

Write all answers on the answer sheets provided. You may refer to source code from any program you have written.

1. (3 pts each) Briefly define each of these terms:
 - (a) ODBC (don't expand the acronym; instead, define it)
 - (b) shadow copy
2. Answer the following question about indexing.
 - (a) (2 pts) What is a primary reason that indexes are used?
 - (b) (2 pts) Why wouldn't a database administrator (DBA) simply index every field?
 - (c) (4 pts) Show the SQL command that would be used to create an index on `title` field of the `album` table of the database specified in problem 6.
3. Answer the following question about B-trees.
 - (a) (2 pts) Are B-trees an implementation of ordered indexing or hashed indexing?
 - (b) (4 pts) Suppose that $t = 2$ for an initially empty B-tree. Show what the B-tree would look like after each of the following elements are inserted: A,B,C,D,E,F.
 - (c) (4 pts) Suppose that $t = 50$. How many values can be stored in a B-tree with three levels if every node is completely full?
4. Answer the following questions about transactions and locking.
 - (a) (3 pts) Give an example of a transaction.
 - (b) (4 pts) Expand the acronym ACID.
 - (c) (3 pts) How do failed transactions and aborted transactions differ?
 - (d) (3 pts) Explain how a software lock works.
 - (e) (2 pts) For what purpose are locks typically used in a DBMS?
5. (4 pts) The following schedule is conflict serializable. Given an example of a schedule that is *not* conflict serializable.

T_1	T_2
write A	
	write A
write C	
read B	
write B	
write D	
	read B
	write B

6. Consider these SQL commands that establishes a database for storing information about albums and their artists.

```
SET SEARCH_PATH TO funalbum;
CREATE TABLE artist (
    id          SERIAL,
    name        VARCHAR(30),
    hometown    VARCHAR(30),
    diversity    INTEGER,
    PRIMARY KEY (id)
);
CREATE TABLE genre (
    id          SERIAL,
    name        VARCHAR(30),
    PRIMARY KEY (id)
);
CREATE TABLE album (
    artist_id    INTEGER NOT NULL,
    title        VARCHAR(50) NOT NULL,
    year         CHAR(4),
    genre_id     INTEGER,
    PRIMARY KEY (artist_id,title),
    FOREIGN KEY (artist_id) REFERENCES artist(id),
    FOREIGN KEY (genre_id) REFERENCES genre(id)
);
```

In the current design each album has a single genre. Suppose you have been asked to redesign the database so that albums can have multiple genres associated with them.

- (a) (6 pts) Describe changes you would need to make to existing tables to support this change. If new tables are needed show the `CREATE TABLE` command(s) that would establish those tables.
- (b) Suppose, for marketing purposes, a genre called “New” will be added to each album automatically when a new album is added to the database.
 - i. (2 pts) Write an SQL command that will add a genre with the name “New” having an id number of 0 to the database.
 - ii. (4 pts) Write an SQL trigger command that will call a trigger function called `add_genre_to_new_album` whenever a new album is inserted into the `album` table.
 - iii. (6 pts) Write the PL/PGSQL trigger function mentioned in problem 6(b)ii that when called will add the genre named “New” to the newly inserted album.
- (c) (8 pts) Suppose the `diversity` field in the `artists` table has just been added for the purpose of storing the number of distinct genres attributed to the artist in the database. For example, if the artist “Dog Man” has three albums one of which is “Rock” and the other two are considered both “Rock” and “Funk” then his

diversity field should be set to 2 (because he is associated with two genres). Write a stored procedure that could be run once and would will correctly populate the **diversity** field for all artists.

7. Answer the following questions about NoSQL DBMSs.

- (a) (6 pts) Explain the CAP theorem.
- (b) (2 pts) What type of NoSQL DMBS is MongoDB?
- (c) Write MongoDB commands to accomplish these tasks:
 - i. (6 pts) Suppose you wanted to represent information about albums in a way that is similar to what was done in problem 6. Create a collection called **album** that stores the album title, year, and a list of genres. Add three entries to your collection. At least one of the documents should have multiple genres associated with it.
 - ii. (6 pts) Produce a query, or if necessary, write a function that will display all albums that have a genre of “Pop” in their list of genres. The list should not show the built-in `_id` field.