

Question Sheet: Recitation #1 DD2471

Fill in all of your answers on a 1 page answer sheet. **Write clearly!** and stay within the frame for an answer. There are no penalties for wrong answers.

True/False

1. The ACID rules are a set of properties whose primary purpose is to guarantee that database indexing is processed reliably.
2. A phantom read occurs when two identical queries are executed and the collection of rows returned by the queries are identical.
3. $R \bowtie R = R$ under bag semantics.
4. XML as a data representation language does not meet first normal form.
5. In SQL, `SELECT * FROM R UNION SELECT * FROM R;` always yields exactly the same results as `SELECT DISTINCT * FROM R;`.
6. In SQL, `SELECT b FROM R WHERE a < 10 OR a >= 10;` always yields the same results as `SELECT b FROM R;`.
7. A primary key value may not be NULL.
8. For $R[AB]$, $S[AC]$, $R \bowtie S \equiv R \bowtie_{R.A=S.A} S$.
9. Projection is more efficient under set semantics than bag semantics.
10. The core SQL database modifications commands are INSERT, UPDATE, and DELETE.
11. For $R[ABCD]$ and FDs $A \rightarrow BC$, $C \rightarrow D$, R is not in 3NF.
12. The chase test is commonly used to determine whether a set of functional dependencies is preserved over a decomposition.
13. One can enforce 3NF in SQL via PRIMARY KEY and UNIQUE.
14. For ECA rules, E stands for Event.
15. An XML file must adhere to a DTD if we want to use XPATH to access it.
16. There is the possibility of a dirty read under isolation level READ COMMITTED.
17. Different transactions may run under different isolation levels.
18. All SQL transaction levels preserve durability.
19. In ODBC connections are always constructed within an environment.
20. Extended relational algebra is less expressive than Tuple Calculus.

Short Answers

21. $\{1, 2, 1, 1, 1, 3\} - \{1, 2, 2\} = ?$ under bag semantics.
22. In SQL's AVG aggregate function $AVG(1, NULL, 1, 0) = ?$
23. In SQL the expression $(NULL \text{ OR } TRUE) \text{ AND NOT } (NULL \text{ AND } FALSE) = ?$
24. What does FLWR stand for?
25. Give a single sentence explanation to why we should not just build indexes for all combinations of attributes of a table.
26. In a single sentence, why would you do `SQLPrepare` followed by `SQLExecute` rather than just `SQLExecuteDirect`?

Normalization

27. Using the top down algorithm presented in the book or lecture, decompose R[BSTCA] under the given functional dependencies into a lossless BCNF decomposition. It suffices to simply show the resulting decomposition.

B → S
ST → B
AT → B
B → C
S → A
A → S

28. Is your decomposition dependency preserving?

29. Using the synthesis algorithm presented in the book and lecture, present a 3NF decomposition of the above schema under the given functional dependencies. It suffices to simply show the resulting decomposition.

Conceptual Modeling

30. Write the relation schema corresponding to the ER diagram on the last page of this exam. Underline keys. Use the *ER method* for translating IS-A hierarchies to relations.

Queries

You have accepted a position with a wholesale seller who keeps a database stored in the following schema:

C(cid, name)
I(iid, name, price)
B(cid, iid, eid)
E(eid, name, boss)

Where C stands for Customers, I stands for Items, and B stands for bought, E for employee. In B, cid, iid, eid are foreign keys out to other relations, and boss is a foreign key to E, representing an employees direct boss.

31. In tuple calculus give the items that were bought by exactly one customer.

32. In SQL using aggregate operators write the query for the average price of items of each customer.

33. Explain in 20 words or less what the following query computes:

```
with recursive R(eid, boss, k) as (  
  select eid, boss, 1 from E where name = 'Ignatius'  
  union  
  select r.eid, e.boss, r.k+1 from E e, R r  
  where r.boss = e.eid)  
select name, k from R natural join E order by k desc;
```

34. The previous database developer apparently quit in frustration and you are now left with the implementation of a request from marketing. Fortunately, you find the solution to the request scabbled on a whiteboard:

$$\Pi_{name}(C \bowtie (\Pi_{cid} - \Pi_{cid}(\Pi_{cid, iid}(C \times I) - B)))$$

Explain, in 15 words or less, what question the request answers.

35. Take the RA expression and rewrite it in SQL without using aggregation operators.

36. Marketing contacts you and explains they are really only interested in this for items that cost more than 1000. Amend your solution from 35 accordingly.

37.

The XML database worldtowns.xml contains the following sort of entities:

```
<town inpolity='France'>Paris</town>  
...  
<pop forpolity='France'>60</pop>  
...
```

Using XQuery, write a query against the worldtowns database that returns the number of cities in each country with a population above 50. Format the output like:

```
<country>  
  <name>France</name>  
  <num>1</num>  
</country>
```

You don't have to preserve output formatting, but you must preserve proper XML tag nestling.

Transactions

Consider transactions over the simple database `Account(id, amount)` where the initial database state has only the two account tuples: `Account(id:1, amount:5)` and `Account(id:2, amount:10)`. Consider the following transaction (T_{ij} is transaction i 's j -th operation):

```
T11: add 5 to value in account 1
T12: subtract 5 from amount in account 2
```

```
T21: read amount of account 1
      into local variable X
T22: if X is greater than 5 then
      subtract 1 from account 1
```

```
T31: read amount of account 2
      into local variable X
T32: if X is greater than 5 then
      subtract 1 from account 2
```

38.

What are the possible values of `SELECT SUM(amount) FROM ACCOUNT;` after these three transactions run with isolation level `SERIALIZABLE`. Remember that transactions can always abort at any step.

39.

What are the possible values of `SELECT SUM(amount) FROM ACCOUNT;` after these three transactions run with isolation level `READ UNCOMMITTED`. Again, remember that transactions can always abort at any step.

40.

How would you rewrite transaction T2 above to only use one operation (i.e. perform T21 and T22 together)?

