# Notes Files and Databases



Bipin C. Desai

BytePress

# **Notes**

# for

# Files and Databases



Concordia University Montreal

#### Limit of Liability/Disclaimer of Warranty:

The authors and the publishers have taken care to prepare this book. However, there is no warranty of the accuracy, completeness or presentation of the latest version/generation of any system discussed in this book. The reader must be aware of the fact that software systems often have multiple bugs and are not well thought out, and are usually suitable for limited situations and/or data combinations. Hence the user must be responsible for the appropriate application of any technique and use of any software or code examples.

Furthermore, there is no assurance whatsoever of the possible usefulness or commercialization of any programs, scripts and examples given in this book.

Any references given are based on their existence at the time of writing and the authors and the publishers do not endorse them or imply any usefulness of the information found therein. The reader must be aware that any web site cited may change, disappear or change their terms of service.

This document in electronic form, bearing a CopyForward permission, could be used for personal use and/or study, free of charge Anyone could use it to derive updated versions. The derived version must be published under CopyForward. All authors of the version used to derive the new version must be included in the updated version in the existing order, followed by name(s) of author(s) producing the derived work. Such derived version must be made available free of charge in electronic form under CopyForward. Any other means of reproduction requires that annual profits(income minus the actual production costs) should be shared with established charitable organizations for children. This annual share must be at least 25% of the profits and the organization being supported must have a very modest administrative charges(20-30% of their annual budget and this sharing amount must be at least 15% of the gross annual revenue). The 25% of the profits is the minimum and the original creator of the digital content may increase it to up to 40%. The derived contents would be governed by the term of the original creator of contents.

Readers who found a CopyForward content or any derived work useful are encouraged to also make a donation to their favourite children charity. Make sure to choose charity which has very modest administrative charges or give directly to some deserving children in your community.

This children's charity contribution requirement of CopyForward is civil and moral! It would be judged in the court of public opinion and the author and all author(s) of this and any derived works allow(s) interested party/parties to take legal actions against the violator of the spirit of sharing.

Published by: Electronic Publishing BytePress.com Inc.

ISBN: 978-1-988392-16-5

#### Notes

# Files and Databases Introduction



Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

Bipin C Desai

#### **Objectives**

- Concepts and use data models(E-R, RDM, unstructured)
- Intro. To the Relational database management systems (RDBMS)
- Query languages (Relational Algebra & Calculus, SQL)
- Concepts of checks, assertions, and triggers.
- Database design and web-programming including: HTML, Javascript, PhP design/implementation of a real application.

#### Data

Where does data come from?

Records of status

Operation of organizations

Data + action → more data

What to record

Legal and traditional commitments

cost of recording and preserving the records

Data + Algorithm → Programs +++

Bipin C Desai

3

# Computers

Jacquard machine - 1804
Textile weaving machine
Pattern controlled by punched cards (first input device)

Charles Babbage:

Difference Engine – 1832

Analytical Engine - no funds

Processor, storage, input device, output device

Partially completed in 1910

Fully programmable

First programmer: Lady Lovelace -Ada (daughter of Byron- the poet)

**US Census-**

Hollerith and the tabulating machine – used punched cards 1911 -Computing-Tabulating-Recording Company ► IBM

# Computers

1930 Vennevar Bush ▶ an analog computer ▶Differential Analyzer

Programming an analog computer



⊕ Bipin C Desai

5

# Computers

1934 James Hilton wrote GoodBye Mr. Chips in 4 days-£50 a royalty!

1939-1843 Howard Aiken and IBM

- ► Mark I (mechanical) Mark IV (vacuum tube)
- 1939 Atanasoff/Berry Iowa State
- 1943 Digital computer- Colossus (UK)
- 1945 Eniac (USA)
- 1949 Manchester Mark I

https://www.britannica.com/technology/computer/The-first-computer

1945

Feb. 4-11 Yalta conference: Decide Europe's re-organization

April 30 Hitler commits suicide

May 8 Armistice day

March-July: Interim committee for the deployment of the Atomic bomb,

Members of Committee

Dr. Vannevar Bush
Dr. Karl T. Compton
Dr. James B. Conant
Mr. George L. Harrison, Acting Chairman

July: Final report of the interim cmt. & a press statement on dropping the bomb(s)

July: Publication(The Atlantic, July, 1945) of *As We May Think* by VANNEVAR BUSH concept of linked documents

https://www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/

August: Most of the Japanese forces have been defeated

August 6 US drops atomic bomb on Hiroshima August 9 US drops atomic bomb on Nagasaki

Bipin C Desai

7

Late 1940s

End of colonization, and new ones -digital and internet based? https://en.wikipedia.org/wiki/Decolonisation\_of\_Asia - /Nakba 1950s

Development of the digital computers

1960

Concept of wide area network, ARPANET, packet transmission Early databases

1970s

TCP/IP protocol, Relational Databases, SQL – IBM a late R-DBMS starter Time sharing, multi-tasking. IBM and the effective end of Usain anti-monopoly law applications. Birth of PC and DOS -drop-out kid wonder! 1980s

Commercialization of the internet, ISPs,

Hypertext, HTML, HTTP, OODB

1990s

Web browsers, search engines, massive data collection, tracking 21st Century

NoSQL, OSN, more drop-out billionaires, AI, LLM, Internet colonization



Allan Gondeck on an IBM1620 1964



IBM 026 keypunches, IBM 1403 line printer and IBM 729 tape drives https://commons.wikimedia.org/w/index.php?curid=17267381



https://en.wikipedia.org/wiki/IBM 729



IBM 7090 computer https://commons.wikimedia.org/w/index.php?curid=2878809

Bipin C Desai

0

# ©BCDesai, NASA, Purdue 1965

#### Modelling:

Represent (approximate)
-physical thing,
-conceptual thing

Broomstick stability control

Application: Stability of rockets such as: Saturn V



Bipin C Desai

Saturn V - The Real thing - 1969



wikimedia.org/w/index.php?curid=6448924

Bipin C Desai





# Data deluge and exploitation: The Beginning!



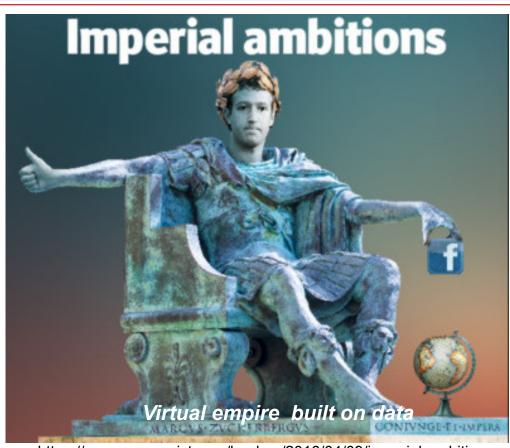
http://www94.web.cern.ch/WWW94/Images/ClosingPanel/Closingpanel1.html

Bipin C Desai



Constantine the Great https://www.flickr.com/photos/yorkminster/5390106900/

Bipin C Desai



https://www.economist.com/leaders/2016/04/09/imperial-ambitions

# Recent Disasters in Database Related Projects

#### **HORIZON**

https://en.wikipedia.org/wiki/British\_Post\_Office\_scandar

#### **PHOENIX:**

https://pipsc.ca/news-issues/press-releases/press-release-phoenix-pay-system-turns-nine-billion-

dollar-breakdown

https://spectrum.ieee.org/canadian-governments-phoenix-pay-system-an-incomprehensible-failure

https://en.wikipedia.org/wiki/Phoenix\_pay\_system

https://ottawacitizen.com/opinion/desousa-9-years-of-phoenix-the-payroll-disaster-thats-burning-workers

https://ottawacitizen.com/news/feds-spending-1-billion-on-maintaining-public-servant-pay-system-over-two-years -as-it-tries-to-fix-phoenix-issues

#### **SAAQClick**

https://globalnews.ca/news/11153545/quebecs-saaqclic-scandal-500-million/

#### Unity

Lack of knowledge, trasperancy, over hype by marketing people, ignorance of elected officers, fear of challenging authority and its reults

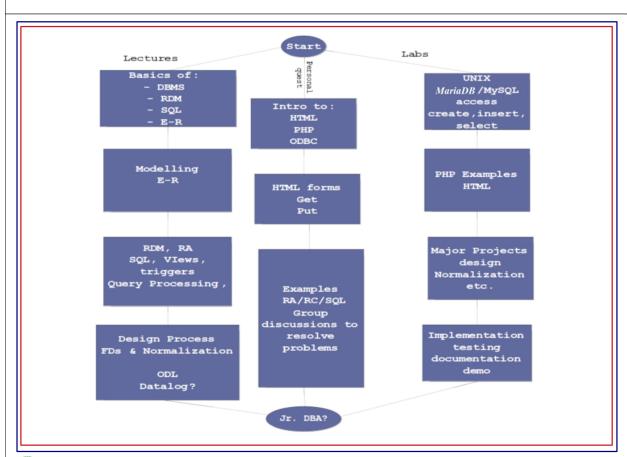
Bipin C Desai

# Files and Databases Introduction

Bipin C. DESAI

Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

de Bipin C. DESAI



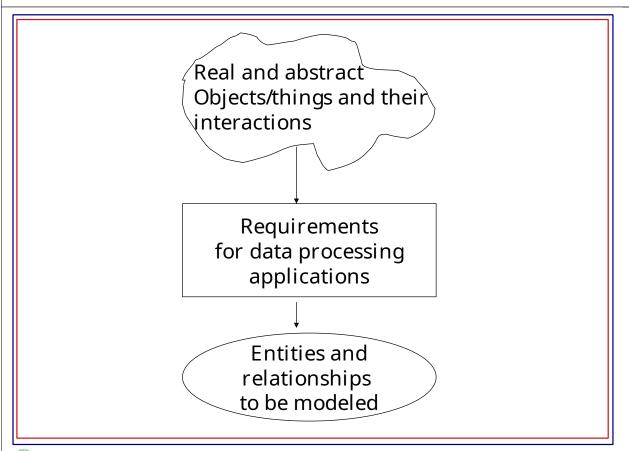
Modeling techniques (E-R, ODL, UML)

Basic relational model

Design of database applications

Database programming (MySQL, SQL, PHP, HTML, CSS, Javascript)

di Bipin C. DESAI



# DBMS! What is it?

Database is an integrated data collection (Logically consistent and persistent)

It is derived from the model of a set of applications for a real world enterprise.

