



# **Dnyansagar Arts and Commerce College Balewadi Pune-45**

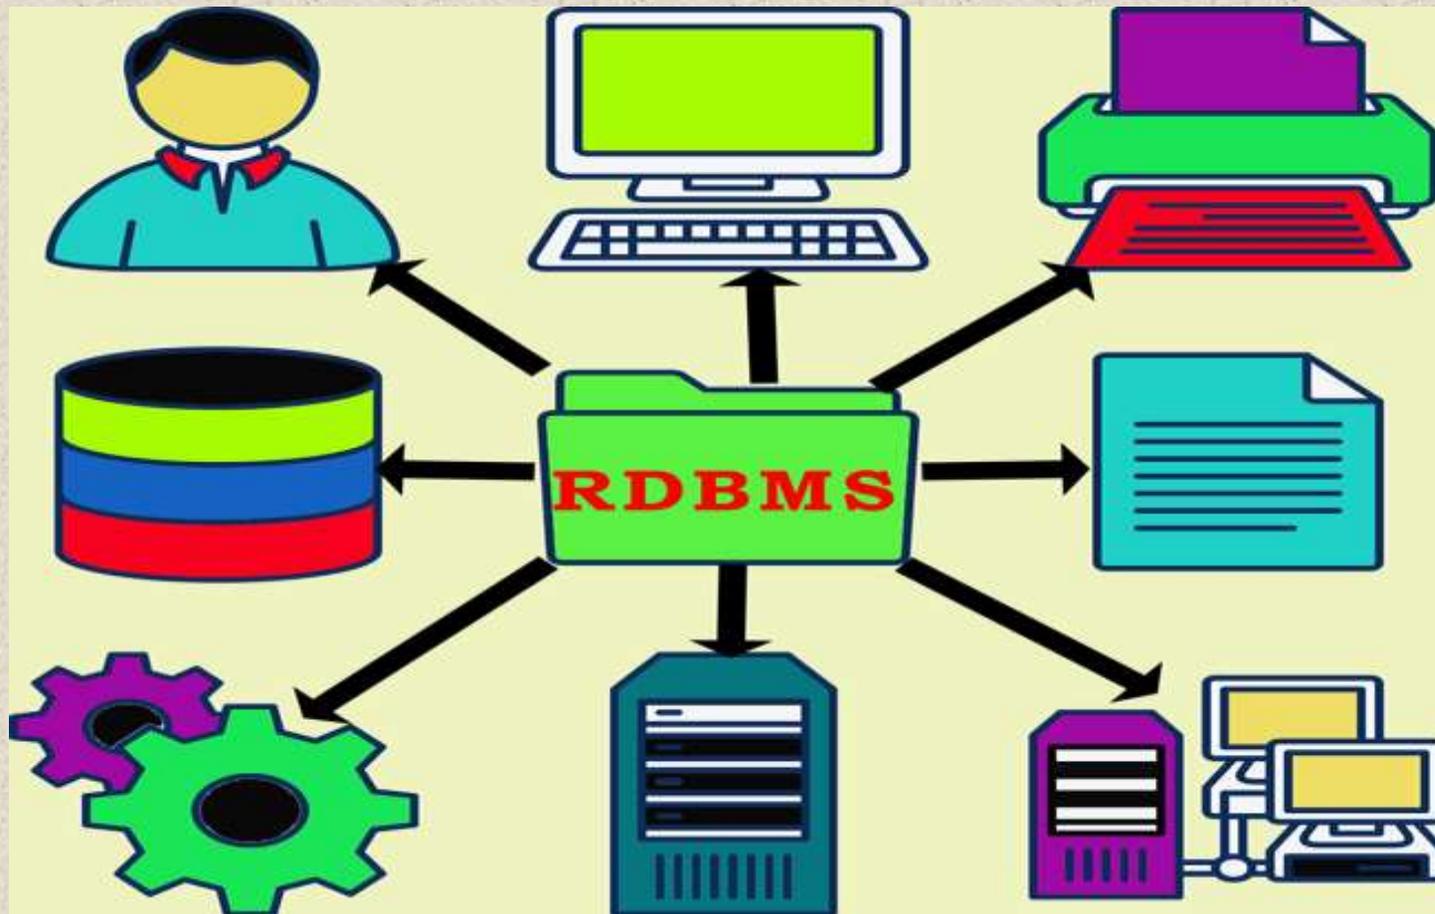
**Subject: (Relational Database Management System)**

**SYBBA(CA) – IV Sem (2019 Pattern)**

**Presented By: Prof.Gayatri A Amate**



## Unit 1: Introduction To RDBMS





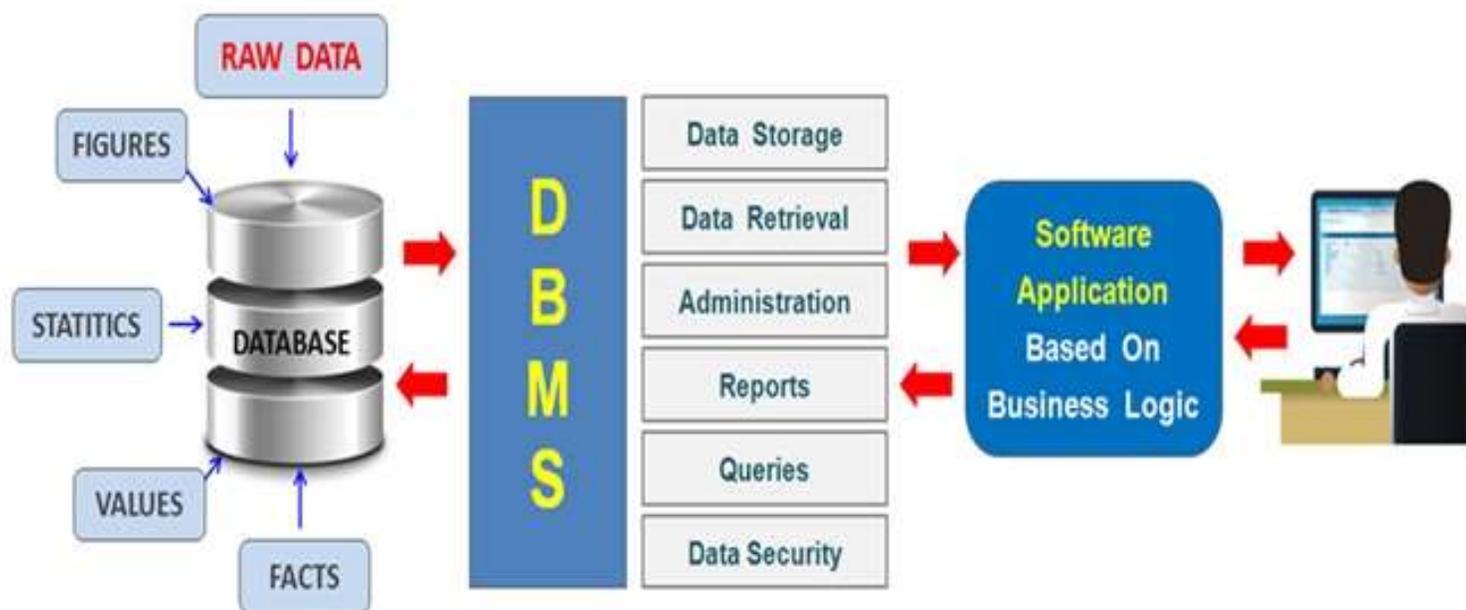
A Relational database management system (RDBMS) is a collection of programs and capabilities that enable IT teams and others to create, update, administer and otherwise interact with a relational database.

RDBMSes store data in the form of tables, with most commercial relational database management systems using [Structured Query Language \(SQL\)](#) to access the database. However, since SQL was invented after the initial development of the relational model, it is not necessary for RDBMS use.

The RDBMS is the most popular database system among organizations across the world. It provides a dependable method of storing and retrieving large amounts of data while offering a combination of system performance and ease of implementation.



## RDBMS - What Is RDBMS ?





## RDBMS vs. DBMS

In general, databases store sets of data that can be queried for use in other applications. A database management system supports the development, administration and use of database platforms.

An RDBMS is a type of [database management system](#) (DBMS) that stores data in a row-based table structure which connects related data elements. An RDBMS includes functions that maintain the security, accuracy, integrity and consistency of the data. This is different than the file storage used in a DBMS.

Other differences between database management systems and relational database management systems include:

- Number of allowed users. While a DBMS can only accept one user at a time, an RDBMS can operate with multiple users.
- Hardware and software requirements. A DBMS needs less software and hardware than an RDBMS.
- Amount of data. RDBMSes can handle any amount of data, from small to large, while a DBMS can only manage small amounts.
- Database structure. In a DBMS, data is kept in a hierarchical form, whereas an RDBMS utilizes a table where the headers are used as column names and the rows contain the corresponding values.
- ACID implementation. DBMSes do not use the atomicity, consistency, isolation and durability ([ACID](#)) model for storing data. On the other hand, RDBMSes base the structure of their data on the ACID model to ensure consistency.



- Distributed databases. While an RDBMS offers complete support for [distributed databases](#), a DBMS will not provide support.
- Types of programs managed. While an RDBMS helps manage the relationships between its incorporated tables of data, a DBMS focuses on maintaining databases that are present within the computer network and system [hard disks](#).
- Support of database normalization. An RDBMS can be [normalized](#), but a DBMS cannot.

What is a field?

Every table is broken up into smaller entities called fields. The fields in the CUSTOMERS table consist of ID, NAME, AGE, ADDRESS and SALARY.

A field is a column in a table that is designed to maintain specific information about every record in the table.

What is a Record or a Row?

A record is also called as a row of data is each individual entry that exists in a table. For example, there are 7 records in the above CUSTOMERS table. Following is a single row of data or record in the CUSTOMERS table –

```
+----+-----+----+-----+-----+
| 1 | Ramesh | 32 | Ahmedabad | 2000.00 |
```



+-----+-----+-----+-----+-----+

A record is a horizontal entity in a table.

What is a column?

A column is a vertical entity in a table that contains all information associated with a specific field in a table.

For example, a column in the CUSTOMERS table is ADDRESS, which represents location description and would be as shown below –

+-----+

| ADDRESS |

+-----+

| Ahmedabad |

| Delhi |

| Kota |

| Mumbai |

| Bhopal |

| MP |



| Indore |

+-----+

What is a NULL value?

A NULL value in a table is a value in a field that appears to be blank, which means a field with a NULL value is a field with no value.

It is very important to understand that a NULL value is different than a zero value or a field that contains spaces. A field with a NULL value is the one that has been left blank during a record creation.

SQL Constraints



Constraints are the rules enforced on data columns on a table. These are used to limit the type of data that can go into a table. This ensures the accuracy and reliability of the data in the database.

Constraints can either be column level or table level. Column level constraints are applied only to one column whereas, table level constraints are applied to the entire table.

Following are some of the most commonly used constraints available in SQL –

- [NOT NULL Constraint](#) – Ensures that a column cannot have a NULL value.
- [DEFAULT Constraint](#) – Provides a default value for a column when none is specified.
- [UNIQUE Constraint](#) – Ensures that all the values in a column are different.



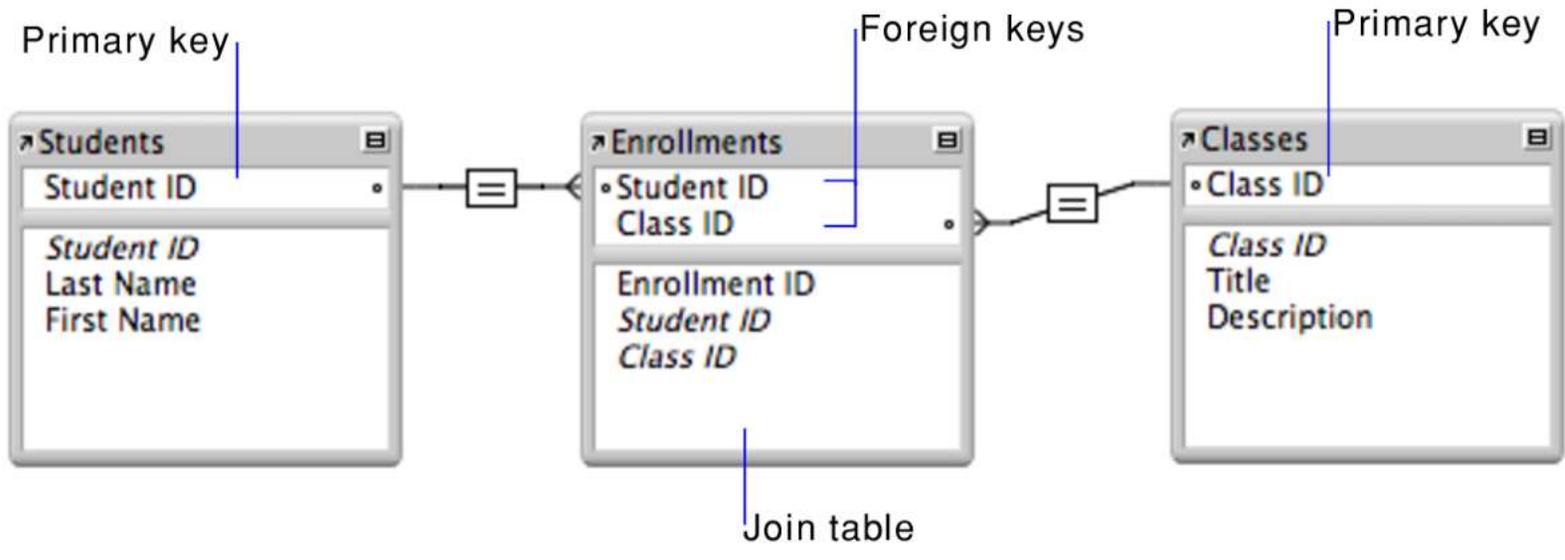
- PRIMARY Key – Uniquely identifies each row/record in a database table.
- FOREIGN Key – Uniquely identifies a row/record in any another database table.
- CHECK Constraint – The CHECK constraint ensures that all values in a column satisfy certain conditions.
- INDEX – Used to create and retrieve data from the database very quickly.

Primary Key



Employee Id must be Unique

Employee Id	Name	Salary
101	Hitesh	250000
102	Gaurav	40000
103	Anjali	25000
104	Gaurav	60000
105	Poonam	25000





*Composite Primary Key*      *Unique*      *Not Null*      *Check*

<u>CUST_ID</u>	<u>PAN</u>	<u>ACCT_NO</u>	<u>LOC</u>	<u>GENDER</u>	<u>ADDRESS</u>
101	ABCDE1234F	9234567892	BANGALORE	M	BELLANDUR
102	ABCDE2234F	8234567892	CHENNAI	M	BELLANDUR
103	ABCDE3234F		BANGALORE	F	MARATHALLI
104	ABCDE4234F	6234567892	HYDERABAD	F	ADARSH NAGAR
105	ABCDE5234F	5234567892	BANGALORE	M	MARATHALLI
106	ABCDE6234F		CHENNAI	M	BRIGADE ROAD
107	ABCDE7234F	3234567892	HYDERABAD	M	M.G ROAD
108	ABCDE8234F	2234567892	BANGALORE	F	BTM
109	ABCDE9234F	1234567892	BANGALORE	F	HSR LAYOUT
110	ABCDE0234F	3434567892	BANGALORE	M	FRAZER ROAD

**Sample Database for Constraints: Customer Table**

## Data Integrity

The following categories of data integrity exist with each RDBMS –

- Entity Integrity – There are no duplicate rows in a table.
- Domain Integrity – Enforces valid entries for a given column by restricting the type, the format, or the range of values.
- Referential integrity – Rows cannot be deleted, which are used by other records.



- User-Defined Integrity – Enforces some specific business rules that do not fall into entity, domain or referential integrity.

## Database Normalization

Database normalization is the process of efficiently organizing data in a database. There are two reasons of this normalization process –

- Eliminating redundant data, for example, storing the same data in more than one table.
- Ensuring data dependencies make sense.

Both these reasons are worthy goals as they reduce the amount of space a database consumes and ensures that data is logically stored. Normalization consists of a series of guidelines that help guide you in creating a good database structure.

Normalization guidelines are divided into normal forms; think of a form as the format or the way a database structure is laid out. The aim of normal forms is to organize the database structure, so that it complies with the rules of first normal form, then second normal form and finally the third normal form.

It is your choice to take it further and go to the fourth normal form, fifth normal form and so on, but in general, the third normal form is more than enough.

- [First Normal Form \(1NF\)](#)
- [Second Normal Form \(2NF\)](#)
- [Third Normal Form \(3NF\)](#)



## **Unit 2 : PL/SQL**

The PL/SQL programming language was developed by Oracle Corporation in the late 1980s as procedural extension language for SQL and the Oracle relational database. Following are certain notable facts about PL/SQL

- PL/SQL is a completely portable, high-performance transaction-processing language.
- PL/SQL provides a built-in, interpreted and OS independent programming environment.
- PL/SQL can also directly be called from the command-line SQL\*Plus interface.
- Direct call can also be made from external programming language calls to database.



- PL/SQL's general syntax is based on that of ADA and Pascal programming language.
- Apart from Oracle, PL/SQL is available in TimesTen in-memory database and IBM DB2.

### Features of PL/SQL

PL/SQL has the following features –

- PL/SQL is tightly integrated with SQL.
- It offers extensive error checking.
- It offers numerous data types.
- It offers a variety of programming structures.
- It supports structured programming through functions and procedures.
- It supports object-oriented programming.
- It supports the development of web applications and server pages.

### Advantages of PL/SQL

PL/SQL has the following advantages –

- SQL is the standard database language and PL/SQL is strongly integrated with SQL. PL/SQL supports both static and dynamic SQL. Static SQL supports DML operations and transaction control from PL/SQL block. In Dynamic SQL, SQL allows embedding DDL statements in PL/SQL blocks.



- PL/SQL allows sending an entire block of statements to the database at one time. This reduces network traffic and provides high performance for the applications.
- PL/SQL gives high productivity to programmers as it can query, transform, and update data in a database.
- PL/SQL saves time on design and debugging by strong features, such as exception handling, encapsulation, data hiding, and object-oriented data types.
- Applications written in PL/SQL are fully portable.
- PL/SQL provides high security level.
- PL/SQL provides access to predefined SQL packages.
- PL/SQL provides support for Object-Oriented Programming.
- PL/SQL provides support for developing Web Applications and Server Pages.

In this chapter, we will discuss the Basic Syntax of PL/SQL which is a block-structured language; this means that the PL/SQL programs are divided and written in logical blocks of code. Each block consists of three sub-parts –

S.No	Sections & Description
------	------------------------



1	<p><b>Declarations</b></p> <p>This section starts with the keyword <b>DECLARE</b>. It is an optional section and defines all variables, cursors, subprograms, and other elements to be used in the program.</p>
2	<p><b>Executable Commands</b></p> <p>This section is enclosed between the keywords <b>BEGIN</b> and <b>END</b> and it is a mandatory section. It consists of the executable PL/SQL statements of the program. It should have at least one executable line of code, which may be just a <b>NULL</b> command to indicate that nothing should be executed.</p>
3	<p><b>Exception Handling</b></p> <p>This section starts with the keyword <b>EXCEPTION</b>. This optional section contains exception(s) that handle errors in the program.</p>

Every PL/SQL statement ends with a semicolon (;). PL/SQL blocks can be nested within other PL/SQL blocks using **BEGIN** and **END**. Following is the basic structure of a PL/SQL block –

**DECLARE**

<declarations section>



**BEGIN**

<executable command(s)>

**EXCEPTION**

<exception handling>

**END;**

The 'Hello World' Example

**DECLARE**

message varchar2(20):= 'Hello, World!';

**BEGIN**

dbms\_output.put\_line(message);

**END;**

/

The end; line signals the end of the PL/SQL block. To run the code from the SQL command line, you may need to type / at the beginning of the first blank line after the last line of the code. When the above code is executed at the SQL prompt, it produces the following result –

Hello World



PL/SQL procedure successfully completed.

