Databases Final Exam Practice

ER-Relational Mapping, SQL, Relational Design, Physical DMBS Design

Name: _____

GT account (gtg, gth, msmith3, etc): _____ Section (e.g., B1): _____

Signature: _____

- Failure to properly fill in the information on this page will result in a deduction of up to 4 points from your exam score.
- Signing signifies that you agree to comply with the **Academic Honor Code of Georgia Tech**.
- Calculators and cell phones are NOT allowed.

Completely fill in the box corresponding to your answer choice for each question.

1.	[A]	[B]	[C]	[D]
2.	[A]	[B]	[C]	[D]
3.	[A]	[B]	[C]	[D]
4.	[A]	[B]	[C]	[D]
5.	[A]	[B]	[C]	[D]
6.	[A]	[B]	[C]	[D]
7.	[A]	[B]	[C]	[D]
8.	[A]	[B]	[C]	[D]
9.	[A]	[B]	[C]	[D]
10.	[A]	[B]	[C]	[D]
11.	[A]	[B]	[C]	[D]
12.	[A]	[B]	[C]	[D]
13.	[A]	[B]	[C]	[D]
14.	[A]	[B]	[C]	[D]
15.	[A]	[B]	[C]	[D]
16.	[A]	[B]	[C]	[D]
17.	[A]	[B]	[C]	[D]
18.	[A]	[B]	[C]	[D]
19.	[A]	[B]	[C]	[D]
20.	[A]	[B]	[C]	[D]
21.	[A]	[B]	[C]	[D]
22.	[A]	[B]	[C]	[D]
23.	[A]	[B]	[C]	[D]
24.	[A]	[B]	[C]	[D]
25.	[A]	[B]	[C]	[D]

