

Introduction to Databases

Exercise.

The AirQ company monitors and analyzes the quality of air in cities and wants to design a database to manage its activities.

- AirQ analyzes the concentration of various pollutants such as carbon monoxide and nitrogen dioxide. Each pollutant is identified by a code and it is characterized by its name, a brief description, the measurement unit and a list of the main causes of production of the pollutant (e.g. vehicular traffic, heating systems, industrial activity).
 - The employees of the AirQ company are identified by their social security number (SSN). For each employee, the name, the hire date, the mobile phone number and the e-mail address (if available) are known. Employees are classified as administrative staff, technical staff and analysts. For analysts, the qualification is recorded. The database stores the holiday periods (start date, end date) for each employee. Note that different employees may be on holiday at the same time.
 - A network of fixed stations is exploited to monitor the concentration of pollutants. Each station is characterized by a unique code and its geographical position, expressed in terms of latitude and longitude. Each station includes different sensors. Each sensor is identified by a unique code and it is characterized by the station where it is placed and the monitored pollutant.
 - The sensor measurements are stored in the database. Each measurement has a value and it is identified by the sensor, the date and the time when it was collected. Each month a report summarizing measured concentrations for various pollutants is generated. Each report is identified by a unique code and is characterized by the release date and the list of measurements included in the report itself. Each measurement is used in at most a report. For each report, the analyst who performed the verification and validation of the report is also recorded.
 - In case of a station failure, a maintenance operation for the station is carried out. For each maintenance operation the database stores the date, the start and end time of the intervention, the station involved, the technician who carried out the maintenance and the issues that caused the failure. Please consider that different maintenance operations can be performed for the same station on the same date. However, a technician cannot perform two or more maintenance operations simultaneously.
1. Describe the conceptual schema of a database for the above application by means of an ER diagram.
 2. Derive a normalized relational logical schema for the same database.
 3. Define referential integrity constraints for 3 relations of your choice among those defined in the conceptual schema.



