

University of Technology
الجامعة التكنولوجية



Computer Science Department
قسم علوم الحاسوب

Database
قواعد البيانات

Lect. Nada Najel & Lect. Osama Younis
م. ندى نجيل و م. أسامة يونس



cs.uotechnology.edu.iq

Preface

(Fundamentals of Database Systems, Third Edition)



© Copyright 2000 by Ramez Elmasri and Shamkant B. Navathe

Contents of This Edition
Guidelines for Using This Book
Acknowledgment

s

This book introduces the fundamental concepts necessary for designing, using, and implementing database systems and applications. Our presentation stresses the fundamentals of database modeling and design, the languages and facilities provided by database management systems, and system implementation techniques. The book is meant to be used as a textbook for a one- or two-semester course in database systems at the junior, senior, or graduate level, and as a reference book. We assume that readers are familiar with elementary programming and data-structuring concepts and that they have had some exposure to basic computer organization.

We start in Part 1 with an introduction and a presentation of the basic concepts from both ends of the database spectrum—conceptual modeling principles and physical file storage techniques. We conclude the book in Part 6 with an introduction to influential new database models, such as active, temporal, and deductive models, along with an overview of emerging technologies and applications, such as data mining, data warehousing, and Web databases. Along the way—in Part 2 through Part 5—we provide an in-depth treatment of the most important aspects of database fundamentals.

The following key features are included in the third edition:

- The entire book has a self-contained, flexible organization that can be tailored to individual needs.
- Complete and updated coverage is provided on the relational model—including new material on Oracle and Microsoft Access as examples of relational systems—in Part 2.
- A comprehensive new introduction is provided on object databases and object-relational systems in Part 3, including the ODMG object model and the OQL query language, as well as an overview of object-relational features of SQL3, INFORMIX, and ORACLE 8.
- Updated coverage of EER conceptual modeling has been moved to Chapter 4 to follow the basic ER modeling in Chapter 3, and includes a new section on notation for UML class diagrams.
- Two examples running throughout the book—called COMPANY and UNIVERSITY—allow the reader to compare different approaches that use the same application.
- Coverage has been updated on database design, including conceptual design, normalization techniques, physical design, and database tuning.

Part 1: Basic Concepts

(Fundamentals of Database Systems, Third Edition)



Chapter 1: Databases and Database Users

Chapter 2: Database System Concepts and Architecture

Chapter 3: Data Modeling Using the Entity-Relationship Model

Chapter 4: Enhanced Entity-Relationship and Object Modeling

Chapter 5: Record Storage and Primary File Organizations

Chapter 6: Index Structures for Files

Chapter 1: Databases and Database Users

1.1 Introduction

1.2 An Example

1.3 Characteristics of the Database Approach

1.4 Actors on the Scene

1.5 Workers behind the Scene

1.6 Advantages of Using a DBMS

1.7 Implications of the Database Approach

1.8 When Not to Use a DBMS

1.9 Summary

Review Questions

Exercises

Selected Bibliography

Footnotes

Databases and database systems have become an essential component of everyday life in modern society. In the course of a day, most of us encounter several activities that involve some interaction with a database. For example, if we go to the bank to deposit or withdraw funds; if we make a hotel or airline reservation; if we access a computerized library catalog to search for a bibliographic item; or if we order a magazine subscription from a publisher, chances are that our activities will involve someone accessing a database. Even purchasing items from a supermarket nowadays in many cases involves an automatic update of the database that keeps the inventory of supermarket items.

The above interactions are examples of what we may call **traditional database applications**, where most of the information that is stored and accessed is either textual or numeric. In the past few years, advances in technology have been leading to exciting new applications of database systems.

Multimedia databases can now store pictures, video clips, and sound messages. **Geographic information systems (GIS)** can store and analyze maps, weather data, and satellite images. **Data warehouses** and **on-line analytical processing (OLAP)** systems are used in many companies to extract and analyze useful information from very large databases for decision making. **Real-time and active database technology** is used in controlling industrial and manufacturing processes. And database search techniques are being applied to the World Wide Web to improve the search for information that is needed by users browsing through the Internet.

To understand the fundamentals of database technology, however, we must start from the basics of traditional database applications. So, in Section 1.1 of this chapter we define what a database is, and then we give definitions of other basic terms. In Section 1.2, we provide a simple UNIVERSITY database example to illustrate our discussion. Section 1.3 describes some of the main characteristics of database systems, and Section 1.4 and Section 1.5 categorize the types of personnel whose jobs involve using and interacting with database systems. Section 1.6, Section 1.7, and Section 1.8 offer a more thorough discussion of the various capabilities provided by database systems and of the implications of using the database approach. Section 1.9 summarizes the chapter.

The reader who desires only a quick introduction to database systems can study Section 1.1 through Section 1.5, then skip or browse through Section 1.6, Section 1.7 and Section 1.8 and go on to Chapter 2.

1.1 Introduction

Databases and database technology are having a major impact on the growing use of computers. It is fair to say that databases play a critical role in almost all areas where computers are used, including business, engineering, medicine, law, education, and library science, to name a few. The word *database* is in such common use that we must begin by defining a database. Our initial definition is quite general.

A **database** is a collection of related data (Note 1). By **data**, we mean known facts that can be recorded and that have implicit meaning. For example, consider the names, telephone numbers, and addresses of the people you know. You may have recorded this data in an indexed address book, or you may have stored it on a diskette, using a personal computer and software such as DBASE IV or V, Microsoft ACCESS, or EXCEL. This is a collection of related data with an implicit meaning and hence is a database.

The preceding definition of *database* is quite general; for example, we may consider the collection of words that make up this page of text to be related data and hence to constitute a database. However, the common use of the term *database* is usually more restricted. A database has the following implicit properties:

- A database represents some aspect of the real world, sometimes called the **miniworld** or the **Universe of Discourse (UoD)**. Changes to the miniworld are reflected in the database.
- A database is a logically coherent collection of data with some inherent meaning. A random assortment of data cannot correctly be referred to as a database.
- A database is designed, built, and populated with data for a specific purpose. It has an intended group of users and some preconceived applications in which these users are interested.

In other words, a database has some source from which data are derived, some degree of interaction with events in the real world, and an audience that is actively interested in the contents of the database.

A database can be of any size and of varying complexity. For example, the list of names and addresses referred to earlier may consist of only a few hundred records, each with a simple structure. On the other hand, the card catalog of a large library may contain half a million cards stored under different categories—by primary author's last name, by subject, by book title—with each category organized in alphabetic order. A database of even greater size and complexity is maintained by the Internal Revenue Service to keep track of the tax forms filed by U.S. taxpayers. If we assume that there are 100 million tax-payers and if each taxpayer files an average of five forms with approximately 200 characters of information per form, we would get a database of $100 \times (10^8) \times 200 \times 5$ characters (bytes) of information. If the IRS keeps the past three returns for each taxpayer in addition to the current return, we would get a database of $4 \times (10^8) \times 200 \times 5$ bytes (400 gigabytes). This huge amount of information must be organized and managed so that users can search for, retrieve, and update the data as needed.

A database may be generated and maintained manually or it may be computerized. The library card catalog is an example of a database that may be created and maintained manually. A computerized database may be created and maintained either by a group of application programs written specifically for that task or by a database management system.

A **database management system (DBMS)** is a collection of programs that enables users to create and maintain a database. The DBMS is hence a *general-purpose software system* that facilitates the processes of *defining*, *constructing*, and *manipulating* databases for various applications. **Defining** a database involves specifying the data types, structures, and constraints for the data to be stored in the database. **Constructing** the database is the process of storing the data itself on some storage medium that is controlled by the DBMS. **Manipulating** a database includes such functions as querying the database to retrieve specific data, updating the database to reflect changes in the miniworld, and generating reports from the data.

It is not necessary to use general-purpose DBMS software to implement a computerized database. We could write our own set of programs to create and maintain the database, in effect creating our own *special-purpose* DBMS software. In either case—whether we use a general-purpose DBMS or not—we usually have to employ a considerable amount of software to manipulate the database. We will call the database and DBMS software together a **database system**. Figure 01.01 illustrates these ideas.

1.2 An Example

Let us consider an example that most readers may be familiar with: a UNIVERSITY database for maintaining information concerning students, courses, and grades in a university environment. Figure 01.02 shows the database structure and a few sample data for such a database. The database is organized as five files, each of which stores data records of the same type (Note 2). The STUDENT file stores data on each student; the COURSE file stores data on each course; the SECTION file stores data on each section of a course; the GRADE_REPORT file stores the grades that students receive in the various sections they have completed; and the PREREQUISITE file stores the prerequisites of each course.

To *define* this database, we must specify the structure of the records of each file by specifying the different types of **data elements** to be stored in each record. In Figure 01.02, each STUDENT record includes data to represent the student's Name, StudentNumber, Class (freshman or 1, sophomore or 2, . . .), and Major (MATH, computer science or CS, . . .); each COURSE record includes data to represent the CourseName, CourseNumber, CreditHours, and Department (the department that offers the course); and so on. We must also specify a **data type** for each data element within a record. For example, we can specify that Name of STUDENT is a string of alphabetic characters, StudentNumber of STUDENT is an integer, and Grade of GRADE_REPORT is a single character from the set {A, B, C, D, F, I}. We may also use a coding scheme to represent a data item. For example, in Figure 01.02 we represent the Class of a STUDENT as 1 for freshman, 2 for sophomore, 3 for junior, 4 for senior, and 5 for graduate student.

Databases store data to represent each student, course, section, grade report, and prerequisite as a record in the appropriate file. Notice that records in the various files may

be related. For example, the record for "Smith" in the `STUDE` file is related to two records in the `GRADE_REPORT` file that specify Smith's grades in two sections. Similarly, each record in the

`PREREQUISITE` file relates two course records: one representing the course and the other representing the prerequisite. Most medium-size and large databases include many types of records and have

many relationships among the records.

Database *manipulation* involves querying and updating. Examples of queries are "retrieve the transcript—a list of all courses and grades—of Smith"; "list the names of students who took the section of the Database course offered in fall 1999 and their grades in that section"; and "what are the prerequisites of the Database course?" Examples of updates are "change the class of Smith to Sophomore"; "create a new section for the Database course for this semester"; and "enter a grade of A for Smith in the Database section of last semester." These informal queries and updates must be specified precisely in the database system language before they can be processed.

1.3 Characteristics of the Database Approach

- 1.3.1 Self-Describing Nature of a Database System
- 1.3.2 Insulation between Programs and Data, and Data Abstraction
- 1.3.3 Support of Multiple Views of the Data
- 1.3.4 Sharing of Data and Multiuser Transaction Processing

A number of characteristics distinguish the database approach from the traditional approach of programming with files. In traditional **file processing**, each user defines and implements the files needed for a specific application as part of programming the application. For example, one user, the *grade reporting office*, may keep a file on students and their grades. Programs to print a student's transcript and to enter new grades into the file are implemented. A second user, the accounting office, may keep track of students' fees and their payments. Although both users are interested in data about students, each user maintains separate files—and programs to manipulate these files—because each requires some data not available from the other user's files. This redundancy in defining and storing data results in wasted storage space and in redundant efforts to maintain common data up-to-date.

In the database approach, a single repository of data is maintained that is defined once and then is accessed by various users. The main characteristics of the database approach versus the file-processing approach are the following.

1.3.1 Self-Describing Nature of a Database System

A fundamental characteristic of the database approach is that the database system contains not only the database itself but also a complete definition or description of the database structure and constraints. This definition is stored in the system **catalog**, which contains information such as the structure of each file, the type and storage format of each data item, and various constraints on the data. The information stored in the catalog is called **meta-data**, and it describes the structure of the primary database (Figure 01.01).

The catalog is used by the DBMS software and also by database users who need information about the database structure. A general purpose DBMS software package is not written for a specific database application, and hence it must refer to the catalog to know the structure of the files in a specific database, such as the type and format of data it will access. The DBMS software must work equally well with *any number of database applications*—for example, a university database, a banking database, or a company database—as long as the database definition is stored in the catalog.

In traditional file processing, data definition is typically part of the application programs themselves. Hence, these programs are constrained to work with only *one specific database*, whose structure is declared in the application programs. For example, a PASCAL program may have record structures declared in it; a C++ program may have "struct" or "class" declarations; and a COBOL program has Data Division statements to define its files. Whereas file-processing software can access only specific databases, DBMS software can access diverse databases by extracting the database definitions from the catalog and then using these definitions.

In the example shown in Figure 01.02, the DBMS stores in the catalog the definitions of all the files shown. Whenever a request is made to access, say, the Name of a STUDENT record, the DBMS software refers to the catalog to determine the structure of the STUDENT file and the position and size of the Name data item within a STUDENT record. By contrast, in a typical file-processing application, the file structure and, in the extreme case, the exact location of Name within a STUDENT record are already coded within each program that accesses this data item.

1.3.2 Insulation between Programs and Data, and Data Abstraction

In traditional file processing, the structure of data files is embedded in the access programs, so any changes to the structure of a file may require *changing all programs* that access this file. By contrast, DBMS access programs do not require such changes in most cases. The structure of data files is stored in the DBMS catalog separately from the access programs. We call this property **program-data independence**. For example, a file access program may be written in such a way that it can access only STUDENT records of the structure shown in Figure 01.03. If we want to add another piece of data to each STUDENT record, say the Birthdate, such a program will no longer work and must be changed. By contrast, in a DBMS environment, we just need to change the description of STUDENT records in the catalog to reflect the inclusion of the new data item Birthdate; no programs are changed. The next time a DBMS program refers to the catalog, the new structure of STUDENT records will be accessed and used.

In object-oriented and object-relational databases (see Part III), users can define operations on data as part of the database definitions. An **operation** (also called a *function*) is specified in two parts. The *interface* (or *signature*) of an operation includes the operation name and the data types of its arguments (or parameters). The *implementation* (or *method*) of the operation is specified separately and can be changed without affecting the interface. User application programs can operate on the data by invoking these operations through their names and arguments, regardless of how the operations are implemented. This may be termed **program-operation independence**.

The characteristic that allows program-data independence and program-operation independence is called **data abstraction**. A DBMS provides users with a **conceptual representation** of data that does not include many of the details of how the data is stored or how the operations are implemented. Informally, a **data model** is a type of data abstraction that is used to provide this conceptual representation. The data model uses logical concepts, such as objects, their properties, and their interrelationships, that may be easier for most users to understand than computer storage concepts. Hence, the data model *hides* storage and implementation details that are not of interest to most database users.

For example, consider again Figure 01.02. The internal implementation of a file may be defined by its record length—the number of characters (bytes) in each record—and each data item may be specified

by its starting byte within a record and its length in bytes. The STUDENT

record would thus be represented as shown in Figure 01.03. But a typical database user is not concerned with the location of each data item within a record or its length; rather the concern is that, when a reference is made to Name of STUDENT, the correct value is returned. A conceptual representation of the STUDENT records is shown in Figure 01.02.

Many other details of file-storage organization—such as the access paths specified on a file—can be hidden from database users by the DBMS;

we will discuss storage details in Chapter 5 and Chapter 6.

In the database approach, the detailed structure and organization of each file are stored in the catalog. Database users refer to the conceptual representation of the files, and the DBMS extracts the details of file storage from the catalog when these are needed by the DBMS software. Many data models can be used to provide this data abstraction to database users. A major part of this book is devoted to presenting various data models and the concepts they use to abstract the representation of data.

With the recent trend toward object-oriented and object-relational databases, abstraction is carried one level further to include not only the data structure but also the operations on the data. These operations provide an abstraction of miniworld activities commonly understood by the users. For example, an operation CALCULATE_GPA can be applied to a student object to calculate the grade point average. Such operations can be invoked by the user queries or programs without the user knowing the details of how they are internally implemented. In that sense, an abstraction of the miniworld activity is made available to the user as an **abstract operation**.

1.3.3 Support of Multiple Views of the Data

A database typically has many users, each of whom may require a different perspective or **view** of the database. A view may be a subset of the database or it may contain **virtual data** that is derived from the database files but is not explicitly stored. Some users may not need to be aware of whether the data they refer to is stored or derived. A multiuser DBMS whose users have a variety of applications must provide facilities for defining multiple views. For example, one user of the database of Figure 01.02 may be interested only in the transcript of each student; the view for this user is shown in Figure 01.04(a). A second user, who is interested only in checking that students have taken all the prerequisites of each course they register for, may require the view shown in Figure 01.04(b).

1.3.4 Sharing of Data and Multiuser Transaction Processing

A multiuser DBMS, as its name implies, must allow multiple users to access the database at the same time. This is essential if data for multiple applications is to be integrated and maintained in a single database. The DBMS must include **concurrency control** software to ensure that several users trying to update the same data do so in a controlled manner so that the result of the updates is correct. For example, when several reservation clerks try to assign a seat on an airline flight, the DBMS should ensure that each seat can be accessed by only one clerk at a time for assignment to a passenger. These types of applications are generally called **on-line transaction processing (OLTP)** applications. A fundamental role of multiuser DBMS software is to ensure that concurrent transactions operate correctly.

The preceding characteristics are most important in distinguishing a DBMS from traditional file-processing software. In Section 1.6 we discuss additional functions that characterize a DBMS. First, however, we categorize the different types of persons who work in a database environment

DATABASE SYSTEM CONCEPTS

SIXTH EDITION

Abraham Silberschatz

Yale University

Henry F. Korth

Lehigh University

S. Sudarshan

Indian Institute of Technology, Bombay

Published by McGraw-Hill, a business unit of The McGraw-Hill Companies, Inc., 1221 Avenue of the Americas, New York, NY 10020. Copyright © 2011 by The McGraw-Hill Companies, Inc. All rights reserved. Previous editions © 2006, 2002, and 1999. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of The McGraw-Hill Companies, Inc., including, but not limited to, in any network or other electronic storage or transmission, or broadcast for distance learning.

4 Chapter 1 Introduction

- **Data redundancy and inconsistency.** Since different programmers create the files and application programs over a long period, the various files are likely to have different structures and the programs may be written in several programming languages. Moreover, the same information may be duplicated in several places (files). For example, if a student has a double major (say, music and mathematics) the address and telephone number of that student may appear in a file that consists of student records of students in the Music department and in a file that consists of student records of students in the Mathematics department. This redundancy leads to higher storage and access cost. In addition, it may lead to **data inconsistency**; that is, the various copies of the same data may no longer agree. For example, a changed student address may be reflected in the Music department records but not elsewhere in the system.