DBMS is a software package designed to make managing almost any database.

DBMS offers: data independence, efficiency, integrity, security, concurrency, recovery

Bipin C. DESAI

-

# Why Database?

Information Age: 30-40% of world trade and growing

Web(Unstructured data) and .com Email,

Digital Library Entertainment

OSN,

Human Genome Project Shopping

Day to day operation of Mama/Papa Store

List of titles, artist, and download site of shared files.

Information about employees, departments, projects, etc. in an organization

Information about students, courses, enrollments, professors, etc. in an educational institute

Information about books, videos, albums, members, etc. in a library

DBMS is a complex set of software packages:

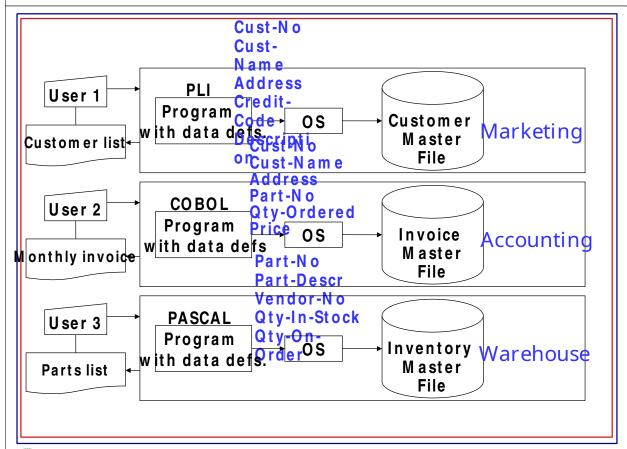
- create new databases, store and manage data provide application development and support environment

**Application Support**: Gives developers tools to build applications for using the data. Allows easy method for users to query and modify the data

Persistent storage: Support the storage of data
Transaction management: Controls concurrent access
to data from many users
Supports the ACID properties.

Atomicity Concurrency Integrity Durability

di Bipin C. DESAI



# Pros & Cons of file based system

Sharing not possible data definition is "locked" in application programs which "owns" the file and the data in it

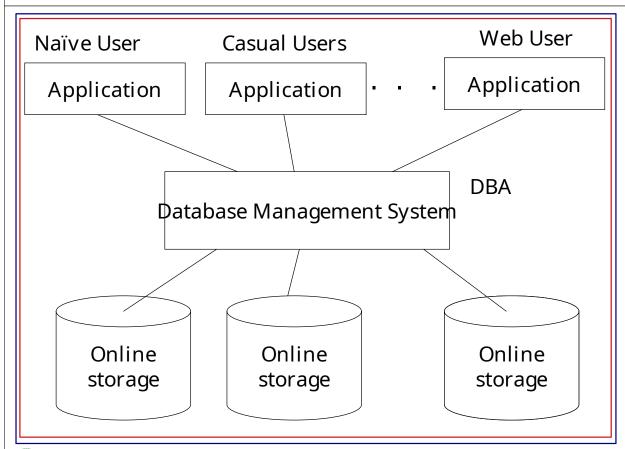
Redundancy of data: Same data is duplicated perhaps in slightly different format over various files

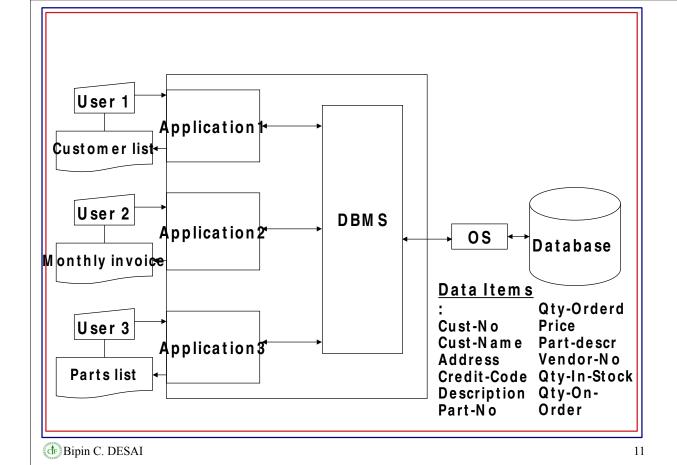
Multiple updates: Changes have to be made to all files containing the same data. *Possibility of inconsistency* 

Waste of storage space:

Reliability and better local control

de Bipin C. DESAI





#### Pros & Cons of DBMS

Reduce data redundancy and avoiding inconsistency Provide Concurrent access Offer Centralized control

- security(appropriate authorization and its control),
  - integrity(constraints and their enforcements)
  - reliability(backups and replication)

Data abstraction and independence

# First Step: Data Models

- <sup>♯</sup>Data Model: concept to describe data
- <sup>★</sup>Schema: description of a collection of data using a specific data model
- \*Relational Model: Based on the concept of relation(table with rows and columns)

d Bipin C. DESAI

#### A Data Model is a collection of concepts for describing

Entities(objects) and relationships among them

Expressing the semantics and constraints from the real world

Object-Based Modeling Techniques

Entity-Relationship (ER) Model

Object-Oriented (OO) Model

**Record-Based Models** 

Hierarchical Model: used by earliest DBMS – IBM's IMS

Network Model: second generation DBMS - DBTG

Relational Model: the first based on theory - relations

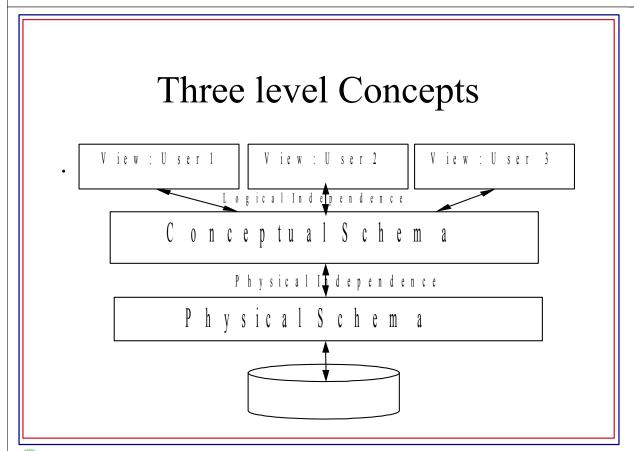
(RA, RC, Datalog)

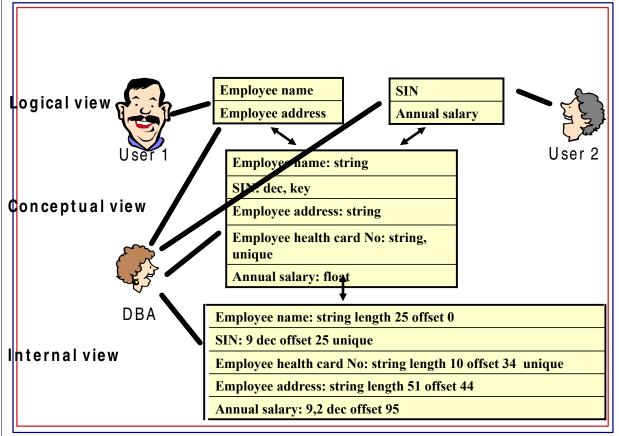
Employee Name Employee Phone Number Employee Name Employee SIN Employee Salary

Employee Name
Employee Phone Number
Employee SIN
Employee Address
Employee Annual Salary
Employee YTD Salary

Employee Name string
Employee Phone Number digits
Employee SIN digits
Employee Address string
Employee Annual Salary money units
Employee YTD Salary money units

de Bipin C. DESAI





(d) Bipin C. DESAI

# Three levels & Independence

- User View: How users view data derived from conceptual view-
- Conceptual Schema: Logical structure of the database
- Physical Schema: The actual files and indices used
- Schema defined using DDL

**Data Independence**: modify definition of schema at one level without affecting a schema definition at a higher level.

**Logical Data Independence**: modify logical schema without causing application programs to be rewritten

adding new fields to a record or changing the type of a field

**Physical Data Independence**: modify physical schema without causing logical schema or applications to be rewritten

changing file structure from sequential to direct access

d Bipin C. DESAI

# University Database

**External Schema:** 

Course\_Enrol(C#:char, Number:int);

