

ANNAMACHARYA UNIVERSITY:: RAJAMPET

DEPARTMENT OF ARTIFICIAL INTELLIGENCE AND DATA SCIENCE

LECTURE NOTES



Name of the Faculty: M.NANDINI

Class: II B.Tech II Sem.

Branch: AI&DS

Name of the Course: DATABASE MANAGEMENT SYSTEMS

Subject Code: 24ACSE41T

Academic Year: 2025-2026

UNIT-1:INTRODUCTION TO DBMS AND BASICS OF SQL

INTRODUCTION:Database-System Applications, Characteristics, Purpose Of Database Systems, View Of Data, Database Languages, Data Storage And Querying, Transaction Management, Data Base Architecture, Database Users And Administrators.

Basic Sql: Simple Database Schema , Data Types, Table Definitions(Create, Alter), Diferent DML Operations(Insert,Delete,Update)

Introduction to Database Management System

As the name suggests, the database management system consists of two parts. They are:

1. Database and
2. Management System

What is a Database?

To find out what database is, we have to start from data, which is the basic building block of any DBMS.

Data: Facts, figures, statistics etc. having no particular meaning (e.g. 1, ABC, 19 etc).

Record: Collection of related data items, e.g. in the above example the three data items had no meaning. But if we organize them in the following way, then they collectively represent meaningful information.

Roll	Name	Age
1	ABC	19

Table or Relation: Collection of related records.

Roll	Name	Age
1	ABC	19
2	DEF	22
3	XYZ	28

The columns of this relation are called **Fields, Attributes or Domains**. The rows are called **Tuples or Records**.

Database: Collection of related relations. Consider the following collection of tables:

Roll	Name	Age
1	ABC	19
2	DEF	22
3	XYZ	28

Year	Hostel
I	H1
II	H2

T1

Roll	Address
1	KOL
2	DEL
3	MUM

T2

T3

T4

Roll	Year
1	I
2	II
3	I

Now we have a collection of 4 tables. They can be called a “related collection” because we can clearly find out that there are some common attributes existing in a selected pair of tables. Because of these common attributes we may combine the data of two or more tables together to find

DEPARTMENT OF AIDS

DBMS

M. NANDINI

out the complete details of a student. Questions like “Which hostel does the youngest student live in?” can be answered now, although *Age* and *Hostel* attributes are in different tables.

A database in a DBMS could be viewed by lots of different people with different responsibilities.

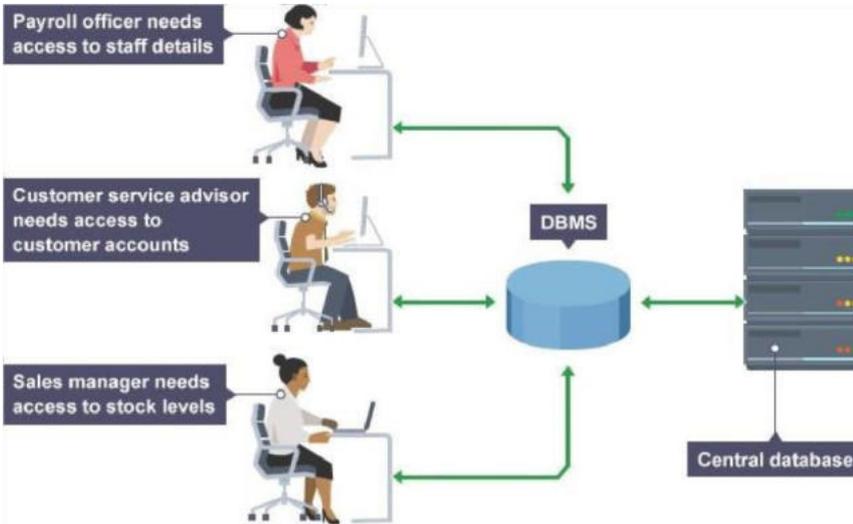


Figure 1.1: Employees are accessing data through DBMS

For example, within a company there are different departments, as well as customers, who each need to see different kinds of data. Each employee in the company will have different levels of access to the database with their own customized **front-end** application.

In a database, data is organized strictly in row and column format. The rows are called **Tuple** or **Record**. The data items within one row may belong to different data types. On the other hand, the columns are often called **Domain** or **Attribute**. All the data items within a single attribute are of the same data type.

What is a Management System?

A **database-management system (DBMS)** is a collection of interrelated data and a set of programs to access that data. This is a collection of related data with an implicit meaning and hence is a database. The collection of data, usually referred to as the **database**, contains information relevant to an enterprise. The primary goal of a DBMS is to provide a way to store and retrieve database information that is both *convenient* and *efficient*. By **data**, we mean known facts that can be recorded and that have implicit meaning.

The management system is important because without the existence of some kind of rules and regulations it is not possible to maintain the database. We have to select the particular attributes which should be included in a particular table; the common attributes to create a relationship between two tables; if a new record has to be inserted or deleted then which tables should have to be handled etc. These issues must be resolved by having some kind of rules to follow in order to maintain the integrity of the database.

Database systems are designed to manage large bodies of information. Management of data involves both defining structures for storage of information and providing mechanisms for the manipulation of information. In addition, the database system must ensure the safety of the information stored, despite system crashes or attempts at unauthorized access. If data are to be shared among several users, the system must avoid possible anomalous results.

Because information is so important in most organizations, computer scientists have developed a large body of concepts and techniques for managing data. These concepts and techniques form the focus of this book. This chapter briefly introduces the principles of database systems.

Database Management System (DBMS) and Its Applications:

A Database management system is a computerized record-keeping system. It is a repository or a container for collection of computerized data files. The overall purpose of DBMS is to allow the users to define, store, retrieve and update the information contained in the database on demand. Information can be anything that is of significance to an individual or organization.

Databases touch all aspects of our lives. Some of the major areas of application are as follows:

1. Banking
2. Airlines
3. Universities
4. Manufacturing and selling

5. Human resources

Enterprise Information

- *Sales*: For customer, product, and purchase information.
- *Accounting*: For payments, receipts, account balances, assets and other accounting information.
- *Human resources*: For information about employees, salaries, payroll taxes, and benefits, and for generation of paychecks.
- *Manufacturing*: For management of the supply chain and for tracking production of items in factories, inventories of items in warehouses and stores, and orders for items.
- *Online retailers*: For sales data noted above plus online order tracking, generation of recommendation lists, and maintenance of online product evaluations.

Banking and Finance

- *Banking*: For customer information, accounts, loans, and banking transactions.
- *Credit card transactions*: For purchases on credit cards and generation of monthly statements.
- *Finance*: For storing information about holdings, sales, and purchases of financial instruments such as stocks and bonds; also for storing real-time market data to enable online trading by customers and automated trading by the firm.
- *Universities*: For student information, course registrations, and grades (in addition to standard enterprise information such as human resources and accounting).
- *Airlines*: For reservations and schedule information. Airlines were among the first to use databases in a geographically distributed manner.
- *Telecommunication*: For keeping records of calls made, generating monthly bills, maintaining balances on prepaid calling cards, and storing information about the communication networks.

Purpose of Database Systems

Database systems arose in response to early methods of computerized management of commercial data. As an example of such methods, typical of the 1960s, consider part of a university organization that, among other data, keeps information about all instructors, students, departments, and course offerings. One way to keep the information on a computer is to store it in operating system files. To allow users to manipulate the information, the system has a number of application programs that manipulate the files, including programs to:

- ✓ Add new students, instructors, and courses
 - ✓ Register students for courses and generate class rosters
 - ✓ Assign grades to students, compute grade point averages (GPA), and generate transcripts
- System programmers wrote these application programs to meet the needs of the university.

New application programs are added to the system as the need arises. For example, suppose that a university decides to create a new major (say, computer science). As a result, the university creates a new department and creates new permanent files (or adds information to existing files) to record information about all the instructors in the department, students in that major, course offerings, degree requirements, etc. The university may have to write new application programs to deal with rules specific to the new major. New application programs may also have to be written to handle new rules in the university. Thus, as time goes by, the system acquires more files and more application programs.

This typical **file-processing system** is supported by a conventional operating system. The system stores permanent records in various files, and it needs different application programs to extract records from, and add records to, the appropriate files. Before database management systems (DBMSs) were introduced, organizations usually stored information in such systems. Keeping organizational information in a file processing system has a number of major disadvantages:

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 database 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

DBMS

M. NANDINI

\$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.

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.

Advantages of DBMS:

Controlling of Redundancy: Data redundancy refers to the duplication of data (i.e. storing same data multiple times). In a database system, by having a centralized database and centralized control of data by the DBA the unnecessary duplication of data is avoided. It also eliminates the extra time for processing the large volume of data. It results in saving the storage space.

Improved Data Sharing: DBMS allows a user to share the data in any number of application programs.

Data Integrity: Integrity means that the data in the database is accurate. Centralized control of the data helps in permitting the administrator to define integrity constraints to the data in the database. For example: in customer database we can enforce an integrity that it must accept the customer only from Noida and Meerut city.

Security : Having complete authority over the operational data, enables the DBA in ensuring that the only mean of access to the database is through proper channels. The DBA can define authorization checks to be carried out whenever access to sensitive data is attempted.

Data Consistency : By eliminating data redundancy, we greatly reduce the opportunities for inconsistency. For example: if a customer address is stored only once, we cannot have disagreement on the stored values. Also updating data values is greatly simplified when each value is stored in one place only. Finally, we avoid the wasted storage that results from redundant data storage.

Efficient Data Access: In a database system, the data is managed by the DBMS and all access to the data is through the DBMS providing a key to effective data processing

Enforcement of Standards: With the centralized data, DBA can establish and enforce the data standards which may include the naming conventions, data quality standards etc.

Data Independence : In a database system, the database management system provides the interface between the application programs and the data. When changes are made to the data representation, the meta data obtained by the DBMS is changed but the DBMS continues to provide the data to application program in the previously used way. The DBMS handles the task of transformation of data wherever necessary.

Reduced Application Development and Maintenance Time : DBMS supports many important functions that are common to many applications, accessing data stored in the DBMS, which facilitates the quick development of application.

Disadvantages of DBMS

- 1) It is bit complex. Since it supports multiple functionality to give the user the best, the underlying software has become complex. The designers and developers should have thorough knowledge about the software to get the most out of it.
- 2) Because of its complexity and functionality, it uses large amount of memory. It also needs large memory to run efficiently.
- 3) DBMS system works on the centralized system, i.e.; all the users from all over the world access this database. Hence any failure of the DBMS, will impact all the users.
- 4) DBMS is generalized software, i.e.; it is written work on the entire systems rather specific one. Hence some of the application will run slow.

View of Data



A database system is a collection of interrelated data and a set of programs that allow users to access and modify this 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.

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:

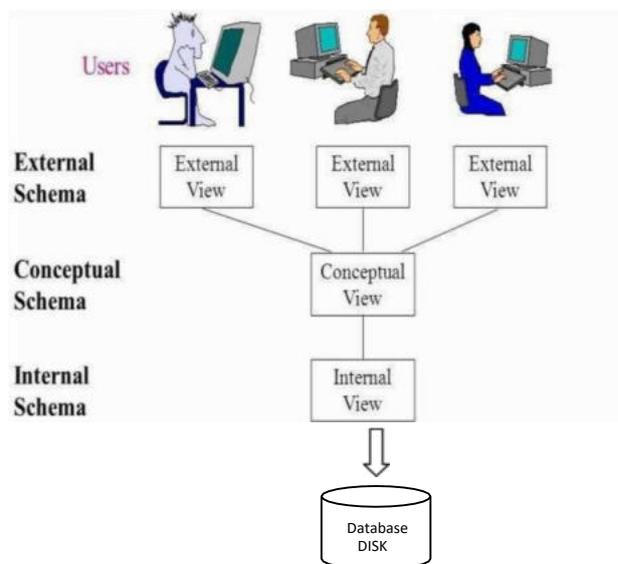


Figure 1.2: Level of Abstraction in a DBMS

- **Physical level (or Internal View / Schema):** 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 (or Conceptual View / Schema):** 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 (or External View / Schema):** 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 1.2 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 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



DBMS

M. NANDINI

name and 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.

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.

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**, which 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.

Data Models



M. NANDINI

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.

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.

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.

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.

Semi-structured Data Model. The semi-structured 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 semi-structured data.

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.

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 updates. In practice, the data definition and data-manipulation languages are not two separate languages; instead they simply form parts of a single database language, such as the widely used SQL language.

Data-Manipulation Language

A **data-manipulation language (DML)** is a language that enables users to access or manipulate data as organized by the appropriate data model. The types of access are:

- Retrieval of information stored in the database
- Insertion of new information into the database
- Deletion of information from the database
- Modification of information stored in the database



There are basically two types:

- **Procedural DMLs** require a user to specify *what* data are needed and *how* to get those data.
- **Declarative DMLs** (also referred to as **nonprocedural DMLs**) require a user to specify *what* data are needed *without* specifying how to get those data.

Declarative DMLs are usually easier to learn and use than are procedural DMLs. However, since a user does not have to specify how to get the data, the database system has to figure out an efficient means of accessing data. A **query** is a statement requesting the retrieval of information. The portion of a DML that involves information retrieval is called a **query language**. Although technically incorrect, it is common practice to use the terms *query language* and *data-manipulation language* synonymously.

Data-Definition Language (DDL)

We specify a database schema by a set of definitions expressed by a special language called a **data-definition language (DDL)**. The DDL is also used to specify additional properties of the data.

We specify the storage structure and access methods used by the database system by a set of statements in a special type of DDL called a **data storage and definition language**. These statements define the implementation details of the database schemas, which are usually hidden from the users.

The data values stored in the database must satisfy certain **consistency constraints**.

For example, suppose the university requires that the account balance of a department must never be negative. The DDL provides facilities to specify such constraints. The database system checks these constraints every time the database is updated. In general, a constraint can be an arbitrary predicate pertaining to the database. However, arbitrary predicates may be costly to test. Thus, database systems implement integrity constraints that can be tested with minimal overhead.

- **Domain Constraints.** A domain of possible values must be associated with every attribute (for example, integer types, character types, date/time types). Declaring an attribute to be of a particular domain acts as a constraint on the values that it can take. Domain constraints are the most elementary form of integrity constraint. They are tested easily by the system whenever a new data item is entered into the database.
- **Referential Integrity.** There are cases where we wish to ensure that a value that appears in one relation for a given set of attributes also appears in a certain set of attributes in another relation (referential integrity). For example, the department listed for each course must be one that actually exists. More precisely, the *deptname* value in a *course* record must appear in the *deptname* attribute of some record of the *department* relation. Database modifications can cause violations of referential integrity. When a referential-integrity constraint is violated, the normal procedure is to reject the action that caused the violation.
- **Assertions.** An assertion is any condition that the database must always satisfy. Domain constraints and referential-integrity constraints are special forms of assertions. However, there are many constraints that we cannot express by using only these special forms. For example, “Every department must have at least five courses offered every semester” must be expressed as an assertion. When an assertion is created, the system tests it for validity. If the assertion is valid, then any future modification to the database is allowed only if it does not cause that assertion to be violated.
- **Authorization.** We may want to differentiate among the users as far as the type of access they are permitted on various data values in the database. These differentiations are expressed in terms of **authorization**, the most common being: **read authorization**, which allows reading, but not modification, of data; **insert authorization**, which allows insertion of new



M. NANDINI

data, but not modification of existing data; **update authorization**, which allows modification, but not deletion, of data; and **delete authorization**, which allows deletion of data. We may assign the user all, none, or a combination of these types of authorization.

The DDL, just like any other programming language, gets as input some instructions (statements) and generates some output. The output of the DDL is placed in the **data dictionary**, which contains **metadata**—that is, data about data. The data dictionary is considered to be a special type of table that can only be accessed and updated by the database system itself (not a regular user). The database system consults the data dictionary before reading or modifying actual data.

Data Dictionary

We can define a data dictionary as a DBMS component that stores the definition of data characteristics and relationships. You may recall that such “data about data” were labeled metadata. The DBMS data dictionary provides the DBMS with its self-describing characteristic. In effect, the data dictionary resembles an X-ray of the company’s entire dataset, and is a crucial element in the data administration function.

Two main types of data dictionary exist, integrated and stand-alone. An integrated data dictionary is included with the DBMS. For example, all relational DBMSs include a built-in data dictionary or system catalog that is frequently accessed and updated by the RDBMS. Other DBMSs, especially older types, do not have a built-in data dictionary; instead, the DBA may use third-party stand-alone data dictionary systems. Data dictionaries can also be classified as active or passive. An active data dictionary is automatically updated by the DBMS with every database access, thereby keeping its access information up-to-date. A passive data dictionary is not updated automatically and usually requires a batch process to be run. Data dictionary access information is normally used by the DBMS for query optimization purposes.

The data dictionary’s main function is to store the description of all objects that interact with the database. Integrated data dictionaries tend to limit their metadata to the data managed by the DBMS. Stand-alone data dictionary systems are more usually more flexible and allow the

DBA to describe and manage all the organization’s data, whether or not they are computerized. Whatever the data dictionary’s format, its existence provides database designers and end users with a much-improved ability to communicate. In addition, the data dictionary is the tool that helps the DBA to resolve data conflicts.

Although there is no standard format for the information stored in the data dictionary, several features are common. For example, the data dictionary typically stores descriptions of all:

- Data elements that are defined in all tables of all databases. Specifically, the data dictionary stores the name, data types, display formats, internal storage formats, and validation rules. The data dictionary tells where an element is used, by whom it is used, and so on.
- Tables defined in all databases. For example, the data dictionary is likely to store the name of the table creator, the date of creation, access authorizations, the number of columns, and so on.
- Indexes defined for each database table. For each index, the DBMS stores at least the index name, the attributes used, the location, specific index characteristics, and the creation date.
- Database definitions: who created each database, the date of creation, where the database is located, who the DBA is, and so on.
- End users and the administrators of the database.
- Programs that access the database, including screen formats, report formats, application formats, SQL queries, and so on.
- Access authorization for all users of all databases.



- Relationships among data elements which elements are involved: whether the relationship are mandatory or optional, the connectivity and cardinality and so on.

Database Administrators and Database Users

A primary goal of a database system is to retrieve information from and store new information in the database. People who work with a database can be categorized as database users or database administrators.

Database Users and User Interfaces

There are four different types of database-system users, differentiated by the way they expect to interact with the system. Different types of user interfaces have been designed for the different types of users. **Naive users** are unsophisticated users who interact with the system by invoking one of the application programs that have been written previously. For example, a bank teller who needs to transfer \$50 from account *A* to account *B* invokes a program called *transfer*. This program asks the teller for the amount of money to be transferred, the account from which the money is to be transferred, and the account to which the money is to be transferred.

As another example, consider a user who wishes to find her account balance over the World Wide Web. Such a user may access a form, where she enters her account number. An application program at the Web server then retrieves the account balance, using the given account number, and passes this information back to the user. The typical user interface for naive users is a forms interface, where the user can fill in appropriate fields of the form. Naive users may also simply read *reports* generated from 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 without writing a program. There are also special types of programming languages that combine imperative control structures (for example, for loops, while loops and if-then-else statements) with statements of the data manipulation language. These languages, sometimes called *fourth-generation languages*, often include special features to facilitate the generation of forms and the display of data on the screen. Most major commercial database systems include a fourth generation language.

