

M358 October 2003

Assignment Booklet IV

Please send all your answers to each tutor-marked assignment (TMA), together with an appropriately completed assignment form (TMA form), to reach your tutor on or before the cut-off date.

If you wish your tutor to acknowledge receipt of your TMA, complete an assignment acknowledgment card, affix the appropriate postage and send it to your tutor together with your TMA.

You will find instructions on how to fill in the TMA form in the current Student Handbook. Remember to fill in the correct Assignment Number as listed above and allow sufficient time in the post for each assignment to reach its destination on or before the cut-off date.

Please note that there will be no extension for TMA04 and you should also make sure all marked TMAs are returned before the examination; check with your tutor or registry if you suspect any missing TMA.

Question 1 (13 marks)

Before attempting this part you will need to set the SQL. Anywhere terminator character, normally a semicolon (;), to an exclamation mark (!). Details of how and why this is done are given at the beginning of the filed SQL statements for Section 6. Do not forget to change it back to a semicolon after you have answered this question.

- (i) Using the **create_city_pop** example as a model, complete the following definition of a new SQL function. The new function, **births**, takes a country name as a parameter and returns the number of births recorded in the **country** table for that country.

```
CREATE FUNCTION births (country_name VARCHAR(16))
    RETURNS DECIMAL (4,1) [8]
```

- (ii) Give SQL function calls that determine the birth rates for Italy and Monaco. [5]

Question 2 (35 marks)

The relation *Vetyield* has the following heading:

Vetyield (PlotNo, VarietyName, PlotSize, MemberNo, MemberName, VegetableType, ClubRecord, Yield)

The relation *Vetyield* represents facts about the vegetables grown on the plots of an allotment club. It records details of plots, the members who work them and the vegetables grown on the plots during the previous growing season. This includes the club record for an individual type of vegetable and the yield for each plot. More specifically, *Vetyield* represents the following single-valued facts (SVFs).

- 1 A plot is of a particular size.
- 2 A plot is worked by a member.
- 3 A member has a name.
- 4 A named variety grown on a plot produces a yield.
- 5 A named variety is of a particular vegetable type.
- 6 A vegetable type has a club record.

A representative extension of *Vetyield* is shown below.

PlotNo	VarietyName	PlotSize	MemberNo	MemberName	VegetableType	ClubRecord	Yield
12	King Edward	200	m1	Smith	Potato	2.82	350
12	Favourite	200	m1	Smith	Carrot	2.04	180
12	Pioneer	200	m1	Smith	Pea	2.48	95
15	Favourite	180	m2	Archer	Carrot	2.04	128
15	Meteor	180	m2	Archer	Pea	2.48	114
24	Foremost	150	m2	Archer	Potato	2.82	401
24	Baton Vert	150	m2	Archer	Cucumber	3.31	88
17	Pioneer	64	m3	Jones	Pea	2.48	97
28	Autumn King	180	m4	Smith	Carrot	2.04	95
28	St Valery	180	m4	Smith	Carrot	2.04	115
28	Snow White	180	m4	Smith	Beetroot	9.54	128

An extension of *Vetyield*

- (a) Write down the six FDs that correspond to the six SVFs given above. [9]

- (b) Write down any additional FDs that can be derived from those given in your answer to part (a), with an explanation of the derivation. [4]
- (c) Produce a set of relations in BCNF that is equivalent to *Vetyield*, showing clearly how you obtained your answer. [12]
- (d) Draw an E-R diagram for the conceptual data model expressed by the BCNF version of *Vetyield*. You are not expected to show participation conditions. [10]

Question 3 (40 marks)

A manufacturing company has a number of sites at which it makes products that are precision-engineered to meet the special needs of each of its customers. The data requirements for managing the manufacture and distribution of these products are as follows.

A site is identified by a site code, which is recorded together with the site's address and telephone number. All sites include a factory that makes many different kinds of product. In addition, some sites have a warehouse holding stocks of products, which may have been made in a factory in a different site, and are responsible for delivering products to customers.


Each kind of product has a unique part number, a short name and a longer description, and is manufactured in just one factory for a single customer. When it is manufactured, each product is kept in stock at one of the warehouse sites, from which it is delivered according to the instructions of the customer for whom it was made. It is essential to know where and how many of each kind of product is held in stock.

A customer is assigned a customer identifier, which is recorded together with their name, address and telephone number. Many kinds of product may be made for each customer.

The company has a fleet of vehicles for moving and delivering products. Each vehicle is assigned to one of the sites responsible for deliveries, and its registration number, make, model and capacity need to be recorded.

Given the incomplete E-R data model, you are required to complete this E-R data model to represent those data requirements as specified in the question. (Hints: consider the relevant part(s) of the E-R data model such as E-R diagram, entity types, constraints and any assumptions that you consider necessary. You should also include the participation condition(s) in your E-R diagram.)

[40]



Entity types

- Customer (...)
- Factory (...)
- DeliverySite (...)
- Vehicle (...)
- Product (...)
- Stock (...)

Constraints

.... ...

Assumptions


- A delivery site does not have to have any vehicles assigned to it.
- A factory may not be making any products.
-

The incomplete E-R data model for Question 3

Question 4 (12 marks)

It is possible for a relationship in a conceptual data model to have mandatory participation of both the entity types involved.

Consider the fragment of a conceptual data model given in Figure 1.



Entity types

A (a1, a2)
B (b1, b2)

Figure 1 Fragment of a data model

Assume that an SQL database is designed as a direct relational representation of the given data model – that is, there is a table **a**, having columns **a1** and **a2**, and a table **b**, having columns **b1**, **b2** and **a1**, each defined with appropriate properties, such as its primary key.

Explain how you would express each of the two mandatory participation conditions of this data model as an SQL constraint. [12]

End