

**Test Pearl 010 — Databases**  
Pearls of Computer Science (201300070)  
Bachelor module 1.1, Technical Computer Science, EWI

22 september 2016, 08:45–09:45  
Module coordinator: Maurice van Keulen  
Instructor: Maurice van Keulen

- You may use 1 A4 document with your own notes for this exam and a *simple* calculator.
- Scientific or graphical calculators, laptops, mobile phones, books etc. are not allowed.  
***Put those in your bag now!***

Total number of points: 40.  
Total number of pages: 5.

**1 SQL (20 points)**

For this question, we use the tables belonging to a ‘shared notebook’ database. A description of the table structure can be found in Figure 2. The database stores notes of users (who each have an account). The notes are organized in notebooks. Moreover, individual notes as well as whole notebooks can be shared with other users.

The sharing mechanism perhaps needs some more explanation. If a user shares a whole notebook with another user, a row in ‘shared\_nb’ is created with the id of the notebook (nb\_id) and the id of the account (account\_id) the notebook is shared with. Moreover, the attribute ‘access’ is a string that either contains “edit” or “read” which determines whether the other user is given read/write or read-only access to the notebook, respectively.

Alternatively, it is also possible to share an individual note with another user. In that case, a row in shared\_note is created. This row additionally contains the notebook\_id of the notebook where the other user has placed this shared note.

Figure 3 contains some example data. There are 3 accounts for 3 pearl teachers you now know. Maurice has two notebooks, “Private” and “Teaching” of which the latter is shared in full with Pieter-Tjerk. The notebook “Private” contains a note “Quotes”. The notebook “Teaching” contains 2 notes: “TA Planning” and “Sign off list”, which are both shared with Arend who keeps them in his notebook “From Maurice”. Pieter-Tjerk has one notebook “My Notes” with two notes “Test” and “4ArendAndMaurice”. The latter note is shared with read access with both Arend (who keeps it in notebook “From Other”) and Maurice (who keeps it in notebook “Teaching”).

**Tip:** Figure 4 contains an informal description of the syntax of SQL.

- (a) Write an SQL-query that produces the names of all notebooks owned by user Maurice (i.e., disregard sharing).
- (b) Maurice (id 1) wants to share his note “Quotes” (id 100) with read access with Arend (id 3) who likes to keep it in notebook “From Maurice” (id 13). Write an SQL-statement that inserts the row(s) needed to establish this sharing. You may directly use the mentioned id’s.
- (c) Write an SQL-query that shows per user, how many notebooks are shared with the user.
- (d) Write an SQL-query that shows for user Arend the title of all the notes he has access to (so the ones he owns as well as the ones that are shared with him).



- (e) Write an SQL-query that shows per notebook name how many shared notes are in the notebook.
- (f) Pieter-Tjerk would like to change the access to 'edit' for everyone he shared his note "4ArendAndMaurice" with. Write an SQL-statement that makes this change in the database.

□

## 2 Databases (10 points)

For this questions, please use correct database terminology as much as possible.

- (a) Give an important reason why the ANSI/SPARC architecture distinguishes between the conceptual schema and the physical schema.
- (b) SQL is at the same time a QL, DML, DDL, and SDL. To which of these kinds of languages does the CREATE-statement belong?
- (c) In one transaction we execute the following four statements with the intent to transfer 200 euro from account A to B:
- a:=read account A
  - b:=read account B
  - put b+200 in account B
  - put a-200 in account A

Initially, A=100 and B=1000. Suppose an integrity constraint has been defined that the value in every account should be positive, i.e., above zero. The fourth statement, however, violates this constraint, i.e., it tries to put the value -100 in account A.

Explain what happens and give the values of both accounts after execution of this transaction.

- (d) What is the difference between *data distribution* and *data replication*?

□

## 3 Database design (10 points)

Figure 1 contains part of the datamodel belonging to the database of a social media provider. The main entity is 'Account' with an attribute 'name' for the user's name and 'fullname' contain his or her full name. Below the entity 'Account' you see a line going from 'Account' back to 'Account' with a '\*' at each end and a dashed line to 'Connection'. This represents the connections between accounts with their 'role' (e.g., "family" for connections between family members, or "business" for business relations). Every account obviously has messages (with a title and body) and a message may have zero or 1 picture associated with it. Finally, an account may also have a profile picture.

