Question Sheet: Recitation #1

True/False

1. In SQL databases, no two whole tuples in a relation can ever be identical.

2. Phantoms occur under the READ COMMITTED isolation level.

3. Normalization to 3NF is guaranteed to preserve all functional dependencies of the original universal relation.

4. Many to many relations in an E/R model are dissolved and their attributes are transferred to one or the other of the connected entities when creating a relational database of the model.

5. In SQL, `SELECT * FROM R UNION SELECT * FROM R` may yield more tuples than `SELECT DISTINCT * FROM R`.

6. Running all transactions under `SERIALIZABLE` isolation level has much better throughput than repeatable `READ COMMITTED` level.

7. Under the default case, we can insert a NULL value into a foreign key attribute.

8. For R[AB], S[CD], $R \bowtie S \equiv R \times S$.

9. if $A \rightarrow BC$ then $A \rightarrow B$ and $A \rightarrow C$ and vice versa.

10. Under the default case, a primary key value for a tuple can be `NULL`.

11. For R[ABCD] and FDs $AB \rightarrow CD$, $C \rightarrow A$, R is in 3NF.

12. The chase test can be used to determine whether a decomposition is lossless.

13. One can enforce BCNF in SQL via `PRIMARY KEY` and `UNIQUE`.

14. Using XPath on an XML database, `//*/mountain` and `//mountain` will always return the same set of entities.

15. `<delivery><car id="23"></delivery>` is well formed XML.

16. There is the possibility of a dirty read under isolation level REPEATABLE READ.

17. We can replace a many-to-many relationships in the E/R model with two oppositely directed one-to-many relationships.

18. All SQL isolation levels preserve atomicity of transactions.

19. In ODBC one must possess a connection handle before one obtains an environment handle.

20. Given relations R(A,B) and S(B,C), the output of R natural join S vs. R cross join S where R.B = S.B is identical.

Short Answers

21. What are the candidate keys of $R(ABCDE)$ where $D \to A$, $A \to B$, $C \to D$.

22.

```
CREATE TABLE Monoid (
  id INT PRIMARY KEY,
  mother INT REFERENCES monoid
    ON DELETE CASCADE
    DEFERRABLE INITIALLY DEFERRED
);
INSERT INTO Monoid VALUES(1,NULL);
INSERT INTO Monoid VALUES(2,1);
INSERT INTO Monoid VALUES(3,2);
INSERT INTO Monoid VALUES(4,3);
INSERT INTO Monoid VALUES(5,4);
```

What is the value of the count query here.

```
DELETE FROM Monoid WHERE id = 4;
SELECT COUNT(*) FROM Monoid;
```

23. If instead of `VALUES(1,NULL)` in the above, we had `VALUES(1,5)`, and assuming that we could get it to run without error, what would be the value of the count query here.

```
DELETE FROM Monoid WHERE id = 4;
SELECT COUNT(*) FROM Monoid;
```

24. As it stands here, the inserts would fail for problem 23. How could we actually perform the inserts in question 23?

25. What does ECA stand for in ECA rules?

26. Why does ODBC have `SQLPrepare` and `SQLBindParameter`? Why not just `SQLExecute`?

Normalization

27. Boats have ids (B) and every night (N) they may be parked at most one boat place (P) which is located in a marina (M). Each boat has a length (L) and a width (W). Naturally only one boat can occupy a given place in a given marina on a given night. Give the functional dependencies that govern this example.

28. Use the 3NF synthesis algorithm to decompose $R(ABCDE)$ into 3NF under $A \to B$, $B \to C$, $A \to CD$, $D \to B$. Present the resulting decomposition.

29. Use the decomposition algorithm to decompose $R(ABCDE)$ into BCNF under $ABC \to DE$, $D \to ABC$, $D \to E$. Present the resulting decomposition.

Conceptual Modeling

30. Show the relation schema obtained from the ER diagram on the last page of this exam. Underline keys and show foreign keys.

Relational Queries

```
CREATE TABLE Show (
  id INT PRIMARY KEY,
  title VARCHAR(30),
  genre VARCHAR(10),
  description VARCHAR(1000)
);

CREATE TABLE Showing (
  channel VARCHAR(10),
  start TIMESTAMP,
  stop TIMESTAMP,
  show INT REFERENCES Show,
  PRIMARY KEY (channel, start),
  UNIQUE (start, stop),
  CHECK (start < stop)
);
```

31. In SQL give the show titles starting between '2016-01-01 18:00:00' and '2016-01-01 22:00:00' on TV4

32. In tuple calculus give the channels not showing any show of genre 'news'

33. In relational algebra, give the ids for shows that overlap in time. Note '$<$' works over timestamps.

34. In SQL, give the average length of shows for each channel. Note '$-$' works over timestamps to yield an `INTERVAL` type value which can be averaged.

Semi-structured Queries

An XML database contains, among other things, entities like these:

```
<customer id="55">Hillary Cruz</customer>
...
<orderdate id="00471" customer="55"
   product="tp3-004">
  2016-01-02
</orderdate>
...
<product id="tp3-004" price="39">
  Melodifestivalen Mouse Pad
</product>
```

Assume that you can access this database through the variable `$db`.

