

FINAL EXAM: DATABASES ("BASES DE DATOS") – 4/9/2006 – SCHEMA

Consider the following relational schema, which will be referred to as WORKING SCHEMA, which maintains information on individual assessments which are performed to several applicants for a job.

APPLICANT(**app_code**: app_code_dom, **name**: name_dom, **NIF**: NIF_dom, **telephone**: tel_dom, **age**: age_dom)

PK: {app_code}

NNV: {name, NIF, telephone}

Unique: {NIF}

TOPIC(**topic_code**: topic_code_dom, **title**: title_dom, **difficulty**: diff_dom)

PK: {topic_code}

NNV: { title}

diff_dom: {'high', 'medium', 'low'}

CORRECTOR(**cor_code**: dom_cor_code, **name**: name_dom, **NIF**: NIF_dom, **telephone**: tel_dom, **age**: age_dom)

PK: {cor_code}

NNV: {name, NIF, telephone, age}

Unique: {NIF}

CORRECTION(**cor_code**: dom_cor_code, **topic_code**: topic_code_dom, **level**: level_dom)

PK: {cor_code, topic_code}

FK: {cor_code} → CORRECTOR

On cascade delete. On cascade update. Deferrable initially immediate

FK: {topic_code} → TOPIC

On cascade delete. On cascade update. Deferrable initially immediate

NNV: {level}

ASSESSMENT(**assessment_code**: assessment_code_dom, **time**: time_code, **app_code**: app_code_dom, **cor_code**: dom_cor_code, **topic_code**: topic_code_dom, **qualification**: qualification_dom)

PK: {assessment_code}

FK: {app_code} → APPLICANT

Restrictive delete. On cascade update. Deferrable initially immediate

FK: {cor_code, topic_code} → CORRECTION

Partial referential integrity.

On cascade delete. On cascade update. Deferrable initially immediate

NNV: {topic_code}

CREATE ASSERTION R1 CHECK (
NOT EXISTS (SELECT * FROM ASSESSMENT WHERE qualification is not null AND
(cor_code is null OR app_code is null)));

where the attributes and tables have the following meaning

APPLICANT: information of the people who apply for the job.

app_code, *name*, *NIF*, *telephone* and *age* represent the code, name, national identification number, telephone and age of the applicant respectively.

TOPIC: information on the several topics over which the assessments deal about.

topic_code, *title*, *difficulty* represent the code and the title of the topic, as well as the degree of difficulty ('high', 'medium' or 'low').

CORRECTOR: information of the people who will correct the assessments.

cor_code, *name*, *NIF*, *telephone* and *age* represent the code, name, NIF, telephone and age of the corrector respectively.

CORRECTION: information on which correctors correct each topic.

The corrector with code *cor_code* corrects the topic of code *topic_code* and his/her level of expertise is *level*.

ASSESSMENT: information of several assessments performed (or to be performed) by the applicants.

assessment_code, *time* and *qualification* represent the code of the assessment, the time which is assigned to perform the assessment and the obtained qualification; *app_code* represents the applicant who has to do the assessment; *topic_code* represents the topic which is questioned in the assessment; and, finally, *cor_code* represents the corrector which corrects the assessment.

And consider the following extension of the previous schema. We will refer to this extension as database (DB). The blank cells represent null values.

APPLICANT					TOPIC		
app_code	name	NIF	telephone	age	topic_code	title	difficulty
1	Marta	X25	146484	32	t1	Law	high
2	Juan	Y78	465418	28	t2	Marketing	high
3	Mateo	Z98	165481	25	t3	History	low
4	Alodia	Ñ97	44664	30	t4	Economy	medium

CORRECTOR					CORRECTION		
cor_code	name	NIF	telephone	age	cor_code	topic_code	level
1	Ana	G45	784646	49	1	t1	5
2	Eva	I45	741519	56	1	t2	3
3	Olga	B71	4896776	32	1	t3	2
					2	t1	2
					2	t3	5
					3	t1	4

ASSESSMENT					
assessment_code	time	app_code	cor_code	topic_code	qualification
1	2	1		t1	
2	2:30	1	1	t2	
3	3	2	1	t3	7
4	2	3	1	t3	6
5	3:30	3		t2	
6	2			t1	
7	3			t3	

Final exam: "databases" – 4/9/2006 – QUESTIONNAIRE TYPE A

This questionnaire has 14 questions; for each one we propose four possible answers. Only one of them is correct. The answer must be included in the answer sheet which has been handed with the exam. The maximum mark for the questionnaire is 3.5 points. The result is obtained through the formula: $(\text{Right} - \text{Wrong}/3) \times 0.25$.