- **Difficulty in accessing data.** Suppose that one of the university clerks needs to find out the names of all students who live within a particular postal-code area. The clerk asks the data-processing department to generate such a list. Because the designers of the original system did not anticipate this request, there is no application program on hand to meet it. There is, however, an application program to generate the list of *all* students. The university clerk has now two choices: either obtain the list of all students and extract the needed information manually or ask a programmer to write the necessary application program. Both alternatives are obviously unsatisfactory. Suppose that such a program is written, and that, several days later, the same clerk needs to trim that list to include only those students who have taken at least 60 credit hours. As expected, a program to generate such a list does not exist. Again, the clerk has the preceding two options, neither of which is satisfactory.

The point here is that conventional file-processing environments do not allow needed data to be retrieved in a convenient and efficient manner. More responsive data-retrieval systems are required for general use.

- **Data isolation.** Because data are scattered in various files, and files may be in different formats, writing new application programs to retrieve the appropriate data is difficult.

- **Integrity problems.** The data values stored in the database must satisfy certain types of **consistency constraints**. Suppose the university maintains an account for each department, and records the balance amount in each account. Suppose also that the university requires that the account balance of a department may never fall below zero. Developers enforce these constraints in the system by adding appropriate code in the various application programs. However, when new constraints are added, it is difficult to change the programs to enforce them. The problem is compounded when constraints involve several data items from different files.

- **Atomicity problems.** A computer system, like any other device, is subject to failure. In many applications, it is crucial that, if a failure occurs, the data

be restored to the consistent state that existed prior to the failure. Consider a program to transfer \$500 from the account balance of department *A* to the account balance of department *B*. If a system failure occurs during the execution of the program, it is possible that the \$500 was removed from the balance of department *A* but was not credited to the balance of department *B*, resulting in an inconsistent database state. Clearly, it is essential to database consistency that either both the credit and debit occur, or that neither occur. That is, the funds transfer must be *atomic*—it must happen in its entirety or not at all. It is difficult to ensure atomicity in a conventional file-processing system.

- **Concurrent-access anomalies.** For the sake of overall performance of the system and faster response, many systems allow multiple users to update the data simultaneously. Indeed, today, the largest Internet retailers may have millions of accesses per day to their data by shoppers. In such an environment, interaction of concurrent updates is possible and may result in inconsistent data. Consider department *A*, with an account balance of \$10,000. If two department clerks debit the account balance (by say \$500 and \$100, respectively) of department *A* at almost exactly the same time, the result of the concurrent executions may leave the budget in an incorrect (or inconsistent) state. Suppose that the programs executing on behalf of each withdrawal read the old balance, reduce that value by the amount being withdrawn, and write the result back. If the two programs run concurrently, they may both read the value \$10,000, and write back \$9500 and \$9900, respectively. Depending on which one writes the value last, the account balance of department *A* may contain either \$9500 or \$9900, rather than the correct value of \$9400. To guard against this possibility, the system must maintain some form of supervision. But supervision is difficult to provide because data may be accessed by many different application programs that have not been coordinated previously. As another example, suppose a registration program maintains a count of students registered for a course, in order to enforce limits on the number of students registered. When a student registers, the program reads the current count for the courses, verifies that the count is not already at the limit, adds one to the count, and stores the count back in the database. Suppose two students register concurrently, with the count at (say) 39. The two program executions may both read the value 39, and both would then write back 40, leading to an incorrect increase of only 1, even though two students successfully registered for the course and the count should be 41. Furthermore, suppose the course registration limit was 40; in the above case both students would be able to register, leading to a violation of the limit of 40 students.

- **Security problems.** Not every user of the database system should be able to access all the data. For example, in a university, payroll personnel need to see only that part of the database that has financial information. They do not need access to information about academic records. But, since application programs are added to the file-processing system in an ad hoc manner, enforcing such security constraints is difficult.

6

These difficulties, among others, prompted the development of database systems. In what follows, we shall see the concepts and algorithms that enable database systems to solve the problems with file-processing systems. In most of this book, we use a university organization as a running example of a typical data-processing application.

1.3 View of Data

A database system is a collection of interrelated data and a set of programs that allow users to access and modify these data. A major purpose of a database system is to provide users with an *abstract* view of the data. That is, the system hides certain details of how the data are stored and maintained.

1.3.1 Data Abstraction

For the system to be usable, it must retrieve data efficiently. The need for efficiency has led designers to use complex data structures to represent data in the database. Since many database-system users are not computer trained, developers hide the complexity from users through several levels of abstraction, to simplify users' interactions with the system:

- **Physical level.** The lowest level of abstraction describes *how* the data are actually stored. The physical level describes complex low-level data structures in detail.
- **Logical level.** The next-higher level of abstraction describes *what* data are stored in the database, and what relationships exist among those data. The logical level thus describes the entire database in terms of a small number of relatively simple structures. Although implementation of the simple structures at the logical level may involve complex physical-level structures, the user of the logical level does not need to be aware of this complexity. This is referred to as **physical data independence**. Database administrators, who must decide what information to keep in the database, use the logical level of abstraction.
- **View level.** The highest level of abstraction describes only part of the entire database. Even though the logical level uses simpler structures, complexity remains because of the variety of information stored in a large database. Many users of the database system do not need all this information; instead, they need to access only a part of the database. The view level of abstraction exists to simplify their interaction with the system. The system may provide many views for the same database.

Figure 1.1 shows the relationship among the three levels of abstraction. An analogy to the concept of data types in programming languages may clarify the distinction among levels of abstraction. Many high-level programming

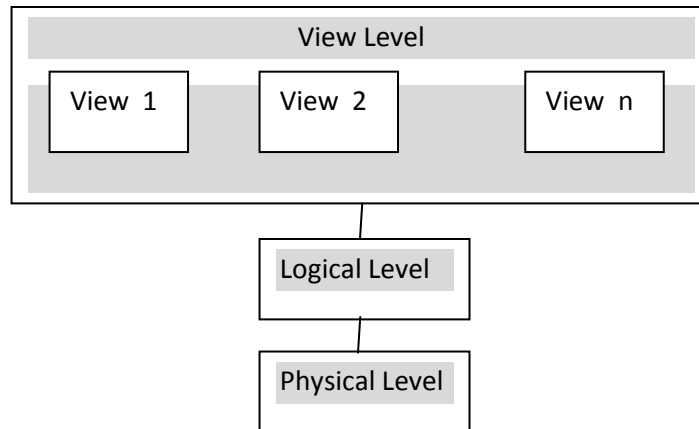


Figure 1.1 The three levels of data abstraction.

languages support the notion of a structured type. For example, we may describe a record as follows:

```

type instructor = record
ID : char (5);
name : char (20);
dept name : char (20);
salary : numeric (8,2);
end;
  
```

This code defines a new record type called *instructor* with four fields. Each field has a name and a type associated with it. A university organization may have several such record types, including

- *department*, with fields *dept name*, *building*, and *budget*
- *course*, with fields *course id*, *title*, *dept name*, and *credits*
- *student*, with fields *ID*, *name*, *dept name*, and *tot cred*

At the physical level, an *instructor*, *department*, or *student* record can be described as a block of consecutive storage locations. The compiler hides this level of detail from programmers. Similarly, the database system hides many of the lowest-level storage details from database programmers. Database administrators, on the other hand, may be aware of certain details of the physical organization of the data.

The actual type declaration depends on the language being used. C and C++ use **struct** declarations. Java does not have such a declaration, but a simple class can be defined to the same effect.

8

At the logical level, each such record is described by a type definition, as in the previous code segment, and the interrelationship of these record types is defined as well. Programmers using a programming language work at this level of abstraction. Similarly, database administrators usually work at this level of abstraction.

Finally, at the view level, computer users see a set of application programs that hide details of the data types. At the view level, several views of the database are defined, and a database user sees some or all of these views. In addition to hiding details of the logical level of the database, the views also provide a security mechanism to prevent users from accessing certain parts of the database. For example, clerks in the university registrar office can see only that part of the database that has information about students; they cannot access information about salaries of instructors.

1.3.2 Instances and Schemas

Databases change over time as information is inserted and deleted. The collection of information stored in the database at a particular moment is called an **instance** of the database. The overall design of the database is called the database **schema**. Schemas are changed infrequently, if at all.

The concept of database schemas and instances can be understood by analogy to a program written in a programming language. A database schema corresponds to the variable declarations (along with associated type definitions) in a program. Each variable has a particular value at a given instant. The values of the variables in a program at a point in time correspond to an *instance* of a database schema. Database systems have several schemas, partitioned according to the levels of abstraction. The **physical schema** describes the database design at the physical level, while the **logical schema** describes the database design at the logical level. A database may also have several schemas at the view level, sometimes called **subschemas**, that describe different views of the database.

Of these, the logical schema is by far the most important, in terms of its effect on application programs, since programmers construct applications by using the logical schema. The physical schema is hidden beneath the logical schema, and can usually be changed easily without affecting application programs. Application programs are said to exhibit **physical data independence** if they do not depend on the physical schema, and thus need not be rewritten if the physical schema changes.

We study languages for describing schemas after introducing the notion of data models in the next section.

1.3.3 Data Models

Underlying the structure of a database is the **data model**: a collection of conceptual tools for describing data, data relationships, data semantics, and consistency constraints. A data model provides a way to describe the design of a database at the physical, logical, and view levels.

There are a number of different data models that we shall cover in the text. The data models can be classified into four different categories:

• **Relational Model.** The relational model uses a collection of tables to represent both data and the relationships among those data. Each table has multiple columns, and each column has a unique name. Tables are also known as **relations**. The relational model is an example of a record-based model. Record-based models are so named because the database is structured in fixed-format records of several types. Each table contains records of a particular type. Each record type defines a fixed number of fields, or attributes. The columns of the table correspond to the attributes of the record type. The relational data model is the most widely used data model, and a vast majority of current database systems are based on the relational model.

Chapters 2 through 8 cover the relational model in detail.

• **Entity-Relationship Model.** The entity-relationship (E-R) data model uses a collection of basic objects, called *entities*, and *relationships* among these objects. An entity is a “thing” or “object” in the real world that is distinguishable from other objects. The entity-relationship model is widely used in database design, and Chapter 7 explores it in detail.

• **Object-Based Data Model.** Object-oriented programming (especially in Java, C++, or C#) has become the dominant software-development methodology. This led to the development of an object-oriented data model that can be seen as extending the E-R model with notions of encapsulation, methods (functions), and object identity. The object-relational data model combines features of the object-oriented data model and relational data model. Chapter 22 examines the object-relational data model.

• **Semistructured Data Model.** The semistructured data model permits the specification of data where individual data items of the same type may have different sets of attributes. This is in contrast to the data models mentioned earlier, where every data item of a particular type must have the same set of attributes. The **Extensible Markup Language (XML)** is widely used to represent semistructured data. Chapter 23 covers it.

Historically, the **network data model** and the **hierarchical data model** preceded the relational data model. These models were tied closely to the underlying implementation, and complicated the task of modeling data. As a result they are used little now, except in old database code that is still in service in some places. They are outlined online in Appendices D and E for interested readers.

1.4 Database Languages

A database system provides a **data-definition language** to specify the database schema and a **data-manipulation language** to express database queries and up

information is updated without taking precautions to update all copies of the information. For example, different offerings of a course may have the same course identifier, but may have different titles. It would then become unclear what the correct title of the course is. Ideally, information should appear in exactly one place.

2. Incompleteness: A bad design may make certain aspects of the enterprise difficult or impossible to model. For example, suppose that, as in case (1) above, we only had entities corresponding to course offering, without having an entity corresponding to courses. Equivalently, in terms of relations, suppose we have a single relation where we repeat all of the course information once for each section that the course is offered. It would then be impossible to represent information about a new course, unless that course is offered. We might try to make do with the problematic design by storing null values for the section information. Such a work-around is not only unattractive, but may be prevented by primary-key constraints.

Avoiding bad designs is not enough. There may be a large number of good designs from which we must choose. As a simple example, consider a customer who buys a product. Is the sale of this product a relationship between the customer and the product? Alternatively, is the sale itself an entity that is related both to the customer and to the product? This choice, though simple, may make an important difference in what aspects of the enterprise can be modeled well. Considering the need to make choices such as this for the large number of entities and relationships in a real-world enterprise, it is not hard to see that database design can be a challenging problem. Indeed we shall see that it requires a combination of both science and “good taste.”

7.2 The Entity-Relationship Model

The entity-relationship (E-R) data model was developed to facilitate database design by allowing specification of an *enterprise schema* that represents the overall logical structure of a database.

The E-R model is very useful in mapping the meanings and interactions of real-world enterprises onto a conceptual schema. Because of this usefulness, many database-design tools draw on concepts from the E-R model. The E-R data model employs three basic concepts: entity sets, relationship sets, and attributes, which we study first. The E-R model also has an associated diagrammatic representation, the E-R diagram, which we study later in this chapter.

7.2.1 Entity Sets

An entity is a “thing” or “object” in the real world that is distinguishable from all other objects. For example, each person in a university is an entity. An entity has a set of properties, and the values for some set of properties may uniquely identify an entity. For instance, a person may have a *person id* property whose

value uniquely identifies that person. Thus, the value 677-89-9011 for *person id* would uniquely identify one particular person in the university. Similarly, courses can be thought of as entities, and *course id* uniquely identifies a course entity in the university. An entity may be concrete, such as a person or a book, or it may be abstract, such as a course, a course offering, or a flight reservation.

An entity set is a set of entities of the same type that share the same properties, or attributes. The set of all people who are instructors at a given university, for example, can be defined as the entity set *instructor*. Similarly, the entity set *student* might represent the set of all students in the university.

In the process of modeling, we often use the term *entity set* in the abstract, without referring to a particular set of individual entities. We use the term extension of the entity set to refer to the actual collection of entities belonging to the entity set. Thus, the set of actual instructors in the university forms the extension of the entity set *instructor*. The above distinction is similar to the difference between a relation and a relation instance, which we saw in Chapter 2.

Entity sets do not need to be disjoint. For example, it is possible to define the entity set of all people in a university (*person*). A *person* entity may be an *instructor* entity, a *student* entity, both, or neither.

An entity is represented by a set of attributes. Attributes are descriptive properties possessed by each member of an entity set. The designation of an attribute for an entity set expresses that the database stores similar information concerning each entity in the entity set; however, each entity may have its own value for each attribute. Possible attributes of the *instructor* entity set are *ID*, *name*, *dept name*, and *salary*. In real life, there would be further attributes, such as street number, apartment number, state, postal code, and country, but we omit them to keep our examples simple. Possible attributes of the *course* entity set are *course id*, *title*, *dept name*, and *credits*.

Each entity has a value for each of its attributes. For instance, a particular *instructor* entity may have the value 12121 for *ID*, the value Wu for *name*, the value Finance for *dept name*, and the value 90000 for *salary*.

The *ID* attribute is used to identify instructors uniquely, since there may be more than one instructor with the same name. In the United States, many enterprises find it convenient to use the *social-security* number of a person² as an attribute whose value uniquely identifies the person. In general the enterprise would have to create and assign a unique identifier for each instructor.

A database thus includes a collection of entity sets, each of which contains any number of entities of the same type. Figure 7.1 shows part of a university database that consists of two entity sets: *instructor* and *student*. To keep the figure simple, only some of the attributes of the two entity sets are shown.

A database for a university may include a number of other entity sets. For example, in addition to keeping track of instructors and students, the university also has information about courses, which are represented by the entity set *course*

²In the United States, the government assigns to each person in the country a unique number, called a social-security number, to identify that person uniquely. Each person is supposed to have only one social-security number, and no two people are supposed to have the same social-security number.

Instructor

student

Figure 7.1 Entity sets *instructor* and *student*.

with attributes *course id*, *title*, *dept name* and *credits*. In a real setting, a university database may keep dozens of entity sets.

7.2.2 Relationship Sets

A relationship is an association among several entities.

For example, we can define a relationship *advisor* that associates instructor Katz with student Shankar. This relationship specifies that Katz is an advisor to student Shankar.

A relationship set is a set of relationships of the same type.

Formally, it is a mathematical relation on $n \geq 2$ (possibly nondistinct) entity sets. If E_1, E_2, \dots, E_n are entity sets, then a relationship set R is a subset of $\{(e_1, e_2, \dots, e_n) \mid e_1 \in E_1, e_2 \in E_2, \dots, e_n \in E_n\}$ where (e_1, e_2, \dots, e_n) is a relationship.

Consider the two entity sets *instructor* and *student* in Figure 7.1. We define the relationship set *advisor* to denote the association between instructors and students. Figure 7.2 depicts this association.

As another example, consider the two entity sets *student* and *section*. We can define the relationship set *takes* to denote the association between a student and the course sections in which that student is enrolled.

The association between entity sets is referred to as participation; that is, the entity sets E_1, E_2, \dots, E_n participate in relationship set R . A relationship instance in an E-R schema represents an association between the named entities in the real-world enterprise that is being modeled. As an illustration, the individual *instructor* entity Katz, who has instructor *ID* 45565, and the *student* entity Shankar, who has student *ID* 12345, participate in a relationship instance of *advisor*. This relationship instance represents that in the university, the instructor Katz is advising student Shankar.

