

GANAPATI INSTITUTE OF ENGINEERING AND TECHNOLOGY

(POLYTECHNIC)

Mathasahi , Jagatpur , Cuttack - 754200, Odisha

"DATA BASE MANAGEMENT SYSTEM"

Prepared by: Mr Ananta Bikram Basudev Mohapatra.

CONTENTS

LECTURE-1:IntroductiontoData **LECTURE-2: DBMS** LECTURE-3:3levelArchitectureofDBMS **LECTURE-4: Elements of DBMS** LECTURE-5: ER-MODEL LECTURE-6: ER-DIAGRAM: LECTURE-7:AdvancedER-Diagram: LECTURE-8:ConversionofER-DiagramtoRelationalDatabase LECTURE-9: Record Based Logical Model LECTURE-10:RELATIONALMODEL LECTURE-11: CONSTRAINTS LECTURE-12:FILEORGANISATION LECTURE-13: INDEX **LECTURE-14: Clustering Index** LECTURE-15: B+ Tree Index LECTURE-16: Hash File Organization **LECTURE-17:** Query Processing LECTURE-18: Evaluation of Expressions LECTURE-19Relational Algebra **LECTURE-20** Additional Operations **LECTURE-21Tuple Relational Calculus** LECTURE-22StructuredQueryLanguage(SQL) **LECTURE-23** Nested Sub queries **LECTURE-24Integrity Constraints** LECTURE-25 Query by Example (QBE) LECTURE-26RelationalDatabaseDesign LECTURE-27ClosureofasetofFunctionalDependencies **LECTURE-28 Loss less Decomposition LECTURE-29** Normalization LECTURE-30Boyce-CodeNormalForm(BCNF) **LECTURE-31Query Processing** LECTURE-32QueryOptimization **LECTURE-33Transaction** LECTURE-34Problemsduetolocking LECTURE-35MultiversionTechniqueBasedonTimestampOrdering LECTURE-36Serializability LECTURE-37ObjectOrientedDatabases **LECTURE-38Parallel Database**

Module-1:

LECTURE-1:IntroductiontoData

Introduction:

In computerized information system data are the basic resource of the organization. So, proper organization and management for data is required for organization to run smoothly. Database management system deals the knowledge of how data stored and managed on a computerized information system. In any organization, it requires accurate and reliable data for better decision making, ensuring privacy of data and controlling data efficiently.

The examples include deposit and/or withdrawal from a bank, hotel, airline or railway reservation, purchase items from supermarkets in all cases, a database is accessed.

Whatis data?

Data are the known facts or figures that have implicit meaning. It can also be defined as it is the representation of facts, concepts or instructions in a formal manner, which is suitable for understanding and processing. Data can be represented in alphabets (A-Z, a-z), digits (0-9) and using special characters (+,-.#,\$, etc)

e.g:25,"ajit"etc.

Information:

Informationistheprocesseddataonwhichdecisionsandactionsarebased.Informationcanbe defined as the organized and classified data to provide meaningful values. Eg: "TheageofRaviis25"

File:

Fileisa collectionofrelateddata storedinsecondarymemory.

FileOrientedApproach:

The traditional file oriented approach to information processing each application has a separate master file and its own set of personal file. In file oriented approach the program dependent on the files and files dependent upon the programs.

Disadvantagesoffileorientedapproach:

1) Dataredundancyandinconsistency:

The same information may be written in several files. This redundancy leads to higher storage and access cost. It may lead data inconsistencythat is the various copies of the same data may present at multiple places for example a changed customer address may be reflected in single file but not else where in the system.

2) Difficultyinaccessingdata:

The conventional file processing system do not allow data to be retrieved in a convenient and efficient manner according to user choice.

3) Dataisolation:

Because data are scattered in various files and files may be in different formats withnew application programs to retrieve the appropriate data is difficult.

4) IntegrityProblems:

Developers enforce data validation in the system by adding appropriate code in the various application program. How ever when new constraints are added, it is difficult to change the programs to enforce them.

5) Atomicity:

Itisdifficult to ensureatomicityina fileprocessingsystemwhentransactionfailureoccurs due to power failure, networking problems etc. (atomicity: either all operations of the transaction are reflected properly in the database or non are)

6) Concurrentaccess:

In the file processing systemit is not possible to access the same file for transaction at same the time.

7) Securityproblems:

There is no security provided in file processing system to securethedatafrom unauthoriz- ed user access.

LECTURE-2: DBMS

Database:

Adatabase is organized collection of related dataof an organization stored in formatted way which is shared by multiple users.

Themainfeatureofdatainadatabaseare:

- 1. Itmustbewellorganized
- 2. Itisrelated
- 3. Itisaccessibleinalogicalorderwithoutanydifficulty
- 4. Itisstoredonlyonce

For example consider the roll no, name, address of a student stored in a student file. It is collection of related data with an implicit meaning. Data in the database may be persistent, integrated and shared.

Persistent:

Ifdataisremovedfromdatabasedueto someexplicitrequest fromuserto remove.

Integrated:

Adatabase can be a collection of data from differentfiles and when any redundancy among those files are removed from database is said to be integrated data.

SharingData:

Thedatastoredinthedatabasecanbeshared bymultipleusers simultaneouslywithout affectingthe correctness of data.

WhyDatabase:

Inordertoovercomethe limitationofa file system, anewapproachwasrequired. Henceadatabase approachemerged. Adatabase isapersistent collectionoflogicallyrelateddata. The initial attempts were to provide a centralized collection of data. Adatabase has a self describing nature. It contains not only the data sharing and integration of data of an organization in a single database.

Asmall database can be handledmanually butfor a large database and havingmultiple users it is difficult to maintain it. In that case a computerized database is useful.

Theadvantagesofdatabasesystemovertraditional, paper basedmethodsofrecordkeeping are:

- Compactness:Noneed forlargeamountofpaper files
- Speed: The machine can retrieve and modify the data more faster way then human being
- Lessdrudgery: Muchofthemaintenanceoffiles by handis eliminated
- Accuracy: Accurate, up-to-date information is fetched as per requirement of the user at any time.

DatabaseManagementSystem(DBMS):

Adatabase management system consists of collection of related data and refers to aset of programs for defining, creation, maintenance and manipulation of a database.

FunctionofDBMS:

- 1. **Definingdatabaseschema**:itmustgivefacilityfordefiningthedatabasestructurealso specifies access rights to authorized users.
- 2. Manipulation of the database: The dbms must have functions like insertion of record into database, updation of data, deletion of data, retrieval of data
- **3.** Sharing of database: The DBMS must share data items for multiple users by maintaining consistency of data.
- 4. Protectionofdatabase:Itmustprotectthedatabaseagainstunauthorizedusers.
- 5. Databaserecovery: Ifforany reason the system fails DBMS must facilitate database recovery.

Advantages of DBMS:

Reductionofredundancies:

Centralized control of data by the DBA avoids unnecessary duplication of data and effectively reduces the total amount of data storage required avoiding duplication in the elimination of the inconsistencies that tend to be present in redundant data files.

SharingofData:

A database allows the sharing of data under its control by any number of application programs or users.

DataIntegrity:

Data integrity means that the data contained in the database is both accurate and consistent. Therefore data values being entered for storage could be checked to ensure that they fall with in a specified range and are of the correct format.

DataSecurity:

The DBAwho has the ultimate responsibility for the data in the dbms can ensure that proper access procedures are followed including proper authentication to access to the DataBase System and additional check before permitting access to sensitive data.

ConflictResolution:

DBA resolve the conflict on requirements of various user and applications. The DBA chooses the best file structure and access method to get optionalperformance for the application.

DataIndependence:

Data independence is usually considered from two points of views; physically data independence and logical data independence.

Physical Data Independence allows changes in the physical storage devices ororganization of the files to be made without requiring changes in the conceptual view or any of the external views and hence in the application programs using the data base.

Logical Data Independence indicates that the conceptual schema can be changed without affecting the existing external schema or any application program.

DisadvantageofDBMS:

- 1. DBMSsoftwareandhardware(networkinginstallation)cost ishigh
- 2. The processing overhead by the dbms for implementation of security, integrity and sharing of the data.
- 3. Centralizeddatabasecontrol
- 4. Setupofthedatabasesystemrequires moreknowledge, money, skills, and time.
- 5. The complexity of the database may result in poor performance.

LECTURE-3:3levelArchitectureofDBMS

Database Basics:

Data Item:

Thedataitemisalsocalledasfieldindataprocessingandisthesmallestunitofdatathathas meaning to its users. Eg:"e101","sumit"

28. erer , emin

Entitiesand attributes:

Anentityis athing orobject intherealworld that is distinguishable from all other objects Eg: Bank, employee, student

Attributesarepropertiesarepropertiesofanentity. Eg: Empcode, ename, rolno, name

Logicaldataandphysicaldata:

Logicaldataarethedataforthetablecreatedbyuser inprimarymemory. Physical datarefers to the datastored in the secondary memory.

Schemaandsub-schema:

Aschema is a logicaldatabasedescription and is drawnas achart ofthetypesofdatathat areused. It gives the names of the entities and attributes and specifythe relationships betweenthem.

Adatabaseschemaincludessuchinformationas:

- > Characteristicsofdataitemssuchasentitiesand attributes.
- Logicalstructuresandrelationshipsamongthesedataitems.
- Formatfor storagerepresentation.
- > Integrityparameterssuchasphysicalauthorizationandbackuppolicies.

A *subschema* is derived schema derived from existing schema as per the user requirement. Theremay be more then one subschema create for a single conceptual schema.

ThreeLevelArchitectureofDBMS:

Externallevel View	Vie	ew	View
Conceptual	Mapping supplie	ed by DBMS	
level	Conc		
Ma	pping supplied	by DBMS/OS	
Internallevel	Inte	ernallevel	

Adatabasemanagementsystemthatprovidesthreelevelofdataissaidtofollowthree-level architecture .

- Externallevel
- Conceptuallevel
- Internallevel

ExternalLevel:

The external level is at the highest level of database abstraction. At this level, there will be many views define for different users requirement. A view will describe only a subset of the database. Any number of user views may exist for a given global schema (coneptual schema).

For example, each student has different view of the time table. the view of a student of BTech(CSE) is different from the view of the student of Btech (ECE). Thus this level of abstraction is concerned with different categories of users.

Eachexternalviewisdescribed by meansofaschemacalledsubschema.

ConceptualLevel:

Atthis level of database abstraction all the databaseentities and the relationships among them are included. One conceptual view represents the entire database. This conceptual view is defined by the conceptual schema.

The conceptual schema hides the details of physical storage structures and concentrate ondescribing entities, data types, relationships, user operations and constraints.

It describes all the records and relationships included in the conceptual view. There is only one conceptual schema per database. It includes feature that specify the checks to relation data consistency and integrity.

Internallevel:

It is the lowestlevelofabstractionclosest to the physicalstoragemethod used. It indicates how the datawillbe stored and describes the data structures and accessmethods to be used by the database. The internal view is expressed by internal schema.

Thefollowingaspectsareconsideredatthis level:

- 1. Storageallocatione.g:B-tree,hashing
- 2. Accesspathseg.specificationofprimaryand secondarykeys,indexesetc
- 3. Miscellaneous eg. Data compression and encryptiontechniques, optimization of the internal structures.

DatabaseUsers:

NaiveUsers:

Users who need not be awareofthe presence of the database systemoranyother systemsupporting their usage are considered naïve users. Auserofanautomatic teller machine fallsonthiscategory.

OnlineUsers:

These are users who may communicate with the database directly via an online terminal or indirectlyvia a user interface and application program. These users are aware of the database system and also know the data manipulation language system.

ApplicationProgrammers:

Professional programmers who are responsible for developing application programs or user interfaces utilized by the naïve and online user falls into this category.

DatabaseAdministration:

Apersonwho hascentralcontrolover the system is called database administrator. The function of DBA are :

- 1. CreationandmodificationofconceptualSchemadefinition
- 2. Implementationofstoragestructureand accessmethod.
- 3. Schemaandphysicalorganizationmodifications.
- 4. Grantingofauthorization for dataaccess.
- 5. Integrityconstraints specification.
- 6. Executeimmediate recoveryprocedure incase of failures
- 7. Ensure physicalsecuritytodatabase

Databaselanguage:

1) **Datadefinitionlanguage(DDL):**

DDL is used to define database objects .The conceptual schema is specified by a set of definitions expressed by this language. It also gives some details about how to implement this schema in the physical devices used to store the data. This definition includes all the entity sets and their associated attributes and their relationships. The result of DDL statements will be set oftables that are stored in special file called data dictionary.

2) DataManipulationLanguage(DML):

ADML is a language that enables users to access or manipulate data stored in the database. Datamanipulation involves retrievalofdatafrom database, insertionofnew datainto the database and deletion of data or modification of existing data.

There are basicallytwotypesofDML:

- **Procedural**: Which requires a user to specify what data is needed and how to get it.
- Non-Procedural:whichrequiresausertospecifywhatdataisneededwithout specifying how to get it.

3) DataControlLanguage(DCL):

This language enables user to grant authorization and canceling authorization of database objects.

LECTURE-4:Elementsof DBMS

Elementsof DBMS:

DMLPre-Compiler:

It converts DML statements embedded in an application program to normal procedure calls in the host language. The pre-complier must interact with the query processor in order to generate the appropriate code.

DDLCompiler:

The DDL compiler converts the data definition statements into a set of tables. These tables contains information concerning the database and are in a form that can be used by other components of the dbms.

File Manager:

File manager manages the allocation of space on disk storage and the data structure used to represent information stored on disk.

Database Manager:

Adatabase manager is a program module which provides the interface between the low level data stored in the database and the application programs and queries submitted to the system.

Theresponsibilities of database managerare:

- 1. Interaction with File Manager: The datais storedonthe disk using the file systemwhich is provided by operating system. The database manager translate the different DML statements into low-levelfile systemcommands so the database manager is responsible for the actual storing, retrieving and updating of data in the database.
- 2. Integrity Enforcement: The data values stored in the database must satisfy certain constraints (eg: the age of a person can't be less then zero). These constraints are specified byDBA. Datamanager checkstheconstraints and ifit satisfiesthenit storesthedatainthe database.
- **3.** Security Enforcement: Data manager checks the security measures for database from unauthorized users.
 - 4. Backup and Recovery: Database manager detects the failures occur due to different causes (like disk failure, power failure, deadlock, software error) and restores the database to original state of the database.
 - 5. Concurrency Control: When several users access the same database file simultaneously, there may be possibilities of data inconsistency. It is responsible of database manager to control the problems occur for concurrent transactions.

QueryProcessor:

The queryprocessorused to interpret to online user's query and convert it into an efficient series of operations in a form capable of being sent to the data manager for execution. The query processor

usesthedatadictionarytofindthedetailsofdatafileandusingthisinformationitcreatequery plan/access plan to execute the query.

DataDictionary:

Datadictionaryisthetablewhichcontainstheinformationaboutdatabaseobjects.Itcontains information like

- 1. external, conceptual and internal databased escription
- 2. descriptionofentities, attributes as well as meaning of data elements
- 3. synonyms, authorizationandsecuritycodes
- 4. databaseauthorization

Thedatastored inthedatadictionary iscalled meta data.

DBMSSTRUCTURE:



 $\label{eq:Que:Listfoursignificant differences between a File-Processing System and a DBMS.$

Ans: Some major differences between a database management systemand a file-processing system are:

• Both systems contain a collection of data and a set of programs which access that data. A database management system coordinates both the physical and the logical access to thedata, whereas a file-processing system coordinates only the physical access.

- A database management system reduces the amount of data duplication by ensuring that a physical piece of data is available to all programs authorized to have access to it, where as data written by one program in a file-processing system may not be readable by another program.
- A database management system is designed to allow flexible access to data (i.e., queries), whereas a file-processing system is designed to allow predetermined access to data (i.e., compiled programs).
- Adatabase management system is designed to coordinate multiple users accessing the same data at the same time. A file-processing system is usually designed to allow one or more programs to access different data files at the same time. In a file-processing system, a file can be accessed by two programs concurrently only if both programs have read-only access to the file.

Que:Explainthedifferencebetweenphysicalandlogicaldata independence.

Ans:

- Physical data independence is the ability to modify the physical scheme without making it necessary to rewrite application programs. Such modifications include changing from unblocked to blocked record storage, or from sequential to random access files.
- Logicaldata independence is the abilityto modifythe conceptualscheme without making it necessaryto rewrite application programs. Such a modification might be adding a field to a record; an application program's view hides this change from the program.

Que: List five responsibilities of a database management system. For each responsibility, explain the problems that would arise if the responsibility were not discharged.

Ans: Ageneral purposed at a basemanager (DBM) has five responsibilities:

- a. interactionwiththefilemanager.
- b. integrityenforcement.
- c. securityenforcement.
- d. backupand recovery.
- e. concurrencycontrol.

If these responsibilities were not met by a given DBM (and the text points out that sometimes a responsibility is omitted by design, such as concurrency control on a single-user DBM for a micro computer) the following problems can occur, respectively:

- a. No DBM can do without this, if there is no file manager interaction then nothing stored in he files can be retrieved.
- b. Consistency constraints may not be satisfied, when account balances could go below the minimum allowed, employees could earn too much overtime (e.g.,hours > 80) or, airline pilots may fly more hours than allowed by law.
- $c. \ \ Unauthorized users may access the data base, or users authorized to access part of the$

database may be able to access parts of the database for which they lack authority. For example, a high school student could get access to national defense secret codes, or employees could find out what their supervisors earn.

- d. Data could be lost permanently, rather than at least being available in a consistent state that existed prior to a failure.
- e. Consistencyconstraints may be violated when intgrityconstraints failed ina transaction. For example, incorrect bank balances might be reflected due to simultaneous withdrawals and deposits, and so on.

Que. What are five main functions of a database a dministrator?

Ans: Fivemainfunctionsofadatabaseadministrator are:

- Tocreatetheschemedefinition
- Todefinethestoragestructureandaccess methods
- Tomodifytheschemeand/orphysicalorganizationwhennecessary
- Tograntauthorizationfordataaccess
- Tospecifyintegrityconstraints

 $\label{eq:Que:Listsix} Que: Listsix majors teps that you would take in setting up a data base for a particular enterprise.$

Ans: Sixmajorstepsinsettingupadatabaseforaparticularenterpriseare:

- Define the highle velre quirements of the enterprise (this step generates a document known as the system requirements specification.)
- Defineamodelcontainingallappropriate typesofdata and datarelationships.
- Define the integrity constraints on the data.
- Definethephysicallevel.
- For each known problem to be solved on a regular basis (e.g., tasks to be carried out by clerks or Web users) define a user interface to carry out the task, and write the necessary application programs to implement the user interface.
- Create/initializethedatabase.

EXERCISE:

- 1. Whatisdatabasemanagementsystem?
- 2. Whatarethedisadvantageoffileprocessingsystem?
- 3. Stateadvantageanddisadvantageofdatabasemanagementsystem.
- 4. Whataredifferenttypesofdatabaseusers?
- 5. Whatisdatadictionaryandwhatareitscontents?
- 6. WhatarethefunctionsofDBA?
- 7. Whatarethedifferentdatabaselanguages?Explainwithexample.
- 8. ExplainthethreelayerarchitectureofDBMS.
- 9. Differentiatebetweenphysicaldataindependenceandlogicaldataindependence.
- 10. Explainthefunctionsofdatabasemanager.
- 11. Explainmetadata.

LECTURE-5:ER-MODEL

DataModel:

The data model describes the structure of a database. It is a collection of conceptual tools for describing data, data relationships and consistency constraints and various types of data models uch as

- 1. Objectbasedlogicalmodel
- 2. Recordbased logical model
- 3. Physicalmodel

Typesofdatamodel:

- 1. Objectbasedlogicalmodel
 - a. ER-model
 - b. Functionalmodel
 - c. Objectoriented model
 - d. Semanticmodel
- 2. Recordbased logical model
 - a. Hierarchicaldatabasemodel
 - b. Networkmodel
 - c. Relationalmodel
- 3. Physicalmodel

Entity RelationshipModel(ERModel)

The entity-relationship data model perceives the real world as consisting of basic objects, called entities and relationships among these objects. It was developed to facilitate database design by allowing specification of an enterprise schema which represents the overall logical structure of a data base.

MainFeaturesofER-MODEL:

- Entityrelationshipmodelisa high levelconceptualmodel
- It allows us to describe the data involved in a real world enterprise in terms of objects and their relationships.
- Itiswidelyused todevelop an initialdesignofa database
- It provides a set of useful concepts that make it convenient for a developer to move from a basic set of information to a detailed and description of information that can be easily implemented in a database system
- Itdescribesdataasacollectionofentities, relationships and attributes.

Basic Concepts:

The E-R data model employs three basic notions: entity sets, relationship sets and attributes.

Entity Sets:

Anentity is a "thing" or "object" in the realworld that is distinguishable fromall other objects. For example, each person in an enterprise is an entity. An entity has a set properties and the values for some set of properties mayuniquely identify an entity. BOOK is entity is properties (called as attributes) bookcode, booktitle, price etc.

Anentitysetisasetofentities of the same type that share the same properties, or attributes. The set

ofallpersonswhoarecustomersatagiven bank.

Attributes:

An entity is represented by a set of attributes. Attributes are descriptive properties possessed by each member of an entity set.

Customer is an entity and its attributes are customerid, custmername, custaddress etc.

Anattributeasused inthe E-Rmodel, can be characterized by the following attribute types.

a) SimpleandCompositeAttribute:

Simple attributes are the attributes which can't be divided into sub parts, e.g. customerid, empno Composite attributes are the attributes which can be divided into subparts, e.g. name consisting of first name, middle name, last name and address consisting of city, pincode, state.

b) Single-ValuedandMulti-ValuedAttribute:

The attribute having unique value is single –valued attribute, e.g. empno, customerid, regdno etc. The attribute having more than one value is multi-valued attribute, eg: phone-no, dependent name, vehicle.

c) DerivedAttribute:

The values for this type of attribute can be derived from the values of existing attributes, e.g. age which can be derived from currentdate – birthdate and experience_in_year can be calculated as currentdate-joindate.

d) NULLValuedAttribute:

TheattributevaluewhichisnotknowntouseriscalledNULLvalued attribute.