Conceptual Schema:

Student(S#, Name, Dept)

Course(C#, Cname, Credits)

Enrollment(C#, S#, grade)

Physical Schema:

files, indexed on S#, C#, etc

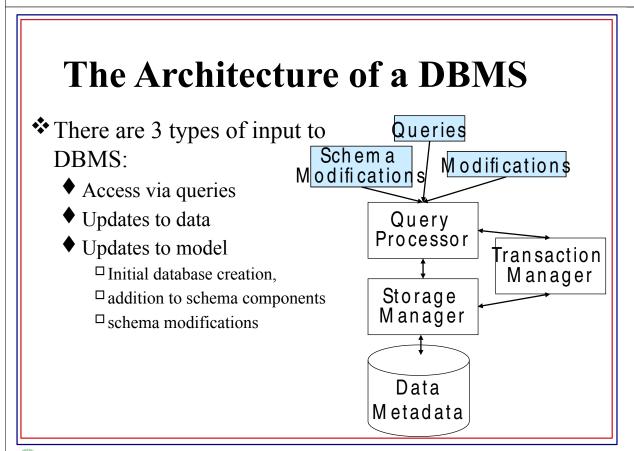
A database schema is a description of a particular collection of data, using a given data model

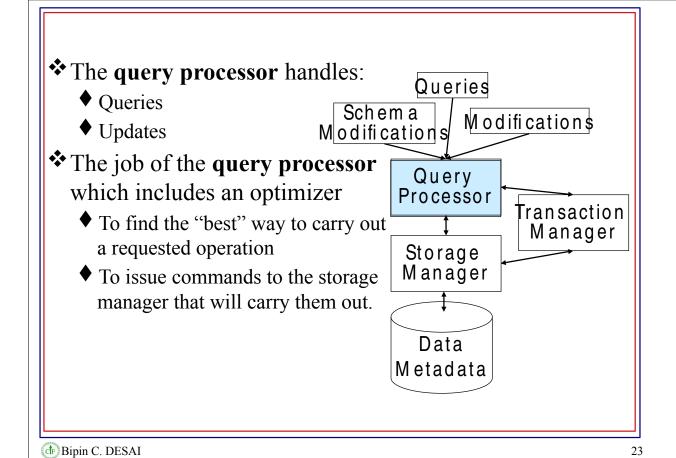
Part of a schema for a university. database in relation model would contain among others, the following:

Students (sid, name, department, dob, address)

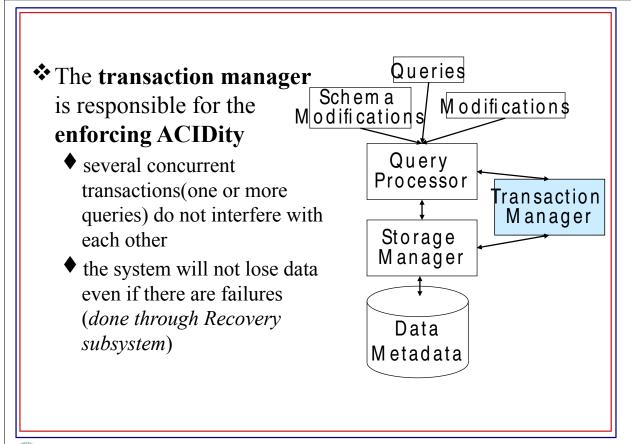
An in	stance o	f a datak	ase sch	em a is th
sid	name	department	dob	address
111222	John Smith		<u>u12-61-82</u> a Si	22 Pine, #1203
22233341C	DAI ABFORPHO I	n <sub>E</sub> te in tim	<b>Q</b> 1-08-73	2000 St. Marc
3334445	Youwong Li	CS	23-11-79	1150 Guy

d Bipin C. DESAI 2

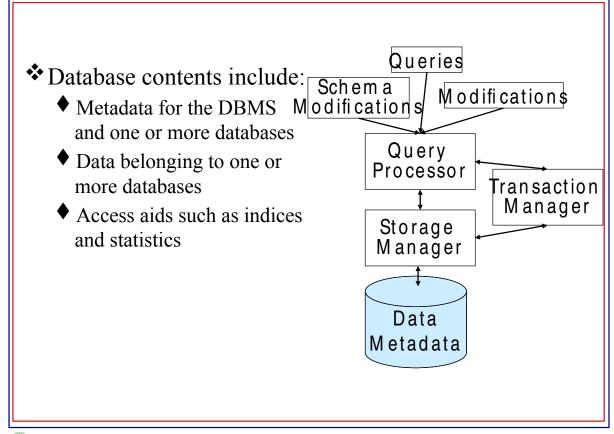


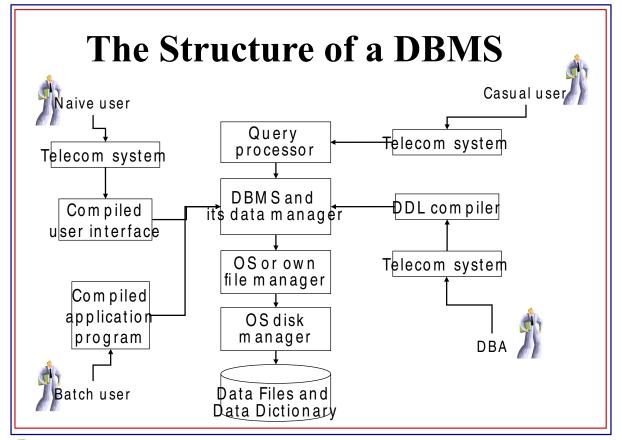


The job of the storage Queries Schema Modifications manager is Modifications To obtain requested Query information **from** the data Processor Transaction storage Manager To modify the information to Storage the data storage when Manager requested. Data Metadata



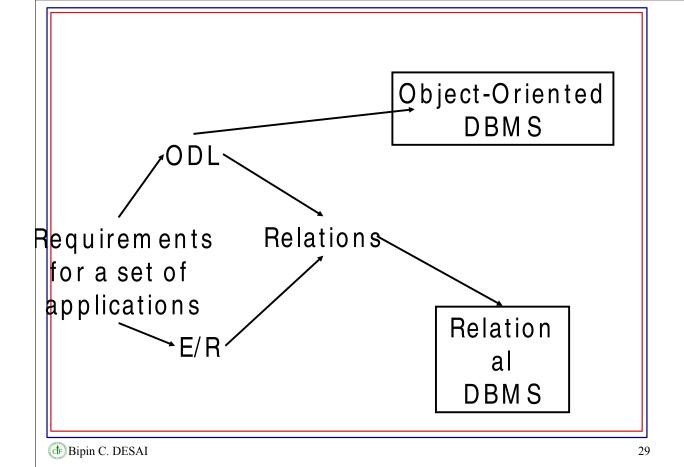
de Bipin C. DESAI 25

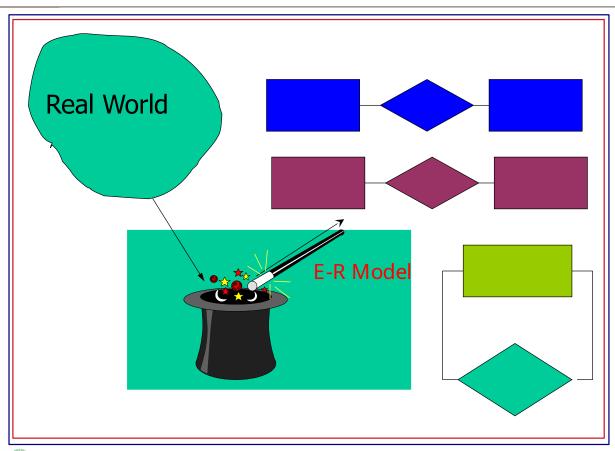


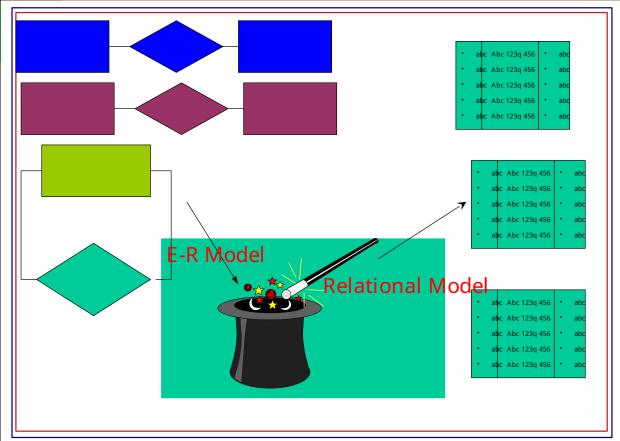


d Bipin C. DESAI 27

# Database Design Process and Conceptual Design







de Bipin C. DESAI

#### Relational Model

In this model, the data is organized in relations (tables)