- (a) According to the ER model of Figure 1, is it possible for the same picture to be associated with a message as well as with an account as profile picture? Explain your answer.
- (b) According to the ER model of Figure 1, is it possible for the same picture to be the profile picture for two different accounts? Explain your answer.
- (c) Give a table structure for the ER-model of Figure 1. Do this in terms of a list of tables and attribute names. The attribute types may be omitted.

□



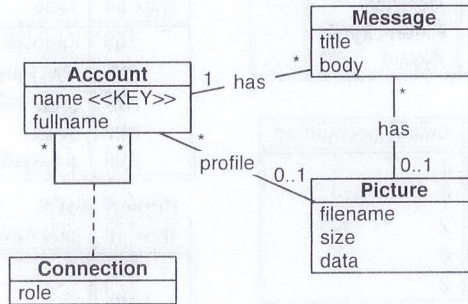


Figure 1: ER model of a social media provider’s database.

Table	Attributes / Description
account	<b>account_id</b> , name, user_name, picture, ... For every account a unique number (account_id; also primary key) and the name of the account, the user’s name, his/her picture, etc..
notebook	<b>nb_id</b> , name, owner_account_id For every notebook a unique number (nb_id; also primary key), the name of the notebook, and the account of the owner of the notebook (i.e., the account_id of the account the notebook belongs to).
note	<b>note_id</b> , title, body, nb_id For every note a unique number (note_id; also primary key), the title and body of the note, and the id of the notebook the note belongs to.
shared_nb	<b>nb_id</b> , <b>account_id</b> , access This table contains all sharing definitions for notebooks. A row specifies that notebook with id nb_id is shared with account account_id and provides either ‘edit’ or ‘read’ access.
shared_note	<b>note_id</b> , <b>account_id</b> , <b>nb_id</b> , access This table contains all sharing definitions for individual notes. A row specifies that note with id note_id is shared with account account_id who has placed it in his/her notebook with id nb_id, and provides either ‘edit’ or ‘read’ access.

Figure 2: Table structure of a “shared notebook” database (Primary keys in **bold**).



account			
account_id	name	user_name	...
1	mvkeulen	Maurice	...
2	ptdeboer	Pieter-Tjerk	...
3	arensink	Arend	...

notebook		
nb_id	name	owner_account_id
10	Private	1
11	Teaching	1
12	My Notes	2
13	From Maurice	3
14	From Other	3

shared_nb		
nb_id	account_id	access
11	2	edit

note			
note_id	title	body	nb_id
100	Quotes	...	10
101	TA Planning	...	11
102	Sign off list	...	11
103	Test	...	12
104	4ArendAndMaurice	...	12

shared_note			
note_id	account_id	nb_id	access
101	3	13	edit
102	3	13	edit
104	3	14	read
104	1	11	read

Figure 3: Example data for the “shared notebook” database

In the informal syntax, we use the following notations

- $A|B$  to indicate a choice between A and B
- $[A]$  to indicate that A is optional
- $A^*$  to indicate that A appears 0 or more times
- $A^+$  to indicate that A appears 1 or more times
- ‘A’ to indicate that the symbol A is literally that symbol

We are not precise in punctuation in the syntax, but this is irrelevant in this exam anyway.

**SQL**

*createtable*: CREATE TABLE *tablename* ‘(‘ *columndef*+ *constraint*\* ‘)’

*columndef*: *colname* *type* [NOT NULL] [UNIQUE] [PRIMARY KEY] [REFERENCES *tablename* (*colname*+)]

*constraint*: PRIMARY KEY (*colname*, ... ) | CHECK ( *condition* )  
 | FOREIGN KEY(*colname*, ... ) REFERENCES *tablename*(*colname*, ... )

*query*: SELECT ( *column* [ AS *colname* ] )+  
 FROM ( *tablename* [ AS *colname* ] )+  
 WHERE *condition* [ GROUP BY *column*+ ] [ ORDER BY *column*+ ]

*column*: [ *tablename* ‘.’ ] *colname* | ‘\*’

Examples of *condition*: *column* = *value* [ (OR | AND) [NOT] *column* <> *value* ]  
 | *column* IS [NOT] NULL  
 | *column* [NOT] IN (*value*, ... ) ...

Figure 4: Informal syntax of SQL