Inf1B Data and Analysis
Tutorial 2 (week 4)

29 January 2010

- Please answer all questions on this worksheet in advance of the tutorial, and bring with you all work, including printouts of any computer work. Tutorials cannot function properly unless you do the work in advance.

- Data & Analysis tutorial exercises are not assessed, but are a compulsory and important part of the course. If you do not do the exercises then you are unlikely to pass the exam.

- Attendance at tutorials is obligatory; please let your tutor know if you cannot attend.

- *Required Reading:* Chapter 3 (The Relational Model) of 'Database Management Systems' (Raghu Ramakrishnan and Johannes Gehrke, 2003). The relevant pages of these chapters can be found on Google Books (`http://books.google.com/`).

Introduction

In the previous tutorial, you designed an ER model for a database, based on a description of a poster exhibition scenario. In this tutorial, you are asked to map the ER model to a relational model, using the techniques described in the lectures.

Question 1 - Relational Schema Notation

There are 5 mistakes in the SQL DDL schema below, what are they?

```
create table Film (
     filmId          integer,
     title           char(50),
```

```
year            integer
director name   char(50) not null,
language        char(30),
primary key (filmId),
foreign key director name references Director,
foreign key (leading_actor) references Actor
```

Question 2 - Mapping ER to relational schemata

For this question, please use the ER model provided at the end of this tutorial sheet, which is a simplification of the solution for last week's exercise. Think about how the entity sets and relationship sets can be mapped to tables in a relational model. Then state the SQL `create table` statements to define *all* the relations required to create this database. Note that in these SQL statements you will also need to define any *key* and *foreign key* constraints. Try to implement as many constraints as possible from the diagram; however, you are not required to implement all constraints, if doing so does not seem possible.

Things to think about: What different approaches are there for mapping relationship sets with key constraints? What are the advantages and disadvantages of each approach? Which approach did you choose for translating the provided ER model and why?

Tutorial Discussion - Thinking around the model

(a) In your schemas, which fields are not allowed to take a *null* value by default? Are there any that you should disallow from taking a null value? What constraints can you establish by preventing the fields from taking a *null* value?

(b) Consider the various foreign key constraints in your relational model design. In each case, what action would it be appropriate to instruct SQL to take on delete?

(c) Suppose each graphic designer has to create exactly one poster. How would you represent this in your ER diagram? Would you be able to represent it in your relational model and if yes, how?

(d) Reflect on your original relational model design. Can you implement the total participation constraint of *Poster* on the relationship sets *Creates* and *Judges*? If yes, how?

(e) According to the original scenario: "Each poster is judged by three different judges." Can you model this in your relational model? If yes, how?

(f) Based on your answers for the above questions, what do you think about the expressivity of the relational model compared to the ER model given a specific scenario?

(g) Based on your answers for the above questions, which is the best strategy to come up with a relational model for a given scenario? By directly translating the corresponding ER model, by directly translating the scenario, or as a combination of the two?