1. Given the working schema, independently of the database DB, what would the result be of executing the instruction "DELETE FROM APPLICANT;"?
 - a) All the applicants will always be deleted.
 - b) Only the applicants without assessments will be deleted.
 - c) All the applicants and all the assessments which are assigned to the applicants will be deleted.
 - d) No applicant will be deleted if there is a tuple in the relation *Assessment* which is assigned to an applicant, since the instruction is atomic.

2. What is the fundamental goal of the multiple level architecture in a database management system?
 - a) to maintain data integrity.
 - b) to provide the property of data independence.
 - c) to ease the structuring of data and the definition of the associated operations.
 - d) to ensure data persistence.

3. Given the database DB, if we perform the following atomic instruction.
INSERT INTO ASSESSMENT VALUES (8, NULL, 3, NULL, t1, 10);
Which of the following statements is TRUE?
 - a) The insertion will be performed without problems.
 - b) The insertion is not performed because the attribute 'time' is null.
 - c) The insertion is not performed because there is no tuple in the relation *Correction* with null *cor_code*.
 - d) The insertion is not performed because a general constraint which is defined in the schema is violated.

4. A database management system:
 - a) is a structured data collection.
 - b) is a software tool which is constructed to apply database techniques.
 - c) is the basis for the definition of a data model.
 - d) always does without the operating system in the disk access, in order to significantly improve the information access.

5. Given the database DB, if we perform the following transaction:
COMMIT;
SET CONSTRAINT ALL DEFERRED;
UPDATE applicant SET app_code = 4 WHERE app_code = 3;
UPDATE applicant SET app_code = 3 WHERE app_code = 4;
COMMIT;
What is the state of the database when this transaction in standard SQL finishes?
 - a) Nothing will have changed because the final state is inconsistent because of the primary key of the relation Applicant.
 - b) Nothing will have changed because the "on cascade update" is deferred and the state is inconsistent because of the foreign keys.

c) The following extension:

APPLICANT				
app_code	name	NIF	telephone	age
1	Marta	X25	146484	32
2	Juan	Y78	465418	28
4	Mateo	Z98	165481	25
3	Alodia	Ñ97	44664	30

ASSESSMENT						
assessment_code	time	app_code	cor_code	topic_code	qualification	
1	2	1		t1		
2	2:30	1	1	t2		
3	3	2	1	t3	7	
4	2	4	1	t3	6	
5	3:30	4		t2		
6	2			t1		
7	3			t3		

The rest of tables doesn't vary.

d) The following extension:


APPLICANT				
app_code	name	NIF	telephone	age
1	Marta	X25	146484	32
2	Juan	Y78	465418	28
4	Mateo	Z98	165481	25
3	Alodia	Ñ97	44664	30

The rest of tables doesn't vary.

6. Assuming that the fundamental use of the relation *Assessment* is to obtain a list of all the qualifications of all the applicants sorted by applicant, which of the following physical implementations is more appropriate to store this relation?

- A hashed file using the code of the assessment as the hash field.
- A file sorted by the field code of the assessment.
- A file sorted by the field code of the applicant.
- A disordered file using the code of applicant as the hash field.

7. According to the following temporal diagram, and assuming a DBMS with immediate update, when a main memory loss occurs:



- T1, T2 and T3 will have to be redone.
- T1 will have to be redone and T2 and T3 will have to be undone.
- Half of T1 will have to be redone (until the checkpoint) and T2 and T3 will have to be undone.
- T1 and T2 will have to be undone and T3 will have to be redone until the checkpoint.

8. About transaction management, which of the following options is FALSE?
- a) The deferred update implies that the updates of a transaction over a database are postponed until this is confirmed.
 - b) The immediate update implies that the updates of a transaction over a database must be copied to disk before the transaction is confirmed.
 - c) The property of persistence implies that the changes which have been performed by confirmed transactions cannot be lost due to failures from other transaction or system failures.
 - d) The atomicity is a property which is not always respected by some systems.

9. U1, U2 and U3 are the users of a database which is conformed by relations R1 and R2. Exclusively taking the following SQL sentences into account:

GRANT INSERT ON R1 TO U1 WITH GRANT OPTION;

GRANT DELETE ON R2 TO U1;

GRANT ALL ON R2 TO U3;

Which of the following statements is FALSE?

