


8-1/2 X 11" crib sheet.

Turn in crib sheet with completed test.

1. Consider this ER diagram.



1. a. Give corresponding tables, attributes, primary keys, and foreign keys. (add or remove lines as needed)

<u>Table</u>	<u>primary key</u>	<u>attributes</u>	<u>foreign keys</u>

1. b. For your answer to 1a, give all foreign key / primary key relationships by filling in the table (add or remove lines as needed).

<u>Foreign key</u>	<u>references</u>	<u>Primary key</u>

2. Consider this database schema and instance.

T1 (A B C)
fd: A B \rightarrow C

primary key: AB

foreign key: T1 . C referencing T2 . C

T2 (C D E)
fd: C D \rightarrow E

primary key: C D

Table T1

A	B	C
1	2	3
10	20	4
10	11	4

Table T2

C	D	E
3	"sea"	"oak"
4	"sea"	"maple"
4	"floor"	"maple"

For 2a -- 2d, give the changed table contents or the error if the database cannot perform the operation. Recall the foreign key / primary key relationship.

2. a. What is the result of inserting tuple (3, "floor", "ash") to T2?

2. b. What is the result of inserting tuple (4, "floor", "ash") to T2?

2. c. What is the result of inserting tuple (5, 10, 3) to T1?

2. d. What is the result of inserting tuple (12, 12, 10) to T1?

3. Using the same database schema and original instance given in question 2.

T1 (A B C)
 fd: A B → C

primary key: AB

foreign key: T1.C referencing T2.C

T2 (C D E)

fd: C D → E

primary key: C D

Table T1

A	B	C
1	2	3
10	20	4
10	11	4

Table T2

C	D	E
3	"sea"	"oak"
4	"sea"	"maple"
4	"floor"	"maple"

3. a. Show the result of

```
SELECT A, E
FROM T1, T2
WHERE T1.C = T2.C
```

3. b. Name the table resulting from the SQL statement in 3.a: $T3 = T1 \bowtie T2$

List the functional dependencies in T3. HINT: the primary key of T3 is A B

3.c. Using the same T3 from 3. b., circle the TRUE or FALSE for each assertion.

T1 is in 3NF	TRUE	FALSE
T2 is in 3NF	TRUE	FALSE
T3 is in 2NF:	TRUE	FALSE
T3 is in 3NF:	TRUE	FALSE

4. Give the result (the tuples) after performing the following SQL statements on the UNIVERSITY DATABASE.

4.a.

```
SELECT SUM (DISTINCT stuId)
FROM Student
WHERE major = "Math"
```


4.b.

```
SELECT      Enroll.stuId
FROM        Class, Enroll
WHERE       Class.classNumber = Enroll.classNumber AND
            Class.room = "H225"
```

5. Give the SQL statement the corresponds to these relational algebra expressions or expression trees. (5a and 5b are not on the same database)

5.a. $\Pi_{T1.a, T2.b} ((\sigma_{T1.c = 3} T1) \bowtie_{T1.a = T2.a} T2)$

5.b.



6. Consider this transaction. Insert locking and unlocking statements to implement two phase locking. (there is more than one correct answer)

Transaction: X1

```
read(x)

x = x + 1

read(y)

y = x

write(y)

read(z)

if (z > 100) abort

x = x + z

write(x)

commit
```

7. Suppose two users, U1 and U2, are running the transaction X1 concurrently (interleaved).

Using your answer to #6 (with the added locking and unlocking statements), give TWO possible schedules with interleaved instructions from U1's transaction and U2's transaction.

Schedule 1:

Schedule 2:

8. Simple logging. Consider the transaction X1 given in # 6 (repeated here for your convenience)

Transaction: X1

```
read(x)
x = x + 1
read(y)
y = x
write(y)
read(z)
if (z > 100) abort
x = x + z
write(x)
commit
```

Before the transaction starts, the database contains:

```
x = 10
y = 10
z = 10.
```

A log contains (1) data update records, and (2) commit records.

A data update record contains:

```
< LSN, transaction id, data name, old image, new image >
```

where LSN is Log Sequence Number.

Give the log for the transaction X1 executing in isolation.

9. Data partitioning. Consider the UNIVERSITY DATABASE.

9. a. Give the horizontal partitioning of `Enroll` into shards where the partitioning is based on `stuId`. (Supply the tuples in each shard)

9. b. Give one possible vertical partitioning of `Student` into two partitions. Note: `stuId` is the primary key.

Short Answers.

10. A DBMS has two phase locking. Which level of granularity will cause the worst performance?

- A. field
- B. record
- C. page
- D. table (i.e., file)
- E. database

11. A DBMS has a recovery protocol, for example the ARIES protocol. Why are CHECKPOINTS useful?
(Circle all that apply)

- A. To confirm consistency
- B. To avoid anomalies
- C. To reduce the number of steps required for REDO and UNDO

12. What is a system catalog?

14. Table T (A, B, C) has primary key A. But the designer chose to have the table sorted on attribute B.

Answer TRUE or FALSE for this assertion: It is not possible to create a primary key for table T.

TRUE

FALSE

15. A particular horizontal partitioning of 10,000 record table has one shard with 100 records and the second shard with 9,900 records. Give a good reason for why the database designer chose unequal sized shards.

16. Suppose you have a schema R that has been normalized to R1 and R2. To re-create R from R1 and R2 you will join R1 and R2. What kind of join will be used ?
(Circle all that apply)

A. inner join

B. natural join

C. full outer join