FINAL EXAM: DATABASES ("DATABASES") – 17/09/2010 – SCHEMA

Consider the following relational schema, which will be referred to as WORKING SCHEMA, which maintains information about art museums, their paintings and expositions.

PAINTINGS (**painting**: integer, **title**: string(40), **author**: integer, **year**: integer, **style**: string(20))
PK: {painting}
Uni: {title, author, year}
NNV: {title}
FK: {author} → AUTHORS
ON DELETE SET TO NULL
ON UPDATE CASCADE

AUTHORS (**author**: integer, **name**: string(40), **century**: string(5), **country**: string(20))
PK: {author}
NNV: {name}

MUSEUMS (**museum**: integer, **name**: string(40), **city**: string(20), **year**: integer, **num_paintings**: integer)
PK: {museum}
Uni: {name, city}

IS_IN (**painting**: integer, **museum**: integer, **on_loan**: boolean)
PK: {painting, museum}
FK: {painting} → PAINTINGS
RESTRICTED DELETION
RESTRICTED UPDATE
FK: {museum} → MUSEUMS
RESTRICTED DELETION
ON UPDATE CASCADE

EXPOSITIONS (**exposition**: integer, **name**: string(40), **museum**: integer, **style**: string(20), **start_date**: date, **closing_date**: date, **num_paintings**: integer)
PK: {exposition, museum}
FK: {museum} → MUSEUMS
RESTRICTED DELETION
ON UPDATE CASCADE

COMPOSED_OF (**painting**: integer, **exposition**: integer, **museum**: integer)
PK: {painting, exposition, museum}
FK: {painting} → PAINTINGS
RESTRICTED DELETION
ON UPDATE CASCADE
FK: {exposition, museum} → EXPOSITIONS
RESTRICTED DELETION
ON UPDATE CASCADE
COMPLETE REFERENTIAL INTEGRITY

where the attributes and tables have the following meaning:

PAINTINGS: for each painting, we store its identifier (painting), its title, its author, the year in which it was painted and the pictorial style.

AUTHORS: for each painter, we store their identifier (author), their name, their century and their country.

MUSEUMS: for each museum, we store its identifier (museum), its name, its city, its opening year and the number of paintings it owns.

IS_IN: indicates that a painting is (or has been) in a museum. The attribute *on_loan* indicates whether the painting is owned by the museum (on_loan = false) or not (on_loan = true).

EXPOSITIONS: for each temporary exposition that is organised by a museum, we store a numerical identifier, the name of the exposition, the museum where it takes place, the style of the paintings in the exposition, the start and closing dates for the exposition, and the number of paintings on display in the exposition.

COMPOSED_OF: indicates whether a painting is on display in an exposition.

And consider the following extension of the previous schema. We will refer to this extension as database (DB). The symbol '?' represents null values:

Paintings

painting	title	author	year	style
1	Las Meninas	2	1656	Baroque
2	El barbero del Papa	2	?	Baroque
3	La maja desnuda	1	1800	Neoclassical
4	Retrato de Isabel Porcel	1	1805	Romantic
5	The birth of Venus	3	1484	Renaissance
6	La Gioconda	5	1506	Renaissance
7	Authorretrato	2	1650	Baroque

Is_in

painting	museum	on_loan
1	1	false
2	1	true
3	1	false
4	2	false
5	4	true
6	3	true
7	5	true
7	3	false

Authors

author	name	century	country
1	Francisco de Goya	XIX	Spain
2	Diego Velázquez	XVII	Spain
3	Sandro Botticelli	XV	Italy
4	Rembrandt	XVII	Holland
5	Leonardo Da Vinci	XV	Italy

Composed_Of

painting	exposition	museum
2	1	3
7	1	3
5	1	2
6	1	2

Museums

museum	name	city	year	num_paintings
1	Museum del Prado	Madrid	1819	2
2	National Gallery	London	1824	1
3	Musée du Louvre	Paris	1793	1
4	Galleria degli Uffizi	Florence	1765	0
5	Museum de Bellas Artes	Valencia	1839	0

