

CS482/CS502 – Final Exam

8:00 – 10:10 am

December 13, 2006

Note:

- There are 5 questions (100 points for 25% of the total grade).
- There are also some bonus questions.
- Be short and precise in your answers.
- There are some questions that are different for **CS482** and **CS502**. Be sure that you do the correct one.

Name:

Register for: CS482 or CS502 (Please circle one)

Signature:

The questions are related to the following information about the LC-airliner, a Las Cruces based airline.

The LC-airliner stores flight information, its fleet of airplanes, and its pilots in the following tables:

- Aircraft (Airplane_ID, Type, Capacity, Base)

This table stores the information about the airplanes owned by the company. Each airplane is assigned a unique Airplane_ID, is of a specific Type (e.g. Boeing 747, Airbus 4000, etc.), and can carry a certain number of passengers. Each airplane is assigned to a base (or an airport) as well. The following is an instance of this table.

Airplane_ID	Type	Capacity	Base
1	Boeing 747	250	Las Cruces
2	Airbus 4000	210	Hatch
3	Boeing 757	360	White Sand
4	Boeing 777	500	White Sand

- Pilot (ID, Name, Year_Of_Experience, Home)

This table stores the information about each pilot worked for the company. It includes the pilot unique identification (ID), his name, and his years of experience. Like each aircraft, each pilot has a home base airport from which he can be reached in case of emergence. The following is an instance of this table.

ID	Name	Year_Of_Experience	Home
1	Son Tran	0	Las Cruces
2	Harry Me	10	Hatch
3	Forget Full	5	White Sand

- Flight (FNo, Airplane_Type, From, To, Time, Duration)

This table stores the information about each flight operated by the airliner. Each flight is identified by a unique number FNo and is associated with an Airplane_Type. The attributes From, To, Time, and Duration represent the origin and destination airport, the departure time, and the duration of the flight, respectively. The following is an instance of this table.

FNo	Airplane_Type	From	To	Time	Duration
001	Boeing 747	Las Cruces	Albuquerque	7am	1.5
002	Boeing 757	Albuquerque	Las Cruces	7pm	1.5
003	Boeing 747	Las Cruces	DC	7am	4

- Certificate (ID, Type)

This table stores the information regarding the types of airplane that a pilot are certified to fly. A pilot can be certified to flight several types of airplanes. The following is an instance of this table.

ID	Type
1	Boeing 747
2	Boeing 747
2	Boeing 757
2	Airbus 4000
3	Airbus 4000
3	Boeing 747

1. (20 points) Given the above information, identify **as many as possible** constraints that you would impose on the tables. For each constraint, give an one sentence explanation for your decision. By constraints, we mean one of the following:

- Key
- Unique constraint (do not list those you already listed as keys).
- Foreign key
- Semantic constraint

The constraints should be *reasonable* and you only need to list the constraints with a simple explanation. It is **not necessary** to give the SQL command to create them.

This part is for CS502 only: Identify one more table that you would add to the list of tables that the company has created to help them in managing their data. Give a brief explanation for your decision. Your explanation should point out the advantages of having the additional table.

2. (20 points) In the following, we say that a pilot is a possible pilot of a flight if he is certified to flight the type of airplane listed with the flight. (In the sample tables, 2 (or Harry Me) is a possible pilot for every flight while 1 (Son Tran) and 3 (Forget Full) are possible pilots for flight number 001 and 003).

Express each of the following request in relational algebra using the LC-airliner's schema (i.e. the set of tables).

- (a) List all pilots (with their id, name, and years of experience).
- (b) List all flight numbers with their possible pilots. (Note that one flight can have more than one possible pilots).
- (c) List all pilot IDs, who can operate flight number x (x is given), whose home base airport is the same as the origin of the flight x .
- (d) **Bonus question (5 points extra):** Suppose that we need to find a pilot for an emergency flight, which uses an airplane of type t , from the airport x . Someone claims that the answer of part (c), with very little modification, can be used for this purpose. Under what condition(s) do you think that this claim is correct? If so, what modification(s) will you make.

3. (20 points) Express the following queries using SQL.

- (a) List all flights of the LC-airliner with their number, origin, destination, and departure time.
- (b) List all types of airplanes owned by the company with the number of pilots certified to flight each of these types of airplanes. (Notice that the company might have some new airplanes which cannot be operated by any of its current pilots.)
- (c) List all pilots who has the highest number of certificates.
- (d) List all pilots who can be a possible pilot for the largest number of flights.
- (e) List all pilots with their years of experience in decending order.
- (f) Compute the average number of years of experience of all pilots.

4. (20 points) Suppose that the `Flight` and `Certificate` tables have the following indices:

- `Flight`: an B^+ -tree index on `FNo`, 4 levels, unintegrated
- `Flight`: an B^+ -tree index on `Airplane_Type`, 4 levels, unintegrated
- `Certificate`: an B^+ -tree index on `ID`, 4 levels, unintegrated
- `Certificate`: a hash index on `Type`

For each of the following queries, specify the index you would use in computing its result. Give an **one sentence explanation** for your choice.

(a) `SELECT f.from, f.to, f.time, f.duration`
 `FROM FLIGHT f`
 `WHERE f.fno = 3`

(b) `SELECT f.from, f.to, f.time, f.duration`
 `FROM FLIGHT f`
 `WHERE f.airplane_type = 'Boeing 747'`

(c) `SELECT c.type`
 `FROM CERTIFICATE c`
 `WHERE c.id >= 10 and c.id <= 20`

(d) `SELECT c.id`
 `FROM CERTIFICATE c`
 `WHERE c.type = 'Boeing 747'`

5. (20 points) Suppose that we would like to execute the following query:

```
SELECT f1.fno, f1.from, f1.to
FROM   flight f1, flight f2
WHERE  f1.to = f2.origin AND
       f1.time + f1.duration < f2.time AND
       f2.destination='DC'
```

Answer the following:

- (a) What does the SQL query compute?
- (b) It is known that the `Flight` table is stored in 1000 pages and the main memory of the computer is also 1000 pages. Assume that the result of the query is about 3 pages and assume that you can process (e.g. sort, search, etc.) data in the main memory without any cost. What will be the minimal number of page IOs required for executing the query? Give an explanation for your answer.
- (c) Which index will you create in order to cut down the number of page IOs required for the above query? Explain how your index help in reducing the number of page IOs.