RelationshipSets:

Arelationship is an association among several entities. Arelationship set is a set of relationships of the same type. Formally, it is a mathematical relation on n>=2 entity sets. If E_1 , E_2 ... E_n are entity sets, then a relation ship set R is a subset of

 $\{(e_1,e_2,\ldots e_n)|e_1\in E_1,e_2\in E_2..,e_n\in E_n\}$ where $(e_1,e_2,\ldots e_n)$ is a relation ship.



Consider the two entitysets customer and loan. We define the relationship set borrow to denote the association between customers and the bank loans that the customers have.

Mapping Cardinalities:

Mapping cardinalities or cardinality ratios, express the number of entities to which another entity can be associated via a relationship set. Mapping cardinalities are most useful in describing binary relationship sets, although they can contribute to the description of relationship sets that involve more than twoentity sets. For a binary relationship set R between entity setsA and B, the mapping

cardinalitiesmustbeoneofthe following:

1. OnetoOne:

An entity in A is associated with at most one entity in B, and an entity in B is associated with at most one entity in A.





2. Oneto Many:

AnentityinAisassociated withany number of entities in B.Anentityin B is associated withat the most one entity in A.

Eg:Relationshipbetweendepartmentandfaculty



3. ManytoOne:

AnentityinAisassociatedwithatmostoneentityinB.AnentityinBisassociatedwithany number in A.



4. ManytoMany:

EntitiesinAand Bareassociated withany numberofentitiesfromeachother.



Moreabout EntitiesandRelationship:

RecursiveRelationships:

When the same entity type participates more than once in a relationship type in different roles, the relationship types are calledrecursive relationships.

Participation Constraints:

Theparticipation constraints specify whether the existence of any entity depends on its being related to another entity via the relationship. There are two types of participation constraints

a) Total : When all the entities from an entity set participate in a relationship type, is called total participation. For example, the participation of the entity set student on the relationship set must 'opts' is said to be total because everystudent enrolled must opt for a course.

b) Partial: When it is not necessary for all the entities from an entity set to particapte in a relationship type, it is called partial participation. For example, the participation of the entity set student in 'represents' is partial, since not everystudent in a class is a class representative.

WeakEntity:

Entitytypes that do not containanykeyattribute, and hence can not be identified independentlyare called weak entity types. Aweak entity can be identified by uniquely only by considering some of its attributes in conjunction with the primary key attribute of another entity, which is called the identifying owner entity.

Generallyapartialkeyisattachedtoaweakentitytypethatisusedforuniqueidentificationof weak entities related to a particular owner type. The following restrictions must hold:

- Theownerentitysetandtheweakentitysetmustparticipateinonetomayrelationshipset. Thisrelationshipsetiscalledtheidentifyingrelationshipsetoftheweakentityset.
- $\bullet \quad The weak entity set must have total participation in the identifying relationship.$

Example:

Consider the entity type Dependent related to Employee entity, which is used to keep track of the dependents of each employee. The attributes of Dependents are: name, birthdate, sex and relationship. Each employee entityset is said to its own the dependent entities that are related to it. However, not that the 'Dependent'entitydoes not exist of its own, it is dependent on the Employee entity.

Keys:

SuperKey:

Asuperkeyisaset of one or more attributes that taken collectively, allowus to identify uniquely an entity in the entity set. For example, customer-id, (cname, customer-id), (cname, telno)

CandidateKey:

In a relation R, a candidate key for R is a subset of the set of attributes of R, which have the following properties:

- 1. Uniqueness: NotwodistincttuplesinR havethesamevaluesforthecandidatekey
- 2. *Irreducible*:Noproper subsetof the candidate key has the uniqueness property thatis the candidate key. Eg: (cname,telno)

PrimaryKey:

The primarykey is the candidate keythat is chosen by the database designer as the principal means of identifying entities within an entityset. The remaining candidate keys if any, are called *Alternate Key*.

LECTURE-6:ER-DIAGRAM:

The overall logical structure of a data base using ER-model graphically with the help of an ER-diagram.

SymbolsuseER-diagram:





Figure 2.1 E-R diagram for a Car-insurance company.



Figure 2.2 E-R diagram for a hospital.

AUniveristyregistrar's office maintains data about the following entities:

- (a) Course, includeing number, title, credits, syllabus and prerequisites
- (b) courseoffering,includingcoursenumber, year, semester, sectionnumber, instructor timings, and class room
- (c) Studentsincludingstudent-id,nameandprogram
- (d) Instructors, including identification number, name, department and title

further, the enrollment of students incourses and grades awarded to students in each course they are enrolled for must be appropriate modeled.

Construct an E-R diagram for the registrar's office. Document all assumptions that you may makeabout the mapping constraints



Figure 2.3 E-R diagram for a university.



Figure 2.4 E-R diagram for marks database.



Figure 2.7 E-R diagram for all teams statistics.

Cosidet a universitydatabase for the scheduling of class rooms for final exams. This database could be modeled as the single entityset exam, with a tributes course-name, section-number, room-number and time, Alternatively, one or more additional entitysets would be defined, along with relationship sets to replae some of the attributes of the exam entity set, as

- coursewithattributesname,departmentand c-number
- section with attributess-number and enroll ment and dependent as a weaken tity set on course
- roomwithattributesr-number,capacityandbuilding



Figure 2.12 E-R diagram for exam scheduling.



Figure 2.12 E-R diagram for exam scheduling.

LECTURE-7:AdvancedER-Diagram:

Abstraction is the simplification mechanism used to hide superfluous details of a set of objects. It allowsoneto concentrateontheproperties that areofinterest to the application. There are two main abstraction mechanism used to model information:

Generalizationandspecialization:

Generalization is the abstracting process of viewing set of objects as a single general class by concentrating on the general characteristics of the constituent sets while suppressing or ignoring their differences. It is the union of a number of lower-levelentity types for the purpose of producing a higher-levelentitytype. For instance, student is a generalization of graduate or undergraduate, full-time or part-time students. Similarly, employee is generalization of the classesof objects cook, waiter, and cashier. Generalization is an IS_A relationship; therefore, manager IS_AN employee, cook IS_AN employee, waiter IS_AN employee, and so forth.

Specialization is the abstracting process of introducing new characteristics to an existing class of objects to create one or more new classes of objects. This involves taking a higher-level, and using additional characteristics, generating lower-level entities. The lower-level entities also inherits the, characteristics of the higher-level entity. In applying the characteristics size to car we can create a full-size, mid-size, compact orsubcompact car.Specialization maybe seenasthe reverse processof generalization addition specific properties are introduced at a lower level in a hierarchyof objects.



EMPLOYEE(empno,name,dob) FULL_TIME_EMPLOYEE(empno,salary) PART_TIME_EMPLOYEE(empno,type) Faculty(**empno**,degree,intrest) Staff(**empno**,hour-rate) Teaching (**empno**,stipend)



Figure 2.19 E-R diagram of motor-vehicle sales company.

Aggregation:

Aggregation is the process of compiling information on an object, there by abstracting a higher level object. The entity person is derived by aggregating the characteristics of name, address, ssn. Another formof the aggregation is abstracting a relationship objects and viewing the relationship as an object.





Figure 2.8 E-R diagram Example 1 of aggregation.



Figure 2.9 E-R diagram Example 2 of aggregation.

<u>ER-DiagramForCollegeDatabase</u>



LECTURE-8: ConversionofER-DiagramtoRelationalDatabase

ConversionofEntitySets:

1. For each strong entity type E in the ER diagram, we create a relation R containing all the single attributes of E. The primarykeyofthe relation R will be one of the keyattribute of R.

STUDENT(rollno (primary key),name, address) FACULTY(id(primary key),name ,address, salary) COURSE(course-id,(primarykey),course_name,duration) DEPARTMENT(dno(primary key),dname)

2. For each weak entity type W in the ER diagram, we create another relation R that contains allsimple attributes of W. If E is anowner entity of W thenkey attribute of E is also include In R. This key attribute of R is set as a foreign key attribute of R. Now the combination of primary key attribute of owner entity type and partial key of the weak entity type will form the key of the weak entity type

GUARDIAN((rollno,name)(primarykey),address,relationship)

ConversionofRelationshipSets:

BinaryRelationships:

• One-to-One Relationship:

For each 1:1 relationship type R in the ER-diagram involving two entities E1 and E2 we choose one of entities(say E1) preferably with total participation and add primary key attributeofanother E asaforeignkeyattribute inthetableofentity(E1).Wewillalso include all the simple attributes of relationship type R in E1 if any, For example, the department relationship has been extended tp include head-id and attribute of the relationship. DEPARTMENT(D_NO,D_NAME,HEAD_ID,DATE_FROM)

• One-to-Many Relationship:

For each1:NrelationshiptypeR involvingtwo entitiesE1andE2, we identifytheentitytype (say E1) at the N-side of the relationship type R and include primary key of the entity on the other side of the relation (sayE2) as a foreign key attribute in the table of E1. We include all simple attribute (or simple components of a composite attribute of R (ifany) in the table E1)

Forexample:

The worksinrelationship between the DEPARTMENT and FACULTY. For this relationship choose the entity at N side, i.e, FACULTY and add primary key attribute of another entity DEPARTMENT i.e., DNO as a foreign key attribute in FACULTY.

FACULTY(CONTAINSWORKS_INRELATIOSHIP) (ID, NAME, ADDRESS, BASIC_SAL, DNO)

• Many-to-ManyRelationship:

For each M:N relationship type R, we create a new table (say S) to represent R, we also include the primary key attributes of both the participating entity types as a foreign key attribute in S.Any simple attributes of the M:N relationship type (or simple components as a composite attribute) is also included as attributes of S.

Forexample:

The M:N relationship taught-by between entities COURSE and FACULTY should be represented as a new table. The structure of the table will include primary key of COURSE and primary key of FACULTY entities.

TAUGHT-BY (ID (primary key of FACULTY table), course-id (primary key of COURSE table)

• N-aryRelationship:

For each N-ary relationship type R where n>2, we create a new table S to represent R, We include as foreign key attributes in S the primary keys of the relations that represent the participatingentitytypes. Wealso includeany simple attributes of the N-ary relationship (or simple components of complete attribute) as attributes of S. The primary key of S is usually a combination of all the foreign keys that reference the relations representing the participating entity types.



LOAN-SANCTION(cusomer-id, loanno,empno,sancdate,loan_amount)

• Multi-ValuedAttributes:

For each multivalued attribute 'A', we create a new relation R that includes an attribute corresponding to plus the primary key attributes k of the relation that represents the entity type or relationship that has as an attribute. The primary key of R is then combination of A and k.

Forexample, ifaSTUDENTentityhasrollno, nameandphone numberwherephone number is a multivalued attribute then we will create table PHONE (rollno, phoneno) where primary key is the combination. In the STUDENT table we need not have phone number, instead if can be simply (rollno, name) only.

PHONE(rollno,phoneno)



• ConvertingGeneralisation/SpecificationHierarchytoTables: A simpleruleforconversionmay betodecomposeall thespecializedentitiesintotablein case theyare disjoint, for example, for the figure we can create the three tables as: Account(account_no,name,branch,balance) Saving_Account (account-no, intrest) Current_Account (account-no, charges)

LECTURE-9:RecordBasedLogicalModel

HierarchicalModel:

- Ahierarchicaldatabaseconsistsofacollectionof*records*whichareconnectedtoone another through *links*.
- Arecord is a collection offields, each of which contains only one datavalue.
- Alinkisanassociationbetweenpreciselytwo records.
- The hierarchical model differs from the network model in that the records are organized as collections of trees rather than as arbitrary graphs.

Tree-StructureDiagrams:

- Theschemafor ahierarchicaldatabaseconsists of
 - o *boxes*, which correspond to record types
 - o *lines*, which correspond to links
- Record types are organized in the form of a root edtree.
 - Nocyclesintheunderlyinggraph.
 - Relationshipsformedinthegraphmustbesuchthatonly one-to-manyorone-to-onerelationshipsexistbetweenaparentanda child.

Databaseschemaisrepresented asacollectionoftree-structurediagrams.

- singleinstanceofadatabasetree
- The rootofthistree is a dummynode
- The childrenofthatnodeareactualinstances of the appropriate record type

WhentransformingE-Rdiagramsto correspondingtree-structurediagrams, we must ensure that the resulting diagrams are in the form of rooted trees.

SingleRelationships:

- ExampleofE-Rdiagramwithtwoentitysets, *customer* and *account*, related through a binary, one-to-many relationship *depositor*.
- Correspondingtree-structurediagramhas
 - therecord type *customer* with three fields: *customer-name*, *customer-street*, and *customer-city*.
 - therecord type *account* with two fields: *account-number* and *balance*
 - o thelink*depositor*, withanarrowpointingto*customer*
- If the relationship *depositor* is one to one, then the link *depositor* has two arrows.

customer-nam	ie custom	ier-street	custo	mer-city	customer	
l	account-nui	mber ba	lance	account		

• Onlyone-to-manyandone-to-onerelationshipscanbedirectlyrepresented in the hierarchical mode.

TransformingMany-To-Many Relationships:

- Must consider thetype of queries expected and the degree to which the database schema fits the given E-R diagram.
- Inallversionsofthistransformation, the underlying database tree (or trees) will have replicated records.



- Createtwotree-structurediagrams, *T*1, with the root *customer*, and *T*2, with the root *account*.
- In*T*1,create*depositor*,amany-to-onelink from*account*to*customer*.
- In*T*2,create*account-customer*,amany-to-one link from*customer*to*account*.



VirtualRecords:

- Formany-to-manyrelationships,recordreplicationisnecessarytopreservethetree- structure organization of the database.
- Datainconsistencymayresultwhenupdating takes place
- Wasteofspaceisunavoidable

- *Virtualrecord*—containsnodatavalue,onlyalogicalpointertoaparticularphysical record.
- Whenarecordistobereplicated inseveral database trees, a single copy of that record is kept in one of the trees and all other records are replaced with a virtual record.
- Let *R* be a record type that is replicated in *T*1, *T*2,...,*Tn*. Create a new virtual record type *virtual*-*R* and replace *R* in each of the *n*-1 trees with a record of type *virtual*-*R*.
- Eliminate data replication in the following diagram ;create *virtual-customer* and *virtual-account*.
- Replace*account*with*virtual-account*inthefirsttree,andreplace*customer*with*virtual- customer* in the second tree.
- Addadashedlinefrom*virtual-customer*to*customer*, and from *virtual-account*to *account*, to specify the association between a virtual record and its corresponding physical record.



NetworkModel:

- Data are represented by collections of *records*.
 - o similartoanentityinthe E-Rmodel
 - Records and their fields are represented as *record type*
 - typecustomer=recordtypeaccount= record typecustomer-name:string;account-number:integer;customer-street:string;balance:integer;customer-city:string;balance:integer;
- end

end

- Relationshipsamong dataarerepresented by links
 - similartoarestricted(binary)formofanE-Rrelationship
 - \circ restrictions on links depend on whether the relationship is many-to-many, many-to-one, or one-to-one.

Data-StructureDiagrams:

- Schemarepresentingthe designofa networkdatabase.
- Adata-structurediagramconsistsoftwobasic components:
 - Boxes, which correspond to record types.
 - o Lines, which correspond to links.
- Specifiesthe overalllogicalstructureofthe database.

 $For every E-R diagram, there is a corresponding \ data-structure diagram.$



Since a link cannot contain any data value, represent an E-R relationship with attributes with a new record type and links.



Torepresentan E-R relationship of degree 3 or higher, connect the participating record typesthrough a new record type that is linked directly to each of the original record types.

- 1. Replace entitysets*account,customer*, and*branch*withrecordtypes*account,customer*, and *branch*, respectively.
- 2. CreateanewrecordtypeRlink(referred toasadummyrecordtype).
- 3. Createthefollowingmany-to-one links:
 - *CustRlink* from*Rlink* recordtype to*customer*record type
 - AcctRlnkfromRlinkrecordtype toaccountrecord type
 - *BrncRlnk*from*Rlink*recordtype to*branch*recordtype



TheDBTGCODASYLModel:

- Alllinksare treated as many-to-one relationships.
- To model many-to-many relationships, a record type is defined to represent the relationship and two links are used.



DBTGSets:

- Thestructureconsisting of two record types that are linked together is referred to in the DBTG model as a *DBTG set*
- In eachDBTG set, one recordtype is designated as the *owner*, and the other is designated as the *member*, of the set.
- EachDBTGsetcanhaveanynumber of *setoccurrences* (actualinstances of linked records).
- Since many-to-many links are disallowed, each set occurrence has precisely one owner, and has zero or more member records.
- o Nomemberrecordofasetcanparticipateinmorethanoneoccurrenceofthesetatany point.
- Amemberrecordcanparticipatesimultaneouslyinseveralsetoccurrencesof*different* DBTGsets.

LECTURE-10:RELATIONALMODEL

RELATIONALMODEL

Relationalmodelissimplemodelinwhichdatabaseisrepresentedasacollectionof"relations" whereeachrelationis representedbytwo-dimensionaltable.

account_number	branch_name	balance
A-101	Downtown	500
A-102	Perryridge	400
A-201	Brighton	900
A-215	Mianus	700
A-217	Brighton	750
A-222	Redwood	700
A-305	Round Hill	350

The relational model was founded by E. F. Coddof the IBM in 1972. The basic concept in the relational model is that of a relation.

Properties:

- It is column homogeneous. In other words, in any given column of a table, all items are of the same kind.
- Eachitem is a simple number or a character string. That is a table must be infirst normal form.
- Allrowsofatablearedistinct.
- Theordering ofrowswith inatableisimmaterial.
- \circ The column of a table are assigned distinct names and the ordering of the secolumn sis immaterial.

Domain, attributestuplesandrelational:

Tuple:

Eachrowinatablerepresentsarecordandiscalledatuple. A tablecontaining 'n'attributesina recordiscallediscalledn-tuple.

Attributes:

The name of each column in a table is used to interpret its meaning and is called an attribute.Each table is called a relation. In the above table, account_number, branch name, balance are the attributes.

Domain:

A domainisasetofvaluesthatcanbegiventoanattributes.Soeveryattributeinatablehasa specific domain. Values to these attributes can not be assigned outside their domains.

Relation:

Arelationconsist of

- Relationalschema
- Relationinstance

RelationalSchema:

Arelationalschemaspecifiestherelation'sname, its attributes and the domain of each attribute. If
R is the name of a relation and A1, A2,...An is a list of attributes representing R then R(A1,A2,...,An) is called a Relational Schema. Each attribute in this relational schema takes avalue from some specific domain called domain(Ai).

Example:

PERSON(PERSON_ID:INTEGER,NAME:STRING, AGE:INTEGER, ADDRESS:STRING)

Total number of attributes in a relation denotes the degree of a relation since the PERSON relation scheme contains four attributes, so this relation is of degree 4.

RelationInstance:

Arelational instance denoted as r is a collection of tuples for a given relational schema at a specific point of time.

Arelationstaterto therelationsschemaR(A1,A2...,An) also denotedbyr(R) isaset ofn-tuples

 $R{t1,t2,...tm}$

Whereeach n-tupleisanordered listofn values

T=<v1,v2,....vn>

Whereeachvibelongstodomain(Ai) or containsnullvalues.

Therelationschema isalso called 'intension' and the relationstate is also called 'extension'.

Eg:Relationschemafor Student

STUDENT(rollno:string,name:string,city:string,age:integer)

Relationinstance:

Student:

uciti			
Rollno	Name	City	Age
101	Sujit	Bam	23
102	kunal	bbsr	22

Keys:

Superkey:

Asuperkeyisanattributeorasetofattributesusedto identifytherecordsuniquelyinarelation. For example, customer-id, (cname, customer-id), (cname,telno)

Candidatekey:

• Uniqueness:

Super keys of a relation can contain extra attributes. Candidate keys are minimal super keys. i.e, such a keycontains no extraneous attribute. An attribute is called extraneous ifeven after removing it from the key, makes the remaining attributes still has the properties of a key(atribute represents entire table).

In a relation R, a candidate key for R is a subset of the set of attributes of R, which have the following properties:

- notwodistincttuplesinR havethesamevalues for
 - thecandidatekey

• *Irreducible:* Nopropersubset of the candidate key has the

uniqueness property that is the candidate key.

- Acandidatekey'svaluesmustexist.Itcan'tbenull.
- Thevaluesofacandidatekeymustbestable.Itsvaluecannotchangeoutsidethe

controlofthesystem.

Eg:(cname,telno)

Primarykey:

The primary key is the candidate key that is chosen by the database designer as the principal means of identifying entities with in an entity set. The remaining candidate keys if any arecalled *alternate key*.

LECTURE-11:CONSTRAINTS

RELATIONALCONSTRAINTS:

Thereare three types of constraints on relational database that include

- DOMAINCONSTRAINTS
- KEYCONSTRAINTS
- INTEGRITYCONSTRAINTS

DOMAINCONSTRAINTS:

Itspecifies thateach attribute in a relation an atomic value from the corresponding domains. The data types associated with commercial RDBMS domains include:

- Standardnumericdatatypes for integer
- Realnumbers
- Characters
- Fixedlengthstringsandvariablelengthstrings

Thus, domain constraints specifies the condition that we to put on each instance of the relation. So the values that appear in each column must be drawn from the domain associated with that column.

Rollno	Name	City	Age
101	Sujit	Bam	23
102	kunal	bbsr	22

KeyConstraints:

This constraints states that the key attribute value each tuplemsube unique .i.e, no two tuples contain the same value for the key attribute.(null values can allowed)

Emp(empcode, name, address). here empcode can be unique

IntegrityCONSTRAINTS:

Therearetwotypes of integrity constraints:

- EntityIntegrityConstraints
- ReferentialIntegrityconstraints

EntityIntegrity Constraints:

Itstatesthatno primarykeyvaluecanbe nulland unique. This isbecausetheprimarykeyisusedto identify individual tuple in the relation. So we will not be able to identify the records uniquely containing null values for the primary key attributes. This constraint is specified on one individual relation.