35. Write an XQuery expression that returns all the customers who have bought something.

36. Write an XQuery expression that returns all the customers who have not bought anything.

37. Write an XQuery expression that returns an entity

```
<customervalue tot="487">
  Donald Sanders
</customervalue>
```

containing the total value of each customer's orders.

Transactions

Consider two transactions over the simple database `Seats(person, seat)`. There are two seats (1,2) and two persons (Alice and Bob). Alice runs transaction 'A'. Bob runs the transaction 'B'.

Transaction A is:

```
   BEGIN;
1. select seat from seats where person = null;
2. update seats set person = 'Alice' where seat = X
3. COMMIT;
```

Transaction B is exactly the same, except 'Alice' is 'Bob'. Also X is the lowest number seat in the read query in operation 1. For example is operation 1 returns $\{1, 2\}$ then $X = 1$.

Here is a possible schedule:

```
A1   // This means operation 1 is executed for A
A2
B1
A3
B ABORT  // This means B aborts his transaction
```

38.

Give an example of a schedule with a lost update.

39.

Give an example of a schedule with a dirty read.

40.

Consider that we have thousands of seats and thousands of people all looking for the lowest number seat as possible. **Argue** for which is the best isolation level to run these transactions under. Your answer must be at least three sentences and constitute 'an argument' (Note answer may be in either Swedish or English).

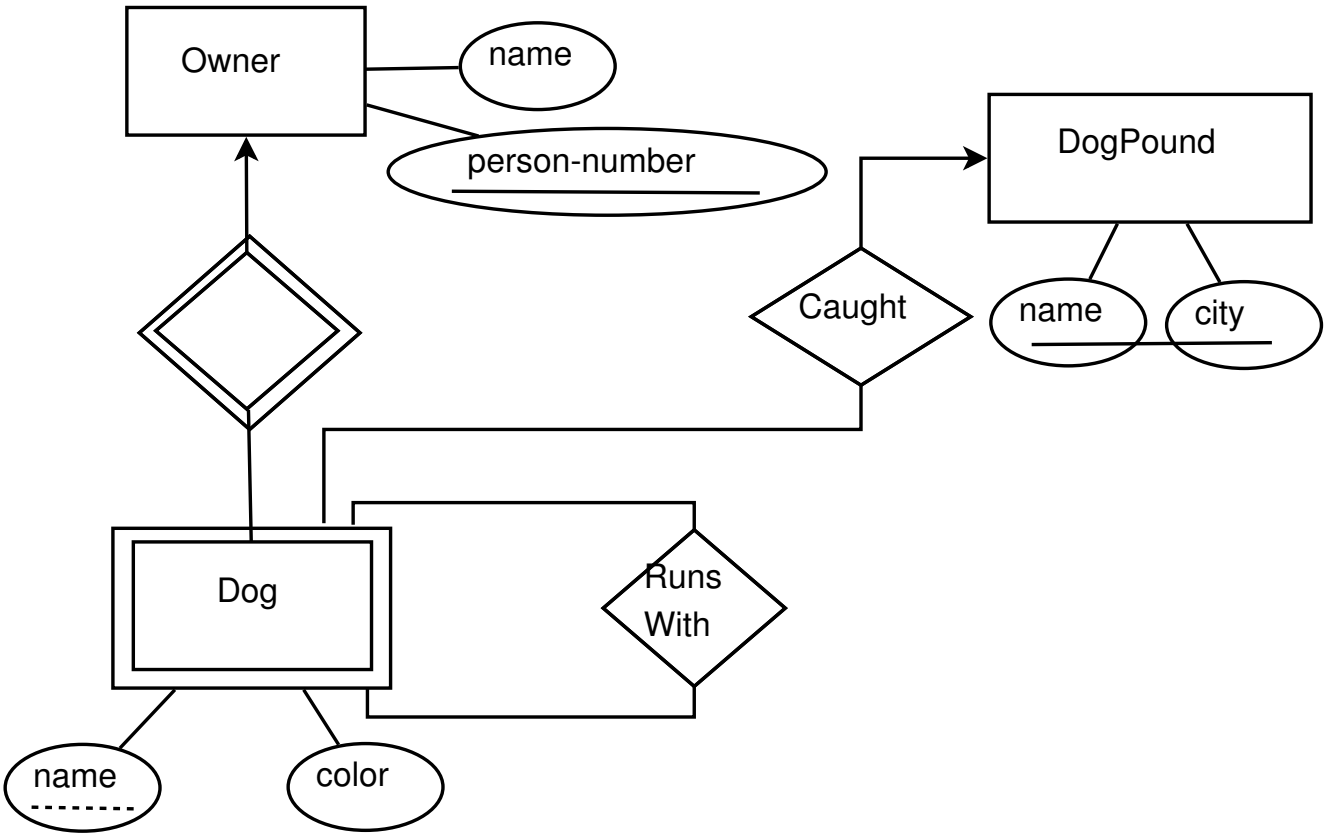
Owner
name
person-number
DogPound
Caught
name
city
Dog
Runs
With
name
color