Expositions

exposition	name	museum	style	start_date	closing_date	num_paintings
1	Velázquez and the baroque	3	Baroque	01/06/2010	15/12/2010	2
1	Renaissance Italy	2	Renaissance	01/09/2010	01/04/2011	2
2	German cubism	2	Cubist	01/02/2011	30/09/2011	0

Deliberate blank page

Deliberate blank page

FINAL EXAM: DATABASES – 17/09/2010 – Problems

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

1. Obtain the identifier and the title of all the paintings of style 'Baroque' for which we do not know the year in which they were painted. **(0.50 points)**

2. Create a general constraint in standard SQL to ensure that all the paintings that compose an exposition are of the same style that the style of the exposition. **(0.75 points)**

3. Obtain the name of the exposition, the name of the museum, and the start and closing dates, for the expositions that contain some paintings of the author with name 'Botticelli' and finish before '31/12/2010'. **(0.50 points)**

4. For all the authors, obtain their identifier, their name and the number of styles of their paintings. **(0.75 points)**

5. Obtain the identifier and the name of the authors of the paintings of style "renaissance" such that their paintings do not take part in any exposition of style "renaissance". **(0.75 points)**

6. Obtain the identifier and the name of the author whose paintings have taken part in expositions such that the overall number of museums of these expositions is the greatest. **(1 point)**

7. Obtain the identifier and the name of the Spanish authors such that all the museums which were inaugurated after 1800 own one or more of his/her paintings. **(1 point)**

8. The attribute *num_paintings* in the relation EXPOSITIONS is a derived attribute, which stores the total number of paintings that compose each exposition.

   a) Apart from the modification of the attributes *museum* and *exposition* in the relation COMPOSED_OF, please indicate the instructions that may affect this derived attribute. **(0.60 points)**

   b) Please write a trigger to handle the instruction: "modification of the attribute *museum* or the attribute *exposition* in the relation COMPOSED_OF". **(0.65 points)**

Final exam: "Databases" **– 17/09/2010** – 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 that has been handed with the exam. The maximum mark for the questionnaire is 3.5 points. The result is obtained by the formula: (Right – Wrong/3) × 0.25.

1. In the database DB of the working schema, if we add the following integrity constraint:

```
CREATE ASSERTION R1
CHECK (NOT EXISTS (SELECT * FROM Museums M
                   WHERE num_paintings <> (SELECT COUNT(*) FROM Is_in E
                                           WHERE E.museum=M.museum
                                             AND E.on_loan=false)));
```

Which is the set of instructions that can violate this integrity constraint?

a) In MUSEUMS: modify the attribute *num_paintings*, insert tuple. In IS_IN: modify the attribute *museum*, modify the attribute *on_loan*, delete tuple, insert tuple.
b) In MUSEUMS: modify the attribute *num_paintings*, delete tuple. In IS_IN: modify the attribute *museum*, insert tuple.
c) In MUSEUMS: modify the attribute *num_paintings*, modify the attribute *museum*, insert tuple. In IS_IN: modify the attribute *on_loan*, delete tuple, insert tuple.
d) In MUSEUMS: modify the attribute *num_paintings*, delete tuple. In IS_IN: modify the attribute *museum*, modify the attribute *on_loan*, insert tuple.

2. In the database DB of the working schema, what information is obtained with the following view?

```
CREATE VIEW V1
AS SELECT M.name, COUNT(DISTINCT E.exposition), SUM(E.num_paintings)
    FROM Museums M, Expositions E, COMPOSED_OF C
    WHERE M.museum=E.museum AND E.museum=C.museum
            AND E.exposition=C.exposition
    GROUP BY M.museum, M.name;
```

a) For all the museums with expositions, it gets their name, the number of expositions that have been organised and the total number of paintings on display.
b) For all the museums with expositions with at least one painting, it gets their name, the number of arranged expositions and the total number of paintings on display.
c) For all the museums with expositions with at least one painting, it gets their name, the number of arranged expositions and the number of paintings which are exhibited in each exposition.
d) The view would not give any information at all since it does not include the security clause “WITH CHECK OPTION”.