Sophisticated users interact with the system without writing programs. Instead, they form their requests in a database query language. They submit each such query to a **query processor**, whose function is to break down DML statements into instructions that the storage manager understands. Analysts who submit queries to explore data in the database fall in this category.

Online analytical processing (OLAP) tools simplify analysts' tasks by letting them view summaries of data in different ways. For instance, an analyst can see total sales by region (for example, North, South, East, and West), or by product, or by a combination of region and product (that is, total sales of each product in each region). The tools also permit the analyst to select specific regions, look at data in more detail (for example, sales by city within a region) or look at the data in less detail (for example, aggregate products together by category). Another class of tools for analysts is **data mining** tools, which help them find certain kinds of patterns in data. **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, knowledge base and expert systems, systems that store data with complex data types (for example, graphics data and audio data), and environment-modeling systems.



Database Architecture:

We are now in a position to provide a single picture (Figure 1.3) of the various components of a database system and the connections among them.

The architecture of a database system is greatly influenced by the underlying computer system on which the database system runs. Database systems can be centralized, or client-server, where one server machine executes work on behalf of multiple client machines. Database systems can also be designed to exploit parallel computer architectures. Distributed databases span multiple geographically separated machines.

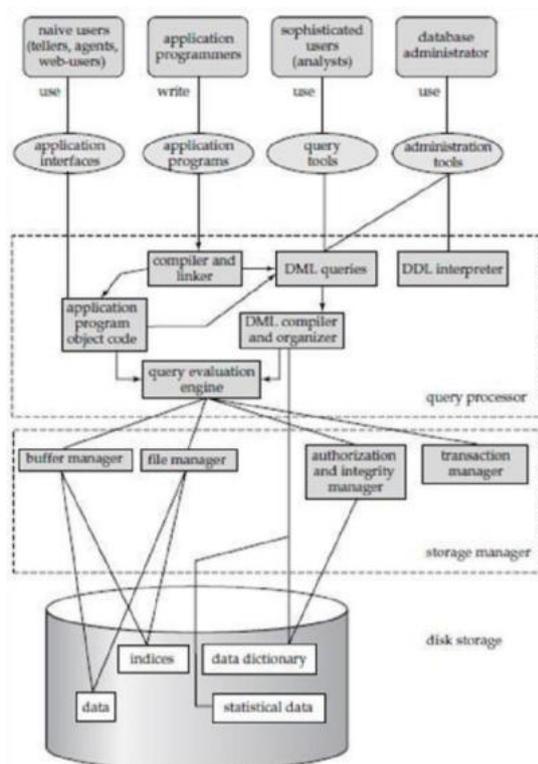


Figure 1.3: Database System Architecture

A database system is partitioned into modules that deal with each of the responsibilities of the overall system. The functional components of a database system can be broadly divided into the **storage manager** and the **query processor** components. The storage manager is important because databases typically require a large amount of storage space. The query processor is important because it helps the database system simplify and facilitate access to data.

It is the job of the database system to translate updates and queries written in a non-procedural language, at the logical level, into an efficient sequence of operations at the physical level.

Database applications are usually partitioned into two or three parts, as in Figure 1.4. In a two-tier architecture, the application resides at the client machine, where it invokes database system functionality at the server machine through query language statements. Application program interface standards like ODBC and JDBC are used for interaction between the client and the server. In contrast, in a three-tier architecture, the client machine acts as merely a front end

and does not contain any direct database calls. Instead, the client end communicates with an application server, usually through a forms interface.

The application server in turn communicates with a database system to access data. The business logic of the application, which says what actions to carry out under what conditions, is embedded in the application server, instead of being distributed across multiple clients. Three-tier applications are more appropriate for large applications, and for applications that run on the World Wide Web.

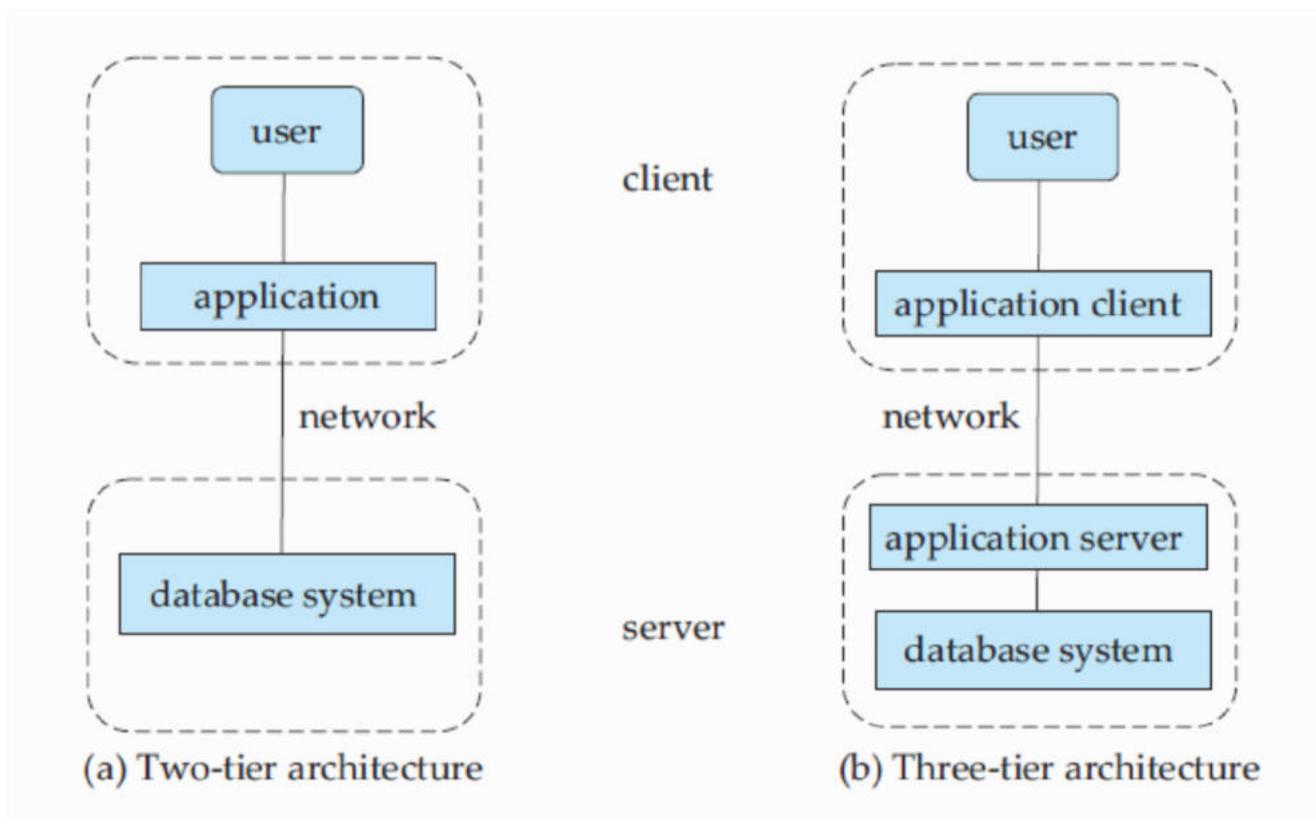


Figure 1.4: Two-tier and three-tier architectures.

Query Processor:

The query processor components include

- **DDL interpreter**, which interprets DDL statements and records the definitions in the data dictionary.
- **DML compiler**, which translates DML statements in a query language into an evaluation plan consisting of low-level instructions that the query evaluation engine understands.

A query can usually be translated into any of a number of alternative evaluation plans that all give the same result. The DML compiler also performs **query optimization**, that is, it picks the lowest cost evaluation plan from among the alternatives.

Query evaluation engine, which executes low-level instructions generated by the DML compiler.

Storage Manager:



A *storage manager* is a program module that provides the interface between the low-level data stored in the database and the application programs and queries submitted to the system. The storage manager is responsible for the interaction with the file manager. The raw data are stored on the disk using the file system, which is usually provided by a conventional operating system. The storage manager translates the various DML statements into low-level file-system commands. Thus, the storage manager is responsible for storing, retrieving, and updating data in the database. The storage manager components include:

- **Authorization and integrity manager**, which tests for the satisfaction of integrity constraints and checks the authority of users to access data.
- **Transaction manager**, which ensures that the database remains in a consistent (correct) state despite system failures, and that concurrent transaction executions proceed without conflicting.
- **File manager**, which manages the allocation of space on disk storage and the data structures used to represent information stored on disk.
- **Buffer manager**, which is responsible for fetching data from disk storage into main memory, and deciding what data to cache in main memory. The buffer manager is a critical part of the database system, since it enables the database to handle data sizes that are much larger than the size of main memory.

Transaction Manager:

A **transaction** is a collection of operations that performs a single logical function in a database application. Each transaction is a unit of both atomicity and consistency. Thus, we require that transactions do not violate any database-consistency constraints. That is, if the database was consistent when a transaction started, the database must be consistent when the transaction successfully terminates. **Transaction - manager** ensures that the database remains in a consistent (correct) state despite system failures (e.g., power failures and operating system crashes) and transaction failures.

A database schema is the structural blueprint of a database, which includes definitions for tables, columns, constraints (like primary and foreign keys), and the **datatypes** assigned to each column. Datatypes define the kind of values a column can hold (e.g., integer, text, date).

Common data types generally fall into these main categories, although specific names and sizes vary between different Database Management Systems (DBMS) like MySQL, PostgreSQL, or SQL Server:

Common Data Type Categories

DBMS

M. NANDINI

□ **Numeric Types:** Used for storing numbers.

○ **INT** (or **INTEGER**): For whole numbers (integers).

○ **DECIMAL(p, s)** or **NUMERIC(p, s)**: For exact precision numbers, such as currency or financial data, where 'p' is the total digits and 's' is the digits after the decimal point.

○ **FLOAT** or **REAL**: For floating-point (approximate) numbers that contain decimals.

○ **BOOLEAN** (**BIT**): For logical values, typically **TRUE** or **FALSE** (or 1 or 0).

□ **String/Character Types:** Used for storing text, letters, numbers, and symbols.

○ **CHAR(size)**: A fixed-length string. It always uses the full specified length, which is efficient for data of a consistent size (e.g., a 2-character state abbreviation).

○ **VARCHAR(size)**: A variable-length string. It only uses the storage needed for the actual data entered, plus

○ a small overhead.

length string. It only uses the storage needed for the actual data entered, plus

○ **TEXT** or **MEMO**: For storing large amounts of text data.

□ **Date and Time Types:** Used for working with dates and

○ **DATE**: Stores a date only.

○ **TIME**: Stores a time only.

○ **DATETIME** or **TIMESTAMP**: Stores both date and time, often used for recording when a row was created or modified.

□ **Binary Types:** Used for storing binary data such as images, audio, or other files.

○ **BLOB** (Binary Large Object) or **BINARY**: Stores raw binary byte strings.

Importance in Schema Design

Choosing the correct data type is crucial during database design (data modeling) for several reasons:

DBMS

M. NANDINI

- **Data Integrity:** Ensures only valid data types can be entered into a column (e.g., preventing text from being stored in a numeric age field).
- **Storage Efficiency:** Using the appropriate data type helps manage disk space effectively.
- **Performance:** Proper data types and indexing improve the speed and efficiency of data retrieval (queries).

DDL AND DML COMMANDS:

DDL (Data Definition Language) commands are used to define or modify the database structure, while DML (Data Manipulation Language) commands are used to manage the data within those structures.

DDL Commands (Data Definition Language)

DDL commands affect the schema or structure of the database and are auto-committed, meaning the changes are permanent and cannot be rolled back.

- **CREATE:** Used to create new database objects, such as tables, indexes, views, or the database itself.

o `CREATE TABLE Students (StudentID int, Name varchar(255));` *Example:*

- **ALTER:** Used to modify the structure of an existing database object,

like adding, deleting, or modifying columns in a table. o `ALTER TABLE Students ADD Email varchar(255);` *Example:*

- **DROP:** Used to permanently delete a database object (e.g., a table, index, or database) and all the data within it.

o *Example:* `DROP TABLE Students;`

- **TRUNCATE:** Used to quickly delete all records from a table but keeps the table structure intact for future use. It is a DDL command because, like DROP, its operations are non-transactional and cannot be rolled back.

o *Example:* `TRUNCATE TABLE Students;`

- **RENAME:** Used to change the name of a database object.

o `RENAME TABLE Students TO Student_Data;` *Example:*

DML Commands (Data Manipulation Language)

DBMS

M. NANDINI

DML commands are used to manage the data stored within the database objects. These commands are transactional, meaning changes can be rolled back (undone) if necessary before being committed permanently.

INSERT: Used to add new rows (records) of data into a table.

σ `INSERT INTO Students (StudentID, Name) VALUES (1, 'John Doe');` *Example:*

□: Used to **UPDATE** modify existing data within a table.

o *Example:*

```
UPDATE Students SET Name = 'Jane Doe' WHERE StudentID = 1;
```

□: **DELETE** Used to remove existing records (rows) from a table.

o `DELETE FROM Students WHERE StudentID = 1;` *Example:*

– **SELECT:** While often categorized separately as Data

Query Language (DQL), SELECT is widely considered part of DML as it interacts with the data within the tables.

o *Example:* `SELECT * FROM Students;`

UNIT II: DATABASE DESIGN

Database Design: Database Design And Er Diagrams, Entities, Attributes And Entity Sets, Relationships And Relationship Sets, Additional Features Of The Er Model, Conceptual Design With The Er Model, Case Study: The Internet Shop.

The Relational Model: Introduction To The Relational Model, Integrity Constraints Over Relations, Enforcing Integrity Constraints, Querying Relational Data, Logical Data Base Design: Er To Relational.

What is ER Modeling?

A graphical technique for understanding and organizing the data independent of the actual database implementation

We need to be familiar with the following terms to go further.

Entity

Anything that has an independent existence and about which we collect data. It is also known as entity type.

In ER modeling, notation for entity is given below.



Entity instance

Entity instance is a particular member of the entity type. Example for entity instance : A particular employee **Regular Entity**

An entity which has its own key attribute is a regular entity. Example for regular entity : Employee. **Weak entity**

An entity which depends on other entity for its existence and doesn't have any key attribute of its own is a weak entity.

Example for a weak entity: In a parent/child relationship, a parent is considered as a strong entity and the child is a weak entity.

In ER modeling, notation for weak entity is given below.



Attributes

Properties/characteristics which describe entities are called attributes. In ER modeling, notation for attribute is given below.



Domain of Attributes

This set of possible values that an attribute can take is called the domain of the attribute. For example, the attribute day may take any value from the set {Monday, Tuesday...Friday}. Hence this set can be termed as the domain of the attribute day. **Key attribute**

The attribute (or combination of attributes) which is unique for every entity instance is called key attribute.

DBMS

M. NANDINI

E.g. the employee_id of an employee, pan_card_number of a person etc. If the key attribute consists of two or more attributes in combination, it is called a composite key.

In ER modeling, notation for key attribute is given below.



Simple attribute

If an attribute cannot be divided into simpler components, it is a simple attribute.

Example for simple attribute: employee_id of an employee.

Composite attribute

If an attribute can be split into components, it is called a composite attribute.

Example for composite attribute : Name of the employee which can be split into First_name, Middle_name, and Last_name.

Single valued Attributes

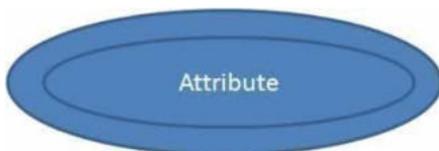
If an attribute can take only a single value for each entity instance, it is a single valued attribute.

example for single valued attribute: age of a student. It can take only one value for a particular student.

Multi-valued Attributes

If an attribute can take more than one value for each entity instance, it is a multi-valued attribute. Multi-valued example for multi valued attribute : telephone number of an employee, a particular employee may have multiple telephone numbers.

In ER modeling, notation for multi-valued attribute is given below.



Stored Attribute

An attribute which needs to be stored permanently is a stored attribute. Example for stored attribute : name of a student

Derived Attribute

An attribute which can be calculated or derived based on other attributes is a derived attribute.

Example for derived attribute: age of an employee which can be calculated from date of birth and current date.

In ER modeling, notation for derived attribute is given below.

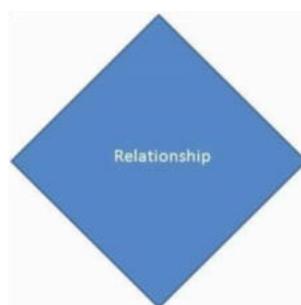


Relationships

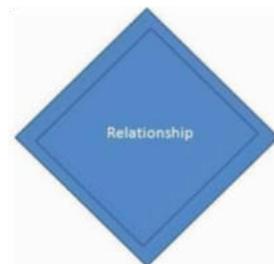
Associations between entities are called relationships

Example: An employee works for an organization. Here "works for" is a relation between the entities employee and organization.

In ER modeling, notation for relationship is given below.



However in ER Modeling, To connect a weak Entity with others, you should use a weak relationship notation as given below



Degree of a Relationship

Degree of a relationship is the number of entity types involved. The n-ary relationship is the general form for degree. Special cases are unary, binary, and ternary, where the degree is 1, 2, and 3, respectively.

Example for unary relationship : An employee is a manager of another employee

Example for binary relationship : An employee works-for department.

Example for ternary relationship: customer purchase item from a shopkeeper

Cardinality of a Relationship

Relationship cardinalities specify how many of each entity type is allowed. Relationships can have four possible connectivities as given below.

1. One to one (1:1) relationship
2. One to many (1:N) relationship

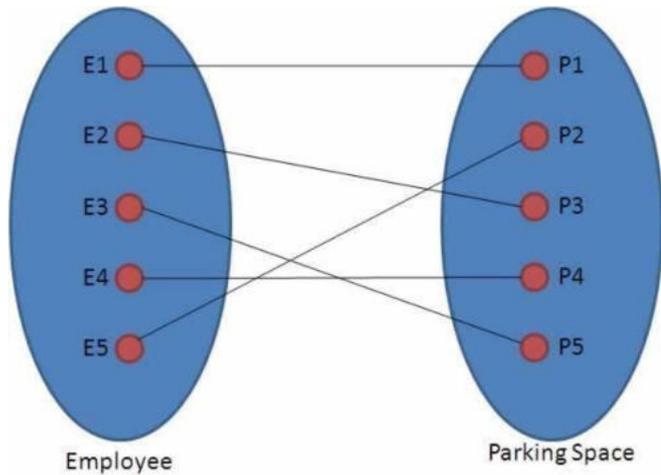
3. Many to one (M:1) relationship

4. Many to many (M:N) relationship

The minimum and maximum values of this connectivity is called the cardinality of the relationship

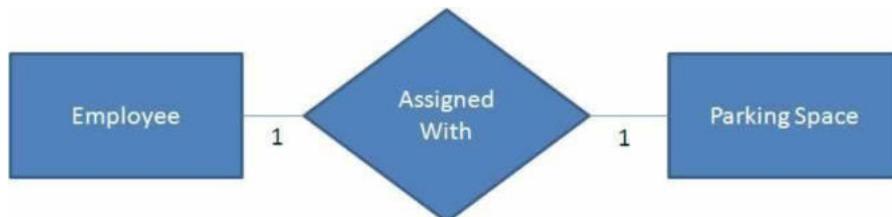
Example for Cardinality – One-to-One (1:1) Employee is

assigned with a parking space.