ReferentialIntegrityConstraints:

It states that the tuple in one relation that refers to another relation must refer to an existing tuple in that relation. This constraints is specified on two relations. If a column is declared as foreign key that must be primary key of another table.

Department(deptcode, dname) Herethedeptcodeistheprimarykey.

Emp(*empcode*,name,city,*deptcode*). Herethedeptcodeis foreignkey.

CODD'SRULES

Rule1:Theinformation Rule.

"All information in a relationaldata base is represented explicitlyat the logical leveland in exactly one way - by values in tables."

Everything within the database exists in tables and is accessed via table accessroutines.

Rule2: GuaranteedaccessRule.

"Each and every datum (atomic value) in a relational data base is guaranteed to be logically accessible by resorting to a combination of table name, primarykey value and column name."

To access any data-item you specify which column within which table it exists, there is no reading of characters 10 to 20 of a 255 byte string.

Rule3:Systematictreatmentofnullvalues.

"Null values (distinct from the empty character string or a string of blank characters and distinct from zero or any other number) are supported in fully relational DBMS for representing missing information and inapplicable information in a systematic way, independent of datatype."

If data does not exist or does not apply then a value of NULL is applied, this is understood by the RDBMS as meaning non-applicable data.

Rule4:Dynamicon-linecatalogbased ontherelationalmodel.

"The data base description is represented at the logical level in the same way as-ordinary data, so that authorized users canapply the same relational language to its interrogation they apply to the regular data."

The Data Dictionary is held within the RDBMS, thus there is no-need for off-line volumes to tell you the structure of the database.

Rule5:Comprehensivedatasub-languageRule.

"A relational system may support several languages and various modes of terminal use (for example, the fill-in-the-blanks mode). However, there must be at least one language whose statements are expressible, per some well-defined syntax, as character strings and that is comprehensive in supporting all the following items

- DataDefinition
- ViewDefinition
- DataManipulation(Interactiveandbyprogram).
- IntegrityConstraints
- Authorization.

Every RDBMS should provide alanguage toallow the user toquery the contents of the RDBMS and also manipulate the contents of the RDBMS.

Rule6:.ViewupdatingRule

"Allviewsthataretheoreticallyupdateablearealso updateablebythesystem."

Notonlycantheusermodifydata, butsocantheRDBMS when the user is not logged-in.

Rule7 :High-levelinsert,updateanddelete.

"The capability of handling a base relation or a derived relation as a single operand applies not only to the retrieval of data but also to the insertion, update and deletion of data."

The user shouldbe able tomodify several tables by modifying the view towhich they actas base tables.

Rule8 : Physical data independence.

"Application programs and terminal activities remain logically unimpaired whenever any changes are made in either storage representations or access methods."

Theuser shouldnot beawareofwhereoruponwhichmediadata-filesarestored

Rule 9 :Logicaldataindependence.

"Application programs and terminal activities remain logically unimpaired when informationpreserving changes of any kind that theoretically permit un-impairment are made to the basetables." User programs and the user should not be aware of any changes to the structure of the tables (such as the addition of extra columns).

Rule10: Integrityindependence.

"Integrity constraints specific to a particular relational data base must be definable in the relational data sub-language and storable in the catalog, not in the application programs."

If a column only accepts certain values, then it is the RDBMS which enforces these constraints and not the user program, this means that an invalid value can never be entered into this column, whilst if the constraints were enforced via programs there is always a chance that a buggyprogram might allow incorrect values into the system.

Rule11:Distributionindependence.

"ArelationalDBMShasdistributionindependence."

The RDBMS mayspread across more thanone systemand across several networks, however to the end-user the tables should appear no different to those that are local.

Rule12:Non-subversion Rule.

"If a relational system has a low-level (single-record-at-a-time) language, that low level cannot be used to subvert or bypass the integrityRules and constraints expressed in the higher levelrelational language (multiple-records-at-a-time)."

LECTURE-12:FILEORGANISATION

FILEORGANISATIONANDITSTYPES:

Afile organization is a technique to organize data in the secondary memory. File organization is a wayofarranging the records in a file when the file is stored on the disk. Data files are organized so as to facilitate access to records and to ensure their efficient storage. ADBMS supports several file organization techniques.



Heapfiles(unorderedfile)

Basicallythese files are unordered file. It is the simplest and most basic type. These files consist of randomlyordered records. The records willhave no particular order. The operations we can perform on the records are insert, retrieve and delete. The features of the heap file organization are:

- Newrecordscanbeinsertedinanyemptyspace that can accommodate them.
- Whenold records are deleted, the occupied space becomes empty and available for any new insertion.
- Ifupdatedrecordsgrow, they mayneed to be relocated to an even ptyspace. This needs to keep a list of emptyspace.

Advantageofheapfiles:

- 1. Thisisasimplefileorganizationmethod
- 2. Insertionissomehow efficient
- 3. Goodforbulk-lading datainto atable.
- 4. Bestiffilescansarecommonorinsertionsarefrequent

Disadvantagesofheapfiles:

1. Retrievalrequiresalinearsearchandis inefficient

2. Deletioncanresultinunusedspace/needforreorganization

Sequentialfileorganization:

The most basic way to organize the collection of records in a file is to use sequential organization. Records ofthe file are stored in sequence by the primarykey field values/ Theyare accessible only intheorder storedi.e, inthe primarykey order. This kind is offile organization works well for tasks which need to access nearly every record in a file. Eg. Payroll..

In a sequentially organized file records are written consecutively when the file is created and must be accessed consecutively when the file later used for input.

Asequential file maintains the records in the logicalsequence of its primarykey values. Sequential file are inefficient for random access. And files can be stored on devices like magnetic tape that allow sequential access. As records are in sorted order it will use binary search technique to search for a record.

Advantagesofsequentialfileorganization:

• Itisfastandefficientwhendealingwithlargevolumesof datathatneedtobeprocessed periodically(batch system)

Disadvantagesofsequentialfileorganization:

- Requires that all new transactions be sorted into the proper sequence for sequential accessprocessing
- Locating, storing, modifying, deleting or adding records in the file requirer earranging the file/
- Thismethodistooslowtohandleapplicationrequiringimmediateupdatingorresponses.

Indexedsequentialfileorganization:

It organized the file like a large dictionary, i.e, records are stored in order of the key, but an index is kept which also permits a type of direct access. The records are stored sequentially by primary key values and there is an index built over the primarykey field.

An index is a set of index value, address pairs. Indexing associates a set of objects to a set pf orderable quantities, that are usually smaller in number or their properties. Thus an index is a mechanism for faster search. Although the indices and the data blocks are kept physically, they are logically distinct.

Asequential file that is indexed on its primary key is called an index sequential file. The index allows for random access to records, while the sequential storage of the records of the file provides easy access to the sequential records. An addition feature of this file system is the overflow area. The overflow area provides additional space for record addition without the need to create.

Advantage of ISAM indexes:

- 1. Because the whole structure is ordered to a large extent, partial (like ty%) and range(between 12 and 29) based retrievals can often benefit from the use of this type of index.
- 2. ISAM isgoodforstatictablesbecausethereareusuallyfewerindexlevelsthanB-tree

3. Ingeneraltherearefewer diskI/Osrequiredtoaccessdata,providedthereisnooverflow.

DisadvantageofISAMindexes:

- **1.** ISAMisstillnotasquickassomeinhashfile organization
- **2.** Overflowcanbearealprobleminhighlyvolatiletable.

Hashedfileorganization:

Hashing is the most common form of purely random access to a file or database. It is also used to access columns that do not have a index as an optimization technique. Hash functions calculate the address of the page in which the record is tobe stored based on one ormore fields in the record. The records in a hash file appear randomly distributed across the available space. It requires some hashing algorithmand the technique. Hashing algorithmconverts a primarykey value into a record address.

Advantageofhashedfileorganization:

- 1. Insertionorsearchonhashkeyisfast.
- 2. Bestifequalitysearch isneeded onhash key

Disadvantageofhashedfile organization:

- **1.** Itisacomplexfileorganization method
- 2. searchisslow
- 3. Itsuffers fromdiskspaceoverhead
- 4. Unbalanced bucketsdegrade performance
- 5. Rangesearchisslow

LECTURE-13:INDEX

TypesofIndexes:

- Indexingmechanismsusedtospeedupaccesstodesired data.
 - E.g., author catalogin library
- SearchKey-attributetoset of attributes used to lookup records in a file.
- Anindexfileconsists of records (called indexentries) of the form
- Indexfilesaretypicallymuchsmaller thantheoriginalfile
- Twobasickindsofindices:
 - $\circ \quad \textbf{Ordered indices:} searchkeys are stored in sorted order$
 - **Hash indices:**search keys are distributed uniformlyacross "buckets" using a "hash function".

OrderedIndices

Indexingtechniquesevaluatedonbasisof:

1. Inan**orderedindex,**indexentriesarestoredsortedonthesearchkeyvalue. author catalog in library.

E.g.,

- 2. **Primaryindex:** inasequentiallyorderedfile, the indexwhosesearchkeyspecifiesthe sequential order of the file.
 - a. Alsocalledclusteringindex
 - **b.** Thesearchkeyofaprimary index isusually butnot necessarilytheprimarykey.
- 3. Secondaryindex:anindexwhosesearchkeyspecifiesanorderdifferent from the sequential order of the file. Also called non-clustering index.
- 4. Index-sequentialfile: orderedsequentialfilewithaprimaryindex.

There aretwotypes of ordered indices that we can use:

Denseindex:

An index record appears for every search-key value in the file. In a dense primary index, the index record contains the search-key value and a pointer to the firstdata record with thatsearch-key value. The rest of the records with the same search key-value would be stored sequentially after the first record, since, because the index is a primary one, records are sorted on the same search key. Dense index implementations may store a list of pointers to all records with the same search-key value; doing so is not essential for primary indices.

Sparseindex:

An index record appears for only some of the search-key values. As is true in dense indices, each index record contains a search-key value and a pointer to the first data record with that search-key value. To locate arecord, we find the indexentry with the largest search-key value that isless than or equal to the search-key value for which we are looking. We start at the record pointed to by that index entry, and follow the pointers in the file until we find the desired record.

PrimaryIndexes:

Aprimary index is an ordered file whose records are of fixed length with two fields. The first field is ofthe same datatypes as the ordering keyfiled of the data file and the second field is a pointerto adisk block- ablockaddress. Theoreming keyfield is called the primary key of the data file. There

is one index entry in the index file for each block in the data file. Each index entryhas the value of the primary key field for the first record in a block and a pointer to other block as its two fields values.

The first recordofeachblockofthedatafile isknownasanchor recordoftheblock. Primary index is an example of non dense index

Abhay			
Amit			
Asit			
Bapi Bikash			
Williem Wood			7
	Abhay Amit Asit Bapi Bikash William Wood	Abhay Amit Asit Bapi Bikash William Wood	Abhay Amit Amit Asit Bapi Bikash William Wood

(Primaryindexesontheorderingkeyfield)

A major problem with primary index as with any ordered file- is insertion and deletion of records. With a primary Index, the problem is compounded because if we attempt to insert a record in its correct position in the data file, we not onlyhave to move records to make space for the new record but also have to change some index entries because moving records willchange the anchor records of some blocks.

LECTURE-14:ClusteringIndex

ClusteringIndexes:

If the records of a file are physically ordered on a non key field that does not have a distinct value for each record, that filed is called the clustering filed of the file. We can create a different type of index called clustering index to speed up retrieval of records that have the same value for the clustering field. This differs from primary index, which requires that the ordering field of the data file have a distinct value for each record.

Aclustering index is also an ordered file with two fields, the first field is of the same type as the clustering field of the data file and the second field is block pointer. There is one entryin the clustering index for each distinct value of the clustering field, containing that value and a pointer to the first block in the data file that has a record with that value for its clustering field.

Aclusteringindexisanother exampleofnondenseindex.



SecondaryIndexes:

A secondary index also is an ordered file with two fields and as in the other indexes, the second filed is a pointerto a disk block. The first field isofthe same type assome nonordering field of the data file. The field on which the secondary index is constructed is called an indexing field of the file. Whether its values are distinct for every record or not.



Therecanbemanysecondaryindexesandhenceindexingfields forthesamefile.

PrimaryandSecondaryIndices:

Secondaryindices havetobe dense.

- 1. Indicesoffersubstantialbenefitswhensearchingforrecords.
- 2. Whenafileismodified, every index on the file must be updated, Updating indices imposes overhead on database modification.
- 3. Sequential scan using primary index is efficient, but a sequential scan using a secondary index is expensive
 - a. eachrecordaccessmayfetchanewblockfromdisk

LECTURE-15:B⁺TreeIndex

B+-TreeIndexFiles:

 $B{+}{-}tree indices are an alternative to indexed-sequential files.\\$

- 1. Disadvantage of indexed-sequential files: performance degrades as file grows, since many overflow blocks get created.Periodic reorganization fentire file is required.
- 2. Advantage of B+-tree index files:automatically reorganizes itself with small, local, changes, in the face of insertions and deletions. Reorganization of entire file is not required to maintain performance.
- 3. DisadvantageofB+-trees: extrainsertionand deletionoverhead, spaceoverhead.
- 4. AdvantagesofB+-treesoutweighdisadvantages, and they are used extensively.

AB+-tree is a rooted tree satisfying the following properties:

- 1. Allpathsfromroottoleafareofthesamelength
- 2. Eachnode that is not a root or a leaf has between [n/2] and *n* children.
- 3. Aleafnodehasbetween [(n-1)/2] and n-1 values
- 4. Specialcases:
 - a. If the root is not a least 2 children.
 - b. If the root is a leaf (that is, there are no other nodes in the tree), it can have between 0 and (n-1) values.

B⁺-TreeNodeStructure

Typicalnode



Piarepointerstochildren(fornon-leafnodes)orpointers to records or buckets of records (for leaf nodes). Thesearch-keysinanodeareordered

 $K_{1}K_{2}K_{3}\dots K_{n-1}$

ExampleofaB+-tree



B+-TreeFileOrganization

- 1. Index file degradation problem is solved by using B+-Tree indices. Data file degradation problem is solved by using B+-Tree File Organization.
- 2. TheleafnodesinaB+-treefileorganizationstorerecords, instead of pointers.
- 3. Since records are larger thanpointers, the maximumnumber of records that can be stored in a leaf node is less than the number of pointers in a non leaf node.

- 4. Leafnodesarestillrequired tobehalffull.
- 5. Insertion and deletion are handled in the same way as insertion and deletion of entries in a B+-tree index.



B-TreeIndexFiles

- **1.** SimilartoB+-tree,butB-treeallowssearch-keyvaluestoappearonlyonce;eliminates redundant storage of search keys.
- 2. Search keys in non leaf nodes appear nowhere else in the B-tree; an additional pointer field for each search key in a non leaf node must be included.
- 3. GeneralizedB-treeleafnode



P_1	B_1	<i>K</i> ₁	<i>P</i> ₂	<i>B</i> ₂	<i>K</i> ₂	 P_{m-1}	B_{m-1}	K_{m-1}	P_m

(b)

4. Nonleafnode-pointersBiarethebucketorfilerecord pointers.

B-TreeIndexFileExample:



AdvantagesofB-Treeindices:

- > Mayuselesstree nodesthana correspondingB+-Tree.
- Sometimespossibletofindsearch-keyvaluebeforereachingleafnode.

DisadvantagesofB-Treeindices:

- Onlysmallfractionofallsearch-keyvaluesare foundearly
- Non-leafnodesarelarger,sofan-outisreduced. Thus,B-Treestypicallyhave greater depth than corresponding B+-Tree

> Insertionanddeletionmorecomplicatedthanin B+-Trees

> ImplementationisharderthanB+-Trees.

Typically,advantagesofB-Treesdonotoutweighdisadvantages.

LECTURE-16:HashFileOrganization

HashFileOrganization

In a **hash file organization**, we obtain the address of the disk block containing a desired record directly by computing a function on the search-key value of the record.In our description of hashing, we shall use the term**bucket** to denote aunit of storage that can store one or more records. Abucket is typically a disk block, but could be chosen to be smaller or larger than a disk block.

Formally, let K denote the set of all search-key values, and let B denote the set of all bucket addresses. Ahash function h is a function from K to B. Let h denote a hash function.

To insert a record with search key Ki, we compute h(Ki), which gives the address of the bucket for that record. Assume for now that there is space in the bucket to store record. Then, the record is stored in that bucket.

To perform a lookup on a search-key value Ki, we simply compute h(Ki), then search the bucket with that address. Suppose that two search keys, K5 and K7, have the same hash value; that is,h(K5) = h(K7). If we perform a lookup on K5, the bucket h(K5) contains records with search-key values K5 and records with search key values K7. Thus, we have to check the search-key value of every record in the bucket to verify that the record is one that we want. Deletion is equally straightforward. If the search-key value of the record to be deleted is Ki, we compute h(Ki), then search the corresponding bucket for that record, and delete the record from the bucket.

Hash Indices

Hashing can be used not only for file organization, but also for index-structure creation. A **hash index** organizes the search keys, with their associated pointers, into a hash file structure. We construct a hash index as follows. We apply a hash function on a search key to identify a bucket, and storethe keyand its associated pointers in the bucket (or in overflow buckets). Figureshows a secondary hash index on the *account* file, for the search key *account-number*. The hash function in the figure computes the sum of the digits of the account number modulo 7. The hashindex has sevenbuckets, eachofsize2 bucket sizes). Oneofthebuckets hasthreekeys mapped it, so it has an overflow bucket. In this example, *account-number* is a primary key for *account*, so each search key has onlyone associated pointer. In general, multiple pointers can be associated with each key.



Figure 12.23 Hash index on search key account-number of account file.

We use the term*hash index* to denote hash file structures as wellas secondaryhash indices. Strictly speaking, hash indices are only secondary index structures. A hash index is never needed as a primaryindexstructure, since, ifa file itselfisorganized by hashing, there is no need for a separate hash index structure on it. However, since hash file organization provides the same direct access to records that indexing provides, we pretend that a file organized by hashing also has a primary hash index on it.

Advantages of hashing:

- 1. Exactkeymatchesareextremelyquick
- **2.** Hashing is verygood long keys, or those with multiple columns, provided the complete key value is provided for the query
- **3.** This organization usually allows for the allocation of disk space so a good deal of disk management is possible
- 4. Nodiskspaceisusedbythis indexingmethod

Disadvantagesof hashing:

- 1. It becomes difficult to predict overflow because the working of the hashing algorithm will not be visible to the DBA
- 2. Nosortingofdataoccurseither physicallyorlogicallysosequentialaccessispoor
- 3. Thisorganizationusuallytakesalotofdiskspacetoensurethatnooverflowoccurs.

LECTURE-17:QueryProcessing

QueryProcessing

Query Processing would mean the entire process or activity which involves query translation into low level instructions, queryoptimization to save resources, cost estimation or evaluation of query, and extraction of data from the database.

Goal: To find an efficient QueryExecution Plan for a given SQLquery which would minimize the cost considerably, especially time.

Cost Factors: Disk accesses [which typically consumes time], read/write operations [which typically needs resources such as memory/RAM].

Themajorsteps involved inquery processing are depicted in the figure below;



Let usdiscuss the whole process with an example. Let us consider the following two relations as the example tables for our discussion;

Employee(Eno, Ename, Phone) Proj_Assigned(Eno,Proj_No,Role,DOP)

where,

Eno is Employee number,

EnameisEmployeename,

Proj_No isProjectNumber inwhichanemployeeisassigned, Role is

the role of an employee in a project,

DOPisdurationoftheprojectin months.

With this information, let us write a query to find the list of all employees who are working in a project which is more than 10 months old.

SELECTEname FROMEmployee,Proj_Assigned WHEREEmployee.Eno =Proj_Assigned.Eno ANDDOP>10;

Input:

A query writtenin SQLisgiven asinputto thequery processor.For our case, letusconsider the SQL query written above.

Step1:Parsing

In this step, the parse roft hequery processor module checks the syntax of the query, the user's privileges to execute the query, the table names and attribute names, etc. The correct table names, and the parse of the table names of ta

attribute names and the privilege of the users can be taken from the system catalog (data dictionary).

Step2:Translation

If we have written availed query, then it is converted from highlevellanguage SQL to low level instruction in Relational Algebra.

Forexample,ourSQLquerycanbeconvertedinto aRelationalAlgebraequivalentas follows;

 $\pi_{\text{Ename}}(\sigma_{\text{DOP}>10\Lambda\text{Employee}.\text{Eno}=\text{Proj}_{\text{Assigned}.\text{Eno}}(\text{Employee}XProf}_{\text{Assigned}}))$

Step3:Optimizer

Optimizer uses the statistical data stored as part of data dictionary. The statistical data are information about the size of the table, the length of records, the indexes created on the table, etc. Optimizer also checks for the conditions and conditional attributes which are parts of the query.

Step4:ExecutionPlan

A query can be expressed in many ways. The query processor module, at this stage, using the information collected in step 3 to find different relational algebra expressions that are equivalentand return the result of the one which we have written already.

Forourexample, the query written in Relational algebra can also be written as the one given below;

 $\pi_{\text{Ename}}(\text{Employee} \Join_{\text{Eno}}(\sigma_{\text{DOP}>10}(\text{Prof}_\text{Assigned})))$

So far, we have got two execution plans. Only condition is that both plans should give the same result.

Step5:Evaluation

Though we got many execution plans constructed through statistical data, though they return same result (obvious), they differ in terms of Time consumption to execute the query, or the Space required executing the query. Hence, it is mandatory choose one plan which obviously consumes less cost.

At this stage, we choose one execution plan of the severalwe have developed. This Execution plan accesses data from the database to give the final result.

In our example, the second plan may be good. In the first plan, we join two relations (costly operation) thenapplythe condition(conditions are considered as filters) on the joined relation. This consumes more time as well as space.

In the second plan, we filter one of the tables (Proj_Assigned) and the result is joined with the Employee table. This join may need to compare less number of records. Hence, the second plan is the best (with the information known, not always).

QueryTree



Usedinqueryrepresentationusedinparsing.

QueryOptimization:

A single query can be executed through different algorithms or re-written in different forms and structures. Hence, the questionofqueryoptimizationcomes into the picture –Whichofthese forms or pathways is the most optimal? The queryoptimizer attempts determine the most efficient way to execute a given query by considering the possible query plans.

