This is a closed book exam: the only resources allowed are blank paper, pens, and your head, but you may use a handwritten A4 page with information that you consider useful for solving the exam exercises. 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 for managing dishes offered by caterers who deliver them to venues, for which the following information is of interest. Of each *caterer*, we are interested in the taxcode (identifier) and the address. Of each *dish*, we are interested in the name and type (normal, vegetarian, vegan, etc.), which together identify the dish, the calories, and the caterers that offer the dish, each with the price at which they offer it. It should be noted that a dish must be offered by at least one caterer and can be offered by multiple caterers. Additionally, of each dish, we are interested in the *set of venues* where it is offered. This set must not be empty but can change over time, at most once per month, and we are interested in the date (i.e., the day, in addition to the month and year) since when the set of venues offers the dish. *Gourmet dishes* are a type of dish with special ingredients, for which one of the offering caterers acts as the main provider, who specifies a minimum price for the dish. Of each gourmet dish we are interested in the number of special ingredients, the main provider, and the minimum price he has specified for the dish. Notice that no caterer might offer a gourmet dish at a lower price than the minimum price set by the main provider. Finally, of each *venue* we are interested in the name (identifier) and the city where it is located.

Problem 2 [40%] Carry out the logical design of the database, producing the complete relational schema with constraints, taking into account the following indications: *(i)* We access gourmet dishes separately from ordinary dishes. *(ii)* Every time we access a gourmet dish, we always want to know who is the main provider, the number of special ingredients, and also the calories.

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. [8%] the restructured relational schema (again with constraints).

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

Problem 3 [15%] Consider a database D containing the two relations:

(i) `Activity(code,type)`, which stores the code and the type (a string) of the activities offered by a resort;
(ii) `Schedule(guest,activitycode,time)`, with a foreign key constraint from `activitycode` to `code` of `Activity`, which stores the activities that guests have scheduled to perform at the resort.

Note that the set of all activity types is given by all types that appear as the type of some activity that is being offered by the resort.

1. Write in **relational algebra** a query over D that computes all pairs (g, t) such that no activity of type t appears in the schedule of guest g.
2. Write in **SQL** a query over D that computes all guests who have in their schedule at least one activity of each type.

Problem 4 [15%] Provide the definition of integrity constraint and of referential integrity constraint.

Then consider the following relational schema S (in which primary keys are underlined and attributes with an asterisk may contain NULL):

$R(\underline{A}, B^*, C)$
tuple constraint: B `IS NOT NULL` implies $B = A$

$Q(D, E, \underline{F})$
foreign key: $Q[D, E, F] \subseteq R[A, B, C]$

For each of the following three items, say whether there exists at least one database instance that is correct with respect to schema S with the corresponding property:

1. the set $\{B\}$ is not a superkey for R;
2. the set $\{C\}$ is not a superkey for R;
3. the set $\{D, E\}$ is not a superkey for Q.

For each of the three items listed above, if the answer is positive, then show the corresponding database instance; if the answer is negative, then describe in detail the reason why a database instance that is correct with respect to schema S does not exist.