Consider the following relational schema, which will be referred to as WORKING SCHEMA, which maintains information about a summer piano school:

TEACHER(**teacher_id**:string(4), **name**:string(40), **level**:integer)
PK: {teacher_id}
NNV: {level}

STUDENT(**stud_id**:string(4), **name**:string(40), **age**:integer, **level**:integer, **teacher_id**:string(4))
PK: {stud_id}
NNV: {name, level}
FK: {teacher_id} → TEACHER
      On delete set to null
      On update cascade

WORK(**work_id**:string(4), **title**:string(20), **level**:integer)
PK: {work_id}
NNV: {level}

CONCERT(**concert_id**:string(4), **concert_date**:date, **place**:string(20), **num_works**:integer)
PK:{concert_id}

STUDIES(**stud_id**:string(4), **work_id**:string(4), **date_of_pass**:date)
PK: {stud_id, work_id}
FK: {stud_id} → STUDENT
      Restrictive delete and modification
FK:{work_id} → WORK
      Restrictive delete and modification

PERFORMS(**stud_id**:string(4), **work_id**:string(4), **concert_id**:string(4))
PK: {stud_id, work_id, concert_id}
FK:{stud_id, work_id} → STUDIES
      Restricted delete
      On update cascade
FK:{concert_id} → CONCERT
      Restricted delete
      On update cascade

where the attributes and tables have the following meaning

TEACHER: contains the information about the teachers who teach in the summer school: their identifier (**teacher_id**), their name and their piano level.

STUDENT: has the information of the students who are inscribed in the summer school: the student's identifier (**stud_id**), their name, their age, their piano level and the teacher who is assigned to the student during the summer school (**teacher_id**).

WORK: has the information of the works which are practised during the summer school: the identifier (**work_id**), the title and the level of difficulty.

CONCERT: has the information of the concerts which are held during the summer school: the identifier for the concert (**concert_id**), the date of the concert, the place and the number of works (not necessarily different) which are performed.

STUDIES: contains the relation of works that each student is working on. Consequently, each tuple in this relation associates a student (**stud_id**) with a work the student is preparing (**work_id**) and the date in which the student has been able to perform the work well (**date_of_pass**).

PERFORMS: contains the information of the works the students perform during the concerts that are held in the summer school: it holds information about the student who performs the work (**stud_id**), the work which is performed (**work_id**) and the concert in which the work is performed (**concert_id**).

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

STUDENT

stud_id	name	age	level	teacher_id
A01	Samuel	5	1	STP
A03	Encarna	12	1	BK
A07	Jordi	9	1	
A09	María	10	2	BK

STUDIES

stud_id	work_id	date_of_pass
A01	O01	12/1/2009
A01	O03	6/6/2009
A03	O04	
A07	O01	
A09	O01	10/9/2008
A09	O05	1/1/2009
A09	O06	

WORK

work_id	title	level
O01	Twinkle	1
O02	Abejas	1
O03	Cuco	1
O04	Remando suave	1
O05	Canción de niños franceses	2
O06	Puente de Londres	2

PERFORMS

stud_id	work_id	concert_id
A01	O01	C01
A09	O01	C02
A09	O05	C01

CONCERT

concert_id	concert_date	place	num_works
C01	1/7/2009	Salón de actos	2
C02	2/7/2009	Sala pequeña	1
C03	7/7/2009		