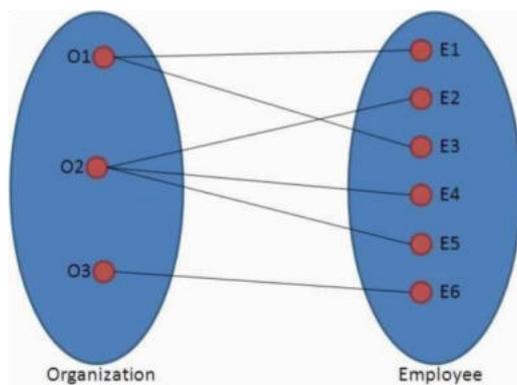
One employee is assigned with only one parking space and one parking space is assigned to only one employee. Hence it is a 1:1 relationship and cardinality is One-To-One (1:1)

In ER modeling, this can be mentioned using notations as given below



Example for Cardinality – One-to-Many (1:N)

Organization has employees



M. NANDINI

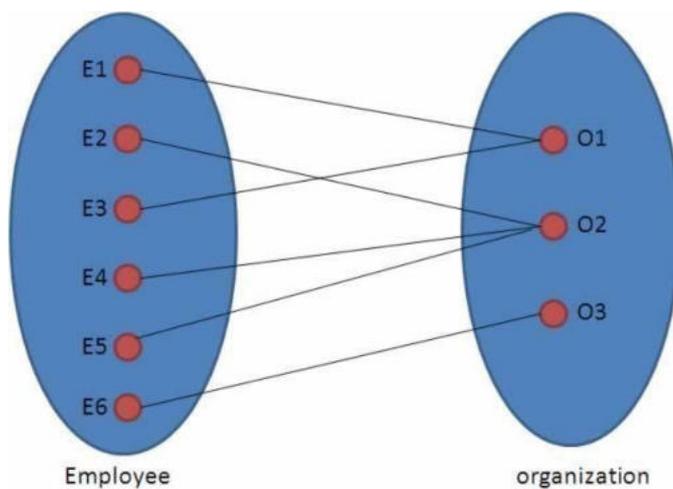
One organization can have many employees, but one employee works in only one organization. Hence it is a 1:N relationship and cardinality is One-To-Many (1:N)

In ER modeling, this can be mentioned using notations as given below



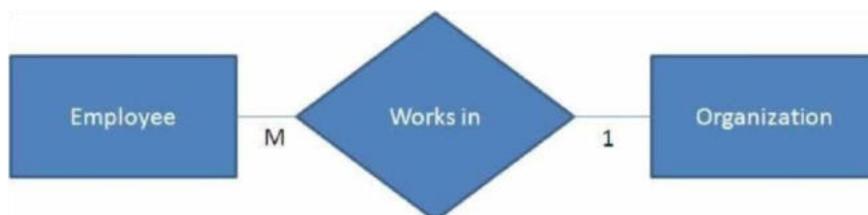
Example for Cardinality–Many-to-One(M:1)

It is the reverse of the One to Many relationship. employee works in organization



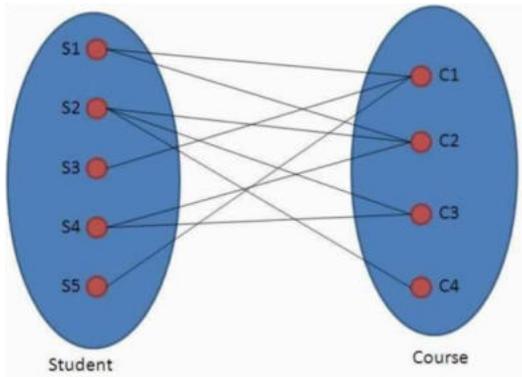
One employee works in only one organization But one organization can have many employees. Hence it is a M:1 relationship and cardinality is Many-to- One (M :1)

In ER modeling, this can be mentioned using notations as given below.

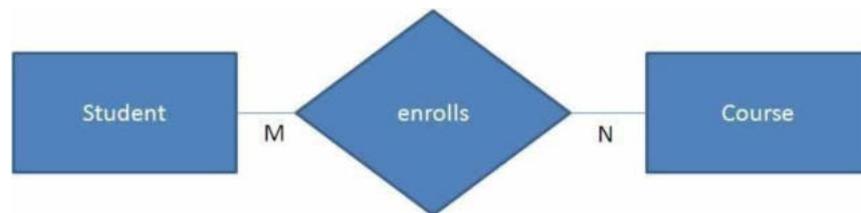


Cardinality–Many-to-Many(M:N)

Students enroll for courses



One student can enroll for many courses and one course can be enrolled by many students. Hence it is a M:N relationship and cardinality is Many-to-Many (M:N). In ER modeling, this can be mentioned using notations as given below



Relationship Participation

1. Total

In total participation, every entity instance will be connected through the relationship to another instance of the other participating entity types.

2. Partial

Example for relationship participation

Consider the relationship - Employee is head of the department.

Here all employees will not be the head of the department. Only one employee will be the head of the department. In other words, only few instances of employee entity participate in the above relationship. So employee entity's participation is partial in the said relationship.

However, each department will be headed by some employee. So department entity's participation is total in the said relationship.



Advantages and Disadvantages of ER Modeling (Merits and

Demerits of ER Modeling) Advantages

1. ER Modeling is simple and easily understandable. It is represented in business users language and it can be understood by non-technical specialist.
2. Intuitive and helps in Physical Database creation.
3. Can be generalized and specialized based on needs.
4. Can help in database design.
5. Gives a higher level description of the system.

Disadvantages

1. Physical design derived from E-R Model may have some amount of ambiguities or inconsistency.
2. Sometimes diagrams may lead to misinterpretations

ADVANCED FEATURES OF ER MODEL:

Beyond basic entities, attributes, and relationships, ER models gain power with Enhanced ER (EER) features, including Specialization/Generalization, Aggregation, Inheritance, Weak Entities, and Union Types (Categories), allowing for more complex data modeling, better abstraction, hierarchy creation, and handling of real-world complexities like IS-A relationships (e.g., "Employee ISA Person") and treating relationships as entities.

Key Additional Features (EER Model)

1. Specialization & Generalization:

- **Specialization:** Creating subclasses (e.g., Student, Professor) from a superclass (e.g., Person) based on specific characteristics.
- **Generalization:** The reverse, abstracting common attributes to form a superclass from multiple subclasses.
- **Attribute Inheritance:** Subclasses automatically inherit attributes from their superclass.

2. Aggregation:

- Treats a relationship set as an abstract entity, allowing relationship to relationships (e.g., a "Sponsor" relationship between "Employee" and "Project" might have an "Amount" attribute, forming an aggregate entity).

3. Weak Entities:

- Entities that depend on a strong entity for their existence (e.g., Dependent relies on Employee and use a partial key).

4. Union Types (Categories):

- Model entities that can belong to one of several different entity sets (e.g., a Registered_User could be a Customer OR an Admin).



5. Enhanced Constraints:

- More detailed rules for generalization/specialization, like Disjointness (a subclass can only belong to one superclass) and Overlap (a subclass can belong to multiple superclasses). These extensions provide a richer, more expressive way to design databases that accurately reflect complex real-world data structures, moving beyond basic ER models.

THE INTERNET SHOP

An ER Diagram (Entity-Relationship Diagram) for an online shop models key components like **Customers**, **Products**, **Orders**, and **Payments**, showing how they connect:

Customers place Orders, Orders contain Products (via an Order_Item/Cart link), and Orders are linked to Payments, with attributes like customer details, product descriptions, order dates, and payment methods defining each entity and relationship, forming the database blueprint.

Key Entities & Attributes

- **Customer:** customer_id (PK), name, email, address, password.
- **Product:** product_id (PK), name, description, price, stock_quantity, image_url.
- **Order:** order_id (PK), customer_id (FK), order_date, total_amount, status, shipping_address.
- **Payment:** payment_id (PK), order_id (FK), amount, payment_method, transaction_id, payment_date.
- **Cart/Cart_Item:** Represents items a customer intends to buy (often a temporary link between Customer and Product).

Key Relationships (Crow's Foot Notation Example)

- **Customer Places Order:** One Customer can place Many Orders (1:M).
- **Order Contains Products:** One Order contains Many Products (M:N, linked through an Order_Item entity).
- **Product Belongs to Category:** One Category has Many Products (1:M).
- **Order Linked to Payment:** One Order has One Payment (1:1 or 1:M if multiple payment attempts).

How It Works in the Diagram

1. Customers register and login.
2. They browse **Products**, which have details and categories.
3. Products are added to a **Cart**.
4. The customer places an **Order**, linking them to the items in their cart.
5. The **Order** is processed, a **Payment** is made, and shipping details are captured.

M. NANDINI

This ER Diagram acts as a blueprint for the database, ensuring efficient storage and retrieval of e-commerce data.

Relational Model

The relational model is today the primary data model for commercial data processing applications. It attained its primary position because of its simplicity, which eases the job of the programmer, compared to earlier data models such as the network model or the hierarchical model. In this, we first study the fundamentals of the relational model. A substantial theory exists for relational databases.

Structure of Relational Databases:

A relational database consists of a collection of **tables**, each of which is assigned a unique name. For example, consider the *instructor* table of Figure 1.5, which stores information about instructors. The table has four column headers: *ID*, *name*, *dept name*, and *salary*. Each row of this table records information about an instructor, consisting of the instructor's *ID*, *name*, *dept name*, and *salary*. Similarly, the *course* table of Figure 1.6 stores information about courses, consisting of a *course id*, *title*, *dept name*, and *credits*, for each course. Note that each instructor is identified by the value of the column *ID*, while each course is identified by the value of the column *course id*.

Figure 1.7 shows a third table, *prereq*, which stores the prerequisite courses for each course. The table has two columns, *course id* and *prereq id*. Each row consists of a pair of course identifiers such that the second course is a prerequisite for the first course.

Thus, a row in the *prereq* table indicates that two courses are *related* in the sense that one course is a prerequisite for the other. As another example, we consider the table *instructor*, a row in the table can be thought of as representing the relationship between a specified *ID* and the corresponding values for *name*, *dept name*, and *salary* values.

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

Figure 1.5: The instructor relation (2.1)

In general, a row in a table represents a *relationship* among a set of values. Since a table is a collection of such relationships, there is a close correspondence between the concept of *table* and the mathematical concept of *relation*, from which the relational data model takes its name. In mathematical terminology, a *tuple* is simply a sequence (or list) of values. A relationship between n values is represented mathematically by an n -*tuple* of values, i.e., a tuple with n values, which corresponds to a row in a table.

course_id	title	dept_name	credits
BIO-101	Intro. to Biology	Biology	4
BIO-301	Genetics	Biology	4
BIO-399	Computational Biology	Biology	3
CS-101	Intro. to Computer Science	Comp. Sci.	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3
CS-319	Image Processing	Comp. Sci.	3
CS-347	Database System Concepts	Comp. Sci.	3
EE-181	Intro. to Digital Systems	Elec. Eng.	3
FIN-201	Investment Banking	Finance	3
HIS-351	World History	History	3
MU-199	Music Video Production	Music	3
PHY-101	Physical Principles	Physics	4

Figure 1.6: The course relation (2.2)

course_id	prereq_id
BIO-301	BIO-101
BIO-399	BIO-101
CS-190	CS-101
CS-315	CS-101
CS-319	CS-101
CS-347	CS-101
EE-181	PHY-101

Figure 1.7: The prereq relation. (2.3)

Thus, in the relational model the term **relation** is used to refer to a table, while the term **tuple** is used to refer to a row. Similarly, the term **attribute** refers to a column of a table.

Examining Figure 1.5, we can see that the relation *instructor* has four attributes:

ID, name, dept name, and salary.

We use the term **relation instance** to refer to a specific instance of a relation, i.e., containing a specific set of rows. The instance of *instructor* shown in Figure 1.5 has 12 tuples, corresponding to 12 instructors.

In this topic, we shall be using a number of different relations to illustrate the various concepts underlying the relational data model. These relations represent part of a university. They do not include all the data an actual university database would contain, in order to simplify our presentation.

The order in which tuples appear in a relation is irrelevant, since a relation is a *set* of tuples. Thus, whether the tuples of a relation are listed in sorted order, as in Figure 1.5, or are unsorted, as in Figure 1.8, does not matter; the relations in the two figures are the same, since both contain the same set of tuples. For ease of exposition, we will mostly show the relations sorted by their first attribute. For each attribute of a relation, there is a set of permitted values, called the **domain** of that attribute.

Thus, the domain of the *salary* attribute of the *instructor* relation is the set of all possible salary values, while the domain of the *name* attribute is the set of all possible instructor names.



We require that, for all relations r , the domains of all attributes of r be atomic. A domain is **atomic** if elements of the domain are considered to be indivisible units.

ID	name	dept_name	salary
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

Figure:1.8: Unsorted display of the instructor relation. (2-4)

For example, suppose the table *instructor* had an attribute *phone number*, which can store a set of phone numbers corresponding to the instructor. Then the domain of *phone number* would not be atomic, since an element of the domain is a set of phone numbers, and it has subparts, namely the individual phone numbers in the set.

The important issue is not what the domain itself is, but rather how we use domain elements in our database. Suppose now that the *phone number* attribute stores a single phone number. Even then, if we split the value from the phone number attribute into a country code, an area code and a local number, we would be treating it as a nonatomic value. If we treat each phone number as a single indivisible unit, then the attribute *phone number* would have an atomic domain.

The **null** value is a special value that signifies that the value is unknown or does not exist. For example, suppose as before that we include the attribute *phone number* in the *instructor* relation. It may be that an instructor does not have a phone number at all, or that the telephone number is unlisted. We would then have to use the null value to signify that the value is unknown or does not exist. We shall see later that null values cause a number of difficulties when we access or update the database, and thus should be eliminated if at all possible. We shall assume null values are absent initially.

Database Schema

When we talk about a database, we must differentiate between the **database schema**, which is the logical design of the database, and the **database instance**, which is a snapshot of the data in the database at a given instant in time. The concept of a relation corresponds to the programming language notion of a variable, while the concept of a **relation schema** corresponds to the programming-language notion of type definition.

In general, a relation schema consists of a list of attributes and their corresponding domains. The concept of a relation instance corresponds to the programming-language notion of a value of a variable. The value of a given variable may change with time;

<i>deptname</i>	<i>building</i>	<i>budget</i>
Biology	Watson	90000
Comp. Sci.	Taylor	100000
Elec. Eng.	Taylor	85000
Finance	Painter	120000
History	Painter	50000
Music	Packard	80000
Physics	Watson	70000

Figure 1.9: The department relation. (2-5)

similarly the contents of a relation instance may change with time as the relation is updated. In contrast, the schema of a relation does not generally change. Although it is important to know the difference between a relation schema and a relation instance, we often use the same name, such as *instructor*, to refer to both the schema and the instance. Where required, we explicitly refer to the schema or to the instance, for example “the *instructor* schema,” or “an instance of the *instructor* relation.” However, where it is clear whether we mean the schema or the instance, we simply use the relation name.

Consider the *department* relation of Figure 1.9. The schema for that relation is

department(deptname, building, budget)

Note that the attribute *dept name* appears in both the *instructor* schema and the *department* schema. This duplication is not a coincidence. Rather, using common attributes in relation schemas is one way of relating tuples of distinct relations.

For example, suppose we wish to find the information about all the instructors who work in the Watson building. We look first at the *department* relation to find the *deptname* of all the departments housed in Watson. Then, for each such department, we look in the *instructor* relation to find the information about the instructor associated with the corresponding *dept name*.

Let us continue with our university database example. Each course in a university may be offered multiple times, across different semesters, or even within a semester. We need a relation to describe each individual offering, or section, of the class. The schema is ***section (course id, sec id, semester, year, building, room number, time slot id)***

Figure 1.10 shows a sample instance of the *section* relation. We need a relation to describe the association between instructors and the class sections that they teach. The relation schema to describe this association is ***teaches (ID, course id, sec id, semester, year)***

course_id	sec_id	semester	year	building	room_number	time_slot_id
BIO-101	1	Summer	2009	Painter	514	B
BIO-301	1	Summer	2010	Painter	514	A
CS-101	1	Fall	2009	Packard	101	H
CS-101	1	Spring	2010	Packard	101	F
CS-190	1	Spring	2009	Taylor	3128	E
CS-190	2	Spring	2009	Taylor	3128	A
CS-315	1	Spring	2010	Watson	120	D
CS-319	1	Spring	2010	Watson	100	B
CS-319	2	Spring	2010	Taylor	3128	C
CS-347	1	Fall	2009	Taylor	3128	A
EE-181	1	Spring	2009	Taylor	3128	C
FIN-201	1	Spring	2010	Packard	101	B
HIS-351	1	Spring	2010	Painter	514	C
MU-199	1	Spring	2010	Packard	101	D
PHY-101	1	Fall	2009	Watson	100	A

Figure 1.10: This section relation. (2-6)

Figure 1.11 shows a sample instance of the *teaches* relation. As you can imagine, there are many more relations maintained in a real university database. In addition to those relations we have listed already, *instructor*, *department*, *course*, *section*, *prereq*, and *teaches*, we use the following relations in this text:

ID	course_id	sec_id	semester	year
10101	CS-101	1	Fall	2009
10101	CS-315	1	Spring	2010
10101	CS-347	1	Fall	2009
12121	FIN-201	1	Spring	2010
15151	MU-199	1	Spring	2010
22222	PHY-101	1	Fall	2009
32343	HIS-351	1	Spring	2010
45565	CS-101	1	Spring	2010
45565	CS-319	1	Spring	2010
76766	BIO-101	1	Summer	2009
76766	BIO-301	1	Summer	2010
83821	CS-190	1	Spring	2009
83821	CS-190	2	Spring	2009
83821	CS-319	2	Spring	2010
98345	EE-181	1	Spring	2009

Figure: 1.11: The teaches relation. (2-7)

- *student*(ID, name, deptname, totcred)
 - *advisor*(sid, iid)
- *takes*(ID, courseid, secid, semester, year, grade)
- *classroom*(building, roomnumber, capacity)
- *timeslot*(timeslotid, day, starttime, endtime)

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

M. NANDINI

tuples must be such that they can *uniquely identify* the tuple. In other words, not 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 $t_1 = t_2$, then $t_1.K = t_2.K$.

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 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 .

DBMS

M. NANDINI

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 t_a , from the *instructor* relation, there must be some tuple, say t_b , in the *department* relation such that the value of the *dept name* attribute of t_a is the same as the value of the primary key, *dept name*, of t_b .

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.

Schema Diagrams

A database schema, along with primary key and foreign key dependencies, can be depicted by **schemadiagrams**. Figure 1.12 shows the schemadiagram 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.