3. In order to obtain the cities that have more than one museum, which of the following expressions in relational algebra is correct?

a) ((MUSEUMS(museum, another_museum) ⋈ MUSEUMS) WHERE museum<>another_museum) [city]
b) ((MUSEUMS((city, another_city),(museum, another_museum)) ⋈ MUSEUMS) WHERE museum<>another_museum ∧ city=another_city) [city]
c) ((MUSEUMS M1 × MUSEUMS M2) WHERE museum<>another_museum ∧ city=another_city) [city]
d) ((MUSEUMS[museum,city](museum, another_museum) ⋈ MUSEUMS) WHERE museum<>another_museum) [city]

4. According to the database DB of the working schema, which of the following statements is **FALSE**?
   a) We can have two paintings with the same title and year, with an unknown author.
   b) We can have two paintings with the same title and year but different authors.
   c) We can have two paintings with the same author and year but different titles.
   d) We can have two paintings with the same author and year, with unknown title.

5. The SQL sentence "DROP VIEW":
   a) Destroys all the tables and views used in the table expression, provided they are in the external schema.
   b) Destroys all the tables and views used in the table expression, but only if the CASCADE option has been included in the view definition.
   c) Destroys all the tables and views used in the table expression, but only if the “WITH CHECK OPTION” clause has been included in the view definition.
   d) It never affects the tables or views used in the table expression.

6. A database is distributed between disks D1 and D2, the logfile is located in disk D2 and the backup copies for the database and for the logfile are located in a tape C3. Assume that the copy of the logfile is more recent than the copy of the database. In front of a fatal failure affecting disk D2, what must be done?
   a) The backup for the database is recovered and all the transactions that were performed since the backup date are repeated.
   b) The backups for the database and the logfile are recovered and all the transactions (confirmed in the logfile) that were performed since the backup date are repeated.
   c) The tables in D1 are kept, and the tables in D2 are recovered from the database backup.
   d) The backups for the database and the logfile are recovered, and all the cancelled transactions which are found in the logfile after the database backup copy. Finally, all the transactions (confirmed in the logfile) that were performed since the backup date are repeated

7. What query is represented by the following expression in relational algebra?

   ((MUSEUMS[museum] – EXPOSITIONS[museum]) ⋈ IS_IN) [painting]

   a) Painting identifiers which are in a museum that has no expositions.
   b) Painting identifiers which are in some museum but in no exposition.
   c) Painting identifiers which are in some exposition but in no museum.
   d) Painting identifiers which are in an exposition arranged by a museum which is not the owner of the painting.

8. If, in the database DB, we execute the SQL instruction:

   DELETE FROM Museums WHERE museum=4;

   a) Nothing will be deleted.
   b) The museum whose code is 4 will be deleted from table MUSEUMS.
   c) The museum whose code is 4 will be deleted from table MUSEUMS. The painting of code 5 will be deleted from table PAINTINGS. The tuple (5,4,true) of IS_IN will also be deleted.
   d) The museum whose code is 4 will be deleted from table MUSEUMS. The tuple (5,4,true) of IS_IN will also be deleted.

9. Please indicate which of the following statements is **FALSE**:
   a) The logical independence guarantees that the logical schema is not affected by changes in the internal schema.
   b) Data independence is stronger the later the binding is.
   c) The external level describes the views of the database for several users.
   d) A change in the physical schema does not force a modification of the application source code, if there is physical independence.