TEACHER

teacher_id	name	level
STP	Stephen	5
BK	Biky	4
MR	María	5
JMS	James	3

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

1. Obtain the identifier and the name of the students who do not have any teacher assigned and do not study any work. **(0.75 points)**
2. Obtain the identifier and the name of the students who do not study any work which is studied by the student with identifier "E18". **(0.75 points)**
3. Obtain the identifier and the name of the oldest students from those who study with a teacher of level 4. **(1 point)**
4. Obtain for all the concerts in the database, the identifier and the date, also indicating how many different works are performed and how many students participate. (Do not use the attribute *num_works* from table Concert). **(0.75 points)**
5. Obtain the identifier and the date of the concerts such that all the works that are studied by the student with identifier ("E18") are performed in the concert. (Assume that the student E18 has studied at least one.). **(1 point)**
6. Obtain the students who are younger than 20 years old who study more than 5 works. Also, show the identifier and the name, indicating the total number of works each student is working on. **(1 point)**
7. Considering that the value of the attribute *num_works* in the relation CONCERT is a derived attribute which indicates how many works are performed in the concert, please:
   a) Indicate the operations that can affect the value of this attribute and how they would affect. **(0.5 points)**.
   b) Implement the trigger which corresponds to the operation Update *concert_id* in PERFORMS. **(0.75 points)**

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 through the formula: (Right − Wrong/3) × 0.25.

1. Given the working schema, which of the following statements is FALSE?
   a) A work can be performed by the same student in more than one concert.
   b) A work can be performed by different students in the same concert.
   c) A work can be performed several times by the same student in the same concert.
   d) A work does not need to be performed by some student in some concert.

2. In the working schema, in case the type of referential integrity of the foreign key {*stud_id, work_id*} in PERFORMS were defined as COMPLETE, which effects would it have to change it to PARTIAL referential integrity?
   a) It would not have any effect because in this case since both types of referential integrity are equivalent.
   b) We could insert the performance of a work in a concert without knowing the student who is performing.
   c) Several works could be performed in a concert, with the works not being in the table WORK.
   d) The referential integrity would not be checked until the student and the work to be performed are known.

3. Which is the reconstruction technique used by a DBMS, with immediate update, to handle those transactions that are interrupted by a system failure that affects main memory?
   a) There is not any technique since for these transactions there is no "end" entry in the log file.
   b) All the changes of the transaction that have been recorded will be undone by using the log file.
   c) There is not any technique since a DBMS only has reconstruction techniques for transactions that appear as confirmed or failed in the log file.
   d) It will try to reconstruct the transaction with the information from the log file.

4. What happens when we write a checkpoint in the log file?
   a) The DBMS write onto disk all the updates that have been produced by all the transactions that appear as confirmed in the log file since the previous checkpoint.
   b) The DBMS confirms the transactions that are previous to the last failure.
   c) The transactions free their locks.
   d) The DBMS undoes all the changes that have already been written for those transactions that are interrupted at that moment.

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

   a) The date in which a student passes the preparation of a work must be previous of any concert where the student performs that work.

   b) Two students who study the same work cannot have different teachers.

   c) A teacher can have a level that is lower than a work that is being studied by one of her/his students.

   d) A student can perform a work that s/he is not studying and s/he has not passed.

6. If we execute the instruction SET CONSTRAINT ALL DEFERRED:

   a) All the deferrable constraints will be deferred until the end of the transaction, when all of them will be set as not deferred.

   b) All the deferrable constraints will be deferred until the end of the transaction, when all of them will be reset to the corresponding state according to their definition.

   c) All the deferrable constraints will be deferred until the end of the transaction, when all of them will be reset as they were just before executing the instruction.

   d) All the constraints will be deferred. Those which are "deferrable" will be deferred indefinitely and those which are "not deferrable" until the end of the transaction.

7. If we want to implement the constraint "the level of a teacher cannot be incremented in more that one unit each time", how could we implement it?

   a) With a uniqueness constraint for the attribute "level" in TEACHER.

   b) With the "CREATE ASSERTION" instruction.

   c) By using TRIGGERs.

   d) It is not possible to implement this constraint in standard SQL.

8. Given the working schema, which of the following statements is FALSE?

   a) A student can study several works.

   b) A student can perform more than a work in a concert.

   c) A work that has been studied by a student must be performed in at least one concert.

   d) A student can perform the same work in two different concerts.

