

Projet de bases de données

You are requested to design and develop an information system that manages the data of several airlines. The system has to provide, as its main functionality, an on-line flight reservation service. It also has to serve as the main data management tool for the airlines.

The design process of an information system always starts with a series of interviews where engineers meet clients to understand which are the requirements for the system to be designed. The outcome of such interviews is usually a set of informal specifications of the data that the system has to model and the functionalities that it should provide. These interviews usually go through higher and higher level of specification and often serve as feedback for the initial design phases.

Assume that you (E) met the client (C) and that the following is the report of your interviews.

E. Good morning.

C. Good morning, please take a seat...

C. We would like to provide an on-line flight reservation service capable of handling uniformly data of several airlines. We would like the system to be used on the one side by customers willing to book their flight, and on the other side by each airline, to keep their data up-to-date.

E. Do you want each airline to be able to access only its own data?

C. Yes, although different airlines may share the same flights. In this case each airline must be able to access the data of all the flights it shares.

E. I would need to ask you some more detailed questions about the nature of the data handled by the airlines.

C. Sure, go ahead!

E. When you say “flight”, what do you mean?...

C. Ehm...

E. I'll be more precise. Today I took the flight AF321 from Paris CDG to Rome FCO, but tomorrow, or in 6 months, there will be an AF321 AirFrance flight from Paris CDG to Rome FCO as well. Are these two different flights? What does the flight code mean then?

- C. Oh, now I understand your question! A flight has a code (321 in your example) which is unique within the airline (AirFrance in your example). The airline code (AF) followed by this unique code, gives you what we call *flight number*. A flight with a given flight number may be scheduled to fly several times in a week.
- E. I see... so today's and tomorrow's flights are two instances of the same flight, the flight AF321. But what can be changed between two different scheduled instances? The departure time I guess...what about the departure and destination airports?
- C. The departure and destination airports do not change with the scheduled instance: for instance all scheduled instances of the flight AF321 fly from Paris CDG to Rome FCO.
- E....let me take a note, this is important, it is what we call functional dependency in database theory...
- C. Functional what??
- E. Never mind, just talking to myself...please continue.
- C. On the other hand different scheduled instances of the same flight may have different departure times, as well as different durations. The way all this changes depends on the airline timetable...
- E. I will ask you later about the organization of the timetable, now I would like to understand more about the flights.
- C. As you prefer!
- E. Is a flight with a given flight number always a direct one?
- C. No. There are flights with stops. For instance the (Lufthansa) LH626 flight from Frankfurt to Doha stops in Bahrain.
- E. What does it mean that it stops? Do passengers have to get off and take another plane?
- C. No. As a general rule, every instance of a single flight always uses a single aircraft. If the flight has stops, passengers that have booked the whole flight just stay in the plane during the stop. But the flight may be booked also partially. For instance, passengers may also book the LH626 flight from Frankfurt to Bahrain, or from Bahrain to Doha. So the stop is also used to make some passengers get off and some other passengers board.
- E. I've never been on such a flight...
- C. Well, indeed they are not many...
- E. So the majority of flights have no stops...but can a flight have several stops?

- C. In principle yes.
- E. Can different instances of the same flight have different stops?
- C. No. The segments the flight consists of, only depend on the flight number, not on the scheduled instance of the flight.
- E. Good. Now, how does the flight sharing works?
- C. A flight has a main operating airline, and its flight number within the operating airline identifies the flight. But the flight can be shared by other airlines. In this case the same flight has different numbers, one for each airline that shares it. As usual the flight number is unique within the airline. So for instance the flight AF5467 is also operated by KLM (whose airline code is KL) where it has flight number 3456. So AF5467 and KL3456 are the same flight.
- E. Ok, things are getting clearer...so let's move to timetables now. How do the airlines schedule their flights?
- C. Usually a flight is scheduled to fly on some given days of the week (monday, wednesday and thursday, for instance), each day with a given departure and arrival time – as well as a given departure and arrival time for each of the possible stops of the flight. This is what we call the timetable of the flight.
- C. Do these departure and arrival times actually change every day of the week where the flight is scheduled?
- C. Usually they do not. but the week may be partitioned into a few blocks, for instance the LH626 flight from Frankfurt to Doha right now is scheduled as follows: from Monday to Friday it departs at 13:25, it arrives in Bahrain at 21:45, it leaves Bahrain at 22:30 and it arrives in Doha at 23:30. While on Saturdays and Sundays it departs at 10:50, it arrives in Bahrain at 20:10, it leaves Bahrain at 21:00 and arrives in Doha at 21:55.
- E. Why do you say “right now”, do you mean that the timetable of a flight may be different in different periods of the year?
- C. Of course. For instance the above time table of the flight LH626 was valid from Jan 12th to March 15th, 2009, but it has been replaced by another timetable from March 16th to Sep 14th, and in the rest of 2009 the flight LH626 was not scheduled at all.
- E. Do you want to keep record also of the past timetables?
- C. Yes, as well as of the past bookings, it can be be useful for statistical analysis.
- E. Can you tell me more about the booking service you want to provide?