The function that an entity plays in a relationship is called that entity's role. Since entity sets participating in a relationship set are generally distinct, roles

ER Model Relationships in general

Relationship: Interaction between entities

Indicator : an attribute of one entity refers to another entity (Represent such references as

relationships not attributes)

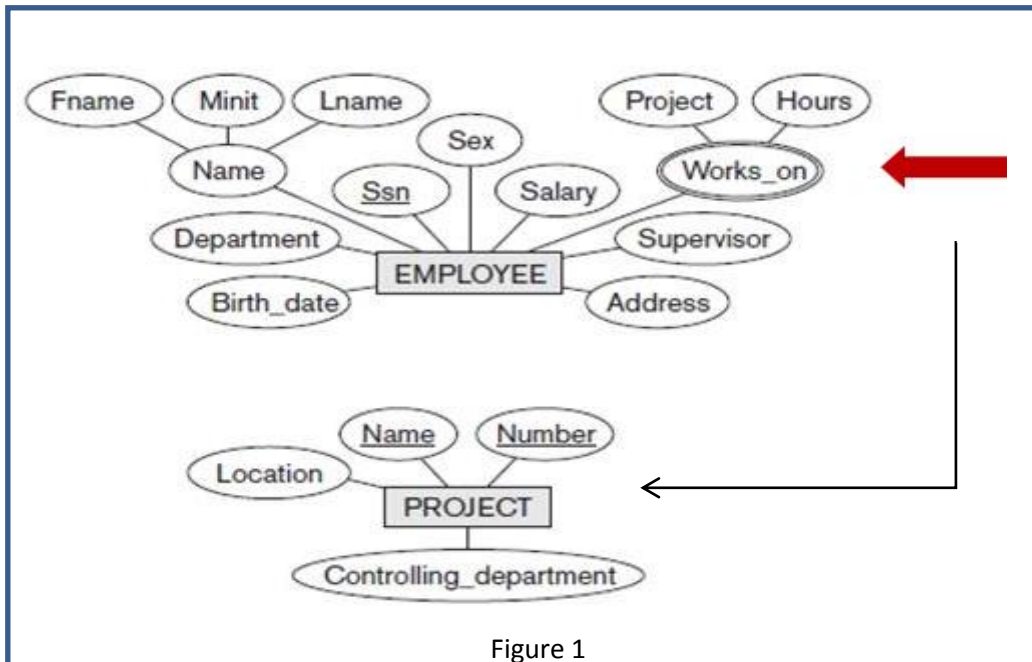


Figure 1

Diagramming Relationship : Diamond for relationship type , connected to each participating entity type

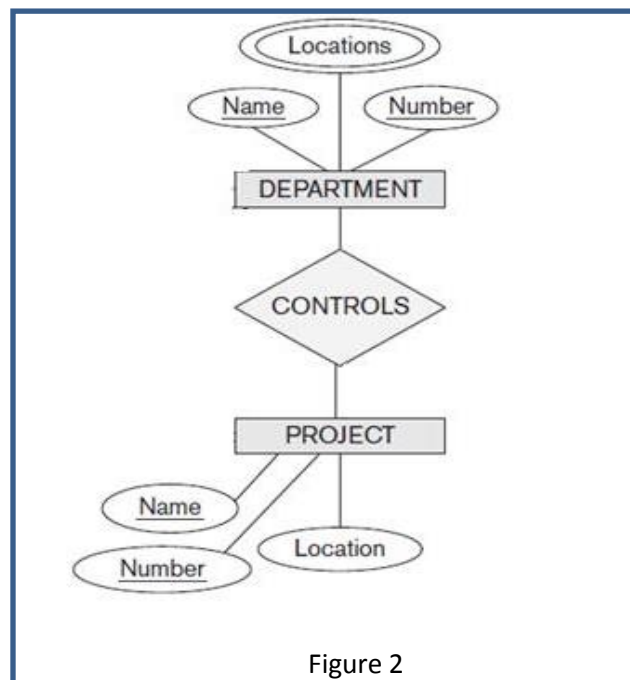
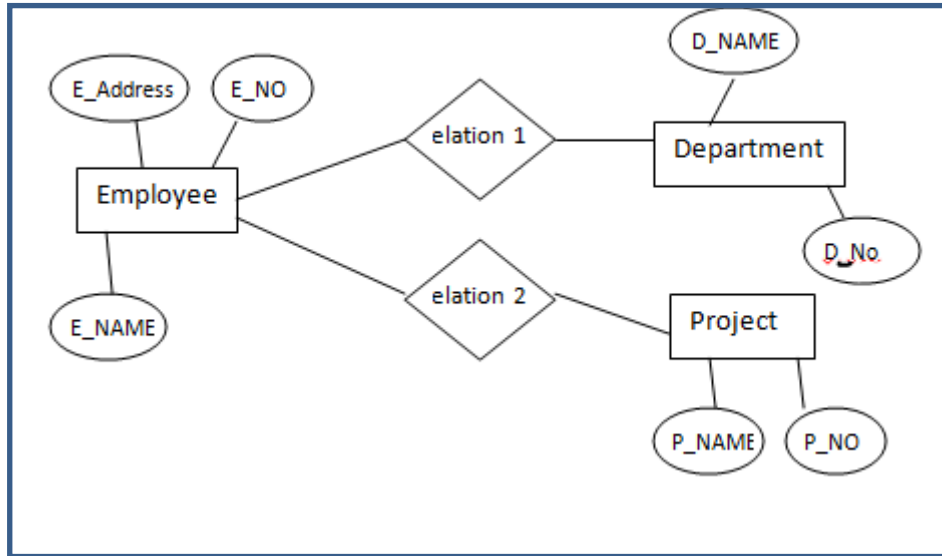


Figure 2

Example 1 : In a company, there are three entity :

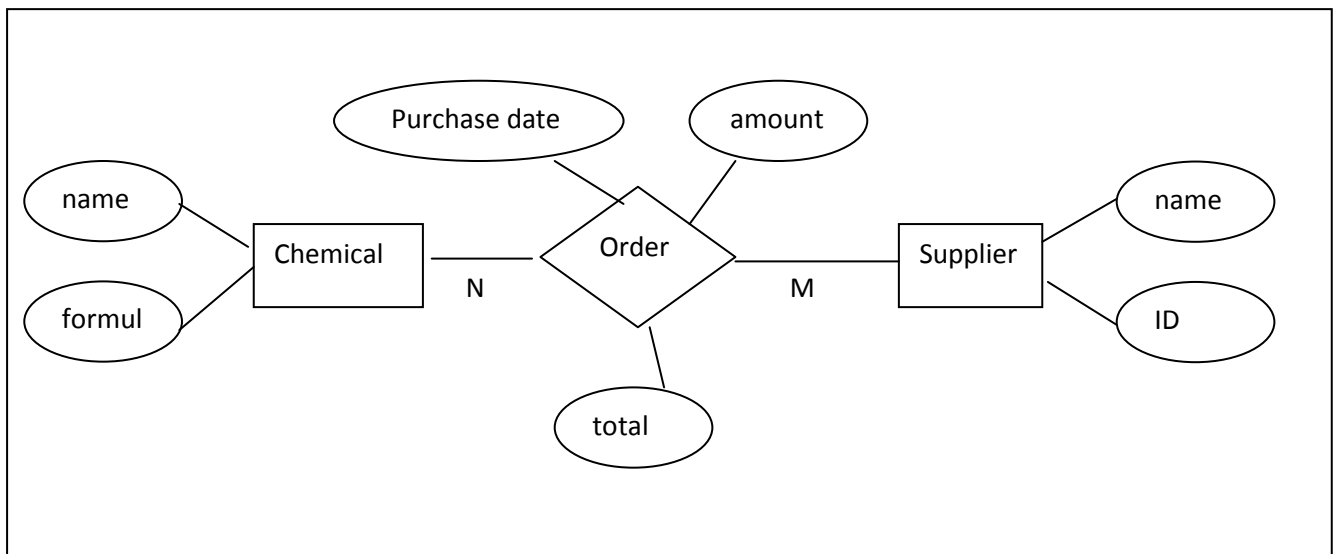
- 1- Employees : E_NO ,E_NAME, E_ADDRESS
- 2- Departments : D_NO ,D_NAME
- 3-Projects : P_NAME , P_NO

Each employee belongs to a department and each project has many employees works on it. There are two relations that connect the three entities



Example 2 : A chemical factory producing chemical materials , each material identified by a name and a formula . The supplier , identified by his name and his ID , purchase from the factory by an order . The order has date , amount and total.

To draw the ER model
 Each supplier can have any material , and any material can go to any supplier .
 The relation is of type **many-to-many**.



Cardinality

<http://ion.uwinnipeg.ca/~rmcfadye/2914/relationships.pdf>(ACS 2914 ERD: Relationships Ron McFadyen)

Cardinality is a constraint on a relationship specifying the number of entity instances that a specific entity may be related to via the relationship. Consider the relationship "works in".



When we ask How many employees can work in a single department? or How many departments can an employee work in? we are asking questions regarding the cardinality of the relationship. The three classifications are: one-to-one, one-to-many, and many-to-many. Below, the ERD shows a relationship between invoice lines and products.



The "n" represents an "arbitrary number of instances", and the "1" represents "at most one instance". We interpret the cardinality specifications with the following business rule statements:

- The "n" indicates that the same Product entity can be specified on "any number of" Invoice Lines.
- The "1" indicates that an Invoice Line entity specifies "at most one" Product entity

Exercise: Consider the University and entity types: Class, Course, Department, Student. What relationship types exist and what cardinalities apply to the various relationships and entities?

One-to-One Relationships

One-to-one relationships have 1 specified for both cardinalities, and do not seem to arise very often. To illustrate a one-to-one, we require very specific business rules.

Suppose we have People and Vehicles. Assume that we are only concerned with the current driver of a vehicle, and that we are only concerned with the current vehicle that a driver is operating. Then, we have a one-to-one relationship between Vehicle and Person (note the role shown for Person in this relationship):



One-to-Many Relationships

This type of relationship has 1 and n specified for cardinalities, and is very common in database designs. Suppose we have customers and orders and the business rules:

- An order is related to one customer, and
- A customer can have any number (zero or more) of orders.

We say there is a one-to-many relationship between customer and order, and we draw this as:



Many-to-Many Relationships

Many-to-many relationships have "many" specified for both cardinalities, and are also very common. However, should you examine a data model in some business, there is a good chance you will not see any many-to-many relationships on the diagram. In those cases, the data modeler has resolved the many-to-many relationships into two one-to-many relationships. Suppose we are interested in courses and students and the fact that students register for courses: Any student may take several courses, A course may be taken by several students. This situation is represented with a many-to-many relationship between Course and Student:



ER Model to Relational Model

https://www.tutorialspoint.com/dbms/er_model_to_relational_model.htm

ER Model, when conceptualized into diagrams, gives a good overview of entity-relationship, which is easier to understand. ER diagrams can be mapped to relational schema, that is, it is possible to create relational schema using ER diagram. We cannot import all the ER constraints into relational model, but an approximate schema can be generated.

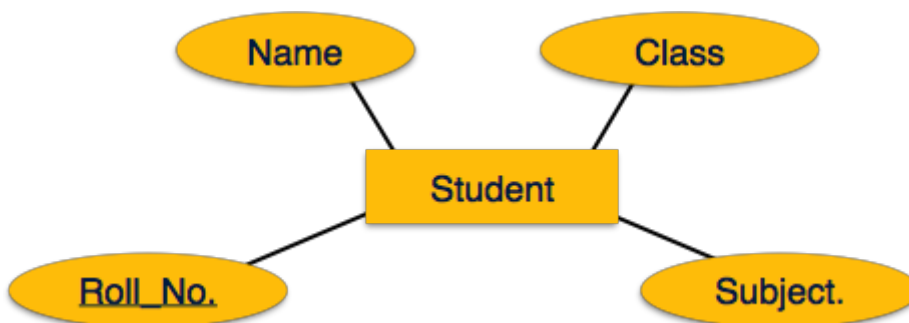
There are several processes and algorithms available to convert ER Diagrams into Relational Schema. Some of them are automated and some of them are manual. We may focus here on the mapping diagram contents to relational basics.

ER diagrams mainly comprise of –

- Entity and its attributes
- Relationship, which is association among entities.

Mapping Entity

An entity is a real-world object with some attributes.

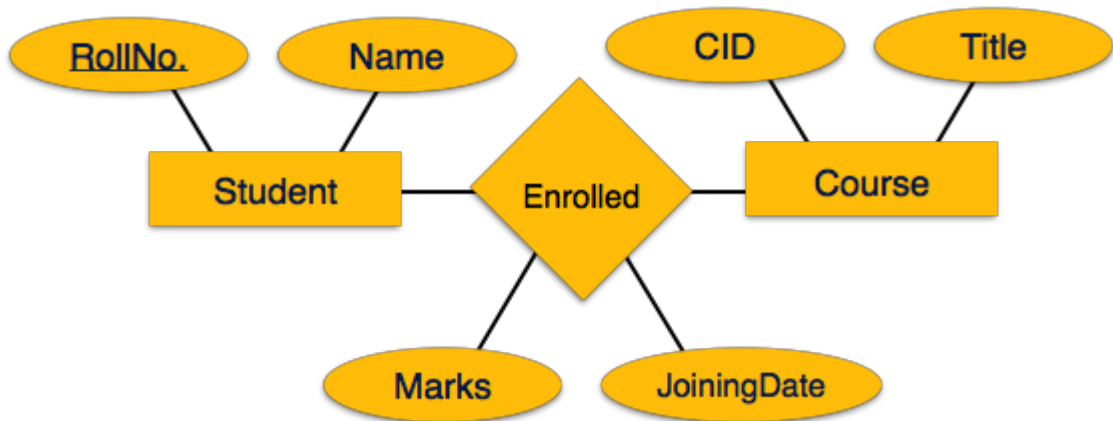


Mapping Process

- Create table for each entity.
- Entity's attributes should become fields of tables with their respective data types.
- Declare primary key.

Mapping Relationship

A relationship is an association among entities.



Mapping Process

- Create table for a relationship.
- Add the primary keys of all participating Entities as fields of table with their respective data types.
- If relationship has any attribute, add each attribute as field of table.
- Declare a primary key composing all the primary keys of participating entities.
- Declare all foreign key constraints.

- *student* (*ID*, *name*, *dept name*, *tot cred*)
- *advisor* (*s id*, *i id*)
- *takes* (*ID*, *course id*, *sec id*, *semester*, *year*, *grade*)
- *classroom* (*building*, *room number*, *capacity*)
- *time slot* (*time slot id*, *day*, *start time*, *end time*)

2.3 Keys

We must have a way to specify how tuples within a given relation are distinguished. This is expressed in terms of their attributes. That is, the values of the attribute values of a tuple must be such that they can *uniquely identify* the tuple. In other words, no two tuples in a relation are allowed to have exactly the same value for all attributes.

A **superkey** is a set of one or more attributes that, taken collectively, allow us to identify uniquely a tuple in the relation. For example, the *ID* attribute of the relation *instructor* is sufficient to distinguish one instructor tuple from another.

Thus, *ID* is a superkey. The *name* attribute of *instructor*, on the other hand, is not a superkey, because several instructors might have the same name.

Formally, let R denote the set of attributes in the schema of relation r . If we say that a subset K of R is a *superkey* for r , we are restricting consideration to instances of relations r in which no two distinct tuples have the same values on all attributes in K . That is, if t_1 and t_2 are in r and $t_1 \neq t_2$, then $t_1.K \neq t_2.K$.

A superkey may contain extraneous attributes. For example, the combination of *ID* and *name* is a superkey for the relation *instructor*. If K is a superkey, then so is any superset of K . We are often interested in superkeys for which no proper subset is a superkey. Such minimal superkeys are called **candidate keys**.

It is possible that several distinct sets of attributes could serve as a candidate key. Suppose that a combination of *name* and *dept name* is sufficient to distinguish among members of the *instructor* relation. Then, both $\{ID\}$ and $\{name, dept name\}$ are candidate keys. Although the attributes *ID* and *name* together can distinguish *instructor* tuples, their combination, $\{ID, name\}$, does not form a candidate key, since the attribute *ID* alone is a candidate key.

We shall use the term **primary key** to denote a candidate key that is chosen by the database designer as the principal means of identifying tuples within a relation.

A key (whether primary, candidate, or super) is a property of the entire relation, rather than of the individual tuples.

Any two individual tuples in the relation are prohibited from having the same value on the key attributes at the same time.

The designation of a key represents a constraint in the real-world enterprise being modeled.

Primary keys must be chosen with care. As we noted, the name of a person is obviously not sufficient, because there may be many people with the same name. In the United States, the social-security number attribute of a person would be a candidate key. Since non-U.S. residents usually do not have social-security

46 Chapter 2 Introduction to the Relational Model

numbers, international enterprises must generate their own unique identifiers. An alternative is to use some unique combination of other attributes as a key. The primary key should be chosen such that its attribute values are never, or very rarely, changed. For instance, the address field of a person should not be part of the primary key, since it is likely to change. Social-security numbers, on the other hand, are guaranteed never to change. Unique identifiers generated by enterprises generally do not change, except if two enterprises merge; in such a case the same identifier may have been issued by both enterprises, and a reallocation of identifiers may be required to make sure they are unique. It is customary to list the primary key attributes of a relation schema before the other attributes; for example, the *dept name* attribute of *department* is listed first, since it is the primary key. Primary key attributes are also underlined.

A relation, say r_1 , may include among its attributes the primary key of another relation, say r_2 . This attribute is called a **foreign key** from r_1 , referencing r_2 .

The relation r_1 is also called the **referencing relation** of the foreign key dependency, and r_2 is called the **referenced relation** of the foreign key. For example, the attribute *dept name* in *instructor* is a foreign key from *instructor*, referencing *department*,

since *dept name* is the primary key of *department*. In any database instance, given any tuple, say ta , from the *instructor* relation, there must be some tuple, say tb , in the *department* relation such that the value of the *dept name* attribute of ta is the same as the value of the primary key, *dept name*, of tb .

Now consider the *section* and *teaches* relations. It would be reasonable to require that if a section exists for a course, it must be taught by at least one instructor; however, it could possibly be taught by more than one instructor.

To enforce this constraint, we would require that if a particular (*course id*, *sec id*, *semester*, *year*) combination appears in *section*, then the same combination must appear in *teaches*. However, this set of values does not form a primary key for *teaches*, since more than one instructor may teach one such section. As a result, we cannot declare a foreign key constraint from *section* to *teaches* (although we can define a foreign key constraint in the other direction, from *teaches* to *section*). The constraint from *section* to *teaches* is an example of a **referential integrity constraint**; a referential integrity constraint requires that the values appearing in specified attributes of any tuple in the referencing relation also appear in specified attributes of at least one tuple in the referenced relation.

2.4 Schema Diagrams

A database schema, along with primary key and foreign key dependencies, can be depicted by **schema diagrams**. Figure 2.8 shows the schema diagram for our university organization. Each relation appears as a box, with the relation name at the top in blue, and the attributes listed inside the box. Primary key attributes are shown underlined. Foreign key dependencies appear as arrows from the foreign key attributes of the referencing relation to the primary key of the referenced relation.

Table Joining

Joining Tables (<https://www.zoho.com/reports/help/table/joining-tables.html>)

In a reporting system often you might require to combine data from two or more tables to get the required information for analysis and reporting. To retrieve data from two or more tables, you have to combine the tables through the operation known as "Joining of tables". Joining is a method of establishing a relationship between tables using a common column.

[https://en.wikipedia.org/wiki/Join_\(SQL\)](https://en.wikipedia.org/wiki/Join_(SQL))

An SQL join clause combines columns from one or more tables in a relational database. It creates a set that can be saved as a table or used as it is. ***A JOIN is a means for combining columns from one (self-join) or more tables by using values common to each.*** ANSI-standard SQL specifies five types of JOIN: INNER, LEFT OUTER, RIGHT OUTER, FULL OUTER and CROSS. As a special case, a table (base table, view, or joined table) can JOIN to itself in a self-join.

Relational databases are usually normalized to ***eliminate duplication*** of information such as when entity types have one-to-many relationships. For example, a Department may be associated with a number of Employees. Joining separate tables for Department and Employee effectively creates another table which combines the information from both tables.

All subsequent explanations on join types in this article make use of the following two tables. The rows in these tables serve to illustrate the effect of different types of joins and join-predicates. In the following tables the DepartmentID column of the Department table (which can be designated as Department.DepartmentID) is the primary key, while Employee.DepartmentID is a foreign key.