Number missed: _____ Written Score: _____

Name: _____ Student account (e.g., msmith3): _____ Section: _____

Refer to the following EER diagram for Questions 1 – 7

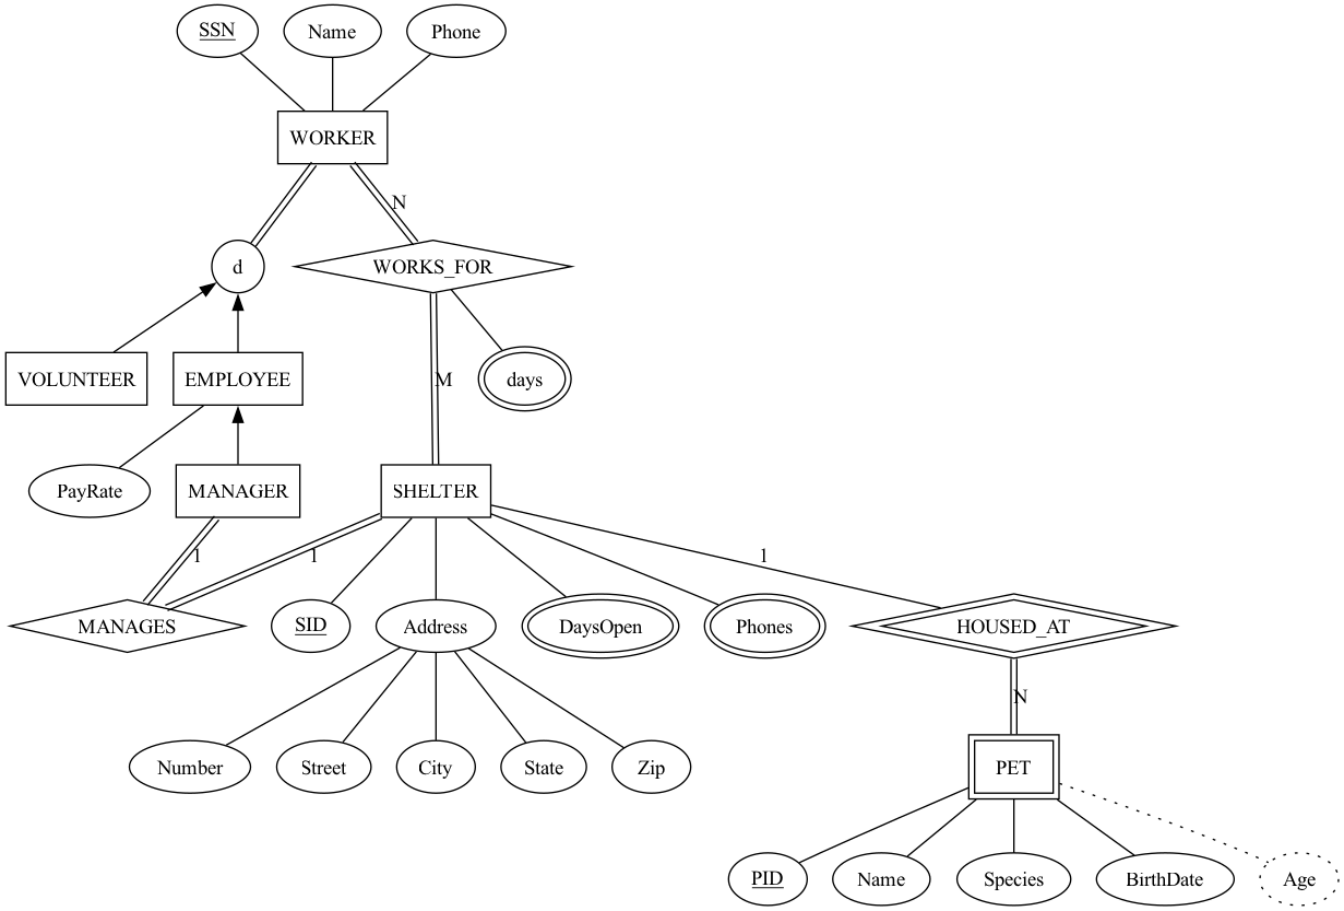

Name: _____ Student account (e.g., msmith3): _____ Section: _____

[4] 1. Which of the following (sets of) relation schemas is a correct mapping of the SHELTER entity type? (Disregard the MANAGES relationship.)

A. SHELTER(SID, Number, Street, City, State, Zip, DaysOpen, Phones)

B. SHELTER(SID, Number, Street, City, State, Zip, Phones), DaysOpen(SID, Day)

C. SHELTER(SID, Number, Street, City, State, Zip), DaysOpen(SID, Day), Phones(SID, Phone)

D. All of the above.

[4] 2. Which of the following relation schemas is a correct mapping of the PET entity type?

A. PET(PID, Name, Species, BirthDate, Age)

B. PET(PID, Name, Species, BirthDate)

C. PET(PID, SID, Name, Species, BirthDate)

D. None of the above

[4] 3. Which of the following sets of relation schemas is a correct mapping of the WORKS_FOR relationship (Disregard multivalued attributes of SHELTER.)?

A. WORKER(SSN, Name, Phone, SID), SHELTER(SID, Number, Street, City, State, Zip)

B. WORKER(SSN, Name, Phone), SHELTER(SID, Number, Street, City, State, Zip, SSN)

C. WORKER_SHELTER(SSN, SID), WORK_DAYS(SSN, SID, Day)

D. WORKER_SHELTER(SSN, SID, Days)

[4] 4. What's the least number of tables necessary to model the WORKER - VOLUNTEER - EMPLOYEE - MANAGER class hierarchy?

A. 1

B. 2

C. 3

D. 4

[4] 5. Which of the following sets of relation schemas acceptably represenents the WORKER - VOLUNTEER - EMPLOYEE - MANAGER class hierarchy?

A. WORKER(SSN, Name, Phone), VOLUNTEER(SSN), EMPLOYEE(SSN,PayRate), MANAGER(SSN)

B. EMPLOYEE(SSN, Name, Phone, PayRate, IsManager), VOLUNTEER(SSN)

C. WORKER(SSN, Name, Phone, PayRate, IsManager)

D. All of the above.

[4] 6. Which of the following create table statements creates a PET table that accurately models the PET entity type?

A. create table pet(PID int primary key, Name varchar(20), Species varchar(20), Birthdate date)

B. create table pet(PID int primary key, Name varchar(20), Species varchar(20), Birthdate date, SID int)

C. create table pet(PID int, Name varchar(20), Species varchar(20), Birthdate date, SID int, primary key (PID, SID), foreign key (SID) references shelter(SID))

D. None of the above.

[4] 7. Which of the following create table statements creates a table that accurately models the WORKS_FOR relationship? (Disregard multivalued attributes.)