- a) U3 can insert a tuple in R1 if he or she is authorised by U1.
 - b) U1 cannot give their permission for deletion to R2.
 - c) U1 is the only one who can delete tuples from R2.
 - d) U1 can insert a tuple in R1, but U2 and U3 cannot.
10. Regarding the physical storage of relational databases, which of the following options is FALSE?
- a) A file is a data structure for secondary memory storage which is conformed of a collection of records.
 - b) In an ordered file, the reading of all the records in the order of the sorting field is extremely efficient.
 - c) In a hashed file, the reading of all the records in the order of the hashing field is extremely efficient.
 - d) An index makes it possible to impose an order in a file without the need of actually sorting it.

11. Given the working schema. Which of the following statements is TRUE?

- a) We can always delete a tuple of *Correction*.
- b) We can delete a tuple of *Correction* if it has no assessments.
- c) We can delete a tuple of *Correction* if it has no topics.
- d) We can delete a tuple of *Correction* if it has no corrector.

12. Given the working schema, which of the following statements is TRUE?

- a) Every assessment has an associated applicant.
- b) Every assessment has an assigned time.
- c) Every assessment has an associated topic.
- d) Every assessment has an associated corrector.

13. Which of the following SQL instructions which try to define the relation CORRECTION is correct? (We assume that the domains have been previously defined)

- a) CREATE TABLE CORRECTION
(cor_code dom_cor_code PRIMARY KEY REFERENCES CORRECTOR
ON UPDATE CASCADE
ON DELETE CASCADE DEFERRABLE INITIALLY IMMEDIATE,
topic_code topic_code_dom PRIMARY KEY REFERENCES TOPIC
ON UPDATE CASCADE
ON DELETE CASCADE DEFERRABLE INITIALLY IMMEDIATE,
level level_dom NOT NULL);
- b) CREATE TABLE CORRECTION
(cor_code dom_cor_code REFERENCES CORRECTOR
ON UPDATE CASCADE
ON DELETE CASCADE DEFERRABLE INITIALLY IMMEDIATE,
topic_code topic_code_dom REFERENCES TOPIC
ON UPDATE CASCADE
ON DELETE CASCADE DEFERRABLE INITIALLY IMMEDIATE,
PRIMARY KEY(cor_code, topic_code),
level level_dom NOT NULL);
- c) CREATE TABLE CORRECTION
(cor_code dom_cor_code,
topic_code topic_code_dom,
level level_dom NOT NULL,
PRIMARY KEY(cor_code, topic_code),
FOREIGN KEY (cor_code, topic_code) REFERENCES (CORRECTOR, TOPIC)
ON UPDATE CASCADE
ON DELETE CASCADE DEFERRABLE INITIALLY IMMEDIATE);
- d) CREATE TABLE CORRECTION
(cor_code dom_cor_code PRIMARY KEY,
topic_code topic_code_dom PRIMARY KEY,
level level_dom NOT NULL,
FOREIGN KEY (cor_code, topic_code) REFERENCES (CORRECTOR, TOPIC)
ON UPDATE CASCADE
ON DELETE CASCADE DEFERRABLE INITIALLY IMMEDIATE);

14. What information is obtained with the following expression in relational algebra over the working schema?

$((\text{ASSESSMENT WHERE qualification} > 5 \wedge \text{topic_code} \neq \text{'t3'} [\text{cor_code}, \text{topic_code}] \bowtie \text{CORRECCION WHERE level} > 3) \bowtie \text{CORRECTOR})[\text{name}]$

- a) Name of the applicants who have obtained a qualification which is greater than 5 in the topics which are different from 't3', and the level of the corrector in the topic being greater than 3.
- b) Name of the correctors who have corrected all the assessments with a qualification which is greater than 5 in topics which are different from 't3', and the level of the corrector in the topic being greater than 3.
- c) Name of the correctors who have corrected some assessment with a qualification which is greater than 5 in topics which are different from 't3', and the level of the corrector in the topic being greater than 3.
- d) Name of the correctors who have only corrected assessments with a qualification which is greater than 5 in topics which are different from 't3', and the level of the corrector in the topic being greater than 3.

