

Data Management for Big Data - Exam

12 June 2020

Exercise 1:

Let us consider the following relational schema about physicians and departments:

PHYSICIAN(*PhysicianId*, *Name*, *Surname*, *Specialization*, *Gender*, *BirthDate*, *Department*);

DEPARTMENT(*Name*, *Building*, *Floor*, *Chief*)

Let every physician be univocally identified by a code and characterized by a name, a surname, a specialization (we assume to record exactly one specialization for each physician), a gender, a birth date, and the relative department (each physician is assigned to one and only one department).

Let every department be univocally identified by a name and characterized by its location (building and floor) and a chief. Let us assume that a physician can be the chief of at most one department (the department he/she belongs to). We do not exclude the possibility for two distinct departments to be located at the same floor of the same building.

Define preliminarily primary keys, other candidate keys (if any), and foreign keys (if any). Then, formulate an SQL query to compute the following data (exploiting aggregate functions only if they are strictly necessary):

- the departments with both male and female physicians, whose physicians are all born after 1955, that is, they are all at most 64 years old.

Exercise 2:

Let us synthesize the ER conceptual schema of a database for a hospital on the basis of the following set of requirements.

- The hospital consists of a number of departments. Each department is characterized by a medical specialty, that univocally identifies it (we assume that there is at most one department for each medical specialty), a location (building and floor), a set of physicians, a chief (one of the physicians), and a set of nurses. We do not exclude the possibility for two distinct departments to be located at the same floor of the same building.
- Each physician is univocally identified by an identifier, and is characterized by a name, a surname, one or more specializations, a gender, a birth date, and the relative department (each physician is assigned to one and only one department).
- Each nurse is univocally identified by an identifier, and is characterized by a name, a surname, a gender, a birth date, the set of training courses he/she attended over the years, and the relative department (each nurse is assigned to one and only one department).
- We would like to record existing wife/husband relationships between pairs of physician/physician, nurse/nurse, and physician/nurse. relationships.

Build an ER schema that describes the above requirements, clearly explaining any assumption you made. In particular, for each entity, identify its possible keys, and carefully specify the constraints associated with each relation.