Questions From Previous Exams

Gathered and organised by units from the exams:

Sept 97, June 98, Sept 98

Questions from Previous Exams (U.1)

1) (Sept 97) An external schema consists of:

A) an implementation of the structures in the logical schema.

B) a definition of the database structures.

C) a definition of several conceptual schemas.

D) a set of derived structures defined from the logical schema.

Questions from Previous Exams (U.1)

2) (June98) Which of the following statements is **FALSE**?

A) Data independence ensures that user applications are independent from changes which might be performed on data which is not used by these applications or performed on the physical representation.

B) The data structures which compose a database are defined independently from the applications which manipulate them.

C) A DBMS must ensure the integrity of the stored information and must only give information access to the authorised users.

D) Every data model is based on a DBMS.

Questions from Previous Exams (U.1)

3) (Sept98) A DBMS is:

A) a collection of structured data.

B) a tool that allows users to manipulate and create databases

C) a set of concepts and rules which makes the representation of reality possible.

D) an information system composed of some data and the programs which manipulate them.

Questions from Previous Exams (U.1)

4) (Sept 98) The logical schema is:

A) The implementation of the structures defined in the DB.

B) The implementation of the partial views defined in the DB.

C) The definition of the data structures which compose a DB.

D) The definition of the partial data views which compose the DB.

Questions from Previous Exams (U.2. R.A.)

5) (June98) Given the relations *R*, *S* and *T,* and the following expression in Relational Algebra:

$$R \triangleright\!\triangleleft (S - T)$$

where the operations are well defined, which is the minimum and maximum cardinality of this expression depending on the cardinalities of *R, S* and *T*?

A) The minimum cardinality is 0 and the maximum is card(R) × (card(S) – card(T)).

B) The minimum cardinality is card(R) and the maximum card(R) × (card(S) – card(T)).

C) The minimum cardinality is 0 and the maximum is card(R) × card(S).

D) The minimum cardinality is card(R) × (card(S) – card(T)) and the maximum is card(R) × card(S).

Questions from Previous Exams (U.2. R.A.)

6) (Sept97) Given the following relation schemas:

E(**code**:code_d, **nom**:name_d, **city**:city_d, **dep_name**:dep_d) /* Employees */

D(**dep_name**:dep_d, **head**:code_d) /* Departments */

What does **((E[code]–D[head](head,code)) ▷◁ E) where city='Alcoi'** mean?

A) Employees who are from Alcoi, but are not Dept Heads.

B) Department Heads who are from Alcoi.

C) Employees who are from Alcoi.

D) It is syntactically incorrect.

Questions from Previous Exams (U.2. R.A.)

7) (Sept97) In the same schema as in the previous question, i.e.:

E(**code**:code_d, **nom**:name_d, **city**:city_d, **dep_name**:dep_d) /* Employees */

D(**dep_name**:dep_d, **boss**:code_d) /* Departments */

Suppose we have 1,000 employees, 10 departments, and 10% of the employees and 10% of the bosses are from Alcoi. How many tuples does the expression **((E[code]–D[boss](boss,cod)) ▷◁ E) where city='Alcoi'** return?

A) 10.000 tuples.

B) 990 tuples.

C) 99 tuples.

D) No tuples.

Questions from Previous Exams (U.2. I.C.)

11) (June98) What is the role of a foreign key in a relational schema?

A) It is a mechanism that associates two relations.

B) It is an integrity constraint which ensures that the association between two relations is adequate.

C) It is useful to accelerate the joins between the two relations involved.

D) It is useful to indicate the DBMS to store the two relations involved in a cluster.

Questions from Previous Exams (U.2. I.C.)

12) (June98) Given the following relational schema:

Department(*code*: string(5), *location*: string(40), *manager*: integer)
- PK: {code}
- Uni: {manager}
- FK: {manager} → Employee

Employee(*id*: integer, *name*: string(50), *address:* string(30), *dep_code*: string(5))
- PK: {id}
- FK: {dep_code} → Department
- NNV: {dep_code}

Project(*code*: string(5), *title*: string(100), *budget*: real, *head*: integer)
- PK: {code}
- FK: {head} → Department f(head) = manager

Which of the following statements is **FALSE**?

A) An employee must belong to a department and a department can have several employees.

B) A department cannot have more than one manager.

C) The head of a project must be department manager and s/he can only be head of one project.

D) The manager of a department can be an employee of another department.

Questions from Previous Exams (U.2. I.C.)