- C. The main services will be to compute one or more itineraries corresponding to the user request, and to allow the user to purchase flight tickets corresponding to the chosen itinerary.
- E. What is an itinerary?
- C. It is a sequence of connected flights. Connected means that a flight in the sequence has to leave from the same city where the preceding flight arrives.
- E. You don't require that the change takes place in the same airport?
- C. No. But in the case the transfer requires an airport change we want the customer to be warned. Moreover in general we would like the computation of the itinerary to be aware of the transfer time.
- E. How?
- C. The proposed itineraries should not have any transfer time of less than two hours. If there is an airport change, than the minimum transfer time goes up to 3 hours.
- E. I guess a stop of a flight is not considered as a flight transfer.
- C. Right, it is not a flight transfer. So no check has to be performed on the stopping time when computing the itinerary.
- E. But how do I know that a change is actually a stop rather than a transfer?
- C. It's easy: in a stop both flights will have the same flight number.
- E. Got it.
- C. And you should also remember that all flight times are shown in the local time zone.
- E. Yes. . . I guess that means I will have to maintain this time zone information somewhere to compute the actual flight duration. . .
- C. Sorry?
- E. Er, nothing, I was just thinking out loud. . . By the way, how do you represent this time zone data?
- C. Time zones are usually given as an offset with respect to the Greenwich Meridian Time.
- E. Good, it's pretty standard. Paris time zone is encoded as +0100: that's one hour and zero minutes after GMT. Similarly, San Francisco is -0800 or New Dehli +0530.
- C. If you say so. . .

- E.** Never mind... Earlier you mentioned you wanted customers to actually purchase flights, I suppose this means you have a pricing system?
- C.** Indeed. The base price of each flight is computed from the distance: 0.20 EUR/km for economy class and 0.50 EUR/km in business class. Additionally, for the economy class, there are a number of discounts. If the ticket is booked at least 28 days before the departure date you have a 60% discount. If you book it between 27 and 14 days before departure it's only 40% discount. Finally, between 13 and 7 days in advance you still get 20% discount. Otherwise you pay the full base price.
- E.** I think I get the idea. And what about return trips? Do I have the same rate on both flights?
- C.** No, each discount is computed for each flight independently.
- E.** I would also need to know how many seats are dedicated to business class travelers.
- C.** True. 15% of the total number of seats of the aircraft are reserved for the business class.
- E.** Now, could you tell me a little bit about the overall user experience you were thinking of?
- C.** Sure, a typical user who wants to book a flight first enters the cities between which he wants to fly, the dates when he wants to travel, whether he wants to book a return trip, and the class (economy or business) he prefers. We only allow one ticket per one passenger to be booked each time. At this point, the system should present the user with a list of possible itineraries, ordered by price or by total trip time at the user choice. If the user selects to purchase one of these flights he will have to identify himself if he has already an account or be proposed to register himself. He can then proceed to the actual booking. Registered users, when they come back, should also have the option to see all the bookings they have made with the system.
- E.** And will the airline companies access the same database?
- C.** Yes, it will be the same database but their accounts will have special permissions. They will be able to change flight schedules, add new flights or cancel others. When a flight schedule is changed, the system should check the itineraries that have been already booked which are affected by the change. If the flight is cancelled, or the new schedule breaks the 2- or 3-hour minimum delay for a change, the system should give the airline the list of all the concerned passengers.
- E.** And could you tell me a little bit more about the data the system will have to manage?

- C. Sure. We manage the schedule of 17 companies which operate flights with 64 distinct types of aircraft between more than 600 different airports all over the world. This results in about 30000 different schedules.
- E. And do you have an idea of the frequency of each of the different operations?
- C. Roughly, we expect schedule changes to concern less than 1% of the flights every week. Itinerary lookups and bookings are much much more frequent operations. We expect dealing with several dozens of lookups per second, 5% of which will turn into actual bookings.
- E. That many?
- C. Well, you know, many flights are running each day, each with 150-200 seats and we do sell most of them.
- E. I see. . . Well, I think I have all the information I need to get me started. Thank you.
- C. Thank you.