McMaster University
SFWRENG 3DB3 Fall 2024
Assignment 1
Due: Sept 27, 2024 at 10:00pm

September 18, 2024

Relational Model (15 marks)

1. (15 marks) Consider the relations PLAYERS and PLAYS given by the schemas below.

   - `PLAYERS (playerID, firstName, lastName, gender, DOB, height, weight, drafted)`
   - `PLAYS (playerID, teamID, teamName, number, position, startYear)`

   `PLAYERS` provides information on all basketball players in the league, giving the playerID, first name and last name of the player, the gender, the date of birth (DOB), the player's height and weight, and the year they were drafted into the league.

   `PLAYS` provides information about which players play on which teams. A player with playerID plays on a team with a teamID and team name. The player has a number, the position they play on the team, and the year they started playing with this team. For example, playerID 5 plays with teamID 1, the Toronto Raptors, with the number 4, in the point guard position, since 2021.

   Given these schemas, answer the following questions:

   a) (9 marks) Identify three candidate keys. For each candidate key, describe the key, and briefly state the assumptions or conditions under which each candidate key would be valid.

   b) (6 marks) List three integrity constraints that should hold over these relations. For each constraint, describe in one sentence why your constraint is necessary.

E-R Modeling (35 marks)

2. You will prepare an E-R diagram describing the schema of airline operations storing information in an airline database. MacAir Aviation manages flight operations, passenger services, fleet maintenance, and staff. The company, henceforth referred to as "MacAir", has hired you to design their database.

MacAir wants to store information about people, where a person is represented with a person ID, name, age, and phone number. There are four types of persons: passenger, pilot, cabin crew, and ground staff. A passenger has a dietary preference (e.g., 'Vegan', 'Gluten-Free', 'Lactose-Free', etc.). A pilot, and a cabin crew both have a position (e.g., 'Captain', 'First Officer', etc.) and a salary. Ground staff have attributes for salary and department (e.g. Billing and invoicing, Information Technology, etc.).

An airline ticket has a 13-digit numeric ticket number, a seat number (e.g., 38A, 2E, etc.), and a class ('E', 'B', or 'F', representing economy, business, and first-class, respectively). Passengers book one or more tickets through a travel website (e.g., 'Expedia', 'SkyScanner', etc.) with an associated price. A ticket is bought by exactly one passenger.

MacAir records an airline with an identifying alias, which is a 2-letter alphabetic code ('AC' for Air Canada), and the airline name (e.g., 'Air Canada'). Airplanes have a serial number, a manufacturer, and a model (e.g. 737MAX). A pilot flies many airplanes, however, an airplane must be flown by at least one pilot. A cabin crew member works for at most one airline, and an airline has to have at least one cabin crew member working for it. An airline must own at least one airplane, but an airplane is owned by exactly one airline.

A country has a code (a 3-letter alphabetic code, e.g., 'CAN' for Canada), a name, and a continent. An airport has an IATA code (International Air Transport Association, 3-letter alphabetic code, e.g., 'YYZ' for Toronto Pearson Airport), a name, and a city. A country has zero or more airports, however, an airport must be in exactly one country. An airline belongs to exactly one country, but a country can have many airlines. Ground staff work for at most one airport but an airport must have at least one ground staff.

A (flight) route is represented with a numeric ID, the number of stops (e.g., 0 for nonstop), and the duration (in hours). A route contains exactly one source airport and exactly one destination airport (e.g., source airport: 'YYZ', destination airport: 'MCO'). However, airports serve as the source or destination on many routes. An airline has many routes around the world, and a route is used by many airlines.

The entity 'Scheduled Flights' contains all flights that serve a route. Scheduled flights are defined via an alpha-numeric flight number, departure date, arrival date, scheduled departure time, scheduled arrival time, actual departure time, and actual arrival time. A scheduled flight contains exactly one route, but a route participates in many (scheduled) flights. For example, the 'YYZ' to 'MCO' route appears in the scheduled flights for `(AC1670, Sept. 13, Sept 13, 17:45, 20:35, 18:00, 20:50)` Airlines use at least one scheduled flight to conduct operations, but a scheduled flight is associated to exactly one airline.

A ticket is bought for exactly one (scheduled) flight, and there must be at least one ticket purchased for a (scheduled) flight. Baggage is associated to exactly one ticket. We record the type of bags (i.e., carry-on, checked, oversized, strollers), total quantity of bags for each type (e.g., 2 carry-on bags, 2 checked bags, 1 stroller, total weight of all bags for a type (e.g., 30kg for carry-on bags, 60kg for checked bags, 5kg for stroller), and whether the bags (per type) are fragile. A ticket is associated to many (types of) bags.

a) `ER.pdf`: Draw the ER diagram capturing the described requirements. You may use any drawing tool of your choice, but please ensure your ER diagram is clearly readable, and the notation you use is clear and consistent (i.e., notation from the lecture slides or textbook).

Handwritten models will not be accepted, this includes pictures of handwritten models (i.e., receive a mark of zero).

b) `ERDesc.pdf`: Give a brief (one sentence) description of each of your entities and relationships, and any constraints that exist. For example, X is a weak entity with attributes (a, b, c), and has a many-one relationship with Y.

c) `airline.ddl`: Provide the corresponding DB2 'CREATE TABLE' statements describing the relational schema. Please include all your statements in an executable script (`airline.ddl`) that can be run on the DB2 command line, in a single command. Ensure that your script runs on the CAS DB2 server. Scripts that do not execute on these servers will not be marked (receive a mark of zero). A sample script template has been provided in the assignments folder to help you get started.

Grading

This assignment is worth 12% towards your final grade.

Submission

All files are to be submitted using the Avenue system. Please ensure you submit all files with the correct names, as described below:

1. For Q1: Please include all your answers in a file `relation.pdf`.
2. For Q2: Submit files `ER.pdf, ERDesc.pdf, airline.ddl`.