13) (June98) Given a relational schema and a database of this schema:

R(*A*: integer, *B*:string(5), *C*:real)
PK: {A}
Uni: {B}

A	B	C
1	Z 1	6 .4
2	Z 6	7 .3
3		6 .4

S(*L*: real, *M*: date, *N*:string(5))
PK: {M}
FK: {N} → R f(N) = B
Weak referential integrity
Modification to nulls
Cascade deletion

L	M	N
6.5	1/1/98	Z1
7.3	2/5/98	
6.5	3/5/89	Z6

Choose the state of the database after executing the following transaction:

T = { **DELETE FROM** *R* **WHERE** *A* = 3, **UPDATE** *R* **SET** *B* = '*Z2*' **WHERE** *A*=2 }

A)

A	B	C
1	Z 1	6 .4
2	Z 2	7 .3

L	M	N
6.5	1/1/98	Z1
7.3	2/5/98	
6.5	3/5/89	

B)

A	B	C
1	Z 1	6 .4
2	Z 2	7 .3

L	M	N
6.5	1/1/98	Z1
7.3	2/5/98	
6.5	3/5/89	Z2

C)

A	B	C
1	Z 1	6 .4
2	Z 2	7 .3

L	M	N
6.5	1/1/98	Z1
7.3	2/5/98	

D) The same, since the transaction cannot be completed because an integrity constraint is violated.

Questions from Previous Exams (U.2. I.C.)

14) (June98) Given the following relational schema and a database of this schema:

R(*A*: string(1), *B*: integer, *C*: string(10))
PK: {A,B}

A	B	C
z	1	Pepe
z	2	Juan
x	1	María
x	2	Luisa

S(*D*: string(2), *A*:string(1), *B*:integer, *E*: *real*)
PK: {D}
FK: {A,B} → R Partial referential integrity

D	A	B	E
FA	z		1.5
DO		1	1.0
RE	x	1	1.6

Which of the following operations would be rejected by the DBMS?

A) **DELETE FROM** *R* **WHERE** *A* =‘x’ AND B= 2

B) **UPDATE** *R* **SET** *A* = ‘w’ **WHERE** *A*= ‘z’ AND B=2

C) **DELETE FROM** *R* **WHERE** *A* =‘z’

D) **DELETE FROM** *R* **WHERE** *A* =‘z’ AND B= 1

Questions from Previous Exams (U.2. I.C.)

15) (Sep98) Given the following relational schema and a database of this schema:

R(*A*: string(1), *B*: integer, *C*: string(10))
PK: {A,B}

A	B	C
z	1	Pepe
z	2	Juan
x	1	María
x	2	Luisa

S(*D*: string(2), *A*:string(1), *B*:integer, *E*: *real*)
PK: {D}
FK: {A,B} → R Weak Referential Integrity

D	A	B	E
FA	z		1.5
DO		1	1.0
RE	x	1	1.6

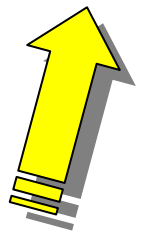

Which of the following operations cannot completed?

A) **DELETE FROM** *R* **WHERE** *A* =‘x’ AND B= 2

B) **UPDATE** *R* **SET** *A* = ‘w’ **WHERE** *A*= ‘z’ AND B=2

C) **INSERT INTO** *S* (D,B,E) **VALUES** (‘MI’, 3, 1.5)

D) **DELETE FROM** *R* **WHERE** *A* =‘x’ AND B= 1

Questions from Previous Exams (U.2. I.C.)

16) (Sep98) Given the following relational schema:

R(*a*: dom1, *b*: dom2)
Uni: {a}
Uni: {b}

S(*c*: dom2, d:dom1)
PK: {c}
FK: {d } → R

What errors can be found in the previous schema?

A) There are no errors.

B) Relation *R* must have a primary key.

C) The foreign key in relation *S* cannot be defined since relation *R*, the one *S* refers to, does not have a primary key.

D) The foreign key in relation *S* requires the bijection to be specified in order to indicate which attribute in *R* it refers to.

Questions from Previous Exams (U.2. I.C.)

17) (Sep98) Given the following relational schema:

S(*a*: dom_a, *b*: dom_b)
PK: {a}

T(*c*: dom_c, a:dom_a)
PK: {c}
FK: {a } → S

And the following expression in Tuple Relational Calculus:

x: S, y: T { x.a, y.c | S(x) ∧ T(y) ∧ x.a = y.a }

(in the logical notation seen since the academic year 2002/2003:

{ x.a, y.c | ∃ x:S, y:T x.a = y.a } or in SQL: **SELECT S.a, T.c FROM S, T WHERE S.a= T.a;)**

What is the maximum and minimum cardinality of the previous expression?
The cardinalities of relations *S* and *T* are represented by *card*(*S*) and *card*(*T*) respectively.

A) Cardinalities: minimum 0 and maximum ∞.

B) Cardinalities: minimum card(T) and maximum card(S) × card(T).

C) Cardinalities: minimum card(S) × card(T) and maximum ∞.

D) Cardinalities: minimum 0 and maximum card(T).

Questions from Previous Exams (U.2. I.C.)