A. create table worker_shelter(SSN int, SID int, days enum (M, Tu, W, Th, F))

B. create table worker_shelter(SSN int, SID int, primary key (SSN, SID), foreign key (SSN) references worker (SSN), foreign key (SID) references shelter (SID))

C. create table worker_shelter(SSN int, SID int, primary key (SSN))

D. None of the above.

Name: _____ Student account (e.g., msmith3): _____ Section: _____

Refer to the following create table statements and table data for Questions 8 – 10.

```
create table dorm (
    dorm_id integer primary key auto_increment,
    name text not null,
    spaces integer
);


create table student (
    student_id integer primary key auto_increment,
    name text,
    gpa float(3,2),
    dorm_id integer not null,
    foreign key (dorm_id) references dorm(dorm_id)
);
```

```
mysql> select * from dorm;
+---------+-----------+--------+
| dorm_id | name      | spaces |
+---------+-----------+--------+
|       1 | Armstrong |    124 |
|       2 | Brown     |    158 |
+---------+-----------+--------+
2 rows in set (0.00 sec)

mysql> select * from student;
+------------+-------+------+---------+
| student_id | name  | gpa  | dorm_id |
+------------+-------+------+---------+
|          1 | Alice | 3.60 |       1 |
|          2 | Bob   | 2.70 |       1 |
+------------+-------+------+---------+
2 rows in set (0.00 sec)
```

[4] 8. Which of the following insert statements will succeed?

A. `insert into dorm (name, spaces) values('Caldwell', 158);`

B. `insert into dorm values('Caldwell', 158);`

C. `insert into dorm (name, spaces) values(null, 158);`

D. All of the above.

[4] 9. Which of the following insert statement is certain to succeed?

A. `insert into student (name, gpa, dorm_id) values ('Cheng', 3.6, 3);`

B. `insert into student (name, gpa, dorm_id) values ('Cheng', 3.6, 1);`

C. `insert into student (name, gpa) values ('Cheng', 3.6);`

D. All of the above.

[4] 10. Which of the following delete statements will fail?

A. `delete from student`

B. `delete from dorm where name = 'Brown';`

C. `delete from dorm where name = 'Armstrong';`

D. None of the above.

Name: _____ Student account (e.g., msmith3): _____ Section: _____

Pubs Database Schema

$author(\underline{author\_id}, first\_name, last\_name)$

$author\_pub(\underline{author\_id}, \underline{pub\_id}, author\_position)$

$book(\underline{book\_id}, title, month, year, editor)$

$pub(\underline{pub\_id}, title, book\_id)$

- *author_id* in *author_pub* is a foreign key referencing *author*
- *pub_id* in *author_pub* is a foreign key referencing *pub*
- *book_id* in *pub* is a foreign key referencing *book*
- *editor* in *book* is a foreign key referencing *author*(*author_id*)
- Primary keys are underlined

Pubs Database State

r(*author*)

author_id	first_name	last_name
1	John	McCarthy
2	Dennis	Ritchie
3	Ken	Thompson
4	Claude	Shannon
5	Alan	Turing
6	Alonzo	Church
7	Perry	White
8	Moshe	Vardi
9	Roy	Batty

r(*author_pub*)

author_id	pub_id	author_position
1	1	1
2	2	1
3	2	2
4	3	1
5	4	1
5	5	1
6	6	1

r(*book*)

book_id	title	month	year	editor
1	CACM	April	1960	8
2	CACM	July	1974	8
3	BST	July	1948	2
4	LMS	November	1936	7
5	Mind	October	1950	NULL
6	AMS	Month	1941	NULL
7	AAAI	July	2012	9
8	NIPS	July	2012	9

r(*pub*)

pub_id	title	book_id
1	LISP	1
2	Unix	2
3	Info Theory	3
4	Turing Machines	4
5	Turing Test	5
6	Lambda Calculus	6

Figure 1: Pubs Database

Name: _____ Student account (e.g., msmith3): _____ Section: _____

For the questions on this page, refer to Figure 1.

[4] 11. Query giving author first name and last name who have published in CACM.

A.

B.

C.

D.

[4] 12. Query giving all authors and the books they have edited, including authors who are not book editors.

A.

B.

C.

D.

[4] 13. Query giving all author first names and last names who have published more than one pub.

A.

B.

C.

D.

[4] 14. Author of the oldest pub in the database.

A.

B.

C.

D.

[4] 15. How many rows returned by `select * from book natural join pub`?

A.

B.

C.

D.