Therearebroadly twowaysa query canbeoptimized:

1. Analyze and transform equivalent relational expressions: Try to minimize the tuple and column counts of the intermediate and final query processes (discussed here).

2. Using different algorithms for each operation: These underlying algorithms determine how tuples are accessed from the data structures they are stored in, indexing, hashing, data retrieval and hence influence the number of disk and block accesses (discussed in query processing).

Analyzeandtransformequivalentrelational expressions

Here, we shall talk about generating minimal equivalent expressions. To analyze equivalent expression, listed are a set of equivalence rules. These generate equivalent expressions for a query written in relational algebra. To optimize a query, we must convert the query into its equivalent form as long as an equivalence rule is satisfied.

1. Conjunctive selection operations can be written as a sequence of individual selections. This is called a sigma-cascade.

Explanation: Applying condition intersection is expensive. Instead, filter outtuples

satisfying condition(inner selection) and then apply condition(outer selection) to the then resulting fewer tuples. This leaves us with less tuples to process the second time. This can be extended for two or more intersecting selections. Since we are breaking a single condition into a series of selections or cascades, it is called a "cascade".

2. Selectioniscommutative.

Explanation:condition is commutative in nature. This means, it does not matter whether we applyfirst orfirst. Inpractice, it is better and more optimalto applythat selection first which yields a fewer number of tuples. This saves time on our outer selection.

3. All following projections can be omitted, only the first projection is required. This is called a pi-cascade.

Explanation: A cascade or a series of projections is meaningless. This is because in the end, we are only selecting those columns which are specified in the last, or the outermost projection. Hence, it is better to collapse all the projections into just one i.e. the outermost projection.

4. SelectionsonCartesian Productscanbere-writtenasTheta Joins.

Equivalence1

Explanation: The cross product operation is known to be very expensive. This is because it matches each tuple of E1 (total m tuples) with each tuple of E2 (total ntuples). This yields m*n entries. If we apply a selection operation after that, we would have to scan through m*n entries to find the suitable tuples which satisfy the condition

. Instead of doing all of this, it is more optimal to use the Theta Join, a join specifically designed to select only those entries in the cross product which satisfy the Theta condition, without evaluating the entire cross product first.

• Equivalence2

Explanation: Theta Join radically decreases the number of resulting tuples, so if we

applyanintersection of both the join conditions i.e. and into the Theta Joinitself,

conditionoutsideunnecessarily

wegetfewerscanstodo.Ontheotherhand,a increases the tuples to scan.

5. ThetaJoinsarecommutative.

Explanation: Theta Joins are commutative, and the query processing time depends to some extent whichtable is used as the outerloop and whichone is used as the inner loopduring the join process (based on the indexing structures and blocks).

6. Joinoperationsare associative.

• NaturalJoin

Explanation: Joins are all commutative as wellas associative, so onemustjoin those two tables first which yield less number of entries, and then apply the other join.

Theta Join

Explanation: ThetaJoinsareassociativeintheabovemanner, where involves attributes from only E2 and E3.

7. Selectionoperationcanbe distributed.

Equivalence1

Explanation: Applying a selection after doing the Theta Join causes all the tuples returned by the Theta Join to be monitored after the join. If this selection contains attributes from only E1, it is better to apply this selection to E1 (hence resulting in a fewer number of tuples) and then join it with E2.

Equivalence2

 $\label{eq:explanation:This can be extended to two selection conditions, and, where$

Theta1 contains the attributes of only E1 and contains attributes of only E2. Hence, we can individually apply the selection criteria before joining, to drastically reduce the number of tuples joined.

8. Projection distributesovertheTheta Join.

Equivalence1

Explanation: The idea discussed for selection can be used for projection as well. Here, if L1 is a projection that involves columns of only E1, and L2 another projection that involves the columns ofonlyE2, then it is betterto individuallyapplythe projections on both the tables before joining. This leaves us with a fewer number of columns on either side, hence contributing to an easier join.

• Equivalence2

Explanation: Here, when applying projections L1 and L2 on the join, where L1contains columns of only E1 and L2 contains columns of only E2, we can introduce another column E3 (which is common between both the tables). Then, we can apply projections L1 and L2 on E1 and E2 respectively, along with the added column L3. L3 enables us to do the join.

9. UnionandIntersectionare commutative.

Explanation: Unionandintersectionarebothdistributive;wecanencloseanytablesin parantheses according to requirement and ease of access.

10. UnionandIntersectionare associative.

Explanation: Unionandintersectionarebothdistributive;wecanencloseanytablesin parantheses according to requirement and ease of access.

11. Selectionoperationdistributesovertheunion,intersection,anddifferenceoperations. Explanation: Insetdifference,weknowthatonlythosetuplesareshownwhichbelongto

table E1 and do not belong to table E2. So, applying a selection condition on the entire set difference is equivalent to applying the selection condition on the individual tables and then applying set difference. This willreduce the number of comparisons in the set difference step.

12. **Projectionoperation distributesovertheunionoperation.**

Explanation: Applying individual projections before computing the union of E1 and E2 is more optimal than the left expression, i.e. applying projection after the union step.



LECTURE-18:EvaluationofExpressions

Evaluation of Expressions

- Sofar:wehaveseenalgorithmsforindividualoperations
- Alternativesforevaluatinganentireexpressiontree

□ Materialization:generate results of an expression whose inputs are relations or are already computed, materialize (store) it on disk.Repeat.

- **Pipelining**: passontuplestoparentoperationsevenasanoperationisbeingexecuted
- Westudyabovealternativesinmoredetail

Materialization

- **Materializedevaluation:**evaluate one operation at time,startingatthelowest-level.Use intermediate results materialized into temporary relations to evaluate next-level operations.
- E.g., infigure below, compute and store

then compute the store its join with *customer*, and finally compute the projections on *customer*name.

σ balance < 2500 σ balance < 2500

account

- Materializedevaluationisalwaysapplicable
- Costofwriting resultstodisk and reading thembackcan bequitehigh
 - □Ourcostformulasforoperationsignore costofwritingresultstodisk,so
 - Overallcost=Sumofcostsofindividualoperations+
 - costofwritingintermediateresultstodisk
- Doublebuffering:usetwooutputbuffersforeach operation, when one is full write itto disk while the other is getting filled

Pipelining

• **Pipelined evaluation :**evaluate several operations simultaneously, passing the results of oneoperation on to the next.

□E.g., in previous expression tree, don't store result of instead, pass tuples directly to the join..Similarly, don't store result of join, pass tuples directly to projection.

- Muchcheaper thanmaterialization:noneedtostoreatemporaryrelationtodisk.
- Pipeliningmaynotalways bepossible-e.g.,sort,hash-join.
- Forpipeliningtobeeffective, use evaluational gorithms that generate output tuples even as tuples are received for inputs to the operation.
- Pipelinescanbeexecuted intwo ways: demanddriven and producerdriven
- Indemanddrivenor lazyevaluation
 - $\label{eq:system} \square \ system repeated ly requests next tuple from top level operation$
 - Each operation requestsnext tuple from children operations as required, in order to output its next tuple

□Inbetweencalls,operationhastomaintain"**state**"soitknowswhattoreturnnext

- $\label{eq:constraint} \Box Each operation is implemented as an iterator implementing the following operations$
 - open()
 - E.g. file scan: initialize file scan, store pointer to beginning of file as state
 - E.g.merge join: sort relations and storepointerstobeginning of sorted relations as state
 - next()
 - E.g.forfilescan:Outputnexttuple,andadvanceandstorefilepointer
 - E.g.formerge join:continuewithmerge fromearlierstatetill next output tuple is found.Save pointers as iterator state.
 - close()
- Inproduce-drivenor **eager**pipelining
 - $\label{eq:construct} \Box Operators produce tuples eagerly and pass the mupt otheir parents$
 - Buffermaintainedbetweenoperators,childputstuplesinbuffer,parent removes tuples from buffer
 - if buffer is full, child waitstill there is space in the buffer, and then generates more tuples
 - System schedules operations that have space in output buffer and can process more input tuples

EvaluationAlgorithmsforPipelining

- Somealgorithmsarenotabletooutputresultsevenastheygetinputtuples
 - □E.g.mergejoin,orhashjoin
 - These resultinintermediate resultsbeingwrittentodiskandthenreadbackalways
- Algorithm variants are possible to generate (at least some) results on the fly, as input tuples are read in
 - \Box E.g. hybrid hash join generates output tuples even as probe relation tuples in the inmemory partition (partition 0) are read in
 - **Pipelined join technique**: Hybrid hash join, modified to buffer partition 0 tuples of both relations in-memory, reading themas they become available, and output results of any matches between partition 0 tuples
 - When a new r_0 tuple is found, match it with existing s_0 tuples, output matches, and save it in r_0
 - Symmetricallyfors₀tuples

ComplexJoins

- Joininvolvingthreerelations: loan depositor customer
- **Strategy1.** Compute*depositor customer*; useresulttocompute*loan* (*depositor customer*)
- **Strategy2.**Computer *loan depositor*first, andthenjointheresultwith*customer*.
- **Strategy3.**Perform the pair of joins at once. Build and index on *loan* for *loan-number*, and on *customer* for *customer-name*.
 - \Box For each tuple *t* in *depositor*, look up the corresponding tuples in *customer* and the corresponding tuples in *loan*.
 - \Box Each tupleof *deposit* is examined exactly once.
 - Strategy 3 combines two operations into one special-purpose operation that is more efficient than implementing two joins of two relations.

Module-2:LECTURE-19

RelationalAlgebra:

Basicoperations:

- 1. Selection (σ) Selects a subset of rows from relation.
- 2. $Projection(\pi)$ Selectsasubsetofcolumnsfrom relation.
- 3. *Cross-product*(×)Allowsusto combinetworelations.
- 4. Set-difference()Tuplesinrelation.1, but not inrelation.2.
- 5. Union(U)Tuples inreln.1andinreln. 2.
- 6. $Rename(\rho)$ UsenewnamefortheTablesorfields.

Additionaloperations:

- 7. Intersection(\cap), *Join*(\bowtie), *Division*(\div): Notessential, but(very!) useful.
- Sinceeachoperationreturnsarelation, operationscanbe*composed*!(Algebrais "closed".)

Projection

- > Deletesattributesthatarenotinprojection list.
- Schemaofresult containsexactly the fields inthe projection list, with the same names that they had in the (only) input relation. (Unary Operation)
- Projectionoperatorhastoeliminate duplicates!(as itreturnsarelationwhichisaset)
 - Note:realsystemstypicallydon't do duplicateeliminationunlesstheuserexplicitly asks for it.(Duplicate values may be representing different real world entityor relationship).

Example:ConsidertheBOOKtable:

Acc-No	Title	Author
100	"DBMS"	"Silbershatz"
200	"DBMS"	"Ramanuj"
300	"COMPILER"	"Silbershatz"
400	"COMPILER"	"Ullman"
500	"OS"	"Sudarshan"
600	"DBMS"	"Silbershatz"

$\pi_{\text{Title}}(\text{BOOK}) =$

Title	
"DBMS"	
"COMPILER"	
"OS"	

Selection

- Selectsrowsthatsatisfyselection condition.
- Noduplicatesinresult
- Schemaofresultidenticaltoschemaof(only)inputrelation.
- *Result* relationcanbethe *input* foranotherrelationalalgebraoperation!(*Operator* composition.)

Example:Fortheexamplegivenabove:

 $\sigma_{Acc-no>300}(BOOK)=$

Acc- No	Title	Author
400	"COMPILER	"Ullman"
500	"OS"	"Sudarshan"

$\sigma_{Title="DBMS"}(BOOK) =$

Acc- No	Title	Author
100	"DBMS"	"Silbershatz"
200	"DBMS"	"Ramanuj"
600	"DBMS"	"Silbershatz"

$$\pi_{\text{Acc-no}}(\sigma_{\text{Title="DBMS"}}(\text{BOOK})) =$$

Acc-
No
100
200
600

Union, Intersection, Set-Difference

- > Alloftheseoperationstaketwo inputrelations, which must be *union-compatible*:
 - Samenumber offields.
 - Corresponding'fieldshavethesame type.
- ➤ What is the schema of result?

Consider:

Borrower		
Cust-	Loan-no	
name		
Ram	L-13	
Shyam	L-30	
Suleman	L-42	

Depositor	
Cust-name	Acc-no
Suleman	A-100
Radheshyam	A-300
Ram	A-401

Listofcustomerswhoareeitherborroweror depositoratbank= $\pi_{Cust-name}(Borrower)U$ $\pi_{Cust-name}(Depositor)=$

Cust-name	
Ram	
Shyam	Customerswho arebothborrowersanddepositors= π_{Cust} -
Suleman	name(Borrower) $\cap \pi_{\text{Cust-name}}(\text{Depositor}) =$
Radeshyam	
Cust-	
name	
Ram	
Suleman	

Customerswhoareborrowersbutnotdepositors= $\pi_{Cust-name}(Borrower)$ (Depositor)=

 $\Box \pi_{\text{Cust-name}}$

$Cartesian-ProductorCross-Product(S1 \times R1)$

- EachrowofS1is pairedwitheachrowof R1.
- ResultschemahasonefieldperfieldofS1andR1, withfieldnames`inherited'ifpossible.
- Consider the borrower and loant ables as follows:

Borrower:						
Cust-name	Loan-no					
Ram	L-13					
Shyam	L-30					
Suleman	L-42					

L	loan:	
	Loan-no	Amount
	L-13	1000
	L-30	20000
	L-42	40000

Crossproduct ofBorrowerandLoan,Borrower×Loan=

Borrower.Cust-	Borrower.Loan-	Loan.Loan-	Loan.Amount
name	no	no	
Ram	L-13	L-13	1000
Ram	L-13	L-30	20000
Ram	L-13	L-42	40000
Shyam	L-30	L-13	1000
Shyam	L-30	L-30	20000
Shyam	L-30	L-42	40000
Suleman	L-42	L-13	1000
Suleman	L-42	L-30	20000
Suleman	L-42	L-42	40000

Therenameoperation can be used to rename the fields to avoid confusion when two field names are same in two participating tables:

For example the statement, $\rho_{\text{Loan-borrower}(\text{Cust-name},\text{Loan-No-1},\text{ Loan-No-2},\text{Amount})}$ (Borrower × Loan) results into- Anew Table named Loan-borrower is created where it has four fields which are renamed as Cust-name, Loan-No-1, Loan-No-2 and Amount and the rows contains the same data as the cross product of Borrower and Loan.

Loan-borrower:

Cust-	Loan-No-1	Loan-	Amount
name		No-2	
Ram	L-13	L-13	1000
Ram	L-13	L-30	20000
Ram	L-13	L-42	40000
Shyam	L-30	L-13	1000
Shyam	L-30	L-30	20000
Shyam	L-30	L-42	40000
Suleman	L-42	L-13	1000
Suleman	L-42	L-30	20000
Suleman	L-42	L-42	40000

RenameOperation:

It can be used in two ways:

- $\rho_x(E)$ return the result of expression Einthetable named x.
- $\rho_{x(A_1,A_2,...,A_n)}(E)$ return the result of expression E in the table named x with the attributes renamed to A₁, A₂,..., A_n.
- It'sbenefit canbeunderstoodbythesolutionofthequery"Findthe largest accountbalance in the bank"

It can be solved by following steps:

- Findout therelationofthosebalanceswhicharenotlargest.
- ConsiderCartesionproductofAccountwithitselfi.e.Account×Account
- Compare the balances of first Account table with balances of second Account table in the product.
- Forthatweshouldrenameoneoftheaccounttablebysomeothernametoavoidthe confusion

Itcanbedonebyfollowingoperation

 $\Pi_{Account.balance}(\sigma_{Account.balance} < d.balance}(Account \times \rho_d(Account))$

- Sotheaboverelationcontainsthebalanceswhicharenotlargest.
- Subtract this relation from the relation containing all the balances i.e. $\Pi_{balance}(Account)$. So

the final statement for solving above query is

 $\Pi_{balance}(Account)$ - $\Pi_{Account,balance}(\sigma_{Account,balance}(Account \times \rho_d(Account))$

LECTURE-20

AdditionalOperations

NaturalJoin($S_1 \bowtie R_1$)

- FormsCartesianproduct of its two arguments, performs selection forcing equality on those attributes that appear in both relations
- For example consider Borrower and Loan relations, the natural join between them Borrower ⋈ Loan willautomaticallyperform the selection on the table returned by Borrower × Loan which force equality on the attribute that appear in both Borrower and Loan i.e. Loan-no and also will have only one of the column named Loan-No.
- That means Borrower \bowtie Loan = $\sigma_{Borrower.Loan-no}$ (Borrower×Loan).
- Thetablereturnedfromthiswillbeasfollows:

Eliminaterowsthatdoesnotsatisfy these lection criteria " $\sigma_{Borrower.Loan-no}$ " from Borrower ×Loan=

Borrower.Cust-	Borrower.Loan-	Loan.Loan-	Loan.Amount	
name	no	no		
Ram	L-13	L-13	1000	
Ram	L-13	L-30	20000	
Ram	L-13	L-42	40000	
Shyam	L-30	L-13	1000	
Shyam	L-30	L-30	20000	
Shyam	L-30	L-42	40000	
Suleman	L-42	L-13	1000	
Suleman	L-42	L-30	20000	
Suleman	L-42	L-42	40000	

Andwillremoveoneofthecolumn namedLoan-no.

ie Borrower ⋈ Loan_

Cust-name	Loan-no	Amount
Ram	L-13	1000
Shyam	L-30	20000
Suleman	L-42	40000

DivisionOperation:

- denotedby÷ isusedfor queriesthatincludethephrase"for all".
- Forexample"Findcustomerswhohasanaccountinallbranchesinbranchcity Agra". This query can be solved by following statement.

 $\Pi_{Customer-name,branch-name}(\overset{Optotic use of total of$

• The division operations can be specified by using only basic operations as follows:

Let r(R) and s(S) be given relations for schema Rand Swith $S \subseteq R_r \div s =$

 $\Pi_{\text{R-S}}(r) \text{ - } \Pi_{\text{R-S}}((\Pi_{\text{R-S}}(r) \times s) \text{ - } \Pi_{\text{R-S},S}(r))$

LECTURE-21

TupleRelationalCalculus

Relationalalgebra isanexampleofprocedurallanguagewhiletuplerelationalcalculus isa nonprocedural query language.

Aqueryisspecifiedas:

 $\{t | P(t)\}, i.eit is the set of all tuples tsuch that predicate Pistrue fort.$

The formulaP(t)is formedusing atoms which uses the relations, tuples of relations and fields of tuples and following symbols

 \in (belongs to),<,>, \leq , \geq , \neq ,=, (comparison operators)

These atoms can then be used to form formulas with following symbols

∀ (universal qualifier generally called "for all")

∃ (existential qualifier generally called "there exists")

 \land (and), \lor (or),7(not)

Forexample: herearesomequeriesand awaytoexpressthemusing tuple calculus:

- FindtheloannumberforeachloanofanamountgreaterthatRs1200.
 {t| ∃ s ∈ Loan(t[loan-number] = s[loan-number] ∧ s[amount] >1200}

 ○ Findthenamesofallthecustomerswhohavealoan fromtheSadarbranch. {t | ∃ s ∈ Borrower (t[customer-name] = s[customer-name] ∧ ∃ u ∈ Loan (u[loan-number] = s[loan-number

 \land u[branch-name] = "Sadar"))}

o Findallcustomerswhohavealoan,anaccount,orbothatthebank
 {t|∃s∈Borrower(t[customer-name] = s[customer-name]))
 V∃u∈Depositor(t[customer-name] = u[customer-name])}

- Findonlythosecustomerswhohavebothanaccountandaloan.
 {t| ∃ s ∈ Borrower (t[customer-name] = s[customer-name])
 ∧ ∃ u ∈ Depositor (t[customer-name] = u[customer-name])}
- o Findallcustomerswhohaveanaccountbutdonothaveloan.
 {t| ∃u ∈ Depositor (t[customer-name] = u[customer-name]) ∧
 ¬∃ s ∈ Borrower (t[customer-name] = s[customer-name])}
- Findallcustomerswhohaveanaccountatallbranches locatedinAgra {t | ∀w ∈ Branch(w[branch-city] = "Agra" ⇒ ∃ s ∈ Depositor (t[customer-name] = s[customer-name]

```
\land \exists u \in Account (u[account-number] = s[account-number]
```

 $\land u[branch-name] = w[branch-name])))$

DomainRelationalCalculus

- 1. Domain relational calculus is another nonprocedural language for expressing database aueries.
- 2. Aqueryisspecifiedas:

 $\{\langle x_1, x_2, \dots, x_n \rangle | P(x_1, x_2, \dots, x_n)\}$ where x_1, x_2, \dots, x_n represents domain variables. Prepresent a predicate formula as in tuple calculus

- Sincethedomainvariables are referred in place of tuples the formuladoesn't refer the fields oftuplesrathertheyreferthe domainvariables.
- Forexamplethequeries indomaincalculusarementionedas follows: •
 - Findthebranch-name, loan-number and amount for loans over Rs1200. $\{<b, l, a > | < b, l, a > \in Loan \land a > 1200\}$
 - FindtheloannumberforeachloanofanamountgreaterthatRs1200. $\{ < l > | \exists b,a(<b, l, a > \in Loan \land a > 1200 \}$
 - Find the names of all the customers who have a loan from the Sadar branch and find the loan amount

 $\{ < c, a > | \exists l (< c, l > \in Borrower \}$ $\exists b (\langle b, l, a \rangle \in Loan \land b = "Sadar")) \}$

- Findnamesofallcustomerswhohavealoan,anaccount,orbothattheSadar Branch $\{ < c > | \exists l (< c, l > \in Borrower \land \exists b, a (< b, l, a > \in Loan \land b = "Sadar") \}$ $\forall \exists a (\langle c, a \rangle \in Depositor \land \exists b, n (\langle b, a, n \rangle \in Account \land b = "Sadar")) \}$
- o Findonlythosecustomerswhohavebothanaccountandaloan. $\{ <c > | \exists l(<c, l > \in Borrower) \land \exists a(<c, a > \in Depositor) \}$
- o Findallcustomerswhohaveanaccountbutdonothaveloan. $\{t \mid \exists a(\langle c, a \rangle \in Depositor) \land \forall \exists l(\langle c, l \rangle \in Borrower)\}$
- o Findallcustomerswhohaveanaccountatallbranches locatedin Agra $\{ < c > | \forall x, y, z (< x, y, z > \in Branch) \land y = "Agra" \Rightarrow \}$

OuterJoin.

