2021

Appendix: Essential Aspects of Physical Design and Implementation of Relational Databases

Tatiana Malyuta
CUNY New York City College of Technology

Ashwin Satyanarayana
CUNY New York City College of Technology

Appendix 1. Case Studies

Each of the suggested case studies has a description of the user requirements[1] and a corresponding logical data model. In the logical models, the primary keys of relations are underlined, and the foreign keys are in italic; the names of foreign key attributes are the same as the names of the respective primary keys of the parent relations.

The cases can be used[2] for designing the physical data model (including the distributed solution), implementing security measures for both the centralized and the distributed approaches, programming either separate queries or transactions for the described tasks, deciding how to improve the performance of these queries and transactions, and maintaining data consistency.

Manufacturing Company

The company has offices in three cities: Boston, New York, and Cleveland. The database has to support data about departments and employees of the company. For the departments, the company needs to keep their codes, names (names are unique), locations, and types; for the employees – IDs, names, and types (full-time, part-time, or consultant). Each employee is assigned to one department. In addition, each employee has a title, which defines the employee's salary; salaries cannot exceed $100,000.

Departments of the company are located in three cities, and the office in each city needs to support local data, that is, data about all departments located in that city and the employees assigned to these departments. Each office has a database server.

The logical model of this database is the following:

Title (titleCode, titleDescription, salary)

Department (deptCode, deptName, location, deptType)

Employee (ID, emplName, emplType, *deptCode*, *titleCode*)

The database is expected to support data about 200 departments and 4000 employees.

Users work with an application to manage (insert, delete, update, and select data) data about the departments and employees of the respective location. Such application performs several modifications of the data about departments and several hundred modifications of the data about employees daily.

This application requests the data about employees of a particular department, including ID, name, and title several thousand times per day. Managers of each office also produce quarterly reports about the local departments and the number of employees in each of them.

An additional application is used by managers of the New York office to produce monthly reports about the

[1] Please note that the case studies are deliberately anecdotal – we are confident that complex concepts of physical design and implementation can be explained and dealt with on such simple examples so that students could focus on these concepts without being overwhelmed with complexities of the case. This nature of the case studies needs to be explained to students.

[2] Assumptions should be made if the supplied information is not sufficient for making decisions.

number of all company's employees of each employee type and the total salary of employees of each title in the company.

Users in each office are allowed to modify and retrieve all relevant data, but they are not allowed to access data about departments and employees of other offices. New York users of the the additional application can retrieve any data for any office. They can modify data about titles only, and perform the transfer of an employee from one department to another.

This case and the following sample tables are used in the book for demonstrations:

Title

Title Code	Title Description	salary
T1	Accountant	10000
T2	Analyst	20000
T3	Programmer	30000
T4	DBA	40000
T5	Manager	50000

Department

Dept Code	Dept Name	location	deptType
001	Computer Center	Boston	IT
002	Budget	New York	Business
003	Marketing	Boston	Marketing
004	Database Support	Cleveland	IT
005	Purchasing	New York	Business

Employee

ID	Empl Name	Empl Type	Dept Code	Title Code
1	John	Full-time	002	T1
2	Adam	Consultant	001	T3
3	Mary	Part-time	004	T4
4	Peter	Full-time	003	T2
5	Scott	Consultant	002	T1
6	Susan	Full-time	005	T5
7	Alex	Part-time	004	T2

Hospital

The hospital has four clinics. The database application has to support data about patients, doctors, and treatments of patients. Each patient is defined by the patient number, SS, name, date of birth, and address. A patient is assigned to one clinic, however, in some special cases, a patient can be transferred to another clinic.; the date when the patient was admitted to a clinic is stored in the database. Doctors of the hospital have an ID, name and pager number; doctors are scheduled for appointments hours in different clinics. Each clinic supports data about the treatments of patients by doctors, including the date of the treatment and the doctor's comments.

Each clinic has a database server. The office of each clinic handles the data about its patients, treatments and doctors' schedules. The central office of the hospital is using the database server of the first clinic; it deals with data about all patients, doctors and treatments of the hospital.

The logical model of the database is the following:

Doctor (doctorID, name, pagerNumb)

Patient (patientNumb, SS, name, dateOfBirth, address, dateAdmitted, clinic)

Treatment (treatmentcode, description)

Pat_Doc_Treat (*patientNumb, doctorID, treatmentCode*, date, comments)

Schedule (*doctorID*, date, clinic, numbOfWorkHours)

There is an application that supports (insert, delete, update and select) data about patients of each clinic, their treatments and the schedules of the doctors. For each clinic, there are several thousand modifications to data about patients and treatments each day. Several hundred times a day users supporting a particular clinic request data about:

- A doctor, including the doctor's ID, name, and number of patients treated by the doctor in the clinic for a given date interval
- A patient, including the patient's number, name, treatment code and number of times the patient had the treatment.