9. Considering the working schema, which information does the following Relational Algebra query returns?

   > STUDENT[stud_id,name]–(((Work(level,work_level)[work_id,work_level] ⋈ PERFORMS) ⋈ STUDENT(level,student_level)) WHERE student_level < work_level) [stud_id,name]

   a) Identifier and name of the students who only perform works of a level that is higher than the piano level they have.

   b) Identifier and name of the students who do not perform any work of a level that is higher than the piano level they have.

   c) Identifier and name of the students who perform a work of a level that is higher than the piano level they have.

   d) Identifier and name of the students who perform a work of a level that is lower than the piano level they have.

10. Assuming that in the working schema all the constraints are defined as INITIALLY IMMEDIATE, consider that we want to insert the tuple ('A02', 'O04', '11/09/2009') into the relation STUDIES. Which of the following statements is FALSE?

   a) This tuple cannot be inserted into STUDIES since there is no student with code 'A02' and the foreign key constraint to STUDENT would be violated.

   b) In order to insert this tuple, we should previously insert the student with code 'A02'.

   c) We could only insert the tuple if we perform the transaction:

      COMMIT;
           SET CONSTRAINT ALL DEFERRED;
           INSERT INTO STUDIES VALUES ('A02', 'O04', '11/09/2009');
      COMMIT;

   d) The tuple could be inserted if we perform the transaction:

      COMMIT;
           INSERT INTO STUDENT VALUES ('A02', 'Pepe', 7, 1, 'STP');
           INSERT INTO STUDIES VALUES ('A02', 'O04', '11/09/2009');
      COMMIT;

11. Regarding the physical database implementation using hash files, which is the goal of a good hash function?

   a) Distribute the records uniformly over the address space in order to minimise the collisions, leaving as few empty buckets as possible.

   b) Guarantee the dynamic expansion of the file as new records are inserted.

   c) Automatically free the empty buckets, so minimising the size of the file.

   d) Group the records according to the values of the hash field in order to minimise the accesses with searches that use the hash field.

12. How would the extension of the working database be affected if the following instruction is executed?

      UPDATE STUDIES SET stud_id='A02' WHERE stud_id='A09' AND work_id='O05'

   a) No modification would take place since it is not allowed to modify the attribute *stud_id* in the table STUDIES.

   b) The tuple in STUDIES would be modified as it is indicated in the instruction but also the table PERFORMS would appear as follows:

PERFORMS

stud_id	work_id	concert_id
A01	O01	C01
A02	O01	C02
A02	O05	C01

   c) The modification cannot be performed because the foreign key constraint to STUDENT would be violated.

   d) The tuple in STUDIES would be modified as it is indicated in the instruction but also the table STUDENT would appear as follows:

STUDENT

stud_id	name	age	level	teacher_id
A01	Samuel	5	1	STP
A03	Encarna	12	1	BK
A07	Jordi	9	1	
A02	María	10	2	BK

13. Which of the following statements is TRUE?

    a) Data independence is worse the later the binding is.

    b) Data independence is lost when the binding between the external and the internal schemas is made.

    c) When the binding between schemas is performed, no update to the stored data can be done any more.

    d) The behaviour of the applications is more efficient the more frequent the binding between schemas takes place.

14. Assume that the database for the working schema is empty and also assume that all the constraints are deferred. Which integrity constraint would be violated by the following transaction?

        INSERT INTO Work VALUES ('So23', 'Appassionata', 3);
        INSERT INTO Work VALUES ('So14', 'Moonlight', 2);
        INSERT INTO Studies VALUES ('JP', 'So23', NULL);
        INSERT INTO Studies VALUES ('AG', 'So14', NULL);
        INSERT INTO Student VALUES ('JP', 'Juan Pérez', NULL, NULL, NULL);
        INSERT INTO Student VALUES ('AG', 'Ana García', NULL, NULL, NULL);
        DELETE FROM Work WHERE level>2;
        UPDATE Student SET level=1;
        COMMIT

    a) FK: {stud_id} in Studies.

    b) FK: {work_id} in Studies.

    c) NNV of the attribute level in Student.

    d) FK: {teacher_id} in Student.