Employee table	
LastName	DepartmentID
Rafferty	31
Jones	33
Heisenberg	33
Robinson	34
Smith	34
Williams	NULL

Department table	
DepartmentID	DepartmentName
31	Sales
33	Engineering
34	Clerical
35	Marketing

Note: In the Employee table above, the employee "Williams" has not been assigned to any department yet. Also, note that no employees are assigned to the "Marketing" department.

1- Cross join

CROSS JOIN returns the Cartesian product of rows from tables in the join. In other words, it will produce rows which combine each row from the first table with each row from the second table.[1]


Example of an implicit cross join:

```
SELECT * FROM employee, department;
```

Employee.LastName	Employee.DepartmentID	Department.DepartmentName	Department.DepartmentID
Rafferty	31	Sales	31
Jones	33	Sales	31
Heisenberg	33	Sales	31
Smith	34	Sales	31
Robinson	34	Sales	31
Williams	NULL	Sales	31
Rafferty	31	Engineering	33
Jones	33	Engineering	33
Heisenberg	33	Engineering	33
Smith	34	Engineering	33
Robinson	34	Engineering	33
Williams	NULL	Engineering	33
Rafferty	31	Clerical	34
Jones	33	Clerical	34
Heisenberg	33	Clerical	34
Smith	34	Clerical	34
Robinson	34	Clerical	34
Williams	NULL	Clerical	34
Rafferty	31	Marketing	35
Jones	33	Marketing	35
Heisenberg	33	Marketing	35
Smith	34	Marketing	35
Robinson	34	Marketing	35
Williams	NULL	Marketing	35

The cross join does not itself apply any predicate to filter rows from the joined table. The results of a cross join can be filtered by using a WHERE clause which may then produce the equivalent of an inner join.

2- Inner join



A Venn Diagram representing an Inner Join SQL statement between the tables A and B.

An inner join requires each row in the *two joined tables to have matching column values*, and is a commonly used join operation in applications but should not be assumed to be the best choice in all situations. Inner join creates a new result table by combining column values of two tables (A and B) based upon the join-predicate. The query compares each row of A with each row of B to find all pairs of rows which satisfy the join-predicate. When the join-predicate is satisfied by matching non-NULL values, column values for each matched pair of rows of A and B are combined into a result row.

The "explicit join notation" uses the JOIN keyword, optionally preceded by the INNER keyword, to specify the table to join, and the ON keyword to specify the predicates for the join, as in the following example:

```
SELECT employee.LastName, employee.DepartmentID, department.DepartmentName
FROM employee
INNER JOIN department ON
employee.DepartmentID = department.DepartmentID
```

Employee.LastName	Employee.DepartmentID	Department.DepartmentName
Robinson	34	Clerical
Jones	33	Engineering
Smith	34	Clerical
Heisenberg	33	Engineering
Rafferty	31	Sales

The "implicit join notation" simply lists the tables for joining, in the FROM clause of the SELECT statement, using commas to separate them. Thus it specifies a cross join, and the WHERE clause may apply additional filter-predicates (which function comparably to the join-predicates in the explicit notation).

The following example is equivalent to the previous one, but this time using implicit join notation:

```
SELECT *
FROM employee, department
WHERE employee.DepartmentID = department.DepartmentID;
```

The queries given in the examples above will join the Employee and Department tables using the DepartmentID column of both tables.

Where the DepartmentID of these tables match (i.e. the join-predicate is satisfied), the query will combine the LastName, DepartmentID and DepartmentName columns from the two tables into a result row. Where the DepartmentID does not match, no result row is generated.


Thus the result of the execution of the query above will be:

Employee.LastName	Employee.DepartmentID	Department.DepartmentName	Department.DepartmentID
Robinson	34	Clerical	34
Jones	33	Engineering	33
Smith	34	Clerical	34
Heisenberg	33	Engineering	33
Rafferty	31	Sales	31

The employee "Williams" and the department "Marketing" do not appear in the query execution results. Neither of these has any matching rows in the other respective table: "Williams" has no associated department, and no employee has the department ID 35 ("Marketing").

Programmers should take special care when joining tables on columns that can contain NULL values, since NULL will never match any other value (not even NULL itself), unless the join condition explicitly uses a combination predicate that first checks that the joins columns are NOT NULL before applying the remaining predicate condition(s).

3- Left outer join



A Venn Diagram representing the Left Join SQL statement between tables A and B.

The result of a *left outer join* (or simply **left join**) for tables A and B always contains all rows of the "left" table (A), even if the join-condition does not find any matching row in the "right" table (B). This means that if the `ON` clause matches 0 (zero) rows in B (for a given row in A), the join will still return a row in the result (for that row), but with NULL in each column from B.

A **left outer join** returns all the values from an inner join plus all values in the left table that do not match to the right table, including rows with NULL (empty) values in the link column.

For example, this allows us to find an employee's department, but still shows employees that have not been assigned to a department (contrary to the inner-join example above, where unassigned employees were excluded from the result).

Example of a left outer join (the `OUTER` keyword is optional), with the additional result row (compared with the inner join) italicized:

```
SELECT * FROM employee LEFT OUTER JOIN department ON
employee.DepartmentID = department.DepartmentID;
```


Employee.LastName	Employee.DepartmentID	Department.DepartmentName	Department.DepartmentID
Jones	33	Engineering	33
Rafferty	31	Sales	31
Robinson	34	Clerical	34
Smith	34	Clerical	34
<i>Williams</i>	NULL	NULL	NULL
Heisenberg	33	Engineering	33

Alternative syntaxes

Oracle supports the deprecated^[8] syntax:

```
SELECT *FROM employee, department
WHERE employee.DepartmentID = department.DepartmentID(+)
```

4- Right outer join



A Venn Diagram representing the Right Join SQL statement between tables A and B.

A **right outer join** (or **right join**) closely resembles a left outer join, except with the treatment of the tables reversed. Every row from the "right" table (B) will appear in the joined table at least once. If no matching row from the "left" table (A) exists, NULL will appear in columns from A for those rows that have no match in B.


A right outer join returns all the values from the right table and matched values from the left table (NULL in the case of no matching join predicate). For example, this allows us to find each employee and his or her department, but still show departments that have no employees.

Below is an example of a right outer join (the **OUTER** keyword is optional), with the additional result row italicized:

```
SELECT *  
FROM employee RIGHT OUTER JOIN department  
ON employee.DepartmentID = department.DepartmentID;
```

Employee.LastName	Employee.DepartmentID	Department.DepartmentName	Department.DepartmentID
Smith	34	Clerical	34
Jones	33	Engineering	33
Robinson	34	Clerical	34
Heisenberg	33	Engineering	33
Rafferty	31	Sales	31
NULL	NULL	<i>Marketing</i>	<i>35</i>

5- Full outer join



A Venn Diagram representing the Full Join SQL statement between tables A and B.

Conceptually, a **full outer join** combines the effect of applying both left and right outer joins. Where rows in the FULL OUTER JOINed tables do not match, the result set will have NULL values for every column of the table that lacks a matching row. For those rows that do match, a single row will be produced in the result set (containing columns populated from both tables).

For example, this allows us to see each employee who is in a department and each department that has an employee, but also see each employee who is not part of a department and each department which doesn't have an employee.

Example of a full outer join (the **OUTER** keyword is optional):

```
SELECT *  
FROM employee FULL OUTER JOIN department  
ON employee.DepartmentID = department.DepartmentID;
```

Employee.LastName	Employee.DepartmentID	Department.DepartmentName	Department.DepartmentID
Smith	34	Clerical	34
Jones	33	Engineering	33
Robinson	34	Clerical	34
Williams	NULL	NULL	NULL
Heisenberg	33	Engineering	33
Rafferty	31	Sales	31
NULL	NULL	Marketing	35

Example: Consider the following three tables :

teachers table

id	name
1	Volker
2	Elke

(2 rows)

Projects table

id	name	duration	teacher
1	compiler	180	1
2	xpaint	120	1
3	game	250	2
4	Perl	80	4

(4 rows)

Assign table

project	stud	percentage
1	2	10
1	4	60
1	1	30
2	1	50
2	4	50
3	2	70
3	4	30

(7 rows)

1. SELECT * FROM teachers, projects *where* teachers.id = projects.id;

id	name	id	name	duration	teacher
1	Volker	1	compiler	180	1
2	Elke	2	xpaint	120	1

2. inner join of tables *teachers* and *project* if the condition is teachers.id != projects.id will be

SELECT * FROM teachers, projects where teachers.id != projects.id;

id	name	id	name	duration	teacher
1	Volker	2	xpaint	180	1
1	Volker	3	game	180	1
1	Volker	4	Perl	180	1
2	Elke	1	compiler	120	1
2	Elke	3	game	120	1
2	Elke	4	Perl	120	1

6 Indexing and Hashing

Many queries reference only a small proportion of the records in a file. For example, the queries “ Find all accounts at the Tech branch” reference only a fraction of the account records. It is inefficient for the system to have to read every record and to check the branch names. The system should be able to locate these records directly. To allow that, we design additional structures with the files.

An index for a file in the system works like a catalog for a book in a library, if we are looking for a book, the catalog of the name of the books tells us where to find the book.

To assist us searching the catalog, the names in the catalog listed in an alphabetic order.

There are two basic kinds of indices :

- **Ordered indices:** such indices are based on a sorted ordering of the values.
- **Hash index :** such indices are based on some values, these values calculated by a function called hash function.

We often want to have more than one index for the file. Return to the example of the library, there can be a catalog for the names of the books and another catalog for the others of the books and third one for the subjects of the books.

6.1 Ordered Indices :

These are used to gain fast random access to records in a file. Each index structure is associated with a particular **search key**. The index stores the values of the search keys in sorted order.

The record in the indexed file may themselves be sorted in some way. The file may have several indices of different search key.

Primary index : if the file containing the record is sequentially ordered, the index whose search key specifies the sequential order of the file, this index is a primary index for that file.

Secondary index : is the index of the file whose search key specifies an order different from the order of the file.

6.1.1 : Ordered Primary Index

Figure 6.1 shows an ordered file for **account** records.

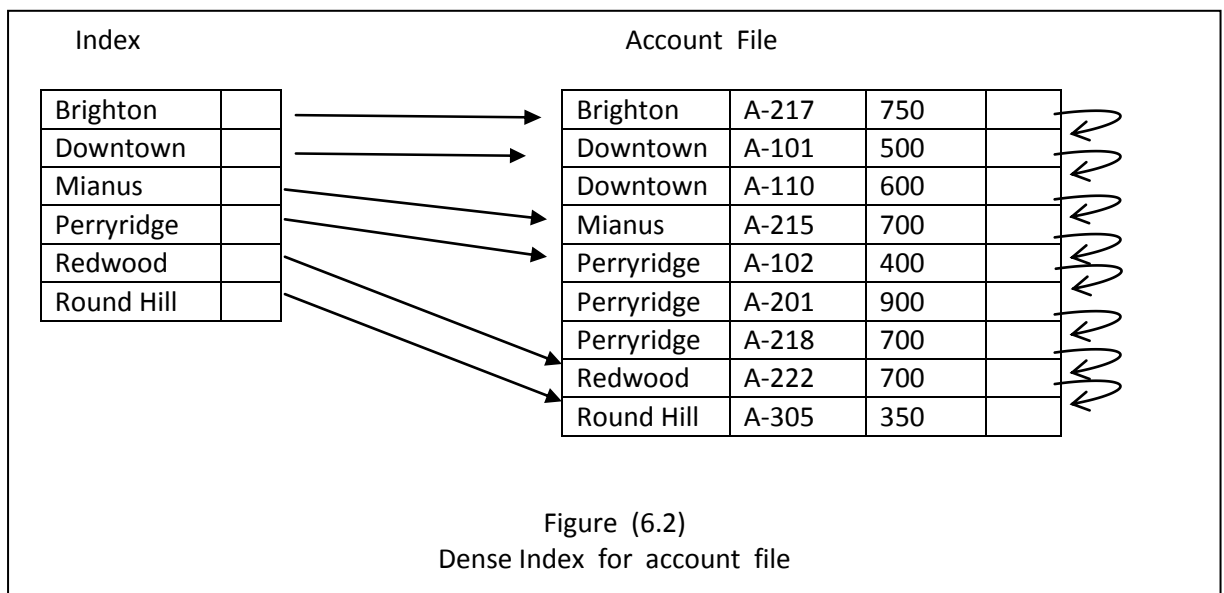
Brighton	A-217	750	
Downtown	A-101	500	
Downtown	A-110	600	
Mianus	A-215	700	
Perryridge	A-102	400	
Perryridge	A-201	900	
Perryridge	A-218	700	
Redwood	A-222	700	
Round Hill	A-305	350	

Figure (6.1) Sequential file for account records

The file of figure 6.1 is sorted on a search key order, with branch name is used as the search key.

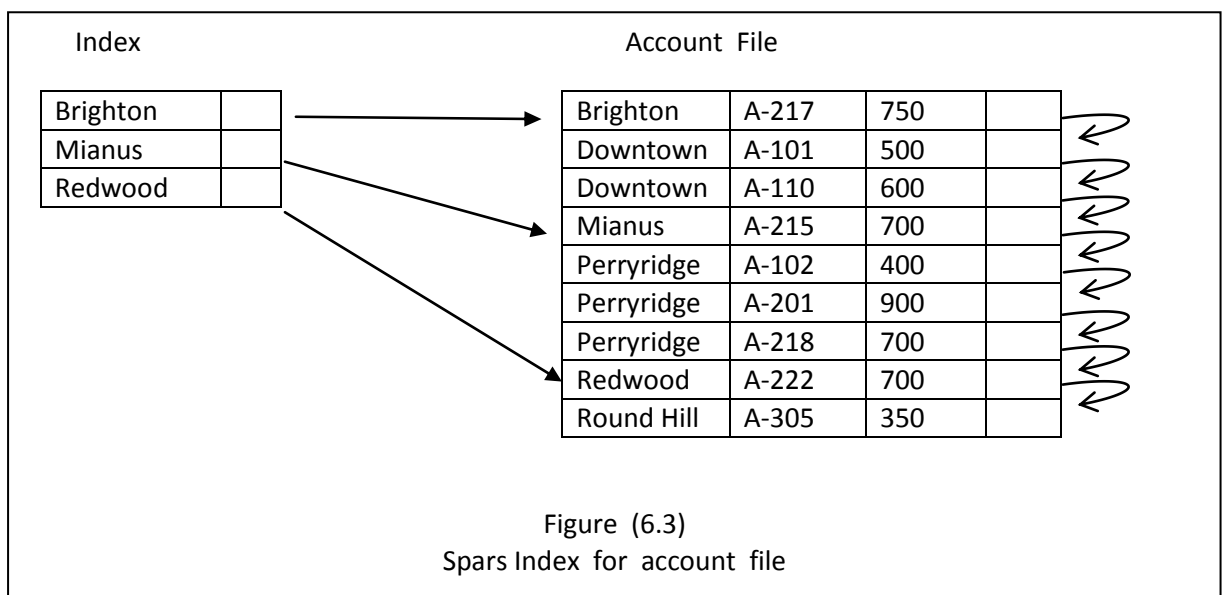
6.1.1.1 : Indices Types : There are two types of ordered indices : **Dense** and **Sparse** indices.

Dense index : an index entry appears for every search key value in the file. The index record contains the search key value and a pointer to the first data record with that search key value, as shown in figure (6.2) for the account file.



Suppose that we are locking up records for the PEERYIDGE branch using the dense index of figure (6.2), we follow the pointer directly to the first PEERYIDGE record . We process this record and **follow the pointer in that record** to locate the next record in search key (branch name) order. We continue processing records until we encounter a record for a branch other than PEERYIDGE.

Spars index : An index record is created for only some of the values. To locate a record we find the index entry with the largest search key value that is less than or equal to the search key value for which we are locking . We start at the record pointed to by that index entry, and follow the pointer in the file until we find the desired record. As shown in figure (6.3) for the account file.



If we are using the spars index of figure (6.3) to find PEERYIDGE , we do not find an index entry for PEERYIDGE in the index. Since the last entry (in alphabetic order) before BEERYIDGE is MIANUS , we follow that pointer.

We then read the account file in sequential order until we find the first PEERYIDEG record and begin processing at that point.

As we have seen , it is faster to locate a record by using a dense index rather than a sparse index., but sparse indices require less space.

6.1.1.2 Index Update

Every index must be updated whenever a record is inserted into the file or deleted from the file .

Deletion : To delete a record we first look up the record to be deleted .

In dense indices ,if the deleted record was the only record with its particular search key value, then the search key value is deleted from the index.

For sparse indices , we delete a key value by replacing its entry (if one exist) in the index with the next search key value . if the next search key value already has an index entry , the entry is deleted instead of being replace.

Insertion : First we perform a lookup using the search key value that appears in the record to be inserted .

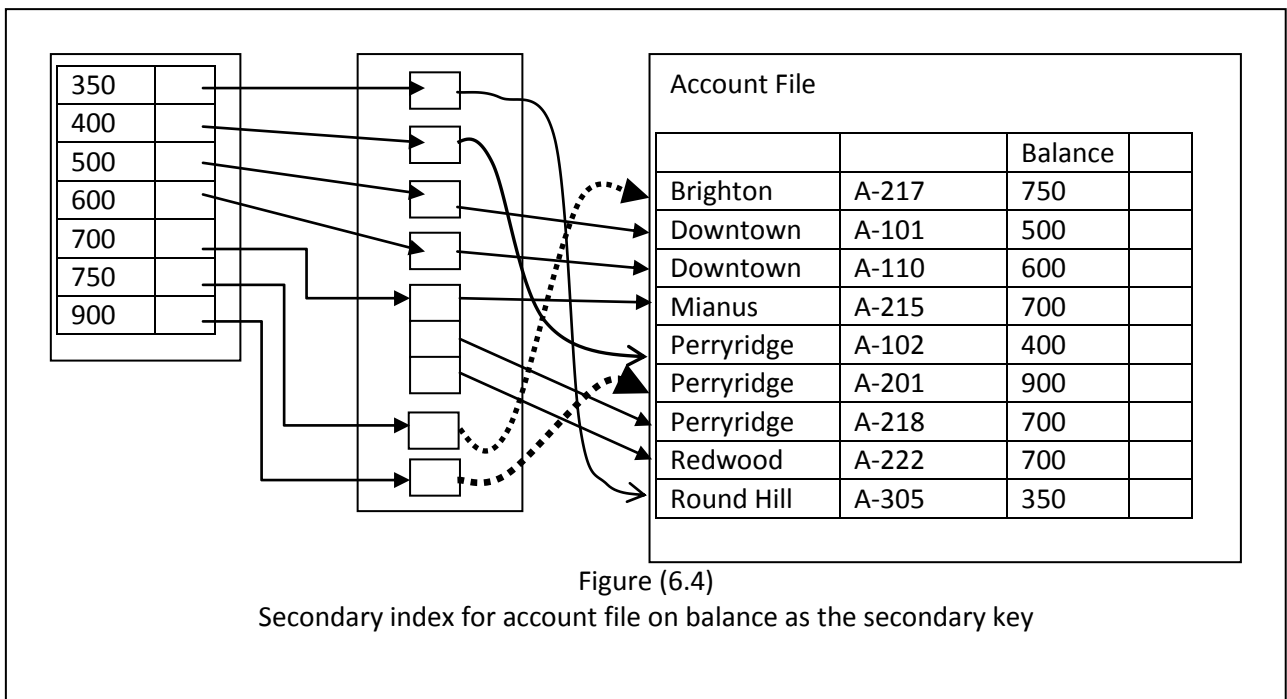
In dense index and the value does not appear in the index, the value is inserted in the index.

