BD2 – June 20, 2017 – V1.1

Please feel free to answer this test in English, Italian, or any mixture

First Part (6 credits and 9 credits):

1. Consider a schema R(A,B,……) S(E,F,……) and the following query

```
SELECT DISTINCT   S.*
FROM              R, S
WHERE             R.A=S.E  and  R.B=0  and  S.F < 10
```

Assume that R and S are stored as heap files. The pair (A,B) is a key for R, and (E,F) is a key for S. Observe that none of the four attributes is a key – the key is the pair. Assume that unclustered RID-sorted index are defined on R.A, R.B, S.E and S.G. Assume that every index on R has 20.000 leaves, every index on S has 4.000 leaves.
Assume the following table for the optimization parameters. If you need Cardenas formula Φ(n,k), approximate it with min(n,k).
Assume that every page is 4.000 bytes long, and that every attribute uses 4 bytes.

	NReg	NPag	NLeaf of Indexes	NKey	Min	Max
R	10.000.000	200.000	20.000			
S	2.000.000	100.000	4.000			
Idx.R.A			See R	100	0	1.000
Idx.R.B			See R	100.000	0	1.000.000
Idx.S.E			See S	1.000.000	0	1.000.000.000
Idx.S.F			See S	1.000.000	0	1.000.000

   a) Knowing that a pair (X,Y) is a key for a table T, and that T has 1.000.000 records, what can be said about NKey of T.X and NKey pf T.Y?
   b) Is DISTINCT redundant in the above query?
   c) Draw an access plan that uses IndexNestedLoop, has R as its outer relation (left hand side), uses tablescan for R, and compute its cost
   d) Could you answer to question (c) in a different way if you had a combined index on (E,F)?
   e) Draw an access plan based on MergeJoin, and compute its cost
   f) When MergeJoin is used with a condition R.X=S.Y where neither attribute is a primary key, MergeJoin may be forced to go back and forth. When does that happen? May that happen in this case?

2. Consider the following log content. Assume that the DB was identical to the buffer before the beginning of this log, and consider a **no-undo** / redo protocol. Beware, it is **no-undo**.

(begin,T1) (W,T1,A,1,20) (begin,T2) (W,T2,B,1,20) (begin-ckp,{T1,T2}) (W,T1,A,20,40) (end-ckp) (begin,T3) (begin,T4) (W,T3,C,1,40) (commit,T2) (W,T3,C,40,60) (W,T4,B,20,80) (commit,T4)

   a) Before starting this log, what was the content of A, B and C in the PS (Persistent Store)?
   b) What happens during the checkpoint that is reported in the log? Keeping in mind that this is no-undo, which pages are written, or may be written, to the PS?

c) Assume there was a crash at the end of the logging period. At crash time, what was the content of A, B and C in the buffer? What can be said about the content of A, B and C in the PS?
d) At restart time, which transactions are undone? Which are redone?
e) List the operations that are redone, in the order in which are redone
f) After restart is finished, what is the content of A, B and C in the buffer?
g) Undo and Redo are executed in the buffer or on the PS?
h) After restart is finished, what is the content of A, B and C in the PS?

Second Part, BSA only (9 credits only):

3. Consider the following document that records data about bank accounts. For each account, it records the holders – a set of clients – and information about money transfers that originate from the account. For each money transfer it records date, money, and issuers: generally speaking, when a set of clients hold an account together, every transfer is issued by a non-empty subset of these holders.

```
Account*
  @IdA
  Bank
  Holder*
    @IdClient
  Transfer*
    @IdTransfer
    Year
    Date
    Amount
    IssuedBy*
        @IdClient
Client*
  @IdClient
  Name
  Country
```

a. Assume that $doc is bound to the above document. Write an XQuery query that, for each client, lists the information about the transfers that have been issued by that client (maybe with others), grouped by account and by year, according the following format

```
Client*
  Name
  Account*
    @IdA
    ByYear*
        Year
        TotalTransfersFromAccountByTheClientInYear (please abbreviate as TTFA)
```

b. Write and XQuery query that returns the IdTransfer of those transfers where at least one of the issuers does not belong to the set of the holders of the account

4. Consider and RDF graph with classes Client, Account, Bank, and predicates

HasHolder $\subseteq$ Account $\times$ Client

HasBank $\subseteq$ Account $\times$ Bank
SharesAccountWith $\subseteq$ Client $\times$ Client
FinecoClient, HasFinecoOnly $\subseteq$ Client

Formalize the following statements in OWL, paying extreme attention to the direction of the implication:

a. Formalize in OWL: If two clients C1 and C2 are both holders of one account (or more than one), then C1 SharesAccountWith C2
b. Define 'reflexive' a relation such that x R y implies that x R x and y R y. Is it possible to deduce from (a) that SharesAccountWith is reflexive?
c. Formalize in OWL: Every Account has exactly one Bank
d. Formalize in OWL the following, assuming (c): a client C belongs to FinecoClient iff it is holder of at least one Account whose Bank is "Fineco", where Fineco is one element of Bank
e. Formalize in OWL the following, assuming (c): a client C belongs to HasFinecoOnly iff every Account of which it is holder has Bank "Fineco"

Consider an arbitrary RDF graph the only contains triples whose predicate is either HasHolder or HasBank, Considering that graph and the five statements above, the fact that OWL has an open world semantics, and considering the directions of the implications, specify which of the following statements about a Client X may be proved and which may never be proved. For those that may be proved give an example of and RDF graph supporting the proof, for those that may never be proved give an extremely brief reason for that

f. X is in FinecoClient
g. X is not in FinecoClient
h. X is in HasFinecoOnly
i. X is not in HasFinecoOnly
j. Would (i) change if you could add a DifferentIndividual statement?
k. Provide an example of a knowledge base that may allow one to prove that X is not in FinecoClient. The knowledge base may contain any OWL statement but should not explictly name the 'FinecoClient' property