Relational database schema: DDL component of SQL

set of table names

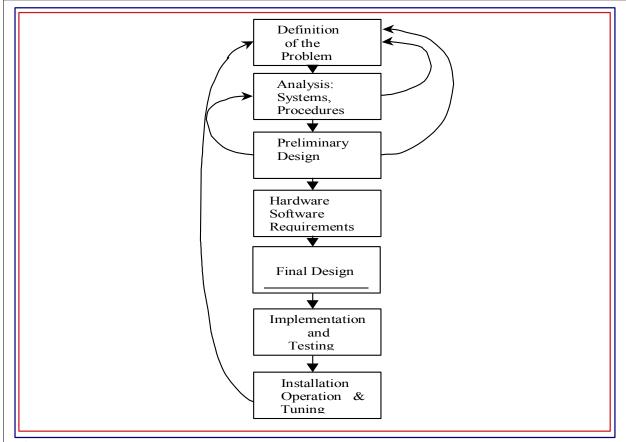
list of attributes for each table and their properties

Examples of tables from a university database:

Student: stud\_num ber, nam e, address, program

Department: name, budget\_code, room, phone

Bipin C. DESAI Se: name number credits



di Bipin C. DESAI

### **Database Design Process**

- ☐ <u>Definition of the problem</u>
- □ <u>Study underlying applications</u>(Procedure Manuals, Interviews etc.)
  - ♣ What are the <u>entities</u> and <u>relationships</u> involved?
  - ♣ What details about them should be in the database?
  - ♣ What are the *procedures, business rules, constraints*?
  - ♣ Who are the users? What do they need?
- ☐ <u>Preliminary Conceptual design</u>:
  - ER Model

#### **Database Design Process**

- □ Software/Hardware Requirements
  - \*UML for software design
- ☐ <u>Final Design: Schema Refinement: (Normalization)</u>
  - \*Check relational schema for redundancies and related anomalies.
  - \* External Schemas, indices, views, access methods
- ☐ Application programs, forms, reports, user interfaces
- ☐ *Implementation and testing*
- ☐ *Installation and Tuning:* 
  - Data Distribution, Physical re-design
  - Performance, Security, Backup & Recovery.

di Bipin C. DESAI 35

sin

#### **ER Model**

- Entity: Real-world object distinguishable from other objects.
  - ◆ An entity is described using a set of *attributes*.
- Entity Set: A collection of similar CREATE TABLE Employees (sin CHAR(9),
  Page CHAR(25)
  - ♦ All entities in an entity set have the same set of attributes.

(SIN CHAR(9), name CHAR(25), grade INTEGER, <u>PRIMARY KEY (sin)</u>)

name

**Employees** 

- ◆ Each entity set has a <u>key.</u>
- ♦ Each attribute has a <u>domain.</u>
- Can map entity set to a relation

de Bipin C. **easily.** 

36

grade

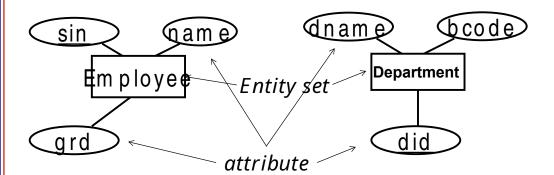
de Bipin C. DESAI 37

The Extra field contains any additional information that is available about a given column.

The value is auto\_increment for columns that have the AUTO\_INCREMENT attribute and empty otherwise.

de Bipin C. DESAI





All employees, and departments have the same set of properties(attributes)

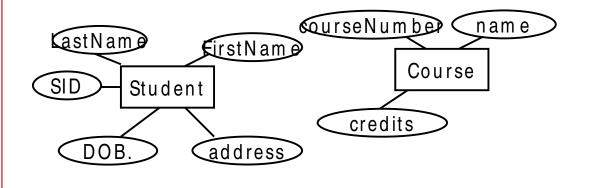
To distinguish one instance of an entity in an entity set from others, we introduce an identifying attribute
This is the primary key and it is underlined

**Entity** – Real world object distinguishable from other objects of the same type

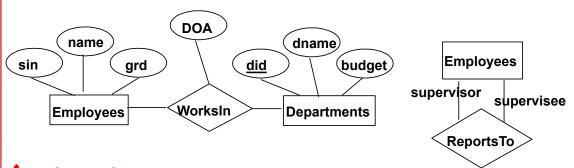
**Entity Set** -- A collection of similar entities: all have same set of properties

ODL:

Object corresponds to entity Class corresponds to entity set

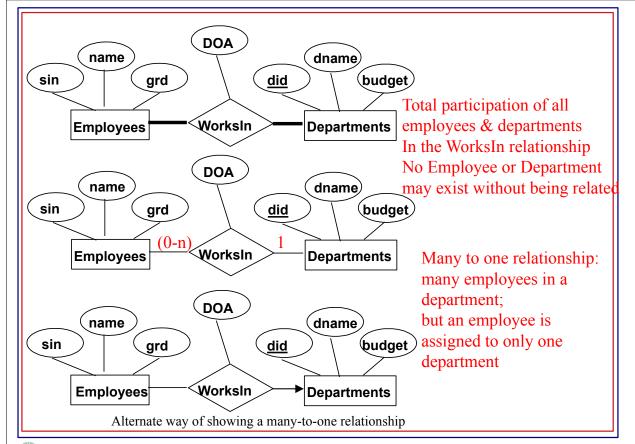


d Bipin C. DESAI 41

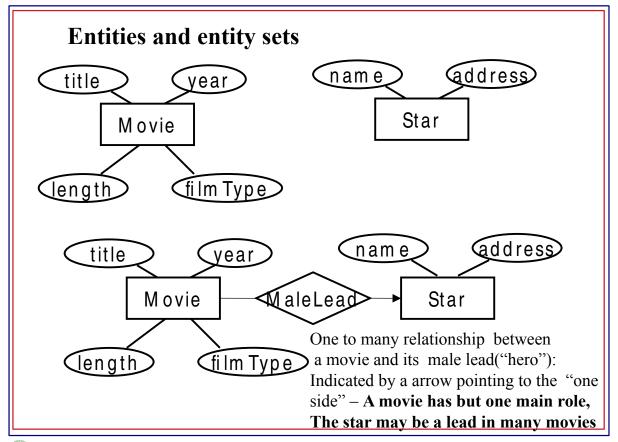


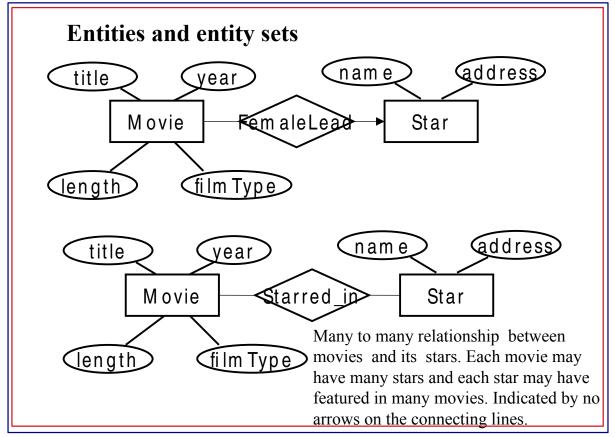
- \* Relationship:
  - ◆ Association among 2 or more entities.
- \* <u>Relationship Set</u>: Collection of similar relationships.
  - An n-ary relationship set R expresses an association among n entity sets E1 ... En; each relationship in R involves entities e1 ∈ E1, ..., en ∈ En

Same entity set could participate in different relationship sets, or in different "roles" in same set.

