

Spring 2007 CS186 Discussion Section:
Week 7, 10/08 - 10/12

Your Friendly TAs

October 8, 2007

1 SQL

Consider the following schema on table 1:

```
Flights(flight_num, source_city, destination_city)
Departures(flight_num, date, plane_type)
Passengers(passenger_id, passenger_name, passenger_address)
Bookings(passenger_id, flight_num, date, seat_number)
```

Table 1: Schema for the supply chain database.

The key fields are underlined. Express the following queries in SQL (feel free to abbreviate relation and attribute names and to use INTERSECT and EXCEPT if you need to):

1. Find the `passenger_id` of all passengers who have a seat booked on a plane of type “747” from San Francisco to Washington. *Do not return any duplicate values.*
2. Find the cities that have direct (non-stop) flights to both Honolulu and Newark.
3. Find the `flight_num` and `date` of all flights for which there are no reservations.
4. Find the `passenger_name` of all passengers who have a seat booked on at least one plane of *every* type.
5. Print an ordered list of all source cities and the number of distinct destination cities that they have direct (non-stop) flights to, with airplanes of type “747”. The list should be ordered in decreasing number of destinations and should contain only those source cities that have flights to 25 or more distinct destinations.

For example, the output should look like:

source_city	NumDestinations
Chicago	228
Atlanta	100
...	
San Francisco	25