

Student ID

Student Name

Total Mark

CMPUT 291
File and Data Management
Midterm Exam
17 February 2000

B2 Tamer Ozsu

- Exam duration: 50 minutes.
- Closed book, and no collaboration. No questions during the exam; if you are unsure, write down your assumptions.
- All questions need to be answered.
- Marked exam will be available on 29th of February (in class). Closing date for appeals is 14th March 2000 at 3PM.
- Answer all questions on this sheet.**

Question 1(20 pts; each 4 points) Given the following database, indicate the results that you would get by executing the queries below. Show the results as tables.

Sailors

SID	SNAME	RATING	AGE
22	Dustin	7	45.0
29	Brutus	1	33.0
31	Lubber	8	55.5
32	Andy	8	25.5
58	Rusty	10	35.0
64	Horatio	9	40.0
71	Zorba	10	16.0
74	Horatio	9	40.0
85	Art	3	25.5
95	Bob	3	63.5

Reserves

SID	BID	DAY
22	101	10/10/98
22	102	10/10/98
22	103	10/8/98
22	104	10/7/98
31	102	11/10/98
31	103	11/6/98
31	104	11/12/98
64	101	9/5/98
64	102	9/8/98
74	103	9/8/98

Boats

BID	BNAME	COLOR
101	Interlake	Blue
102	Interlake	Red
103	Clipper	Green
104	Marine	Red



04751
CMPUT 291 (B2)
OZSU, T.
FEB 00 MIDTERM
PAGES: 9

```
(a) SELECT *  
    FROM Sailors  
    WHERE Rating > 7
```

```
(b) SELECT B.Color  
    FROM Sailors S, Reserves R, Boats B  
    WHERE S.SID = R.SID  
    AND R.BID = B.BID  
    AND S.SNAME = 'Lubber'
```

```
(c) SELECT S.SNAME  
    FROM Sailors S  
    WHERE S.SID IN  
        (SELECT R.SID  
         FROM Reserves R  
         WHERE R.BID = 103)
```

```
(d) SELECT RATING, MIN(AGE)
      FROM Sailors
      WHERE AGE >= 18
      GROUP BY RATING
      HAVING COUNT(*) > 1
```

```
(e) SELECT S.SNAME
      FROM Sailors S
      WHERE S.AGE >
            (SELECT MAX(S2.AGE)
             FROM Sailors S2
             WHERE S2.RATING = 10)
```

Question 2 (23 pts) Consider the ER model given in the following figure. This model represents the operations of a pharmacy chain. Please answer the following questions regarding this model. **NOTE: Solve part (d) on the figure itself.**

(a) Can a pharmaceutical company have multiple phone numbers? If yes, then explain why. If not, what do you need to do to allow this? (2 pts)

- (b) If we delete from the database the pharmaceutical company that manufactures a drug, what happens to the drugs that the company manufactures? Justify (in one or two sentences only) your argument. (3 pts)
- (c) Similar to part (b), but instead of deleting the pharmaceutical company, what if we delete the pharmacy that sells the drug. Do we have to delete the drug too? Why or why not? (3 pts)
- (d) Modify the model so that you can represent the following (15 pts):
- Each patient has to have one and only one primary physician. Each physician has at least one patient. We want to know at least the specialty and the date of entry into the profession of each physician.
 - Instead of modeling only the fact that a patient takes certain drugs, model the fact that a patient takes certain drugs that are prescribed by a physician and the prescription date.
 - Pharmaceutical companies have long-term contracts with pharmacies. A pharmaceutical company can contract with several pharmacies, and a pharmacy can contract with several pharmaceutical companies. For each contract we want to store a start date, an end date.

SOLVE THIS PART ON THE MODEL GIVEN ON THE NEXT PAGE.