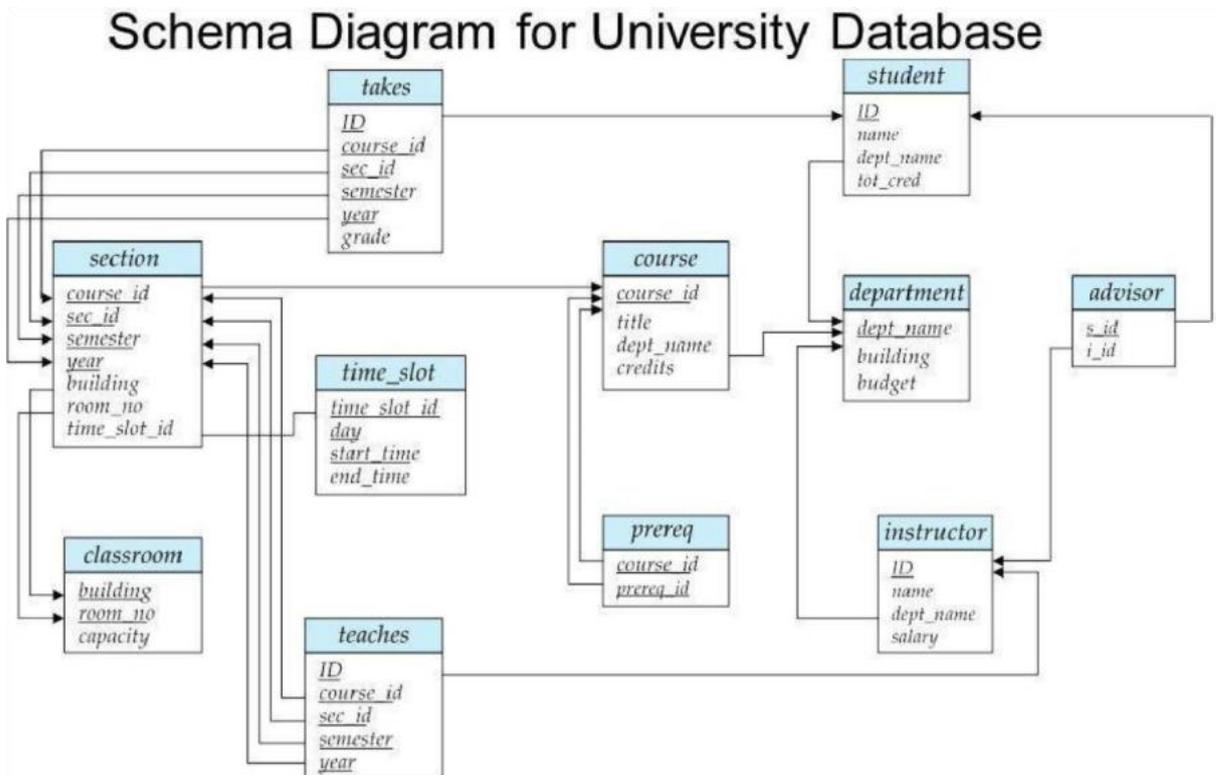


Figure 1.12: Schemadiagramfortheuniversitydatabase.

Referential integrity constraints other than foreign key constraints are not shown explicitly in schema diagrams. We will study a different diagrammatic representation called the entity-relationship diagram.

UNIT III: SOL and PL/SOL

SQL and PL/SQL: Introduction to SQL, Data Definition Commands, Data Manipulation Commands, Basic SQL Querying (Select and project) Select Queries, **Virtual Tables:** Creating View, Altering View, Updating View, Destroying View, Relational Set Operators, Join Operators, Sub Queries and Correlated Queries, implementation of different integrity constraints, Aggregate Functions, Procedural SQL: Stored Procedures, Stored Functions, Triggers, Cursors

DATA DEFINITION COMMANDS, DATA MANIPULATION COMMANDS, BASIC SQL QUERYING (SELECT AND PROJECT) SELECT QUERIES

DDL (Data Definition Language) commands are used to define or modify the database structure, while DML (Data Manipulation Language) commands are used to manage the data within those structures.

DDL Commands (Data Definition Language)

DDL commands affect the schema or structure of the database and are auto-committed, meaning the changes are permanent and cannot be rolled back.

- **CREATE:** Used to create new database objects, such as tables, indexes, views, or the database itself.

```
CREATE TABLE Students (StudentID int, Name varchar(255));
```

Example:

- **ALTER:** Used to modify the structure of an existing database object, like adding, deleting, or modifying columns in a table.

```
ALTER TABLE Students ADD Email varchar(255);
```

Example:

- **DROP:** Used to permanently delete a database object (e.g., a table, index, or database) and all the data within it.

Example: `DROP TABLE Students;`

□ **TRUNCATE** :Usedtoquicklydeleteallrecordsfromatablebut keepsthetablestructureintactforfutureuse.ItisaDDLcommand because,likeDROP,itsoperationsarenon-transactionalandcannot berolledback.

○ *Example:* `TRUNCATETABLEStudents;`

• **RENAME**:Usedtochangethenameofadatabaseobject.

○ *Example:* `RENAMETABLEStudentsTOSTudent_Data;`

DML Commands(Data Manipulation Language)

DML commands are used to manage the data stored within the database objects. These commands are transactional, meaning changes can be rolled back (undone) if necessary before being committed permanently.

• **INSERT**:Usedtoaddnewrows(records)ofdataintoa table.

○ *Example:* `INSERTINTOStudents(StudentID,Name)VALUES(1,'JohnDoe');`

□ **UPDATE** :Usedtomodifyexistingdatawithinatable.

○ `UPDATEStudentsSETName='JaneDoe'WHEREStudentID=1;`

Example:

• :Usedtoremoveexistingrecords(rows)fromatable. ○

DELETE

`DELETEFROMStudentsWHEREStudentID=1;`

Example:

SELECT

SELECT

• : While often categorized separately as Data Query Language (DQL), is widely considered part of DML as it interacts with the data within the tables. ○ *Example:*

```
SELECT * FROM Students;
```

Introduction to PL/SQL

PL/SQL (Procedural Language/Structured Query Language) is an extension of SQL (Structured Query Language) designed by

Oracle. It enables developers to write code that combines the power of SQL queries with the procedural capabilities of a programming

language. PL/SQL is mainly used to interact with Oracle databases, providing a way to execute SQL statements in a procedural manner, which allows for more control over data processing, including logic flow and error handling.

Structure of PL/SQL in DBMS

PL/SQL programs can be structured into different blocks. The basic structure of a PL/SQL block is as follows:

1. Declaration Section (optional):

This section is used to declare variables, constants, cursors, and exceptions that will be used in the PL/SQL block.

DECLARE

variable_name datatype;

CONSTANT_NAME CONSTANT datatype := value; BEGIN

2. ExecutionSection(mandatory):

This is the section where the actual SQL statements and procedural code are written. It is the core of the PL/SQL block and is mandatory. SQL commands like SELECT, INSERT, UPDATE, or DELETE are executed here, along with procedural logic like loops and conditions.

Syntax:

```
BEGIN
    --SQL and PL/SQL code
END;
```

3. ExceptionHandlingSection(optional):

- This section is used to handle runtime errors or exceptions that occur during the execution of the PL/SQL block. It allows programmers to gracefully handle errors instead of having the program crash or stop.

Syntax:

```
EXCEPTION
    WHEN exception_name THEN
        -- Handle error
    WHEN OTHERS THEN
        -- Handle all other errors
END;
```

Example of a Simple PL/SQL Block:

```
DECLARE v_emp_name VARCHAR2(50);
        v_emp_salary NUMBER(10,2);

BEGIN

    SELECT employee_name, salary INTO v_emp_name, v_emp_salary
    FROM employees
    WHERE employee_id=100;

    DBMS_OUTPUT.PUT_LINE('Employee Name: ' || v_emp_name);
    DBMS_OUTPUT.PUT_LINE('Salary: ' || v_emp_salary); EXCEPTION
    WHEN NO_DATA_FOUND THEN
        DBMS_OUTPUT.PUT_LINE('No employee found with the given ID.');
```

```
    WHEN OTHERS THEN
        DBMS_OUTPUT.PUT_LINE('An unexpected error
occurred.');
```

```
END;
```

Topic1:

Decisionmakingstatements Or

Conditionalstatements Or Controlstatements:

Decision making or conditional or Control statements are those statements which are in charge of executing a statement out of multiple given statements based on some condition. The condition will return either true or false. Based on what the condition returns, the associated statement is executed.

Conditional Statements available in PL/SQL are defined below:

1. **IF THEN**
2. **IF THEN ELSE**
3. **NESTED-IF-THEN**
4. **IF THEN ELSE IF-THEN-ELSE**

1. **IF THEN** if then the statement is the most simple decision-making

statement. It is used to decide whether a certain statement or block of statements will be executed or not i.e. if a certain condition is true then a block of statement is executed otherwise not.

Syntax:

```
IF condition THEN
```

```
-- Code to execute if the condition is true END IF; Example:
```

DECLARE

```
v_salary NUMBER := 3000; BEGIN
```

```
IF v_salary > 2500 THEN DBMS_OUTPUT.PUT_LINE('Salary is above  
2500');
```

```
ENDIF; END;
```

2. IF THEN ELSE

The if statement alone tells us that if a condition is true it will execute a block of statements and if the condition is false it won't. But what if we want to do something else if the condition is false. Here comes the else statement.

We can use the else statement with if statement to execute a block of code when the condition is false. **Syntax:**

```
IF condition THEN
```

```
--Code to execute if the condition is true ELSE
```

```
--Code to execute if the condition is false END IF; Example:
```

```
DECLARE v_age NUMBER := 20;
```

```
BEGIN
```

```
IF v_age >= 18 THEN DBMS_OUTPUT.PUT_LINE('Adult');
```

```
ELSE
```

```
DBMS_OUTPUT.PUT_LINE('Minor'); END
```

```
IF;
```

```
END;
```

3. NESTED-IF-THEN

An **nested IF-THEN** statement occurs when an IF-THEN statement is placed inside another IF-THEN statement.

Syntax:

```
IF condition1 THEN
  --Code to execute if condition1 is true IF condition2
  THEN
    --Code to execute if condition2 is true END IF;
END IF;
```

Here, the second IF statement (the nested one) will only be evaluated if the first IF statement is true.

Example:

```
declare
num1 number:=10;
num2 number:=20;
num3 number:=20;

begin if num1 < num2 then
dbms_output.put_line('num1 small num
2');  if  num1  <  num3  then
dbms_output.put_line('num1 small num3 also');
end if;
```

end if;

dbms_output.put_line('afterendif')

; end;

4. if...then...elsif...elsestatement

It is used to check multiple conditions. Sometimes it is required to test more than one condition in that case if...then...else statement cannot be used. For this purpose, if...then...elsif...else statement is suitable in which all the conditions are tested one by one and whichever condition is found to be

TRUE, that block of code is executed. And if all the conditions result in **FALSE** then the else part is executed.

Syntax:

if<test_condition1>then body of action

elsif<test_condition2>then body of

action elsif<test_condition3>then

body of action

...

...

...

else

body of action

```
endif;
```

```
setserveroutputon;
```

Example:

```
DECLARE
```

```
    a int;
```

```
    b int;
```

```
BEGIN
```

```
    a := &a;    b := &b;    if(a>b)then
```

```
    dbms_output.put_line('aisgreaterthanb'); elsif(b>a)
```

```
    then
```

```
        dbms_output.put_line('bisgreaterthana');
```

```
    else dbms_output.put_line('Bothaandbareequal');
```

```
    endif;
```

```
END;
```

Topic2:

PL/SQL Loops:

Loops Statements (or) Iterative Statements

Loops in PL/SQL provides a way of repeating a particular part of any program or any code statement as many times as required. In PL/SQL we have three different loop options to choose from when we want to execute a statement repeatedly in our code block.

They are:

1. Basic Loop
2. While Loop
3. For Loop

1. Basic Loop

Basic loop or simple loop is preferred in PL/SQL code when there is no surety about how many times the block of code is to be repeated. When we use the basic loop the code block will be executed at least once.

While using it, following two things must be considered:

- Simple loop always begins with the keyword LOOP and ends with a keyword END LOOP.
- A basic/simple loop can be terminated at any given point by using the exit statement or by specifying certain condition by using the statement exit when.

Syntax:

LOOP

--Statements to be executed repeatedly

EXIT WHEN condition; --Optional exit condition END LOOP; **Example:**

```
DECLARE counter NUMBER:=1;
```

```
BEGIN
```

```
LOOP
```

```
DBMS_OUTPUT.PUT_LINE('Counter:'||counter); counter := counter  
+ 1;
```

```
EXIT WHEN counter>5; END
```

```
LOOP;
```

```
END;
```

```
/
```

2. While Loop

It is an entry controlled loop which means that before entering in a while loop first the condition is tested, if the condition is **TRUE** the statement or a group of statements get executed and if the condition is **FALSE** the control will move out of the while loop.

Syntax:

```
WHILE <test_condition> LOOP
```

```
    <action> E
```

```
END LOOP;
```

Example:

```
setserveroutputon; DECLARE numint:=1;
```

```
BEGIN
```

```
    while(num<=10)LOOP
```

```
        dbms_output.put_lin
```

```
        e("||no); num
```

```
        := num+2;
```

```
    ENDLOOP;
```

```
END;
```

3. For Loop

This loop is used when some statements in PL/SQL code block are to be repeated for a fixed number of times.

When we use the for loop we are supposed to define a counter variable which decides how many times the loop will be executed based on a starting and ending value provided at the beginning of the loop. The for loop automatically increments the value of the counter variable by 1 at the end of each loop cycle.

Syntax:

```
FOR counter_variable IN start_value..end_value LOOP statement to be executed
```

```
ENDLOOP;
```

Example:

```
setserveroutputon; DECLARE inumber(2);
```

```
BEGIN
```

```
FORiIN1..10 LOOP
```

```
dbms_output.put_line(i); END
```

```
LOOP;
```

```
END;
```

Topic3:

Triggersin PL/SQL

A **trigger** in PL/SQL is a type of stored procedure that automatically executes (or "fires") in response to a specific event occurring in the database. Triggers are used to enforce business rules, validate data, maintain audit trails, or automatically modify data in a table based on changes made to that table.

Triggers are defined to execute either **before** or **after** a specific event on a database object such as an **INSERT, UPDATE, or**

DELETE operation. The event can be a change in a table or view, and triggers are bound to that table or view.

TypesofTriggers

1. BEFORETrigger
2. AFTERTrigger
3. INSTEADOFTrigger
4. Level Triggers

Thereare2differenttypesofleveltriggers,theyare:

1. ROWLEVEL TRIGGERS

2. STATEMENTLEVELTRIGGERS

TriggerSyntax:

```
CREATE[ORREPLACE]TRIGGERtrigger_name
  {BEFORE|AFTER|INSTEADOF}
  {INSERT|UPDATE|DELETE}
  ON table_name
  [FOREACHROW]
  DECLARE
    --Variabledeclarations(optional) BEGIN
    --Triggercode END;
```

1. BEFORETrigger: This trigger fires before the triggering event (INSERT, UPDATE, DELETE) is executed on the table.

Syntax:

```
CREATEORREPLACETRIGGERtrigger_name BEFORE
  {INSERT | UPDATE | DELETE}
  ON table_name
  [FOREACHROW]
  DECLARE
    --Variabledeclarations(optional) BEGIN
    --Triggercode END; Example:
```

```
CREATEORREPLACETRIGGERbefore_employee_insert BEFORE INSERT
  ON employees
  FOREACHROW
```

BEGIN

```
--Set the hire_date to the current system date before the new record is inserted  
:NEW.hire_date:=SYSDATE; END;
```

2. AFTER trigger in PL/SQL is a type of trigger that executes **after** a specified event (such as INSERT, UPDATE, or DELETE) has been performed on a table. This trigger is typically used for actions that need to occur after the data has been modified, such as logging changes, auditing data, or updating related tables.

Syntax:

```
CREATE OR REPLACE TRIGGER trigger_name  
AFTER {INSERT|UPDATE|DELETE}  
ON table_name  
[FOR EACH ROW]  
DECLARE  
--Variable declarations (optional) BEGIN  
--Trigger code END;
```

Example:

```
CREATE OR REPLACE TRIGGER after_employee_insert AFTER  
INSERT ON employees  
FOR EACH ROW  
BEGIN  
--Insert a record into the audit_log table after a new employee is inserted  
INSERT INTO audit_log(log_id, action_type, employee_id, action_date)
```

```
VALUES (audit_log_seq.NEXTVAL, 'INSERT',
:NEW.employee_id,SYSDATE); END;
```

3. INSTEADOFTriggerin PL/SQL

An **INSTEADOF** trigger in PL/SQL is used to define the behavior that should occur **instead of** the usual DML (Data Manipulation

Language) operations such as INSERT, UPDATE, or DELETE. This type of trigger is typically used on **views** rather than tables. The purpose of an INSTEAD OF trigger is to replace the default action with a custom action when an operation is performed on a view.

Syntax:

```
CREATE OR REPLACE TRIGGER trigger_name INSTEAD OF
{INSERT | UPDATE | DELETE}
ON view_name
[FOR EACH ROW]
DECLARE
--Variable declarations (optional) BEGIN
--Trigger code          END;
```

Example:

```
CREATE OR REPLACE TRIGGER update_employee_department
INSTEAD OF UPDATE ON Employee_department_view FOR EACH
ROW
BEGIN
```

```
--Updatetheemployeestable
```

```
UPDATE employees
```

```
SETemployee_name=:NEW.employee_name WHERE employee_id  
=:OLD.employee_id;
```

```
--Updatethedepartmentstable
```

```
UPDATE departments
```

```
SETdepartment_name=:NEW.department_name WHERE department_id  
=:OLD.department_id;
```

```
END;
```

4. Level Triggers in PL/SQL

Triggers are special procedures in PL/SQL that automatically execute in response to certain events on a particular table or view. Based on the scope of their operation, triggers are categorized into **Row-Level Triggers** and **Statement-Level Triggers**.

1. Row-Level Triggers

- A **Row-Level Trigger** is executed once for each row affected by a DML (INSERT, UPDATE, DELETE) statement.
- They are used when you need to work with or track changes to specific rows in a table.

Example of a Row-Level Trigger:

```

CREATE OR REPLACE TRIGGER row_level_trigger_example AFTER INSERT
OR UPDATE ON employees
FOR EACH ROW
BEGIN
    DBMS_OUTPUT.PUT_LINE('Row-Level Trigger fired for
Employee ID: ' || :NEW.employee_id);
END;
/

```

2. Statement-Level Triggers

- A **Statement-Level Trigger** is executed once for the entire DML statement, regardless of the number of rows affected.
- These triggers are used when the operation affects the table as a whole, not specific rows.

Example of a Statement-Level Trigger:

```

CREATE OR REPLACE TRIGGER statement_level_trigger_example
AFTER DELETE ON employees
BEGIN
    DBMS_OUTPUT.PUT_LINE('Statement-Level Trigger fired after a
DELETE operation on employees. ');
END;
/

```

Topic 4:

Cursor in PL/SQL

A **cursor** in PL/SQL is a pointer to a context area that stores the result set of a query. It allows row-by-row processing of query results.

CursorAttributes

PL/SQL provides cursor attributes for obtaining information about the cursor state:

- **%FOUND**: Returns TRUE if a row was fetched, otherwise FALSE.
- **%NOTFOUND**: Returns TRUE if no row was fetched, otherwise FALSE.
- **%ISOPEN**: Returns TRUE if the cursor is open, otherwise FALSE.
- **%ROWCOUNT**: Returns the number of rows fetched so far.

PL/SQL provides two types of cursors:

1. ImplicitCursors

2. ExplicitCursors

1. ImplicitCursor

- Automatically created by PL/SQL for SELECT INTO, INSERT, UPDATE, and DELETE statements that affect a single row or multiple rows.
- You don't need to declare or control it.