10. Given a transaction T1 that is executed over database DB:

```
START T1
INSERT INTO Paintings VALUES (10, "The scream", 15, 1893, "Expressionism");
INSERT INTO Authors VALUES (15, "Munch", "XX", "Norway");
COMMIT WORK;
```

Which of the following statements is **TRUE**?
   a) T1 will not fail if the referential integrity in PAINTINGS is defined as DEFERRABLE INITIALLY DEFERRED.
   b) T1 will not fail because the processing of the transaction meets the atomicity property.
   c) T1 will not fail if the clause WITH NO CHECK OPTION is included in the definition of the transaction.
   d) T1 will not fail if the referential integrity in PAINTINGS is defined as DEFERRABLE INITIALLY IMMEDIATE.

11. Given the database DB, what happens if, in the relation IS_IN, we change the value of the attribute museum=2 to museum=3 in the tuple {(painting, 4), (museum, 2), (on_loan, false)}?
   a) The change will not be performed because the referential integrity is violated.
   b) The change will be performed because the referential integrity is not violated.
   c) The change will be performed and will be propagated in cascade to the table MUSEUMS.
   d) The change will be performed and will be propagated in cascade to the table PAINTINGS.

12. Given the working schema, which of the following statements is **TRUE**?
   a) We can delete a museum if it has no expositions.
   b) We can delete a museum if it has no paintings.
   c) We can delete a museum if it has neither paintings nor expositions.
   d) We can delete a museum in any case.

13. Which SQL instructions would you use to define an external schema for a user?
   a) CREATE SCHEMA
   b) CREATE VIEW and CREATE DOMAIN
   c) CREATE VIEW and GRANT
   d) There are no such instructions in SQL. The external schema is created in the applications which connect to the database.

14. Regarding the lock protocols which are used to solve the problems of concurrent database access, which of the following statements is **FALSE**?
   a) If a transaction locks a datum for reading, no other transaction can access that datum until the first transaction releases it.
   b) The lock protocol in two phases implies that a transaction must perform all the data locks before the first release.
   c) If a transaction locks a datum for writing, no other transaction can access that datum until the first transaction releases it.
   d) The “deadlock” between transactions means that, in a set of two or more transactions, each transaction is waiting for data that other transactions in the set have previously locked.

Problems

1. Obtain the identifier and the title of all the paintings of style 'Baroque' for which we do not know the year in which they were painted. **(0.50 points)**

```
SELECT painting, title
FROM Paintings
WHERE style = 'Baroque' AND year IS NULL;
```

2. Create a general constraint in standard SQL to ensure that all the paintings that compose an exposition are of the same style that the style of the exposition. **(0.75 points)**

```
CREATE ASSERTION same_styles
CHECK (NOT EXISTS (SELECT * FROM Expositions E, Composed_of C
                   WHERE E.exposition = C.exposition
                         AND E.museum = C.museum
                         AND EXISTS (SELECT * FROM Paintings P
                                     WHERE C.painting = P.painting
                                           AND E.style <> P.style)));
```

or:

```
CREATE ASSERTION same_styles
CHECK (NOT EXISTS (SELECT * FROM Expositions E, Composed_of C, Paintings P
                   WHERE E.exposition = C.exposition AND E.museum = C.museum
                         AND C.painting = P.painting AND E.style <> P.style));
```

3. Obtain the name of the exposition, the name of the museum, and the start and closing dates, for the expositions that contain some paintings of the author with name ‘Botticelli’ and finish before '31/12/2010'. **(0.50 points)**

```
SELECT DISTINCT E.name, M.name, E.start_date, E.closing_date
FROM Museums M, Expositions E, Composed_of C, Paintings P, Authors A
WHERE E.museum = M.museum AND E.exposition = C.exposition
      AND E.museum = C.museum AND C.painting = P.painting
      AND P.author = A.author AND A.name LIKE '%Botticelli%'
      AND E.closing_date < '31/12/2010';
```