Outerjoinoperationisanextensionofjoinoperationtodealwithmissinginformation

Suppose that we have following relational schemas: Employee(employee-name, street, city) Fulltime-works(employee-name,branch-name,salary)

Asnapshot of these relations is as follows:

employee-	street	city
name		
Ram	M GRoad	Agra
Shyam	New Mandi	Mathura
	Road	
Suleman	BhagatSingh	Aligarh
	Road	

Fulltime-works

Employee:

employee-	branch-	salary
name	name	
Ram	Sadar	30000
Shyam	SanjayPlace	20000
Rehman	Dayalbagh	40000

 $Suppose we want \ complete information of the full time employees.$

- The natural join (^{Employee ⋈ Fulltime-works})will result into the loss of information for Suleman and Rehman because they don't have record in both the tables (left and right relation). The outer join will solve the problem.
- Threeformsofouterjoin:
 - Left outer join(^{¬™}): the tuples which doesn't match while doing natural join from left relation are also added in the result putting null values in missing field of right relation.
 - **Right outer join** (**C**): the tuples which doesn't match while natural join from right relation are also added in the result putting null values in missing field of left relation.
 - **Full outer join** $(\square M \square)$: include both of the left and right outer joins i.e. adds the tuples which did not match either in left relation or right relation and put null inplace of missing values.
- Theresultforthreeformsofouterjoinareasfollows:

<u>-J</u>				
employee- street		City	branch-	salary
name			name	
Ram	M GRoad	Agra	Sadar	30000
Shyam	New Mandi	Mathura	Sanjay	20000
	Road		Place	
Suleman	BhagatSingh	Aligarh	Null	Null
	Road			

Leftioin: Employee ⊐⊠ Fulltime-works_

Rightjoin: Employee ⋈⊏ Fulltime-works_

employee-	street		city	branch-	salary
name				name	
Ram	M GRo	ad	Agra	Sadar	30000
Shyam	New	Mandi	Mathura	Sanjay	20000
	Road	*		Place	
Rehman	null		null	Dayalbagh	40000

Fulljoin: Employee $\Box \bowtie \Box$ Fulltime-works_

employee-	street	city	branch-	salary
name			name	
Ram	M GRoad	Agra	Sadar	30000
Shyam	New Mandi	Mathura	Sanjay	20000
	Road		Place	
Suleman	BhagatSingh	Aligarh	null	null
	Road			
Rehman	null	null	Dayalbagh	40000



LECTURE-22

StructuredQueryLanguage(SQL)

Introduction

Commercialdatabasesystemsuse moreuserfriendlylanguageto specifythequeries. SQL is

the most influential commercially marketed product language.

Othercommercially used languages are QBE, Quel, and Datalog.

BasicStructure

- ThebasicstructureofanSQLconsistsofthreeclauses: select, from and where.
- **select:** it corresponds to the projection operation of relational algebra. Used to list the attributes desired in the result.
- **from**:correspondstotheCartesianproductoperationofrelationalalgebra.Usedtolist the relationstobescannedintheevaluationoftheexpression
- where:correspondstotheselectionpredicateoftherelationalalgebra. Itconsistsofa predicate involving attributes of the relations that appear in the **from** clause.
- AtypicalSQLqueryhastheform:

selectA₁,A₂,...,A_nfro

m *r*₁, *r*₂,..., *r*_m**where**

Р

- \circ A_irepresentsanattribute
- \circ r_jrepresentsa relation
- Pisa predicate
- Itisequivalenttofollowingrelationalalgebraexpression:
- $\prod_{A_1,A_2,\ldots,A_n} \left(\sigma_{\mathsf{P}} \left(\mathbf{r}_1 \times \mathbf{r}_2 \times \ldots \times \mathbf{r}_m \right) \right)$

[Note: The words marked in dark in this text work as keywords in SQL language. For example "select", "from" and "where" in the above paragraphare shown in bold font to indicate that they are keywords]

SelectClause

 $Let us see some simple queries and use of {\it select} clause to express them in SQL.$

- Findthenamesofall branchesintheLoanrelation
 - selectbranch-name

fromLoan

• Bydefault the **select** clause includesduplicatevalues. If we want to force the elimination of duplicates the **distinct** keyword is used as follows:

selectdistinctbranch-name

fromLoan

• Theallkeywordcanbeusedtospecify*explicitly* thatduplicatesarenotremoved. Evenif we not use all it means the same so we don't require all to use in select clause.

selectallbranch-name

fromLoan

• Theasterisk"*"canbeusedtodenote"allattributes".ThefollowingSQLstatementwill selectandallthe attributesofLoan.

select*

fromLoan

• The arithmetic expressions involving operators, +, -, *, and / are also allowed in **select** clause. The followingstatement willreturn the amount multiplied by 100 for the rows in Loan table.

selectbranch-name,loan-number,amount*10fromLoan.

WhereClause

- Findallloannumbers forloans madeat"Sadar"branchwithloanamountsgreaterthanRs 1200.
 - selectloan-number
 - **from**Loan
 - wherebranch-name="Sadar"andamount>1200
- where clause uses logival connectives and, or, and not
- operandsofthelogicalconnectivescanbeexpressionsinvolvingthecomparisonoperators <,<=,>,>=,=,and<>.
- **between**canbeused tosimplifythecomparisons
 - selectloan-number
 - **from**Loan
 - where amount between 90000 and 100000

FromClause

- The **from** clause by itself defines a Cartesian product of the relations in the clause.
- Whenanattribute is present inmore than one relation they can be referred as *relationname.attribute-name* to avoid the ambiguity.
- Forallcustomerswhohaveloanfromthebank,findtheirnamesandloannumbers selectdistinctcustomer-name,Borrower.loan-number fromBorrower,Loan
 - whereBorrower.loan-number=Loan.loan-number

TheRenameOperation

- Usedfor renamingbothrelationsbothrelationsand attributes in SQL
- Useasclause:old-nameasnew-name
- Findthenamesandloannumbersofthecustomers whohavealoanatthe"Sadar" branch.

selectdistinctcustomer-name, borrower.loan-numberasloan-id

fromBorrower,Loan

whereBorrower.loan-number=Loan.loan-numberand

branch-name= "Sadar"

wecannowrefertheloan-numberinsteadbythenameloan-id.

 $\bullet \quad For all customers who have a loan from the bank, find their names and loan - numbers.$

selectdistinctcustomer-name,T.loan-number

fromBorrowerasT,LoanasS

whereT.loan-number=S.loan-number

• Findthenamesofallbranchesthathaveassetsgreaterthanatleastonebranch locatedin "Mathura".

selectdistinctT.branch-name

frombranchasT, branchasS

where T.assets>S.assets and S.branch-city= "Mathura"

StringOperation

- Twospecialcharactersareusedforpatternmatching instrings:
 - Percent(%):The%charactermatchesanysubstring
 - Underscore(_):The _ charactermatchesanycharacter
- "%Mandi":willmatchwiththestringsendingwith"Mandi" viz. "RajaKimandi","Peepal Mandi"
- "___"matchesanystring ofthree characters.
- Findthenamesofallcustomerswhosestreetaddress includesthesubstring "Main"

selectcustomer-name

fromCustomer

where customer-streetlike "% Main%"

SetOperations

- union, intersect and except operations are set operations available in SQL.
- Relations participating in any of the set operation mustbe compatible; i.e.they musthave the same set of attributes.
- UnionOperation:
 - Findallcustomershavinga loan,anaccount,orbothatthebank (select customer-namefrom Depositor)

union

- (selectcustomer-name fromBorrower)
- It will automatically eliminate duplicates.
- If we want to retain duplicates unional can be used (select customer-name from Depositor) unional
 - (selectcustomer-namefromBorrower)
- IntersectOperation
 - $\circ \ \ \, \mbox{Findall customers who have both an account and a loan at the bank}$
 - (select customer-name from Depositor)
 - intersect
 - (selectcustomer-name fromBorrower)
 - If we want to retail all the duplicates (selectcustomer-namefromDepositor) intersect all
 - $({\bf select} customer-name {\bf from} Borrower)$
- ExceptOpeartion
 - Findallcustomerswho haveanaccount butno loanatthebank (select customer-name from Depositor)
 - except

(selectcustomer-namefromBorrower)

- If we want to retain the duplicates:
 - (selectcustomer-namefromDepositor) except all (selectcustomer-namefromBorrower)

AggregateFunctions

- Aggregate functionsarethose functionswhichtakeacollectionofvaluesas input andreturn a single value.
- SQLoffers5builtinaggregatefunctions-
 - Average:avg
 - Minimum:**min**
 - Maximum:**max**
 - Total:**sum**
 - Count:count
- Theinput to **sum** and **avg** must be a collection of numbers but others may have collections of non-numeric data types as input as well
- FindtheaverageaccountbalanceattheSadar branch selectavg(balance)
$from {\rm Account}$

wherebranch-name="Sadar"

Theresultwillbeatablewhichcontainssinglecell(onerowandonecolumn)having numerical value corresponding to average balance of all account at sadar branch.

- **groupby**clauseisusedtoformgroups,tupleswiththesamevalueonallattributesinthe **groupby**clauseareplacedinonegroup.
- Findtheaverageaccountbalanceateachbranch
 - selectbranch-name,avg(balance)
 - **from**Account
 - groupbybranch-name
- Bydefaulttheaggregatefunctionsincludethe duplicates.
- **distinct**keywordisused toeliminateduplicates inanaggregate functions:
- Findthenumberofdepositorsfor eachbranch
 - selectbranch-name,count(distinctcustomer-name)
 - fromDepositor,Account
 - whereDepositor.account-number=Account.account-number
 - groupbybranch-name
- having clause is used to state condition that applies to groups rather than tuples.
- Findtheaverageaccountbalanceateachbranchwhereaverageaccountbalanceismore than Rs. 1200
 - selectbranch-name,avg(balance)
 - from Account
 - groupbybranch-name
 - havingavg(balance) >1200
- Countthenumber of tuples in Customertable
 - select count(*)
 - fromCustomer
- SQLdoesn'tallow**distinct**with**count**(*)
- When where and having are both present in a statement where is applied before having.

NestedSubqueries

Asubqueryisaselect-from-where expression that is nested within another query.

Set Membership

Theinand not inconnectives are used for this type of subquery.

``Find all customers who have both a loan and an account at the bank", this query can be written using nested subquery form as follows

selectdistinctcustomer-name
fromBorrower
wherecustomer-namein(selectcustomer-name

fromDepositor)

• Selectthenamesofcustomerswhohavealoanatthebank,andwhosenamesareneither "Smith"nor "Jones"

selectdistinctcustomer-name

fromBorrower

```
where customer-name not in ("Smith", "Jones")
```

SetComparison

Find the names of all branches thathave assets greater than those of at least one branch located in Mathura

selectbranch-name
fromBranch
whereasstets>some(select assets

$\mathbf{from} Branch$

wherebranch-city="Mathura")

- 1. Apartfrom>someotherscomparisoncould be<some,<=some,=some,=some,<>some.
- **2.** Findthenames of allbranches thathave assets greater than thatof each branchlocatedin Mathura

selectbranch-name

fromBranch

whereasstets>all(selectassets

fromBranch

wherebranch-city= "Mathura")

Apartfrom>allotherscomparisoncouldbe <all,<= all,=all,=all,

<>all.

Views

InSQLcreateviewcommandisusedtodefineaviewasfollows:

create viewvas<queryexpression>

 $where <\!\! query expression \!\!>\!\! is any legal query expression and v is the view name.$

Theviewconsistingofbranchnamesandthenamesofcustomerswho haveeitheran account or a loan at the branch. This can be defined as follows:

createviewAll-customeras
(selectbranch-name,customer-name
fromDepositor,Account
whereDepositor.account-number=account.account-number)
union

(selectbranch-name,customer-name fromBorrower,Loan whereBorrower.loan-number=Loan.loan-number)

- Theattributesnames maybespecified explicitly within a set of roundbracket after then ame of view.
- Theviewnamesmaybeusedasrelationsinsubsequent queries. Using theview Allcustomer

FindallcustomersofSadarbranch

```
selectcustomer-name
```

```
fromAll-customer
```

```
wherebranch-name= "Sadar"
```

Acreate-viewclausecreatesaview definition inthedatabasewhichstaysuntila command - dropviewview-name- isexecuted.

ModificationofDatabase

Deletion

✤ InSQLwecandeleteonlywholetupleand notthevaluesonanyparticular attributes. The command is as follows:

deletefromrwhereP.

wherePisapredicateandr is a relation.

- deletecommand operatesononlyonerelationatatime.Examplesareasfollows:
- ✤ Delete alltuplesfromtheLoanrelation

deletefromLoan

o DeletealloftheSmith'saccountrecord

delete from Depositor

wherecustomer-name="Smith"

- Deleteallloanswith loanamountsbetweenRs1300 and Rs1500.
 - **deletefrom**Loan
 - whereamountbetween1300and 1500
- o Deletetherecordsofallaccountswith balancesbelowtheaverageat the bank

deletefromAccount

wherebalance<(selectavg(balance)

from Account)

Insertion

InSQL we either specify a tuple to be inserted or write a query whose result is a set of tuples to be inserted. Examples are as follows:

Insert anaccountofaccount numberA-9732attheSadarbranchhaving balance of Rs 1200

insertintoAccount

values("Sadar", "A-9732", 1200) thevalues are specified in the order in which the corresponding attributes are listed in the relation schema.

SQLallowstheattributesto bespecifiedaspartofthe insertstatement

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

values("A-9732", "Sadar", 1200)

 $insert into {\it Account} (branch-name, account-number, balance)$

values("Sadar","A-9732",1200)

Provide for all loan customers of the Sadar branch a new Rs 200 saving account foreachloanaccounttheyhave. Where loan-numberserveastheaccount number for these accounts.

insertintoAccount selectbranch-name,loan-number,200 fromLoan wherebranch-name="Sadar"

Updates

Used to change a value in a tuple without changing all values in the tuple. Suppose that annual interest payments are being made, and all balances are to be increased by 5 percent. **update**Account setbalance=balance*1.05 Suppose that accounts with balances over Rs10000 receive 6 percent interest, whereas all others receive 5 percent. **update**Account setbalance=balance*1.06 where balance > 10000 update Account setbalance=balance*1.05 wherebalance<=10000 DataDefinitionLanguage **Data Types in SQL char**(n): fixed lengthcharacterstring, lengthn. varchar(n):variablelengthcharacterstring, maximumlengthn. int:aninteger. smallint:asmallinteger. numeric(p,d):fixedpoint number,pdigits(plusasign),anddofthepdigitsare to right of the decimal point. real,doubleprecision:floatingpointanddoubleprecisionnumbers. float(n):afloatingpointnumber, precisionat leastndigits. date:calendardate;fourdigits foryear,twoformonthandtwofordayofmonth. time:timeofdaynhoursminutesandseconds. Domainscan bedefined as createdomainperson-namechar(20). thedomainnameperson-namecanbeusedto definethetypeofanattribute just like builtin domain.

SchemaDefinition in SQL

 $\label{eq:createtablecommand} createtable command is used to define relations.$

createtabler(A1D1,A2D2,...,AnDn,

<integrityconstraint1>,

· · · ,

<integrityconstraintk>)

where risrelation name, each Aiisthename of attribute, Diisthedomain type of

values of Ai. Several types of integrity constraints are available to define in SQL.

Integrity Constraints which are allowed in SQL are

primarykey(Aj1,Aj2,...,Ajm) and **check**(P)wherePisthepredicate.

drop table command is used to remove relations from database. **altertable**command isusedto addattributestoanexistingrelation **alter table** r **add** A D itwilladdattributeAofdomaintypeDinrelationr. **altertabler drop**A itwillremovetheattributeAofrelationr.

IntegrityConstraints

- IntegrityConstraintsguard againstaccidentaldamagetothedatabase.
- Integrityconstraintsare predicatespertainingtothe database.
- DomainConstraints:
- PredicatesdefinedonthedomainsareDomainconstraints.
- Simplest Domainconstraintsaredefinedbydefiningstandarddatatypesoftheattributes like Integer, Double, Float, etc.
- We can define domain sby **createdomain** clause also we can define the constraints on such domains as follows:

create domainhourly-wage **numeric**(5,2)

- **constraint**wage-value-testcheck(value>=4.00)
- Sowecanusehourly-wageasdatatypeforanyattributewhere DBMS willautomatically allow onlyvaluesgreaterthanorequalto4.00.
- OtherexamplesfordefiningDomainconstraintsareasfollows:

createdomainaccount-numberchar(10)

constraintaccount-number-null-testcheck(valuenotnull)

create domain account-type char(10)

constraint account-type-test

check(valuein("Checking", "Saving"))

By using the later domain of two above the DBMS will allow only values for any attribute having type as account-type i.e. Checking and Saving.

- ReferentialIntegrity:
- Foreign Key: If two table R and S are related to each other, K1 and K2 are primarykeys of the two relations also K1 is one of the attribute in S. Suppose we want that every row in S must have a corresponding row inR, then we define the K1 inS as foreign key. Example in our original database of librarywe had a table for relationBORROWEDBY, containing two fields Card No. and Acc. No. . Every row of BORROWEDBY relation must have corresponding row in USER Table having same Card No. and a row in BOOK table having sameAcc. No. . Then we will define the Card No. and Acc. No. inBORROWEDBY relation as foreign keys.
- Inother waywe cansaythat everyrowofBORROWEDBYrelation must refer to some row in BOOK and also in USER tables.
- Suchreferentialrequirementinonetableto anothertableiscalled ReferentialIntegrity.

QuerybyExample(QBE)

Query by Example (QBE) is amethod of querycreation that allows the user tosearch for documents based on an example in the form of a selected text string or in the form of a document name or a list of documents. Because the QBE systemformulates the actual query, QBE is easier to learn than formal query languages, such as the standard Structured Query Language (<u>SQL</u>), while still enabling powerful searches.

Selectionsin QBE

QBE uses skelet on tables to represent table name and field names like:

Tablename	Field1	Field2	

 $\label{eq:Forselection} For selection, Poperator along with variable name/constant name is used to display one or more fields.$

Example1:

Consider the relation: student (name, <u>roll</u>, marks)

The following query can be represented as:

SQL:selectnamefromstudentwheremarks>50;

student	name	roll	marks
	P.X		>50

HereX isa constant;alternativelywecanuse _X asavariable.

Example2:

Fortherelation given above

Thefollowing querycanberepresented as:

SQL:select*fromstudentwhere marks>50and marks<80;

student	name	roll	marks
Р.			_X



We can use condition box to represent complex conditions.

Example3:

Fortherelationgiven above

Thefollowing querycanberepresented as:

SQL:selectname, rollfromstudentwheremarks<50ormarks>80;

student	name	roll	marks
	P.A	P.B	<50
	P.A	P.B	>80

ORoperationretrievesresultsinmultiplerows.

Example4:

(JoinsinQBE)

Consider the following tables:

Student (name, <u>rol</u>l, marks)

Grades (roll, grade)

Thefollowing querycanberepresented as:

SQL:select s.name,g.gradefromStudents,Gradesgwheres.roll=g.rollands.marks>50; Uses two skeleton tables:

Student	name	roll	marks
	P.A	_X	>50
And			
Grades	roll	grade	
	_X	P.B	

Insertionsin QBE:

Uses operator I.onthetable. Example:Consider thefollowingqueryonStudenttable SQL:insertintostudentvalues('abc',10,60);

Student	name	roll	marks
I.	abc	10	60

Multipleinsertionscanberepresented byseparaterowsinskeletontable.

DeletionsinQBE:

UsesoperatorD.onthetable.

Example:ConsiderthefollowingqueryonStudenttable SQL: delete from student where marks=0;

Student	name	roll	marks
D.			0

Multipledeletionscanberepresentedbyseparaterowsinskeletontable. Deletions without any condition can truncate the entire table.

UpdationinQBE:

UsesoperatorU.onthetable.

 $Example: Consider the following query on Studenttable \ SQL:$

update student set mark=50 where roll=40;

Student	name	roll	marks
U.			50
		40	

RELATIONALDATABASEDEGIN

Database design is a process in which you create a logical data model for a database, which store data of a company. It is performed after initial database study phase in the database life cycle. You use normalization technique to create the logical data model for a database and eliminate data redundancy.

Normalization also allows you to organize data efficiently in a database and reduce anomalies during dataoperation.Variousnormal forms, suchas first, second and third canbe applied to create a logical data model for a database. The second and third normal forms are based on partial dependencyand transitivitydependency. Partialdependencyoccurs when a rowoftable is uniquely identified by one column that is a part of a primary key. A transitivity dependency occurs when a non keycolumn is uniquely identified by values in another non-keycolumn of a table.

DatabaseDesignProcess:

We can identify six main phases of the databased esign process:

- 1. Requirementcollectionandanalysis
- 2. Conceptualdatabase design
- 3. Choice of aDBMS
- 4. Datamodelmapping(logicaldatabasedesign)
- 5. Physicaldatabasedesign
- 6. Databasesystemimplementationand tuning

1. RequirementCollectionandAnalysis

Before we can effectively design a data base we must know and analyze the expectation of the users and the intended uses of the database in as much as detail.

2. ConceptualDatabaseDesign

Thegoalforthisphaseistoproduceaconceptualschemaforthedatabasethatis independent of a specific DBMS.

- Weoftenuseahigh leveldatamodelsuchER-modelduring thisphase
- We specify as manyofknowndatabase application on transactions as possible using a notation that is independent of any specific dbms.
- Often the dbms choice already made for the organization the intent of conceptual designstill to keep, it as free as possible from implementation consideration.

3. Choice of aDBMS

The choice of dbms is governed by a no. of factors some technical other economic and still other concerned with the politics of the organization.