### The PL/SQL Identifiers

PL/SQL identifiers are constants, variables, exceptions, procedures, cursors, and reserved words. The identifiers consist of a letter optionally followed by more letters, numerals, dollar signs, underscores, and number signs and should not exceed 30 characters.

By default, identifiers are not case-sensitive. So you can use integer or INTEGER to represent a numeric value. You cannot use a reserved keyword as an identifier.

### The PL/SQL Delimiters

A delimiter is a symbol with a special meaning. Following is the list of delimiters in PL/SQL –

Delimiter	Description
+, -, *, /	Addition, subtraction/negation, multiplication, division
%	Attribute indicator
'	Character string delimiter



.	Component selector
(,)	Expression or list delimiter
:	Host variable indicator
,	Item separator
"	Quoted identifier delimiter
=	Relational operator
@	Remote access indicator
;	Statement terminator
:=	Assignment operator
=>	Association operator



	Concatenation operator
**	Exponentiation operator
<<, >>	Label delimiter (begin and end)
/*, */	Multi-line comment delimiter (begin and end)
--	Single-line comment indicator
..	Range operator
<, >, <=, >=	Relational operators
<>, !=, ~=, ^=	Different versions of NOT EQUAL

### The PL/SQL Comments

Program comments are explanatory statements that can be included in the PL/SQL code that you write and helps anyone reading its source code. All programming languages allow some form of comments.



The PL/SQL supports single-line and multi-line comments. All characters available inside any comment are ignored by the PL/SQL compiler. The PL/SQL single-line comments start with the delimiter -- (double hyphen) and multi-line comments are enclosed by /\* and \*/.

DECLARE

```
-- variable declaration
```

```
message varchar2(20):= 'Hello, World!';
```

BEGIN

```
/*
```

```
* PL/SQL executable statement(s)
```

```
*/
```

```
dbms_output.put_line(message);
```

END;

```
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Hello World



PL/SQL procedure successfully completed.

## PL/SQL Program Units

A PL/SQL unit is any one of the following –

- PL/SQL block
- Function
- Package
- Package body
- Procedure
- Trigger
- Type
- Type body

Each of these units will be discussed in the following chapters.

## Data Types

he PL/SQL variables, constants and parameters must have a valid data type, which specifies a storage format, constraints, and a valid range of values. We will focus on the SCALAR and the LOB data types in this chapter. The other two data types will be covered in other chapters.



S.No	Category & Description
1	<p>Scalar</p> <p>Single values with no internal components, such as a NUMBER, DATE, or BOOLEAN.</p>
2	<p>Large Object (LOB)</p> <p>Pointers to large objects that are stored separately from other data items, such as text, graphic images, video clips, and sound waveforms.</p>
3	<p>Composite</p> <p>Data items that have internal components that can be accessed individually. For example, collections and records.</p>
4	<p>Reference</p> <p>Pointers to other data items.</p>

PL/SQL Scalar Data Types and Subtypes



PL/SQL Scalar Data Types and Subtypes come under the following categories –

S.No	Date Type & Description
1	<p><b>Numeric</b></p> <p>Numeric values on which arithmetic operations are performed.</p>
2	<p><b>Character</b></p> <p>Alphanumeric values that represent single characters or strings of characters.</p>
3	<p><b>Boolean</b></p> <p>Logical values on which logical operations are performed.</p>
4	<p><b>Datetime</b></p> <p>Dates and times.</p>

PL/SQL provides subtypes of data types. For example, the data type NUMBER has a subtype called INTEGER. You can use the subtypes in your PL/SQL program to make the data types compatible with data types in other programs while embedding the PL/SQL code in another program, such as a Java program.



## PL/SQL Numeric Data Types and Subtypes

Following table lists out the PL/SQL pre-defined numeric data types and their sub-types –

S.No	Data Type & Description
1	<b>PLS_INTEGER</b> Signed integer in range -2,147,483,648 through 2,147,483,647, represented in 32 bits
2	<b>BINARY_INTEGER</b> Signed integer in range -2,147,483,648 through 2,147,483,647, represented in 32 bits
3	<b>BINARY_FLOAT</b> Single-precision IEEE 754-format floating-point number
4	<b>BINARY_DOUBLE</b> Double-precision IEEE 754-format floating-point number



5	<p>NUMBER(prec, scale)</p> <p>Fixed-point or floating-point number with absolute value in range 1E-130 to (but not including) 1.0E126. A NUMBER variable can also represent 0</p>
6	<p>DEC(prec, scale)</p> <p>ANSI specific fixed-point type with maximum precision of 38 decimal digits</p>
7	<p>DECIMAL(prec, scale)</p> <p>IBM specific fixed-point type with maximum precision of 38 decimal digits</p>
8	<p>NUMERIC(pre, scale)</p> <p>Floating type with maximum precision of 38 decimal digits</p>
9	<p>DOUBLE PRECISION</p> <p>ANSI specific floating-point type with maximum precision of 126 binary digits (approximately 38 decimal digits)</p>



10	<b>FLOAT</b> ANSI and IBM specific floating-point type with maximum precision of 126 binary digits (approximately 38 decimal digits)
11	<b>INT</b> ANSI specific integer type with maximum precision of 38 decimal digits
12	<b>INTEGER</b> ANSI and IBM specific integer type with maximum precision of 38 decimal digits
13	<b>SMALLINT</b> ANSI and IBM specific integer type with maximum precision of 38 decimal digits
14	<b>REAL</b> Floating-point type with maximum precision of 63 binary digits (approximately 18 decimal digits)



Following is a valid declaration –

DECLARE

num1 INTEGER;

num2 REAL;

num3 DOUBLE PRECISION;

BEGIN

null;

END;

/

When the above code is compiled and executed, it produces the following result –

PL/SQL procedure successfully completed

PL/SQL Character Data Types and Subtypes

Following is the detail of PL/SQL pre-defined character data types and their sub-types –

S.No	Data Type & Description
------	-------------------------



1	<b>CHAR</b> Fixed-length character string with maximum size of 32,767 bytes
2	<b>VARCHAR2</b> Variable-length character string with maximum size of 32,767 bytes
3	<b>RAW</b> Variable-length binary or byte string with maximum size of 32,767 bytes, not interpreted by PL/SQL
4	<b>NCHAR</b> Fixed-length national character string with maximum size of 32,767 bytes
5	<b>NVARCHAR2</b> Variable-length national character string with maximum size of 32,767 bytes



6	<b>LONG</b> Variable-length character string with maximum size of 32,760 bytes
7	<b>LONG RAW</b> Variable-length binary or byte string with maximum size of 32,760 bytes, not interpreted by PL/SQL
8	<b>ROWID</b> Physical row identifier, the address of a row in an ordinary table
9	<b>UROWID</b> Universal row identifier (physical, logical, or foreign row identifier)

#### PL/SQL Boolean Data Types

The **BOOLEAN** data type stores logical values that are used in logical operations. The logical values are the Boolean values **TRUE** and **FALSE** and the value **NULL**.

However, SQL has no data type equivalent to **BOOLEAN**. Therefore, Boolean values cannot be used in –

- SQL statements



- Built-in SQL functions (such as TO\_CHAR)
- PL/SQL functions invoked from SQL statements

### PL/SQL Datetime and Interval Types

The DATE datatype is used to store fixed-length datetimes, which include the time of day in seconds since midnight. Valid dates range from January 1, 4712 BC to December 31, 9999 AD.

The default date format is set by the Oracle initialization parameter NLS\_DATE\_FORMAT. For example, the default might be 'DD-MON-YY', which includes a two-digit number for the day of the month, an abbreviation of the month name, and the last two digits of the year. For example, 01-OCT-12.

Each DATE includes the century, year, month, day, hour, minute, and second. The following table shows the valid values for each field –

Field Name	Valid Datetime Values	Valid Interval Values
YEAR	-4712 to 9999 (excluding year 0)	Any nonzero integer
MONTH	01 to 12	0 to 11



DAY	01 to 31 (limited by the values of MONTH and YEAR, according to the rules of the calendar for the locale)	Any nonzero integer
HOUR	00 to 23	0 to 23
MINUTE	00 to 59	0 to 59
SECOND	00 to 59.9(n), where 9(n) is the precision of time fractional seconds	0 to 59.9(n), where 9(n) is the precision of interval fractional seconds
TIMEZONE_HOUR	-12 to 14 (range accommodates daylight savings time changes)	Not applicable
TIMEZONE_MINUTE	00 to 59	Not applicable



TIMEZONE_REGION	Found in the dynamic performance view V\$TIMEZONE_NAMES	Not applicable
TIMEZONE_ABBR	Found in the dynamic performance view V\$TIMEZONE_NAMES	Not applicable

### PL/SQL Large Object (LOB) Data Types

Large Object (LOB) data types refer to large data items such as text, graphic images, video clips, and sound waveforms. LOB data types allow efficient, random, piecewise access to this data. Following are the predefined PL/SQL LOB data types –

Data Type	Description	Size
BFILE	Used to store large binary objects in operating system files outside the database.	System-dependent. Cannot exceed 4 gigabytes (GB).
BLOB	Used to store large binary objects in the database.	8 to 128 terabytes (TB)



CLOB	Used to store large blocks of character data in the database.	8 to 128 TB
NCLOB	Used to store large blocks of NCHAR data in the database.	8 to 128 TB

### PL/SQL User-Defined Subtypes

A subtype is a subset of another data type, which is called its base type. A subtype has the same valid operations as its base type, but only a subset of its valid values.

PL/SQL predefines several subtypes in package STANDARD. For example, PL/SQL predefines the subtypes CHARACTER and INTEGER as follows –

```
SUBTYPE CHARACTER IS CHAR;
```

```
SUBTYPE INTEGER IS NUMBER(38,0);
```

You can define and use your own subtypes. The following program illustrates defining and using a user-defined subtype –

```
DECLARE
```

```
    SUBTYPE name IS char(20);
```

```
    SUBTYPE message IS varchar2(100);
```



```
salutation name;  
greetings message;  
BEGIN  
salutation := 'Reader ';  
greetings := 'Welcome to the World of PL/SQL';  
dbms_output.put_line('Hello ' || salutation || greetings);  
END;  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

```
Hello Reader Welcome to the World of PL/SQL
```

PL/SQL procedure successfully completed.

### NULLs in PL/SQL

PL/SQL NULL values represent missing or unknown data and they are not an integer, a character, or any other specific data type. Note that NULL is not the same as an empty data string or the null character value '\0'. A null can be assigned but it cannot be equated with anything, including itself.



variable

A variable is nothing but a name given to a storage area that our programs can manipulate. Each variable in PL/SQL has a specific data type, which determines the size and the layout of the variable's memory; the range of values that can be stored within that memory and the set of operations that can be applied to the variable.

The name of a PL/SQL variable consists of a letter optionally followed by more letters, numerals, dollar signs, underscores, and number signs and should not exceed 30 characters. By default, variable names are not case-sensitive. You cannot use a reserved PL/SQL keyword as a variable name.

PL/SQL programming language allows to define various types of variables, such as date time data types, records, collections, etc. which we will cover in subsequent chapters. For this chapter, let us study only basic variable types.

### Variable Declaration in PL/SQL

PL/SQL variables must be declared in the declaration section or in a package as a global variable. When you declare a variable, PL/SQL allocates memory for the variable's value and the storage location is identified by the variable name.

The syntax for declaring a variable is –

```
variable_name [CONSTANT] datatype [NOT NULL] [:= | DEFAULT initial_value]
```

Where, *variable\_name* is a valid identifier in PL/SQL, *datatype* must be a valid PL/SQL data type or any user defined data type which we already have discussed in the last chapter. Some valid variable declarations along with their definition are shown below –



```
sales number(10, 2);
```

```
pi CONSTANT double precision := 3.1415;
```

```
name varchar2(25);
```

```
address varchar2(100);
```

When you provide a size, scale or precision limit with the data type, it is called a constrained declaration. Constrained declarations require less memory than unconstrained declarations. For example –

```
sales number(10, 2);
```

```
name varchar2(25);
```

```
address varchar2(100);
```

### Initializing Variables in PL/SQL

Whenever you declare a variable, PL/SQL assigns it a default value of NULL. If you want to initialize a variable with a value other than the NULL value, you can do so during the declaration, using either of the following –

- The DEFAULT keyword
- The assignment operator

For example –

```
counter binary_integer := 0;
```



```
greetings varchar2(20) DEFAULT 'Have a Good Day';
```

You can also specify that a variable should not have a NULL value using the NOT NULL constraint. If you use the NOT NULL constraint, you must explicitly assign an initial value for that variable.

It is a good programming practice to initialize variables properly otherwise, sometimes programs would produce unexpected results. Try the following example which makes use of various types of variables –

```
DECLARE
```

```
  a integer := 10;
```

```
  b integer := 20;
```

```
  c integer;
```

```
  f real;
```

```
BEGIN
```

```
  c := a + b;
```

```
  dbms_output.put_line('Value of c: ' || c);
```

```
  f := 70.0/3.0;
```

```
  dbms_output.put_line('Value of f: ' || f);
```

```
END;
```



/

When the above code is executed, it produces the following result –

Value of c: 30

Value of f: 23.333333333333333333

PL/SQL procedure successfully completed.

### Variable Scope in PL/SQL

PL/SQL allows the nesting of blocks, i.e., each program block may contain another inner block. If a variable is declared within an inner block, it is not accessible to the outer block. However, if a variable is declared and accessible to an outer block, it is also accessible to all nested inner blocks. There are two types of variable scope –

- Local variables – Variables declared in an inner block and not accessible to outer blocks.
- Global variables – Variables declared in the outermost block or a package.

Following example shows the usage of Local and Global variables in its simple form –

DECLARE

-- Global variables

num1 number := 95;



```
num2 number := 85;

BEGIN

dbms_output.put_line('Outer Variable num1: ' || num1);
dbms_output.put_line('Outer Variable num2: ' || num2);

DECLARE

-- Local variables

num1 number := 195;

num2 number := 185;

BEGIN

dbms_output.put_line('Inner Variable num1: ' || num1);
dbms_output.put_line('Inner Variable num2: ' || num2);

END;

END;

/
```

When the above code is executed, it produces the following result –



Outer Variable num1: 95

Outer Variable num2: 85

Inner Variable num1: 195

Inner Variable num2: 185

PL/SQL procedure successfully completed.

#### Assigning SQL Query Results to PL/SQL Variables

You can use the `SELECT INTO` statement of SQL to assign values to PL/SQL variables. For each item in the `SELECT` list, there must be a corresponding, type-compatible variable in the `INTO` list. The following example illustrates the concept. Let us create a table named `CUSTOMERS` –

```
CREATE TABLE CUSTOMERS(  
  ID INT NOT NULL,  
  NAME VARCHAR (20) NOT NULL,  
  AGE INT NOT NULL,  
  ADDRESS CHAR (25),
```



```
SALARY DECIMAL (18, 2),  
PRIMARY KEY (ID)  
);
```

Table Created

Let us now insert some values in the table –

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)  
VALUES (1, 'Ramesh', 32, 'Ahmedabad', 2000.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)  
VALUES (2, 'Khilan', 25, 'Delhi', 1500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)  
VALUES (3, 'kaushik', 23, 'Kota', 2000.00 );
```



```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (4, 'Chaitali', 25, 'Mumbai', 6500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (5, 'Hardik', 27, 'Bhopal', 8500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (6, 'Komal', 22, 'MP', 4500.00 );
```

The following program assigns values from the above table to PL/SQL variables using the SELECT INTO clause of SQL –

```
DECLARE
```

```
  c_id customers.id%type := 1;
```

```
  c_name customers.name%type;
```

```
  c_addr customers.address%type;
```

```
  c_sal customers.salary%type;
```



```
BEGIN
```

```
    SELECT name, address, salary INTO c_name, c_addr, c_sal
```

```
    FROM customers
```

```
    WHERE id = c_id;
```

```
    dbms_output.put_line
```

```
    ('Customer ' || c_name || ' from ' || c_addr || ' earns ' || c_sal);
```

```
END;
```

```
/
```

When the above code is executed, it produces the following result –

Customer Ramesh from Ahmedabad earns 2000

PL/SQL procedure completed successfully

constants and literals in PL/SQL

A constant holds a value that once declared, does not change in the program. A constant declaration specifies its name, data type, and value, and allocates storage for it. The declaration can also impose the NOT NULL constraint.



## Declaring a Constant

A constant is declared using the **CONSTANT** keyword. It requires an initial value and does not allow that value to be changed. For example –

```
PI CONSTANT NUMBER := 3.141592654;
```

```
DECLARE
```

```
-- constant declaration
```

```
pi constant number := 3.141592654;
```

```
-- other declarations
```

```
radius number(5,2);
```

```
dia number(5,2);
```

```
circumference number(7, 2);
```

```
area number (10, 2);
```

```
BEGIN
```

```
-- processing
```

```
radius := 9.5;
```



```
dia := radius * 2;
circumference := 2.0 * pi * radius;
area := pi * radius * radius;
-- output
dbms_output.put_line('Radius: ' || radius);
dbms_output.put_line('Diameter: ' || dia);
dbms_output.put_line('Circumference: ' || circumference);
dbms_output.put_line('Area: ' || area);
END;
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Radius: 9.5

Diameter: 19

Circumference: 59.69

Area: 283.53



PL/SQL procedure successfully completed.

### The PL/SQL Literals

A literal is an explicit numeric, character, string, or Boolean value not represented by an identifier. For example, TRUE, 786, NULL, 'tutorialspoint' are all literals of type Boolean, number, or string. PL/SQL, literals are case-sensitive. PL/SQL supports the following kinds of literals –

- Numeric Literals
- Character Literals
- String Literals
- BOOLEAN Literals
- Date and Time Literals

The following table provides examples from all these categories of literal values.

S.No	Literal Type & Example
1	Numeric Literals



050 78 -14 0 +32767

6.6667 0.0 -12.0 3.14159 +7800.00

6E5 1.0E-8 3.14159e0 -1E38 -9.5e-3

Character Literals

2

'A' '%' '9' ' ' 'z' '('

String Literals

3

'Hello, world!'

'Tutorials Point'

'19-NOV-12'

BOOLEAN Literals

4

TRUE, FALSE, and NULL.



## Date and Time Literals

```
5  DATE '1978-12-25';  
   TIMESTAMP '2012-10-29 12:01:01';
```

To embed single quotes within a string literal, place two single quotes next to each other as shown in the following program –

```
DECLARE
```

```
    message varchar2(30):= 'That"s tutorialspoint!';
```

```
BEGIN
```

```
    dbms_output.put_line(message);
```

```
END;
```

```
/
```

When the above code is executed at the SQL prompt, it produces the following result –

```
That's tutorialspoint.com!
```

```
PL/SQL procedure successfully completed.
```



## operators in PL/SQL

An operator is a symbol that tells the compiler to perform specific mathematical or logical manipulation. PL/SQL language is rich in built-in operators and provides the following types of operators –

- Arithmetic operators
- Relational operators
- Comparison operators
- Logical operators
- String operators

Here, we will understand the arithmetic, relational, comparison and logical operators one by one. The String operators will be discussed in a later chapter – PL/SQL - Strings.

### Arithmetic Operators

Following table shows all the arithmetic operators supported by PL/SQL. Let us assume variable A holds 10 and variable B holds 5, then –

[Show Examples](#)

Operator	Description	Example
----------	-------------	---------



+	Adds two operands	$A + B$ will give 15
-	Subtracts second operand from the first	$A - B$ will give 5
*	Multiplies both operands	$A * B$ will give 50
/	Divides numerator by de-numerator	$A / B$ will give 2
**	Exponentiation operator, raises one operand to the power of other	$A ** B$ will give 100000

### Relational Operators

Relational operators compare two expressions or values and return a Boolean result. Following table shows all the relational operators supported by PL/SQL. Let us assume variable A holds 10 and variable B holds 20, then –

[Show Examples](#)

Operator	Description	Example
=	Checks if the values of two operands are equal or not, if yes then condition becomes true.	$(A = B)$ is not true.