4. For all the authors, obtain their identifier, their name and the number of styles of their paintings. **(0.75 points)**

```
SELECT A.author, A.name, COUNT(DISTINCT P.style)
FROM Authors A LEFT JOIN Paintings P ON A.author = P.author
GROUP BY A.author, A.name;
```

5. Obtain the identifier and the name of the authors of the paintings of style “renaissance” such that their paintings do not take part in any exposition of style “renaissance”. **(0.75 points)**

```
SELECT A.author, A.name
FROM Authors A
WHERE EXISTS (SELECT * FROM Painting P
                    WHERE A.author = P.author AND P.style = 'Renaissance')
        AND A.author NOT IN (SELECT P.author
                                FROM Paintings P, Composed_of C, Expositions E
                                WHERE P.painting = C.painting
                                        AND E.exposition = C.exposition
                                        AND E.museum = C.museum
                                        AND E.style = 'Renaissance');
```

6. Obtain the identifier and the name of the author whose paintings have taken part in expositions such that the overall number of museums of these expositions is the greatest. **(1 point)**

```
SELECT A.author, A.name
FROM Authors A, Paintings P, Composed_of C
WHERE A.author = P.author AND P.painting = C.painting
GROUP BY A.author, A.name
HAVING COUNT(DISTINCT C.museum) >=
                                        ALL (SELECT COUNT(DISTINCT C.museum)
                                                FROM Paintings P, Composed_of C
                                                WHERE P.painting = C.painting
                                                GROUP BY P.author);
```

7. Obtain the identifier and the name of the Spanish authors such that all the museums which were inaugurated after 1800 own one or more of his/her paintings. **(1 point)**

```
SELECT A.author, A.name
FROM Authors A
WHERE country = 'Spain'
        AND NOT EXISTS (SELECT * FROM Museums M
                                WHERE M.year > 1800
                                AND NOT EXISTS (SELECT *
                                                FROM Is_in E, Paintings P
                                                WHERE M.museum = E.museum
                                                        AND E.painting = P.painting
                                                        AND P.author = A.author
                                                        AND E.on_loan = FALSE))
        AND EXISTS (SELECT * FROM Museums M
                        WHERE M.year > 1800);
```

8. The attribute *num_paintings* in the relation EXPOSITIONS is a derived attribute, which stores the total number of paintings that compose each exposition.

   a) Apart from the modification of the attributes *museum* and *exposition* in the relation COMPOSED_OF, please indicate the instructions that may affect this derived attribute. **(0.60 points)**

   - Insert into Expositions: num_paintings = 0
   - Update num_paintings in Expositions: forbidden
   - Insert into Composed_of: increment num_paintings (by one) for the {exposition, museum} which is specified by the inserted tuple.
   - Delete from Composed_of: decrement num_paintings (by one) for the {exposition, museum} which is specified by the deleted tuple.

   b) Please write a trigger to handle the instruction: "modification of the attribute *museum* or the attribute *exposition* in the relation COMPOSED_OF". **(0.65 points)**

```
CREATE OR REPLACE TRIGGER counts_paintings
    AFTER UPDATE OF museum OR UPDATE OF exposition ON Composed_of (*)
    FOR EACH ROW
    WHEN (old.exposition <> new.exposition OR old.museum <> new.museum)
    BEGIN
            UPDATE Expositions SET num_paintings= num_paintings+1
            WHERE exposition = :new.exposition AND museum = :new.museum;
            UPDATE Expositions SET num_paintings= num_paintings-1
            WHERE exposition = :old.exposition AND museum = :old.museum;
     END;
```

(*) Alternative: AFTER UPDATE OF museum, exposition ON Composed_of

---

Questionnaire

1	**A**
2	**B**
3	**D**
4	**D**
5	**D**
6	**B**
7	**A**
8	**A**
9	**A**
10	**A**
11	**B**
12	**C**
13	**C**
14	**A**

The result is obtained by the formula: $(\text{Right} - \text{Wrong}/3) \times 0.25$