The economics and organizational factors that offer the choice of the dbms are:

Softwarecost, maintenance cost, hardwarecost, databasecreation and conversion cost, personnel cost, training cost, operating cost.

4. Datamodelmapping(logicaldatabasedesign)

Duringthisphase, we map the conceptual schema from the highle veldat a model used on

phase2intoadatamodelofthechoice dbms.

5. Physicaldatabsedesign

During this phase we design the specification for the database in terms of physical storagestructure ,record placement and indexes.

6. **Databasesystemimplementationandtuning** Duringthisphase,thedatabaseandapplicationprogramsareimplemented,testedand eventually deployed for service.

InformalGuidelinesforRelationDesign

Want to keep the semantics of the relation attributes clear. The information in a tuple should represent exactly one fact or an entity. The hidden or buried entities are what we want to discover and eliminate.

- Designarelationschemasothatitiseasytoexplainitsmeaning.
- Do notcombineattributesfrommultipleentitytypesandrelationshiptypesinto asingle relation. Use a view if you want to present a simpler layout to the end user.
- Arelationschemashouldcorrespond toonentitytypeorrelationship type.
- Minimizeredundantinformationintuples, thus reducing update anomalies
- If anomalies are present, try to decompose the relation into two ormore to represent the separate facts, or document the anomalies well for management in the applications programs.

Minimizetheuseof**null**values.Nulls havemultipleinterpretations:

- Theattributedoesnotapplyto thistuple
- Theattributevalueisunknown
- Theattributevalue isabsent
- Theattributevaluemightrepresentanactualvalue

If nulls are likely (non-applicable) then consider decomposition of the relation into two or more relations that hold only the non-null valued tuples.

• Donotpermitthecreationofspurioustuples

Toomuchdecomposition of relations into smaller one smay also lose information or generate erroneous information

• Besure that relations can be logically joined using natural join and the result doesn't generate relationships that don't exist

FunctionalDependencies

FD'sareconstraintsonwell-formed relations and represent a formalismon the infrastructure of relation.

Definition: A *functional dependency* (FD) on a relation schema \mathbf{R} is a <u>constraint</u> $X \rightarrow Y$, where X and Y are subsets of attributes of \mathbf{R} .

Definition:anFD isarelationshipbetweenanattribute"Y"andadeterminant (1ormoreother attributes) "X"suchthat for agivenvalueofadeterminantthevalueoftheattributeisuniquely defined.

- Xisa determinant
- XdeterminesY
- Yis functionallydependentonX
- X→Y
- X→YistrivialifY⊆X

Definition: AnFD $X \rightarrow Y$ is *satisfied* in an instance **r** of **R** iffor <u>every</u> pair of tuples, *t* and *s*: if *t* and *s* agree on all attributes in *X* then they must agree on all attributes in *Y*

Akeyconstraint is a specialkind offunctional dependency: all attributes of relation occur on the right-hand side of the FD:

• $SSN \rightarrow SSN, Name, Address$

ExampleFunctionalDependencies

LetRbe NewStudent(*stuId*,*lastName*,*major*,*credits*,*status*,*socSecNo*)

FDs in R include

- *{stuId}→{lastName}*, butnotthereverse
- {stuId}→{lastName,major,credits,status,socSecNo,stuId}
- {socSecNo}→{stuId,lastName,major,credits,status,socSecNo}
- {credits} →{status}, butnot{status} →{credits}

$ZipCode \rightarrow AddressCity$

• 16652isHuntingdon'sZIP

$ArtistName \rightarrow BirthYear$

• Picassowasbornin1881

 $Autobrand \rightarrow Manufacturer, Enginetype$

• PontiacisbuiltbyGeneralMotorswithgasoline engine

Author, Title \rightarrow PublDate

• Shakespeare'sHamletwaspublishedin1600

TrivialFunctionalDependency

TheFDX \rightarrow Yis*trivial* ifset {Y}isasubset ofset{X}

Examples: IfAand B are attributes of R,

- $\{A\} \rightarrow \{A\}$
- $\{A,B\} \rightarrow \{A\}$
- $\{A,B\} \rightarrow \{B\}$
- $\{A,B\} \rightarrow \{A,B\}$

arealltrivial FDs and will not contribute to the evaluation of normalization.

FDAxioms

Understanding: FunctionalDependenciesarerecognized by analysis of the real world; no automation or algorithm. Finding or recognizing them are the database designer's task.

FDmanipulations:

- Soundness--noincorrectFD'saregenerated
- **Completeness**-- allFD'scanbe generated

AxiomName	Axiom	Example
Reflexivity	ifaissetofattributes,b⊆a, thena→b	$SSN, Name \rightarrow SSN$
Augmentation	ifa \rightarrow bholdsandc isasetof attributes, then ca \rightarrow cb	SSN→Namethen SSN,Phone→Name,Phone
Transitivity	ifa \rightarrow bholdsandb \rightarrow cholds, then a \rightarrow c holds	$SSN \rightarrow ZipandZip \rightarrow City$ then $SSN \rightarrow City$
UnionorAdditivity *	ifa \rightarrow banda \rightarrow choldsthen a \rightarrow bc holds	$SSN \rightarrow NameandSSN \rightarrow ZipthenSSN \rightarrow Name, Zip$
Decompositionor Projectivity*	ifa \rightarrow bc holdsthena \rightarrow b and a \rightarrow c holds	$SSN \rightarrow Name, Zipthen SSN \rightarrow Name and SSN \rightarrow Zip$
Pseudotransitivity*	Pseudotransitivity *ifa \rightarrow bandcb \rightarrow dholdthen ac \land Address \rightarrow ProjectandP Address,Date \rightarrow Amount	
(NOTE)	$ab \rightarrow cdoesNOT$ implya $\rightarrow c$ and $b \rightarrow c$	

*Armstrong'sAxioms(basicaxioms)

CLOSUREOFASETOFFUNCTIONALDEPEDENCIES

GivenarelationalschemaR, afunctionaldependencies fon Rislogically implied by a set of functional dependencies F on R if every relation instance r(R) that satisfies F also satisfies f.

The closure of F, denoted by F⁺, is the set of all functional dependencies logically implied by F. The closure of F can be found by using a collection of rules called **Armstrong axioms**. **Reflexivityrule:** If A *is a set of attributes and B is subset or equal to A, then A* \rightarrow Bholds. **Augmentation rule:** If A \rightarrow B holds and C is a set of attributes, then CA \rightarrow CB holds **Transitivity rule:** If A \rightarrow B holds and B \rightarrow C holds, then A \rightarrow C holds. **Unionrule:** If A \rightarrow B holds and A \rightarrow Cthen A \rightarrow BCholds **Decomposition rule:** If A \rightarrow BC holds, then A \rightarrow C holds. **Pseudotransitivityrule:** If A \rightarrow Bholds and BC \rightarrow Dholds, then AC \rightarrow Dholds.

Suppose we are given a relation schema R=(A, B, C, G, H, I) and the set of function dependencies

 $\{A \rightarrow B, A \rightarrow C, CG \rightarrow H, CG \rightarrow I, B \rightarrow H\}$

WelistseveralmembersofF⁺here:

1. $A \rightarrow H$, since $A \rightarrow Band B \rightarrow H$ hold, we apply the transitivity rule.

2. CG \rightarrow HI.SinceCG \rightarrow HandCG \rightarrow I,theunionruleimpliesthatCG \rightarrow HI

3. AG \rightarrow I, since A \rightarrow CandCG \rightarrow I, the pseudotransitivity rule implies that AG \rightarrow I holds

AlgorithmofcomputeF+:

To compute the closure of a set of functional dependencies F: F+=F

repeat

foreachfunctionaldependencyfinF+
 applyreflexivityand augmentationrulesonf
 addtheresultingfunctionaldependenciestoF+
 foreachpairoffunctionaldependenciesf1 and f2inF+
 iff1 andf2canbecombined usingtransitivity
 thenaddtheresultingfunctionaldependencytoF+

until*F*+ doesnot change anyfurther

large.

LOSSLESS DECOMPOSITION

Adecomposition of a relationscheme R < S, F > into the relationschemes Ri(1 <= i <= n) is said to be a lossless join decomposition or simply lossless if for every relation R that satisfies the FDs in F, the natural join of the projections or R gives the original relation R, i.e,

DEPEDENCYPRSERVATION:

GivenarelationschemeR<S,F>whereFistheassociatedsetof functional dependencieson the attributesinS,R is decomposed into the relationschemesR1,R2,...Rnwith the fdsF1,F2...Fn, then this decomposition of R is dependency preserving if the closure of F'(where F'=F1 U F2 U ... Fn) Example:

 $Let R(A,B,C) ANDF = \{A \rightarrow B\}. Then the decomposition of RintoR1(A,B) and R2(A,C) is lossless because the FD \{A \rightarrow B\} is contained in R1 and the common attribute A is a key of R1.$

Example:

Let R(A,B,C) AND $F=\{A\rightarrow B\}$. Then the decomposition of R into R1(A,B) and R2(B,C) is not lossless because the commonattribute B does not functionally determine either Aor C. i.e, it is not a key of R1 or R 2.

Example:

Let R(A,B,C,D) and F={A \rightarrow B,A \rightarrow C, C \rightarrow D,}. Then the decomposition of R into R1(A,B,C) with the FD F1={ A \rightarrow B , A \rightarrow C }and R2(C,D) with FD F2={ C \rightarrow D} . In this decomposition all the original FDs can be logically derived from F1 and F2, hence the decomposition is dependency preserving also . the common attribute C forms a keyof R2. The decomposition is lossless.

Example:

Let R(A,B,C,D) and F={A \rightarrow B, A \rightarrow C, A \rightarrow D,}. Then the decomposition of R into R1(A,B,D) with the FD F1={A \rightarrow B, A \rightarrow D } and R2(B,C) with FD F2={ } is lossy because the commonattribute B is not a candidate key of either R1 and R2.

 $In addition, the fds A \rightarrow C is not implied by any fds R1 or R2. Thus the decomposition is not dependency preserving.$

Fullfunctionaldependency:

Given a relational scheme R and an FD $X \rightarrow Y$, Y is fully functional dependent on X if there is no Z, where Z is a proper subset of X such that $Z \rightarrow Y$. The dependency $X \rightarrow Y$ is left reduced, there being no extraneous attributes attributes in the left hand side of the dependency.

Partialdependency:

 $Given a relation dependencies F defined on the attributes of Rand Kasac and idatekey, if X is a proper subset of KandifF \models X \rightarrow A, then A is said to be partial dependent on K$

Primeattributeandnonprime attribute:

A attribute A in a relation scheme R is a **prime attribute** or simply **prime** if A is part of any candidate key of the relation. If A is not a part of any candidate key of R, A is called a nonprime attribute or simply **non prime**.

Trivialfunctionaldependency:

 $AFDX \rightarrow Yissaid$ to be atrivial functional dependency if Yissubset of X.

LECTURE-29

Normalization

While designing a database out of an entity-relationship model, the main problem existing in that "raw" database is redundancy. Redundancy is storing the same data item in more one place. A redundancy creates several problems like the following:

- 1. Extrastoragespace:storingthesamedatainmanyplacestakeslargeamountofdiskspace.
- 2. Enteringsamedatamore than onceduring data insertion.
- 3. Deletingdatafrommore than one placed uring deletion.
- 4. Modifyingdatainmore than one place.
- 5. Anomaliesmayoccurinthedatabaseifinsertion, deletion, modification et carenodone properly. It creates inconsistency and unreliability in the database.

To solve this problem, the "raw" database needs to be normalized. This is a step bystep process of removing different kinds of redundancy and anomaly at each step. At each step a specific rule is followed to remove specific kind of impurity in order to give the database a slim and clean look.

Un-NormalizedForm(UNF)

If a table contains non-atomic values at each row, it is said to be in UNF. An **atomic value** is something that can not be further decomposed. A **non-atomic value**, as the name suggests, can be further decomposed and simplified. Consider the following table:

Emp-Id	Emp-Name	Month	Sales	Bank-Id	Bank-Name
E01	AA	Jan	1000	B01	SBI
		Feb	1200		
		Mar	850		
E02	BB	Jan	2200	B02	UTI
		Feb	2500		
E03	CC	Jan	1700	B01	SBI
		Feb	1800		
		Mar	1850		
		Apr	1725		

In the sample table above, there are multiple occurrences of rows under eachkeyEmp-Id.Although considered to be the primary key, Emp-Id cannot give us the unique identification facility for any single row. Further, each primarykeypoints to a variable length record (3 for E01, 2 for E02 and 4 for E03).

FirstNormalForm(1NF)

A relation is said to be in 1NF if it contains no non-atomic values and each row can provide a unique combination of values. The above table in UNF can be processed to create the following table in 1NF.

	Emp-Name	Month	Sales	Bank-Id	Bank-Name
Emp-Id	_				
E01	AA	Jan	1000	B01	SBI
E01	AA	Feb	1200	B01	SBI
E01	AA	Mar	850	B01	SBI
E02	BB	Jan	2200	B02	UTI
E02	BB	Feb	2500	B02	UTI
E03	CC	Jan	1700	B01	SBI
E03	CC	Feb	1800	B01	SBI
E03	CC	Mar	1850	B01	SBI
E03	CC	Apr	1725	B01	SBI

As you can see now, each row contains unique combination of values. Unlike in UNF, this relation contains only atomic values, i.e. the rows can not be further decomposed, so the relation is now in 1NF.

SecondNormalForm(2NF)

Arelation is said to be in2NF f if it is alreadyin1NFand eachand everyattribute fullydepends on the primary key of the relation. Speaking inversely, if a table has some attributes which is not dependent on the primary key of that table, then it is not in 2NF.

Let us explain. Emp-Id is the primary key of the above relation. Emp-Name, Month, Sales and Bank-Name all depend upon Emp-Id. But the attribute Bank-Name depends on Bank-Id, which is not the primary key of the table. So the table is in 1NF, but not in 2NF. If this position can be removed into another related relation, it would come to 2NF.

Emp-Id	Emp-Name	Month	Sales	Bank-Id
E01	AA	JAN	1000	B01
E01	AA	FEB	1200	B01
E01	AA	MAR	850	B01
E02	BB	JAN	2200	B02
E02	BB	FEB	2500	B02
E03	CC	JAN	1700	B01
E03	CC	FEB	1800	B01
E03	CC	MAR	1850	B01
E03	CC	APR	1726	B01

Bank-Id	Bank-Name
B01	SBI
B02	UTI

After removing the portion into another relation we store lesser amount of data in two relations without any loss information. There is also a significant reduction in redundancy.

ThirdNormalForm(3NF)

Arelation is said to be in 3NF, if it is already in 2NF and there exists no **transitive dependency** in that relation. Speaking inversely, if a table contains transitive dependency, then it is not in 3NF, and the table must be split to bring it into 3NF.

Whatisatransitivedependency?Withinarelationifwesee $A \rightarrow B[BdependsonA]$ And $B \rightarrow C [CdependsonB]$ Thenwemayderive $A \rightarrow C[C dependsonA]$

Suchderiveddependencies holdwellinmostofthesituations. Forexampleifwehave Roll-Marks And

Marks \rightarrow Grade Thenwemaysafelyderive Roll \rightarrow Grade.

Thisthirddependencywasnotoriginallyspecifiedbutwehavederivedit.

The derived dependency is called a transitive dependency when such dependency becomes

improbable. For example we have been given $Roll \rightarrow City$

And City \rightarrow STDCode

If we try to derive Roll \rightarrow STDCode it becomes a transitive dependency, because obviously the STDCode of a city cannot depend on the roll number issued by a school or college. In such a case the relation should be broken into two, each containing one of these two dependencies: Roll \rightarrow City

And City→STDcode

Boyce-CodeNormalForm(BCNF)

A relationship is said to be in BCNF if it is already in 3NF and the left hand side of every dependency is a candidate key. Arelation which is in 3NF is almost always in BCNF. These could besame situation when a 3NF relation may not be in BCNF the following conditions are found true.

- 1. Thecandidatekeysarecomposite.
- 2. Therearemore than one candidate keys in the relation.
- 3. Therearesomecommonattributesintherelation.

ProfessorCode	Department	Head ofDept.	PercentTime
P1	Physics	Ghosh	50
P1	Mathematics	Krishnan	50
P2	Chemistry	Rao	25
P2	Physics	Ghosh	75
P3	Mathematics	Krishnan	100

Consider, as an example, the above relation. It is assumed that:

- 1. Aprofessor canworkinmorethanone department
- 2. Thepercentageofthetimehespendsineachdepartmentis given.
- 3. Eachdepartmenthasonlyone HeadofDepartment.

Therelationdiagramfortheaboverelationisgivenasthefollowing:



The given relation is in 3NF. Observe, however, that the names of Dept. and Head of Dept. are duplicated. Further, if Professor P2 resigns, rows 3 and 4 are deleted. We lose the information that Rao is the Head of Department of Chemistry.

Thenormalization of the relation is done by creating an ewrelation for Dept. and Head of Dept. and deleting Head of Dept. form the given relation. The normalized relations are shown in the following.

ProfessorCode	Department	PercentTime
P1	Physics	50
P1	Mathematics	50
P2	Chemistry	25
P2	Physics	75
P3	Mathematics	100

Department	Head ofDept.
Physics	Ghosh
Mathematics	Krishnan
Chemistry	Rao

See the dependency diagrams for these new relations.



FourthNormalForm (4NF)

Whenattributes ina relation have multi-valued dependency, further Normalization to 4NFand 5NF are required. Let us first find out what multi-valued dependency is.

A**multi-valueddependency** is a typical kind of dependency in which each and every attribute within a relation depends upon the other, yet none of them is a unique primary key.

We will illustrate this with an example. Consider a vendor supplying many items to many projects in an organization. The following are the assumptions:

- 1. Avendoriscapableofsupplyingmany items.
- 2. Aproject usesmanyitems.
- 3. Avendorsuppliestomanyprojects.
- 4. Anitemmaybesuppliedby manyvendors.

Amultivalueddependencyexistshere becausealltheattributesdependupontheotherand yet none of them is a primary key having unique value.

VendorCode	Item Code	ProjectNo.
V1	I1	P1
V1	I2	P1
V1	I1	P3
V1	I2	P3
V2	I2	P1
V2	I3	P1

V3	I1	P2
V3	I1	P3

The given relation has a number of problems. For example:

- 1. IfvendorV1hastosupplyto project P2, buttheitemisnot yet decided, thenarowwitha blank for item code has to be introduced.
- 2. The informationaboutitemI1isstoredtwicefor vendorV3.

Observe that the relation given is in 3NF and also in BCNF. It still has the problem mentioned above. The problem is reduced by expressing this relation as two relations in the Fourth Normal Form (4NF). Arelation is in 4NF if it has no more than one independent multi valued dependency or one independent multi valued dependency with a functional dependency.

The table can be expressed as the two 4NF relations given as following. The fact that vendors are capable of supplying certain items and that they are assigned to supply for some projects in independently specified in the 4NF relation.

Vendor-Supply		
VendorCode	ItemCode	
V1	I1	
V1 I2		
V2	I2	
V2 I3		
V3	I1	
Vendor-Project		
VendorCode ProjectNo		
V1 P1		

V I	PI	
V1	P3	
V2	P1	
V3	P2	

FifthNormalForm(5NF)

These relations still have a problem. While defining the 4NF we mentioned that all the attributes depend upon each other. While creating the two tables in the 4NF, although we have preserved the dependenciesbetweenVendorCodeandItemcode in the first tableandVendorCodeandItemcode in thesecond table, we have lost therelationship betweenItemCodeand Project No. If therewere a primary key then this loss of dependency would not have occurred. In order to revive this relationship we must add a new table like the following. Please note that during the entire processof normalization, this is the only step where a new table is created by joining two attributes, rather than splitting them into separate tables.

ProjectNo.	Item Code
P1	11
P1	12
P2	11
P3	11
P3	13

Let us finally summarize the normalization steps we have discussed so far.

Input Relation	Transformation	Output Relation
All	Eliminatevariablelengthrecord.Removemulti-attributelines intable.	1NF
Relations		
1NF	Removedependencyofnon-keyattributes onpartofamulti-attribute	2NF
Relation	key.	
2NF	Removed ependency of non-key attributes on other non-key attributes.	3NF
3NF	Removedependencyofanattributeofamultiattributekeyonan	BCNF
	attributeofanother(overlapping)multi-attributekey.	
BCNF	Removemore than one independent multi-valued dependency from	4NF
	relationbysplitting relation.	
4NF	Addonerelationrelatingattributeswithmulti-valueddependency.	

QUERYPROCESSING

Query processing includes translation of high-level queries into low-level expressions that can be used at the physical level of the file system, query optimization and actual execution of the query to get the result. It is a three-stepprocessthat consists of parsing and translation, optimization and execution of the query submitted by the user.



Query-processing Steps

Aquery is processed infourgeneral steps:

- 1. ScanningandParsing
- 2. QueryOptimizationorplanningtheexecutionstrategy
- 3. QueryCodeGenerator(interpretedorcompiled)
- 4. Executionintheruntimedatabaseprocessor

1. ScanningandParsing

When a query is first submitted (via anapplicationsprogram), it must be scanned and parsed to determine if the query consists of appropriate syntax.

Scanningis theprocessofconvertingthequerytext intoatokenizedrepresentation.

The tokenized representation is more compact and is suitable for processing bythe parser.

This representationmay be in a tree form.

The Parserchecksthetokenizedrepresentationforcorrectsyntax.

In this stage, checks are made to determine if columns and tables identified in the query exist in the database and if the query has been formed correctly with the appropriate keywords and structure. If the query passes the parsing checks, then it is passed on to the Query Optimizer.

2. QueryOptimizationorPlanning theExecutionStrategy

For any given query, there may be a number of different ways to execute it. Each operation in the query (SELECT, JOIN, etc.) can be implemented using one or more different *AccessRoutines*. For example, anaccess routine that employs anindex to retrieve some rows would be more efficient that an access routine that performs a full table scan.

The goal of the **query optimizer** is to find *reasonably efficient* strategy for executing the query (not quite what the name implies) using the access routines.