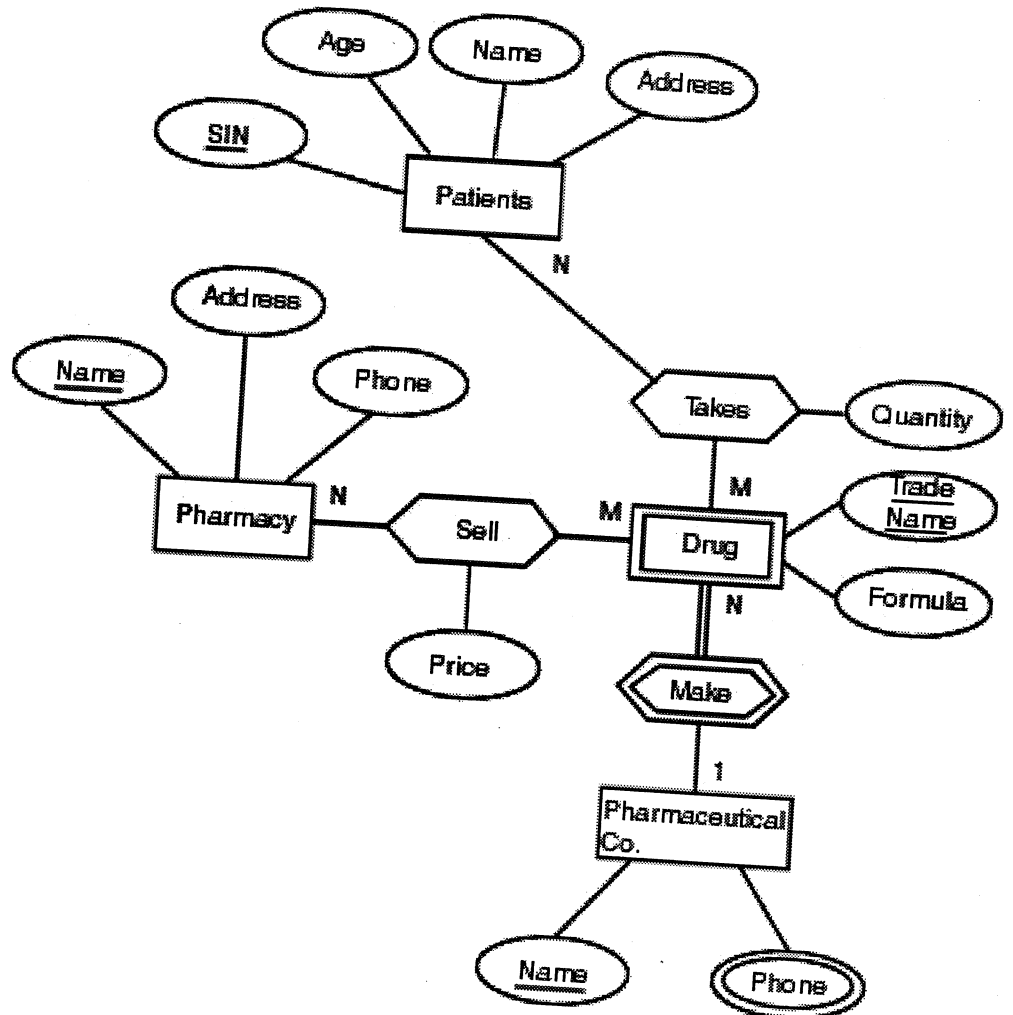


Figure for Question 2.

Question 3 (32 pts; each 4 points; NO PARTIAL MARKS) Consider the following database scheme for a motorcycle database where the underlined attributes constitute the primary keys of each relation.

Manufacturer(Name, Address, Revenue)

Model(ModelNo, EngineSize, Transmission, Type, Price, Manufacturer)

Dealer(DealerId, Name, Owner, CycleCompany, Address)

Sales(Dealer, ModelNo, Buyer, Price, Date)

Customer(CustomerId, Name, Address)

where

- (a) Manufacturer stores information about motorcycle manufacturers (e.g., Honda, Harley-Davidson) whose total revenue for the current year is stored as Revenue attribute.
- (b) Model stores information about motorcycle models that are uniquely identified by their ModelNo (e.g., 'XL883 Sportster' or 'Vulcan 1500 Drifter'). The attributes are straightforward, except that Type indicates the type of the motorcycle (e.g., cruiser or racing bike) and Manufacturer attribute is the foreign key pointing to Manufacturer relation.
- (c) Dealer relation holds information about dealers. Owner is the name of the owner. Note that CycleCompany is a foreign key pointing to Manufacturer relation.
- (d) Sales relation holds information about individual sales for a given year (i.e., the relation is wiped out at the end of each year). The Date attribute holds the sales date within the current year. The price here is the final sales price and may be different than the Price in the Model relation (which represents manufacturer's suggested retail price). Dealer attribute is a foreign key pointing to the Dealer relation, Buyer attribute is a foreign key pointing to the Customer relation, and ModelNo is a foreign key with the obvious semantics.