Database users of each clinic are allowed to modify and retrieve data about patients and their treatments in their clinic, but they cannot read or modify data about patients of other clinics. Users in a clinic can also read data about doctors and their schedules in the clinic; they can modify schedules of the doctors in the clinic of their responsibility.

The application in the central office maintains data about doctors. Several times per month users of this application request reports for all doctors, including ID, name, clinic and number of hours worked in the clinic for a given date interval. With the help of this application, the management of the hospital performs transfers of patients with their treatment histories to another clinic.

Car Rental Company

A car rental company has agencies in three cities: X, Y, and Z. Each agency deals with rentals of its vehicles only. The database of the company supports data about vehicles, clients, and rentals of vehicles to clients. For vehicles, the company stores vehicle ID, make, type (truck or not truck), year, and total mileage. The data about rentals include the ID of the client who rented the vehicle, the number of days, and the mileage increase.

The relational model of the database is the following:

Vehicle (vID, make, type, year, totalMileage, city)

Client (clientID, name, address, licenseNumb)

Rent (*vID, clientID*, date, numberOfDays, mileage)

Each agency maintains data about rentals of its own vehicles with the help of a local application, which supports insert, delete, update, and select operations on the data about vehicles and rentals. Each agency has several hundred vehicles and several thousand rentals per week.

Additionally, the agency in X produces a monthly report about all rentals during that month with information about each rented vehicle and the number of days the vehicle was rented during the month.

The office in Y is tasked to produce once a month a report about the rentals of trucks similar to the monthly report of the agency X. Once a month this agency updates the data about the total mileage of trucks based on the information from truck rentals during the month. The office in Z performs similar accounting and updating of the total mileage for all other vehicles.

Database users in each agency are in charge of the support of data about local vehicles and rentals; they cannot retrieve or modify data about vehicles of other agencies (except modifications of total mileage provided in Y and Z). The managers from Y can retrieve data about trucks and their rentals, and can modify the total mileage of trucks. Managers from Z can perform similar functions with the data about all other vehicles.

Department of Health

The Department of Health in New York has four offices for processing immunization records of children. The first office is in charge of records of children from the Bronx; the second office works with records of children from Manhattan; the third office processes records of children from Queens; and the fourth office maintains records of children from Brooklyn and Staten Island. In addition to the basic information about the children, the database includes the data for all immunizations and immunization events, including the immunization code, when it was given, and to which child.

The relational model of the database is the following:

Child (childID, firstName, lastName, dateOfBirth, street, city, ZIP, phone, borough)

Immunization (immunizationCode, description)

Immunization_Event (*childID, immunizationCode*, date)

Each office supports the data for respective children and their immunization events. In addition to all the necessary data modifications, the office daily processes several thousand requests for immunization events of a child (including immunization code and description, and date of the immunization event), given the child's ID or name and date of birth.

The fourths office, in addition to supporting the local immunization information, produces weekly reports about the immunization events of a particular type for all children in New York. This office also maintains (updates, deletes, inserts) the data about immunizations.

The database users in can retrieve and modify all respective data, but cannot access data of other offices. A group of database users in the fourths office can retrieve data about all children and their immunization events.

Financial Company

The financial company ABC provides tax and audit services to its clients. One group of the company's

employees is in performing tax services, and the other is performing audit services. Each group provides corresponding services, supports data about them, and produces various reports on them. The database of the company contains data about services (including service type: tax or audit), employees, including the type of service the employee provides, and the company's activities, including service provided, the client served, the date of the service, and the amount charged for the service.

The company has two database servers. The logical model of the database is the following:

Service (serviceCode, description, serviceType)

Employee (employeeID, name, serviceType)

Client (clientID, name, address)

Activity (*clientID, employeeID, serviceCode*, serviceDate, amountCharged)

Employees of each group are allowed to read the data about services provided by their group only. Each of employees supports data about his or her activities and is not allowed to see or modify data about activities provided by other employees. In addition, each employee produces quarterly reports about clients he/she served during the quarter.

Several users of the database are in charge of support of data about clients and services; they can insert, delete, and update data about clients. These employees are also maintaining the changes when an employee is assigned to perform another service type, e.g. an employee who has been performing tax services is assigned to perform audit services.

The company's managers produce monthly reports about activities of the employees, including the number of services the employee provided during the month and the total amount charged for the services. Managers can read any data; however, they cannot modify data.