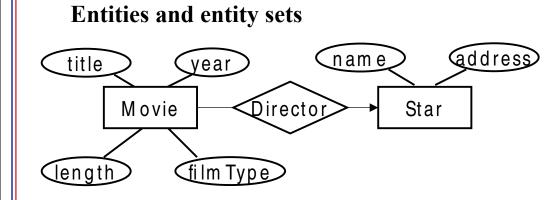


di Bipin C. DESAI 43

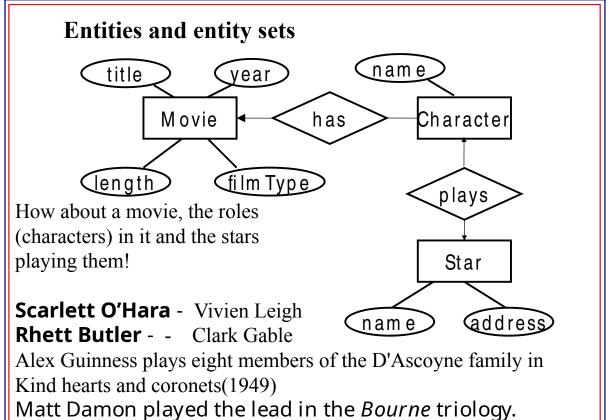




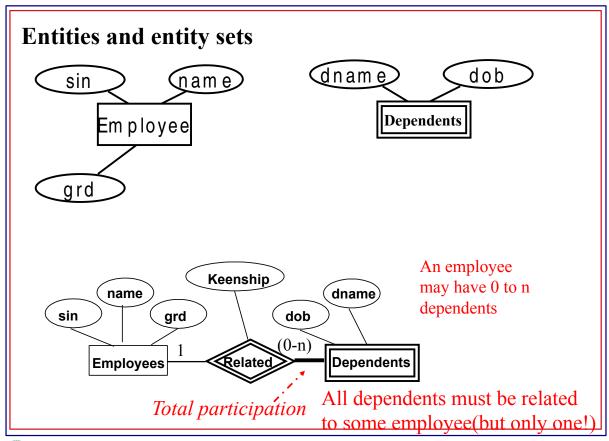
di Bipin C. DESAI 45

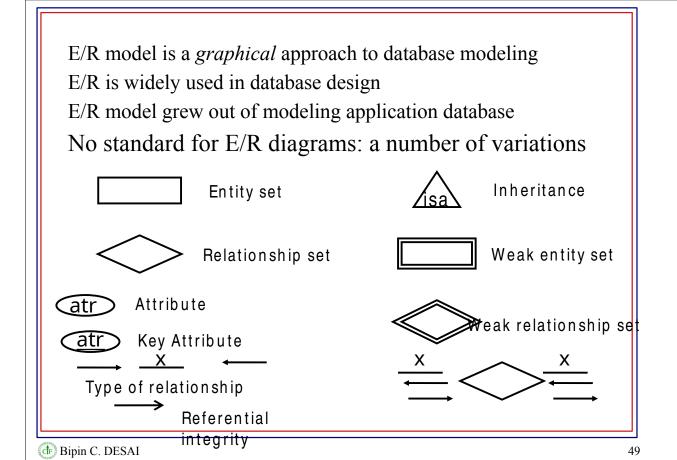


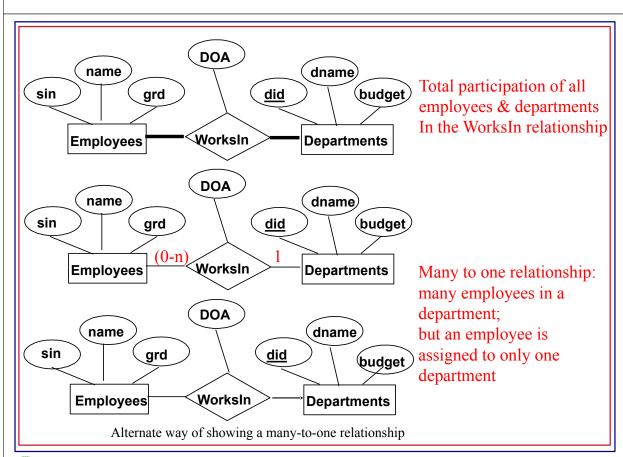
How about a movie and the roles(characters) in it and the stars playing them!

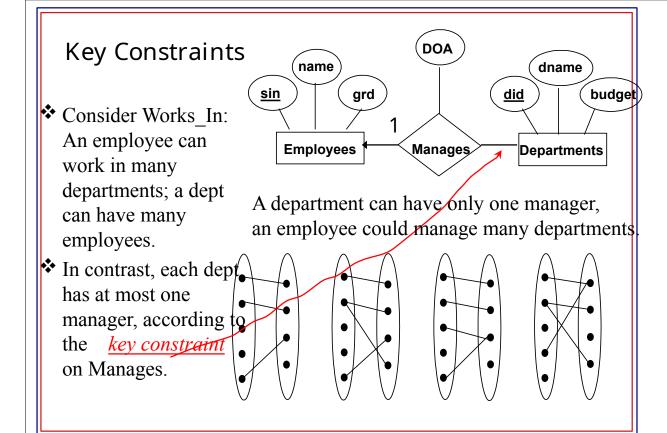


(d) Bipin C. DESAI









di Bipin C. DESAI 51

Relationship sets can have <u>attributes</u>

In translating a relationship set to a relation, attributes of the relation must include:

> Keys for each participating entity set (as foreign keys).

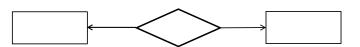
> > □ This set of attributes forms **superkey** for the relation.

♦ All descriptive attributes.

CREATE TABLE WorksIn

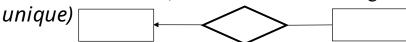
(sin CHAR(9),
did INTEGER,
DOA DATE,
PRIMARY KEY (sin, did),
FOREIGN KEY (sin)
REFERENCES Employees,
FOREIGN KEY (did)
REFERENCES Departments)

If a relationship is 1-to-1 primary key is from *either* of the '1' side, the other side is a *foreign* key



If a binary relationship between two entity sets is 1-to-1,

- the primary key of the relationship is the key of the entity from either of the '1' side, the other side is a foreign key(would be



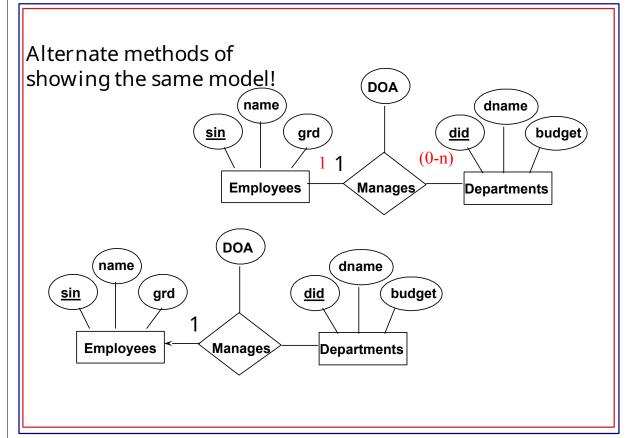
If a binary relationship between two entity sets is 1-to-many,

- the primary key of the relationship is from the 'm' side, the '1' side is the foreign key (would be unique)



- the primary key is composite, consisting of the primary key of the entities from each side of the relationship

d Bipin C. DESAI 53



CREATE TABLE Manages

Map relationship to a table:
(sin CHAR(9), did INTEGER, DOA DATE.

◆ Note that did is the PRIMARY KEY (did), key now! FOREIGN KEY (sin) R

FOREIGN KEY (sin) REFERENCES Employees, FOREIGN KEY (did) REFERENCES Departments)

♦ Separate tables for Employees and

Departments. CREATE TABLE DeptMgr

Since each department has a unique manager, we could instead combine Manages and Departments.

(did INTEGER, dname CHAR(20), budget REAL, sin CHAR(9), DOA DATE, PRIMARY KEY (did),

FOREIGN KEY (sin) REFERENCES Employees)

d Bipin C. DESAI

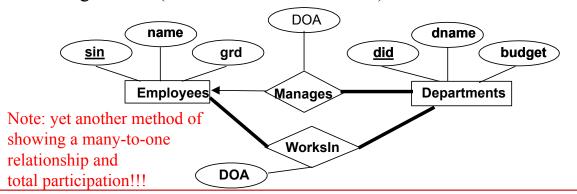
55

#### **Participation Constraints**

- ❖ Every department has a manager (a business rule) ⇒ participation constraint:
  - ◆ The participation of Departments in Manages is *total*

(all instances of Department must have a manager; participation of Employees is partial i.e., not all employees are managers).

Every *did* value in Departments table must appear in a row of the Manages table (with a non-null *sin* value!)

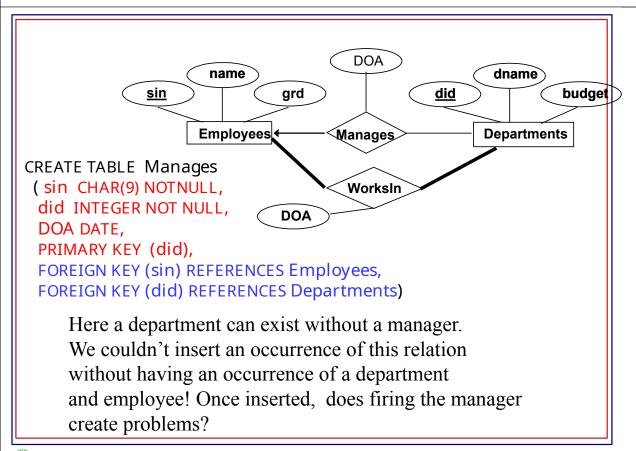


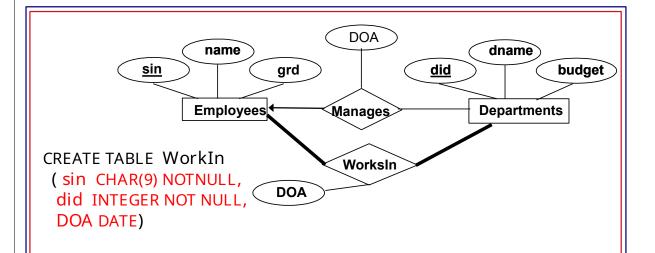
# Participation Constraints: SQL

A participation constraints involving one entity set in a binary relationship, can be expressed as follows without resorting to CHECK constraints.

```
CREATE TABLE DeptMgr Every department must (did INTEGER, have a manager! dname CHAR(20), budget REAL, sin CHAR(9) NOT NULL, DOA DATE, PRIMARY KEY (did), FOREIGN KEY (sin) REFERENCES Employees, ON DELETE NO ACTION)
```

de Bipin C. DESAI 57

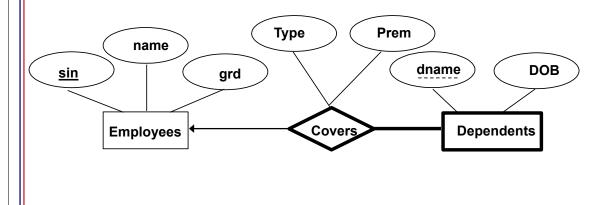




To ensure total participation of department in WorkIn, each did value must be in at least one tuple of WorkIn: enforced by assertion

db Bipin C. DESAI 59

- A weak entity can be identified uniquely only by considering the primary key of another strong-owner entity.
  - ◆ Owner entity set and weak entity set must participate in a one-tomany relationship set (1 owner, many weak entities).
  - ♦ Weak entity set must have total participation in this *identifying* relationship set.



