purchase(ssn,item,year)`, where null values are not allowed. The first table stores the ssn (primary key), age, and department of a set of workers. Each tuple in the second table represents a purchase made by a worker for an item in a certain year. Note that a purchase p might have a "twin", i.e., a purchase done by the same worker of the same item in the same year as those of p. We know that the database satisfies the foreign key constraint from `Purchase[ssn]` to `Worker[ssn]`.

1. Express *in SQL* a query that calculates how many purchases in the `Purchase` table have no twins.
2. Express in *relational algebra* a query that computes the ssn of workers who have done at least one purchase starting from 2020, and for whom if the department is not known, the age is known.

Problem 4 [10%] Consider the relational schema S consisting of two relations `P(X,Y)` and `Q(Z,W)`, where:

1. `P` has `X` as primary key and `Y` as additional key;
2. the default value for attribute `Y` of `P` is `1`;
3. `Q` has `Z` as primary key;
4. `P` and `Q` have mutual foreign keys with delete policies as follows:
   `P(X,Y) UNIQUE Y, DEFAULT Y=1, FOREIGN KEY P[Y] ⊆ Q[Z] ON DELETE SET DEFAULT`
   `Q(Z,W) FOREIGN KEY Q[W] ⊆ P[X] ON DELETE CASCADE`

Answer the following question:

Is there a database D consistent with S such that the deletion from the database D of a tuple t of `P` causes an error? If the answer is negative, motivate your answer. If the answer is positive, show such a database D and the corresponding tuple t and describe the reason for the error.