Example of Implicit Cursor:

```
DECLARE emp_count NUMBER;
```

```
BEGIN
```

```
SELECT COUNT(*) INTO emp_count FROM employees;
```

```
DBMS_OUTPUT.PUT_LINE('Total Employees: ' || emp_count);
```

END;

/

2. ExplicitCursor

- Defined and controlled explicitly by the programmer.
- Useful for queries that return multiple rows.
- Consists of four main steps:
 1. **Declaration:** Declare the cursor.
 2. **Opening:** Open the cursor to execute the query and establish the result set.
 3. **Fetching:** Retrieve rows from the result set one by one.
 4. **Closing:** Close the cursor to release resources.

Example of Explicit Cursor:

```
DECLARE
```

```
CURSOR emp_cursor IS
```

```
SELECT employee_id, name FROM employees; v_emp_id  
employees.employee_id%TYPE;           v_name  
employees.name%TYPE;
```

```
BEGIN
```

```
-- Open the cursor
```

```
OPEN emp_cursor;
```

```
-- Loop through rows LOOP
```

```
FETCH emp_cursor INTO v_emp_id, v_name;
```

```

--Exittheloopwhennomorerowsareavailable EXIT
WHEN emp_cursor%NOTFOUND;
-- Process the fetched data DBMS_OUTPUT.PUT_LINE('Employee ID: ' ||
v_emp_id ||
',Name:' ||v_name); END
LOOP;
-- Close the cursor
CLOSEemp_cursor;
END;
/

```

Exception Handling in PL/SQL

An exception is an error. An error accrued during the program execution is called exception.

Types of Exceptions

1. **Pre-defined or System Exceptions:** These are built-in exceptions in PL/SQL, such as NO_DATA_FOUND, TOO_MANY_ROWS, and ZERO_DIVIDE.

2. **User-Defined Exceptions:** These are explicitly defined by the programmer to handle specific business logic errors.

SyntaxforExceptionHandling

Syntax:

```
DECLARE
    <exception_name>EXCEPTION; BEGIN
    ...
    RAISE<exception_name>;
    ...
END;
```

1.System-definedExceptions

System-defined exceptions are predefined in programming languages to handle runtime errors.

```
DECLARE
    --Declarationstatements; BEGIN
    --SQL statements;
    --Proceduralstatements;
EXCEPTION
    --Exceptionhandlingstatements;
END;
```

Exception	Raisedwhen....
------------------	-----------------------

LOGIN_DENIED	At the time when user login is denied.
TOO_MANY_ROWS	When a select query returns more than one the destination variable can take only single
VALUE_ERROR	When an arithmetic, value conversion, trunc constraint error occurs.

Example:

```
set server output on; DECLARE
```

```
    a int;
```

```
    b int;
```

```
    c int;
```

```
BEGIN
```

```
    a := &a;
```

```
    b := &b; c
```

```
    := a/b;
```

```
    dbms_output.put_line('RESULT='||c);
```

```
EXCEPTION
```

```
    when ZERO_DIVIDE then
```

```
        dbms_output.put_line('Division by 0 is not possible');
```

END;

User-defined Exception

In any program, there is a possibility that a number of errors can occur that may not be considered as exceptions by oracle. In that

case, an exception can be defined by the programmer while writing

the codes such type of exceptions are called User-defined exception. User defined

exceptions are in general defined to handle

special cases where our code can generate exception due to our code logic. Also,

in your code logic, you can explicitly specify to

generate an

exception using the RAISE keyword and then handle it using the EXCEPTION

block. **Syntax:**

```
DECLARE
```

```
    <exceptionname>EXCEPTION
```

```
BEGIN
```

```
    <sqlsentence>
```

```
    If <test_condition>
```

```
        THENRAISE<exception_name
```

```
        >;
```

```
    END IF;
```

```
    EXCEPTION
```

```
        WHEN<exception_name>THEN
```

--some action

END;

Example:

DECLARE

```
v_salary NUMBER := 4000; -- Employee salary e_low_salary EXCEPTION;
--User-defined exception
```

BEGIN

```
--Check if salary is below the threshold IF v_salary
< 5000 THEN
```

```
    RAISE e_low_salary; -- Raise the user-defined exception END IF;
```

```
    DBMS_OUTPUT.PUT_LINE('Salary is acceptable. '); EXCEPTION
    WHEN e_low_salary THEN
```

```
        DBMS_OUTPUT.PUT_LINE('Error: Salary is below the minimum
threshold.');
```

```
    WHEN OTHERS THEN
```

```
        DBMS_OUTPUT.PUT_LINE('An unexpected error
occurred.');
```

END;

/

Procedures in PL/SQL

A **procedure** in PL/SQL is a named block of code that performs a specific task. It is similar to a function but does not necessarily return a value. Procedures are

used to encapsulate and reuse business logic, reducing redundancy and enhancing maintainability.

Parameters in Procedures

In PL/SQL, **parameters** in procedures allow you to pass and retrieve values between the calling program and the procedure. There are three types of parameters:

1. IN Parameter:

- Used to pass input values to the procedure.
- Acts as a constant and cannot be modified within the

procedure. **2. OUT Parameter:**

- Used to return a value from the procedure to the calling program.
- Must be assigned a value within the procedure.

3. INOUT Parameter:

- Used to pass a value into the procedure, modify it within the procedure, and return it to the calling program. **Parameters in Procedures**

In PL/SQL, **parameters** in procedures allow you to pass and retrieve values between the calling program and the procedure. There are three types of parameters:

1. IN Parameter:

- Used to pass input values to the procedure.

- Acts as a constant and cannot be modified within the procedure.

2. **OUT Parameter:**

- Used to return a value from the procedure to the calling program.
- Must be assigned a value within the procedure.

3. **INOUT Parameter:**

- Used to pass a value into the procedure, modify it within the procedure, and return it to the calling program.

Syntax for Parameters in Procedures

```
CREATE OR REPLACE PROCEDURE procedure_name( parameter1
IN datatype, -- Input parameter parameter2 OUT datatype,
-- Output parameter
parameter3 INOUT datatype -- Input and output parameter
) IS
-- Declarations
BEGIN
-- Executable statements EXCEPTION
-- Exception handling END procedure_name;
/
```

Example: Procedure with Parameters Creating the Procedure

```

CREATE OR REPLACE PROCEDURE manage_salary( p_emp_id IN
    NUMBER, -- Input parameter

    p_salary IN OUT NUMBER, -- Input and output parameter
    p_bonus OUT
    NUMBER    -- Output parameter

) IS
BEGIN
    -- Calculate bonus as 10% of salary
    p_bonus := p_salary
    * 0.10;

    -- Update salary by adding bonus
    p_salary := p_salary
    + p_bonus;

    DBMS_OUTPUT.PUT_LINE('Salary updated for Employee ID: ' || p_emp_id);

END;
/

```

Dropping a Procedure

To remove an existing procedure, use the **DROP PROCEDURE** statement.

Syntax:

```
DROP PROCEDURE procedure_name;
```

Example:

```
DROP PROCEDURE manage_salary;
```

PL/SQL Functions

PL/SQL functions are reusable blocks of code that can be used to perform specific tasks. They are similar to **procedures** but must always return a value.

A function in PL/SQL contains:

- **Function Header:** The function header includes the function name and an optional parameter list. It is the first part of the function and specifies the name and parameters.
- **Function Body:** The function body contains the executable statements that implement the specific logic. It can include declarative statements, executable statements, and exception-handling statements.

Create Function in PL/SQL

To create a procedure in PL/SQL, use the **CREATE FUNCTION statement**.

Syntax

```
CREATE [OR REPLACE] FUNCTION function_name  
(parameter_name type [, ...])
```

— This statement is used for functions

```
RETURN return_datatype
```

```
{IS | AS}
```

```
BEGIN
```

```
— program code
```

[EXCEPTION
exception_section;

END[function_name]; Example

```
CREATE OR REPLACE FUNCTION factorial(x NUMBER) RETURN  
NUMBER
```

IS

```
f NUMBER; BEGIN
```

```
IF x=0 THEN f :=
```

```
1;
```

```
ELSE
```

```
f:=x*factorial(x-1);
```

```
END IF;
```

```
RETURN f;
```

```
END;
```

```
/
```

Dropping a Function

To remove an existing function, use the DROP FUNCTION statement.

Syntax:

```
DROP FUNCTION function_name;
```

Example:

```
DROP FUNCTION calculate_bonus;
```

UNIT IV: INTRODUCTION TO SCHEMA REFINEMENT

introduction to Schema Refinement; Purpose of Normalization or Schema refinement Problems Caused

by Redundancy, Decompositions, Problems Related to Decomposition, Functional Dependencies,

Reasoning about FDS, Normal Forms: 1NF, 2NF, 3NF, BCNF, Properties of Decomposition: Lossy Join

Decomposition, Dependency Preserving Decomposition, Multivalued Dependencies, 4 NF.

SCHEMA REFINEMENT

We now present an overview of the problems that schema refinement is intended to address

and a refinement approach based on decompositions. Redundant storage of information is the

root cause of these problems. Although decomposition can eliminate redundancy, it can lead to problems of its own and should be used with caution.

Problems Caused by Redundancy

Storing the same information redundantly, that is, in more than one place within a database, can lead to several problems:

■

Redundant storage: Some information is stored repeatedly.

■

Update anomalies: If one copy of such repeated data is updated, an inconsistency is created unless all copies are similarly updated.

■

Insertion anomalies: It may not be possible to store some information unless some other information is stored as well.

Deletion anomalies: It may not be possible to delete some information without losing some other information as well.

Consider a relation obtained by translating a variant of the Hourly Emps entity set from Chapter 2:

HourlyEmps(*ssn*, *name*, *lot*, *rating*, *hourly wages*, *hours worked*)

In this chapter we will omit attribute type information for brevity, since our focus is on the grouping of attributes into relations. We will often abbreviate an attribute name to a single letter and refer to a relation schema by a string of letters, one per attribute. For example, we will refer to the Hourly Emps schema as *SNLRWH* (*W* denotes the *hourly wages* attribute).

The key for HourlyEmps is *ssn*. In addition, suppose that the *hourly wages* attribute is determined by the *rating* attribute. That is, for a given *rating* value, there is only one permissible *hourly wages* value. This IC is an example of a *functional dependency*. It leads to possible redundancy in the relation HourlyEmps, as illustrated in Figure 15.1.

If the same value appears in the *rating* column of two tuples, the IC tells us that the same value must appear in the *hourly wages* column as well. This redundancy has several negative consequences:

<i>ssn</i>	<i>name</i>	<i>lot</i>	<i>rating</i>	<i>hourlywages</i>	<i>hoursworked</i>
			8		

123-22-3666	Attishoo	48	8	10	40
231-31-5368	Smiley	22	8	10	30
131-24-3650	Smethurst	35	5	7	30
434-26-3751	Guldu	35	5	7	32
612-67-4134	Madayan	35	8	10	40

An Instance of the HourlyEmpsRelation

Some information is stored multiple times. For example, the rating value 8 corresponds to the hourly wage 10, and this association is repeated three times. In addition to wasting space by storing the same information many

times, redundancy leads to potential inconsistency. For example, the *hourly wages* in the first tuple could be updated without making a similar change in the second tuple, which is an example of an *update anomaly*. Also, we cannot insert a tuple for an employee unless we know the hourly wage for the employee's rating value, which is an example of an *insertion anomaly*.

If we delete all tuples with a given rating value (e.g., we delete the tuples for Smethurst and Guldu) we lose the association between that *rating* value and its *hourly wage* value (a *deletion anomaly*).

Let us consider whether the use of *null* values can address some of these problems. Clearly, *null* values cannot help eliminate redundant storage or update anomalies. It appears that they can address insertion and deletion anomalies.

Ideally, we want a schema that does not permit redundancy, but at the very least we want to be able to identify a schema that does allow redundancy. Even if we choose to accept a schema with some of these drawbacks, perhaps owing to performance considerations, we want to make an informed decision.

Use of Decompositions

Intuitively, redundancy arises when a relational schema forces an association between attributes that is not natural. Functional dependencies (and, for that matter, other ICs) can be used to identify such situations and to suggest refinements to the schema. The essential idea is that many problems arising from redundancy can be addressed by replacing a relation with a collection of 'smaller' relations. Each of the smaller relations contains a (strict) subset of the attributes of the original relation. We refer to this process as *decomposition* of the larger relation into the smaller relations.

We can deal with the redundancy in *HourlyEmps* by decomposing it into two relations:

HourlyEmps2(*ssn,name,lot,rating,hoursworked*)Wages(*rating, hourlywages*)

—

The instances of these relations corresponding to the instance of HourlyEmps relation in Figure 15.1 is shown in Figure 15.2.

<i>ssn</i>	<i>name</i>	<i>lot</i>	<i>rating</i>	<i>hoursworked</i>
123-22-3666	Attishoo	48	8	40

231-31-5368	Smiley	22	8	30
131-24-3650	Smethurst	35	5	30
434-26-3751	Guldu	35	5	32
612-67-4134	Madayan	35	8	40

<i>rating</i>	<i>hourlywages</i>
8	10
5	7

InstancesofHourlyEmps2 andWages

Notethatwecaneasilyrecordthehourlywageforanyratingsimplybyaddingatuple to Wages, even if no employee with that rating appears in the current instance of

HourlyEmps.Changingthewageassociatedwitharatinginvolvesupdatingasingle Wages tuple. This is more efficient than updating several tuples (as in the original

design),anditalsoeliminatesthepotentialforinconsistency.Noticethattheinsertion and deletion anomalies have also been eliminated.

ProblemsRelatedtoDecomposition

Unlesswearecareful,decomposingarelationschemacancreatemoreproblemsthan it solves. Two important questions must be asked repeatedly:

1. Doweneedtodecomposearelation?

2. What problems (if any) does a given decomposition cause?

To help with the first question, several *normal forms* have been proposed for relations. If a relation schema is in one of these normal forms, we know that certain kinds of

problems cannot arise. Considering the normal form of a given relation schema can help us to decide whether or not to decompose it further. If we decide that a relation

schema must be decomposed further, we must choose a particular decomposition (i. e., a particular collection of smaller relations to replace the given relation).

With respect to the second question, two properties of decompositions are of particular interest. The *lossless join* property enables us to recover any instance of the decomposed relation from corresponding instances of the smaller relations.

The *dependency preservation* property enables us to enforce any constraint on the original relation by simply enforcing some constraints on each of the smaller relations.

That is, we need not perform joins of the smaller relations to check whether a constraint on the original relation is violated.

FUNCTIONAL DEPENDENCIES

A functional dependency (FD) is a kind of IC that generalizes the concept of a *key*.

Let R be a relation schema and let X and Y be nonempty sets of attributes in R . We

say that an instance r of R

satisfies the FD $X \twoheadrightarrow Y$ if the following holds for every pair of tuples t_1 and t_2 in r : If $t_1[X]$

$$= t_2[X], \text{ then } t_1[Y] = t_2[Y].$$

We use the notation $t1.X$ to refer to the projection of tuple $t1$ onto the attributes in X , in a natural extension of our TRC notation (see Chapter 4): $t.a$ for referring to attribute a of tuple t . An FD $X \rightarrow Y$ essentially says that if two tuples agree on the values in attributes X , they must also agree on the values in attributes Y .

Figure 15.3 illustrates the meaning of the FD $AB \rightarrow C$ by showing an instance that satisfies this dependency. The first two tuples show that an FD is not the same as a key constraint: Although the FD is not violated, AB is clearly not a key for the relation. The third and fourth tuples illustrate that if two tuples differ in either the A field or the B field, they can differ in the C field without violating the FD. On the other hand, if we add a tuple $ha1; b1; c2; d1$ to the instance shown in this figure, the resulting instance would violate the FD; to see this violation, compare the first tuple in the figure with the new tuple.

A	B	C	D
a1	b1	c1	d1
a1	b1	c1	d2
a1	b2	c2	d1
a2	b1	c3	d1

An Instance that Satisfies $AB \rightarrow C$

Recall that a *legal* instance of a relation must satisfy all specified ICs, including all specified FDs. As noted in Section 3.2, ICs must be identified and specified based on the semantics of the real-world enterprise being modeled. By looking at an instance of a relation, we might be able to tell that a certain FD does *not* hold. However, we can never deduce that an FD *does* hold by looking at one or more instances of the

relation because an FD, like other ICs, is a statement about *all* possible legal instances of the relation.

A primary key constraint is a special case of an FD. The attributes in the key play the role of X , and the set of all attributes in the relation plays the role of Y . Note, however, that the definition of an FD does not require that the set X be minimal; the additional minimality condition must be met for X to be a key. If $X \twoheadrightarrow Y$ holds, where Y is the set of all attributes, and there is some subset V of X such that $V \twoheadrightarrow Y$ holds, then X is a *super key*; if V is a strict subset of X , then X is not a key.

In the rest of this chapter, we will see several examples of FDs that are not key constraints.

REASONING ABOUT FUNCTIONAL DEPENDENCIES

The discussion up to this point has highlighted the need for techniques that allow us to carefully examine and further refine relations obtained through ER design (or, for that matter, through other approaches to conceptual design).

Before proceeding with the main task at hand, which is the discussion of such schema refinement techniques, we digress to examine FDs in more detail because they play such a central role in schema analysis and refinement.

Given a set of FDs over a relation schema R , there are typically several additional FDs that hold over R whenever all of the given FDs hold. As an example, consider:

$Workers(\underline{ssn}, name, lot, did, since)$

We know that $ssn \twoheadrightarrow did$ holds, since ssn is the key, and FD $did \twoheadrightarrow lot$ is given to hold. Therefore, in any legal instance of $Workers$, if two tuples have the same ssn value, they must have the same did value (from the first FD), and because they have the same did value, they must also have the same lot value (from the second FD). Thus, the FD $ssn \twoheadrightarrow lot$ also holds on $Workers$.

We say that an FD f is implied by a given set F of FDs if f holds on every relation instance that satisfies all dependencies in F , that is, f holds whenever all FDs in F hold. Note that it is not sufficient for f to hold on some instance that satisfies all dependencies in F ; rather, f must hold on every instance that satisfies all dependencies in F .

Closure of a Set of FDs

The set of all FDs implied by a given set F of FDs is called the closure of F and is denoted as F^+ . An important question is how we can infer, or compute,

the closure of a given set F of FDs. The answer is simple and elegant. The following three rules, called Armstrong's Axioms, can be applied repeatedly to infer all FDs implied by a set F of FDs. We use X , Y , and Z to denote *sets* of attributes over a relation schema R :

Reflexivity: If XY , then $X \twoheadrightarrow Y$.

Augmentation: If $X \twoheadrightarrow Y$, then $XZ \twoheadrightarrow YZ$ for any Z . Transitivity:
If $X \twoheadrightarrow Y$ and $Y \twoheadrightarrow Z$, then $X \twoheadrightarrow Z$.

+
Armstrong's Axioms are sound in that they generate only FDs in F when applied to a set F of FDs. They are complete in that repeated application of these rules will

+ generate all FDs in the closure F . (We will not prove these claims.) It is convenient

+
to use some additional rules while reasoning about F :

Union: If $X \twoheadrightarrow Y$ and $X \twoheadrightarrow Z$, then $X \twoheadrightarrow YZ$.

■
Decomposition: If $X \twoheadrightarrow YZ$, then $X \twoheadrightarrow Y$ and $X \twoheadrightarrow Z$.