18) (Sep98) Given the following relational schema:

Prisoner(number: integer, name: string(40), age: integer)
PK: {number}
NNV: {name}
1stDegreePrisoner(number: integer, cell_code: string(5))
PK: {number}
NNV: {cell_code}
FK: {cell_code} → Cell
FK: {number} → Prisoner
Cell(code: string(5), room: integer, location: string(50))
PK: {code }
Job(id: integer, place: string(50), company: string(50))
PK: { id }
Assigned(id: integer, number: integer)
PK: { id, number }
FK: {id} → Job
FK: {number} → Prisoner

Choose which of the following statements is **FALSE**:

A) Every 1stDegreePrisoner is a Prisoner.

B) Every 1stDegreePrisoner has only a unique Cell assigned.

C) A Job can be assigned to several Prisoners.

D) Any Prisoner can be Assigned to several Jobs.

¬∃x:1stDegreePrisoner ∃y:Assigned (x.number = y.number)

Questions from Previous Exams (U.2. I.C.)

20) (Sept97) In the following schema:

E(**code**:code_d, **name**:name_d, **city**:city_d, **dep_name**:dep_d) /* Employees */

D(**dep_name**:dep_d, **head**:code_d) /* Departments */

Which of the following statements is TRUE?

A) Whenever a new department is inserted, a new employee must be inserted, which will be the head of the department.

B) Whenever a new employee is inserted, the "dep_name" attribute must have one of the values appearing in the "dep_name" attribute in Department.

C) Whenever a new employee is inserted, the "dep_name" attribute can be null.

D) Whenever a department is deleted, the employees in that department must be deleted.

Questions from Previous Exams (U.3)

22) (Sept98) A B^+ tree:

A) Allows access to the records ordered by the search field.

B) Does not allow access to the records ordered by the search field.

C) Contains (in the internal nodes) the physical address of the block where the record is.

D) Does not have all the leaf nodes at the same level.

Questions from Previous Exams (U.3)

23) (Sept97) Choose which of following statements is **FALSE**:

A) The records in a primary index are composed by the ordering key and the disk block address.

B) The insertion of records in the data file with a primary index does not require a physical ordering.

C) In a file organisation with a primary index, binary search is efficient.

D) An index sets an ordering on the file with no need of sorting it physically.

Questions from Previous Exams (U.3)

24) (June98) Which of the following statements is **FALSE**?

A) From the three organisations we have seen, the hash technique provides the quickest access to locate an arbitrary record given the value of the hash field, if there are no overflow buckets.

B) In a hash file, the search for a record through a field which is not the hash field is as expensive as in disordered files.

C) In a hash file, we cannot define secondary indexes over fields other than the hash field.

D) “open addressing” is a mechanism to solve collisions in a hash file.

Questions from Previous Exams (U.3)

26) (June98) One of the key traits of database techniques is the definition of partial views for different users. In which of the following levels of the ANSI/SPARC architecture are these partial views defined?

A) External level.

B) Physical level.

C) Logical level.

D) Conceptual level.

Questions from Previous Exams (U.3)

27) (June98) For the reconstruction of a database using the log file (journal) and the most recent backup, how would you proceed?

A) The database must be loaded from the backup and all the cancelled transactions must be redone from the last checkpoint in the log file.

B) The database must be loaded from the backup and all the confirmed transactions in the log file must be redone from the backup date.

C) The database must be loaded from the backup and all the cancelled transactions must be redone from the backup date.

D) The database must be loaded from the backup and all the confirmed transactions in the log file must be redone from the last checkpoint in the log file.

Questions from Previous Exams (U.3)

28) (June98) In a DBMS the evaluation of a query defined over a structure in the external schema means:

A) A translation of the query according to the access methods defined for the kind of file which implements the structure.

B) A translation of the query into a new query defined over the structures in the logical schema and the subsequent translation of this latter query according to the access methods defined for the kind of file which implements the structure.

C) The query is evaluated directly.

D) It will depend on the underlying data model the DBMS is using.

Questions from Previous Exams (U.3)

29) (Sept97) The physical independence is the level of independence which is established:

A) Between the user application programs in order to avoid them from affecting each other.

B) Between the internal schema and the external schemas so that the latter would not be affected by changes concerning the logical data structures.

C) Between the logical schema and the internal schema so that the changes performed on the logical schema do not imply a modification on the organisations chosen for the files implementing the database.

D) Between the logical schema and the internal schema so that the logical schema is not affected by changes in the internal schema concerning implementation details.

Questions from Previous Exams (U.3)

30) (Sept98) In the following transaction timing diagram, what should the DBMS do when the system reboots after a system failure?