!=	Checks if the values of two operands are equal or not, if values are not equal then condition becomes true.	(A != B) is true.
>	Checks if the value of left operand is greater than the value of right operand, if yes then condition becomes true.	(A > B) is not true.
<	Checks if the value of left operand is less than the value of right operand, if yes then condition becomes true.	(A < B) is true.
>=	Checks if the value of left operand is greater than or equal to the value of right operand, if yes then condition becomes true.	(A >= B) is not true.
<=	Checks if the value of left operand is less than or equal to the value of right operand, if yes then condition becomes true.	(A <= B) is true.

### Comparison Operators



Comparison operators are used for comparing one expression to another. The result is always either TRUE, FALSE or NULL.

[Show Examples](#)

Operator	Description	Example
LIKE	The LIKE operator compares a character, string, or CLOB value to a pattern and returns TRUE if the value matches the pattern and FALSE if it does not.	If 'Zara Ali' like 'Z% A_i' returns a Boolean true, whereas, 'Nuha Ali' like 'Z% A_i' returns a Boolean false.
BETWEEN	The BETWEEN operator tests whether a value lies in a specified range. x BETWEEN a AND b means that $x \geq a$ and $x \leq b$ .	If $x = 10$ then, x between 5 and 20 returns true, x between 5 and 10 returns true, but x between 11 and 20 returns false.



IN	The IN operator tests set membership. x IN (set) means that x is equal to any member of set.	If x = 'm' then, x in ('a', 'b', 'c') returns Boolean false but x in ('m', 'n', 'o') returns Boolean true.
IS NULL	The IS NULL operator returns the BOOLEAN value TRUE if its operand is NULL or FALSE if it is not NULL. Comparisons involving NULL values always yield NULL.	If x = 'm', then 'x is null' returns Boolean false.

### Logical Operators

Following table shows the Logical operators supported by PL/SQL. All these operators work on Boolean operands and produce Boolean results. Let us assume variable A holds true and variable B holds false, then –

[Show Examples](#)

Operator	Description	Examples
----------	-------------	----------



and	Called the logical AND operator. If both the operands are true then condition becomes true.	(A and B) is false.
or	Called the logical OR Operator. If any of the two operands is true then condition becomes true.	(A or B) is true.
not	Called the logical NOT Operator. Used to reverse the logical state of its operand. If a condition is true then Logical NOT operator will make it false.	not (A and B) is true.

### PL/SQL Operator Precedence

Operator precedence determines the grouping of terms in an expression. This affects how an expression is evaluated. Certain operators have higher precedence than others; for example, the multiplication operator has higher precedence than the addition operator.

For example,  $x = 7 + 3 * 2$ ; here,  $x$  is assigned 13, not 20 because operator  $*$  has higher precedence than  $+$ , so it first gets multiplied with  $3*2$  and then adds into 7.

Here, operators with the highest precedence appear at the top of the table, those with the lowest appear at the bottom. Within an expression, higher precedence operators will be evaluated first.

The precedence of operators goes as follows:  $=$ ,  $<$ ,  $>$ ,  $<=$ ,  $>=$ ,  $<>$ ,  $!=$ ,  $\sim=$ ,  $\wedge=$ , IS NULL, LIKE, BETWEEN, IN.

[Show Examples](#)



Operator	Operation
**	exponentiation
+, -	identity, negation
*, /	multiplication, division
+, -,	addition, subtraction, concatenation
comparison	
NOT	logical negation
AND	conjunction
OR	inclusion



## Creating a Function

A standalone function is created using the CREATE FUNCTION statement. The simplified syntax for the CREATE OR REPLACE PROCEDURE statement is as follows –

```
CREATE [OR REPLACE] FUNCTION function_name
```

```
[(parameter_name [IN | OUT | IN OUT] type [, ...])]
```

```
RETURN return_datatype
```

```
{IS | AS}
```

```
BEGIN
```

```
< function_body >
```

```
END [function_name];
```

Where,

- *function-name* specifies the name of the function.
- [OR REPLACE] option allows the modification of an existing function.



- The optional parameter list contains name, mode and types of the parameters. IN represents the value that will be passed from outside and OUT represents the parameter that will be used to return a value outside of the procedure.
- The function must contain a return statement.
- The *RETURN* clause specifies the data type you are going to return from the function.
- *function-body* contains the executable part.
- The AS keyword is used instead of the IS keyword for creating a standalone function.

### Example

The following example illustrates how to create and call a standalone function. This function returns the total number of CUSTOMERS in the customers table.

We will use the CUSTOMERS table, which we had created in the [PL/SQL Variables](#) chapter –

Select \* from customers;

```
+-----+-----+-----+-----+
| ID | NAME   | AGE | ADDRESS | SALARY |
+-----+-----+-----+-----+
```



```
| 1 | Ramesh | 32 | Ahmedabad | 2000.00 |  
| 2 | Khilan | 25 | Delhi | 1500.00 |  
| 3 | kaushik | 23 | Kota | 2000.00 |  
| 4 | Chaitali | 25 | Mumbai | 6500.00 |  
| 5 | Hardik | 27 | Bhopal | 8500.00 |  
| 6 | Komal | 22 | MP | 4500.00 |
```

```
+-----+-----+-----+-----+-----+
```

```
CREATE OR REPLACE FUNCTION totalCustomers
```

```
RETURN number IS
```

```
total number(2) := 0;
```

```
BEGIN
```

```
SELECT count(*) into total
```

```
FROM customers;
```

```
RETURN total;
```



```
END;
```

```
/
```

When the above code is executed using the SQL prompt, it will produce the following result –

Function created.

### Calling a Function

While creating a function, you give a definition of what the function has to do. To use a function, you will have to call that function to perform the defined task. When a program calls a function, the program control is transferred to the called function.

A called function performs the defined task and when its return statement is executed or when the last end statement is reached, it returns the program control back to the main program.

To call a function, you simply need to pass the required parameters along with the function name and if the function returns a value, then you can store the returned value. Following program calls the function `totalCustomers` from an anonymous block –

```
DECLARE
```

```
    c number(2);
```

```
BEGIN
```

```
    c := totalCustomers();
```



```
dbms_output.put_line('Total no. of Customers: ' || c);  
END;  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Total no. of Customers: 6

PL/SQL procedure successfully completed.

### Example

The following example demonstrates Declaring, Defining, and Invoking a Simple PL/SQL Function that computes and returns the maximum of two values.

```
DECLARE
```

```
  a number;
```

```
  b number;
```

```
  c number;
```

```
FUNCTION findMax(x IN number, y IN number)
```



RETURN number

IS

z number;

BEGIN

IF  $x > y$  THEN

z:= x;

ELSE

Z:= y;

END IF;

RETURN z;

END;

BEGIN

a:= 23;

b:= 45;

c := findMax(a, b);



```
dbms_output.put_line(' Maximum of (23,45): ' || c);  
END;  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Maximum of (23,45): 45

PL/SQL procedure successfully completed.

### PL/SQL Recursive Functions

We have seen that a program or subprogram may call another subprogram. When a subprogram calls itself, it is referred to as a recursive call and the process is known as recursion.

To illustrate the concept, let us calculate the factorial of a number. Factorial of a number  $n$  is defined as –

$$\begin{aligned} n! &= n*(n-1)! \\ &= n*(n-1)*(n-2)! \\ &\dots \\ &= n*(n-1)*(n-2)*(n-3)\dots 1 \end{aligned}$$



The following program calculates the factorial of a given number by calling itself recursively –

```
DECLARE
```

```
    num number;
```

```
    factorial number;
```

```
FUNCTION fact(x number)
```

```
RETURN number
```

```
IS
```

```
    f number;
```

```
BEGIN
```

```
    IF x=0 THEN
```

```
        f := 1;
```

```
    ELSE
```

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

```
    END IF;
```



```
RETURN f;
```

```
END;
```

```
BEGIN
```

```
    num:= 6;
```

```
    factorial := fact(num);
```

```
    dbms_output.put_line(' Factorial ' || num || ' is ' || factorial);
```

```
END;
```

```
/
```

When the above code is executed at the SQL prompt, it produces the following result –

```
Factorial 6 is 720
```

PL/SQL procedure successfully completed.

. A subprogram is a program unit/module that performs a particular task. These subprograms are combined to form larger programs. This is basically called the 'Modular design'. A subprogram can be invoked by another subprogram or program which is called the calling program.



A subprogram can be created –

- At the schema level
- Inside a package
- Inside a PL/SQL block

At the schema level, subprogram is a standalone subprogram. It is created with the `CREATE PROCEDURE` or the `CREATE FUNCTION` statement. It is stored in the database and can be deleted with the `DROP PROCEDURE` or `DROP FUNCTION` statement.

A subprogram created inside a package is a packaged subprogram. It is stored in the database and can be deleted only when the package is deleted with the `DROP PACKAGE` statement. We will discuss packages in the chapter 'PL/SQL - Packages'.

PL/SQL subprograms are named PL/SQL blocks that can be invoked with a set of parameters. PL/SQL provides two kinds of subprograms –

- Functions – These subprograms return a single value; mainly used to compute and return a value.
- Procedures – These subprograms do not return a value directly; mainly used to perform an action.

This chapter is going to cover important aspects of a PL/SQL procedure. We will discuss PL/SQL function in the next chapter.

Parts of a PL/SQL Subprogram



Each PL/SQL subprogram has a name, and may also have a parameter list. Like anonymous PL/SQL blocks, the named blocks will also have the following three parts –

S.No	Parts & Description
1	<p><b>Declarative Part</b></p> <p>It is an optional part. However, the declarative part for a subprogram does not start with the DECLARE keyword. It contains declarations of types, cursors, constants, variables, exceptions, and nested subprograms. These items are local to the subprogram and cease to exist when the subprogram completes execution.</p>
2	<p><b>Executable Part</b></p> <p>This is a mandatory part and contains statements that perform the designated action.</p>
3	<p><b>Exception-handling</b></p> <p>This is again an optional part. It contains the code that handles run-time errors.</p>



## Creating a Procedure

A procedure is created with the **CREATE OR REPLACE PROCEDURE** statement. The simplified syntax for the **CREATE OR REPLACE PROCEDURE** statement is as follows –

```
CREATE [OR REPLACE] PROCEDURE procedure_name
```

```
[(parameter_name [IN | OUT | IN OUT] type [, ...])]
```

```
{IS | AS}
```

```
BEGIN
```

```
< procedure_body >
```

```
END procedure_name;
```

Where,

- *procedure-name* specifies the name of the procedure.
- [OR REPLACE] option allows the modification of an existing procedure.
- The optional parameter list contains name, mode and types of the parameters. IN represents the value that will be passed from outside and OUT represents the parameter that will be used to return a value outside of the procedure.



- *procedure-body* contains the executable part.
- The AS keyword is used instead of the IS keyword for creating a standalone procedure.

### Example

The following example creates a simple procedure that displays the string 'Hello World!' on the screen when executed.

```
CREATE OR REPLACE PROCEDURE greetings
```

```
AS
```

```
BEGIN
```

```
    dbms_output.put_line('Hello World!');
```

```
END;
```

```
/
```

When the above code is executed using the SQL prompt, it will produce the following result –

Procedure created.

### Executing a Standalone Procedure

A standalone procedure can be called in two ways –



- Using the EXECUTE keyword
- Calling the name of the procedure from a PL/SQL block

The above procedure named 'greetings' can be called with the EXECUTE keyword as –  
EXECUTE greetings;

The above call will display –  
Hello World

PL/SQL procedure successfully completed.

The procedure can also be called from another PL/SQL block –

```
BEGIN
```

```
    greetings;
```

```
END;
```

```
/
```

The above call will display –

Hello World



PL/SQL procedure successfully completed.

### Deleting a Standalone Procedure

A standalone procedure is deleted with the `DROP PROCEDURE` statement. Syntax for deleting a procedure is –

```
DROP PROCEDURE procedure-name;
```

You can drop the greetings procedure by using the following statement –

```
DROP PROCEDURE greetings;
```

### Parameter Modes in PL/SQL Subprograms

The following table lists out the parameter modes in PL/SQL subprograms –

S.No	Parameter Mode & Description
------	------------------------------



## IN

1 An IN parameter lets you pass a value to the subprogram. It is a read-only parameter. Inside the subprogram, an IN parameter acts like a constant. It cannot be assigned a value. You can pass a constant, literal, initialized variable, or expression as an IN parameter. You can also initialize it to a default value; however, in that case, it is omitted from the subprogram call. It is the default mode of parameter passing. Parameters are passed by reference.

## OUT

2 An OUT parameter returns a value to the calling program. Inside the subprogram, an OUT parameter acts like a variable. You can change its value and reference the value after assigning it. The actual parameter must be variable and it is passed by value.



## IN OUT

3

An IN OUT parameter passes an initial value to a subprogram and returns an updated value to the caller. It can be assigned a value and the value can be read.

The actual parameter corresponding to an IN OUT formal parameter must be a variable, not a constant or an expression. Formal parameter must be assigned a value. Actual parameter is passed by value.

### IN & OUT Mode Example 1

This program finds the minimum of two values. Here, the procedure takes two numbers using the IN mode and returns their minimum using the OUT parameters.

```
DECLARE
```

```
  a number;
```

```
  b number;
```

```
  c number;
```

```
PROCEDURE findMin(x IN number, y IN number, z OUT number) IS
```

```
BEGIN
```



```
IF x < y THEN
    z:= x;
ELSE
    z:= y;
END IF;
END;
BEGIN
    a:= 23;
    b:= 45;
    findMin(a, b, c);
    dbms_output.put_line(' Minimum of (23, 45) : ' || c);
END;
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Minimum of (23, 45) : 23



PL/SQL procedure successfully completed.

### IN & OUT Mode Example 2

This procedure computes the square of value of a passed value. This example shows how we can use the same parameter to accept a value and then return another result.

```
DECLARE
```

```
    a number;
```

```
PROCEDURE squareNum(x IN OUT number) IS
```

```
BEGIN
```

```
    x := x * x;
```

```
END;
```

```
BEGIN
```

```
    a:= 23;
```

```
    squareNum(a);
```

```
    dbms_output.put_line(' Square of (23): ' || a);
```



END;

/

When the above code is executed at the SQL prompt, it produces the following result –

Square of (23): 529

PL/SQL procedure successfully completed.

Methods for Passing Parameters

Actual parameters can be passed in three ways –

- Positional notation
- Named notation
- Mixed notation

Positional Notation

In positional notation, you can call the procedure as –

findMin(a, b, c, d);



In positional notation, the first actual parameter is substituted for the first formal parameter; the second actual parameter is substituted for the second formal parameter, and so on. So, a is substituted for x, b is substituted for y, c is substituted for z and d is substituted for m.

### Named Notation

In named notation, the actual parameter is associated with the formal parameter using the arrow symbol ( => ). The procedure call will be like the following –

```
findMin(x => a, y => b, z => c, m => d);
```

### Mixed Notation

In mixed notation, you can mix both notations in procedure call; however, the positional notation should precede the named notation.

The following call is legal –

```
findMin(a, b, c, m => d);
```

However, this is not legal:

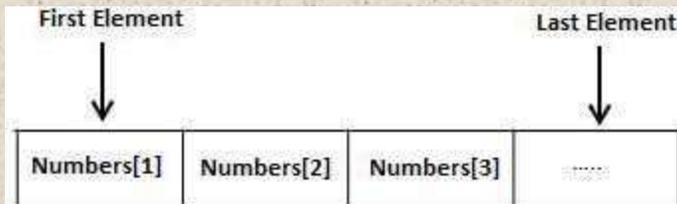
```
findMin(x => a, b, c, d);
```

In this chapter, we will discuss arrays in PL/SQL. The PL/SQL programming language provides a data structure called the VARRAY, which can store a fixed-size sequential collection of elements of the same type. A varray is



used to store an ordered collection of data, however it is often better to think of an array as a collection of variables of the same type.

All varrays consist of contiguous memory locations. The lowest address corresponds to the first element and the highest address to the last element.



An array is a part of collection type data and it stands for variable-size arrays. We will study other collection types in a later chapter 'PL/SQL Collections'.

Each element in a varray has an index associated with it. It also has a maximum size that can be changed dynamically.

### Creating a Varray Type

A varray type is created with the CREATE TYPE statement. You must specify the maximum size and the type of elements stored in the varray.

The basic syntax for creating a VARRAY type at the schema level is –

```
CREATE OR REPLACE TYPE varray_type_name IS VARRAY(n) of <element_type>
```

Where,



- *varray\_type\_name* is a valid attribute name,
- *n* is the number of elements (maximum) in the varray,
- *element\_type* is the data type of the elements of the array.

Maximum size of a varray can be changed using the ALTER TYPE statement.

For example,

```
CREATE Or REPLACE TYPE namearray AS VARRAY(3) OF VARCHAR2(10);  
/
```

Type created.

The basic syntax for creating a VARRAY type within a PL/SQL block is –

```
TYPE varray_type_name IS VARRAY(n) of <element_type>
```

For example –

```
TYPE namearray IS VARRAY(5) OF VARCHAR2(10);
```

```
Type grades IS VARRAY(5) OF INTEGER;
```

Let us now work out on a few examples to understand the concept –



### Example 1

The following program illustrates the use of varrays –

**DECLARE**

```
type namesarray IS VARRAY(5) OF VARCHAR2(10);
```

```
type grades IS VARRAY(5) OF INTEGER;
```

```
names namesarray;
```

```
marks grades;
```

```
total integer;
```

**BEGIN**

```
names := namesarray('Kavita', 'Pritam', 'Ayan', 'Rishav', 'Aziz');
```

```
marks:= grades(98, 97, 78, 87, 92);
```

```
total := names.count;
```

```
dbms_output.put_line('Total ' || total || ' Students');
```

```
FOR i in 1 .. total LOOP
```

```
    dbms_output.put_line('Student: ' || names(i) || '
```



```
Marks: ' || marks(i));
```

```
END LOOP;
```

```
END;
```

```
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Total 5 Students

Student: Kavita Marks: 98

Student: Pritam Marks: 97

Student: Ayan Marks: 78

Student: Rishav Marks: 87

Student: Aziz Marks: 92

PL/SQL procedure successfully completed.

Please note –

- In Oracle environment, the starting index for varrays is always 1.



- You can initialize the varray elements using the constructor method of the varray type, which has the same name as the varray.
- Varrays are one-dimensional arrays.
- A varray is automatically NULL when it is declared and must be initialized before its elements can be referenced.

### Example 2

Elements of a varray could also be a %ROWTYPE of any database table or %TYPE of any database table field. The following example illustrates the concept.