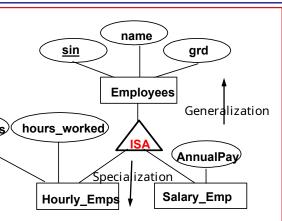
- Weak entity set and identifying relationship set are translated into a single relation.
  - ♦ Weak entity ⇒ total participation
  - When the owner entity is deleted, all owned weak entities must also be deleted.

CREATE TABLE Covers (dname CHAR(20), DOB DATE, Type INTEGER, Cost FLOAT, sin CHAR(9) NOT NULL, PRIMARY KEY (dname, sin), FOREIGN KEY (sin) REFERENCES Employees, ON DELETE CASCADE)

Bipin C. DESAI

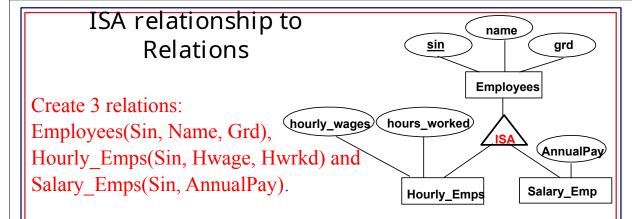
#### Generalization, Specialization

- \* If we declare A ISA B, every A entity is also considered to be a B entity. However, not implemented always: hourly\_wages
- \* Overlap constraints: Can two subclasses contain the same instance of an entity? Can Carole be an Hourly\_Emps as Covering constraints: Does every well as a Salary Emps? (Allowed/disallowed)



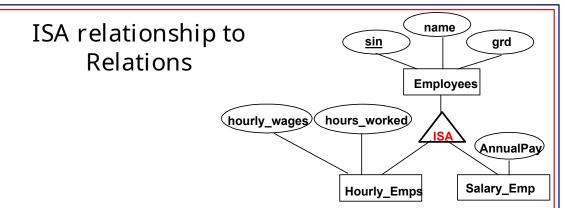
Employees entity also have to be an Hourly\_Emps or a Contract\_ Emps entity? (Yes/no)

- Reasons for using ISA relationship:
  - To add attributes specific to a subclass.
  - ◆ To identify subset of an entity set that participate in a relationship.



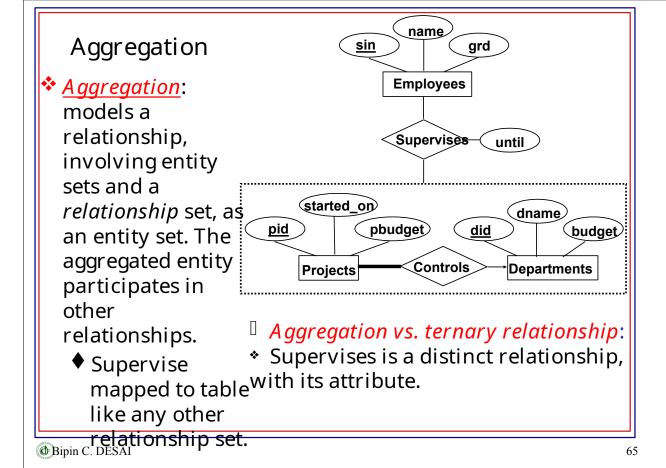
- \* Every employee is recorded in Employees. For hourly employees, value for additional attribute are recorded in Hourly\_Emps; if referenced Employees tuple is deleted, Hourly\_Emps tuple must also be deleted.
- Queries involving all employees easy, those involving just Hourly\_Emps require a join with Employee to get inherited attributes.

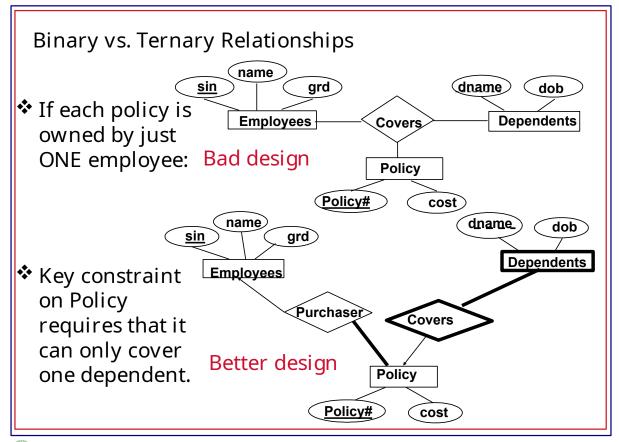
d Bipin C. DESAI 63



- Create two relations:
- Hourly\_Emps(Sin, Name, Grd, HWrkd, Hwages) and Salary\_Emps (Sin, Name, Grd, AnnualPay).

Each employee must be in one of these two subclasses. *All employees require accessing Two relations* 





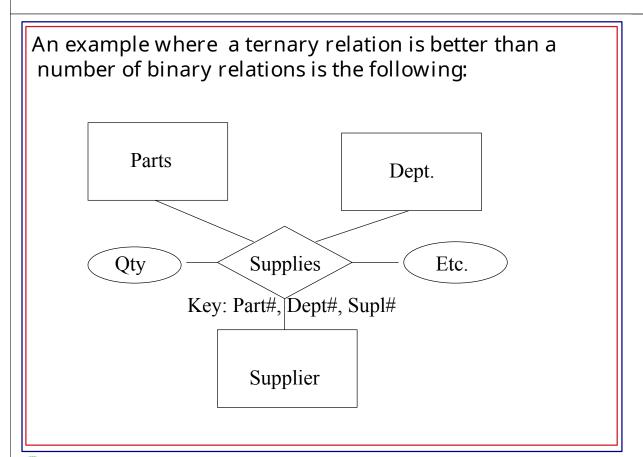
The key constraints allow us to combine Purchaser with Policy and Covers with Dependents.

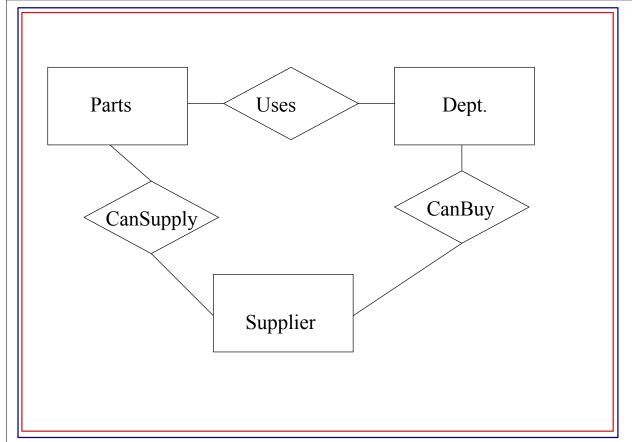
Participation constraints lead to NOT NULL constraints.

CREATE TABLE Policy (
policy# INTEGER,
cost REAL,
sin CHAR(9) NOT NULL,
PRIMARY KEY (policy#).
FOREIGN KEY (sin) REFERENCES
Employees,
ON DELETE CASCADE)

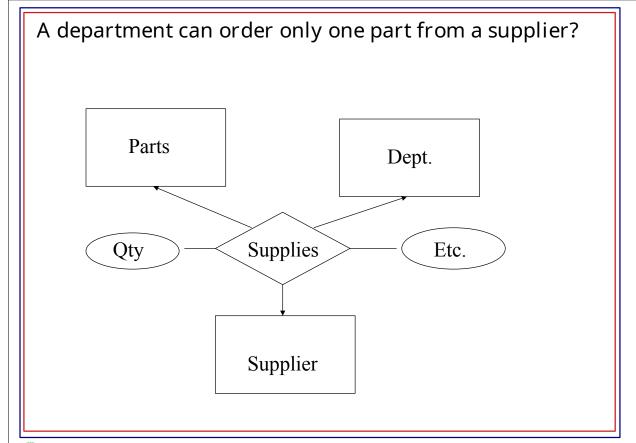
CREATE TABLE Dependents (
dname CHAR(20),
dob DATE,
policy# INTEGER,
PRIMARY KEY (dname, policy#).
FOREIGN KEY (policy#)
REFERENCES Policy,
ON DELETE CASCADE)

© Bipin C. DESAI

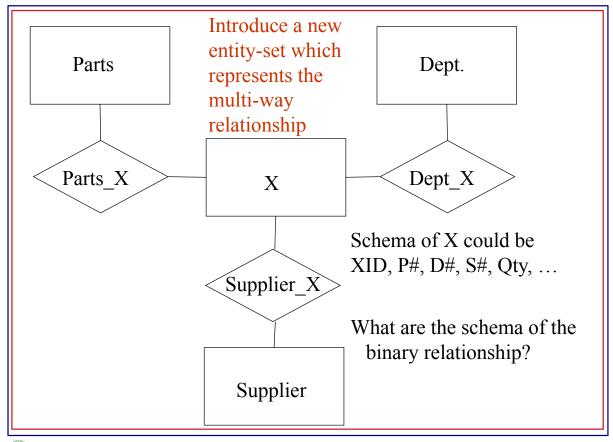




- Supplies relates entity sets Parts, Departments and Suppliers, and has descriptive attributes price, quantity etc.
- ❖ A number of binary relationships may not convey the semantics
- Supplier ``CanSupply" Part, Dept. ``Uses" Part, and Dept. ``Can Buy" from Supplier does not imply that Dept has a PO to buy Part from S.
  - ♦ How do we record the following: which part, quantity price?