6.1.2 : Ordered Secondary Index

A secondary index looks like dense primary index except that the records pointed to by successive values in the index are not stored sequentially.

It is not enough to point to just the first record with each search key value because the remaining records with the same search key value could be anywhere in the file, since the records are ordered by a search key of the primary index rather than by the search key of the secondary index. Therefore a secondary index must contain pointers to all the records.

The pointers in the secondary index do not point directly to the file, instead each pointer points to a bucket that contains pointers to the file. Figure (6.4) shows the account file sorted by the balance field, which is not a primary key.



6.2 : Hash Index

One disadvantage of sequential file is that we must access an index structure to locate data or must use binary search.

File organization based on technique of hashing allow us to avoid accessing an index structure.

Hashing also provide a way of constructing indices.

6.2.1 Hash File Organization

In a hash file organization, we obtain the address of the disk block containing the desired record directly by computing a function on the search key value of the record.

Let (K) denote the set of all search key values. Let (B) denote the set of all blocks addresses.

A **hash function (H)** is a function from (K) to (B).

To insert a record with search key (K_i), we compute [$H(K_i)$] which gives the address of the block for that record and the record stores in that block.

To perform a lockup on the search key value (K_i), we compute [$H(K_i)$], then search the block with that address. Suppose that to search keys , K_5 and K_7 , have the same hash value; that is $H(K_5)=H(K_7)$, thus we have to check the search key value of every record in the block to find the desired record. Figure (6.5) shows Hash organization for the account file.

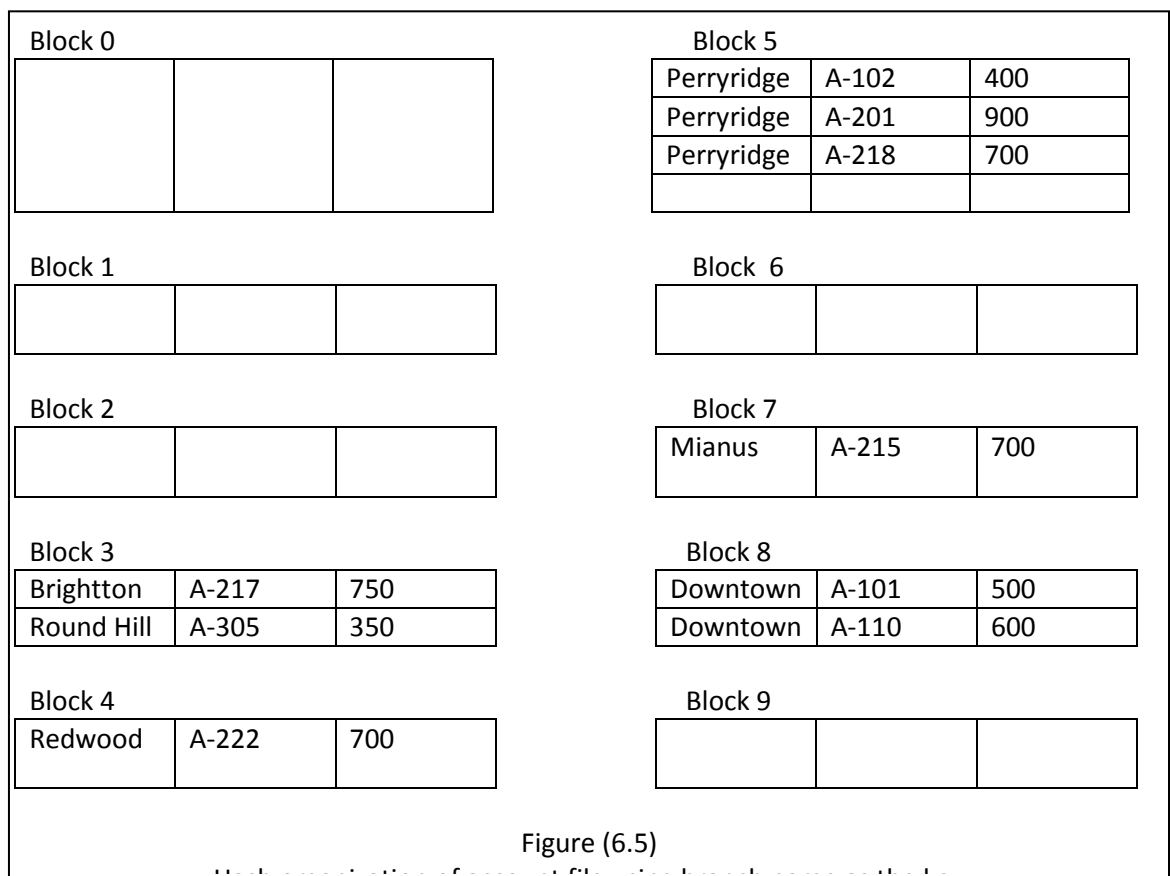


Figure (6.5)

6.2.2 Hash Index

Hashing can be used for index structure creation. A hash index organizes the search keys with their associated pointers into a hash file structure.

We construct a hash index as follows, we apply a hash function on a search key value to identify the block, and store the key and its pointer in that block as shown in figure (6.6).

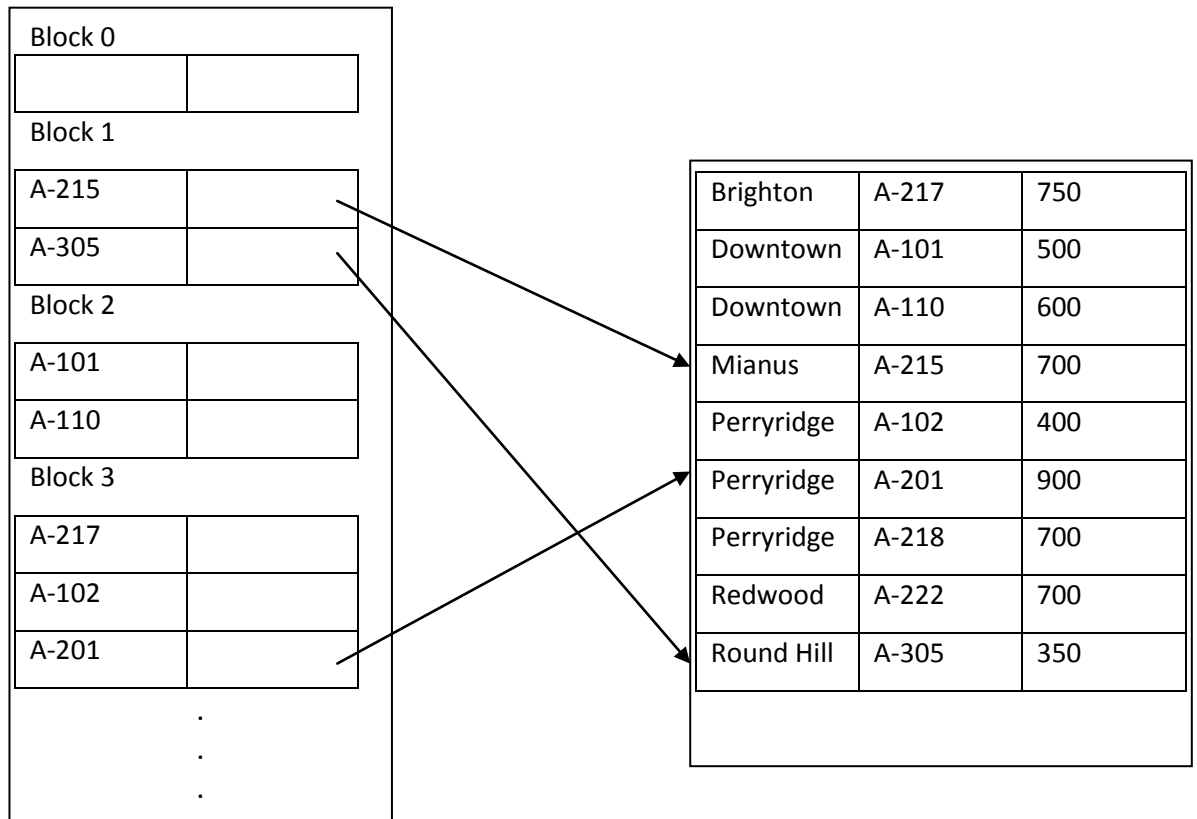


Figure (6.6)
Hash index on search key account-number of account file

28 Chapter 1 Introduction (Database System Concepts)

department *A* invokes a program called *new hire*. This program asks the clerk for the name of the new instructor, her new *ID*, the name of the department (that is, *A*), and the salary.

The typical user interface for naïve users is a forms interface, where the user can fill in appropriate fields of the form. Naïve users may also simply read *reports* generated from the database.

As another example, consider a student, who during class registration period, wishes to register for a class by using a Web interface. Such a user connects to a Web application program that runs at a Web server. The application first verifies the identity of the user, and allows her to access a form where she enters the desired information. The form information is sent back to the Web application at the server, which then determines if there is room in the class (by retrieving information from the database) and if so adds the student information to the class roster in the database.

- **Application programmers** are computer professionals who write application programs. Application programmers can choose from many tools to develop user interfaces. **Rapid application development (RAD)** tools are tools that enable an application programmer to construct forms and reports with minimal programming effort.

- **Sophisticated users** interact with the system without writing programs. Instead, they form their requests either using a database query language or by using tools such as data analysis software. Analysts who submit queries to explore data in the database fall in this category.

- **Specialized users** are sophisticated users who write specialized database applications that do not fit into the traditional data-processing framework. Among these applications are computer-aided design systems, knowledgebase and expert systems, systems that store data with complex data types (for example, graphics data and audio data), and environment-modeling systems. Chapter 22 covers several of these applications.

1.12.2 Database Administrator

One of the main reasons for using DBMSs is to have central control of both the data and the programs that access those data. A person who has such central control over the system is called a **database administrator (DBA)**. The functions of a DBA include:

- **Schema definition.** The DBA creates the original database schema by executing a set of data definition statements in the DDL.

- **Storage structure and access-method definition.**

- **Schema and physical-organization modification.** The DBA carries out changes to the schema and physical organization to reflect the changing needs of the organization, or to alter the physical organization to improve performance.

(Database System Concepts)

- **Granting of authorization for data access.** By granting different types of authorization, the database administrator can regulate which parts of the database various users can access. The authorization information is kept in a special system structure that the database system consults whenever someone attempts to access the data in the system.

- **Routine maintenance.** Examples of the database administrator's routine maintenance activities are:

- Periodically backing up the database, either onto tapes or onto remote servers, to prevent loss of data in case of disasters such as flooding.
- Ensuring that enough free disk space is available for normal operations, and upgrading disk space as required.
- Monitoring jobs running on the database and ensuring that performance is not degraded by very expensive tasks submitted by some users.

1.13 History of Database Systems

Information processing drives the growth of computers, as it has from the earliest days of commercial computers. In fact, automation of data processing tasks predates computers. Punched cards, invented by Herman Hollerith, were used at the very beginning of the twentieth century to record U.S. census data, and mechanical systems were used to process the cards and tabulate results. Punched cards were later widely used as a means of entering data into computers.

Techniques for data storage and processing have evolved over the years:

- **1950s and early 1960s:** Magnetic tapes were developed for data storage. Data processing tasks such as payroll were automated, with data stored on tapes.

Processing of data consisted of reading data from one or more tapes and writing data to a new tape. Data could also be input from punched card decks, and output to printers. For example, salary raises were processed by entering the raises on punched cards and reading the punched card deck in synchronization with a tape containing the master salary details. The records had to be in the same sorted order. The salary raises would be added to the salary read from the master tape, and written to a new tape; the new tape would become the new master tape.

Tapes (and card decks) could be read only sequentially, and data sizes were much larger than main memory; thus, data processing programs were forced to process data in a particular order, by reading and merging data from tapes and card decks.

- **Late 1960s and 1970s:** Widespread use of hard disks in the late 1960s changed the scenario for data processing greatly, since hard disks allowed direct access to data. The position of data on disk was immaterial, since any location on disk could be accessed in just tens of milliseconds. Data were thus freed from

The Design Process (http://en.wikipedia.org/wiki/Database_design#cite_note-3 “Database design – Wikipedia”)

The design process consists of the following steps

- 1- **Determine the purpose of your database** - This helps prepare you for the remaining steps.
- 2- **Find and organize the information required** - Gather all of the types of information you might want to record in the database, such as product name and order number.
- 3- **Divide the information into tables** - Divide your information items into major entities or subjects, such as Products or Orders. Each subject then becomes a table.
- 4- **Turn information items into columns** - Decide what information you want to store in each table. Each item becomes a field, and is displayed as a column in the table. For example, an Employees table might include fields such as Last Name and Hire Date.
- 5- **Specify primary keys** - Choose each table’s primary key. The primary key is a column that is used to uniquely identify each row. An example might be Product ID or Order ID.
- 6- **Set up the table relationships** - Look at each table and decide how the data in one table is related to the data in other tables. Add fields to tables or create new tables to clarify the relationships, as necessary.
- 7- **Refine your design** - Analyze your design for errors. Create the tables and add a few records of sample data. See if you can get the results you want from your tables. Make adjustments to the design, as needed.

Normalization

Normalization In relational database , is the process of organizing data to minimize redundancy . The goal of database normalization is to decompose complex relations(tables) in order to produce smaller, well-structured relations(tables).

Normalization usually involves dividing large, badly-formed tables into smaller, well-formed tables and defining relationships between them. The objective is to isolate data so that additions, deletions, and modifications of a field can be made in just one table and then propagated through the rest of the database via the defined relationships

Free the database of modification anomalies

When an attempt is made to modify (update, insert into, or delete from) a table, undesired side-effects may follow. Not all tables can suffer from these side-effects; rather, *the side-effects can only arise in tables that have not been sufficiently normalized*. An insufficiently normalized table might have one or more of the following characteristics:

- The same information can be expressed on multiple rows; therefore updates to the table may result in logical inconsistencies. For example, each record in an "Employees' Skills" table might contain an Employee ID, Employee Address, and Skill; thus a change of address for a particular employee will potentially need to be applied to multiple records (one for each of his skills). If the update is not carried through successfully—if, that is, the employee's address is updated on some records but not others—then the table is left in an inconsistent state. Specifically, the table provides conflicting answers to the question of what this particular employee's address is. This phenomenon is known as an **update anomaly**.
- There are circumstances in which certain facts cannot be recorded at all. For example, each record in a "Faculty and Their Courses" table might contain a Faculty ID, Faculty Name, Faculty Hire Date, and Course Code—thus we can record the details of any faculty member who teaches at least one course, but we cannot record the details of a newly-hired faculty member who has not yet been assigned to teach any courses except by setting the Course Code to null. This phenomenon is known as an **insertion anomaly**.
- There are circumstances in which the deletion of data representing certain facts necessitates the deletion of data representing completely different facts. The "Faculty and Their Courses" table described in the previous example suffers from this type of anomaly, for if a faculty member temporarily ceases to be assigned to any courses, we must delete the last of the records on which that faculty member appears, effectively also deleting the faculty member. This phenomenon is known as a **deletion anomaly**.

Employees' Skills

Employee ID	Employee Address	Skill
426	87 Sycamore Grove	Typing
426	87 Sycamore Grove	Shorthand
519	94 Chestnut Street	Public Speaking
519	96 Walnut Avenue	Carpentry

An update anomaly. Employee 519 is shown as having different addresses on different records.

Faculty and Their Courses

Faculty ID	Faculty Name	Faculty Hire Date	Course Code
389	Dr. Giddens	10-Feb-1985	ENG-206
407	Dr. Saperstein	19-Apr-1999	CMP-101
407	Dr. Saperstein	19-Apr-1999	CMP-201

424
Dr. Newsome
29-Mar-2007
?

An insertion anomaly. Until the new faculty member, Dr. Newsome, is assigned to teach at least one course, his details cannot be recorded.

Faculty and Their Courses

Faculty ID	Faculty Name	Faculty Hire Date	Course Code
389	Dr. Giddens	10-Feb-1985	ENG-206
407	Dr. Saperstein	19-Apr-1999	CMP-101
407	Dr. Saperstein	19-Apr-1999	CMP-201

DELETE

A deletion anomaly. All information about Dr. Giddens is lost when he temporarily ceases to be assigned to any courses.

→ : Levels of normalization

To normalize a database, there are three levels :

- First Normal Form (1NF)
- Second Normal Form (2NF)
- Third Normal Form (3NF)

1: First Normal Form (1NF)

The role of (1NF) is: No repeating elements or groups of elements, this mean:

- Eliminate duplicative columns from the same table.
- Create separate tables for each group of related data and identify each row with a unique column or set of columns (the primary key).

The first rule dictates that we must not duplicate data within the same row of a table. For example consider the table of figure (7.4) of a bank.

Person ID	Name	City Name	City Number	Account Type	Balance	Account Notes
123	Nader	Baghdad	1	A	4556 \$	*****
123	Nader	Baghdad	1	B	7654 \$	#####
123	Nader	Baghdad	1	C	1287 \$	&&&&&
150	Muna	Basra	2	A	654 \$	*****
150	Muna	Basra	2	B	66743 \$	#####

Figure (7.4) Accounts table of a bank

The table of figure (7.4) consists of repeated information.

The meaning of repeated information is all the information belong for a single person. For person "NADER" there are three records repeated for him.

For "Muna" there are two record repeated for her.