We will use the CUSTOMERS table stored in our database as –

Select \* from customers;

```
+---+-----+---+-----+-----+
| ID | NAME   | AGE | ADDRESS | SALARY |
+---+-----+---+-----+-----+
| 1 | Ramesh | 32 | Ahmedabad | 2000.00 |
| 2 | Khilan | 25 | Delhi    | 1500.00 |
```



```
| 3 | kaushik | 23 | Kota    | 2000.00 |
| 4 | Chaitali | 25 | Mumbai  | 6500.00 |
| 5 | Hardik   | 27 | Bhopal  | 8500.00 |
| 6 | Komal    | 22 | MP      | 4500.00 |
+----+-----+-----+-----+-----+
```

Following example makes the use of cursor, which you will study in detail in a separate chapter.

```
DECLARE
```

```
    CURSOR c_customers is
```

```
    SELECT name FROM customers;
```

```
    type c_list is varray (6) of customers.name%type;
```

```
    name_list c_list := c_list();
```

```
    counter integer :=0;
```

```
BEGIN
```

```
    FOR n IN c_customers LOOP
```

```
        counter := counter + 1;
```



```
name_list.extend;  
name_list(counter) := n.name;  
dbms_output.put_line('Customer('||counter ||'):'||name_list(counter));  
END LOOP;  
END;  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Customer(1): Ramesh

Customer(2): Khilan

Customer(3): kaushik

Customer(4): Chaitali

Customer(5): Hardik

Customer(6): Komal

PL/SQL procedure successfully completed.



A cursor is a pointer to this context area. PL/SQL controls the context area through a cursor. A cursor holds the rows (one or more) returned by a SQL statement. The set of rows the cursor holds is referred to as the active set.

You can name a cursor so that it could be referred to in a program to fetch and process the rows returned by the SQL statement, one at a time. There are two types of cursors –

- Implicit cursors
- Explicit cursors

### Implicit Cursors

Implicit cursors are automatically created by Oracle whenever an SQL statement is executed, when there is no explicit cursor for the statement. Programmers cannot control the implicit cursors and the information in it.

Whenever a DML statement (INSERT, UPDATE and DELETE) is issued, an implicit cursor is associated with this statement. For INSERT operations, the cursor holds the data that needs to be inserted. For UPDATE and DELETE operations, the cursor identifies the rows that would be affected.

In PL/SQL, you can refer to the most recent implicit cursor as the SQL cursor, which always has attributes such as %FOUND, %ISOPEN, %NOTFOUND, and %ROWCOUNT. The SQL cursor has additional attributes, %BULK\_ROWCOUNT and %BULK\_EXCEPTIONS, designed for use with the FORALL statement. The following table provides the description of the most used attributes –

S.No	Attribute & Description
------	-------------------------



1	<p><b>%FOUND</b></p> <p>Returns TRUE if an INSERT, UPDATE, or DELETE statement affected one or more rows or a SELECT INTO statement returned one or more rows. Otherwise, it returns FALSE.</p>
2	<p><b>%NOTFOUND</b></p> <p>The logical opposite of %FOUND. It returns TRUE if an INSERT, UPDATE, or DELETE statement affected no rows, or a SELECT INTO statement returned no rows. Otherwise, it returns FALSE.</p>
3	<p><b>%ISOPEN</b></p> <p>Always returns FALSE for implicit cursors, because Oracle closes the SQL cursor automatically after executing its associated SQL statement.</p>
4	<p><b>%ROWCOUNT</b></p> <p>Returns the number of rows affected by an INSERT, UPDATE, or DELETE statement, or returned by a SELECT INTO statement.</p>

Any SQL cursor attribute will be accessed as `sql%attribute_name` as shown below in the example.



## Example

We will be using the CUSTOMERS table we had created and used in the previous chapters.

Select \* from customers;

```
+----+-----+----+-----+-----+
| ID | NAME   | AGE | ADDRESS | SALARY |
+----+-----+----+-----+-----+
| 1 | Ramesh | 32 | Ahmedabad | 2000.00 |
| 2 | Khilan | 25 | Delhi     | 1500.00 |
| 3 | kaushik | 23 | Kota      | 2000.00 |
| 4 | Chaitali | 25 | Mumbai   | 6500.00 |
| 5 | Hardik | 27 | Bhopal    | 8500.00 |
| 6 | Komal | 22 | MP        | 4500.00 |
+----+-----+----+-----+-----+
```



The following program will update the table and increase the salary of each customer by 500 and use the SQL%ROWCOUNT attribute to determine the number of rows affected –

DECLARE

total\_rows number(2);

BEGIN

UPDATE customers

SET salary = salary + 500;

IF sql%notfound THEN

dbms\_output.put\_line('no customers selected');

ELSIF sql%found THEN

total\_rows := sql%rowcount;

dbms\_output.put\_line( total\_rows || ' customers selected ');

END IF;

END;

/



When the above code is executed at the SQL prompt, it produces the following result –

6 customers selected

PL/SQL procedure successfully completed.

If you check the records in customers table, you will find that the rows have been updated –

Select \* from customers;

```
+---+-----+---+-----+-----+
| ID | NAME   | AGE | ADDRESS | SALARY |
+---+-----+---+-----+-----+
| 1 | Ramesh | 32 | Ahmedabad | 2500.00 |
| 2 | Khilan | 25 | Delhi     | 2000.00 |
| 3 | kaushik | 23 | Kota      | 2500.00 |
| 4 | Chaitali | 25 | Mumbai   | 7000.00 |
| 5 | Hardik | 27 | Bhopal    | 9000.00 |
```



| 6 | Komal | 22 | MP | 5000.00 |

+-----+-----+-----+-----+-----+

### Explicit Cursors

Explicit cursors are programmer-defined cursors for gaining more control over the context area. An explicit cursor should be defined in the declaration section of the PL/SQL Block. It is created on a SELECT Statement which returns more than one row.

The syntax for creating an explicit cursor is –

```
CURSOR cursor_name IS select_statement;
```

Working with an explicit cursor includes the following steps –

- Declaring the cursor for initializing the memory
- Opening the cursor for allocating the memory
- Fetching the cursor for retrieving the data
- Closing the cursor to release the allocated memory

### Declaring the Cursor

Declaring the cursor defines the cursor with a name and the associated SELECT statement. For example –

```
CURSOR c_customers IS
```



```
SELECT id, name, address FROM customers;
```

### Opening the Cursor

Opening the cursor allocates the memory for the cursor and makes it ready for fetching the rows returned by the SQL statement into it. For example, we will open the above defined cursor as follows –

```
OPEN c_customers;
```

### Fetching the Cursor

Fetching the cursor involves accessing one row at a time. For example, we will fetch rows from the above-opened cursor as follows –

```
FETCH c_customers INTO c_id, c_name, c_addr;
```

### Closing the Cursor

Closing the cursor means releasing the allocated memory. For example, we will close the above-opened cursor as follows –

```
CLOSE c_customers;
```

### Example

Following is a complete example to illustrate the concepts of explicit cursors & minus;

```
DECLARE
```



```
c_id customers.id%type;  
c_name customer.name%type;  
c_addr customers.address%type;  
CURSOR c_customers is  
    SELECT id, name, address FROM customers;  
BEGIN  
    OPEN c_customers;  
    LOOP  
    FETCH c_customers into c_id, c_name, c_addr;  
        EXIT WHEN c_customers%notfound;  
        dbms_output.put_line(c_id || ' ' || c_name || ' ' || c_addr);  
    END LOOP;  
    CLOSE c_customers;  
END;  
/
```



When the above code is executed at the SQL prompt, it produces the following result –

1 Ramesh Ahmedabad

2 Khilan Delhi

3 kaushik Kota

4 Chaitali Mumbai

5 Hardik Bhopal

6 Komal MP

PL/SQL procedure successfully completed.

In this chapter, we will discuss Exceptions in PL/SQL. An exception is an error condition during a program execution. PL/SQL supports programmers to catch such conditions using EXCEPTION block in the program and an appropriate action is taken against the error condition. There are two types of exceptions –

- System-defined exceptions
- User-defined exceptions

Syntax for Exception Handling



The general syntax for exception handling is as follows. Here you can list down as many exceptions as you can handle. The default exception will be handled using *WHEN others THEN* –

DECLARE

<declarations section>

BEGIN

<executable command(s)>

EXCEPTION

<exception handling goes here >

WHEN exception1 THEN

exception1-handling-statements

WHEN exception2 THEN

exception2-handling-statements

WHEN exception3 THEN

exception3-handling-statements

.....



WHEN others THEN

exception3-handling-statements

END;

Example

Let us write a code to illustrate the concept. We will be using the CUSTOMERS table we had created and used in the previous chapters –

DECLARE

c\_id customers.id%type := 8;

c\_name customerS.Name%type;

c\_addr customers.address%type;

BEGIN

SELECT name, address INTO c\_name, c\_addr

FROM customers

WHERE id = c\_id;

DBMS\_OUTPUT.PUT\_LINE ('Name: ' || c\_name);



```
DBMS_OUTPUT.PUT_LINE ('Address: ' || c_addr);
```

```
EXCEPTION
```

```
WHEN no_data_found THEN
```

```
    dbms_output.put_line('No such customer!');
```

```
WHEN others THEN
```

```
    dbms_output.put_line('Error!');
```

```
END;
```

```
/
```

When the above code is executed at the SQL prompt, it produces the following result –

No such customer!

PL/SQL procedure successfully completed.

The above program displays the name and address of a customer whose ID is given. Since there is no customer with ID value 8 in our database, the program raises the run-time exception `NO_DATA_FOUND`, which is captured in the `EXCEPTION` block.



## Raising Exceptions

Exceptions are raised by the database server automatically whenever there is any internal database error, but exceptions can be raised explicitly by the programmer by using the command RAISE. Following is the simple syntax for raising an exception –

```
DECLARE
```

```
    exception_name EXCEPTION;
```

```
BEGIN
```

```
    IF condition THEN
```

```
        RAISE exception_name;
```

```
    END IF;
```

```
EXCEPTION
```

```
    WHEN exception_name THEN
```

```
        statement;
```

```
END;
```

You can use the above syntax in raising the Oracle standard exception or any user-defined exception. In the next section, we will give you an example on raising a user-defined exception. You can raise the Oracle standard exceptions in a similar way.



## User-defined Exceptions

PL/SQL allows you to define your own exceptions according to the need of your program. A user-defined exception must be declared and then raised explicitly, using either a RAISE statement or the procedure DBMS\_STANDARD.RAISE\_APPLICATION\_ERROR.

The syntax for declaring an exception is –

DECLARE

```
my-exception EXCEPTION;
```

### Example

The following example illustrates the concept. This program asks for a customer ID, when the user enters an invalid ID, the exception `invalid_id` is raised.

DECLARE

```
c_id customers.id%type := &cc_id;
```

```
c_name customerS.Name%type;
```

```
c_addr customers.address%type;
```

```
-- user defined exception
```

```
ex_invalid_id EXCEPTION;
```



```
BEGIN
```

```
IF c_id <= 0 THEN
```

```
    RAISE ex_invalid_id;
```

```
ELSE
```

```
    SELECT name, address INTO c_name, c_addr
```

```
    FROM customers
```

```
    WHERE id = c_id;
```

```
    DBMS_OUTPUT.PUT_LINE ('Name: ' || c_name);
```

```
    DBMS_OUTPUT.PUT_LINE ('Address: ' || c_addr);
```

```
END IF;
```

```
EXCEPTION
```

```
    WHEN ex_invalid_id THEN
```

```
        dbms_output.put_line('ID must be greater than zero!');
```

```
    WHEN no_data_found THEN
```



```
dbms_output.put_line('No such customer!');
```

```
WHEN others THEN
```

```
dbms_output.put_line('Error!');
```

```
END;
```

```
/
```

When the above code is executed at the SQL prompt, it produces the following result –

```
Enter value for cc_id: -6 (let's enter a value -6)
```

```
old 2: c_id customers.id%type := &cc_id;
```

```
new 2: c_id customers.id%type := -6;
```

```
ID must be greater than zero!
```

PL/SQL procedure successfully completed.

### Pre-defined Exceptions

PL/SQL provides many pre-defined exceptions, which are executed when any database rule is violated by a program. For example, the predefined exception `NO_DATA_FOUND` is raised when a `SELECT INTO` statement returns no rows. The following table lists few of the important pre-defined exceptions –



Exception	Oracle Error	SQLCODE	Description
ACCESS_INTO_NULL	06530	-6530	It is raised when a null object is automatically assigned a value.
CASE_NOT_FOUND	06592	-6592	It is raised when none of the choices in the WHEN clause of a CASE statement is selected, and there is no ELSE clause.
COLLECTION_IS_NULL	06531	-6531	It is raised when a program attempts to apply collection methods other than EXISTS to an uninitialized nested table or varray, or the program attempts to assign values



			to the elements of an uninitialized nested table or varray.
DUP_VAL_ON_INDEX	00001	-1	It is raised when duplicate values are attempted to be stored in a column with unique index.
INVALID_CURSOR	01001	-1001	It is raised when attempts are made to make a cursor operation that is not allowed, such as closing an unopened cursor.
INVALID_NUMBER	01722	-1722	It is raised when the conversion of a character string into a number fails because the string does



			not represent a valid number.
LOGIN_DENIED	01017	-1017	It is raised when a program attempts to log on to the database with an invalid username or password.
NO_DATA_FOUND	01403	+100	It is raised when a SELECT INTO statement returns no rows.
NOT_LOGGED_ON	01012	-1012	It is raised when a database call is issued without being connected to the database.
PROGRAM_ERROR	06501	-6501	It is raised when PL/SQL has an internal



			problem.
ROWTYPE_MISMATCH	06504	-6504	It is raised when a cursor fetches value in a variable having incompatible data type.
SELF_IS_NULL	30625	-30625	It is raised when a member method is invoked, but the instance of the object type was not initialized.
STORAGE_ERROR	06500	-6500	It is raised when PL/SQL ran out of memory or memory was corrupted.
TOO_MANY_ROWS	01422	-1422	It is raised when a SELECT INTO statement returns more



			than one row.
VALUE_ERROR	06502	-6502	It is raised when an arithmetic, conversion, truncation, or sizeconstraint error occurs.
ZERO_DIVIDE	01476	1476	It is raised when an attempt is made to divide a number by zero.

## Triggers

Triggers are stored programs, which are automatically executed or fired when some events occur. Triggers are, in fact, written to be executed in response to any of the following events –

- A database manipulation (DML) statement (DELETE, INSERT, or UPDATE)
- A database definition (DDL) statement (CREATE, ALTER, or DROP).
- A database operation (SERVERERROR, LOGON, LOGOFF, STARTUP, or SHUTDOWN).



Triggers can be defined on the table, view, schema, or database with which the event is associated.

### Benefits of Triggers

Triggers can be written for the following purposes –

- Generating some derived column values automatically
- Enforcing referential integrity
- Event logging and storing information on table access
- Auditing
- Synchronous replication of tables
- Imposing security authorizations
- Preventing invalid transactions

### Creating Triggers

The syntax for creating a trigger is –

```
CREATE [OR REPLACE ] TRIGGER trigger_name
```

```
{BEFORE | AFTER | INSTEAD OF }
```

```
{INSERT [OR] | UPDATE [OR] | DELETE}
```



[OF col\_name]

ON table\_name

[REFERENCING OLD AS o NEW AS n]

[FOR EACH ROW]

WHEN (condition)

DECLARE

Declaration-statements

BEGIN

Executable-statements

EXCEPTION

Exception-handling-statements

END;

Where,

- CREATE [OR REPLACE] TRIGGER trigger\_name – Creates or replaces an existing trigger with the *trigger\_name*.



- {BEFORE | AFTER | INSTEAD OF} – This specifies when the trigger will be executed. The INSTEAD OF clause is used for creating trigger on a view.
- {INSERT [OR] | UPDATE [OR] | DELETE} – This specifies the DML operation.
- [OF col\_name] – This specifies the column name that will be updated.
- [ON table\_name] – This specifies the name of the table associated with the trigger.
- [REFERENCING OLD AS o NEW AS n] – This allows you to refer new and old values for various DML statements, such as INSERT, UPDATE, and DELETE.
- [FOR EACH ROW] – This specifies a row-level trigger, i.e., the trigger will be executed for each row being affected. Otherwise the trigger will execute just once when the SQL statement is executed, which is called a table level trigger.
- WHEN (condition) – This provides a condition for rows for which the trigger would fire. This clause is valid only for row-level triggers.

### Example

To start with, we will be using the CUSTOMERS table we had created and used in the previous chapters –

Select \* from customers;

+-----+-----+-----+-----+-----+



ID	NAME	AGE	ADDRESS	SALARY
1	Ramesh	32	Ahmedabad	2000.00
2	Khilan	25	Delhi	1500.00
3	kaushik	23	Kota	2000.00
4	Chaitali	25	Mumbai	6500.00
5	Hardik	27	Bhopal	8500.00
6	Komal	22	MP	4500.00

The following program creates a row-level trigger for the customers table that would fire for INSERT or UPDATE or DELETE operations performed on the CUSTOMERS table. This trigger will display the salary difference between the old values and new values –

```
CREATE OR REPLACE TRIGGER display_salary_changes
BEFORE DELETE OR INSERT OR UPDATE ON customers
FOR EACH ROW
WHEN (NEW.ID > 0)
```



DECLARE

```
sal_diff number;
```

BEGIN

```
sal_diff := :NEW.salary - :OLD.salary;
```

```
dbms_output.put_line('Old salary: ' || :OLD.salary);
```

```
dbms_output.put_line('New salary: ' || :NEW.salary);
```

```
dbms_output.put_line('Salary difference: ' || sal_diff);
```

END;

/

When the above code is executed at the SQL prompt, it produces the following result –

Trigger created.

The following points need to be considered here –

- OLD and NEW references are not available for table-level triggers, rather you can use them for record-level triggers.



- If you want to query the table in the same trigger, then you should use the AFTER keyword, because triggers can query the table or change it again only after the initial changes are applied and the table is back in a consistent state.
- The above trigger has been written in such a way that it will fire before any DELETE or INSERT or UPDATE operation on the table, but you can write your trigger on a single or multiple operations, for example BEFORE DELETE, which will fire whenever a record will be deleted using the DELETE operation on the table.

### Triggering a Trigger

Let us perform some DML operations on the CUSTOMERS table. Here is one INSERT statement, which will create a new record in the table –

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (7, 'Kriti', 22, 'HP', 7500.00 );
```

When a record is created in the CUSTOMERS table, the above create trigger, display\_salary\_changes will be fired and it will display the following result –

Old salary:

New salary: 7500

Salary difference:



Because this is a new record, old salary is not available and the above result comes as null. Let us now perform one more DML operation on the CUSTOMERS table. The UPDATE statement will update an existing record in the table –

```
UPDATE customers
```

```
SET salary = salary + 500
```

```
WHERE id = 2;
```

When a record is updated in the CUSTOMERS table, the above create trigger, display\_salary\_changes will be fired and it will display the following result –

Old salary: 1500

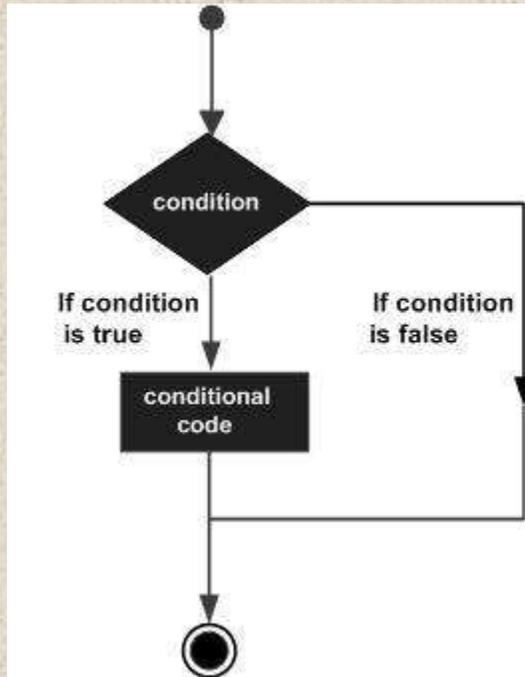
New salary: 2000

Salary difference: 500

conditions in PL/SQL

Decision-making structures require that the programmer specify one or more conditions to be evaluated or tested by the program, along with a statement or statements to be executed if the condition is determined to be true, and optionally, other statements to be executed if the condition is determined to be false.

Following is the general form of a typical conditional (i.e., decision making) structure found in most of the programming languages –



PL/SQL programming language provides following types of decision-making statements. Click the following links to check their detail.

S.No	Statement & Description
------	-------------------------



### IF - THEN statement

1

The IF statement associates a condition with a sequence of statements enclosed by the keywords THEN and END IF. If the condition is true, the statements get executed and if the condition is false or NULL then the IF statement does nothing.