These additional rules are not essential; their soundness can be proved using Armstrong's Axioms.

use a more elaborate version of the Contracts relation:

Contracts(contractid, supplierid, projectid, deptid, partid, qty, value)

We denote the schema for Contracts as $CSJDPQV$. The meaning of a tuple in this relation is that the contract with *contractid* C is an agreement that supplier S (*supplierid*) will supply Q items of part P (*partid*) to project J (*projectid*)

associated with department D (*deptid*); the value V of this contract is equal to *value*.

The following ICs are known to hold:

1. The contract id C is a key: $C! CSJDPQV$.
2. A project purchases a given part using a single contract: $JP!C$.
3. A department purchases at most one part from a supplier: $SD!P$.

Several additional FDs hold in the closure of the set of given FDs:

From $JP!C, C!CSJDPQV$ and transitivity, we infer $JP!CSJDPQV$.

From $SD!P$ and augmentation, we infer $SDJ!JP$.

From $SDJ!JP, JP!CSJDPQV$ and transitivity, we infer $SDJ!CSJDPQV$. (Incidentally, while it may appear tempting to do so, we *cannot* conclude $SD!CSDPQV$, canceling J on both sides. FD inference is not like arithmetic multiplication!)

We can infer several additional FDs that are in the closure by using augmentation

or

decomposition. For example, from $C!CSJDPQV$, using decomposition we can infer

:

$C!C, C!S, C!J, C!D$, etc.

Finally, we have an number of trivial FDs from the reflexivity rule.

NORMAL FORMS

Given a relation schema, we need to decide whether it is a good design or whether we need to decompose it into smaller relations. Such a decision must be guided by an understanding of what problems, if any, arise from the current schema. To provide such guidance, several normal forms have been proposed. If a relation schema is in one of these normal forms, we know that certain kinds of problems cannot arise.

The normal forms based on FDs are *first normal form (1NF)*, *second normal form (2NF)*, *third normal form (3NF)*, and *Boyce-Codd normal form (BCNF)*. These forms have increasingly restrictive requirements: Every relation in BCNF is also in 3NF, every relation in 3NF is also in 2NF, and every relation in 2NF is in 1NF. A relation is in first normal form if every field contains only atomic values, that is, not lists or sets. This requirement is implicit in our definition of the relational model. Although some of the newer database systems are relaxing this requirement, in this chapter we will assume that it always holds. 2NF is mainly of historical interest. 3NF and BCNF are important from a database design standpoint.

While studying normal forms, it is important to appreciate the role played by FDs. Consider a relation schema R with attributes ABC . In the absence of any ICs, any set of ternary tuples is a legal instance and there is no potential for redundancy. On the other hand, suppose that we have the FD $A \twoheadrightarrow B$. Now if several tuples have the same A value, they must also have the same B value. This potential redundancy can be predicted using the FD information. If more detailed ICs are specified, we may be able to detect more subtle redundancies as well.

We will primarily discuss redundancy that is revealed by FD information. In Section 15.8, we discuss more sophisticated ICs called *multivalued dependencies* and *join dependencies* and normal forms based on them.

15.5.1 Boyce-Codd Normal Form

Let R be a relation schema, X be a subset of the attributes of R , and let A be an attribute of R . R is in Boyce-Codd normal form iff for every FD $X \twoheadrightarrow A$ that holds over R , one of the following statements is true:

$A \in X$; that is, it is a trivial FD, or

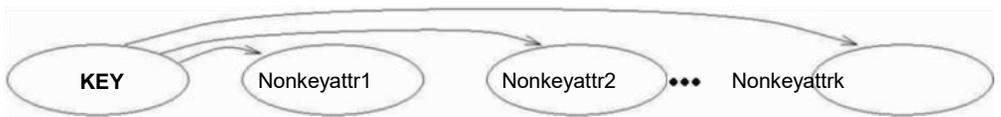
■ X is a super key.

Note that if we are given a set F of FDs, according to this definition, we must

+ consider each dependency $X \twoheadrightarrow A$ in the closure F^+ to determine whether R is in BCNF. However, we can prove that it is sufficient to check whether the left side of each dependency in F is a super key (by computing the attribute closure and seeing if it includes all attributes of R).

Intuitively, in a BCNF relation the only nontrivial dependencies are those in which a key determines some attribute(s). Thus, each tuple can be thought of as an entity or relationship, identified by a key and described by the remaining attributes. Kent puts this colorfully, if a little loosely: "Each attribute must describe [an entity or relationship identified by] the key, the whole key, and nothing but the key." If we use

ovals denote attributes or sets of attributes and draw arcs to indicate FDs, a relation in BCNF has the structure illustrated in Figure, considering just one key for simplicity. (If there are several candidate keys, each candidate key can play the role of KEY in the figure, with the other attributes being the ones not in the chosen candidate key.)



FDs in a BCNF Relation

BCNF ensures that no redundancy can be detected using FD information alone. It is thus the most desirable normal form (from the point of view of redundancy) if we take into account only FD information. This point is illustrated in Figure 15.8.

X	Y	A
-----	-----	-----

x	$y1$	a
x	$y2$	$?$

Instance Illustrating BCNF

This figure shows (two tuples in) an instance of a relation with three attributes X , Y , and A . There are two tuples with the same value in the X column. Now suppose that

we know that this instance satisfies an FD $X \rightarrow A$. We can see that one of the tuples has the value a in the A column. What can we infer about the value in the A column in the second tuple? Using the FD, we can conclude that the second tuple also has the value a in this column. (Note that this is really the only kind of inference we can make about values in the fields of tuples by using FDs.)

But isn't this situation an example of redundancy? We appear to have stored the value a twice. Can such a situation arise in a BCNF relation? No! If this relation is in BCNF, because A is distinct from X it follows that X must be a key. (Otherwise, the FD $X \rightarrow A$ would violate BCNF.) If X is a key, then $y_1 = y_2$, which means that the two tuples are identical. Since a relation is defined to be a set of tuples, we cannot have two copies of the same tuple and the situation shown in Figure 15.8 cannot arise.

Thus, if a relation is in BCNF, every field of every tuple records a piece of information that cannot be inferred (using only FDs) from the values in all other fields in (all tuples of) the relation instance.

Third Normal Form

Let R be a relation schema, X be a subset of the attributes of R , and A be an attribute of R . R is in third normal form if for every FD $X \rightarrow A$ that holds over R , one of the following statements is true:

$A \in X$; that is, it is a trivial FD, or

-
- X is a superkey, or

A is part of some key for R.

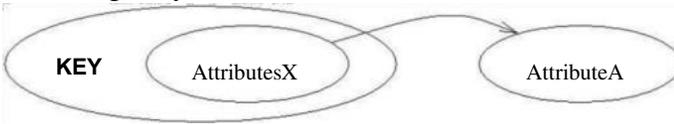
The definition of 3NF is similar to that of BCNF, with the only difference being the third condition. Every BCNF relation is also in 3NF. To understand the third condition, recall that a key for a relation is a *minimal* set of attributes that uniquely determines all other attributes.

A must be part of a key (any key, if there are several).

It is not enough for A to be part of a superkey, because the latter condition is satisfied by each and every attribute! Finding all keys of a relation schema is known to be an NP-complete problem, and so is the problem of determining whether a relation schema is in 3NF.

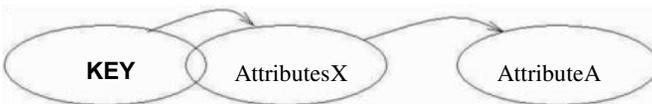
Partial dependencies are illustrated in Figure 15.9, and transitive dependencies are

illustrated in Figure 15.10. Note that in Figure 15.10, the set X of attributes may or may not have some attributes in common with KEY; the diagram should be interpreted as indicating only that X is not a subset of KEY.

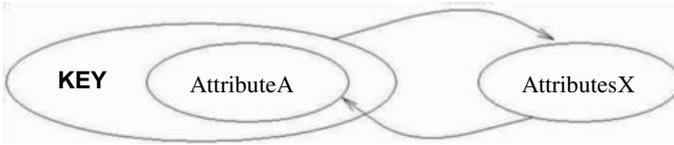


Case 1: A not in KEY

Figure 15.9 Partial Dependencies



Case 1: A not in KEY



Case 2: A is in KEY

Transitive Dependencies

The motivation for 3NF is rather technical. By making an exception for certain dependencies involving key attributes, we can ensure that every relation schema can

be decomposed into a collection of 3NF relations using only decomposition that have

certain desirable properties (Section 15.6). Such a guarantee does not exist for BCNF relations; the 3NF definition weakens the BCNF requirements just enough to make this guarantee possible. We may therefore compromise by settling for a 3NF design. As we shall see in Chapter 16, we may sometimes accept this compromise (or even settle for a non-3NF schema) for other reasons as well.

DECOMPOSITIONS

As we have seen, a relation in BCNF is free of redundancy (to be precise, redundancy that can be detected using FD information), and a relation schema in 3NF comes close.

If a relation schema is not in one of these normal forms, the FDs that cause a violation can give us insight into the potential problems. The main technique for addressing such redundancy-related problems is decomposing a relation schema into relation schemas with fewer attributes.

A decomposition of a relation schema R consists of replacing the relation schema by two (or more) relation schemas that each contain a subset of the attributes of R and

together include all attributes in R . Intuitively, we want to store the information in any given instance of R by storing projections of the instance. This section examines the use of decompositions through several examples.

We begin with the Hourly Emps example from Section 15.3.1. This relation has attributes $SNLRWH$ and two FDs: $S \twoheadrightarrow SNLRWH$ and $R \twoheadrightarrow W$. Since R is not a key and W is not part of any key, the second dependency causes a violation of 3NF.

The alternative design consisted of replacing Hourly Emps with two relations having attributes $SNLRH$ and RW . $S \twoheadrightarrow SNLRH$ holds over $SNLRH$, and S is a key. $R \twoheadrightarrow W$ holds over RW , and R is a key for RW . The only other dependencies that hold over these schemas are those obtained by augmentation. Thus both schemas are in BCNF.

Our decision to decompose $SNLRWH$ into $SNLRH$ and RW , rather than, say, $SNLR$ and $LRWH$, was not just a good guess. It was guided by the observation that the

dependency $R \twoheadrightarrow W$ caused the violation of 3NF; the most natural way to deal with this violation is to remove the attribute W from this schema. To compensate for removing W from the main schema, we can add a relation RW , because each R value is associated with at most one W value according to the FD $R \twoheadrightarrow W$.

A very important question must be asked at this point: If we replace a legal instance r of relation schema $SNLRWH$ with its projections on $SNLRH$ (r_1) and RW (r_2), can we recover r from r_1 and r_2 ? The decision to decompose $SNLRWH$ into $SNLRH$ and RW is equivalent to saying that we will store instances r_1 and r_2 instead of r . However

it is the instance r that captures the intended entities or relationships. If we cannot compute r from r_1 and r_2 , our attempt to deal with redundancy has effectively thrown out the baby with the bathwater. We consider this issue in more detail below.

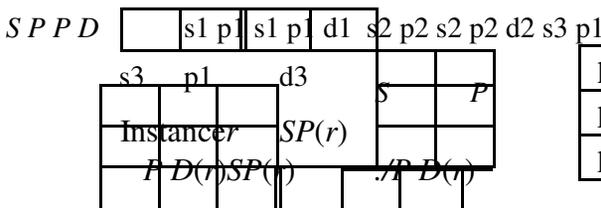
Lossless-Join Decomposition

Let R be a relation.

A decomposition of R into two schemas with attribute sets X and Y is said to be a lossless join decomposition with respect to F if for every instance r of R that satisfies the dependencies in F , $X(r) \bowtie Y(r) = r$.

This definition can easily be extended to cover a decomposition of R into more than two relations. It is easy to see that $X(r) \bowtie Y(r)$ always holds. In general, though, the other direction does not hold. If we take projections of a relation and recombine them using natural join, we typically obtain some tuples that were not in the original relation. This situation is illustrated in Figure 15.11.

S	P	D
-----	-----	-----



p1	d1
p2	d2
p1	d3

D

s1	p1	d1
s2	p2	d2
s3	p1	d3
s1	p1	d3
s3	p1	d1

By replacing the instances shown in Figure 15.11 with the instances $SP(r)$ and $PD(r)$, we lose some information. In particular, suppose that the tuples in r denote relationships. We can no longer tell that the relationships $(s1;p1;d3)$ and $(s3;p1;d1)$ do not hold. The decomposition of schema SPD into SP and PD is therefore a 'lossy' decomposition if the instance r shown in the figure is legal, that is, if this instance could arise in the enterprise being modeled. (Observe the similarities between this example and the Contracts relationship set in Section 2.5.3.)

All decompositions used to eliminate redundancy must be lossless. The following simple test is very useful:

Let R be a relation and F be a set of FDs that hold over R . The decomposition of

R into relations with attribute sets R_1 and R_2 is lossless if

+ and only if F contains either the FD $R_1 \setminus R_2 \twoheadrightarrow R_1$ or the FD $R_2 \setminus R_1 \twoheadrightarrow R_2$.

In other words, the attributes common to R_1 and R_2 must contain a key for either R_1 or R_2 . If a relation is decomposed into two relations, this test is a necessary and

2 sufficient

condition for the decomposition to be lossless-join. If a relation is decomposed into more than two relations, an efficient (time polynomial in the size of the dependency

set) algorithm is available to test whether or not the decomposition is lossless, but we will not discuss it.

Consider the Hourly Emps relation again. It has attributes $SNLRWH$, and the FD $R \twoheadrightarrow !$

W causes a violation of 3NF. We dealt with this violation by decomposing the relation into $SNLRH$ and RW . Since R is common to both decomposed relations, and $R \twoheadrightarrow ! W$ holds, this decomposition is lossless-join.

This example illustrates a general observation:

If an FD $X \rightarrow Y$ holds over a relation R and $X \setminus Y$ is empty, the decomposition of R into $R - Y$ and XY is lossless.

X appears in both $R - Y$ (since $X \setminus Y$ is empty) and XY , and it is a key for XY . Thus, the above observation follows from the test for a lossless-join decomposition.

Another important observation has to do with repeated decompositions. Suppose that a relation R is decomposed into $R1$ and $R2$ through a lossless-join decomposition, and that $R1$ is decomposed into $R11$ and $R12$ through another lossless-join decomposition. Then the decomposition of R into $R11$, $R12$, and $R2$ is lossless-join; by joining $R11$ and $R12$ we can recover $R1$, and by then

Let R be a relation and F be a set of FDs that hold over R . The decomposition of joining $R1$ and $R2$, we can recover R .

Dependency-Preserving Decomposition

Consider the Contracts relation with attributes $CSJDPQV$ from Section 15.4.1. The given FDs are $C \twoheadrightarrow CSJDPQV$, $JP \twoheadrightarrow C$, and $SD \twoheadrightarrow P$. Because SD is not a key the dependency $SD \twoheadrightarrow P$ causes a violation of BCNF.

We can decompose Contracts into two relations with schemas $CSJDPQV$ and SDP to address this violation; the decomposition is lossless join. There is one subtle problem, however. We can enforce the integrity constraint $JP \twoheadrightarrow C$ easily when a tuple is inserted into Contracts by ensuring that no existing tuple has the same JP values (as the inserted tuple) but different C values. Once we decompose Contracts into $CSJDPQV$ and SDP , enforcing this constraint requires an expensive join of the two relations whenever a tuple is inserted into $CSJDPQV$. We say that this decomposition is not dependency-preserving.

Intuitively, a dependency preserving decomposition allows us to enforce all FDs by examining a single relation instance on each insertion or modification of a tuple. (Note that deletions cannot cause violation of FDs.) To define dependency-preserving decompositions precisely, we have to introduce the concept of a projection of FDs.

Let R be a relation schema that is decomposed into two schemas with attribute sets X and Y , and let F be a set of FDs over R . The projection of F on X is the set of FDs in the closure F^+ (not just F !) that involve only attributes in X . We will denote the projection of F on attributes X as F_X . Note that a dependency $U \twoheadrightarrow V$ in F is in F_X only if *all* the attributes in U and V are in X .

The decomposition of relation schema R with FDs F into schemas with attribute sets

+ +

X and Y is dependency-preserving if $(FX \sqcup FY) = F$. That is, if we take the
Consider the Contracts relation with attributes $CSJDPQV$ from Section 15.4.1. The
dependencies in FX and FY and compute the closure of their union, we get back
all dependencies in the closure of F . Therefore, we need to enforce only the
+dependencies in FX and FY ; all FDs in F
are then sure to be satisfied. To enforce FX , we need to examine only relation
 X (on inserts to that relation). To enforce FY , we need to examine only relation
 Y .

NORMALIZATION

Having covered the concepts needed to understand the role of normal forms and decompositions in database design, we now consider algorithms for converting relations to BCNF or 3NF. If a relation schema is not in BCNF, it is possible to obtain a lossless-join decomposition into a collection of BCNF relation schemas. Unfortunately, there may not be any dependency-preserving decomposition into a collection of BCNF relation schemas

Decomposition into BCNF

We now present an algorithm for decomposing a relation schema R into a collection of BCNF relation schemas:

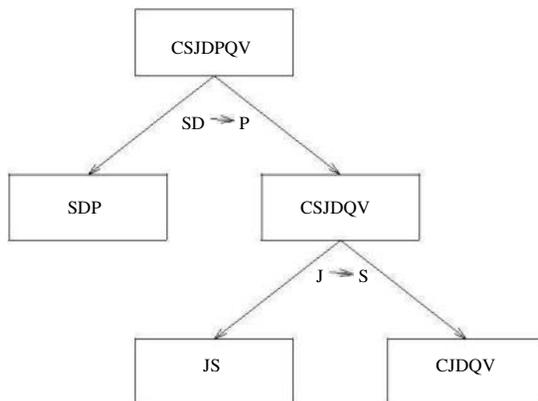
1. Suppose that R is not in BCNF. Let $X \rightarrow A$ be a single attribute in R , and X/A be a FD that causes a violation of BCNF. Decompose R into $R - \rightarrow A$ and $X \rightarrow A$.
2. If either $R - \rightarrow A$ or $X \rightarrow A$ is not in BCNF, decompose them further by recursive application of this algorithm.

$R - \rightarrow A$ denotes the set of attributes other than A in R , and $X \rightarrow A$ denotes the union of attributes in X and A . Since $X \rightarrow A$ violates BCNF, it is not a trivial dependency; further, A is a single attribute. Therefore, A is not in X ; that is, $X \cap A$ is empty. Thus, each decomposition carried out in Step 1 is lossless-join.

This set of dependencies associated with $R - \rightarrow A$ and $X \rightarrow A$ is the projection of F onto their attributes. If one of the new relations is not in BCNF, we decompose it further in Step 1. Since a decomposition results in relations with strictly fewer attributes, this process will terminate, leaving us with a collection of relation schemas that are all in BCNF.

Consider the Contracts relation with attributes $CSJDPQV$ and key C . We are given FDs $JP \neq C$ and $SD \neq P$. By using the dependency $SD \neq P$ to guide the decomposition, we get the two schemas SDP and $CSJDQV$. SDP is in BCNF. Suppose that we also have the constraint that each project deals with a single supplier: $J \neq S$. This means that the schema $CSJDQV$ is not in BCNF. So we decompose it further into JS and $CJDQV$. $C \neq JDQV$ holds over $CJDQV$; the only other FDs that hold are those obtained from this FD by augmentation, and therefore all FDs contain a key in the left side. Thus, each of the schemas SDP , JS , and $CJDQV$ is in BCNF, and this collection of schemas also represents a lossless-join decomposition of $CSJDQV$.