Bipin C. DESAI 7



## Constraints Beyond the ER Model

# \* Functional dependencies:

- ◆ A department can order only one part from a given supplier.

  □ Can't express this in ternary Supplies relationship.
- ◆ Normalization refines ER design by considering FDs.

## \* Inclusion dependencies:

- ◆ Special case: Foreign keys (ER model can express these).
- ◆ e.g., At least 1 person must report to each manager. (Set of sin values in Manages must be subset of supervisor\_ssn values in Reports To.)

#### General constraints:

• e.g., Manager's discretionary budget less than 10% of the combined budget of all departments he or she manages.

# Databases – the generations

#### **Notes**



Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

🕩 Bipin C Desai

#### FIRST GENERATION

1950s – Refinement of storage media, magnetic tape, drums, disks

Early 1960s: Disk access method based on

Index Sequential Access Method(ISAM)

Mid 1960s: Emergence - Information Management System(IMS)-IBM

developed in 1966 along with NASA(Rockwell and Caterpillar)

to support the Apollo/Saturn V program

Current version is IMS 15.4 and runs on IBM z platform

It is still being marketed, used in banking etc.

promises  $> 250*10^9$  transactions per day

1959 : CODASYL(Conf./Committee on Data Systems Languages)

later to become **Database Task Group** (DBTG),

**DBTG** developed the network model and its implementation

Integrated Data Store (IDS),

Integrated Database Management System (IDMS )

both still marketed and supported.

#### **SECOND GENERATION**

1970 Codd's paper about relations

1973/1974 Ingres(UC Berkley, M. Stonebraker, E. Wong)

System R(IBM), Berkley/DB (Sleepy Cat Software, Oracle)

QUEL, SEQUEL(Ingres) and SQL(System R)

1978 Oracle

1981 Informix (IBM)

1984 System R(IBM)

1987 Postgres

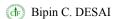
1993 mSQL (mini SQL by D. Hughes)

mSQL used in the development of early dynamic Web applications including CrsMgr and ConfSys

1995 MySQL - bought by Sun in 2008 price- \$1billion

- Sun was taken over by Oracle

2009 Mariadb – a fork of MySQL



3

#### THIRD GENERATION

2004 MapReduce paradigm shift to lower level!

Map(distribute tasks to nodes to filter local data) and then

Reduce(process result in parallel)

2005 Hadoop (Apache)

2008 Cassandra, Hbase,

2009 MongoDB

## Simple SQLPlus & SQL

Bipin C. DESAI

Bipin C. DESAI

2

## Getting & Installing {Apache, Oracle, PHP} or, XAMPP

#### Consult:

http://www.oracle.com/technology/tech/php/htdocs/inst\_php\_apache\_windows.html or whatever is the currrent URL For Oracle you need to register with OTN

MySQL/Mariadb

https://www.apachefriends.org/download.php

The projects are to be demonstrated on one of the systems in our labs. So if you develop the projects on your own systems, make sure you could:

- Upload all the code to CrsMgr
- Have it run on one of AITS systems which has one of the above configurations
- These notes uses Oracle, MySQL, MariaDB

(h) Bipin C. DESAI

## Connecting to SQLPlus

SQLPlus is a "user friendly interface" to ORACLE SQL to be used interactively.

You need Oracle USERID/PASSWORD and appropriate permission to a Oracle DB.

May connect remotely using a secure shell (e.g., Putty)

```
[alpha:bcdesai] 101 => sqlplus

SQL*Plus: Release 9.0.1.0.0 - Production on Mon Sep 20 10:04:53 2004

(c) Copyright 2001 Oracle Corporation. All rights reserved.

Enter user-name: bcd_orcl
Enter password:

Connected to:
Oracle9i Release 9.2.0.3.0 - Production
JServer Release 9.2.0.3.0 - Production
```

(d) Bipin C. DESAI

Tablespace created.

SQL> \_

7

Typically - start database (unless it has been installed as service which starts on boot)

From Start select RunSQL command line

Connect to oracle:

Run SQL Command Line

SQL\*Plus: Release 10.2.0.1.0 - Production on Fri May 23 16:46:37 2008

Copyright (c) 1982, 2005, Oracle. All rights reserved.

SQL> connect bcdesai
Enter password:
Connected.
SQL> create tablespace bcd
2 logging
3 datafile 'c:\Oracle\oradata\bcd'
4 size 32m
5 autoextend on
6 next 32m massize 512m
7 extent management local;

Download and install Oracle (the version changes over time)

Cap Birjin C. DESAI

create table student (SID NUMBER(7) primary key not null, SNAME VARCHAR2(20), To execute a text file containing sql MAJOR CHAR(4), statements interactively from the sql YEAR NUMBER(1), prompt use @ followed by the full BDATE DATE) path to file tablespace bcd pctfree 2; sql>@student.sql 🧬 sunset.cs.concordia.ca - PuTTY SQL> create table student (SID NUMBER(7) primary key not null, SNAME VARCHAR2(20), MAJOR CHAR(4), YEAR NUMBER(1), BDATE DATE); Table created. SQL>

Bipin C. DESAI

9

## Connecting to MySQL/MariaDB

MySQl/Mariadb has a simpler text based interface used for connecting to the database running locally or on a server accessed using a terminal emulator Putty is one used in WinX

Again the DB server must be running and one needs a user ID and password for the database to be used

shell> mysql –u username –p password

If the ID/PW are correct, one gets the prompt from the database

de Bipin C. DESAI

10

Enter password:

Welcome to the MariaDB monitor. Commands end with ; or \g.

Your MariaDB connection id is 96773

Server version: 10.3.17-MariaDB MariaDB Server

Copyright (c) 2000, 2018, Oracle, MariaDB Corporation Ab and others.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

MariaDB [(none)]> connect test;

Reading table information for completion of table and column names You can turn off this feature to get a quicker startup with -A

Connection id: 29348 Current database: test

Bipin C. DESAI

11

mysql> create table student

(SID DECIMAL(7) primary key not null,

SNAME VARCHAR (20), MAJOR CHAR(4), YEAR DEC(1),

BDATE DATE);

To execute a text file containing sql statements interactively from the sql prompt use @ followed by the full path to file

sql>@student.sql in MySQL use "source student.sql"

mysql> desc student;

+		-+-		+-		-+-		+-	+
					_		Default		
+		-+-		+-		+-		+-	+
SID	decimal(7,0)		NO		PRI		NULL		1
SNAME	varchar(20)		YES				NULL		1
MAJOR	char(4)		YES				NULL		
YEAR	decimal(1,0)		YES				NULL		
BDATE	date		YES				NULL		
+		-+-		+-		+-		+	+

5 rows in set (0.00 sec)

(f) Bipin C. DESAI

```
Inserting Data in a table — table must exist!

| Substance | Subst
```

G Bipin C. DESAI

```
MariaDB [test] > \! tcsh -- escape to interative shell (tcsh)
101 => emacs -nw students.sql
104 => more students.sql
insert into student values(10, "Dupont", 'ENGL', 1, '1980-05-13');
insert into student values(13, 'Kelly', 'SENG', 4,'1980-08-12');
insert into student values(14, 'Jack', 'CSAP', 1, '1970-02-12');
105 \Rightarrow exit
exit
MariaDB [test]>system cat students.sql;
create table student
(SID DECIMAL(7) primary key not null,
SNAME VARCHAR (20),
MAJOR CHAR(4),
YEAR DEC(1),
BDATE DATE):
insert into student values(10, "Dupont", 'ENGL', 1, '1980-05-13');
insert into student values(13, 'Kelly', 'SENG', 4,'1980-08-12');
insert into student values(14, 'Jack', 'CSAP', 1, '1970-02-12');
MariaDB [test]>
```

(dr) Bipin C. DESAI

```
students.sql-GNU Emacs at Confsys (on Confsys)

File Edit Options Buffers Tools SQL Help

Create table student

(SID DECIMAL(7) primary key not null,

SNAME VARCHAR (20),

MAJOR CHAR(4),

YEAR DEC(1),

BDATE DATE);

insert into student values(10, "Dupont", 'ENGL', 1, '1980-05-13');

insert into student values(13, 'Kelly', 'SENG', 4,'1980-08-12');

insert into student values(14, 'Jack', 'CSAP', 1, '1970-02-12');

-:--- students.sql All L8 (SQL[ANSI] +2)

Use +,-,0 for further adjustment
```