### IF-THEN-ELSE statement

2

IF statement adds the keyword ELSE followed by an alternative sequence of statement. If the condition is false or NULL, then only the alternative sequence of statements get executed. It ensures that either of the sequence of statements is executed.

### IF-THEN-ELSIF statement

3

It allows you to choose between several alternatives.



#### Case statement

4

Like the IF statement, the CASE statement selects one sequence of statements to execute.

However, to select the sequence, the CASE statement uses a selector rather than multiple Boolean expressions. A selector is an expression whose value is used to select one of several alternatives.

#### Searched CASE statement

5

The searched CASE statement has no selector, and its WHEN clauses contain search conditions that yield Boolean values.

#### nested IF-THEN-ELSE

6

You can use one IF-THEN or IF-THEN-ELSIF statement inside another IF-THEN or IF-THEN-ELSIF statement(s).

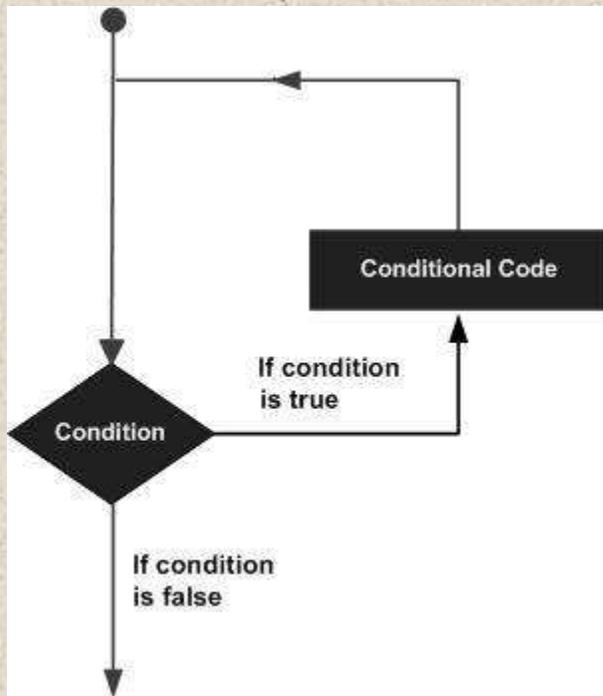
Loops in PL/SQL.

here may be a situation when you need to execute a block of code several number of times. In general, statements are executed sequentially: The first statement in a function is executed first, followed by the second, and so on.

Programming languages provide various control structures that allow for more complicated execution paths.



A loop statement allows us to execute a statement or group of statements multiple times and following is the general form of a loop statement in most of the programming languages –



PL/SQL provides the following types of loop to handle the looping requirements. Click the following links to check their detail.

S.No	Loop Type & Description
------	-------------------------



### PL/SQL Basic LOOP

1

In this loop structure, sequence of statements is enclosed between the LOOP and the END LOOP statements. At each iteration, the sequence of statements is executed and then control resumes at the top of the loop.

### PL/SQL WHILE LOOP

2

Repeats a statement or group of statements while a given condition is true. It tests the condition before executing the loop body.

### PL/SQL FOR LOOP

3

Execute a sequence of statements multiple times and abbreviates the code that manages the loop variable.

### Nested loops in PL/SQL

4

You can use one or more loop inside any another basic loop, while, or for loop.

Labeling a PL/SQL Loop



PL/SQL loops can be labeled. The label should be enclosed by double angle brackets (<< and >>) and appear at the beginning of the LOOP statement. The label name can also appear at the end of the LOOP statement. You may use the label in the EXIT statement to exit from the loop.

The following program illustrates the concept –

```
DECLARE
```

```
  i number(1);
```

```
  j number(1);
```

```
BEGIN
```

```
  << outer_loop >>
```

```
  FOR i IN 1..3 LOOP
```

```
    << inner_loop >>
```

```
    FOR j IN 1..3 LOOP
```

```
      dbms_output.put_line('i is: ' || i || ' and j is: ' || j);
```

```
    END loop inner_loop;
```

```
  END loop outer_loop;
```

```
END;
```



/

When the above code is executed at the SQL prompt, it produces the following result –

i is: 1 and j is: 1

i is: 1 and j is: 2

i is: 1 and j is: 3

i is: 2 and j is: 1

i is: 2 and j is: 2

i is: 2 and j is: 3

i is: 3 and j is: 1

i is: 3 and j is: 2

i is: 3 and j is: 3

PL/SQL procedure successfully completed.

The Loop Control Statements



Loop control statements change execution from its normal sequence. When execution leaves a scope, all automatic objects that were created in that scope are destroyed.

PL/SQL supports the following control statements. Labeling loops also help in taking the control outside a loop. Click the following links to check their details.

S.No	Control Statement & Description
1	<a href="#"><u>EXIT statement</u></a> The Exit statement completes the loop and control passes to the statement immediately after the END LOOP.
2	<a href="#"><u>CONTINUE statement</u></a> Causes the loop to skip the remainder of its body and immediately retest its condition prior to reiterating.
3	<a href="#"><u>GOTO statement</u></a> Transfers control to the labeled statement. Though it is not advised to use the GOTO statement in your program.

string in PL/SQL



The string in PL/SQL is actually a sequence of characters with an optional size specification. The characters could be numeric, letters, blank, special characters or a combination of all. PL/SQL offers three kinds of strings –

- Fixed-length strings – In such strings, programmers specify the length while declaring the string. The string is right-padded with spaces to the length so specified.
- Variable-length strings – In such strings, a maximum length up to 32,767, for the string is specified and no padding takes place.
- Character large objects (CLOBs) – These are variable-length strings that can be up to 128 terabytes.

PL/SQL strings could be either variables or literals. A string literal is enclosed within quotation marks. For example,

'This is a string literal.' Or 'hello world'

To include a single quote inside a string literal, you need to type two single quotes next to one another. For example,

'this isn't what it looks like'

Declaring String Variables

Oracle database provides numerous string datatypes, such as CHAR, NCHAR, VARCHAR2, NVARCHAR2, CLOB, and NCLOB. The datatypes prefixed with an 'N' are 'national character set' datatypes, that store Unicode character data.



If you need to declare a variable-length string, you must provide the maximum length of that string. For example, the VARCHAR2 data type. The following example illustrates declaring and using some string variables –

**DECLARE**

```
name varchar2(20);
```

```
company varchar2(30);
```

```
introduction clob;
```

```
choice char(1);
```

**BEGIN**

```
name := 'John Smith';
```

```
company := 'Infotech';
```

```
introduction := ' Hello! I'm John Smith from Infotech.';
```

```
choice := 'y';
```

```
IF choice = 'y' THEN
```

```
    dbms_output.put_line(name);
```

```
    dbms_output.put_line(company);
```



```
dbms_output.put_line(introduction);  
END IF;  
END;  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

John Smith

Infotech

Hello! I'm John Smith from Infotech.

PL/SQL procedure successfully completed

To declare a fixed-length string, use the CHAR datatype. Here you do not have to specify a maximum length for a fixed-length variable. If you leave off the length constraint, Oracle Database automatically uses a maximum length required. The following two declarations are identical –

```
red_flag CHAR(1) := 'Y';
```

```
red_flag CHAR := 'Y';
```

PL/SQL String Functions and Operators



PL/SQL offers the concatenation operator (||) for joining two strings. The following table provides the string functions provided by PL/SQL –

S.No	Function & Purpose
1	<code>ASCII(x);</code> Returns the ASCII value of the character x.
2	<code>CHR(x);</code> Returns the character with the ASCII value of x.
3	<code>CONCAT(x, y);</code> Concatenates the strings x and y and returns the appended string.
4	<code>INITCAP(x);</code> Converts the initial letter of each word in x to uppercase and returns that string.
5	<code>INSTR(x, find_string [, start] [, occurrence]);</code>



Searches for find\_string in x and returns the position at which it occurs.

6 INSTRB(x);  
Returns the location of a string within another string, but returns the value in bytes.

7 LENGTH(x);  
Returns the number of characters in x.

8 LENGTHB(x);  
Returns the length of a character string in bytes for single byte character set.

9 LOWER(x);  
Converts the letters in x to lowercase and returns that string.

10 LPAD(x, width [, pad\_string] );  
Pads x with spaces to the left, to bring the total length of the string up to



width characters.

11	<p>LTRIM(x [, trim_string]);</p> <p>Trims characters from the left of x.</p>
12	<p>NANVL(x, value);</p> <p>Returns value if x matches the NaN special value (not a number), otherwise x is returned.</p>
13	<p>NLS_INITCAP(x);</p> <p>Same as the INITCAP function except that it can use a different sort method as specified by NLSSORT.</p>
14	<p>NLS_LOWER(x) ;</p> <p>Same as the LOWER function except that it can use a different sort method as specified by NLSSORT.</p>
15	<p>NLS_UPPER(x);</p>



Same as the UPPER function except that it can use a different sort method as specified by NLSSORT.

16 NLSSORT(x);  
Changes the method of sorting the characters. Must be specified before any NLS function; otherwise, the default sort will be used.

17 NVL(x, value);  
Returns value if x is null; otherwise, x is returned.

18 NVL2(x, value1, value2);  
Returns value1 if x is not null; if x is null, value2 is returned.

19 REPLACE(x, search\_string, replace\_string);  
Searches x for search\_string and replaces it with replace\_string.

20 RPAD(x, width [, pad\_string]);  
Pads x to the right.



21 RTRIM(x [, trim\_string]);

Trims x from the right.

22 SOUNDEX(x) ;

Returns a string containing the phonetic representation of x.

23 SUBSTR(x, start [, length]);

Returns a substring of x that begins at the position specified by start. An optional length for the substring may be supplied.

24 SUBSTRB(x);

Same as SUBSTR except that the parameters are expressed in bytes instead of characters for the single-byte character systems.

25 TRIM([trim\_char FROM] x);

Trims characters from the left and right of x.

26 UPPER(x);



Converts the letters in x to uppercase and returns that string.

Let us now work out on a few examples to understand the concept –

Example 1

DECLARE

```
greetings varchar2(11) := 'hello world';
```

BEGIN

```
dbms_output.put_line(UPPER(greetings));
```

```
dbms_output.put_line(LOWER(greetings));
```

```
dbms_output.put_line(INITCAP(greetings));
```

```
/* retrieve the first character in the string */
```

```
dbms_output.put_line ( SUBSTR (greetings, 1, 1));
```



```
/* retrieve the last character in the string */  
dbms_output.put_line ( SUBSTR (greetings, -1, 1));  
  
/* retrieve five characters,  
   starting from the seventh position. */  
dbms_output.put_line ( SUBSTR (greetings, 7, 5));  
  
/* retrieve the remainder of the string,  
   starting from the second position. */  
dbms_output.put_line ( SUBSTR (greetings, 2));  
  
/* find the location of the first "e" */  
dbms_output.put_line ( INSTR (greetings, 'e'));  
END;  
/
```



When the above code is executed at the SQL prompt, it produces the following result –

HELLO WORLD

hello world

Hello World

h

d

World

ello World

2

PL/SQL procedure successfully completed.

Example 2

DECLARE

greetings varchar2(30) := '.....Hello World.....';

BEGIN



```
dbms_output.put_line(RTRIM(greetings, '.'));  
dbms_output.put_line(LTRIM(greetings, '.'));  
dbms_output.put_line(TRIM( '.' from greetings));  
END;  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

.....Hello World

Hello World.....

Hello World

PL/SQL procedure successfully completed.

### **Packages in PL/SQL**

Packages are schema objects that groups logically related PL/SQL types, variables, and subprograms.

A package will have two mandatory parts –

- Package specification



- Package body or definition

### Package Specification

The specification is the interface to the package. It just **DECLARES** the types, variables, constants, exceptions, cursors, and subprograms that can be referenced from outside the package. In other words, it contains all information about the content of the package, but excludes the code for the subprograms.

All objects placed in the specification are called public objects. Any subprogram not in the package specification but coded in the package body is called a private object.

The following code snippet shows a package specification having a single procedure. You can have many global variables defined and multiple procedures or functions inside a package.

```
CREATE PACKAGE cust_sal AS  
  
    PROCEDURE find_sal(c_id customers.id%type);  
  
END cust_sal;  
  
/
```

When the above code is executed at the SQL prompt, it produces the following result –

Package created.

Package Body



The package body has the codes for various methods declared in the package specification and other private declarations, which are hidden from the code outside the package.

The CREATE PACKAGE BODY Statement is used for creating the package body. The following code snippet shows the package body declaration for the *cust\_sal* package created above. I assumed that we already have CUSTOMERS table created in our database as mentioned in the [PL/SQL - Variables](#) chapter.

```
CREATE OR REPLACE PACKAGE BODY cust_sal AS
```

```
    PROCEDURE find_sal(c_id customers.id%TYPE) IS
        c_sal customers.salary%TYPE;
    BEGIN
        SELECT salary INTO c_sal
        FROM customers
        WHERE id = c_id;
        dbms_output.put_line('Salary: ' || c_sal);
    END find_sal;
END cust_sal;
```



/

When the above code is executed at the SQL prompt, it produces the following result –

Package body created.

Using the Package Elements

The package elements (variables, procedures or functions) are accessed with the following syntax –

`package_name.element_name;`

Consider, we already have created the above package in our database schema, the following program uses the *find\_sal* method of the *cust\_sal* package –

DECLARE

`code customers.id%type := &cc_id;`

BEGIN

`cust_sal.find_sal(code);`

END;

/

When the above code is executed at the SQL prompt, it prompts to enter the customer ID and when you enter an ID, it displays the corresponding salary as follows –



Enter value for cc\_id: 1

Salary: 3000

PL/SQL procedure successfully completed.

### Example

The following program provides a more complete package. We will use the CUSTOMERS table stored in our database with the following records –

Select \* from customers;

```
+----+-----+----+-----+-----+
| ID | NAME   | AGE | ADDRESS | SALARY |
+----+-----+----+-----+-----+
| 1 | Ramesh | 32  | Ahmedabad | 3000.00 |
| 2 | Khilan | 25  | Delhi     | 3000.00 |
| 3 | kaushik | 23  | Kota      | 3000.00 |
```



| 4 | Chaitali | 25 | Mumbai | 7500.00 |

| 5 | Hardik | 27 | Bhopal | 9500.00 |

| 6 | Komal | 22 | MP | 5500.00 |

+-----+-----+-----+-----+-----+

### The Package Specification

CREATE OR REPLACE PACKAGE c\_package AS

-- Adds a customer

PROCEDURE addCustomer(c\_id customers.id%type,

c\_name customerS.No.ame%type,

c\_age customers.age%type,

c\_addr customers.address%type,

c\_sal customers.salary%type);

-- Removes a customer

PROCEDURE delCustomer(c\_id customers.id%TYPE);



--Lists all customers

```
PROCEDURE listCustomer;
```

```
END c_package;
```

```
/
```

When the above code is executed at the SQL prompt, it creates the above package and displays the following result

—

Package created.

Creating the Package Body

```
CREATE OR REPLACE PACKAGE BODY c_package AS
```

```
  PROCEDURE addCustomer(c_id customers.id%type,
```

```
    c_name customerS.No.ame%type,
```

```
    c_age customers.age%type,
```

```
    c_addr customers.address%type,
```

```
    c_sal customers.salary%type)
```



```
IS
```

```
BEGIN
```

```
    INSERT INTO customers (id,name,age,address,salary)
```

```
        VALUES(c_id, c_name, c_age, c_addr, c_sal);
```

```
END addCustomer;
```

```
PROCEDURE delCustomer(c_id customers.id%type) IS
```

```
BEGIN
```

```
    DELETE FROM customers
```

```
        WHERE id = c_id;
```

```
END delCustomer;
```

```
PROCEDURE listCustomer IS
```

```
CURSOR c_customers is
```

```
    SELECT name FROM customers;
```



```
TYPE c_list IS TABLE OF customers.Name%TYPE;  
name_list c_list := c_list();  
counter INTEGER := 0;  
  
BEGIN  
  
  FOR n IN c_customers LOOP  
    counter := counter + 1;  
    name_list.extend;  
    name_list(counter) := n.name;  
    dbms_output.put_line('Customer(' || counter || ') ' || name_list(counter));  
  END LOOP;  
  
END listCustomer;  
  
END c_package;  
  
/
```

The above example makes use of the nested table. We will discuss the concept of nested table in the next chapter.



When the above code is executed at the SQL prompt, it produces the following result –

Package body created.

Using The Package

The following program uses the methods declared and defined in the package *c\_package*.

DECLARE

```
code customers.id%type:= 8;
```

BEGIN

```
c_package.addcustomer(7, 'Rajnish', 25, 'Chennai', 3500);
```

```
c_package.addcustomer(8, 'Subham', 32, 'Delhi', 7500);
```

```
c_package.listcustomer;
```

```
c_package.delcustomer(code);
```

```
c_package.listcustomer;
```

END;

/

When the above code is executed at the SQL prompt, it produces the following result –



Customer(1): Ramesh

Customer(2): Khilan

Customer(3): kaushik

Customer(4): Chaitali

Customer(5): Hardik

Customer(6): Komal

Customer(7): Rajnish

Customer(8): Subham

Customer(1): Ramesh

Customer(2): Khilan

Customer(3): kaushik

Customer(4): Chaitali

Customer(5): Hardik

Customer(6): Komal

Customer(7): Rajnish PL/SQL procedure successfully completed



## Unit 3. Transaction Management

In this chapter, we will discuss the transactions in PL/SQL. A database **transaction** is an atomic unit of work that may consist of one or more related SQL statements. It is called atomic because the database modifications brought about by the SQL statements that constitute a transaction can collectively be either committed, i.e., made permanent to the database or rolled back (undone) from the database.

A successfully executed SQL statement and a committed transaction are not same. Even if an SQL statement is executed successfully, unless the transaction containing the statement is committed, it can be rolled back and all changes made by the statement(s) can be undone.

### Starting and Ending a Transaction

A transaction has a **beginning** and an **end**. A transaction starts when one of the following events take place –

- The first SQL statement is performed after connecting to the database.
- At each new SQL statement issued after a transaction is completed.

A transaction ends when one of the following events take place –

- A **COMMIT** or a **ROLLBACK** statement is issued.



- A **DDL** statement, such as **CREATE TABLE** statement, is issued; because in that case a **COMMIT** is automatically performed.
- A **DCL** statement, such as a **GRANT** statement, is issued; because in that case a **COMMIT** is automatically performed.
- User disconnects from the database.
- User exits from **SQL\*PLUS** by issuing the **EXIT** command, a **COMMIT** is automatically performed.
- **SQL\*Plus** terminates abnormally, a **ROLLBACK** is automatically performed.
- A **DML** statement fails; in that case a **ROLLBACK** is automatically performed for undoing that **DML** statement.