Optimizationtypicallytakesoneof

twoforms:*HeuristicOptimization*or*CostBasedOptimization*In**Heuristic Optimization**,the query execution is refined based on*heuristic rules* for reordering the individual operations.

With **Cost Based Optimization**, the overall cost of executing the query is systematically reduced by estimating the costs of executing several different execution plans.

3. QueryCodeGenerator(interpretedorcompiled)

Once the query optimizer has determined the execution plan(the specific ordering of access routines), the code generator writes out the actual access routines to be executed.

With an interactive session, the query code is interpreted and passed directly to the runtime database processor for execution.

It is also possible to compile the access routines and store them for later execution.

4. Executionintheruntimedatabaseprocessor

At this point, the query has been scanned, parsed, planned and (possibly)compiled. The runtime database processor then executes the access routines against the database. The results are returned to the application that made the query in the first place. Anyruntimeerrors are also returned.

Lecture-32

QueryOptimization

To enable the system toachieve (or improve) acceptable performance by choosing a better (if not the best) strategy during the processof a query. One of the great strengths to the relational database.

AutomaticOptimizationvs.HumanProgrammer

- 1. Agood automatic optimizer will have a wealth of information available toit that human programmers typically do not have.
- 2. An automatic optimizer can easily reprocess the original relational request when the organization of the database is changed. For a human programmer, reorganization would involve rewriting the program.
- 3. The optimizer is a program, and therefore is capable of considering literally hundreds of different implementation strategies for a given request, which is much more than a human programmer can.
- 4. The optimizerisavailabletoawiderangeofusers, in an efficient and cost-effective manner.

TheOptimizationProcess

- 1. Castthequeryintosomeinternalrepresentation, such as aquery trees tructure.
- 2. Converttheinternalrepresentationtocanonical form.

*Asubset (sayC) of a set of queries (sayQ) is said to be a set of canonicalforms for Q if and only if every query Q is equivalent to just one query in C.

During this step, some optimization is already achieved bytransforming the internal representation to a better canonical form.

Possibleimprovements

- a. Doingtherestrictions(selects)beforethe join.
- b. Reduce the amount of comparisons by converting a restriction condition to an equivalent condition in **conjunctivenormal form** that is, a condition consisting of a set of restrictions that areANDed together, where each restriction in turn consists of a set of simple comparisons connected only by OR's.
- c. Asequenceofrestrictions(selects)beforethejoin.
- d. Inasequenceofprojections, all butthelastcanbeignored.
- e. Arestrictionofprojectionisequivalenttoaprojectionofarestriction.
- f. Others
- 3. Choosecandidatelow-levelproceduresbyevaluatethe transformedquery.

*Access path selection: Consider the queryexpression as a series of basic operations (join, restriction, etc.), then the optimizer choose from a set of pre-defined, low-level implementation procedures. These procedures may involve the user of primary key, foreign key or indexes and other information about the database.

4. Generate query plans and choose the cheapest by constructing a set of candidate query plans first, then choose the best plan. To pick the best plan can be achieved by assigning cost to each given plan. The costs is computed according to the number of disk I/O's involved.

<u>Module-</u> <u>3:Lecture-33</u>

Transaction

Atransactionisthesmallest unitofoperationdoneonadatabase. It can

be basically of two types:

→Read Transaction

→WriteTransaction

3.2. ACIDProperties of transaction:

Atomicity:(allornothing)

Atransaction is said to be atomic if a transaction always executes all its actions in one step or not executes anyactions at allIt means either allor none of the transactions operations are performed.

Consistency:(Noviolationofintegrityconstraints)

Atransaction must preserve the *consistency* of a database after the execution. The DBMS assumes that this property holds for each transaction. Ensuring this property of a transaction is the responsibility of the user.

Isolation:(concurrentchangesinvisibles)

The transactions must behave as if they are executed in isolation. It means that if several transactions are executed concurrently the results must be same as if they were executed serially in some order. The data used during the execution of a transaction cannot be used by a second transaction until the first one is completed.

Durability:(committedupdatepersist)

The effect of committed transactions should persist even after a crash. It means once a transaction commits, the systemmust guarantee that the result of its operations will never be lost, in spite of subsequent failures.

3.3. Statesofatransaction:

Atransactionmustbeinoneofthefollowingstates:

- Active: the initial state, the transaction stays in this state while it is executing.
- **Partiallycommitted**:afterthefinalstatementhas been executed.
- Failed:whenthenormalexecutioncannolongerproceed.
- Aborted: after the transaction has been rolled back and the database has been restored to its state prior to the start of the transaction.
- **Committed**:aftersuccessfulcompletion.

ThestatediagramcorrespondingtoatransactionisshowninFigure.



We say that a transaction has committed only if it has entered the committed state. Similarly, we say that a transaction has aborted only if it has entered the aborted state. A transaction is said to have terminated if has either committed or aborted.

A transaction starts in the active state. When it finishes its final statement, it enters the partially committedstate. Atthispoint, the transaction has completed its execution, but it is still possible that it may have to be aborted, since the actual output may still be temporarily hiding in main memory and thus a hardware failure may preclude its successful completion

The database system writes out enough <u>information</u> to disk that, even in the event of a failure, the updates performed by the transaction can be recreated when the system estarts after the failure. When the last of this information is written out, the transaction enters the committed state.

Problemsdue toconcurrency:

- 1) Lostupdateproblem
- 2) DirtyRead problem
- 3) Incorrectsummaryproblem
 - 1. The lost updateproblem: Asecondtransactionwrites asecond value of a data-item(datum) ontop of a first value written by a first concurrent transaction, and the first value is lost to other transactions running concurrently which need, by their precedence, to read the first value. The transactions that have read the wrong value end with incorrect results.
 - 2. Thedirtyreadproblem: Transactionsreadavalue writtenbyatransactionthat hasbeenlater aborted. This value disappears from the database upon abort, and should not have been read by any transaction ("dirty read"). The reading transactions end with incorrect results.
 - 3. The incorrect summaryproblem: While one transaction takes a summaryover the values of allthe instances of a repeated data-item, a second transactionupdates some instances ofthat data-item. The resulting summarydoes not reflect a correct result for any(usually neededforcorrectness)precedenceorderbetweenthetwotransactions(ifone isexecutedbeforethe other), but rather some random result, depending on the timing of the updates, and whether certain update results have been included in the summary or not.

ConcurrencyControlTechniques:

- 1) Locking
- 2) Timestampordering
- 3) Multi-versionconcurrencycontrol
- 4) Optimistic concurrency control

Locking(2-PhaseLocking)

2PLprotocollocksareapplied and removed in two phases:

- 1. Expandingphase:locksareacquiredandnolocksarereleased.
- 2. Shrinkingphase:locksarereleasedandnolocksareacquired.

Locks areoftwotypes:

- 1. BinaryLock
- 2. Share/Exclusive(Read/Write)Lock

→BinaryLock

Abinarylockcanhave 2Statesorvalues

- Locked(or1)and
- Unlocked (or0)

We represent the current state (or value) of the lock associated with data item X as LOCK (X).

Operationsused with Binary Locking

- 1. **lock_item:**Atransactionrequest accesstoanitembyfirst issuinga lock_item(X) operation.
 - IfLOCK(X)=1 orL(X): thetransactionisforcedtowait.
 - **IfLOCK(X)=0orU(x):** it issetto 1(thetransaction**locks**the item)andthe transaction is a load to access item X.
- 2. **unlock_item:** Afterusing the data item the transaction issues an operation **unlock(X)**, which sets the operation **LOCK(X)** to 0 i.e. unlocks the data item so that X may be accessed by another transactions.

TransactionRulesforBinaryLocks

Everytransactionmustobeythefollowingrules:

 \rightarrow AtransactionT mustissuethelock(X)operation beforeanyread(X)orwrite(X)operationsinT. \rightarrow AtransactionT mustissuetheunlock(X)operationafterallread(X)andwrite(X)operations inT.

 \rightarrow IfatransactionTalreadyholdsthelock onitemX,thenTwillnotissuealock(X)operation.

 \rightarrow If a transaction does not hold sthe lock on item X, then T will not is sue an unlock (X) operation.

lock_item (X):

```
B: if LOCK (X)=0 (* item is unlocked *)
then LOCK (X) ←1 (* lock the item *)
else begin
wait (until lock (X)=0 and
the lock manager wakes up the transaction);
go to B
end;
```

unlock_item (X):

```
LOCK (X)←0; (* unlock the item *)
if any transactions are waiting
then wakeup one of the waiting transactions;
```

→Shared/Read and Exclusive/WriteLock

The binarylock is too restrictive for data items because at most one transaction canhold on a given itemwhether the transaction is reading or writing. To improve it we have shared and exclusive locks inwhichmore than one transaction can access the same item for reading purposes. i.e. the read operations on the same item by different transactions are not conflicting.

Inthistypesoflock, system supports two kinds of lock:

- Exclusive(orWrite)Locks and
- Shared(or Read)Locks.

SharedLocks

If a transaction Tihas locked the data item Ainshared mode, then are quest from another transaction Tj on A for :

- WriteoperationonA:Denied.TransactionTjhastowaituntiltransactionTiunlockA.
- Read operationonA:Allowed.

ExclusiveLocks

IfatransactionTihas lockedadataitema inexclusive modethenrequest fromsomeanother transaction Tj for

- Read operationonA: Denied
- Writeoperation onA:Denied

OperationsUsedwithSharedandExclusiveLocks

- 1. Read_lock(A)or s(A)
- 2. Write_lock(A)orX(A)
- 3. Unlock(X)orU(A)

ImplementationofSharedandExclusiveLocks

Sharedandexclusive locksareimplementedusing4fields:

- 1. Data_item_name
- 2. LOCK
- 3. NumberofRecords and
- 4. Locking_transaction(s)

Againtosavespace, itemsthat arenot in he locktableareconsidered to be unlocked. The system maintains only those records for the items that are currently locked in the lock table.

ValueofLOCK(A):ReadLockedorWriteLocked

- **IfLOCK(A)=write-locked**–The valueoflockingtransactionisasingletransactionthat holds the exclusive(write) Lock on A.
- **IfLOCK(A)=read-locked** –Thevalueoflockingtransactionisa listofoneormore transactions that hold the Shared(read) onA.

TransactionRulesforSharedandExclusiveLocks

Everytransactionmustobeythefollowingrules:

- 1. AtransactionT mustissuetheoperation s(A)orread_lock(A)orx(A)orwrite_lock(A) before any read(A) operation is performed in T.
- 2. AtransactionTmust issuetheoperationx(A)orwrite_lock(A)beforeanywrite(A) operation is performed in T.
- 3. Aftercompletionofallread(A)and write(A)operationsinT,atransactionTmust issuean unlock(A) operation.
- 4. If a transactional ready holds are ad (shared) lock or a write (exclusive) lock on item A, then T will not issue an unlock (A) operation.
- 5. Atransaction that already holds a lockonitemA, is allowed to convert the lock from one locked state to another under certain conditions.
 - UpgradingtheLockbyIssuingawrite_lock(A)OperationorConversionof read_lock() to write_lock() :
 - **Case1–WhenConversionNotPossible:**AtransactionTwillnot issuea write_lock(A) operation if it already holds a read (shared) lock or write (exclusive) lock on item A.
 - Case2–WhenConversionPossible: IfTistheonlytransactionholdinga read lock onAat the time it issues the

write_lock(A)operation,thelockcanbeupgraded;

 DowngradingtheLockbyIssuingaread_lock(A)orConversionofwrite_lock() to read_lock() :

AtransactionTdowngradefromthewrite locktoareadlockbyacquiringthe write_lock(A) or x(A), then the read_lock(A) or s(A) and then releasing the write_lock(A) or x(A).

read_lock (X):

```
B: if LOCK (X)="unlocked"

then begin LOCK (X)← "read-locked";

no_of_reads(X)← 1

end

else if LOCK(X)="read-locked"

then no_of_reads(X)← no_of_reads(X) + 1

else begin wait (until LOCK (X)="unlocked" and

the lock manager wakes up the transaction);

go to B

end;
```

write_lock (X):

B: if LOCK (X)='unlocked" then LOCK (X)'- "write-locked" else begin wait (until LOCK(X)="unlocked" and the lock manager wakes up the transaction); go to B end;

unlock_item (X):

```
if LOCK (X)="write-locked"
then begin LOCK (X)← "unlocked;"
wakeup one of the waiting transactions, if any
end
else if LOCK(X)="read-locked"
then begin
no_of_reads(X)← no_of_reads(X) – 1;
if no_of_reads(X)=0
then begin LOCK (X)="unlocked";
wakeup one of the waiting transactions, if any
end
end;
```

Lecture-34

Problemsduetolocking

1) deadlock

2) starvation

Deadlock:

Itisanindefinitewait situationinwhichaseriesoftransactionswait foreachotherforunknown amount of time, it obeys the following conditions:

→Hold&Wait

→NoPreemption

→ Mutual Exclusion

→CircularWait

Example:

For example, assume a set oftransactions $\{T_0, T_1, T_2, ..., T_n\}$. Toneeds a resource X to complete its task.ResourceX isheld by T_1 , and T_1 is waiting for a resourceY, which isheld by T_2 . To resource Z, which is held by T_0 . Thus, all the processes wait for each other to release resources. In this situation, none of the processes can finish their task. This situation is known as a deadlock.

Starvation:

Itisasituationinwhichatransactionhas lockedadatabase forunfair means&allother transactions are in indefinite waiting.

Starvation	Deadlock
Starvationhappensifsametransactionis always chosen as victim.	Adeadlock is a condition in which two or more transaction is waiting for each other.
It occursif the waiting schemeforlocked itemsinunfair,givingprioritytosome transactionsoverothers.	A situationwheretwoormoretransactionsareunableto proceedbecauseeachiswaitingforoneoftheothertodo something.
Starvationisalso knownaslivedlock.	Deadlockisalsoknownascircularwaiting.
Avoidance: ->switch priorities so that every thread has a chance to have high priority. ->UseFIFOorderamongcompeting request.	Avoidance: ->Acquire locksarepredefinedorder. ->Acquirelocksatoncebeforestarting.
Itmeansthattransactiongoesinastate where transaction never progress.	Itisasituationwheretransactionsarewaitingforeach other.

TimestampOrdering

The timestamp-ordering protocolensures serializability among transactions in their conflicting read and write operations. This is the responsibility of the protocol system that the conflicting pair of tasks should be executed according to the timestamp values of the transactions.

- The timestampoftransaction T_i is denoted as TS(T_i).
- Readtime-stampofdata-itemXisdenotedbyR-timestamp(X).
- Writetime-stampofdata-itemXisdenotedbyW-timestamp(X).

Timestamporderingprotocolworksasfollows-

- IfatransactionTiissuesaread(X)operation-
 - \circ IfTS(Ti)<W-timestamp(X)
 - Operationrejected.
 - IfTS(Ti)>=W-timestamp(X)
 - Operationexecuted.
 - Alldata-itemtimestamps updated.
- IfatransactionTiissuesawrite(X)operation-
 - IfTS(Ti)<R-timestamp(X)
 - Operationrejected.
 - IfTS(Ti)<W-timestamp(X)
 - OperationrejectedandTirolledback.
 - Otherwise, operation executed.

Avoidingdeadlock:

Amajordisadvantageoflocking isdeadlockwhichcanbeavoidedusingtimestamporderingas follows:

Therearetwoalgorithms fordeadlockavoidance.

- Wait/Die
- Wound/Wait

Here isthetablerepresentationofresourceallocationforeachalgorithm. Bothofthesealgorithms takeprocessage intoconsiderationwhiledeterminingthebest possiblewayofresourceallocation for deadlock avoidance.

	Wait/Die	Wound/Wait
Olderprocessneedsaresourceheldbyyounger process	Olderprocess waits	Youngerprocess dies
Youngerprocessneedsaresourceheldbyolder process	Youngerprocess dies	Youngerprocess waits

Wait-DieScheme

Inthisscheme, if a transaction requests to lock are source (dataitem), which is already held with a conflicting lock by another transaction, then one of the two possibilities may occur –

- If $TS(T_i) < TS(T_j)$ -that is T_i , which is requesting a conflicting lock, isolder than T_j -then T_i is allowed to wait until the data-item is available.
- If $TS(T_i) > TS(t_j)$ -that is T_i is younger than T_j -then T_i dies. T_i is restarted later with a random delay but with the same timestamp.

This scheme allows the older transaction to wait but kills the youngerone.

Wound-WaitScheme

Inthisscheme, if a transaction requests to lock are source (dataitem), which is already held with conflicting lock by some another transaction, one of the two possibilities may occur –

- If $TS(T_i) < TS(T_j)$, then T_i forces T_j to be rolled back that $is T_i$ wounds T_j . T_j is restarted later with a random delay but with the same timestamp.
- $IfTS(T_i)>TS(T_j)$, then T_i is forced to wait until the resource is available.

Thisscheme, allows the younger transaction to wait; but when an older transaction requests an item held by a younger one, the older transaction forces the younger one to abort and release the item.

Inboththe cases, the transaction that enters the systemata later stage is aborted.

Multi-versionConcurrencyControlTechniques:

This concurrency control technique keeps the old values of a data item when the item is updated. These are known as multiversion concurrency control, because several versions (values) of an item are maintained.

 \rightarrow When a transaction requires access to an item, an *appropriate* version is chosen to maintain the serializability of the currently executing schedule, if possible. The idea is that some read operations that would be rejected in other techniques can still be accepted by reading an *older version* of the item to maintain serializability. When a transaction writes an item, it writes a *new version* and the old version of the item is retained. Some multiversion concurrency control algorithms use the concept of view serializability rather than conflict serializability.

 \rightarrow An obvious drawback of multiversion techniques is that more storage is needed to maintain multiple versions of the database items. However, older versions may have to be maintained anyway—for example, for recoverypurposes. In addition, some database applications require older versions to be kept to maintain a history of the evolution of data item values.

 \rightarrow Theextremecase isa *temporaldatabase*, whichkeepstrackofallchangesandthetimesat which they occurred. In such cases, there is no additional storage penalty for multiversion techniques, since older versions are already maintained.

Lecture-35

Multiversion Technique Based on Timestamp Ordering

Inthismethod, several versions,, of each data item *X* are maintained. For *each version*, the value of version and the following two timestamps are kept:

 $1. \ read_TS: The read timestampofis the largest of all the timestamps of transactions that have successfully read version .$

2. write_TS:Thewritetimestampofisthetimestampofthetransactionthatwrotethevalueof version

Whenever a transaction T is allowed to execute a write_item(X) operation, a new version of item X is created, with both the write_TS and the read_TS set to TS(T). Correspondingly, when a transactionT is allowed to readthe value of *Xi*, the value of read_TS() is set to the larger of the current read_TS() and TS(T).

Toensureserializability, thefollowingtworulesare used:

1. If transaction T issues a write_item(X) operation, and version *i*of X has the highest write_TS() of all versions of X that is also *less than or equalto* TS(T), and read_TS() >TS(T), then abort and roll back transaction T; otherwise, create a new version of X with read_TS() = write_TS() = TS(T).

2. If transaction T issues a read_item(X) operation, find the version *i* of X that has the highest write_TS() of all versions of X that is also *less than or equal to* TS(T); then return the value of to transaction T, and set the value of read_TS() to the larger of TS(T) and the current read_TS().

As we can see in case 2, a read_item(X) is always successful, since itfinds the appropriate version to read based on the write_TS of the various existing versions of X. In case 1, however, transaction T may be aborted and rolled back. This happens if T is attempting to write a version of X thatshould have been read by another transaction T whose timestamp is read_TS(); however, T has already readversion Xi, which was written by the transaction with timestamp equal towrite_TS(). If this conflict occurs, T is rolled back; otherwise, a new version of X, written by transaction T, is created. Notice that, if T is rolled back, cascading rollback may occur. Hence, to ensure recoverability, a transaction T should not be allowed to commit until after all the transactions that have written some version that T has read have committed.

MultiversionTwo-PhaseLockingUsingCertifyLocks

In this multiple-mode locking scheme, there are *three locking modes* for an item: read, write, and certify, instead of just the two modes (read, write). Hence, the state of LOCK(X) for an item X can be one of read-locked, write-locked, certify-locked, or unlocked.

In the standard locking scheme, once a transaction obtains a write lock on an item, no other transactions canaccess that item. The idea behind multiversion2PL is to allow other transactions T to read an item X while a single transaction T holds a write lock on X. This is accomplished by allowing *two versions* for each item X; one version must always have been written by some committed transaction. The second version X is created when a transaction T acquires a write lock on the item. Other transactions can continue to read the *committed version* of X while T holds the writelock.TransactionTcanwritethevalueofXasneeded,withoutaffectingthevalueofthe

committed version X. However, once Tisreadyto commit, it must obtain certify lock onallitems that it currently holds write locks on before it can commit. The certify lock is not compatible with read locks, so the transaction may have to delay its commit until all its write-locked items are released by any reading transactions in order to obtain the certify locks. Once the certify locks— which are exclusive locks—are acquired, the committed version X of the data item is set to the value of version X, version X is discarded, and the certify locks are then released. In this multiversion 2PL scheme, reads can proceed concurrently with a single write operation—an arrangement not permitted under the standard 2PL schemes.

3. 6.4. Optimistic Concurrency Control Techniques:

Basic idea: all transactions consist of three phases:

- 1. Read.Here,allwritesaretoprivatestorage(shadowcopies).
- 2. Validation.Makesurenoconflictshaveoccurred.
- 3. Write.IfValidationwassuccessful, makewritespublic.(Ifnot,abort!)

Useful in the following cases:

- 1. Alltransactionsarereaders.
- 2. Lotsoftransactions, eachaccessing/modifyingonlyasmallamountofdata, largetotal amount of data.
- 3. Fractionoftransactionexecutioninwhichconflicts"reallytakeplace" issmallcomparedto total pathlength.

TheValidationPhase

- Goal:toguaranteethatonlyserializableschedulesresult.
- Technique:actuallyfindanequivalent serializableschedule.Inparticular,
- 1. AssigneachtransactionaTNduringexecution.
- 2. Ensure that if your untransactions in order induced by "<" on TNs, you get an equivalent serial schedule.