The table must be **separated** from the repeated information into **two tables** as shown in figure (7.5). The two tables are joined by the **Person ID** as a key. For table 2 the primary key is the combination of the first two fields 'Person Id + Account Type'

Primary Key				The Primary Key			
Person ID	Name	City Name	City Number	Person ID	Account Type	Balance	Account Notes
123	Nader	Baghdad	1	123	A	4556 \$	*****
150	Muna	Basra	2	123	B	7654 \$	#####
				123	C	1287 \$	&&&&&
				150	A	654 \$	*****
				150	B	66743 \$	#####

Figure (7.5) the new two tables of (1NF)

2: Second Normal Form (2NF)

A table that has a concatenated primary key, each column in the table that is not part of the primary key must depend upon the entire concatenated key for its existence. If any column only depends upon one part of the concatenated key, then we say that the entire table has failed Second Normal Form and we must create another table to rectify the failure

We look at the tables with a primary key made from many fields, for each non key fields that do not depend on the primary key (with all the fields of the primary key), this non key fields must be separated in another table.

Table 1 of figure (7.5) have a primary key with single fields, this mean it is in (2NF).

Table 2 of figure (7.5) , field BALANCE depends on both the fields of the primary key, because each balance we must know the person and the account for it.

But “ACCOUNT NOTES” depends only on the second field (Account Type) because the not of the account is the same for all the persons .

Figure (7.6) shows the (2NF)

Person ID	Name	City Name	City Number
123	Nader	Baghdad	1
150	Muna	Basra	2

Table 1 with no change

Person ID	Account Type	Balance
123	A	4556 \$
123	B	7654 \$
123	C	1287 \$
150	A	654 \$
150	B	66743 \$

Table 2

Account Type	Account Notes
A	*****
B	#####
C	&&&&&

Table 3

Table 2 of figure (7.5) is separated into two tables

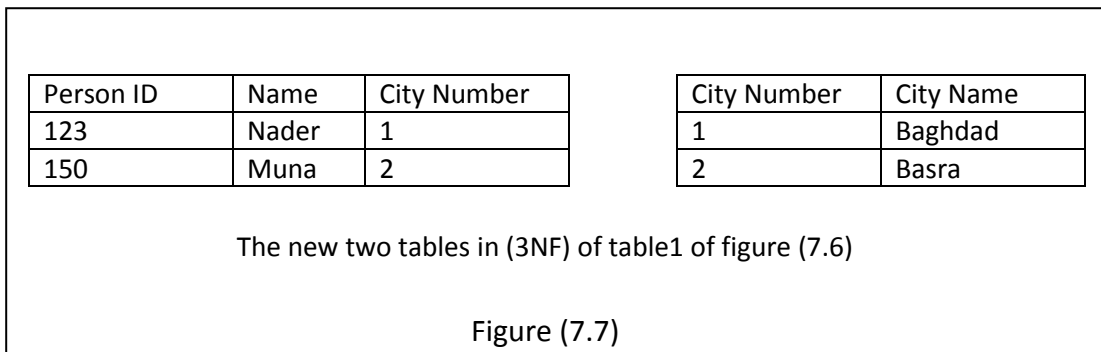
Figure (7.6) the Second Normal Form (2NF)

3 : Third Normal Form (3NF)

Third normal form (3NF) is to remove columns that are not dependent upon the primary key. This means we look at the non key fields and try to find a relation between them.

Table 2 and table 3 of figure (7.6) are in the (3NF). But table 1 is not.

In table 1 we find that City Number depends on City Name , not on the key. We separate these two fields on a table and link this new table with the exist one as shown in figure (7.7)



Example: Find the appropriate tables in 3NF for the following inventory form

Store Inventory form Store No : Store Address : Date of inventory :			
Part No.	Description	Location	Quantity

Solution:

First we must find the unnormalized table :

Store no	Store address	Date of inventory	Location	Part No	Desc	Qty

1- The 1NF will be:

Table 1

Store no	Store address	Date of inventory
----------	---------------	-------------------

Primary Key

Table 2

Store no	Location	Part No	Desc	Qty
----------	----------	---------	------	-----

Primary Key

2- The 2NF : (non-key fields depends on Key-field)

Table 1 of 1Nf is the 2NF with no change.

Table 2 has the following problem: when we look at the non-key fields, we find that two fields **(Part No) and (Qty)** depends on the entire primary key, but **(Desc) don't depend on any part of the primary key**. instead (Desc) depends on non-key field (Part No).

The 2NF will be

Table 1

Store no	Store address	Date of inventory
----------	---------------	-------------------

Primary Key

(table 2 of 1NF will be)

Table 2 on 2NF

Store no	Location	Part No	Qty
----------	----------	---------	-----

Primary Key

Table 3 of 2NF

Part No	Desc
---------	------

Primary Key

When we splits (Desc) on a new table 3 , we put with it (Part No) to connect it with table 2

3- The 3NF : (Non-key field depends on non-key field)

3NF is the same as 2NF.

Problem : look at table 1 of 3NF (or 2NF) , if any store have many inventory bill, each bill with a deferent date. In table 1 , we must sore a record for each date and repeat the store no **and the address** . This is a problem.

It is better to split table 1 into two tables as bellow :

Table 1-1

Store no	Store address
----------	---------------

Table 1-2

Store no	Date of inventory
----------	-------------------

Another problem arise : in table 2 how we will know each record belong to which date, i.e. how we shall link table 2 with table 1-2 ?

CHAPTER17 (Database System Concepts)

Database-System Architectures

The architecture of a database system is greatly influenced by the underlying computer system on which it runs, in particular by such aspects of computer architecture as networking, parallelism, and distribution:

- Networking of computers allows some tasks to be executed on a server system and some tasks to be executed on client systems. This division of work has led to *client–server database systems*.
- Parallel processing within a computer system allows database-system activities to be speeded up, allowing faster response to transactions, as well as more transactions per second. Queries can be processed in a way that exploits the parallelism offered by the underlying computer system. The need for parallel query processing has led to *parallel database systems*.
- Distributing data across sites in an organization allows those data to reside where they are generated or most needed, but still to be accessible from other sites and from other departments. Keeping multiple copies of the database across different sites also allows large organizations to continue their database operations even when one site is affected by a natural disaster, such as flood, fire, or earthquake. *Distributed database systems* handle geographically or administratively distributed data spread across multiple database systems.

We study the architecture of database systems in this chapter, starting with the traditional centralized systems, and covering client–server, parallel, and distributed database systems.

17.1 Centralized and Client–Server Architectures

Centralized database systems are those that run on a single computer system and do not interact with other computer systems. Such database systems span a range from single-user database systems running on personal computers to high-performance database systems running on high-end server systems. Client–server systems, on the other hand, have functionality split between a server system and multiple client systems.

17.1.1 Centralized Systems

A modern, general-purpose computer system consists of one to a few processors and a number of device controllers that are connected through a common bus that provides access to shared memory (Figure 17.1).

The processors have local cache memories that store local copies of parts of the memory, to speed up access to data. Each processor may have several independent **cores**, each of which can execute a separate instruction stream. Each device controller is in charge of a specific type of device (for example, a disk drive, an audio device, or a video display).

The processors and the device controllers can execute concurrently, competing for memory access. Cache memory reduces the contention for memory access, since it reduces the number of times that the processor needs to access the shared memory.

We distinguish two ways in which computers are used: as single-user systems and as multiuser systems. Personal computers and workstations fall into the first category. A typical **single-user system** is a desktop unit used by a single person, usually with only one processor and one or two hard disks, and usually only one person using the machine at a time. A typical **multiuser system**, on the other hand, has more disks and more memory and may have multiple processors. It serves a large number of users who are connected to the system remotely.

Database systems designed for use by single users usually do not provide many of the facilities that a multiuser database provides. In particular, they may not support concurrency control, which is not required when only a single user can generate updates.

Provisions for crash recovery in such systems are either absent or primitive—for example, they may consist of simply making a backup of the database before any update. Some such systems do not support SQL, and they provide a simpler query language, such as a variant of QBE. In contrast,

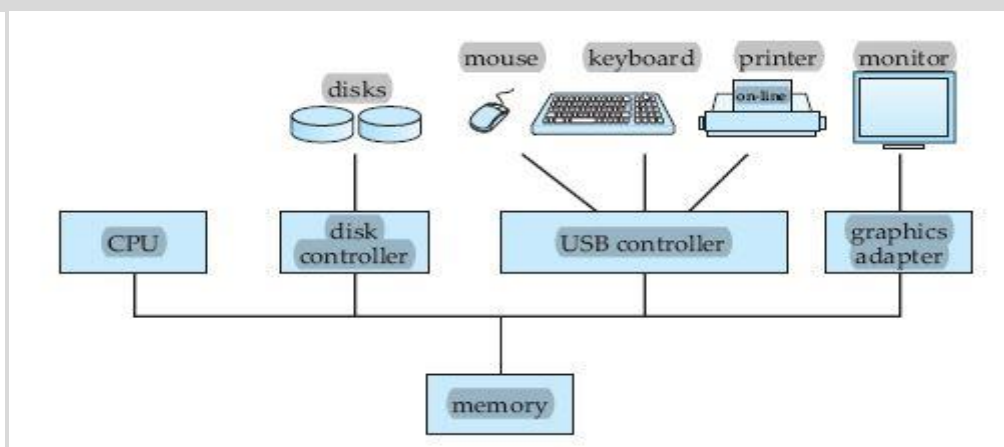


Figure 17.1 A centralized computer system.

17.1 Centralized and Client–Server Architectures 771

database systems designed for multiuser systems support the full transactional features that we have studied earlier.

Although most general-purpose computer systems in use today have multiple processors, they have **coarse-granularity parallelism**, with only a few processors (about two to four, typically), all sharing the main memory. Databases running on such machines usually do not attempt to partition a single query among the processors; instead, they run each query on a single processor, allowing multiple queries to run concurrently. Thus, such systems support a higher throughput; that is, they allow a greater number of transactions to run per second, although individual transactions do not run any faster.

Databases designed for single-processor machines already provide multitasking, allowing multiple processes to run on the same processor in a time-shared manner, giving a view to the user of multiple processes running in parallel. Thus, coarse-granularity parallel machines logically appear to be identical to single processor machines, and database systems designed for time-shared machines can be easily adapted to run on them.

In contrast, machines with **fine-granularity parallelism** have a large number of processors, and database systems running on such machines attempt to parallelize single tasks (queries, for example) submitted by users. We study the architecture of parallel database systems in Section 17.3.

Parallelism is emerging as a critical issue in the future design of database systems. Whereas today those computer systems with multicore processors have only a few cores, future processors will have large numbers of cores.¹ As a result, parallel database systems, which once were specialized systems running on specially designed hardware, will become the norm.

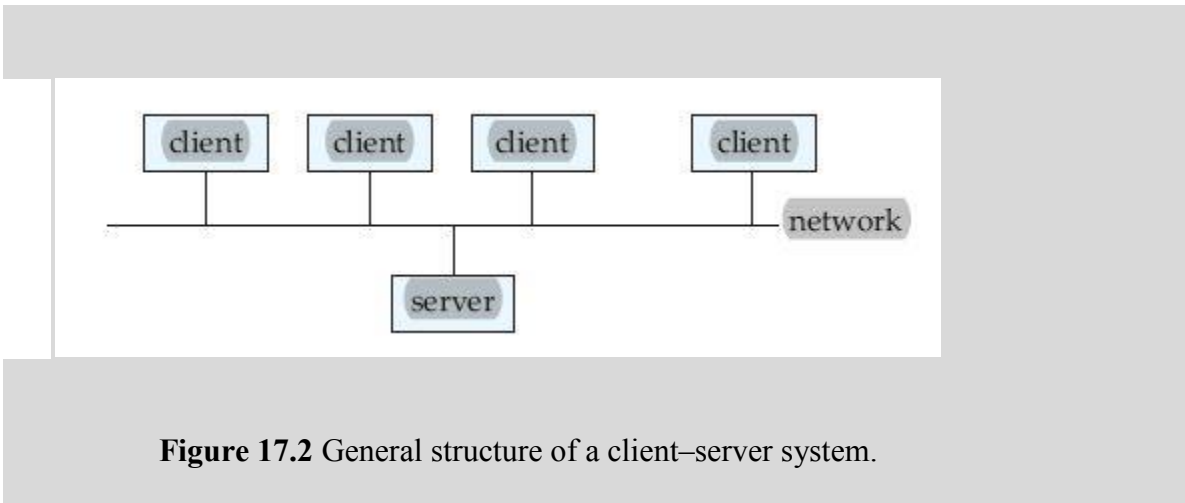
17.1.2 Client–Server Systems

As personal computers became faster, more powerful, and cheaper, there was a shift away from the centralized system architecture. Personal computers supplanted terminals connected to centralized systems.

Correspondingly, personal computers assumed the user-interface functionality that used to be handled directly by the centralized systems. As a result, centralized systems today act as **server systems** that satisfy requests generated by *client systems*.

Figure 17.2 shows the general structure of a client–server system.

Functionality provided by database systems can be broadly divided into two parts—the front end and the back end. The back end manages access structures, query evaluation and optimization, concurrency control, and recovery. The front end of a database system consists of tools such as the SQL user interface, forms interfaces, report generation tools, and data mining and analysis tools (see Figure 17.3). The interface between the front end and the back end is through SQL, or through an application program.

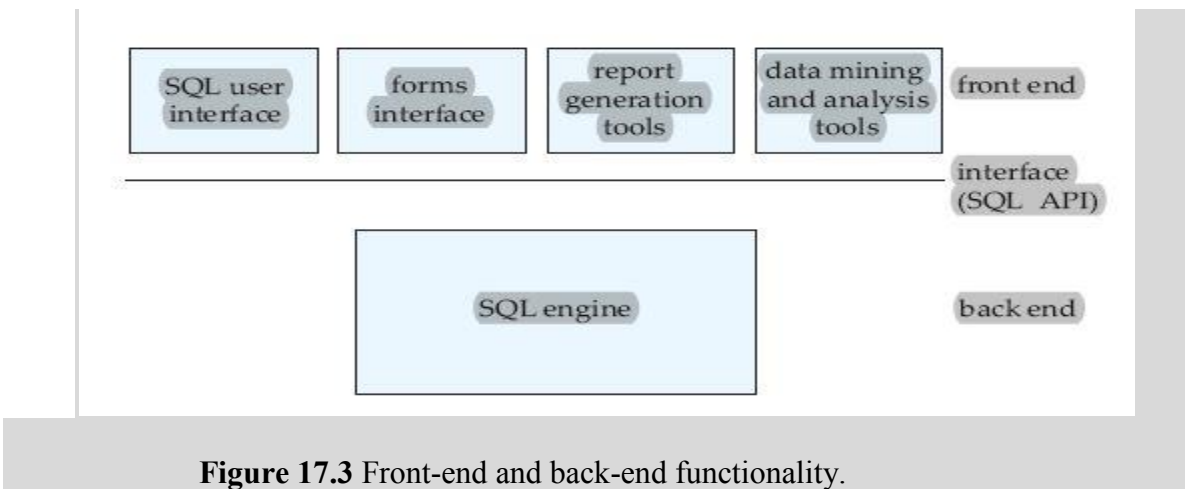


Standards such as *ODBC* and *JDBC*, which we saw in Chapter 3, were developed to interface clients with servers. Any client that uses the ODBC or JDBC interface can connect to any server that provides the interface. Certain application programs, such as spreadsheets and statistical-analysis packages, use the client–server interface directly to access data from a back-end server. In effect, they provide front ends specialized for particular tasks. Systems that deal with large numbers of users adopt a three-tier architecture, which we saw earlier in Figure 1.6 (Chapter 1), where the front end is a Web browser that talks to an application server. The application server, in effect, acts as a client to the database server.

Some transaction-processing systems provide a **transactional remote procedure call** interface to connect clients with a server. These calls appear like ordinary procedure calls to the programmer, but all the remote procedure calls from a client are enclosed in a single transaction at the server end. Thus, if the transaction aborts, the server can undo the effects of the individual remote procedure calls.

17.2 Server System Architectures

Server systems can be broadly categorized as transaction servers and data servers.



17.2.3 Cloud-Based Servers

Servers are usually owned by the enterprise providing the service, but there is an increasing trend for service providers to rely at least in part upon servers that are owned by a “third party” that is neither the client nor the service provider. One model for using third-party servers is to outsource the entire service

to another company that hosts the service on its own computers using its own software. This allows the service provider to ignore most details of technology and focus on the marketing of the service.

Another model for using third-party servers is **cloud computing**, in which the service provider runs its own software, but runs it on computers provided by another company. Under this model, the third party does not provide any of the application software; it provides only a collection of machines. These machines are not “real” machines, but rather simulated by software that allows a single real computer to simulate several independent computers. Such simulated machines are called **virtual machines**. The service provider runs its software (possibly including a database system) on these virtual machines. A major advantage of cloud computing is that the service provider can add machines as needed to meet demand and release them at times of light load. This can prove to be highly cost-effective in terms of both money and energy.

A third model uses a cloud computing service as a data server; such *cloud-based data storage* systems are covered in detail in Section 19.9. Database applications using cloud-based storage may run on the same cloud (that is, the same set of machines), or on another cloud. The bibliographical references provide more information about cloud-computing systems.

17.3 Parallel Systems

Parallel systems improve processing and I/O speeds by using multiple processors and disks in parallel. Parallel machines are becoming increasingly common, making the study of parallel database systems correspondingly more important. The driving force behind parallel database systems is the demands of applications that have to query extremely large databases (of the order of terabytes—that is, 10^{12} bytes) or that have to process an extremely large number of transactions per second

(of the order of thousands of transactions per second). Centralized and client-server database systems are not powerful enough to handle such applications.

In parallel processing, many operations are performed simultaneously, as opposed to serial processing, in which the computational steps are performed sequentially.

A **coarse-grain** parallel machine consists of a small number of powerful processors; a **massively parallel** or **fine-grain parallel** machine uses thousands of smaller processors.

Virtually all high-end machines today offer some degree of coarse-grain parallelism: at least two or four processors.

Massively parallel computers can be distinguished from the coarse-grain parallel machines by the much larger degree of parallelism that they support. Parallel computers with hundreds of processors and disks are available commercially.

There are two main measures of performance of a database system:

(1) **throughput**, the number of tasks that can be completed in a given time interval, and

(2) **response time**, the amount of time it takes to complete a single task from the time it is submitted. A system that processes a large number of small transactions can improve throughput by processing many transactions in parallel. A system that processes large transactions can improve response time as well as throughput by performing subtasks of each transaction in parallel.

17.3.1 Speedup and Scaleup

Two important issues in studying parallelism are speedup and scaleup. Running a given task in less time by increasing the degree of parallelism is called **speedup**.

Handling larger tasks by increasing the degree of parallelism is called **scaleup**.