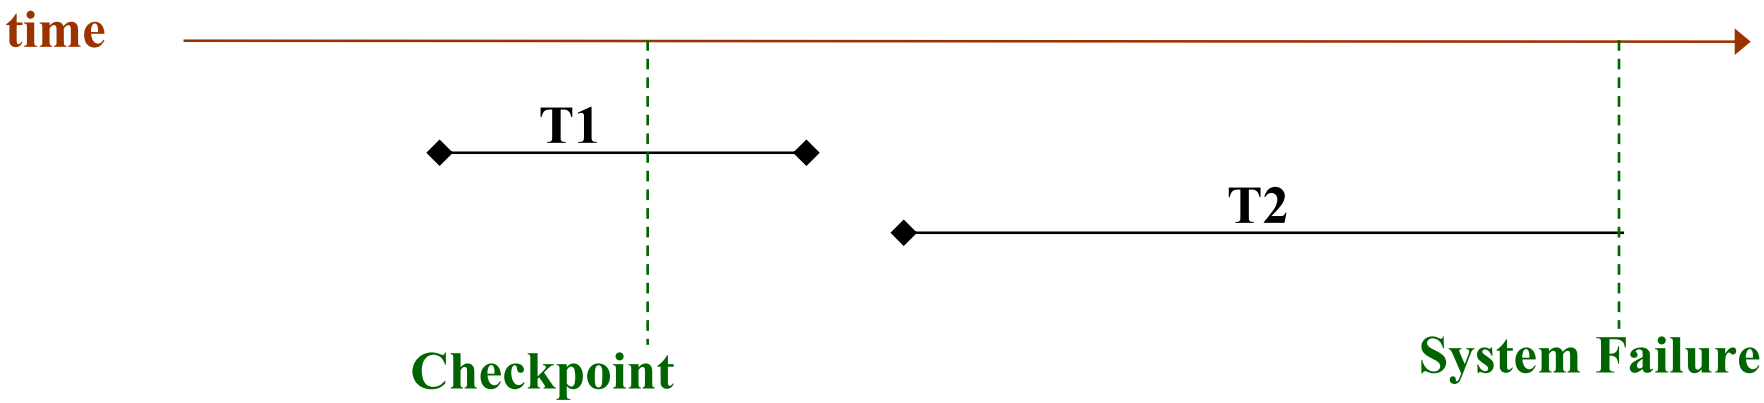

A) Redo T2 and undo the changes performed by T1.

B) Redo T1 and undo the changes performed by T2.

C) Undo the changes performed by T2.

D) Undo the changes performed by T1 and undo the changes performed by T2.

Questions from Previous Exams (U.3)

31) (Sept98) In transaction handling of a relational system complying with standard SQL, the checking of the integrity constraints in the schema is performed:

A) After the transaction has finished.

B) After the transaction has been confirmed.

C) Depends on the DEFERRABLE option chosen for each constraint.

D) After each instruction in the transaction that may violate a constraint.

Questions from Previous Exams (U.3)

32) (Sept98) Would it be necessary to have a log file (journal) during the normal operation of a DBMS (no system failure or external storage damage)?

A) No.

B) Depends on the independence levels the DBMS provides.

C) Yes, if the DBMS is relational.

D) Yes, in order to undo the changes performed by a transaction which is cancelled.

Questions from Previous Exams (SQL)

33) (June98) The clause “WITH GRANT OPTION”:

A) Is a compulsory clause for giving access permission to users who do not own the database.

B) Is a compulsory clause which allows a user who has been granted permissions to grant permissions to third persons.

C) Is an optional clause which allows a user who has been granted permissions to grant a subset of these permissions to third persons.

D) Is a clause for giving permissions to views and not only to tables.

Questions from Previous Exams (SQL)

34) (Sept97) U1, U2 and U3 are users of a database composed of relations R1 and R2. Exclusively taking into account the following SQL statements:

GRANT SELECT ON R1 TO U1 WITH GRANT OPTION;

GRANT INSERT ON R2 TO U2;

Which of the following statements is **FALSE**?

A) U3 could query R1 if s/he is authorised by U1.

B) U3 could insert a tuple into R2 if s/he authorised by U2.

C) U1 can only query R1.

D) Only U2 can insert a tuple into R2.

Questions from Previous Exams (SQL)

35) (Sept97) With the following schema:

E (*code*: code_d, *nom*: name_d, *city*: city_d, *nom_dep*: dep_d)
PK: {code}
FK: {nom_dep} → D

D (*dep_name*: dep_d, *boss:* code_d)
PK: {dep_name}
FK: {boss} → E Uni: {boss} NNV: {boss}

What does the following requirement mean?

```
SELECT dep_name FROM D
WHERE NOT UNIQUE (SELECT dep_name FROM E WHERE E.dep_name = D.dep_name)
```

A) The relation D has more than one employee.

B) It is syntactically incorrect: it must be NOT EXISTS.

C) Returns the name of the departments with more than one employee.

D) Returns the name of the departments without employees.