### Committing a Transaction

A transaction is made permanent by issuing the **SQL** command **COMMIT**. The general syntax for the **COMMIT** command is –

**COMMIT;**

For example,

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (1, 'Ramesh', 32, 'Ahmedabad', 2000.00 );
```



```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (2, 'Khilan', 25, 'Delhi', 1500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (3, 'kaushik', 23, 'Kota', 2000.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (4, 'Chaitali', 25, 'Mumbai', 6500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (5, 'Hardik', 27, 'Bhopal', 8500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (6, 'Komal', 22, 'MP', 4500.00 );
```



COMMIT;

### Rolling Back Transactions

Changes made to the database without COMMIT could be undone using the ROLLBACK command.

The general syntax for the ROLLBACK command is –

```
ROLLBACK [TO SAVEPOINT < savepoint_name>];
```

When a transaction is aborted due to some unprecedented situation, like system failure, the entire transaction since a commit is automatically rolled back. If you are not using **savepoint**, then simply use the following statement to rollback all the changes –

```
ROLLBACK;
```

### Savepoints

Savepoints are sort of markers that help in splitting a long transaction into smaller units by setting some checkpoints. By setting savepoints within a long transaction, you can roll back to a checkpoint if required. This is done by issuing the **SAVEPOINT** command.

The general syntax for the SAVEPOINT command is –

```
SAVEPOINT < savepoint_name >;
```

For example

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```



