

## Coursework Instructions Milestone 2: Constructing and Populating a Database

The goal of the second part of your IDA project is to

1. translate the ER-diagram of your project into a relational schema;
2. implement the schema as a PostgreSQL database; and
3. fill your PostgreSQL database with data.

For this milestone, you will have an occasion to apply the material of the lectures on (1) the relational model, (2) the translation of ER models into relational schemas, and (3) the basic data definition and manipulation facilities of SQL.

In the following, we will describe the single steps in more detail.

### Translation into Relational Schema

Write up a relational schema that is equivalent to your ER diagram. Make sure that you translate not only entities but also relationships and that you also capture, whenever possible, the integrity constraints of your ER diagram.

Explain briefly how the various entities and relations are expressed in the relational schema. Whenever there was a choice for the translation, explain the reasons for the choice you made.

During the translation you may find out that the relational schema resulting from the translation does not make a lot of sense. In this case, revise your ER diagram and the data requirements. It is imperative that all documents are consistent with each other.

**Note:** If you revise your data requirements and the ER diagram, submit the revised version with your deliverables for this milestone and do not change the original submission.

## Implementation of the Schema in PostgreSQL

Write an SQL script with the commands to create the tables of your database. An SQL script is simply a file with SQL commands. You find an example for a script in the file `solutions.sql` that you can reach from the Web page with the exercises of this course.

The `CREATE TABLE` statements must specify

1. appropriate types for the attributes;
2. the primary key;
3. constraints such as `NOT NULL` and `UNIQUE` whenever appropriate;
4. default values where appropriate;
5. `FOREIGN KEY` constraints, together with the policy for reacting to changes (remember that the default is `ON DELETE NO ACTION`)

Write comments into the script that explain the rationale behind the definition of your constraints. Comments in SQL can be written in two ways:

- over several lines enclosed in `/*...*/`, i.e., in the form `/*Comment*/`
- in a single line after two hyphens, i.e., in the form `--Comment`

## Loading Data

Load data into the database. Write a second script that contains `INSERT` statements for your relations. You may have to pay attention to insert records in such an order that the foreign key constraints are not violated.

## Deliverables

The deliverable will be a document consisting of three parts:

1. a file with the relational schema and an explanation how it has been derived from the ER diagram;
2. an SQL script that creates the tables of your database and contains comments on the constraints etc. you have chosen;
3. a second script with SQL `insert` statements.

## **Deadline and Submission**

The work is to be submitted by publishing it on the web site of your IDA group. The deadline is

Tuesday, 9 November, 11:30 pm.