Consider a database application running on a parallel system with a certain number of processors and disks. Now suppose that we increase the size of the system by increasing the number of processors, disks, and other components of the system. The goal is to process the task in time inversely proportional to the number of processors and disks allocated. Suppose that the execution time of a task on the larger machine is TL , and that the execution time of the same task on the smaller machine is TS . The speedup due to parallelism is defined as TS/TL . The parallel system is said to demonstrate **linear speedup** if the speedup is N when the larger system has N times the resources (processors, disk, and so on) of the smaller system. If the speedup is less than N , the system is said to demonstrate **sublinear speedup**. Figure 17.5 illustrates linear and sublinear speedup.

Scaleup relates to the ability to process larger tasks in the same amount of time by providing more resources. Let Q be a task, and let QN be a task that is N times bigger than Q . Suppose that the execution time of task Q on a given machine

linear speedup

sublinear speedup

resources

speed

17.3.3.4 Hierarchical

The **hierarchical architecture** combines the characteristics of shared-memory, shared-disk, and shared-nothing architectures. At the top level, the system consists of nodes that are connected by an interconnection network and do not share disks or memory with one another. Thus, the top level is a shared-nothing architecture.

Each node of the system could actually be a shared-memory system with a few processors. Alternatively, each node could be a shared-disk system, and each of the systems sharing a set of disks could be a shared-memory system. Thus, a system could be built as a hierarchy, with shared-memory architecture with a few processors at the base, and a shared-nothing architecture at the top, with possibly a shared-disk architecture in the middle. Figure 17.8d illustrates a hierarchical architecture with shared-memory nodes connected together in a shared-nothing architecture. Commercial parallel database systems today run on several of these architectures.

Attempts to reduce the complexity of programming such systems have yielded **distributed virtual-memory** architectures, where logically there is a single shared memory, but physically there are multiple disjoint memory systems; the virtual memory-mapping hardware, coupled with system software, allows each processor to view the disjoint memories as a single virtual memory. Since access speeds differ, depending on whether the page is available locally or not, such an architecture is also referred to as a **nonuniform memory architecture (NUMA)**.

17.4 Distributed Systems

In a **distributed database system**, the database is stored on several computers. The computers in a distributed system communicate with one another through various communication media, such as high-speed private networks or the Internet. They do not share main memory or disks.

The computers in a distributed system may vary in size and function, ranging from workstations up to mainframe systems.

The computers in a distributed system are referred to by a number of different names, such as **sites** or **nodes**, depending on the context in which they are mentioned. We mainly use the term **site**, to emphasize the physical distribution of these systems. The general structure of a distributed system appears in Figure 17.9.

The main differences between shared-nothing parallel databases and distributed databases are that distributed databases are typically geographically separated, are separately administered, and have a slower interconnection. Another major difference is that,

in a distributed database system, we differentiate between local and global transactions. A **local transaction** is one that accesses data only from sites where the transaction was initiated. A **global transaction**, on the other hand, is one that either accesses data in a site different from the one at which the transaction was initiated, or accesses data in several different sites.

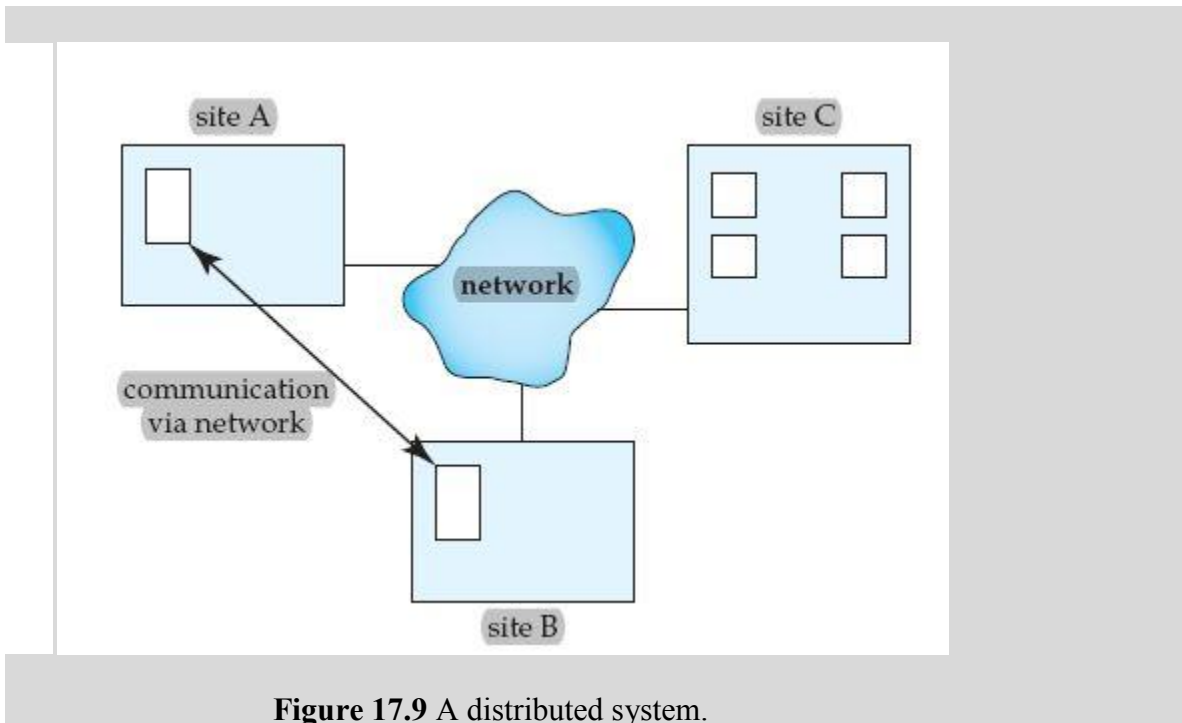


Figure 17.9 A distributed system.

There are several reasons for building distributed database systems, including sharing of data, autonomy, and availability.

- **Sharing data.** The major advantage in building a distributed database system is the provision of an environment where users at one site may be able to access the data residing at other sites. For instance, in a distributed university system, where each campus stores data related to that campus, it is possible for a user in one campus to access data in another campus. Without this capability, the transfer of student records from one campus to another campus would have to resort to some external mechanism that would couple existing systems.

- **Autonomy.** The primary advantage of sharing data by means of data distribution is that each site is able to retain a degree of control over data that are stored locally. In a centralized system, the database administrator of the central site controls the database. In a distributed system, there is a global database administrator responsible for the entire system. A part of these responsibilities is delegated to the local database administrator for each site.

Depending on the design of the distributed database system, each administrator may have a different degree of **local autonomy**. The possibility of local autonomy is often a major advantage of distributed databases.

- **Availability.** If one site fails in a distributed system, the remaining sites may be able to continue operating. In particular, if data items are **replicated** in several sites, a transaction needing a particular data item may find that item in any of several sites. Thus, the failure of a site does not necessarily imply the shutdown of the system.

9 Transactions

9.1 : Definition :A transaction is unit of program execution that accesses and possibly updates various data items.

Often a collection of several operations on the database is considered to be a single unit from the point of view of the user. For example a transfer of funds from a checking account to a saving account is a single operation for the user , but for the database it comprise several operations. A database system must ensure proper execution of transaction, either the entire transaction is executed or none of it does.

To ensure integrity of the data, the database system maintains the following properties of the transaction:

- 1- Atomicity : either all operations of the transaction are executed or none are.
- 2- Consistency : execution of the transaction in isolation to preserve the consistency of the database
- 3- Isolation : if there are two transaction T_a and T_b ; it appear for T_a that T_b either finished before T_a started , or T_b start execution after T_a finished.
Thus each transaction unaware of other transactions executing concurrently in the system.
- 4- Durability : after a transaction completes successfully , the change it has made to the database persist even if there are a system failures.

For example , access to the database is accomplished by the two operation:

- Read (X) : is to read X from the database to a local buffer belonging to the transaction.
- Write (X) : is to write X to the database from the local buffer.

Let T_a be a transaction that transfer 50\$ from account A to account B as follow:

```
Ta : Read(A)
     A=A - 50;
     Write (A)
     Read(B)
     B=B + 50;
     Write(B)
```

*Let us consider **Atomicity*** : the database system keeps track (on disk) of the old values of any data on which a transaction performs a write, and if the transaction does not complete its execution, the old value is restored to make it appear as though the transaction never executed.

*Let us consider **Consistency*** : the consistency required here is the sum of A and B be unchanged.

Without this consistency, money could be created or destroyed by the transaction. If the database is consistent before the execution of the transaction then the database must remain consistent after the execution of the transaction.

Let us consider **Durability** : we assume that system failure may result of losing data in main memory but data written to disk are never lost. We can guarantee durability by:

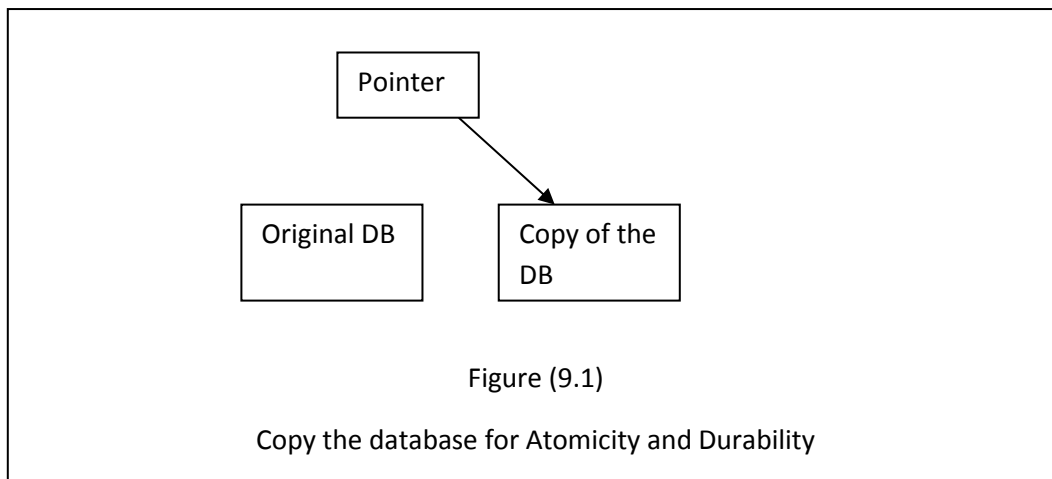
- 1- The update carried by the transaction have been written to disk before the transaction completes.
- 2- There are sufficient Information about the update to enable the database to reconstruct the update when the system is restarted after the failure.

Let us consider **Isolation** : while executing the T_a transaction , the T_a transaction execute the third statement (Write (A)) and before executing the last statement (Write (B)), another transaction T_b read the values of A and B , and calculate the sum ($A+B$), then the consistency of transaction T_b will discover an error of the values of A and B.

The solution of the problem of concurrently execution is to execute the transaction in sequence, one after another.

9.2 : Implementation of atomicity and durability

The recovery-management system is the one responsible of atomicity and durability. One way to implement them is to use a copy of the database, and a pointer that point at the current copy as shown in figure (9.1).



If a transaction want to update the database , all updates are done on the copy database. If an error occur during the execution of the transaction then the copy is deleted and the system return to the original database.

9.3 : Concurrent Execution

Transaction processing system usually allow multiple transaction to run concurrently. When several transaction run concurrently, database consistency can be destroyed.

The database system must control the interaction among the concurrent transaction to prevent them from destroying the consistency of the database.

As an example , consider a banking system which has several accounts, and a set of transactions that access and updates these accounts.

Let T1 and T2 be two transactions that transfer funds from one account to another, they can be executed one after another T1 then T2 as shown in figure (9.2) or T2 then T1 as shown in figure (9.3).

T1	T2
Read(A) A=A-50 Write(A) Read (B) B=B+50 Write (B)	
	Read (A) temp=A * 0.1 A=A-temp Write (A) Read(B) B=B+temp Write(B)

Figure (9.2)
Execution of T1 then T2

T1	T2
	Read (A) temp=A * 0.1 A=A-temp Write (A) Read(B) B=B+temp Write(B)
Read(A) A=A-50 Write(A) Read (B) B=B+50 Write (B)	

Figure (9.3)
Execution of T2 then T1

When several transactions are executed concurrently , the system may execute one transaction for a little while , then switch to another transaction and execute it for a while then switch back to the first one, and so on as shown in figure (9.4)

T1	T2
Read(A) A=A-50 Write(A)	
	Read (A) temp=A * 0.1 A=A-temp Write (A)
Read (B) B=B+50 Write (B)	
	Read(B) B=B+temp Write(B)

Figure (9.4)
Execution of two transaction

10 : Database Security

10.1 : Introduction

There is a need to secure computer systems, and securing data must be part of an overall computer security plan. Growing amounts of sensitive data are being retained in databases and more of these databases are being made accessible via the Internet. As more data is made available electronically, it can be assumed that threats and vulnerabilities to the integrity of that data will increase as well.

10.2 : Security objective

The primary objectives of database security are:

- Confidentiality : access control
- Integrity : data corruption
- Availability :

To preserve the data confidentiality, enforcing access control policies on the data, these policies are defined on the database management system (DBMS). Access control is to insure that only authorized users perform authorized activities at authorized time.

There are two points of concern in access control:

- Authentication
- Authorization

To preserve data integrity, we must guaranty that the data cannot be corrupted in an invisible way.

Availability property is to ensure timely and reliable access to the database.

10.3 Access control

Access control is the process by which rights and privileges are assigned to users and database objects. Database objects include tables, views, rows and columns. To ensure proper access to the data, authentication and authorization are applied.

Authentication is the process by which you verify that someone is who they claim they are. This usually involves a user name and a password, but can include any other method of demonstrating identity, such as a smart card, voice recognition or fingerprints.

Authentication is equivalent to showing your driver license or your ID.

Authorization is finding out that the person, once identified, is permitted to have the resource. This is usually determined by finding out if that person has paid admission or has a particular level of security clearness.

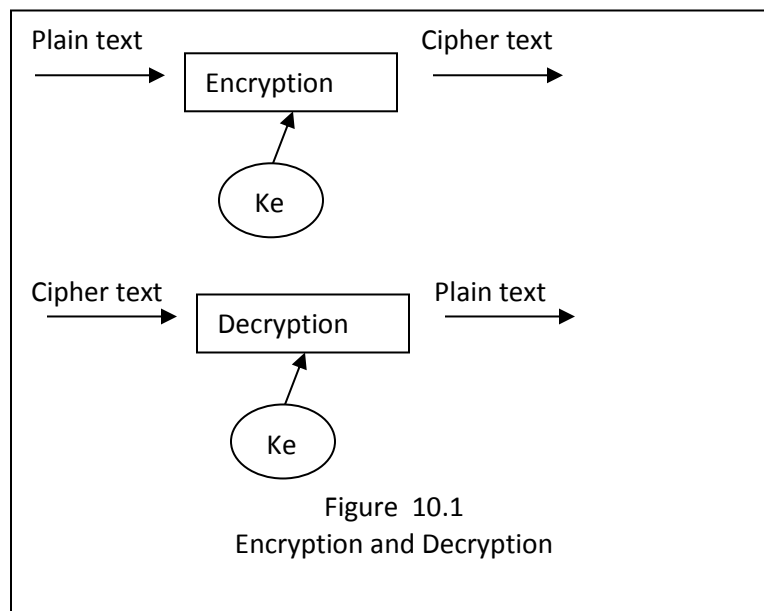
Authorization is equivalent to checking your name in a list of names, or checking your ticket for something.

For Example, student A may be given login rights to the university database with authorization includes read only for the course listing data table.

10.4 Database encryption

Application data security has to deal with several security threats and issues beyond those handled by SQL authorization. For example, data must be protected while they are being transmitted, data may need to be protected from intruders.

Encryption is the process of transferring a clear text (plain text) into disguised text (cipher text) by using a key. Decryption is the process to transfer the cipher text back to plain text as shown in figure 10.1



Database encryption refers to the use of encryption techniques to transform a plain text into encrypted database, thus making it unreadable to anyone except those who possess the knowledge of the encryption.

The purpose of database encryption is to ensure the database opacity by keeping the information hidden to any unauthorized persons. Even if someone gets through and bypasses the access control policies, he/she is still unable to read the encrypted database.

There are a vast number of techniques for the encryption of data. An example for simple encryption techniques is to substitute each character with the next character in the alphabet. ***For example 'Perryridge' becomes 'Qfsszsjehf'.***

A good encryption techniques has the following properties:

- 1- It is relatively simple for authorized users to encrypt and decrypt data.
- 2- It depends not on the security of the algorithm, but rather on a parameter of the algorithm called the encryption key.
- 3- Its encryption key is extremely difficult for an intruder to determine.

10.5 Database encryption level

10.5.1 Application-Level Encryption:

When you encrypt information at the application level, you can protect sensitive data. In many ways the application is the obvious place to encrypt and decrypt data because the application knows exactly which data is sensitive and can apply protection selectively.

You can task a given application with encrypting its own data. This encryption capability is designed into the application itself. By the time the database receives the data, it has already been encrypted and then stored in the database in this encrypted state. As the traffic travels from the application to the database, the data can also be encrypted across the network.

10.5.2 Database Encryption Level:

In this case the information is encrypted in the database. As an example, we'll discuss *Oracle Advanced Security's transparent data encryption (TDE)*, which automatically encrypts and decrypts the data stored in the database and provides this capability without having to write additional code.

With TDE, the encryption process and associated encryption keys are created and managed by the database. This is transparent to database users who have authenticated to the database. At the operating system, however, attempts to access database files return data in an encrypted state. Therefore, for any operating system level users, the data remains inaccessible. Additionally, because the database is doing the encryption, there is no need to change the application(s), and there is a minimal performance overhead when changes occur in the database. TDE is designed into the database itself: Oracle has integrated the TDE syntax with its data definition language (DDL).

If you encrypt on the database, that means the data is sent to and from the database in unencrypted form. This potentially allows for snooping/tampering between the application and the encryption routines on the database.

11 : Fundamental of relational algebra:

Relational algebra consists of a set of operations that takes one or two relations as input and produce a new relation as their result. In the relational algebra, symbols are used to denote an operation.

- For SELECT we use the sigma letter σ . The relation(table name) is written in parentheses:
Select * from Loan where B_name="Perryridge"
will be : $\sigma_{B_name="Perryridge"}(\text{Loan})$

- For projection we use the pi letter π . projection mean select some fields from the table , not all the fields.

If we have : Student (S-Id,S-Name,S-Address)

To display only the names → Select S-Name from student

will be : $\pi_{S_Name}(\text{Student})$

- Selection (σ) and Projection (π) can be used together to select some of the fields with a condition.