SupposeTN(Ti)<TN(Tj).Thenifoneofthefollowingthreeconditionsholds, it's serializable:

- 1. TicompletesitswritephasebeforeTjstartsitsread phase.
- 2. WS(Ti) intersect RS(Tj)=*emptyset*andTicompletesitswritephasebeforeTjstarts its write phase.
- 3. WS(Ti)intersectRS(Tj)=WS(Ti)*intersect*WS(Tj)=*emptyset*andTicompletesits read phase before Tj completes its read phase.

Recovery

TypesofFailure

Failures maybe

Transaction	Causedbyerrorswithinthetransactionprocesses.
System	Caused byfailureofnetworkoroperatingsystemorphysicalthreatstothe system as a whole.
Media	Failureofhard disk,outofmemoryerrors,outofdiskspace errors.

ReasonsforFailure

Failuremaybecausedbyanumber ofthings.

ASystemCrash	Ahardware, software or network error causes the transaction to fail.
TransactionorSystem error	Someoperationinthetransactionmaycausethe failureortheusermay interrupt the transaction.
LocalErrorsor Exceptions	Conditionsoccurduring the transaction that results in transaction cancellation.
ConcurrencyControl Enforcement	Severaltransactionsmaybe indeadlocksothetransactionmaybeaborted to be restarted later.
DiskFailure	ReadWriteerroronthephysicaldisk.
PhysicalProblems	Thiscanbeanyrangeofphysicalproblems, suchaspowerfailure, mounting wrong disk or tape by operator, wiring problems etc
CatastropheSituations	Largescalethreatstothesystemandthedataforexample fire,cyclone, security breaches etc.

Transaction errors, system errors, system crashes, concurrency problems and local errors or exceptions are the more common causes of system failure. The system mustbe able to recoverfrom such failures without loss of data.

Log-BasedRecovery

COMMIT

Signalsthe successfulendofatransaction

- Anychangesmadebythetransactionshouldbe saved
- These changes are now visible to other transactions

ROLLBACK

Signalsthe unsuccessfulendofatransaction

- Anychangesmadebythetransactionshouldbe undone
- Itisnowas ifthetransactionnever existed

Log-BasedRecovery

The most widely used structure for recording database modifications is the *log*. The log is a sequence of *log records* and maintains a history of all update activities in the database. There are several types of log records.

Anupdatelogrecorddescribesasingledatabasewrite:

- Transactionsidentifier.
- Data-item identifier.
- Old value.
- Newvalue.

Other special log records exist to record significant events during transaction processing, such as the start of a transaction and the commit or abort of a transaction. We denote the various types of log records as:

- <Tistart>.TransactionTihasstarted.
- <Ti,Xj, V1, V2> Transaction Tihas performeda write on data item Xj.Xj hadvalue V1 before write, and will have value V2 after the write.
- <Ticommit>TransactionTihascommitted.
- <Tiabort>TransactionTihasaborted.
Whenever a transaction performs a write, it is essential that the log record for that write be created before the database is modified. Once a log record exists, we can output the modification that has alreadybeenoutputtothedatabase. Also we have the ability to *undo* amodification that has already been output to the database, by using the old-value field in the log records.

For log records to be useful for recovery from system and disk failures, the log must reside onstable storage. However, since the log contains a complete record of all database activity, the volume of data stored in the log may become unreasonable large.

DeferredDatabaseModification

The deferred-modification technique ensures transaction atomicity by recording all database modifications in the log, but deferring all write operations of a transaction until the transaction partially commits (i.e., once the final action of the transaction has been executed). Then the information in the logs is used to execute the deferred writes. If the system crashes or if the transaction aborts, then the information in the logs is ignored.

LetT0betransactionthattransfers\$50 fromaccountAto account B: T0:

read (A); A:=A-50; Write(A); Read(B); B:=B+50; Write(B).

ImmediateDatabaseModification

The immediate-update technique allows database modifications to be output to the database while thetransactionisstillintheactivestate. These modifications are called *uncommittedmodifications*. In the event of a crash or transaction failure, the system must use the old-value field of the log records to restore the modified data items.

Transactions T0 and T1 executed one after the other in the order T0 followed byT1. The portionofthe logcontaining there levant information concerning these two transactions appears in the following,

PortionofthesystemlogcorrespondingtoT0andT1

<T0start>

<T0,A,1000,950> <T0,B,2000,2050>

<T0commit> <T1 start > <T1,C,700,600> <T0commit>

Checkpoints

When a system failure occurs, we must consult the log to determine those transactions that need to be redone and those that need to be undone. Rather than reprocessing the entire log, which is time-consuming and much of it unnecessary, we can use *checkpoints*:

- $\bullet \quad Output onto stable storage all the log records currently residing in main memory.$
- Outputtothediskallmodified buffer blocks.
- Outputontostablestoragealog record, < checkpoint >.

Lecture-36

Serializability:

When several concurrent transactions are trying to access the same data item, the instructions within these concurrent transactions must be ordered in some way so as there are no problem in accessing and releasing the shared data item. There are two aspects of serializability which are described here:

ConflictSerializability

Two instructions of two different transactions may want to access the same data item in order to perform a read/write operation. Conflict Serializability deals with detecting whether the instructions are conflicting in any way, and specifying the order in which these two instructions will be executed in case there is any conflict. A **conflict** arises if at least one (or both) of the instructions is a write operation. The following rules are important in Conflict Serializability:

- 1. If two instructions of the two concurrent transactions are both for read operation, then they are not in conflict, and can be allowed to take place in anyorder.
- 2. If one of the instructions wants to perform a read operation and the other instruction wants to perform a write operation, then they are in conflict, hence their ordering is important. If the read instruction is performed first, then it reads the old value of the data item and after the reading is over, the new value of the data item is written. It the write instruction is performed first, then updates the data itemwith the new value and the read instruction reads the newly updated value.
- 3. If both the transactions are for write operation, then they are in conflict but can be allowed to take place in anyorder, because the transaction do not read the value updated by each other. However, the value that persists in the data itemafter the schedule is over is the one written by the instruction that performed the last write.

It mayhappenthat wemaywant to execute the same set of transaction in a different schedule on another day. Keeping in mind these rules, we may sometimes alter parts of one schedule (S1) to create another schedule (S2) by swapping only the non-conflicting parts of the first schedule. The conflicting parts cannot be swapped in this way because the ordering of the conflicting instructions is important and cannot be changed in any other schedule that is derived from the first. If these two schedules are made of the same set of transactions, then both S1 and S2 would yield the same result if the conflict resolution rules are maintained while creating the new schedule. In that case the schedule S1 and S2 would be called **Conflict Equivalent**.

ViewSerializability:

This is another type of serializability that can be derived by creating another schedule out of an existing schedule, involving the same set of transactions. These two schedules would be called View Serializable if the following rules are followed while creating the second schedule out of the first. Let us consider that the transactions T1 and T2 are being serialized to create two different schedules S1 and S2 which we want to be **View Equivalent** and bothT1 and T2 wants to access the same data item.

- 1. If inS1, T1reads theinitial value of the data item, then in S2also, T1 should read the initial value of that same data item.
- 2. IfinS1,T1writesavalueinthedataitemwhichisreadbyT2,theninS2also,T1 should write the value in the data item before T2 reads it.
- 3. IfinS1,T1performsthefinalwriteoperationonthatdataitem,theninS2also,T1 should perform the final write operation on that data item.

Exceptinthesethreecases, any alteration can be possible while creating S2 by modifying S1.

Lecture-37

ObjectOriented Databases

Object oriented databases are also called Object Database Management Systems (ODBMS). Object databases storeobjects ratherthandata suchas integers, strings orrealnumbers. Objects are used in object oriented languages such as Smalltalk, C++, Java, and others. Objectsbasically consist of the following:

- Attributes Attributes are data which defines the characteristics of an object. This data may be simple such as integers, strings, and real numbers or it may be a reference to a complex object.
- Methods Methods define the behavior of an object and are what was formally called procedures or functions.

Thereforeobjectscontainbothexecutablecodeanddata

ObjectPersistence

With traditional databases, data manipulated by the application is transient and data in the database is persisted (Stored on a permanent storage device). In object databases, the application can manipulate both transient and persisted data.

WhentoUseObjectDatabases

Object databases should be used whenthere is complexdata and/or complexdata relationships. This includes a many to many object relationship. Object databases should not be used when there would be few join tables and there are large volumes of simple transactional data.

Objectdatabasesworkwellwith:

- CAS Applications (CASE-computer aided software engineering,CAD-computer aided design, CAM-computer aided manufacture)
- MultimediaApplications
- Objectprojectsthatchangeovertime.
- Commerce

ObjectDatabaseAdvantagesoverRDBMS

- Objectsdon't requireassemblyanddisassemblysavingcodingtimeandexecutiontimeto assemble or disassemble objects.
- Reducedpaging
- Easiernavigation
- Better concurrencycontrol-Ahierarchyofobjectsmaybelocked.
- Datamodelisbasedonthe realworld.
- Workswellfordistributedarchitectures.
- Lesscoderequiredwhenapplicationsareobjectoriented.

ObjectDatabaseDisadvantagescomparedtoRDBMS

- Lowerefficiencywhendataissimpleandrelationshipsare simple.
- Relationaltables are simpler.
- Latebindingmayslowaccessspeed.
- MoreusertoolsexistforRDBMS.
- StandardsforRDBMSaremorestable.

• Support for RDBMS is more certain and change is less likely to be required.

HowDataisStored

Two basic methods are used to store objects by different database vendors.

- Eachobject hasauniqueIDand isdefinedasasubclassofabaseclass, using inheritanceto determine attributes.
- Virtualmemorymapping isused forobjectstorageand management.

DataWarehouse

- A data warehouse is a subject-oriented, integrated, time-variant and non-volatile collection of data in support of management's decision making process.
- **Subject-Oriented**: A data warehouse can be used to analyze a particular subject area. For example, "sales" can be a particular subject.
- **Integrated**: Adata warehouse integrates data from multiple data sources. For example, source A and source B may have different ways of identifying a product, but in a data warehouse, there will be only a single way of identifying a product.
- **Time-Variant**: Historical data is kept in a data warehouse. For example, one canretrieve data from 3 months, 6 months, 12 months, or even older data from a data warehouse. This contrasts with a transactions system, where often only the most recent data is kept. For example, a transaction system may hold the most recent address of a customer, where a data warehouse can hold all addresses associated with a customer.
- Non-volatile: Once data is in the data warehouse, it will not change. So, historical datain a data warehouse should never be altered.

Data Warehousing Architecture & Components

Followingdiagramdepictsdifferent componentsofDataWarehousearchitecture.



OperationalSourceSystem

It'sthetraditionalOLTPsystemswhichstorestransactiondataoftheorganizationsbusiness. Its generally used one record at any time not necessarily stores history of the organizations information's. Operational sourcesystems generally not used for reporting likedatawarehouse.

DataStagingArea

Datastagingarea isthestorageareaaswellassetofETLprocessthatextract datafromsource system. It is everything between source systems and Data warehouse.

Datastagingareneverbeused forreportingpurpose.Dataisextractedfromsourcesystemand stored, cleansed, transformed in staging area to load into data warehouse.

StagingarenotnecessarilytheDBMS.Itcould beflat filesalso.Stagingareacanbe structuredlike normalized source systems. It totally depends on choice and need of development process.

DataPresentationArea

Data presentation area is generally called as data warehouse. It's the place where cleaned, transformeddataisstoredinadimensionallystructuredwarehouseand madeavailable foranalysis purpose

DataAccessTools

oncedataisavailable inpresentationarea it isaccessedusingdataaccesstoolslikeBusiness Objects.

Schema: (Starschema)

The star schema architecture is the simplest data warehouse schema. It is called a star schema because the diagram resembles a star, with points radiating from center. The center of the star consists of fact table and the points of the star are the dimensiontables. Usually the fact tables in a star schema are inthird normalform(3NF) whereas dimensional tables are de-normalized. Despite the fact that the star schema is the simplest architecture, it is most commonly used now adays and is recommended by Oracle.

→FactTables

Afact table typically has two types of columns: foreign keys to dimension tables and measures those that contain numeric facts. Afact table can contain fact's data on detailor aggregated level.

→DimensionTables

Adimension is a structure usually composed of one or more hierarchies that categorizes data. If a dimensionhasn't gotahierarchiesand levelsit iscalled **<u>flatdimensionorlist</u>**. Theprimarykeysof each of the dimension tables are part of the composite primary keyof the fact table. Dimensional attributes help to describe the dimensional value. They are normally descriptive, textual values. Dimensiontablesaregenerallysmallin sizethen fact table.

Typicalfacttablesstoredataaboutsaleswhiledimensiontablesdataaboutgeographic region(markets, cities), clients, products, times, channels.

Themaincharacteristicsofstar schema:

- 1. Simplestructure->easytounderstandschema
- 2. Greatqueryeffectives ->smallnumber oftablesto join
- 3. Relativelylongtimeofloadingdataintodimensiontables ->de-normalization, redundancy data caused that size of the table could be large.
- 4. Themostcommonly usedinthedatawarehouseimplementations->widely supported by a large number of business intelligence tools



DataMining:

 \rightarrow Generally, data mining (sometimes called data or knowledge discovery) is the process of analyzing data from different perspectives and summarizing it into useful information - information that can be used to increase revenue, cuts costs, or both. Data mining software is one of a number of analytical tools for analyzing data. It allows users to analyze data from manydifferentdimensionsorangles, categorizeit, and summarize the relationships identified. Technically, data mining is the process of finding correlations or patterns among dozens of fields in large relational databases.

 \rightarrow While large-scale information technology has been evolving separate transaction and analytical systems, data mining provides the link between the two. Data mining software analyzes relationships and patterns in stored transaction data based on open-ended user queries.Severaltypesofanalyticalsoftwareareavailable:statistical, machine learning, and neural networks. Generally, any of four types of relationships are sought:

- **Classes:** Stored data is used to locate data in predetermined groups. For example, a restaurant chaincould mine customer purchase data to determine whencustomersvisit and what theytypicallyorder. This information could be used to increase traffic by having daily specials.
- **Clusters**: Data items are grouped according to logical relationships or consumer preferences.Forexample,datacanbeminedto identifymarket segmentsorconsumer affinities.
- Associations:Datacanbe minedto identifyassociations.The beer-diaperexample isan example of associative mining.

• **Sequentialpatterns**:Dataisminedtoanticipatebehaviorpatternsandtrends.Forexample, an outdoor equipment retailer could predict the likelihood of a backpack being purchased based on a consumer's purchase of sleeping bags and hiking shoes.

Dataminingconsistsoffive majorelements:

- Extract, transform, and load transaction data on to the data warehouse system.
- Storeandmanagethedatainamultidimensionaldatabasesystem.
- Providedataaccesstobusinessanalystsandinformationtechnologyprofessionals.
- Analyze thedatabyapplicationsoftware.
- Presentthedatainauseful format, such as a graphor table.

Techniquesusedindatamining:

Association

Associationisoneofthebest-knowndataminingtechnique. In association, apatternis discoveredbasedonarelationship betweenitemsin thesametransaction. That'sisthe reason why association technique is also known as *relation technique*. The association technique is used in *market basket analysis* to identify a set of products that customers frequently purchase together.

Retailers are using association technique to research customer's buying habits. Based on historicalsaledata, retailersmight findoutthat customersalwaysbuycrispswhentheybuy beers, and, therefore, theycan put beers and crisps next to each other to save time for customer and increase sales.

Classification

Classification is a classic data mining technique based on machine learning. Basically, classification is used to classify each item in a set ofdata into one of a predefined set of classes or groups. Classification method makes use of mathematical techniques such as decision trees, linear programming, neural network and statistics. In classification, we develop thesoftware that can learn how to classify the data items into groups. For example, we can apply classification in the application that "given all records of employees who left the company, predict who will probably leave the company in future period." In this case, we divide the records of employees into two groups that named "leave" and "stay". And then we can ask our data mining software to classify the employees into separate groups.

• Clustering

Clustering is a data mining technique that makes a meaningful or useful cluster of objects which have similar characteristics using the automatic technique. The clustering technique defines the classes and puts objects in each class, while in the classification techniques, objectsareassigned into predefined classes. To make the concept clearer, we can take book management in the library as an example. In a library, there is a wide range of books on various topics available. The challenge is how to keep those books in a wide readers can take several books on a particular topic without hassle. By using the clustering technique, we cankeep books that have some kinds of similarities inone cluster or one shelf and label it with a meaningfulname. If readers want tograbbooks in that topic, they would only have to go to that shelf instead of looking for the entire library.

• Prediction

The prediction, as its name implied, is one of a data mining techniques that discovers the relationship between independent variables and relationship between dependent and independent variables. For instance, the prediction analysistechnique can be used in the sale to predict profit for the future if we consider the sale is an independent variable, profit could be a dependent variable. Then based on the historical sale and profit data, we can draw a fitted regression curve that is used for profit prediction.

• SequentialPatterns

Sequential patterns analysis is one of data mining technique that seeks to discover or identifysimilar patterns, regular eventsortrends intransactiondataover abusinessperiod.

In sales, with historical transaction data, businesses can identify a set of items that customersbuytogetherdifferent timesinayear. Thenbusinesses can use this information to recommend customers buy it with better deals based on their purchasing frequency in the past.

• Decisiontrees

TheAdecision tree is one of the most common used data mining techniques because its modeliseasyto understand for users.Indecisiontreetechnique, therootofthedecisiontree is a simple question or condition that has multiple answers. Each answer then leads to a set of questions or conditions that help us determine the data so that we can make the final decision based on it. For example, We use the following decision tree to determine whether or not to play tennis:



Startingattherootnode, if the outlook is overcast then we should definitely playtennis. If it is rainy, we should only play tennis if the wind is the week. And if it is sunnythen we should play tennis in case the humidity is normal.

Lecture-38

ParallelDatabase

Aparalleldatabase systemperformsparalleloperations, such as loadingdata, building indexes and evaluating queries.

Paralleldatabasescanbe roughlydividedinto twogroups,

a) Multiprocessorarchitecture:

→ Shared memory architecture

Where multiple<u>processors</u>sharethe<u>mainmemory</u>space.



→ Shareddiskarchitecture

Whereeachnodehasitsownmainmemory, but allodessharemassstorage, usually a <u>storage area</u> <u>network</u>. In practice, each node usually also has multiple processors.



→ Shared nothing architecture

Whereeachnodehas itsownmassstorageaswellasmainmemory.



Shared-Nothing Architecture

b) Theotherarchitecturegroupiscalledhybridarchitecture, which includes:

• Non-UniformMemoryArchitecture(NUMA), whichinvolvesthenon-uniformmemory access.



• Cluster(sharednothing+shareddisk:SAN/NAS),which is formed by a group of connected computers.

DistributedDatabase:

 \rightarrow Acentralized distributed database management system (DDBMS) manages the database as if it were all stored on the same computer. The DDBMS synchronizes all the data periodically and, in cases where multiple users must access the same data, ensures that updates and deletes performed on the data at one location will be automatically reflected in the data stored elsewhere.

 \rightarrow The users and administrators of a distributed system, should, with proper implementation, interact with the system as if the system was centralized. This transparency allows for the functionality desired in such a structured system without special programming requirements, allowing for any number of local and/or remote tables to be accessed at a given time across the network.

ThedifferenttypesoftransparencysoughtafterinaDDBMSare

- 1. datadistributiontransparency,
- 2. heterogeneitytransparency,
- 3. transactiontransparency, and
- 4. performancetransparency.

 \rightarrow Data distribution transparency requires that the user of the database should not have to know how the data is fragmented (fragmentation transparency), know where the data they access is actually located (location transparency), or be aware of whether multiple copies of the data exist (replication transparency).

 \rightarrow Heterogeneity transparency requires that the user should not be aware of the fact that they are using a different DBMS if they access data from a remote site. The user should be able to use the same language that they would normally use at their regular access point and the DDBMS should handle query language translation if needed.

 \rightarrow Transaction transparency requires that the DDBMS guarantee that concurrent transactions do not interfere with each other (concurrency transparency) and that it must also handled at abase recovery

(recoverytransparency).

 \rightarrow PerformancetransparencymandatesthattheDDBMSshouldhaveacomparablelevelof performance to a centralized DBMS. Queryoptimizers can be used to speed up response time.

TypesofDDBDesign

Non-Partitioned, Non-Replicated Partitioned, Non-Replicated Non-Partitioned, Replicated Partitioned, Replicated

Advantages of DDBMS's

- Reflectsorganizationalstructure

- Improvedshareability
- Improved availability
- Improved reliability
- Improvedperformance
- Dataarelocatednearestthegreatestdemandsiteandaredispersedtomatchbusiness requirements.
- FasterDataAccessbecauseusersonlywork withalocallystoredsubsetofthe data.
- Fasterdataprocessingbecausethedataisprocessedatseveraldifferentsites.
- -GrowthFacilitation:Newsitescanbeaddedwithout compromisingtheoperationsofothersites.
- -Improved communications because local sites are smaller and closer to customers.
- Reduced operating costs: It ismore cost-effective toadd workstations to a network rather than update a mainframe system.
- UserFriendlyinterfaceequippedwithaneasy-to-useGUI.
- Lessinstancesofsingle-pointfailurebecausedataandworkloadaredistributedamongother

workstations.

- Processorindependence: The enduser is able to access any available copy of data.

Disadvantagesof DDBMS

- IncreasedCost

-Integritycontrol moredifficult,

-Lackofstandards,

-Databasedesignmore complex.

- Complexityofmanagement and control. Applications must recognized at a location and they must be able to stitch together data from various sites.

- Technologically difficult: Data integrity, transaction management, concurrencycontrol, security, backup, recovery, query optimization, access path selection are all issues that must be addressed and resolved

- Securitylapseshave increasedinstanceswhendataareinmultiplelocations.

- Lack of standards due to the absence of communication protocols can make the processing and distribution of data difficult.

- Increased storage and infrastructure requirements because multiple copies of data are required at various separate locations which would require more disk space.

- Increasedcostsduetothe highercomplexityoftraining.

- Requires duplicate infrastructure (personnel, software and licensing, physical location/environment) and these can sometimes offset any operational savings.