Formulate the following queries in SQL:

- (a) Retrieve the names of manufacturers who have motorcycles with 800cc engines.

(b) Find the total sales of each dealer for each day and display the dealer id together with this value.

(c) List the motorcycle models whose price is greater than the price of all the models of cruiser type.

(d) For each model, list the model number, the manufacturer, and the total sales (in dollars) made for that motorcycle model. Assume that at least one sale has been realized for each model.

(e) Find the names of manufacturers who have 50 or more dealers.

(f) Find the names of dealers who have made no sales.

(g) For each dealer who has made at least two sales, retrieve the dealer id, dealer name, and the total number of sales (not price, just the number of sales) that the dealer has made.

- (h) Define a view called DEALER_SALES each tuple of which will contain , for each dealer, the dealer name, the total number of unique models that the dealer has sold and the total income that the dealer has secured.

Question 4 (25 pts) Consider the following ER model. Convert it to a relational model. Show your steps.

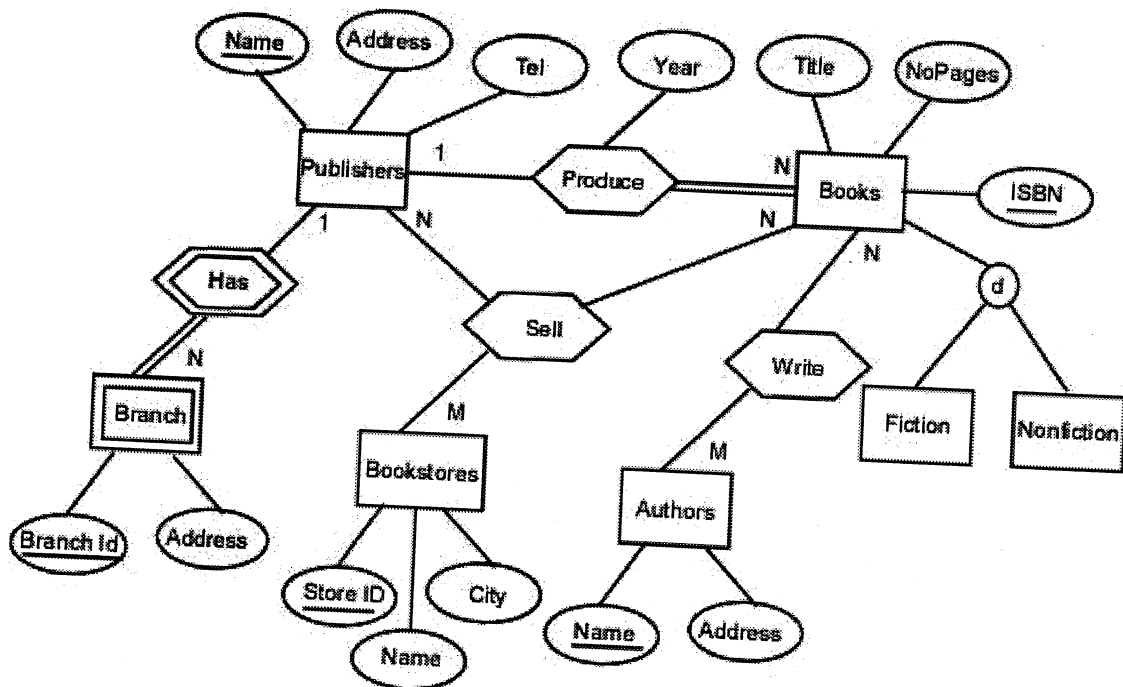


Figure for Question 4