The steps in this decomposition process can be visualized as a tree, as shown in Figure. The root is the original relation $CSJDPQV$, and the leaves are the BCNF relations that are the result of the decomposition algorithm, namely, SDP , JS , and $CJDQV$. Intuitively, each internal node is replaced by its children through a single decomposition step that is guided by the FD shown just below the node.



Decomposition of $CSJDQV$ into SDP , JS , and $CJDQV$

Redundancy in BCNF Revisited

The decomposition of $CSJDQV$ into SDP , JS , and $CJDQV$ is not dependency preserving. Intuitively, dependency $JP \twoheadrightarrow C$ cannot be enforced without a join. One way to deal with this situation is to add a relation with attributes CJP . In effect, this solution amounts to storing some information redundantly in order to make the dependency enforcement cheaper.

This is a subtle point: Each of the schemas CJP , SDP , JS , and $CJDQV$ is in BCNF, yet there is some redundancy that can be predicted by FD information. In particular, if we join the relation instances for SDP and $CJDQV$ and project onto the attributes CJP , we must get exactly the instance stored in the relation with schema CJP . We saw in Section 15.5.1 that there is no such redundancy within a single BCNF relation. The current example shows that redundancy can still occur across relations, even though there is no redundancy within a relation.

Minimal Cover for a Set of FDs

A minimal cover for a set F of FDs is a set G of FDs such that:

1. Every dependency in G is of the form $X \twoheadrightarrow A$, where A is a single attribute.
2. The closure F^+ is equal to the closure G^+ .
3. If we obtain a set H of dependencies from G by deleting one or more dependencies, then $H^+ \neq G^+$.

Intuitively, a minimal cover for a set F of FDs is an equivalent set of dependencies that is *minimal* in two respects: (1) Every dependency is as small as possible; that is, each attribute on the left side is necessary and the right side is a single attribute.

(2) Every dependency in it is required in order for the closure to be equal to F^+ . As an example, let F be the set of dependencies:

$A!B, ABCD!E, EF!G, EF!H,$ and $ACDF!EG.$

First, let us rewrite $ACDF!EG$ so that every right side is a single attribute:

$ACDF!E$ and $ACDF!G.$

Next consider $ACDF!G.$ This dependency is implied by the following FDs:

$A!B, ABCD!E,$ and $EF!G.$

Therefore, we can delete it. Similarly, we can delete $ACDF!E.$ Next consider $ABCD!E.$ Since $A!B$ holds, we can replace it with $ACD!E.$ (At this point, the readers should verify that each remaining FD is minimal and required.) Thus, a minimal cover for F is the set:

$A!B, ACD!E, EF!G,$ and $EF!H.$

The preceding examples suggest a general algorithm for obtaining a minimal cover of a set F of FDs:

1. Put the FDs in a standard form: Obtain a collection G of equivalent FDs with a single attribute on the right side (using the decomposition axiom).
2. Minimize the left side of each FD: For each FD in G , check each attribute in the left side to see if it can be deleted while preserving equivalence to F .
3. Delete redundant FDs: Check each remaining FD in G to see if it can be deleted while preserving equivalence to F .

Note that the order in which we consider FDs while applying these steps could produce different minimal covers; there could be several minimal covers for a given set of FDs.

More important, it is necessary to minimize the left sides of FDs *before* checking for redundant

FDs. If these two steps are reversed, then a set of FDs could still contain some redundant FDs (i.e., not be a minimal cover), as the following example illustrates. Let F be the set of dependencies, each of which is already in the standard form:

$ABCD \rightarrow E, E \rightarrow D, A \rightarrow B, \text{ and } AC \rightarrow D.$

Observe that none of these FDs is redundant; if we checked for redundant FDs first, we would get the same set of FDs F . The left side of $ABCD \rightarrow E$ can be replaced by AC while preserving equivalence to F , and we would stop here if we checked for redundant FDs in F before minimizing the left sides. However, the set of FDs we have is not a minimal cover:

$AC \rightarrow E$, $E \rightarrow D$, $A \rightarrow B$, and $AC \rightarrow D$.

From transitivity, the first two FDs imply the last FD, which can therefore be deleted + while preserving equivalence to F . The important point to note is that $AC \rightarrow D$ becomes redundant only after we replace $AC \rightarrow E$ with $AC \rightarrow D$. If we minimize left sides of FDs first and then check for redundant FDs, we are left with the first three FDs in the preceding list, which is indeed a minimal cover for F .

Dependency-Preserving Decomposition into 3NF

Returning to the problem of obtaining a lossless-join, dependency-preserving decomposition into 3NF relations, let R be a relation with a set F of FDs that is a minimal cover, and let $R_1; R_2; \dots; R_n$ be a lossless-join decomposition of R . For $1 \leq i \leq n$, suppose that each R_i is in 3NF and let F_i denote the projection of F onto the attributes of R_i . Do the following:

Identify the set N of dependencies in F that are not preserved, that is, not included in the closure of the union of F_i s.

For each FD $X \rightarrow A$ in N , create a relation schema XA and add it to the decomposition of R .

Obviously, every dependency in F is preserved if we replace R by the R_i s plus the schemas of the form XA added in this step. The R_i s are given to be in 3NF. We can show that each of these schemas XA is in 3NF as follows: Since $X \rightarrow A$ is in the minimal cover F , $Y \rightarrow A$ does not hold for any Y that is a strict subset of X . Therefore, X is a key for XA . Further, if any other dependencies hold over XA , the right side can involve only attributes in X because A is a single attribute (because $X \rightarrow A$ is an FD in a minimal cover). Since X is a key for XA , none of these additional dependencies causes a violation of 3NF (although they might cause a violation of BCNF).

As an optimization, if these N contains several FDs with the same left side, say, $X_1 \rightarrow A_1, X_1 \rightarrow A_2, \dots, X_1 \rightarrow A_n$, we can replace them with a single equivalent FD $X_1 \rightarrow A_1 \dots A_n$. Therefore, we produce one relation schema $X_1 \rightarrow A_1 \dots A_n$, instead of several schemas $X_1 \rightarrow A_1; \dots; X_1 \rightarrow A_n$, which is generally preferable.

Consider the *Contracts* relation with attributes $CSJDPQV$ and FDs $JP \rightarrow C, SD \rightarrow P$, and $J \rightarrow S$. If

we decompose $CSJDPQV$ into SDP and $CSJDQV$, then SDP is in BCNF, but $CSJDQV$ is not even in 3NF. So we decompose it further into JS and $CJDQV$.

The relation schemas SDP , JS , and $CJDQV$ are in 3NF (in fact, in BCNF), and the decomposition is lossless-join. However, the dependency $JP \rightarrow C$ is not preserved. This problem can be addressed by adding a relation schema CJP to the decomposition.

This set of FDs is not a minimal cover, and so we must do more. We first replace $CSJDPQV$ with the FDs:

$C \rightarrow S, C \rightarrow J, C \rightarrow D, C \rightarrow P, C \rightarrow Q$, and $C \rightarrow V$.

The FD $C \rightarrow P$ is implied by $C \rightarrow S, C \rightarrow D$, and $SD \rightarrow P$; so we can delete it. The FD $C \rightarrow S$ is implied by $C \rightarrow J$ and $J \rightarrow S$; so we can delete it. This leaves us with a minimal cover:

$C \rightarrow J, C \rightarrow D, C \rightarrow Q, C \rightarrow V, JP \rightarrow C, SD \rightarrow P$, and $J \rightarrow S$.

Using the algorithm for ensuring dependency-preservation, we obtain the relational schema CJ, CD, CQ, CV, CJP, SDP , and JS . We can improve this schema by combining relations for which C is the key into $CDJPQV$. In addition, we have SDP and JS in our decomposition. Since one of these relations ($CDJPQV$) is a super key, we are done.

Comparing this decomposition with the one that we obtained earlier in this section, we find that they are quite close, with the only difference being that one of them has

CDJPQV instead of *CJP* and *CJDQV*. In general, however, there could be significant differences.

Database designers typically use a conceptual design methodology (e.g., ER design) to arrive at an initial database design. Given this, the approach of repeated decomposition to rectify instances of redundancy is likely to be the most natural use of FDs and normalization techniques. However, a designer can also consider the alternative designs suggested by the synthesis approach.

OTHER KINDS OF DEPENDENCIES*

Ds are probably the most common and important kind of constraint from the point of view

of database design. However, there are several other kinds of dependencies. In particular, there is a well-developed theory for database design using *multivalued dependencies* and *join*

dependencies. By taking such dependencies into account, we can identify potential redundancy problems that cannot be detected using FDs alone.

This section illustrates the kinds of redundancy that can be detected using multivalued dependencies. Our main observation, however, is that simple guidelines (which can be checked using only FD reasoning) can tell us whether we even need to worry about complex constraints such as multivalued and join dependencies. We also comment on the role of *inclusion dependencies* in database design.

Multivalued Dependencies

Suppose that we have a relation with attributes *course*, *teacher*, and *book*, which we denote as *CTB*. The meaning of a tuple is that teacher *T* can teach course *C*, and book *B* is a recommended text for the course. There are no FDs; the key is *CTB*. However, the recommended texts for a course are independent of the instructor. The instance shown in Figure 15.13 illustrates this situation.

<i>course</i>	<i>teacher</i>	<i>book</i>
Physics101	Green	Mechanics
Physics101	Green	Optics
Physics101	Brown	Mechanics

Physics101	Brown	Optics
Math301	Green	Mechanics
Math301	Green	Vectors
Math301	Green	Geometry

BCNF Relation with Redundancy That Is Revealed by MVDs

There are three points to note here:

The relation schema CTB is in BCNF; thus we would not consider decomposing it further if we looked only at the FDs that hold over CTB .

There is redundancy. The fact that Green can teach Physics101 is recorded once per recommended text for the course. Similarly, the fact that Optics is a text for Physics101 is recorded once per potential teacher.

■ The redundancy can be eliminated by decomposing CTB into CT and CB .

The redundancy in this example is due to the constraint that the texts for a course are independent of the instructors, which cannot be expressed in terms of FDs. This constraint is an example of a *multivalued dependency*, or MVD.

Ideally, we should model this situation using two binary relationship sets, $Instructors$ with attributes CT and $Text$ with attributes CB . Because these are two essentially independent relationships, modeling them with a single ternary relationship set with attributes CTB is inappropriate. (See Section 2.5.3 for a further discussion of ternary versus binary relationships.) Given the subjectivity of ER design, however, we might create a ternary relationship. A careful analysis of the MVD information would then reveal the problem.

Let R be a relation schema and let X and Y be subsets of the attributes of R . Intuitively, the multivalued dependency $X \twoheadrightarrow Y$ is said to hold over R if, in every legal instance r of R , each X value is associated with a set of Y values and this set is independent of the values in the other attributes.

Formally, if the MVD $X \twoheadrightarrow Y$ holds over R and $Z = R - XY$, the following must be true for every legal instance r of R :

If $t_1, t_2 \in r$ and $t_1[X] = t_2[X]$, then there must be some $t_3 \in r$ such that $t_1[Y] = t_3[Y]$ and $t_2[Z] = t_3[Z]$.

Figure illustrates this definition. If we are given the first two tuples and told that the MVD $X \twoheadrightarrow Y$ holds over this relation, we can infer that the relation instance must also contain the third tuple. Indeed, by interchanging the roles of the first two tuples treating

the first tuple as t_2 and the second tuple as t_1 | we can deduce that the tuple t_4 must also be in the relation instance.

X	Y	Z	
a	b 1	c1	tuple t_1
a	b2	c2	tuple t_2
a	b 1	c2	tuple t_3
a	b 2	c1	tuple t_4

Illustration of MVD Definition

This table suggests another way to think about MVDs: If $X \twoheadrightarrow Y$ hold over R , then $Y \perp Z \mid X=x(R) = Y \perp Z \mid X=x(R)$ in every legal instance of R , for any value x that appears in the X column of R . In other words, consider groups of tuples in R with the same X -value, for each X -value. In each such group consider the projection onto the attributes YZ . This projection must be equal to the crossproduct of the projections onto Y and Z . That is, for a given X -value, the Y values and Z -values are independent.

(From this definition it is easy to see that $X \twoheadrightarrow Y$ must hold whenever $X \rightarrow Y$ holds. If the FD $X \rightarrow Y$ holds, there is exactly one Y -value for a given X -value, and the conditions in the MVD definition hold trivially. The converse does not hold, as Figure 15.14 illustrates.)

Returning to our CTB example, the constraint that course texts are independent of instructors can be expressed as $C \parallel T$. In terms of the definition of MVDs, this constraint can be read as follows:

\forall (there is a tuple showing that) C is taught by teacher T ,
 and (there is a tuple showing that) C has book B as text,
 then (there is a tuple showing that) C is taught by T and has text B .
 Given a set of FDs and MVDs, in general we can infer that several additional FDs and MVDs hold. A sound and complete set of inference rules consists of the three Armstrong Axioms plus several additional rules. Three of the additional rules involve only MVDs:

MVD Complementation: If $X \parallel Y$, then $X \parallel R - XY$.

■

MVD Augmentation: If $X \parallel Y$ and WZ , then $WX \parallel YZ$.

MVD Transitivity: If $X \parallel Y$ and $Y \parallel Z$,

then $X \parallel (Z - Y)$.

As an example of the use of these rules, since we have $C \parallel T$ over CTB , MVD complementation allows us to infer that $C \parallel CTB - CT$ as well, that is, $C \parallel B$. The remaining two rules relate FDs and MVDs:

Replication: If $X \parallel Y$, then $X \parallel Y$.

■

Coalescence: If $X \parallel Y$ and there is a W such that $W \cap Y$ is empty, $W \parallel Z$, and $Y \parallel Z$, then $X \parallel Z$.

Observe that replication states that every FD is also an MVD.

Fourth Normal Form

Fourth normal form is a direct generalization of BCNF. Let R be a relation schema, X and Y be nonempty subsets of the attributes of R , and F be a set of dependencies that includes both FDs and MVDs. R is said to be in fourth normal form (4NF) if for every MVD $X \twoheadrightarrow Y$ that holds over R , one of the following statements is true:

$Y \subseteq X$ or $XY = R$, or

■ X is a Superkey.

In reading this definition, it is important to understand that the definition of a key has not changed; the key must uniquely determine all attributes through FDs alone. $X \twoheadrightarrow Y$ is a trivial MVD if $Y \subseteq X$ or $XY = R$; such MVDs always hold.

The relation CTB is not in 4NF because $C \twoheadrightarrow T$ is a nontrivial MVD and C is not a key. We can eliminate the resulting redundancy by decomposing CTB into CT and CB ; each of these relations is then in 4NF.

To use MVD information fully, we must understand the theory of MVDs. However, the following result due to Date and Fagin identifies conditions detected using only FD information under which we can safely ignore MVD information. That is, using MVD information in addition to the FD information will not reveal any redundancy. Therefore, if these conditions hold, we do not even need to identify all MVDs.

If a relation schema is in BCNF, and at least one of its keys consists of a single attribute, it is also in 4NF.

An important assumption is implicit in any application of the preceding result: *The set of FDs identified thus far is indeed the set of all FDs that hold over the relation.*

This assumption is important because the result relies on the relation being in BCNF, which in turn depends on the set of FDs that hold over the relation.

Figure shows three tuples from an instance of $ABCD$ that satisfies the given MVDB $A!C$. From the definition of an MVD, given tuples t_1 and t_2 , it follows

B	C	A	D	
-----	-----	-----	-----	--

b	c_1	a_1	d_1	tuple t_1
b	c_2	a_2	d_2	tuple t_2
b	c_1	a_2	d_2	tuple t_3

Three Tuples from a Legal Instance of $ABCD$

that tuple t_3 must also be included in the instance. Consider tuples t_2 and t_3 . From the given FDA/BCD and the fact that these tuples have the same A -value, we can deduce that $c_1 = c_2$. Thus, we see that the FD $B!C$ must hold over $ABCD$ whenever the FD $A!BCD$ and the MVDB $A!C$ hold. If $B!C$ holds, the relation $ABCD$ is not in BCNF (unless additional FDs hold that make B a key)!

Join Dependencies

A join dependency is a further generalization of MVDs. A join dependency (JD) $\Join R_1; \dots; R_n$ is said to hold over a relation R if $R_1; \dots; R_n$ is a lossless-join decomposition of R .

An MVD $X \twoheadrightarrow Y$ over a relation R can be expressed as the join dependency $\Join XY, X(R-Y)$. As an example, in the CTB relation, the MVD $C \twoheadrightarrow T$ can be expressed as the join dependency $\Join CT, CB$.

Unlike FDs and MVDs, there is no set of sound and complete inference rules for JDs.

Fifth Normal Form

A relation schema R is said to be in fifth normal form (5NF) if for every JD $\Join R_1; \dots; R_n$ that holds over R , one of the following statements is true:

■ $R_i = R$ for some i , or

The JD is implied by the set of those FDs over R in which the left side is a key for R .

The second condition deserves some explanation, since we have not presented inference rules for FDs and JDs taken together. Intuitively, we must be able to show that the decomposition of R into $R_1; \dots; R_n$ is lossless join whenever the key dependencies (FDs in which the left side is a key for R) hold. $\Join R_1; \dots; R_n$ is a trivial JD if $R_i = R$ for some i ; such a JD always holds.

The following result, also due to Date and Fagin, identifies conditions again, detected using only FD information under which we can safely ignore JD information.

If a relation schema is in 3NF and each of its keys consists of a single attribute, it is also in 5NF.

The conditions identified in this result are sufficient for a relation to be in 5NF, but not necessary. The result can be very useful in practice because it allows us to conclude that a relation is in 5NF *without ever identifying the MVDs and JDs that may hold over the relation.*

UNITV:TRANSACTIONSANDCONCURRENCYCONTROL

ACIDProperties:ConsistencyandIsolation,Atomicityand Durability, Transactions and Schedules

Serializability, Recoverability, Testing for serializability,

Concurrent Execution of Transactions, Lock-Based

ConcurrencyControl,PerformanceofLocking,Transaction Support in SQL. Deadlocks, Recovery and Atomicity..

Transactions are a set of operations used to perform a logical set of work. A transaction usually means that the data in the database has changed. One of the major uses of DBMS is to protect the user's data from system failures. It is done by ensuring that all the data is restored to a consistent state when the computer is restarted after a crash. The transaction is any one execution of the user program in a DBMS. Executing the same program multiple times will generate multiple transactions.

Example –

Transaction to be performed to withdraw cash from an ATM vestibule.

Example: Suppose an employee of bank transfers Rs 800 from X's account to Y's account.

This small transaction contains several low-level tasks:

X'account:

Open_Account(X)

Old_Balance = X.balance

New_Balance=Old_Balance-800 X.balance =

New_Balance Close_Account(X)

Y'sAccount

DepartmentOfAIDS

PreparedBy:MNANDINI DBMS

Open_Account(Y)

Old_Balance=Y.balance

New_Balance=Old_Balance+800 Y.balance =

New_Balance Close_Account(Y)

Operations of Transaction:

Following are the main operations of transaction:

Read(X): Read operation is used to read the value of X from the database and stores it in a buffer in main memory.

Write(X): Write operation is used to write the value back to the database from the buffer.