```
VALUES (7, 'Rajnish', 27, 'HP', 9500.00 );
```

```
INSERT INTO CUSTOMERS (ID,NAME,AGE,ADDRESS,SALARY)
```

```
VALUES (8, 'Riddhi', 21, 'WB', 4500.00 );
```

```
SAVEPOINT sav1;
```

```
UPDATE CUSTOMERS
```

```
SET SALARY = SALARY + 1000;
```

```
ROLLBACK TO sav1;
```

```
UPDATE CUSTOMERS
```

```
SET SALARY = SALARY + 1000
```

```
WHERE ID = 7;
```

```
UPDATE CUSTOMERS
```

```
SET SALARY = SALARY + 1000
```



WHERE ID = 8;

COMMIT;

**ROLLBACK TO sav1** – This statement rolls back all the changes up to the point, where you had marked savepoint sav1.

After that, the new changes that you make will start.

Automatic Transaction Control

To execute a **COMMIT** automatically whenever an **INSERT, UPDATE** or **DELETE** command is executed, you can set the **AUTOCOMMIT** environment variable as –

SET AUTOCOMMIT ON;

You can turn-off the auto commit mode using the following command –

SET AUTOCOMMIT OFF;

Transaction

- The transaction is a set of logically related operation. It contains a group of tasks.
- A transaction is an action or series of actions. It is performed by a single user to perform operations for accessing the contents of the database.



**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's Account**

1. Open\_Account(X)
2. Old\_Balance = X.balance
3. New\_Balance = Old\_Balance - 800
4. X.balance = New\_Balance
5. Close\_Account(X)

### **Y's Account**

1. Open\_Account(Y)
2. Old\_Balance = Y.balance
3. New\_Balance = Old\_Balance + 800
4. Y.balance = New\_Balance
5. 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:

1. 1. R(X);
2. 2.  $X = X - 500$ ;
3. 3. 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.

Transaction property

The transaction has the four properties. These are used to maintain consistency in a database, before and after the transaction.

Property of Transaction

1. Atomicity
2. Consistency
3. Isolation
4. Durability



## Atomicity

means either all successful or none.

## Consistency

ensures bringing the database from one consistent state to another consistent state.  
ensures bringing the database from one consistent state to another consistent state.

## Isolation

ensures that transaction is isolated from other transaction.

## Durability

means once a transaction has been committed, it will remain so, even in the event of errors, power loss etc.



## Atomicity

- It states that all operations of the transaction take place at once if not, the transaction is aborted.
- There is no midway, i.e., the transaction cannot occur partially. Each transaction is treated as one unit and either run to completion or is not executed at all.

Atomicity involves the following two operations:

**Abort:** If a transaction aborts then all the changes made are not visible.

**Commit:** If a transaction commits then all the changes made are visible.

**Example:** Let's assume that following transaction T consisting of T1 and T2. A consists of Rs 600 and B consists of Rs 300. Transfer Rs 100 from account A to account B.

T1	T2
Read(A) A:= A-100 Write(A)	Read(B) Y:= Y+100 Write(B)

After completion of the transaction, A consists of Rs 500 and B consists of Rs 400.



If the transaction T fails after the completion of transaction T1 but before completion of transaction T2, then the amount will be deducted from A but not added to B. This shows the inconsistent database state. In order to ensure correctness of database state, the transaction must be executed in entirety.

### Consistency

- The integrity constraints are maintained so that the database is consistent before and after the transaction.
- The execution of a transaction will leave a database in either its prior stable state or a new stable state.
- The consistent property of database states that every transaction sees a consistent database instance.
- The transaction is used to transform the database from one consistent state to another consistent state.

**For example:** The total amount must be maintained before or after the transaction.

1. Total before T occurs =  $600+300=900$
2. Total after T occurs =  $500+400=900$

Therefore, the database is consistent. In the case when T1 is completed but T2 fails, then inconsistency will occur.

### Isolation

- It shows that the data which is used at the time of execution of a transaction cannot be used by the second transaction until the first one is completed.
- In isolation, if the transaction T1 is being executed and using the data item X, then that data item can't be accessed by any other transaction T2 until the transaction T1 ends.



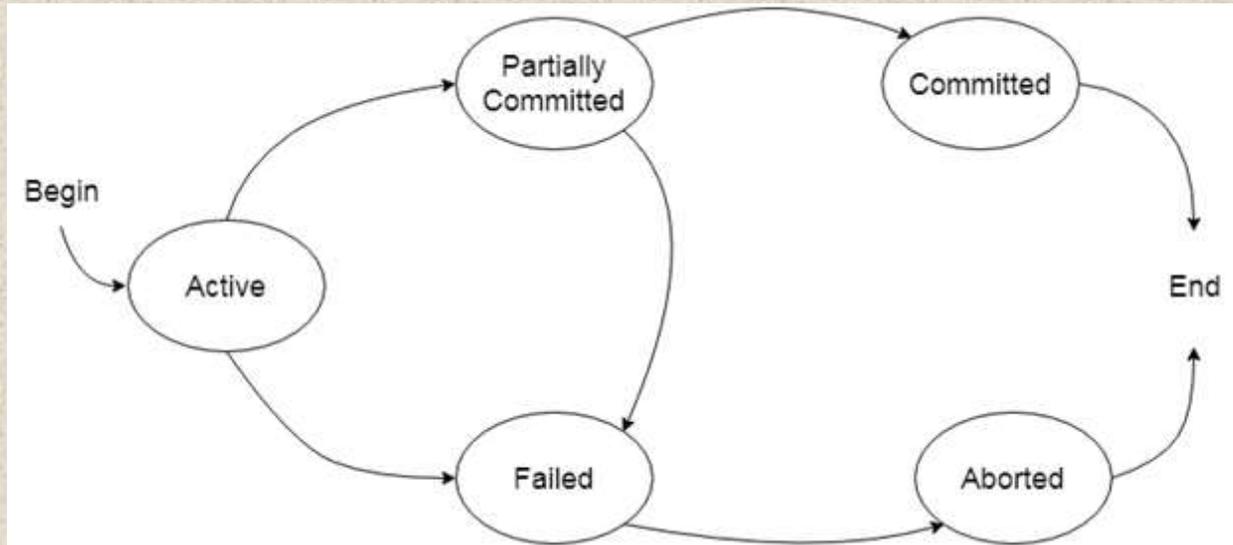
- The concurrency control subsystem of the DBMS enforced the isolation property.

### Durability

- The durability property is used to indicate the performance of the database's consistent state. It states that the transaction made the permanent changes.
- They cannot be lost by the erroneous operation of a faulty transaction or by the system failure. When a transaction is completed, then the database reaches a state known as the consistent state. That consistent state cannot be lost, even in the event of a system's failure.
- The recovery subsystem of the DBMS has the responsibility of Durability property.

### States of Transaction

In a database, the transaction can be in one of the following states -



### Active state

- The active state is the first state of every transaction. In this state, the transaction is being executed.
- For example: Insertion or deletion or updating a record is done here. But all the records are still not saved to the database.

### Partially committed

- In the partially committed state, a transaction executes its final operation, but the data is still not saved to the database.



- In the total mark calculation example, a final display of the total marks step is executed in this state.

### Committed

A transaction is said to be in a committed state if it executes all its operations successfully. In this state, all the effects are now permanently saved on the database system.

### Failed state

- If any of the checks made by the database recovery system fails, then the transaction is said to be in the failed state.
- In the example of total mark calculation, if the database is not able to fire a query to fetch the marks, then the transaction will fail to execute.

### Aborted

- If any of the checks fail and the transaction has reached a failed state then the database recovery system will make sure that the database is in its previous consistent state. If not then it will abort or roll back the transaction to bring the database into a consistent state.
- If the transaction fails in the middle of the transaction then before executing the transaction, all the executed transactions are rolled back to its consistent state.
- After aborting the transaction, the database recovery module will select one of the two operations:
  1. Re-start the transaction



## 2. Kill the transaction

### DBMS Concurrency Control

Concurrency Control is the management procedure that is required for controlling concurrent execution of the operations that take place on a database.

But before knowing about concurrency control, we should know about concurrent execution.

### Concurrent Execution in DBMS

- In a multi-user system, multiple users can access and use the same database at one time, which is known as the concurrent execution 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 should be done in an interleaved manner, and no operation should affect the other executing operations, thus maintaining the consistency of the database. Thus, on making the concurrent execution of the transaction operations, there occur several challenging problems that need to be solved.

### Problems with Concurrent Execution



In a database transaction, the two main operations are **READ** and **WRITE** operations. So, there is a need to manage these 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:

**Problem 1: Lost Update Problems (W - W Conflict)**

The problem occurs *when two different database transactions perform the read/write operations on the same database items in an interleaved manner (i.e., concurrent execution) that makes the values of the items incorrect hence making the database inconsistent.*

**For example:**

**Consider the below diagram where two transactions  $T_X$  and  $T_Y$ , 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

- At time  $t_1$ , transaction  $T_x$  reads the value of account A, i.e., \$300 (only read).
- At time  $t_2$ , transaction  $T_x$  deducts \$50 from account A that becomes \$250 (only deducted and not updated/write).
- 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 sets to inconsistent.

#### 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  $t_1$ , transaction  $T_x$  reads the value of account A, i.e., \$300.
- At time  $t_2$ , transaction  $T_x$  adds \$50 to account A that becomes \$350.
- At time  $t_3$ , transaction  $T_x$  writes the updated value in account A, i.e., \$350.
- Then at time  $t_4$ , transaction  $T_y$  reads account A that will be read as \$350.
- Then at time  $t_5$ , 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.

#### Unrepeatable Read Problem (W-R Conflict)

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

#### **For example:**

**Consider two transactions,  $T_X$  and  $T_Y$ , 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  $t_1$ , transaction  $T_x$  reads the value from account A, i.e., \$300.
- At time  $t_2$ , transaction  $T_y$  reads the value from account A, i.e., \$300.
- At time  $t_3$ , transaction  $T_y$  updates the value of account A by adding \$100 to the available balance, and then it becomes \$400.
- At time  $t_4$ , transaction  $T_y$  writes the updated value, i.e., \$400.
- After that, at time  $t_5$ , 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.

Thus, in order to maintain consistency in the database and avoid such problems that take place in concurrent execution, management is needed, and that is where the concept of Concurrency Control comes into role.

### Concurrency Control

Concurrency Control is the working concept that is required for controlling and managing the concurrent execution of database operations and thus avoiding the inconsistencies in the database. Thus, for maintaining the concurrency of the database, we have the concurrency control protocols.

### Concurrency Control Protocols

The concurrency control protocols ensure the *atomicity, consistency, isolation, durability* and *serializability* of the concurrent execution of the database transactions. Therefore, these protocols are categorized as:

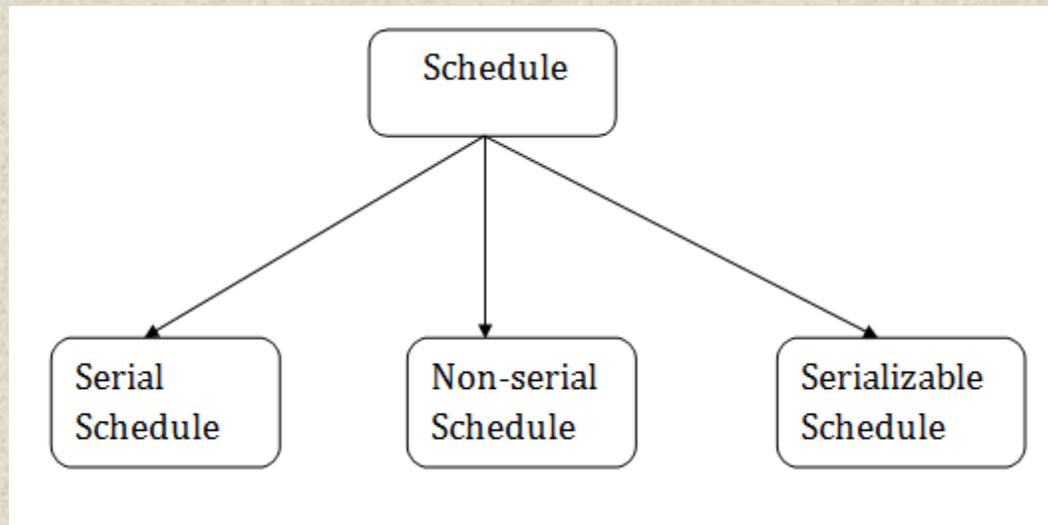
- Lock Based Concurrency Control Protocol
- Time Stamp Concurrency Control Protocol
- Validation Based Concurrency Control Protocol

We will understand and discuss each protocol one by one in our next section

Schedule



A series of operation from one transaction to another transaction is known as schedule. It is used to preserve the order of the operation in each of the individual transaction.



### 1. Serial Schedule

The serial schedule is a type of schedule where one transaction is executed completely before starting another transaction. In the serial schedule, when the first transaction completes its cycle, then the next transaction is executed.

**For example:** Suppose there are two transactions T1 and T2 which have some operations. If it has no interleaving of operations, then there are the following two possible outcomes:



1. Execute all the operations of T1 which was followed by all the operations of T2.
2. Execute all the operations of T1 which was followed by all the operations of T2.
  - In the given (a) figure, Schedule A shows the serial schedule where T1 followed by T2.
  - In the given (b) figure, Schedule B shows the serial schedule where T2 followed by T1.

## 2. Non-serial Schedule

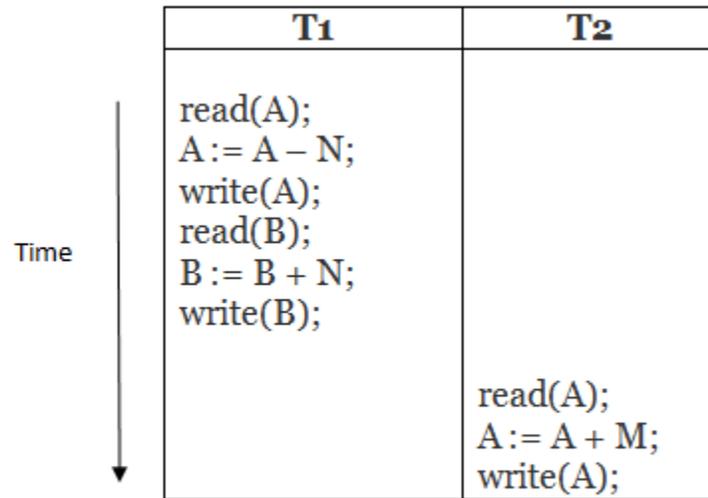
- If interleaving of operations is allowed, then there will be non-serial schedule.
- It contains many possible orders in which the system can execute the individual operations of the transactions.
- In the given figure (c) and (d), Schedule C and Schedule D are the non-serial schedules. It has interleaving of operations.

## 3. Serializable schedule

- The serializability of schedules is used to find non-serial schedules that allow the transaction to execute concurrently without interfering with one another.
- It identifies which schedules are correct when executions of the transaction have interleaving of their operations.
- A non-serial schedule will be serializable if its result is equal to the result of its transactions executed serially.



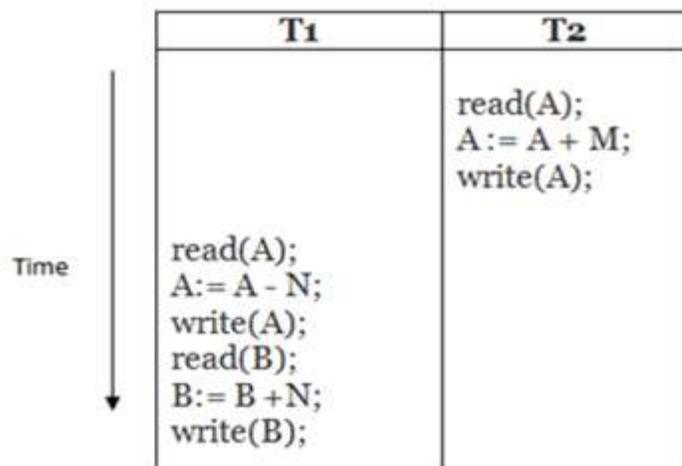
(a)



**Schedule A**



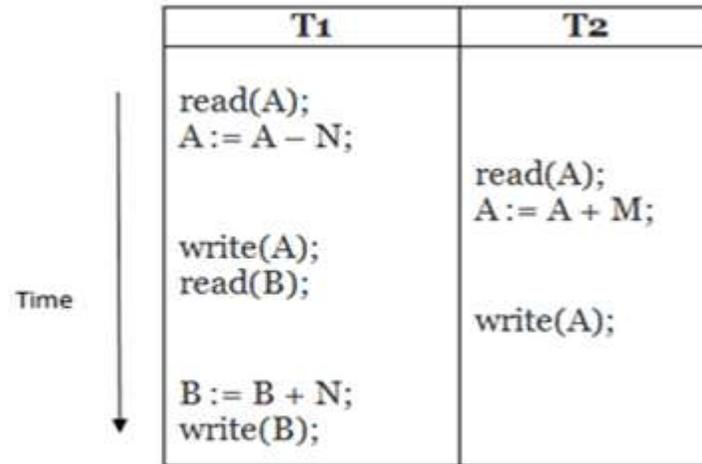
(b)



**Schedule B**



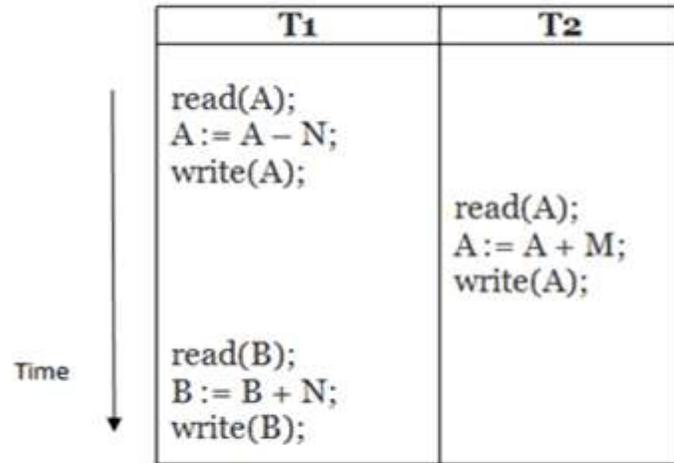
(c)



Schedule C



(d)



**Schedule D**

**Here,**

Schedule A and Schedule B are serial schedule.

Schedule C and Schedule D are Non-serial schedule.

Testing of Serializability

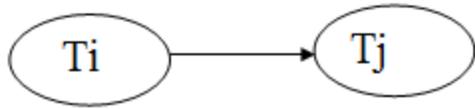
Serialization Graph is used to test the Serializability of a schedule.



Assume a schedule  $S$ . For  $S$ , we construct a graph known as precedence graph. This graph has a pair  $G = (V, E)$ , where  $V$  consists a set of vertices, and  $E$  consists a set of edges. The set of vertices is used to contain all the transactions participating in the schedule. The set of edges is used to contain all edges  $T_i \rightarrow T_j$  for which one of the three conditions holds:

1. Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes write ( $Q$ ) before  $T_j$  executes read ( $Q$ ).
2. Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes read ( $Q$ ) before  $T_j$  executes write ( $Q$ ).
3. Create a node  $T_i \rightarrow T_j$  if  $T_i$  executes write ( $Q$ ) before  $T_j$  executes write ( $Q$ ).

#### Precedence graph for Schedule $S$



- If a precedence graph contains a single edge  $T_i \rightarrow T_j$ , then all the instructions of  $T_i$  are executed before the first instruction of  $T_j$  is executed.
- If a precedence graph for schedule  $S$  contains a cycle, then  $S$  is non-serializable. If the precedence graph has no cycle, then  $S$  is known as serializable.

**For example:**



	T1	T2	T3
Time ↓	Read(A)	Read(B)	
	A := f <sub>1</sub> (A)		Read(C)
		B := f <sub>2</sub> (B) Write(B)	C := f <sub>3</sub> (C) Write(C)
	Write(A)		Read(B)
	Read(C)	Read(A) A := f <sub>4</sub> (A)	
	C := f <sub>5</sub> (C) Write(C)	Write(A)	B := f <sub>6</sub> (B) Write(B)

**Schedule S1**

**Explanation:**



**Read(A):** In T1, no subsequent writes to A, so no new edges

**Read(B):** In T2, no subsequent writes to B, so no new edges

**Read(C):** In T3, no subsequent writes to C, so no new edges

**Write(B):** B is subsequently read by T3, so add edge  $T2 \rightarrow T3$

**Write(C):** C is subsequently read by T1, so add edge  $T3 \rightarrow T1$

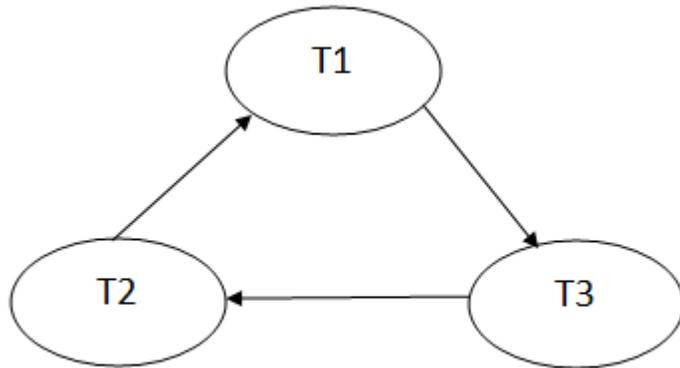
**Write(A):** A is subsequently read by T2, so add edge  $T1 \rightarrow T2$

**Write(A):** In T2, no subsequent reads to A, so no new edges

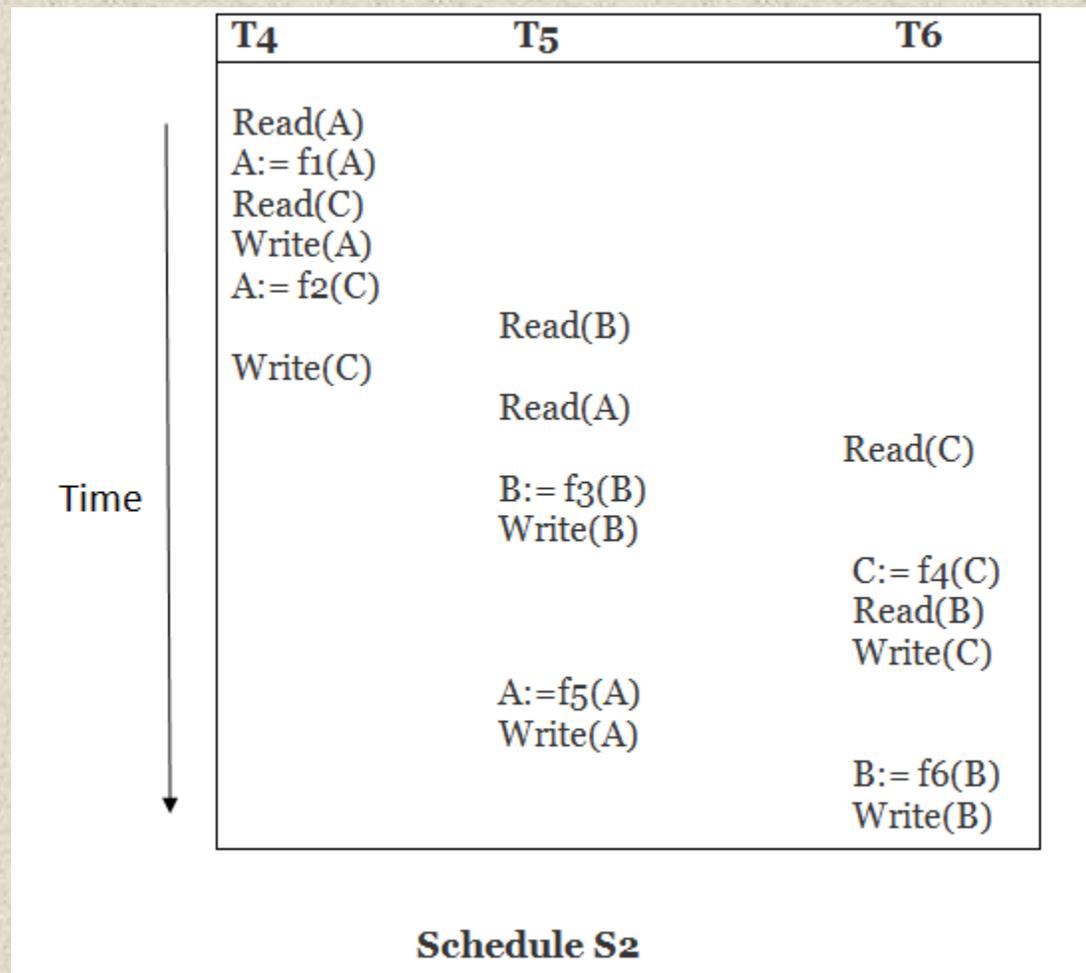
**Write(C):** In T1, no subsequent reads to C, so no new edges

**Write(B):** In T3, no subsequent reads to B, so no new edges

Precedence graph for schedule S1:



The precedence graph for schedule S1 contains a cycle that's why Schedule S1 is non-serializable.



**Explanation:**



**Read(A):** In T4, no subsequent writes to A, so no new edges

**Read(C):** In T4, no subsequent writes to C, so no new edges

**Write(A):** A is subsequently read by T5, so add edge  $T4 \rightarrow T5$

**Read(B):** In T5, no subsequent writes to B, so no new edges

**Write(C):** C is subsequently read by T6, so add edge  $T4 \rightarrow T6$

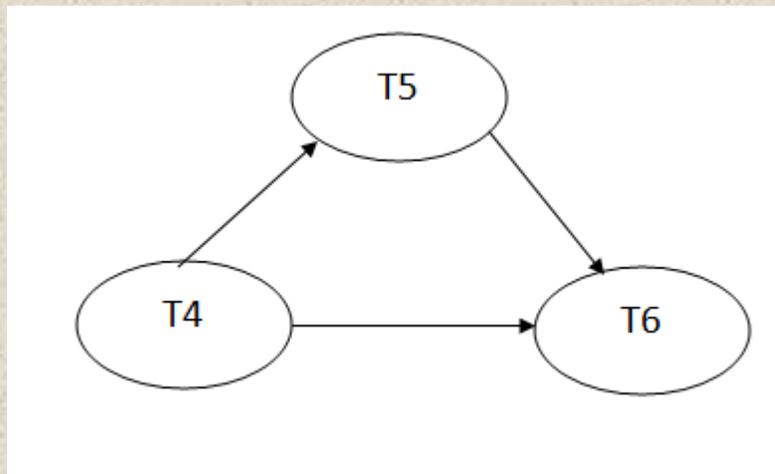
**Write(B):** A is subsequently read by T6, so add edge  $T5 \rightarrow T6$

**Write(C):** In T6, no subsequent reads to C, so no new edges

**Write(A):** In T5, no subsequent reads to A, so no new edges

**Write(B):** In T6, no subsequent reads to B, so no new edges

Precedence graph for schedule S2:





The precedence graph for schedule S2 contains no cycle that's why ScheduleS2 is serializable.

### Conflict Serializable Schedule

- A schedule is called conflict serializability if after swapping of non-conflicting operations, it can transform into a serial schedule.
- The schedule will be a conflict serializable if it is conflict equivalent to a serial schedule.

### Conflicting Operations

The two operations become conflicting if all conditions satisfy:

1. Both belong to separate transactions.
2. They have the same data item.
3. They contain at least one write operation.

Example:

Swapping is possible only if S1 and S2 are logically equal.



**1. T1: Read(A) T2: Read(A)**

T1	T2
Read(A)	
	Read(A)

Swapped



T1	T2
	Read(A)
Read(A)	

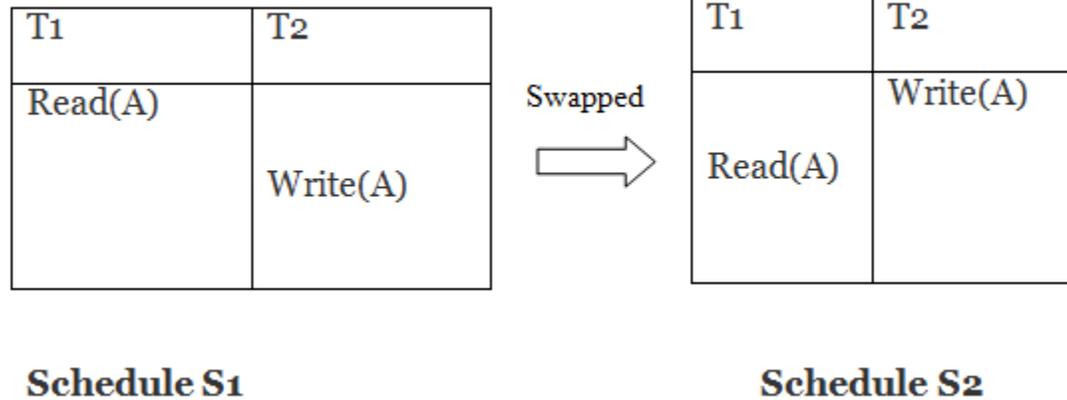
**Schedule S1**

**Schedule S2**

Here,  $S1 = S2$ . That means it is non-conflict.



## 2. T1: Read(A) T2: Write(A)



Here,  $S1 \neq S2$ . That means it is conflict.

### Conflict Equivalent

In the conflict equivalent, one can be transformed to another by swapping non-conflicting operations. In the given example, S2 is conflict equivalent to S1 (S1 can be converted to S2 by swapping non-conflicting operations).

Two schedules are said to be conflict equivalent if and only if:

1. They contain the same set of the transaction.
2. If each pair of conflict operations are ordered in the same way.



Example:

### Non-serial schedule

T1	T2
Read(A) Write(A)	
	Read(A) Write(A)
Read(B) Write(B)	
	Read(B) Write(B)

**Schedule S1**

### Serial Schedule

T1	T2
Read(A) Write(A) Read(B) Write(B)	
	Read(A) Write(A) Read(B) Write(B)

**Schedule S2**

Schedule S2 is a serial schedule because, in this, all operations of T1 are performed before starting any operation of T2. Schedule S1 can be transformed into a serial schedule by swapping non-conflicting operations of S1.

**After swapping of non-conflict operations, the schedule S1 becomes:**



**T1**

**T2**

Read(A)  
Write(A)  
Read(B)  
Write(B)

Read(A)  
Write(A)  
Read(B)  
Write(B)

Since, S1 is conflict serializable.

### View Serializability

- A schedule will view serializable if it is view equivalent to a serial schedule.
- If a schedule is conflict serializable, then it will be view serializable.
- The view serializable which does not conflict serializable contains blind writes.

### View Equivalent

Two schedules S1 and S2 are said to be view equivalent if they satisfy the following conditions:



## 1. Initial Read

An initial read of both schedules must be the same. Suppose two schedule S1 and S2. In schedule S1, if a transaction T1 is reading the data item A, then in S2, transaction T1 should also read A.

T1	T2
Read(A)	Write(A)

**Schedule S1**

T1	T2
Read(A)	Write(A)

**Schedule S2**

Above two schedules are view equivalent because Initial read operation in S1 is done by T1 and in S2 it is also done by T1.

## 2. Updated Read

In schedule S1, if  $T_i$  is reading A which is updated by  $T_j$  then in S2 also,  $T_i$  should read A which is updated by  $T_j$ .



T1	T2	T3
Write(A)	Write(A)	Read(A)

**Schedule S1**

T1	T2	T3
Write(A)	Write(A)	<u>Read(A)</u>