Name: _____ Student account (e.g., msmith3): _____ Section: _____

For the questions on this page, use this relation schema and set of functional dependencies F:

ATL-TRANSIT($DriverSsn, EmpName, RouteNum, BusId, RouteDate, ServiceDate$)

$$\begin{aligned} DriverSsn &\rightarrow RouteNum \\ RouteNum, RouteDate &\rightarrow BusId \\ BusId &\rightarrow ServiceDate \\ RouteNum, RouteDate &\rightarrow DriverSsn \\ DriverSsn &\rightarrow EmpName \end{aligned}$$

[4] 16. Which one of the following functional dependencies is in F^+?

A. $RouteDate \rightarrow BusId$

B. $ServiceDate \rightarrow BusId$

C. $RouteNum \rightarrow BusId$

D. $BusId, DriverSsn, EmpName \rightarrow BusId$

[4] 17. What is $\{RouteNum, RouteDate\}^+$ with respect to F?

A. $\{RouteNum, RouteDate\}$

B. $\{RouteNum, RouteDate, BusId, DriverSsn\}$

C. $\{RouteNum, RouteDate, BusId, DriverSsn, EmpName, ServiceDate\}$

D. the empty set

[4] 18. Which of the following is a key for the ATL-TRANSIT schema?

A. $DriverSsn$

B. $\{RouteNum, RouteDate\}$

C. $\{DriverSsn, RouteDate\}$

D. Both B and C

[4] 19. What is the highest normal form that the ATL-TRANSIT schema satisfies?

A. 1NF

B. 2NF

C. 3NF

D. BCNF

[4] 20. Suppose we decompose the ATL-TRANSIT schema into

$ATL1(DriverSsn, RouteNum, BusId, RouteDate, ServiceDate)$
$ATL2(DriverSsn, EmpName)$

Does that decomposition have the lossless join property?

A. Yes

B. No

[4] 21. Suppose we decompose the ATL-TRANSIT schema into

$ATL1(RouteNum, RouteDate, BusId)$
$ATL2(DriverSsn, RouteNum, EmpName, ServiceDate)$

Does that decomposition have the lossless join property?

A. Yes

B. No

Name: _____ Student account (e.g., msmith3): _____ Section: _____

For the questions on this page, use this relation schema and set of functional dependencies F:

$ATL-TRANSIT(DriverSsn, EmpName, RouteNum, BusId, RouteDate, ServiceDate)$

$$\begin{aligned} DriverSsn &\rightarrow RouteNum \\ RouteNum, RouteDate &\rightarrow BusId \\ BusId &\rightarrow ServiceDate \\ RouteNum, RouteDate &\rightarrow DriverSsn \\ DriverSsn &\rightarrow EmpName \end{aligned}$$

[4] 22. Which attribute is fully functionally dependent on the set of attributes $\{RouteNum, RouteDate\}$?

A. $BusId$

B. $DriverSsn$

C. $EmpName$

D. all of the above

[4] 23. Which of the following attributes are prime attributes?

A. Only DriverSsn

B. Only RouteNum

C. RouteNum and RouteDate

D. DriverSsn, RouteNum and RouteDate

[4] 24. Suppose we decompose the ATL-TRANSIT schema into

$ATL1(RouteNum, RouteDate, BusId, DriverSsn)$
$ATL2(DriverSsn, RouteDate, EmpName, ServiceDate)$

Which of those schemas is in 3NF?

A. ATL1

B. ATL2

C. Both ATL1 and ATL2

D. None of the above

[4] 25. Consider the current state for our ATL-TRANSIT schema as shown below. What values could be inserted for the two missing column values, RouteNum and ServiceDate, without violating any of the FDs that have been defined for the ATL-TRANSIT schema. The domain for RouteNum is $\{10, 11, 12, 13, 14\}$ and the domain for ServiceDate is any valid date

DriverSsn	**EmpName**	**RouteNum**	**BusId**	**RouteDate**	**ServiceDate**
111-22-3333	Brown	11	101	07-07-2007	06-06-2006
333-33-4444	Smith		202	07-11-2007	07-12-2005
222-44-5555	Green	12	101	07-12-2007	
333-33-4444	Smith	10	203	07-12-2007	08-22-2006

A. The values 11 for RouteNum and '07-12-2005' for ServiceDate

B. The values 10 for RouteNum and '06-06-2006' for ServiceDate

C. The values 13 for RouteNum and '09-01-2006' for ServiceDate

D. None of the above

Physical DMBS Design

[4] 26. Storage

A.

B.

C.

D.

[4] 27. Indexing

A.

B.

C.

D.

[4] 28. Query processing

A.

B.

C.

D.

[4] 29. Transaction processing

A.

B.

C.

D.

[4] 30. ?

A.

B.

C.

D.