FINAL EXAM: DATABASES – 4/9/2006 – Problems

Given the working schema presented before, solve the following exercises in standard SQL:

- 1) Obtain the average qualification of the assessments whose correctors are 50 years old or more. (0.5 points)
- 2) Obtain the NIF, name and telephone of the applicants such that we know his/her telephone and haven't performed any assessment. (0.5 points)
- 3) Obtain the NIF and name of the applicants who have obtained a qualification of more than 7 in at least two assessments. (0.75 points)
- 4) Obtain, for every applicant, his/her NIF, name and the number of assessments he/she has performed. (0.75 points)
- 5) Obtain the NIF and the name of the applicants who have an average qualification in his/her assessments which is greater than 8. (0.75 points)
- 6) Obtain the NIF and the name of the applicants who have performed at least one assessment, and also that all his/her assessments are of topics of 'high' difficulty. (1 point)
- 7) Obtain the NIF and the name of the applicant who has the highest average qualification. Also show the number of assessments he/she has performed. (1 points)
- 8) Assume that we add to the relation *Applicant* a derived attribute *total_qualification* which has to be obtained by summing up all the qualifications of the assessments which have been performed by the applicant:
 - a. Apart from the insertion in *Assessment*, write down which operations affect the value of the derived attribute *total_qualification* of the relation *Applicant*. (0.5 points)
 - b. Write a trigger to handle the insertion operation in *Assessment*. (0.75 points)

SOLUTIONS TO THE QUESTIONNAIRE:

1	2	3	4	5	6	7	8	9	10	11	12	13	14
d	b	d	b	a	c	b	b	c	c	a	c	b	c

SOLUTIONS TO THE PROBLEMS:

- 1)

```
SELECT AVG(P.qualification)
FROM Assessment P
WHERE P.cor_code IN
      (SELECT C.cor_code
       FROM Corrector C
       WHERE C.age>50);
```
- 2)

```
SELECT NIF, name, telephone
FROM Applicant
WHERE telephone IS NOT NULL
      AND app_code NOT IN (SELECT app_code FROM Assessment)
```
- 3)

```
SELECT C.NIF, C.name
FROM Applicant C
WHERE 2 < (SELECT COUNT(*)
           FROM Assessment P
           WHERE C.app_code =P.app_code AND P.qualification >7)
```
- 4)

```
SELECT C.NIF, C.name, COUNT(P.assessment_code)
FROM Applicant C LEFT JOIN Assessment P ON P.app_code=C.app_code
GROUP BY C.NIF, C.name
```
- 5)

```
SELECT C.NIF, C.name
FROM Applicant C, Assessment P
WHERE C.app_code =P.app_code
GROUP BY C.NIF, C.name
HAVING 8 < AVG(P.qualification)
```
- 6)

```
SELECT C.NIF, C.name
FROM Applicant C
WHERE C.app_code IN (SELECT app_code FROM Assessment)
      AND
      NOT EXISTS (SELECT * FROM Assessment P,
                  WHERE P.app_code=C.app_code AND
                        NOT EXISTS (SELECT T.topic_code
                                   FROM TOPIC T
                                   WHERE P.topic_code=T.topic_code
                                         AND T.dificultad='High'))
```
- 7)

```
SELECT C.name, COUNT(*)
FROM Assessment P, Applicant C
WHERE P.app_code = C.app_code
GROUP BY C.app_code, C.name
HAVING AVG(P.qualification) = (SELECT MAX(AVG(P2.qualification))
                              FROM Assessment P2
                              GROUP BY P2.app_code)
```
- 8)

a) In this question, the student had to list the operations which might affect the derived attribute. In the following table, we also show the condition (if any) which must hold in order to react to the operation. We also show the action to be performed.

Operation	Condition	Action
Insert into <i>Applicant</i>	<i>total_qualification</i> <> 0	Reject the insertion
Modify <i>total_qualification</i> in <i>Applicant</i>		Reject the update
Insert into <i>Assessment</i>	<i>app_code</i> and <i>qualification</i> NOT NULL	Increment, for the applicant who made the assessment, the value <i>total_qualification</i> with the value <i>qualification</i>
Update of <i>Assessment</i>	<i>app_code</i> and <i>qualification</i> NOT NULL	Decrement, for the applicant who made the assessment, the value <i>total_qualification</i> with the value <i>qualification</i>
Update <i>qualification</i> en <i>Assessment</i>	old value for <i>qualification</i> <> new value for <i>qualification</i>	Modify, for the applicant who made the assessment, the value <i>total_qualification</i> with the different between the values (old and new) in <i>qualification</i>
Modify <i>app_code</i> in <i>Assessment</i>	old value for <i>app_code</i> <> old value for <i>app_code</i>	Increment, for the new applicant, the value <i>total_qualification</i> with the value <i>qualification</i> and decrement, for the old applicant, the value <i>total_qualification</i> with the value <i>qualification</i>

b) CREATE TRIGGER INSERT_INTO_assessment
AFTER INSERT ON Assessment
FOR EACH ROW
WHEN (new.app_code is not null and new.qualification is not null)
BEGIN
UPDATE Applicant
SET total_qualification=total_qualification+:new.qualification
WHERE app_code= :new.app_code;
END