**Schedule S2**

Above two schedules are not view equal because, in S1, T3 is reading A updated by T2 and in S2, T3 is reading A updated by T1.

### 3. Final Write

A final write must be the same between both the schedules. In schedule S1, if a transaction T1 updates A at last then in S2, final writes operations should also be done by T1.



T1	T2	T3
Write(A)	Read(A)	
		Write(A)

**Schedule S1**

T1	T2	T3
Write(A)	Read(A)	
		Write(A)

**Schedule S2**

Above two schedules is view equal because Final write operation in S1 is done by T3 and in S2, the final write operation is also done by T3.

**Example:**

T1	T2	T3
Read(A)		
Write(A)	Write(A)	
		Write(A)



## Schedule S

With 3 transactions, the total number of possible schedule

1.  $= 3! = 6$
2.  $S1 = \langle T1 T2 T3 \rangle$
3.  $S2 = \langle T1 T3 T2 \rangle$
4.  $S3 = \langle T2 T3 T1 \rangle$
5.  $S4 = \langle T2 T1 T3 \rangle$
6.  $S5 = \langle T3 T1 T2 \rangle$
7.  $S6 = \langle T3 T2 T1 \rangle$

**Taking first schedule S1:**

<b>T1</b>	<b>T2</b>	<b>T3</b>
Read(A) Write(A)	Write(A)	Write(A)



## **Schedule S1**

**Step 1:** final updation on data items

In both schedules S and S1, there is no read except the initial read that's why we don't need to check that condition.

**Step 2:** Initial Read

The initial read operation in S is done by T1 and in S1, it is also done by T1.

**Step 3:** Final Write

The final write operation in S is done by T3 and in S1, it is also done by T3. So, S and S1 are view Equivalent.

The first schedule S1 satisfies all three conditions, so we don't need to check another schedule.

**Hence, view equivalent serial schedule is:**

1.  $T1 \rightarrow T2 \rightarrow T3$



## Unit 4 Concurrency Control & Recovery System

### **RDBMS Concurrency Control**

Concurrency Control is the management procedure that is required for controlling concurrent execution of the operations that take place on a database.

But before knowing about concurrency control, we should know about concurrent execution.

### Concurrent Execution in RDBMS

- In a multi-user system, multiple users can access and use the same database at one time, which is known as the concurrent execution 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 should be done in an interleaved manner, and no operation should affect the other executing operations, thus maintaining the consistency of the database. Thus, on making the concurrent execution of the transaction operations, there occur several challenging problems that need to be solved.

### Problems with Concurrent Execution

In a database transaction, the two main operations are **READ** and **WRITE** operations. So, there is a need to manage these 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:

#### Problem 1: Lost Update Problems (W - W Conflict)

The problem occurs *when two different database transactions perform the read/write operations on the same database items in an interleaved manner (i.e., concurrent execution) that makes the values of the items incorrect hence making the database inconsistent.*

#### **For example:**

**Consider the below diagram where two transactions  $T_X$  and  $T_Y$ , 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

- At time  $t_1$ , transaction  $T_x$  reads the value of account A, i.e., \$300 (only read).
- At time  $t_2$ , transaction  $T_x$  deducts \$50 from account A that becomes \$250 (only deducted and not updated/write).
- 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 sets to inconsistent.

#### 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  $t_1$ , transaction  $T_x$  reads the value of account A, i.e., \$300.
- At time  $t_2$ , transaction  $T_x$  adds \$50 to account A that becomes \$350.
- At time  $t_3$ , transaction  $T_x$  writes the updated value in account A, i.e., \$350.
- Then at time  $t_4$ , transaction  $T_y$  reads account A that will be read as \$350.
- Then at time  $t_5$ , 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.

#### Unrepeatable Read Problem (W-R Conflict)

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

#### **For example:**

**Consider two transactions,  $T_X$  and  $T_Y$ , 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  $t_1$ , transaction  $T_x$  reads the value from account A, i.e., \$300.
- At time  $t_2$ , transaction  $T_y$  reads the value from account A, i.e., \$300.
- At time  $t_3$ , transaction  $T_y$  updates the value of account A by adding \$100 to the available balance, and then it becomes \$400.
- At time  $t_4$ , transaction  $T_y$  writes the updated value, i.e., \$400.
- After that, at time  $t_5$ , 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.

Thus, in order to maintain consistency in the database and avoid such problems that take place in concurrent execution, management is needed, and that is where the concept of Concurrency Control comes into role.

### Concurrency Control

Concurrency Control is the working concept that is required for controlling and managing the concurrent execution of database operations and thus avoiding the inconsistencies in the database. Thus, for maintaining the concurrency of the database, we have the concurrency control protocols.

### Concurrency Control Protocols

The concurrency control protocols ensure the *atomicity, consistency, isolation, durability* and *serializability* of the concurrent execution of the database transactions. Therefore, these protocols are categorized as:

- Lock Based Concurrency Control Protocol
- Time Stamp Concurrency Control Protocol
- Validation Based Concurrency Control Protocol

We will understand and discuss each protocol one by one in our next sections.

### Lock-Based Protocol



In this type of protocol, any transaction cannot read or write data until it acquires an appropriate lock on it. There are two types of lock:

### **1. Shared lock:**

- It is also known as a Read-only lock. In a shared lock, the data item can only read by the transaction.
- It can be shared between the transactions because when the transaction holds a lock, then it can't update the data on the data item.

### **2. Exclusive lock:**

- In the exclusive lock, the data item can be both reads as well as written by the transaction.
- This lock is exclusive, and in this lock, multiple transactions do not modify the same data simultaneously.

There are four types of lock protocols available:

#### **1. Simplistic lock protocol**

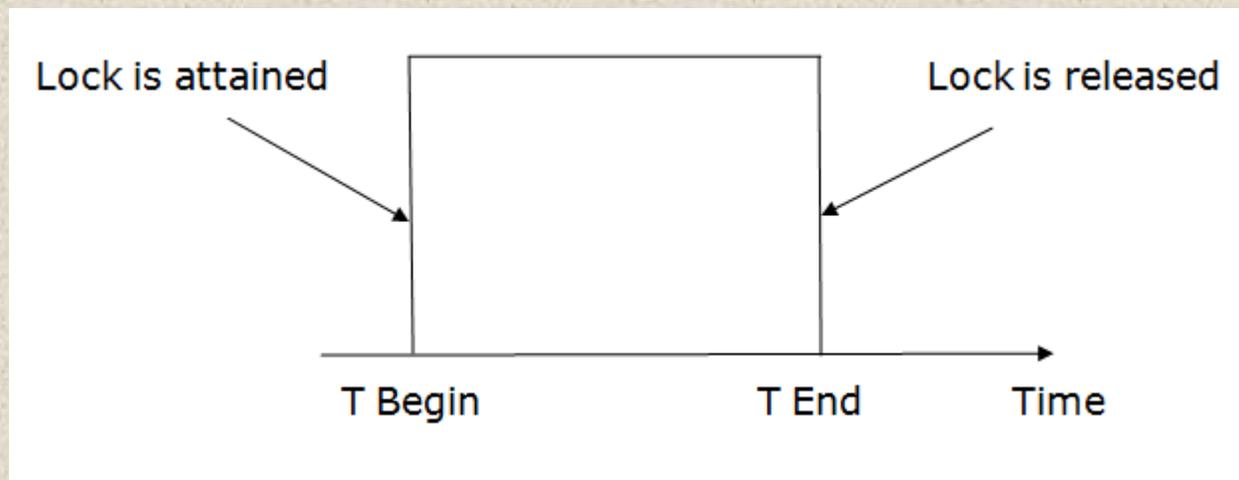
It is the simplest way of locking the data while transaction. Simplistic lock-based protocols allow all the transactions to get the lock on the data before insert or delete or update on it. It will unlock the data item after completing the transaction.

#### **2. Pre-claiming Lock Protocol**

- Pre-claiming Lock Protocols evaluate the transaction to list all the data items on which they need locks.
- Before initiating an execution of the transaction, it requests DBMS for all the lock on all those data items.



- If all the locks are granted then this protocol allows the transaction to begin. When the transaction is completed then it releases all the lock.
- If all the locks are not granted then this protocol allows the transaction to rolls back and waits until all the locks are granted.

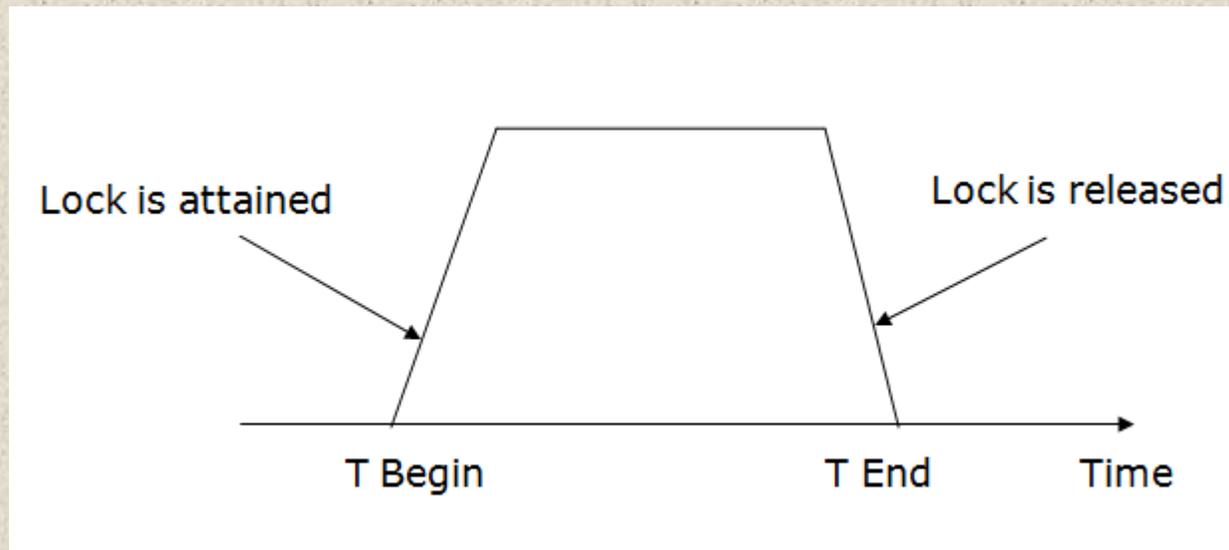


### 3. Two-phase locking (2PL)

- The two-phase locking protocol divides the execution phase of the transaction into three parts.
- In the first part, when the execution of the transaction starts, it seeks permission for the lock it requires.



- In the second part, the transaction acquires all the locks. The third phase is started as soon as the transaction releases its first lock.
- In the third phase, the transaction cannot demand any new locks. It only releases the acquired locks.



There are two phases of 2PL:

**Growing phase:** In the growing phase, a new lock on the data item may be acquired by the transaction, but none can be released.



**Shrinking phase:** In the shrinking phase, existing lock held by the transaction may be released, but no new locks can be acquired.

In the below example, if lock conversion is allowed then the following phase can happen:

1. Upgrading of lock (from S(a) to X (a)) is allowed in growing phase.
2. Downgrading of lock (from X(a) to S(a)) must be done in shrinking phase.

**Example:**



	<b>T1</b>	<b>T2</b>
0	LOCK-S(A)	
1		LOCK-S(A)
2	LOCK-X(B)	
3	—	—
4	UNLOCK(A)	
5		LOCK-X(C)
6	UNLOCK(B)	
7		UNLOCK(A)
8		UNLOCK(C)
9	—	—

The following way shows how unlocking and locking work with 2-PL.



### **Transaction T1:**

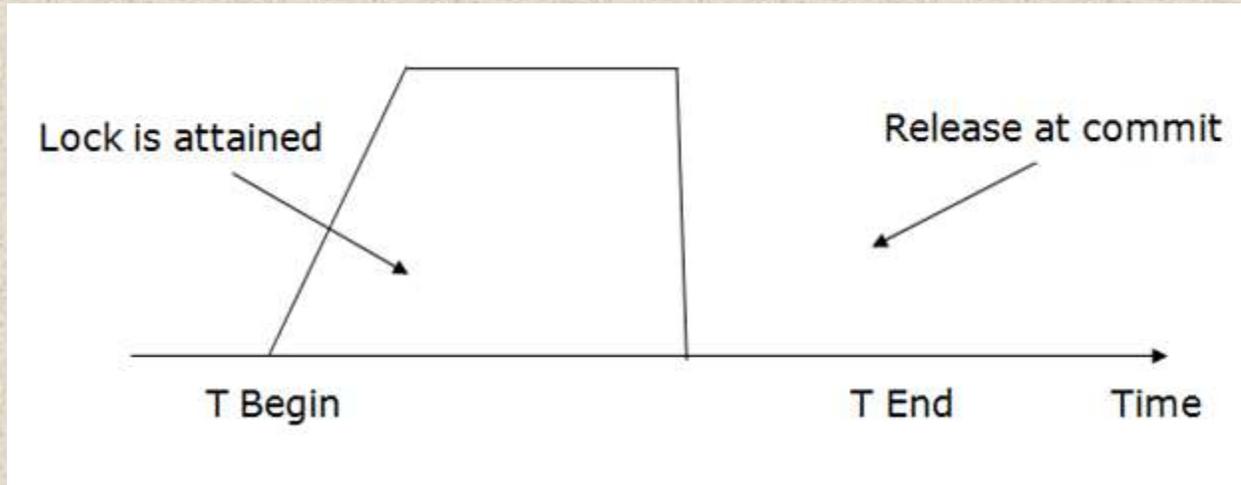
- **Growing phase:** from step 1-3
- **Shrinking phase:** from step 5-7
- **Lock point:** at 3

### **Transaction T2:**

- **Growing phase:** from step 2-6
- **Shrinking phase:** from step 8-9
- **Lock point:** at 6

### 4. Strict Two-phase locking (Strict-2PL)

- The first phase of Strict-2PL is similar to 2PL. In the first phase, after acquiring all the locks, the transaction continues to execute normally.
- The only difference between 2PL and strict 2PL is that Strict-2PL does not release a lock after using it.
- Strict-2PL waits until the whole transaction to commit, and then it releases all the locks at a time.
- Strict-2PL protocol does not have shrinking phase of lock release.



It does not have cascading abort as 2PL does.

### Timestamp Ordering Protocol

- The Timestamp Ordering Protocol is used to order the transactions based on their Timestamps. The order of transaction is nothing but the ascending order of the transaction creation.
- The priority of the older transaction is higher that's why it executes first. To determine the timestamp of the transaction, this protocol uses system time or logical counter.
- The lock-based protocol is used to manage the order between conflicting pairs among transactions at the execution time. But Timestamp based protocols start working as soon as a transaction is created.



- Let's assume there are two transactions T1 and T2. Suppose the transaction T1 has entered the system at 007 times and transaction T2 has entered the system at 009 times. T1 has the higher priority, so it executes first as it is entered the system first.
- The timestamp ordering protocol also maintains the timestamp of last 'read' and 'write' operation on a data.

**Basic Timestamp ordering protocol works as follows:**

1. Check the following condition whenever a transaction  $T_i$  issues a **Read (X)** operation:

- If  $W\_TS(X) > TS(T_i)$  then the operation is rejected.
- If  $W\_TS(X) \leq TS(T_i)$  then the operation is executed.
- Timestamps of all the data items are updated.

2. Check the following condition whenever a transaction  $T_i$  issues a **Write(X)** operation:

- If  $TS(T_i) < R\_TS(X)$  then the operation is rejected.
- If  $TS(T_i) < W\_TS(X)$  then the operation is rejected and  $T_i$  is rolled back otherwise the operation is executed.

**Where,**

**TS(TI)** denotes the timestamp of the transaction  $T_i$ .

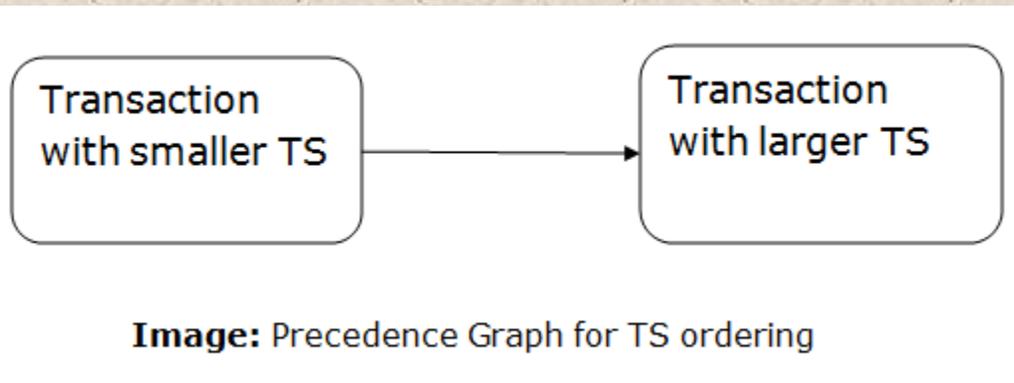
**R\_TS(X)** denotes the Read time-stamp of data-item X.

**W\_TS(X)** denotes the Write time-stamp of data-item X.



Advantages and Disadvantages of TO protocol:

- TO protocol ensures serializability since the precedence graph is as follows:



- TS protocol ensures freedom from deadlock that means no transaction ever waits.
- But the schedule may not be recoverable and may not even be cascade- free.

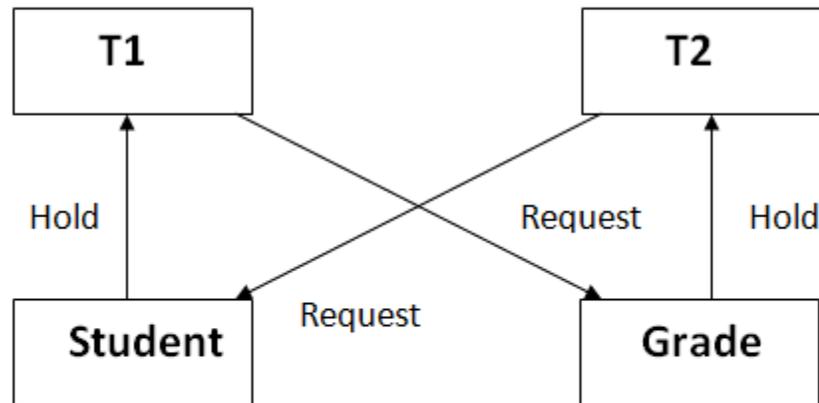
Deadlock in DBMS

A deadlock is a condition where two or more transactions are waiting indefinitely for one another to give up locks. Deadlock is said to be one of the most feared complications in DBMS as no task ever gets finished and is in waiting state forever.



**For example:** In the student table, transaction T1 holds a lock on some rows and needs to update some rows in the grade table. Simultaneously, transaction T2 holds locks on some rows in the grade table and needs to update the rows in the Student table held by Transaction T1.

Now, the main problem arises. Now Transaction T1 is waiting for T2 to release its lock and similarly, transaction T2 is waiting for T1 to release its lock. All activities come to a halt state and remain at a standstill. It will remain in a standstill until the DBMS detects the deadlock and aborts one of the transactions.



**Figure:** Deadlock in DBMS

Deadlock Avoidance



- When a database is stuck in a deadlock state, then it is better to avoid the database rather than aborting or restating the database. This is a waste of time and resource.
- Deadlock avoidance mechanism is used to detect any deadlock situation in advance. A method like "wait for graph" is used for detecting the deadlock situation but this method is suitable only for the smaller database. For the larger database, deadlock prevention method can be used.

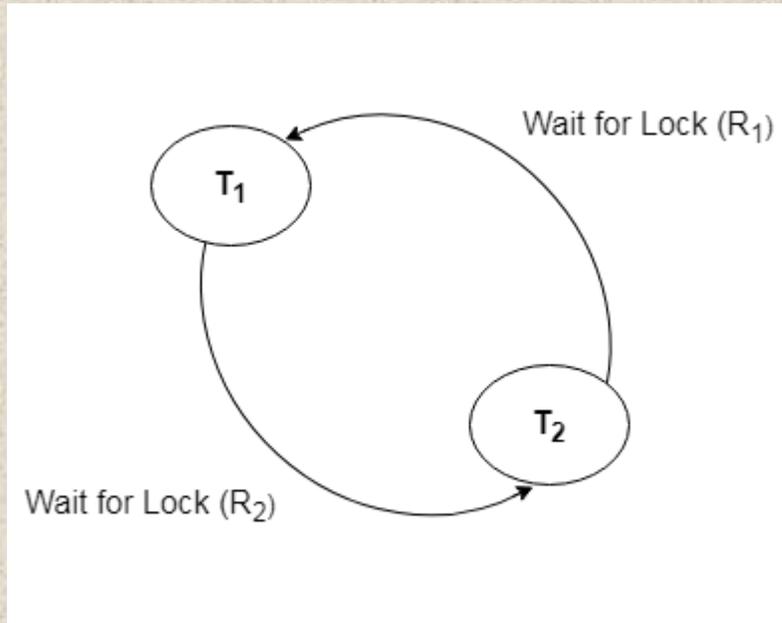
### Deadlock Detection

In a database, when a transaction waits indefinitely to obtain a lock, then the DBMS should detect whether the transaction is involved in a deadlock or not. The lock manager maintains a Wait for the graph to detect the deadlock cycle in the database.

### Wait for Graph

- This is the suitable method for deadlock detection. In this method, a graph is created based on the transaction and their lock. If the created graph has a cycle or closed loop, then there is a deadlock.
- The wait for the graph is maintained by the system for every transaction which is waiting for some data held by the others. The system keeps checking the graph if there is any cycle in the graph.

The wait for a graph for the above scenario is shown below:



### Deadlock Prevention

- Deadlock prevention method is suitable for a large database. If the resources are allocated in such a way that deadlock never occurs, then the deadlock can be prevented.
- The Database management system analyzes the operations of the transaction whether they can create a deadlock situation or not. If they do, then the DBMS never allowed that transaction to be executed.

### Wait-Die scheme



In this scheme, if a transaction requests for a resource which is already held with a conflicting lock by another transaction then the DBMS simply checks the timestamp of both transactions. It allows the older transaction to wait until the resource is available for execution.

Let's assume there are two transactions  $T_i$  and  $T_j$  and let  $TS(T)$  is a timestamp of any transaction  $T$ . If  $T_2$  holds a lock by some other transaction and  $T_1$  is requesting for resources held by  $T_2$  then the following actions are performed by DBMS:

1. Check if  $TS(T_i) < TS(T_j)$  - If  $T_i$  is the older transaction and  $T_j$  has held some resource, then  $T_i$  is allowed to wait until the data-item is available for execution. That means if the older transaction is waiting for a resource which is locked by the younger transaction, then the older transaction is allowed to wait for resource until it is available.
2. Check if  $TS(T_j) < TS(T_i)$  - If  $T_i$  is older transaction and has held some resource and if  $T_j$  is waiting for it, then  $T_j$  is killed and restarted later with the random delay but with the same timestamp.

#### Wound wait scheme

- In wound wait scheme, if the older transaction requests for a resource which is held by the younger transaction, then older transaction forces younger one to kill the transaction and release the resource. After the minute delay, the younger transaction is restarted but with the same timestamp.
- If the older transaction has held a resource which is requested by the Younger transaction, then the younger transaction is asked to wait until older releases it.

#### 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.

Log-based Recovery



Log is a sequence of records, which maintains the records of actions performed by a transaction. It is important that the logs are written prior to the actual modification and stored on a stable storage media, which is failsafe.

Log-based recovery works as follows –

- The log file is kept on a stable storage media.
- When a transaction enters the system and starts execution, it writes a log about it.

$\langle T_n, \text{Start} \rangle$

- When the transaction modifies an item  $X$ , it write logs as follows –

$\langle T_n, X, V_1, V_2 \rangle$

It reads  $T_n$  has changed the value of  $X$ , from  $V_1$  to  $V_2$ .

- When the transaction finishes, it logs –

$\langle T_n, \text{commit} \rangle$

The database can be modified using two approaches –

- **Deferred database modification** – All logs are written on to the stable storage and the database is updated when a transaction commits.
- **Immediate database modification** – Each log follows an actual database modification. That is, the database is modified immediately after every operation.



## Recovery with Concurrent Transactions

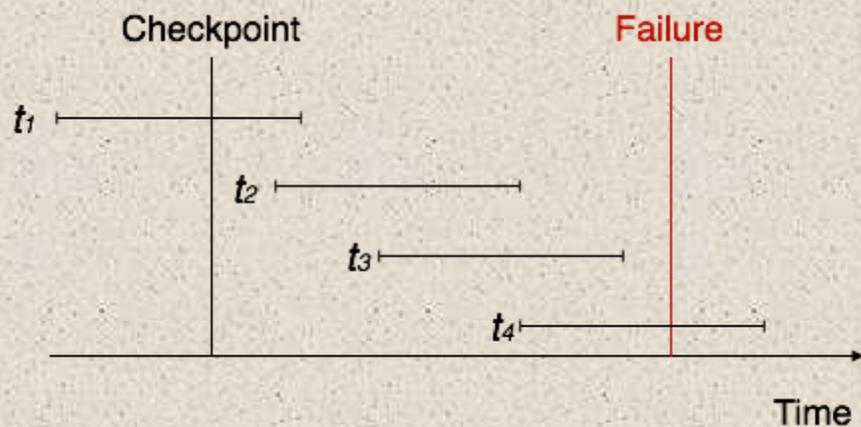
When more than one transaction are being executed in parallel, the logs are interleaved. At the time of recovery, it would become hard for the recovery system to backtrack all logs, and then start recovering. To ease this situation, most modern DBMS use the concept of 'checkpoints'.

### Checkpoint

Keeping and maintaining logs in real time and in real environment may fill out all the memory space available in the system. As time passes, the log file may grow too big to be handled at all. Checkpoint is a mechanism where all the previous logs are removed from the system and stored permanently in a storage disk. Checkpoint declares a point before which the DBMS was in consistent state, and all the transactions were committed.

### Recovery

When a system with concurrent transactions crashes and recovers, it behaves in the following manner –





- The recovery system reads the logs backwards from the end to the last checkpoint.
- It maintains two lists, an undo-list and a redo-list.
- If the recovery system sees a log with  $\langle T_n, \text{Start} \rangle$  and  $\langle T_n, \text{Commit} \rangle$  or just  $\langle T_n, \text{Commit} \rangle$ , it puts the transaction in the redo-list.
- If the recovery system sees a log with  $\langle T_n, \text{Start} \rangle$  but no commit or abort log found, it puts the transaction in undo-list.

All the transactions in the undo-list are then undone and their logs are removed. All the transactions in the redo-list and their previous logs are removed and then redone before saving their logs.

### Recovery with Concurrent Transaction

- Whenever more than one transaction is being executed, then the interleaved of logs occur. During recovery, it would become difficult for the recovery system to backtrack all logs and then start recovering.
- To ease this situation, 'checkpoint' concept is used by most DBMS.

As we have discussed [checkpoint](#) in Transaction Processing Concept of this tutorial, so you can go through the concepts again to make things more clear



### **References:**

- **Database Management System Bipin Desai Galgotia Publications New Delhi**
- **SQL/PLSQL the programming language of oracle Ivan Bayross**
- **Database System Concepts 5th Edition Silberschatz, Korth, Sudershan McGraw-Hill**
- <https://www.javatpoint.com/>
- <https://www.tutorialspoint.com/>