To display only the names whose address in Baghdad from student table:

→ Select S-Name from student where S-Address='Baghdad'

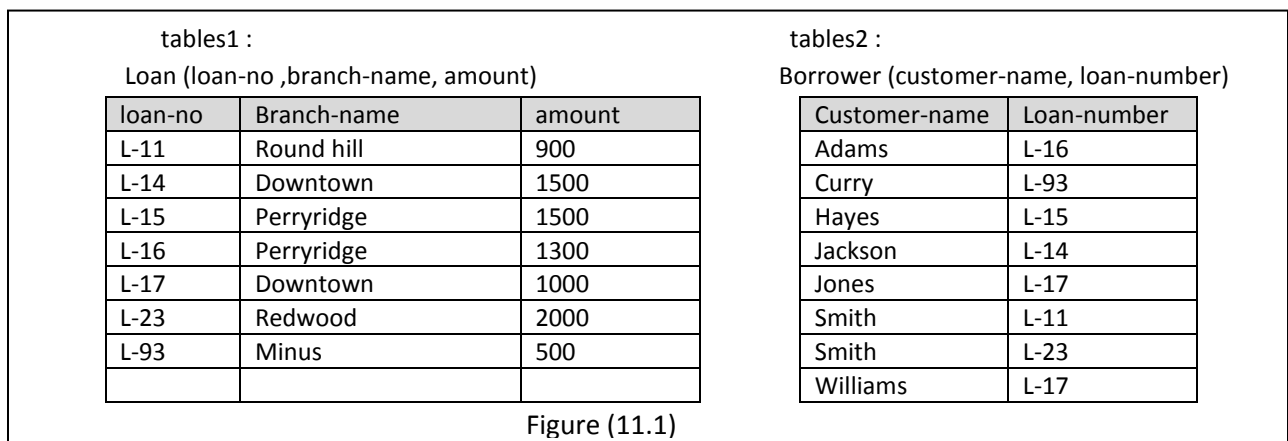
Will be : $\pi_{S_Name}(\sigma_{S_address='Baghdad'}(\text{Student}))$

- Cartesian product (Cross Join) between two tables denoted by X.
It is used to combine information from any two relations. It will produce a tuple from each possible pair of tuples: one from the first table and one from the second.

Ex 1 : To Cross Join between Student and Class ; with the condition only for level 2

→ $\sigma_{LVL=2}(\text{Student X Class})$

Example 2: we have two tables as shown in figure (11.1)



Database Design – Second Course

If we want to know the names of all customers who have a lone at the Perryridge branch. So if we write : $\sigma_{\text{branch-name}='Perry ridge'}(\text{Borrower X Loan})$; the result of this cross join shown in figure (11.2).

customer-name	Borrower. loan-number	Loan. loan-number	branch-name	amount
Adams	L-16	L-15	Perry ridge	1500
Adams	L-16	L-16	Perry ridge	1300
Curry	L-93	L-15	Perry ridge	1500
Curry	L-93	L-16	Perry ridge	1300
Hayes	L-15	L-15	Perry ridge	1500
Hayes	L-15	L-16	Perry ridge	1300
Jackson	L-14	L-15	Perry ridge	1500
Jackson	L-14	L-16	Perry ridge	1300
Jones	L-17	L-15	Perry ridge	1500
Jones	L-17	L-16	Perry ridge	1300
Smith	L-11	L-15	Perry ridge	1500
Smith	L-11	L-16	Perry ridge	1300
Smith	L-23	L-15	Perry ridge	1500
Smith	L-23	L-16	Perry ridge	1300
Williams	L-17	L-15	Perry ridge	1500
Williams	L-17	L-16	Perry ridge	1300

Figure (11.2)
the result of cross join $\sigma_{\text{branch-name}='Perry ridge'}(\text{Borrower X Loan})$

THE RESULT IS NOT RIGHT!!.

The Cross Join links every record from Borrower with all the records of Loan who have Perry ridge in branch-name.

The correct answer will be :

$$\sigma_{\text{borrower. loan-number}=\text{loan.loan-number}} (\sigma_{\text{branch-name}='Perry ridge'}(\text{Borrower X Loan}))$$

Adams	L-16	L-16	Perry ridge	1300
Hayes	L-15	L-15	Perry ridge	1500

- The natural join operation
It allows us to combine certain selections and a Cartesian product into one operation. It is denoted by the join symbol (\bowtie). The natural join operation do the following:
 - Cartesian product of its arguments (ex: two tables)
 - Perform selection forcing equality on those attributes that appear in both tables.
 - Remove duplicate attributes.

Example 11.1: Consider the borrower and loan tables in figure (11. 1), to find the names of all customers who have a loan at the bank, and find the amount of the loan:

$$\Pi_{\text{customer-name,loan-number,amount}}(\text{borrower} \bowtie \text{loan})$$

Because borrower and loan tables both have the attribute loan-number, the natural join operation considers only pairs of tuples from the two tables that have the same value on loan-number. The result will be as shown in figure (11.3) :

Customer-name	Loan-number	amount
Adams	L-16	1300
Curry	L-93	500
Hayes	L-15	1500
Jackson	L-14	1500
Jones	L-17	1000
Smith	L-11	900
Smith	L-23	2000
Williams	L-17	1000

Figure (11.3)
The result of natural join $\Pi_{\text{customer-name,loan-number,amount}}(\text{borrower} \bowtie \text{loan})$

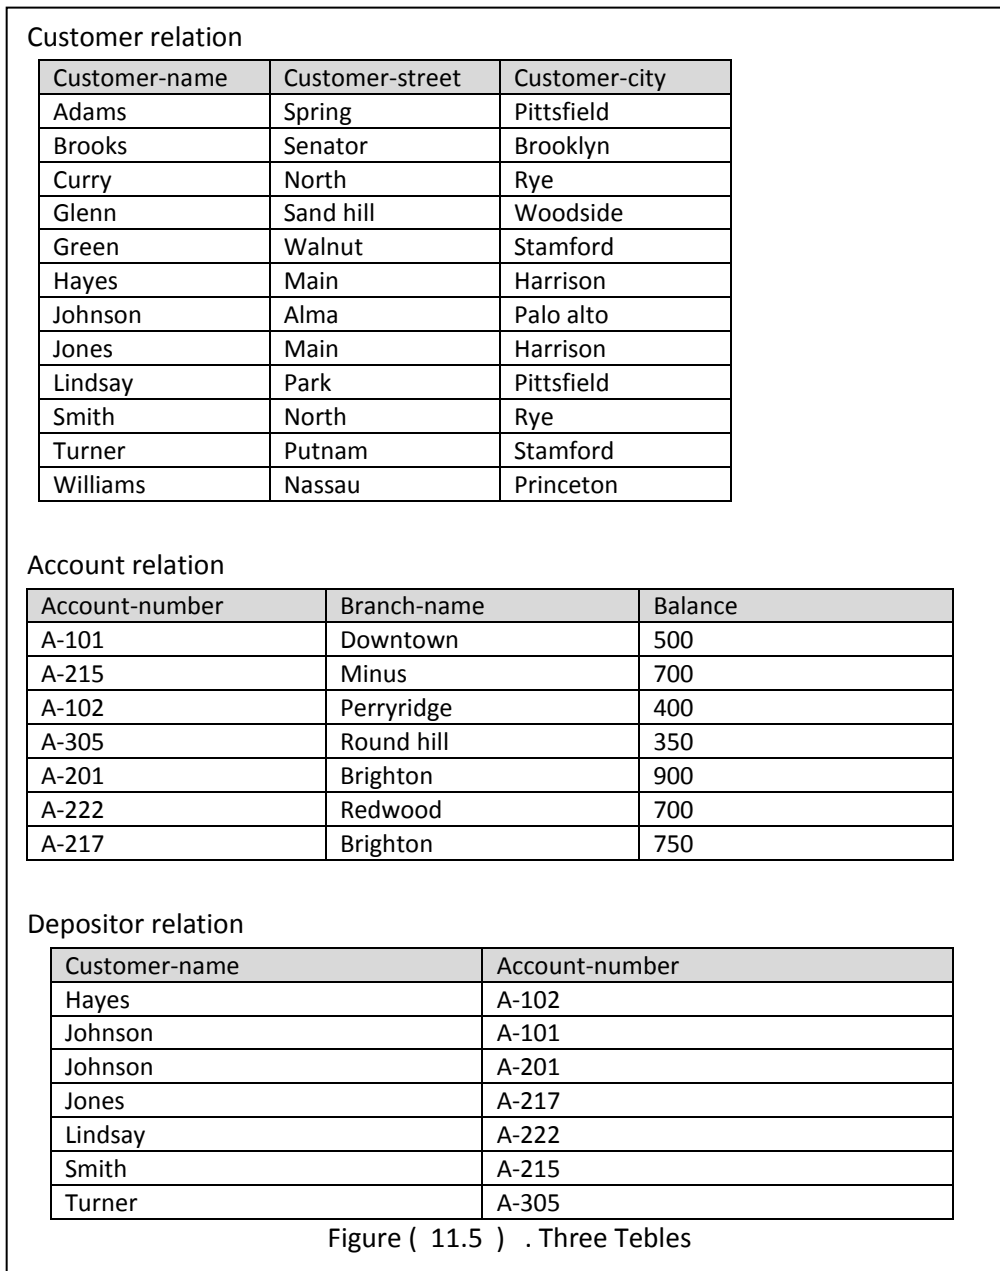
Example 11.2 : Notes: the natural join usually required that the two relations must have at least one common attribute, but if this constraint is omitted, and the two relations have no common attributes, then the natural join becomes exactly the Cartesian product as shown in figure (11.4)

<i>Car</i>		<i>Boat</i>		<i>Car</i> \bowtie <i>Boat</i>			
CarModel	CarPrice	BoatModel	BoatPrice	CarModel	CarPrice	BoatModel	BoatPrice
CarA	20,000	Boat1	10,000	CarA	20,000	Boat1	10,000
CarB	30,000	Boat2	40,000	CarA	20,000	Boat2	40,000
				CarB	30,000	Boat1	10,000
				CarB	30,000	Boat2	40,000

Figure (11.4)

Natural join becomes a cross join because there are no common attribute

Example 11.3 : find the names of all branches with customers who have an account in the bank and who live in Harrison for the relations shown in figure (11.5).



The result will be : $\Pi_{\text{branch-name}}(\sigma_{\text{customer-city}='Harrison'}(\text{customer} \bowtie \text{account} \bowtie \text{depositor}))$

Branch-name
Perryridge
Brighton

12 : Query processing

Query processing refers to the number of activities involved in extracting data from a database.

The steps involved in query processing are :

- 1- Parsing and translation
- 2- Optimization
- 3- Evaluation

12.1- Parsing and translation

Before query processing can begin, the system must translate the query into a usable form. A language such as SQL is suitable for human use , but it is not suitable for the internal representation of the query in the system, thus the query must translated into its internal form, and this is the work of the **PARSER**.

The **PARSER** check for :

- The syntax of the query
- The relation names appearing in the query are exist in the database
- Generate the relational-algebra expression.

12.2- Optimization

The **query optimizer** is the component of a database management system that attempts to determine the most efficient way to execute a query.

The optimizer considers the possible query plans for a given input query, and attempts to determine which of those plans will be the most efficient. [*A **query plan (or query execution plan)** is an ordered set of steps used to access or modify information in a database.*]

Cost-based query optimizers assign an estimated "cost" to each possible query plan, and choose the plan with the smallest cost.

[Costs are used to estimate the runtime cost of evaluating the query, in terms of the number of I/O operations required, the CPU requirements(CPU time to execute a query), the cost of memory used for the query and the cost of communication (in distributed or client-server DB)]

12.2-1 : Equivalent expression

To find the least-costly query evaluation plan, the optimizer generates alternative plan (by generating alternative algebra expression) that produce the same result but with deferent costs.


The rules of equivalence are :

- Rule (1) : Selection operations are commutative

$$\sigma_a(\sigma_b(E)) = \sigma_b(\sigma_a(E))$$


- Rule (2): Natural Join operations are commutative

$$E1 \bowtie E2 = E2 \bowtie E1$$



- Rule(3): Natural join operations are associative


$$(E1 \bowtie E2) \bowtie E3 = E1 \bowtie (E2 \bowtie E3)$$



- Rule(4) :The selection operation distributes over the join operation under the following two condition:

- a- It distribute when all the attributes in selection condition (b) involves only the attributes of one of the expressions (say E1) being joined.

$$\sigma_b (E1 \bowtie E2) = (\sigma_b(E1)) \bowtie E2$$



- b- It distributes when selection condition (b_1) involves only the attributes of E1
And (b_2) involves only the attributes of E2

$$\sigma_{b_1 \wedge b_2} (E1 \bowtie E2) = (\sigma_{b_1} (E1)) \bowtie (\sigma_{b_2} (E2))$$

Example 12.1 : consider the relational algebra

$$\Pi_{\text{customer-name}}(\sigma_{\text{branch-city}='Brooklyn'}(\text{branch} \bowtie (\text{account} \bowtie \text{depositor})))$$

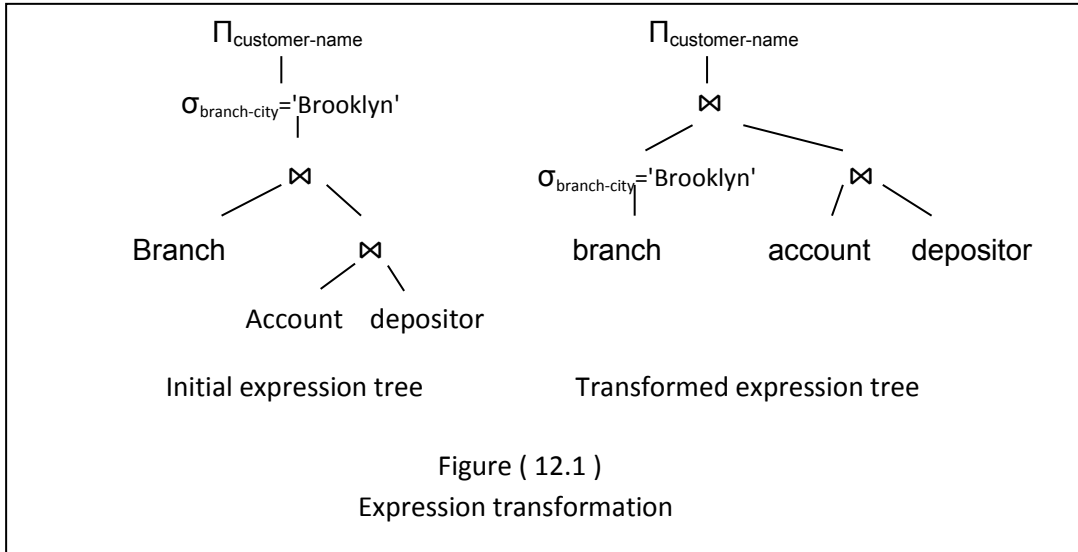
This expression construct a large intermediate relation , $\text{branch} \bowtie \text{account} \bowtie \text{depositor}$.

However we are interesting in only a few tuples of this relation (branch in Brooklyn) and one attribute of the relation (customer-name).

Since we are concerned with only tuples those in Brooklyn in the **branch** relation, we do not need to consider those tuples that do not have $\text{branch-city}='Brooklyn'$ from the **branch** relation.

By reducing the number of tuples of the branch relation that we need to access, we reduce the size of the intermediate result. By using rule (4.a) our query is now represented by the relational expression: (figure 12.1)

$$\Pi_{\text{customer-name}}((\sigma_{\text{branch-city}='Brooklyn'}(\text{branch})) \bowtie (\text{account} \bowtie \text{depositor}))$$



Example12. 2: We have three relations

Branch (branch-name,branch-city,assets)

Account (account-number,branch-name,balance)

Depositor (customer-name,account-number)

To find customer names in Brooklyn and have balance over 1000\$. The relational algebra is:

$$\Pi_{\text{customer-name}}(\sigma_{\text{branch-city}='Brooklyn' \wedge \text{balance}>1000}(\text{branch} \bowtie (\text{account} \bowtie \text{depositor})))$$

We cannot apply rule (4) directly because the condition involves attributes of both the **branch** and **account** relations. We can apply rule (3) to transform the join (branch ⋈ (account ⋈ depositor))

into (branch ⋈ account) ⋈ depositor):

$$\Pi_{\text{customer-name}}(\sigma_{\text{branch-city}='Brooklyn' \wedge \text{balance}>1000}((\text{branch} \bowtie \text{account}) \bowtie \text{depositor}))$$

Then using rule (4.a) to take **depositor relation** out from the condition:

$$\Pi_{\text{customer-name}}((\sigma_{\text{branch-city}='Brooklyn' \wedge \text{balance}>1000}(\text{branch} \bowtie \text{account})) \bowtie \text{depositor})$$

Then using rule (4.b) :

$$\Pi_{\text{customer-name}}(\sigma_{\text{branch-city}='Brooklyn'}(\text{branch} \bowtie \sigma_{\text{balance}>1000}(\text{account})) \bowtie \text{depositor})$$

12.2-2 : Disk I/O cost

In the database, the cost to access data from disk is important, since disk accesses are slow compared to in memory operation. Disk access measured by taking into account:

- * Number of disk seeks (average-seek-cost)
- * Number of blocks transfers from disk (average-block-read-cost)

If the disk subsystem takes an average of

tT – seconds to transfer one block of data

tS – seconds for one seek (block access time)

then the operation of transfers N blocks and performs S seeks would take :

$$N * tT + S * tS \quad \text{seconds}$$

When calculating disk I/O cost, some system need one seek to transfer many block. For example if there are 10 blocks need to be transfer from disk to memory with one seek , then the time will be $tS + 10 * tT$.

12.2-3 :Projection Example

Projections produce a result tuple for every argument tuple. Change in the output size is the change in the length of tuples .

Let's take a relation 'R' : R(a, b, c), the number of tuples in this relation are (20,000 tuples).

Each Tuple (190 bytes size) : header = 24 bytes, a = 8 bytes, b = 8 bytes, c = 150 bytes.

Each Block (1024): header = 24 bytes

We can fit 5 tuples into 1 block

- 5 tuples * 190 bytes(size of the tuple) = 950 bytes can fit into 1 block
- For 20,000 tuples, we would require **4,000** blocks (20,000 / 5 tuples per block)

With a projection resulting in elimination of column c (150 bytes), we could estimate that each tuple would decrease to 40 bytes (190 – 150 bytes)

Now, the new estimate will be 25 tuples in 1 block. (25 tuples * 40 byte= 1000)

- 25 tuples * 40 bytes/tuple = 1000 bytes will be able to fit into 1 block
- With 20,000 tuples, the new estimate is 800 blocks (20,000 tuples / 25 tuples per block = **800** blocks)

Result is reduction by a factor of 5

12.3- Evaluation

Query evaluation is the process of executing the plan for that query and return the result to query.

The **query-execution engine** is the subsystem of the DBMS that execute the query plan.