Let's take an example to debit transaction from an account which consists of following operations:

R(X); X=

X - 500; W(X);

Let's assume the value of X before starting of the transaction is 4000.

- The first operation reads X's value from database and stores it in a buffer.
- The second operation will decrease the value of X by 500. So buffer will contain 3500. ○

The third operation will write the buffer's value to the database. So X's final value will be 3500.

But it may be possible that because of the failure of hardware, software or power, etc. that transaction may fail before finished all the operations in the set.

For example: If in the above transaction, the debit transaction fails after executing operation 2 then X's value will remain 4000 in the database which is not acceptable by the bank.

To solve this problem, we have two important operations:

Commit: It is used to save the work done permanently.

Rollback: It is used to undo the work done.

Topic II:ACID Properties:

A [transaction](#) is a single logical unit of work that accesses and possibly modifies the contents of a database. Transactions access data using read and write operations. In order to maintain consistency in a database, before and after the transaction, certain properties are followed. These are called **ACID** properties.

Atomicity:

By this, we mean that either the entire transaction takes place at once or doesn't happen at all. There is no midway i.e. transactions do not occur partially. Each transaction is considered as one unit and either runs to completion or is not executed at all. It involves the following two operations.

—**Abort:** If a transaction aborts, changes made to the database are not visible.

—**Commit:** If a transaction commits, changes made are visible. Atomicity is also known as the 'All or nothing rule'.

Consider the following transaction **T** consisting of **T1** and **T2**: Transfer of 100 from account **X** to account **Y**.

Before: X : 500	Y: 200
Transaction T	
T1	T2
Read (X) X: = X - 100 Write (X)	Read (Y) Y: = Y + 100 Write (Y)
After: X : 400	Y : 300

If the transaction fails after completion of **T1** but before completion of **T2**. (say, after **write(X)** but before **write(Y)**), then the amount has been deducted from **X** but not added to **Y**. This results in an inconsistent database state. Therefore, the transaction must be executed in its entirety in order to ensure the correctness of the database state.

Consistency:

This means that integrity constraints must be maintained so that the database is consistent before and after the transaction. It refers to the correctness of a database. Referring to the example above,

the total amount before and after the transaction must be maintained.

Total before T occurs = $500 + 200 = 700$.

Total after T occurs = $400 + 300 = 700$. Therefore, the database is **consistent**. Inconsistency occurs in case **T1** completes but **T2** fails. As a result, T is incomplete.

Isolation:

This property ensures that multiple transactions can occur concurrently without leading to the inconsistency of the database state. Transactions occur independently without interference. Changes occurring in a particular transaction will not be visible to any other transaction until that particular change in that transaction is written to memory or has been committed. This property ensures that the execution of transactions

concurrently will result in a state that is equivalent to a state achieved these were executed serially in some order.

Let $X = 500$, $Y = 500$.

Consider two transactions T and T'' .

T	T''
Read (X)	Read (X)
$X := X * 100$	Read (Y)
Write (X)	$Z := X + Y$
Read (Y)	Write (Z)
$Y := Y - 50$	
Write (Y)	

Suppose T has been executed till **Read (Y)** and then T'' starts. As a result, interleaving of operations takes place due to which T'' reads the correct value of X but the incorrect value of

Y and sum computed by

T'' : $(X+Y = 50, 000+500=50, 500)$

is thus not consistent with the sum at end of the transaction:

T : $(X+Y=50,000+450=50,450)$.

This results in database inconsistency, due to a loss of 50 units. Hence,

transactions must take place in isolation and changes should be visible only after they have been made to the main memory.

Durability:

This property ensures that once the transaction has completed execution, the updates and modifications to the database are stored in and written to disk and they persist even if a system failure occurs. These updates now become permanent and are stored in non-volatile memory.

The effects of the transaction, thus, are never lost.

Some important points:

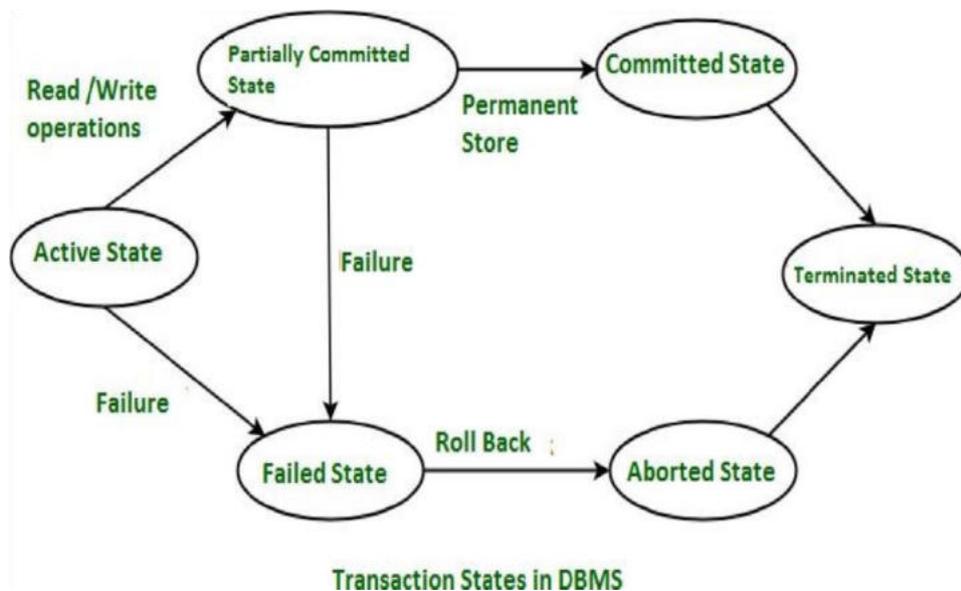
Property	Responsibilityformaintainingproperties
Atomicity	TransactionManager
Consistency	Applicationprogrammer
Isolation	ConcurrencyControlManager
Durability	RecoveryManager

The **ACID** properties, in totality, provide a mechanism to ensure the correctness and consistency of a database in a way such that each transaction is a group of operations that acts as a single unit, produces consistent results, acts in isolation from other operations, and updates that it makes are durably stored.

Topic III: Transactionstates:

States through which a transaction goes during its lifetime. These are the states which tell about the current state of the Transaction and also

tell how we will further do the processing in the transactions. These states govern the rules which decide the fate of the transaction whether it will commit or abort. They also use **Transaction log**. Transaction log is a file maintained by recovery management component to record all the activities of the transaction. After commit is done transaction log file is removed.



These are different types of Transaction States:

1. Active State –

When the instructions of the transaction are running then the transaction is in active state. If all the ‘read and write’ operations are performed without any error then it goes to the “partially committed state”; if any instruction fails, it goes to the “failed state”.

2. Partially Committed –

After completion of all the read and write operation the changes are made in main memory or local buffer. If the changes are made permanent on the DataBase then the state will change to “committed state” and in case of failure it will go to the “failed state”.

3. Failed State –

When any instruction of the transaction fails, it goes to the “failed state” or if failure occurs in making a permanent change of data on Data Base.

4. Aborted State –

After having any type of failure the transaction goes from “failed state” to “aborted state” and since in previous states, the changes are only made to local buffer or main memory and hence these changes are deleted or rolled-back.

5. Committed State –

It is the state when the changes are made permanent on the DataBase and the transaction is complete and therefore terminated in the “terminated state”.

6. Terminated State –

If there is no any roll-back or the transaction comes from the

“committed state”, then the system is consistent and ready for new transactionandtheoldtransactionisterminated.

Topic4:ConcurrentExecution:

ConcurrencyControlisthe managementprocedurethatisrequiredfor controlling concurrent execution of the operations that take place on a database.

Butbeforeknowingaboutconcurrencycontrol,weshouldknowabout concurrent execution.

ConcurrentExecutioninDBMS

Inamulti-usersystem,multipleuserscanaccessandusethe same databaseatonetime,whichisknownastheconcurrentexecution of the database. It means that the same database is executed simultaneously on a multi-user system by different users.

- While working on the database transactions, there occurs the requirement of using the database by multiple users for performing different operations, and in that case, concurrent execution of the database is performed.
- The thing is that the simultaneous execution that is performed shouldbedoneinaninterleavedmanner,andnooperationshould affect the other executing operations, thus maintaining the consistencyofthedatabase.Thus,onmakingtheconcurrent executionofthetransactionoperations,thereoccurseveral challenging problems that need to be solved.

ProblemswithConcurrentExecution:

are**READ**and**WRITE**operations.So,thereisaneedtomanagethese two operations in the concurrent execution of the transactions as if these operations are not performed in an interleaved manner, and the data may become inconsistent. So, the following problems occur with the Concurrent Execution of the operations:

Problem1:LostUpdateProblems(W-WConflict)

Theproblemoccurswhentwodifferentdatabasetransactionsperform theread/writeoperationsonthesamedatabaseitemsinaninterleaved manner (i.e., concurrent execution) that makes the values of the items incorrect hence making the database inconsistent.

Forexample:

ConsiderthebelowdiagramwheretwotransactionsTXandTY,are performed on the same account A where the balance of account A is \$300.

Time	T _x	T _y
t ₁	READ (A)	—
t ₂	A = A - 50	—
t ₃	—	READ (A)
t ₄	—	A = A + 100
t ₅	—	—
t ₆	WRITE (A)	—
t ₇	—	WRITE (A)

LOST UPDATE PROBLEM

- Attimet1,transactionT_xreadsthevalueofaccountA,i.e.,\$300 (only read).
- Attimet2,transactionT_xdeducts\$50fromaccountAthatbecomes \$250 (only deducted and not updated/write).

DepartmentOfAIDS

PreparedBy:MNANDINI DBMS

- Alternately, at time t_3 , transaction T_Y reads the value of account A that will be \$300 only because T_X didn't update the value yet.
- At time t_4 , transaction T_Y adds \$100 to account A that becomes \$400 (only added but not updated/write).
- At time t_6 , transaction T_X writes the value of account A that will be updated as \$250 only, as T_Y didn't update the value yet.
- Similarly, at time t_7 , transaction T_Y writes the values of account A, so it will write as done at time t_4 that will be \$400. It means the value written by T_X is lost, i.e., \$250 is lost.

Hence data becomes incorrect, and database set to inconsistent.

Problem 2: Dirty Read Problems (W-R Conflict)

The dirty read problem occurs *when one transaction updates an item of the database, and somehow the transaction fails, and before the data gets rollback, the updated database item is accessed by another transaction. There comes the Read-Write Conflict between both transactions.*

For example:

Consider two transactions T_X and T_Y in the below diagram performing read/write operations on account A where the available balance in account A is \$300:

Time	T _x	T _y
t ₁	READ (A)	—
t ₂	A = A + 50	—
t ₃	WRITE (A)	—
t ₄	—	READ (A)
t ₅	SERVER DOWN ROLLBACK	—

DIRTY READ PROBLEM

- At time t₁, transaction T_x reads the value of account A, i.e., \$300.
- At time t₂, transaction T_x adds \$50 to account A that becomes \$350.
- At time t₃, transaction T_x writes the updated value in account A, i.e., \$350. ○ Then at time t₄, transaction T_y reads account A that will be read as \$350.
- Then at time t₅, transaction T_x rolls back due to server problem, and the value changes back to \$300 (as initially).
- But the value for account A remains \$350 for transaction T_y as committed, which is the dirty read and therefore known as the Dirty Read Problem.

Problem3:UnrepeatableReadProblem(W-RConflict)

Also known as Inconsistent Retrievals Problem that occurs when in a transaction, two different values are read for the same database item.

Forexample:

Consider two transactions, TX and TY, performing the read/write operations on account A, having an available balance = \$300. The diagram is shown below:

Time	T _x	T _y
t ₁	READ (A)	—
t ₂	—	READ (A)
t ₃	—	A = A + 100
t ₄	—	WRITE (A)
t ₅	READ (A)	—

UNREPEATABLE READ PROBLEM

- At time t₁, transaction T_x reads the value from account A, i.e., \$300.
- At time t₂, transaction T_y reads the value from account A, i.e., \$300. ○
- At time t₃, transaction T_y updates the value of account A by adding \$100 to the available balance, and then it becomes \$400.
- At time t₄, transaction T_y writes the updated value, i.e., \$400.
- After that, at time t₅, transaction T_x reads the available value of account A, and that will be read as \$400.
- It means that within the same transaction T_x, it reads two different values of account A, i.e., \$300 initially, and after updation made by transaction T_y, it reads \$400. It is an unrepeatable read and is therefore known as the Unrepeatable read problem.

Chapter2:RecoverySystem Topic1:

Storage Structure:

A database system provides an ultimate view of the stored data. However, data in the form of bits, bytes get stored in different storage devices.

In this section, we will take an overview of various types of storage devices that are used for accessing and storing data.

StorageStructure

We have already described the storage system. In brief, the storage structure can be divided into two categories –

- **Volatile storage** – As the name suggests, a volatile storage cannot survive system crashes. Volatile storage devices are placed very close to the CPU; normally they are embedded onto the chipset itself. For example, main memory and cache memory are examples of volatile storage. They are fast but can store only a small amount of information.
- **Non-volatile storage** – These memories are made to survive system crashes. They are huge in data storage capacity, but slower in accessibility. Examples may include hard-disks, magnetic tapes, flash memory, and nonvolatile (battery backed up) RAM.

Topic2: Recovery And Atomicity Recovery and Atomicity

When a system crashes, it may have several transactions being executed and various files opened for them to modify the data items. Transactions are made of various operations, which are atomic in nature. But according to ACID properties of DBMS, atomicity of transactions as a whole must be maintained, that is, either all the operations are executed or none.

When a DBMS recovers from a crash, it should maintain the following

–

- It should check the states of all the transactions, which were being executed.

- A transaction may be in the middle of some operation; the DBMS must ensure the atomicity of the transaction in this case.
- It should check whether the transaction can be completed now or it needs to be rolled back.
- No transactions would be allowed to leave the DBMS in an inconsistent state.

There are two types of techniques, which can help a DBMS in recovering as well as

maintaining the atomicity of a transaction –

- Maintaining the logs of each transaction, and writing them onto some stable storage before actually modifying the database.
- Maintaining shadow paging, where the changes are done on a volatile memory, and later, the actual database is updated.

Topic3:LogBasedRecovery

The log is a sequence of records. Log of each transaction is maintained in some stable storage so that if any failure occurs, then it can be recovered from there. If any operation is performed on the database, then it will be recorded in the log. But the process of storing the logs should be done before the actual transaction is applied in the database.

Let's assume there is a transaction to modify the City of a student. The following logs are written for this transaction.

When the transaction is initiated, then it writes 'start' log.

<Tn,Start>

When the transaction modifies the City from 'Noida' to 'Bangalore', then another log is written to the file.

<Tn, City, 'Noida', 'Bangalore'>

When the transaction is finished, then it writes another log to indicate the end of the transaction.

There are two approaches to modify the database:

1. Deferred database modification:

- The deferred modification technique occurs if the transaction does not modify the database until it has committed.
- In this method, all the logs are created and stored in the stable storage, and the database is updated when a transaction commits.

2. Immediate database modification:

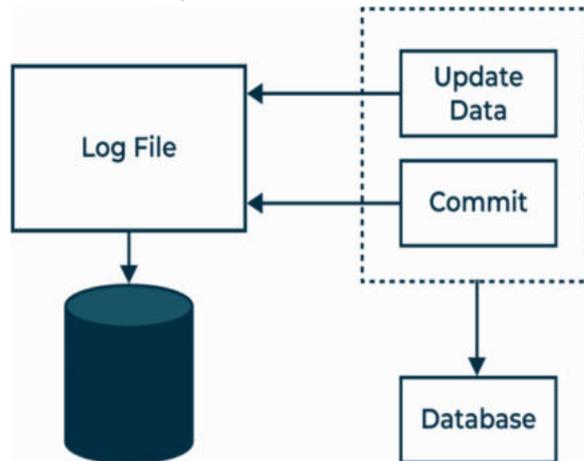
- The Immediate modification technique occurs if database modification occurs while the transaction is still active.

In this technique, the database is modified immediately after every operation. It follows an actual database modification.

Topic 4:

Log-Based Recovery in DBMS

Log-based recovery is a technique used in database management systems (DBMS) to ensure database consistency and durability in case of system failures. It relies on maintaining a **transaction log**, which records all changes made to the database, and uses this log to recover the database to a consistent state.



Phases of Log-Based Recovery

1. Log Analysis:

- Determines which transactions were committed and which were active at the time of the failure.
- Mark transactions as:
 - **Committed:** Changes will be done.
 - **Active/Uncommitted:** Changes will be undone.

2. Redo Phase:

- Reapplies changes of committed transactions to ensure durability.
- Scans the log forward from the last checkpoint to apply the **after image**.

3. Undo Phase:

- Rolls back the changes of uncommitted transactions to maintain atomicity.
- Scans the log backward from the end and applies the **before image**.

ExampleofLog-BasedRecovery Scenario:

- A database contains a data item A with an initial value of 50.
- Two transactions, **T1** and **T2**, update the value of A:
 - **T1** changes A to 60.
 - **T2** changes A to 70.

□

A system crash occurs
 after **T1** commits but before **T2** commits. **Log**

Entries:

Log Sequence Number (LSN)	Transaction ID	Operation	Data Item	Before Image	After Image	Status
1	T1	Update	A	50	60	Committed
2	T2	Update	A	60	70	Uncommitted

Recovery Process:

1. Analysis Phase:

- The log is analyzed to determine that:

- **T1**iscommitted.
- **T2**isactive/uncommittedatthetimeofthecrash.

2. RedoPhase:

- **T1's**changes(A=60)arereappliedbecause**T1**was committed.
- **T2's** changes (A = 70) are not redone because **T2** was not committed.

3. UndoPhase:

- Rollsbak**T2's**changesusingthe**beforeimage**inthelog, restoring A to 60.

FinalState:

- Afterrecovery,A=60,reflectingthecommittedchangesof**T1** andignoringtheuncommittedchangesof**T2**.

Topic4:

RIESRecoveryTechniqueinDBMS

ARIES (Algorithm for Recovery and Isolation Exploiting Semantics) is a widely used recovery technique in database management systems (DBMS) for ensuring database consistency and fault tolerance. It efficiently handles failures by supporting atomicity anddurability,twokeypropertiesofdatabasetransactionsintheACID model.

StepsofARIESRecoveryTechnique ARIESperformsrecoveryinthreephases:

1. AnalysisPhase:

- Scans thelog to identifytransactions thatwere activeat thetime of failure.
- Determinesthepointwhererecoveryshouldbegin.
- Reconstructstransactionanddirtypageinformation.

2. RedoPhase:

- Redoesallchangesmadebycommittedtransactionssincethelast checkpoint.

DepartmentOfAIDS

PreparedBy:MNANDINI DBMS

- Ensuresallcommittedchangesareappliedtothedatabase.

3. UndoPhase:

- Revertstheeffectsofuncommittedtransactionstomaintain database consistency.
- Usesthelogtorollbackuncommittedchangesinreverseorder.

KeyFeaturesofAriesRecoveryAlgorithminDBMS

1. **Write-Ahead Logging (WAL):** This make sure that all changesare logged before they are applied to the database.
2. **Checkpointing:**Tocreateastablepointinthedatabasefromwhich recovery can start.
3. **ThreePhasesofRecovery:**ToensuredatabaserecoveryAnalysis, Redo, and Undo phases are crucial.