

Introduction to Databases, ITU, Fall 2005

Rasmus Pagh

September 27, 2005

Group project – part 2

This is the second part of the **mandatory** group project. It should be handed in **by e-mail** to the teaching assistant associated with your group, no later than:

Monday October 17, 23.59 PM.

Learning outcomes

The first and second group hand-ins develop competences in two areas mentioned in the goals of the course description. After working on these hand-ins you should have improved your ability to:

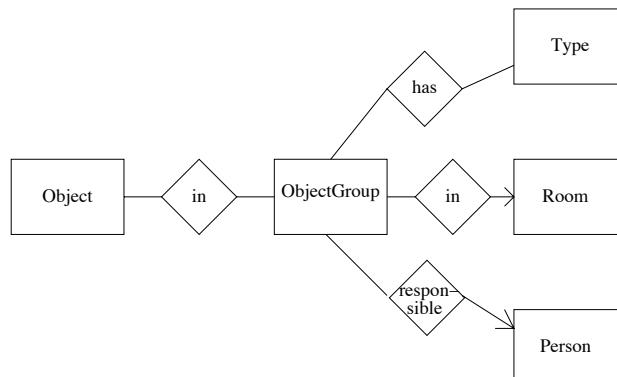
- make a database design in the E/R model [...] and convert it to the relational model
- use theoretical tools such as normalization [...] to improve database design [...]

Case description – continued

The hand-in concerns the continued development of the new database of IT@, for which your initial E/R diagram is now under scrutiny. Note that you will be having a meeting with either Simer or Priya on October 3, some time between 4 PM and 6.30 PM, to get feedback on this E/R diagram (the exact time for your group will be announced on the course home page).

However, there are already lots of things to do. It was decided to integrate an existing database, concerning inventory, with the new system. The aim of the inventory database is to keep track of all objects (e.g., chairs, tables, and computers) that may be moved around IT@. Objects are grouped according to type, according to in what room they are located, and according to who is responsible for the objects in that room. (E.g., all the chairs in a room will typically be of the same type, and hence form a group.)

The database was developed by *Rapid Consulting* (who offered a very advantageous price) and is described in the following. The E/R diagram (omitting attributes of entity sets), and the corresponding relation schema are shown below:



```

Object(ObjectID,buyDate,supplier)
ObjectGroup(ObjectGroupID,inRoom,responsible)
Type(TypeID,description)
Room(RoomID,type,capacity,AVequipment,comments)
Person(PersonID,FirstName,LastName)
ObjectInObjectGroup(ObjectID,ObjectGroupID)
ObjectGroupHasType(ObjectGroupID,TypeID)

```

Having studied *Introduction to Databases* you will wonder if this is really in BCNF. If not, you should perform the normalization, and change the E/R diagram such that it reflects the normalized design. Also, you may consider whether the E/R diagram can be simplified, or improved in some other way.

After the meeting on October 3, you should be able to make a final, complete E/R diagram of the whole database. This should be converted into relations, using the method discussed in the course. Note that for super/subclass relationships you will have to choose among several methods.

In the third hand-in you will be doing SQL programming in Oracle. Groups of 3 or 4 students must create the relations and insert at least 5 tuples for each relation corresponding to an entity set, and at least 10 tuples for each relation corresponding to a relationship. Make sure that there are no references to tuples that do not exist (e.g. an exam result for a student not in the database). In hand-in 3, groups of 1 or 2 will get data to work on by importing data entered by other groups.

What must be handed in

You must hand in a final data model for the system, including:

- **E/R diagram.** An E/R diagram for the system, using the notation of GUW. For clarity you may wish to omit attributes from this diagram. You should explain the process behind any changes you made to the E/R diagram by *Rapid Consulting*.
- **Description of entities, relationships and their attributes.** A short description in words for all parts of the E/R diagram whose meaning is not **completely** obvious.
- **Relational data model.** This should be derived from the E/R diagram. If you made any choices, e.g. in connection with super/subclasses, you should briefly explain your choice. The relations must be **in BCNF**. If you end up with relations that are not, you should modify the E/R diagram such that this is no longer the case. Unless you make special assumptions about the data, you do **not** need to argue that the relations are in BCNF.

Furthermore, groups of 3 and 4 must hand-in documentation for the relation instances created:

- **Relation instances.** The easiest way of doing this is to run `SELECT * FROM ...` queries in Oracle, and record the database activities into a file. This can be done using the command `spool myfile.txt` in SQL*Plus, after which the terminal output is recorded in the file `myfile.txt` in the Unix file system (written when you exit SQL*Plus). To retrieve the file use an SSH client (with file transfer mode), or move it into your `public_html` directory (viewable with a web browser).
- **Oracle user name.** The user name of an Oracle account at ITU where the relations reside. You don't have to provide a password. Instead, you should make the relations viewable by others by issuing the command `GRANT SELECT ON <MyRelation> TO public` a number of times, substituting `<MyRelation>` with the name of each of your relations.

On the first page, clearly specify the members of your group (if someone dropped the project, don't include him/her). The project should be sent as a **single file in PDF format** to your TA, either Simer Sawhney (`simer@aai-tee-yuu.dk`) or Priya Seetharaman (`itspri@aai-tee-yuu.dk`). (Note that the domain name must be changed to `itu.dk`.) Observe that this hand-in will be made available to the other groups, on the internal course pages, in connection with the third group hand-in.