

ISYS1055/1057 Database Concepts S1/2018 Assignment 2

This is an individual assignment. Plagiarism in oral, written or visual presentations is the presentation of the work, idea or creation of another person, without appropriate referencing, as though it's one's own. Plagiarism is not acceptable and may result in charges of academic misconduct which carries a range of penalties. It is also a disciplinary offence for students to allow their work to be plagiarised by another student.

Submission

Final submission is due at 23:59 21 May 2018 Monday in Week 12. Submit via the assignment submission system on Canvas a .zip file named after your student number (e.g., S12345.zip) that **contains** files for answering each question: Q1 (.pdf), Q2 (.pdf), Q3 (.sql, in plain text format), Q4 (.pdf) and Q5 (.pdf). The penalty for late submission is 10 marks per day or part day. After five days, assignments get 0 marks.

- **Your submission is successful only when you see your submitted file in your Canvas account. No confirmation email is sent to you.**
- Never leave submission to the last minute -- you may have difficulty uploading files.
- You can submit multiple times – a new submission will override any earlier submissions.

Assessment

There are five questions of 100 (10+10+40+20+20 =100) marks in total. This assignment is worth 30% of the overall assessment for ISYS1057 and ISYS1055. The assessment components and weights for the course are:

Assignment 1	Assignment 2	Exam
20%	30%	50%

There is not a hurdle for any assessment component.

Question 1. The Relational model (10 marks).

A database is needed to keep data for the booking systems of the ABC Clinic. Consider the below database schema of one relation including attributes for doctors (doc-), patients (pat-) and appointments (app-).

ABC(doc-firstname, doc-surname, doc-gender, doc-rego, doc-qualification, pat-ID, pat-givenname, pat-surname, pat-gender, pat-DOB, pat-addr, pat-phone, app-ID, app-datetime, app-type)

- A doctor has a unique registration number (doc-rego) and is also described by name, gender and qualification.
- A patient is identified by a unique patient ID (pat-ID) and has other information.
- Each appointment by a patient with a doctor is assigned a unique appointment ID (app-ID). An appointment can be of the long or short type.

Answer questions:

- 1.1) (3 marks) Give likely FDs.
- 1.2) (3 marks) Give the candidate keys for the ABC relation. In your working, show how you develop the closure for each candidate key you have discovered..
- 1.3) (4 marks) Is the relation ABC in BCNF or 3NF? Explain your answer.

Question 2. Normalisation (10 marks).

Consider the two relations below. They are in BCNF with primary key attributes underlined:

Customer(custID, firstname, lastname)
Item(itemNo, desc, price)

A Transaction relation as below is proposed to keep data for orders. Each order is by one customer and it can contain multiple items with their quantities.

Transaction(custID, itemNo, orderID, quantity, discount, amount_due)

Given the FDs below:

orderID \rightarrow custID, amount_due, discount
orderID, itemNo \rightarrow quantity
custID, orderID \rightarrow amount_due, discount

Answer questions.

- 2.1) (4 marks) Give the minimal basis for the given FDs.
- 2.2) (2 marks) The Transaction relation is not in BCNF or 3NF. Give the reason.
- 2.3) (4 marks) Follow the BCNF/3NF decomposition algorithm to decompose Transaction into relations in BCNF or 3NF. Give the relations after decomposition and specify the primary key and any foreign keys for each relation.

Question 3. SQL (40 marks).

*In addition to the lecture notes, you should also study by yourself the SQL*Plus tutorial on Canvas (the Oracle section) and other resources for syntax and useful functions.*

The relational schema for the Academics database is as follows:

```
DEPARTMENT(deptnum, descrip, instname, deptname, state, postcode)
ACADEMIC(acnum, deptnum*, famname, givenname, initials, title)
PAPER(panum, title)
AUTHOR(panum*, acnum*)
FIELD(fieldnum, id, title)
INTEREST(fieldnum*, acnum*, descrip)
```

Some notes on the Academics database:

- An academic department belongs to one institution (instname) and often has many academics. An academic only works for one department.
- Research papers (PAPER) are often authored by several academics, and of course an academic often writes several papers (AUTHOR).
- A research field (FIELD) often attracts many academics and an academic can have interest in several research fields (INTEREST).

Primary keys are underlined and foreign keys are marked with *. You should download the SQL script for defining and populating the database *academics.sql* from Canvas (the Oracle section) and run *academics.sql* in your Oracle account to build the database.

Write ONE SQL query for each of questions 3.1)–3.9). Put your answer for Question 3.10) in comments (starting each line with “--”).

Notes for marking:

- Each question is worth 4 marks. For questions with the “You must (not) ...” requirement, queries failing to meet the requirement receive maximum 1 mark. For example, question 3.1) has “You must use the NATURAL JOIN operator”. A query not using the NATURAL JOIN operator receives maximum 1 mark.
- Do not include the result of the query or the script used to create the tables.
- Your query should not output duplicates but use DISTINCT only if necessary.
- Queries are marked in terms of both correctness and efficiency. Unnecessary joins will incur deduction.

- 3.1) List the deptnum and total number of academics for CS departments, in alphabetical order of deptname. CS departments are departments whose deptname contains the phrase "Computer ... Science" or "Computing ... Science" in upper case or lower case letters. You must use the NATURAL JOIN operator.
- 3.2) List research fields where at least one academic is interested in. List the fieldnum, ID and title of these research fields. You must use a subquery.
- 3.3) Find papers that have three or more authors. Give the panum, title and number of authors for these papers.
- 3.4) For EACH academic, compute the total number of papers s/he has written. Output should include the acnum and total number of papers for each academic. In particular, an academic without any papers should have zero(0) as number of papers in the output. You must use a JOIN operator.
- 3.5) Give the total number of academics that do not have research interests. You must use the NOT IN operator.
- 3.6) Are there any research fields where less than 20, including zero, academics are interested in. List the fieldnum, ID, title and number of interested academics for these research fields.
- 3.7) Find the papers whose title contain the string 'data' and where at least one author is from the department with deptnum 100. List the panum and title of these papers. You must use the EXISTS operator. Ensure your query is case-insensitive.
- 3.8) Return the research interest that has the largest number of interested academics. You must not use MAX. Note: An SQL query that lists all research interests in decreasing order of their total number of interested academics is incorrect.
- 3.9) The following SQL query is intended to find academics (acnum) who are ONLY interested in "Data" (descrip) fields. But it is incorrect. Give the correct SQL query.

```

select acnum
from interest
where upper(descrip) like '%DATA%'

```
- 3.10) Consider the SQL query given below, give the English explanation for the output of a) the subquery, and b) the whole SQL query. Literal explanation will receive zero mark.

```

select distinct AC1.givenname, AC1.famname, AC2.givenname, AC2.famname
from academic AC1, author AU1, academic AC2, author AU2
where AC1.acnum=AU1.acnum
and AC2.acnum=AU2.acnum
and AU1.panum=AU2.panum
and AU2.acnum>AU1.acnum
and not exists
(select *
from Interest I1, Interest I2
where I1.acnum =AC1.acnum
and I2.acnum=AC2.acnum

```

and l1.fieldnum=l2.fieldnum);

Question 4. ER model (20 marks).

The A-Star consulting firm has decided to build a database. The company has hired you to design the database. Requirements are as follows.

- Clients are given a unique client-ID. Other client details including name, address and telephone number are also recorded.
- A-Star has several departments. Each department is described by a unique name, address and contact phone number. Each department must have a manager and has many employees. A manager can only manage one department.
- Each employee of A-Star has a unique employee ID, name, and birth date. An employee must belong to a department.
- An employee is assigned as the consultant for a group of clients. But a client has only one consultant.
- Once signed up, clients book consultation sessions with their consultant. All consultation sessions for a client are numbered sequentially as 1, 2, 3, etc. and the details for consultation sessions are also recorded, including date, time and topic.
- A-Star has vehicles for visiting clients. Each vehicle has a registration number, model, and year the vehicle was made. Consultants book vehicles for consultation sessions with clients.

The database is aimed to:

- Keep track of data for departments, employees, vehicles and clients.
- Keep track of the consultation sessions for clients and vehicle usage for booking vehicles for consultation.
- Keep track of scheduled consultation sessions for consultants.

According to the requirements and design aim, give an Entity Relationship (ER) diagram for the database, making assumptions where necessary. You must represent entities, relationships and their attributes, and all applicable constraints in your ER diagram. Explain any constraints that can not be expressed in the ER diagram.

- Your ER diagram must only use notations from the lecture notes and must not be hand drawn. ER diagrams using other notations will receive zero mark.
- You can use the ER diagramming tool Dia, which can be downloaded from Canvas (The ER section) and is also available as an App on mydesktop.rmit.edu.au. When exporting your ER diagram in Dia to a pdf file, you need to first set “page setup” to “Fit to 1 by 1” so that your pdf diagram scales properly. You can also use any other diagramming tool.

Question 5. ER to relational schema mapping. (20 marks).

Consider the Musician database ER diagram as shown in Figure 1.

- 5.1) (10 marks) Give the FDs for the constraints in the ER diagram. You should not include trivial or redundant FDs.
- 5.2) (10 marks) Map the ER diagram to a relational database schema following the ER-to-relational-database-schema mapping rules. Indicate the primary key (underline) and any foreign keys (asterisk) in each relation.

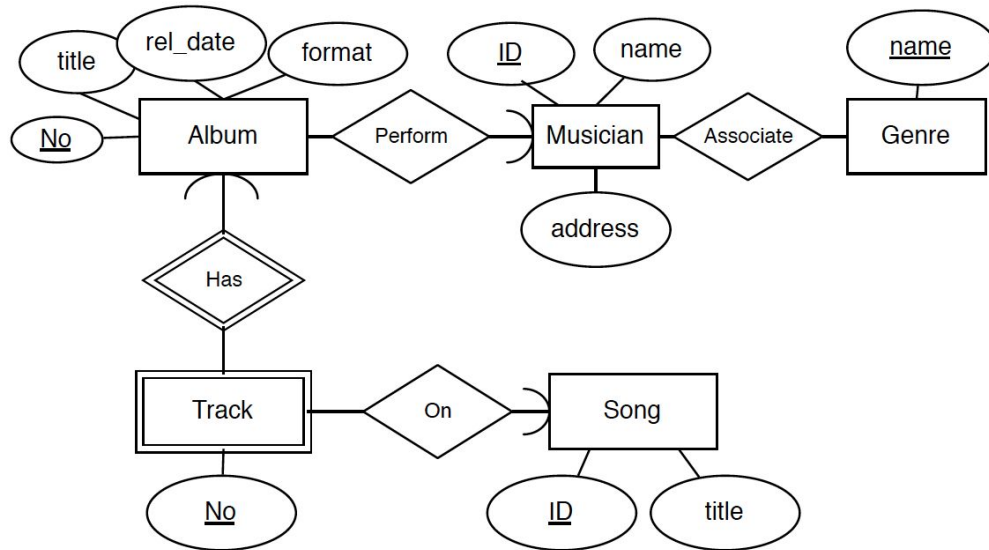


Fig. 1 The Musician database ER diagram