(d) Bipin C. DESAI

```
MariaDB [test]>
MariaDB [test]> source students.sql;
Query OK, 1 row affected (0.028 sec)
Query OK, 1 row affected (0.050 sec)
Query OK, 1 row affected (0.050 sec)
MariaDB [test]> select * from student;
+----+----
 SID | SNAME | MAJOR | YEAR | BDATE
 ____+
                       1 | 1980-05-13 |
  10 | Dupont | ENGL |
                       4 | 1980-08-12 |
  13 | Kelly | SENG
                       1 | 1970-02-12 |
            | CSAP
  14 | Jack
  ---+-----
3 rows in set (0.001 sec)
MariaDB [test]>
```

#### Find all students (ORACLE)

SQL> select \* from student;

SID	SNAME	MAJO	YEAR	BDATE
8	Brenda	COMP	2	13-AUG-77
10	Dupont	ENGL	1	13-MAY-80
13	Kelly	SENG	4	12-AUG-80
14	Jack	CSAP	1	12-FEB-77

SQL>column major format a5

SQL>column sid format 9,9

format not available in MySQL

SQL>column sname format a12 SID SNAME MAJOR YEAR BDATE SQL>column major format a5 8 Brenda SQL>column year format 999 1,0 Dupont 1,3 Kelly SQL>column bdate format a12 1,4 Jack

COMP 2 13-AUG-77 1 13-MAY-80 ENGL 4 12-AUG-80 SENG CSAP 1 12-FEB-77

Bipin C. DESAI

```
MariaDB [test]> select * from student;
+----+----
  sid | sname | major | year | bdate
   8 | Brenda | COMP
                          2 | 1997-08-13 |
  10 | Dupont | ENGL
                          1 | 1980-05-13 |
  13 | Kelly
                          4 | 1980-08-12
               SENG
       Jack
                          1 | 1970-02-12
              | CSAP |
   14
4 rows in set (0.001 \text{ sec})
```

```
select s.sname
from student s
where to_date(s.bdate) like '%13%';

select s.sname
from student s

where s.bdate like '%13%';

+------
Brenda
Dupont
| sname |
+-----+
| sname |
| Dupont |
| Dupont |
| Dupont |
| Towns in set (0.000 sec)
```

(d) Bipin C. DESAI

19

#### Find students born in August

SQL script: month.sql

```
select s.sname

from student s

where s.bdate like '%-08-%';

where to_date(s.bdate) like '%AUG%';

sname |
| sname |
| +-----+
| SNAME | Brenda |
| Kelly |
| Brenda |
| Kelly |
| 2 rows in set (0.000 sec
```

```
select s.sname
from student s
where to_date(s.bdate) like '%77%';

select s.sname from student s

SNAME

SNAME

Sname

| sname |
| sname |
| Left |
| Dupont |
| Kelly |
| SQL script: year.sql |
| 2 rows in set (0.001 sec)
```

di Bipin C. DESAI

```
create table dept
(DEPT CHAR(20) not null,
CODE CHAR(4) primary key not null);

insert into dept values('Computer Science', 'COMP');
insert into dept values('Decision Science', 'DISC');

create table deptmajor
(CODE CHAR(4),
MAJOR CHAR(20),
primary key (CODE, MAJOR))

insert into deptmajor values('COMP', 'COTH');
insert into deptmajor values('COMP', 'SENG');
insert into deptmajor values('COMP', 'CSAP');
insert into deptmajor values('DISC', 'OPRS');
```

```
create table course
(CNAME CHAR(20),
CNUMBER CHAR(8) primary key NOT NULL,
CREDITS NUMBER(2),
ODEPT CHAR(4),
foreign key (ODEPT) references dept(code)
on delete cascade)

insert into course values('C++','COMP248',3,'COMP');
insert into course values('DATA STRUCTURES ','COMP352',3,
'COMP');
insert into course values('OPERATING SYSTEMS','COMP346',4,'COMP');
insert into course values('DATABASE','COMP353',4,'COMP');
insert into course values('Operation Research','DISC253',4,'DISC');
```

de Bipin C. DESAI

```
create table crs_section
(SECID NUMBER(6) primary key NOT NULL,
COURSE_NUM CHAR(8),
SECTION CHAR(2),
SEMESTER CHAR(4),
YEAR CHAR(4),
SCHEDULE CHAR(10),
ROOM CHAR(7));

insert into crs_section values
(85,'COMP352','A','FALL', '1998','TH16001715','H123');
insert into crs_section values
(90,'COMP353','B','FALL','1999','MW08451000','H631');
insert into crs_section values
(95,'DISC253','B','FALL','1999','MW10151130','H631');
```

```
create table prereq
(COURSE_Number CHAR(8),
PREREQ CHAR(8), primary key (course_number, prereq));
insert into prereq values('COMP353','COMP352');

insert into prereq values('COMP353','COMP346');
insert into prereq values('COMP352','COMP248');

create table enrollment
(STUDENT_NUMBER NUMBER(3) not null,
SECTION_ID NUMBER(6) not null, GRADE CHAR(1),
primary key(student_number, section_id));
insert into enrollment values(8,85,null);
insert into enrollment values(10,90,null);
insert into enrollment values(14,90,null);
insert into enrollment values(14,90,null);
insert into enrollment values(14,95,null);
```

Bipin C. DESAI

25

```
Find details of studs. taking a course offered by the "DISC" dept.
```

```
select s.SID, s.SNAME, s.MAJOR, s.YEAR, s.BDATE from student s, dept d, course c, crs_section r, enrolment e where c.ODEPT=d.CODE and r.COURSE_NUM=c.CNUMBER and r.SECID=e.SECTION_ID and e.STUDENT_NUMBER = s.SID and d.CODE= 'DISC';
```

SQL script: ex-select3.sql

#### Find student who are registered in a course offered by their majoring dept.

(d) Bipin C. DESAI

```
Find students who are currently registered.
```

SID SNAME	MAJOR YEA	R BDATE
8 Brenda	COMP	2 13-AUG-80
1,0 Dupont	ENGL	1 13-MAY-80
1,4 Jack	CSAP	1 12-FEB-77

```
select s.SID, s.SNAME, s.MAJOR, s.YEAR, s.BDATE from student s, dept d, course c, crs_section r, enrolment e where c.ODEPT=d.CODE and r.COURSE_NUM=c.CNUMBER and r.SECID=e.SECTION_ID and e.STUDENT_NUMBER = s.SID and d.CODE= 'COMP';
```

SQL>	@ex-select2.s	sql		
SID	SNAME	MAJOR	YEAR	BDATE
8	Brenda	COMP	2	13-AUG-80
1,0	Dupont	ENGL	1	13-MAY-80
8	Brenda	COMP	2	13-AUG-80
1,4	Jack	CSAP	1	12-FEB-77

Bipin C. DESAI

29

#### The DUAL table in Oracle

```
SQL> describe dual;

Name

Null? Type

DUMMY

VARCHAR2 (1)

Contains one row and one column. Can be used to put results

SQL> select power(2,10) from dual;

POWER (2,10)

1024

SQL> select to_date(sysdate) from dual;

TO_DATE (S

-----
29-SEP-02
```

G Bipin C. DESAI

30

Bipin C. DESAI

31

## **Editing SQL Buffer**

Command	<u>abbrev.</u>	Operation on crnt. line/all li
append txt	a text	adds text at the end of a line
change /old/new/	c /old/new/	change old to new in a line
change /txt	c /txt	delete text from a line
<b>clear</b> buffer	cl buff	delete all lines in the buffer
delete	del	delete the current line
<b>delete</b> n	<b>del</b> n	delete line n
<b>delete</b> last	<b>del</b> last	delete the last line of the buffer
<b>delete</b> n,m	<b>del</b> n,m	delete lines n - m from buffer
ed	ed	edit the buffer or a file
<b>get</b> file		load file into buffer
input	i	add one or more lines
input txt	i txt	add text as a line
h <mark>o</mark> st		exit temp to OS, exit back to SQLPlus
li <mark>s</mark> t	1	list all lines of buffer
li <mark>st</mark> n	l n (n)	list line n and make it current
list *	*	list current. line

GF Bipin C. DESAI

# **Editing SQL Buffer**

nnmand <u>abbrev.</u> <u>Operation on crnt. line/all line</u>:

m n lm n list lines m – n

ve file sav file save buffer to file

execute the commands in buffer

#### ther useful commands:

ter user userid identified by newpassword

ool nameoffile

#### omments

for multi-line comments \*/
n for a single line comment

**BDATE DATE)** 

comments that can start anywhere in a line up to the eol

(d) Bipin C. DESAI

create table student -- we will create a table for students
(SID NUMBER(7) primary key not null, --not null is redundant
SNAME VARCHAR2(20), --varchar2 is a variable length string
/\*
We will now define
the student's major and year
\*/
MAJOR CHAR(4),
YEAR NUMBER(1),
rem BDATE is his/her birth date
rem It can be used to compute the age which is not stored.

The editor used for the ed command is the default editor set using

setenv EDITOR {emas| vi | gedit | xemacs | ndedit} for tcsh/csh

export EDITOR={ emas | vi | gedit | xemacs | ndedit} for bash

Alternatively, you can set up your editor using the define command:

SQL> define \_editor=emacs

Bipin C. DESAI

35

```
SQL> define _USER=scott
                              ——Show user defined varaibles
SQL> define _PW=tiger
SQL> define
DEFINE CONNECT IDENTIFIER = "cind" (CHAR)
DEFINE _SQLPLUS_RELEASE = "902000100" (CHAR)
DEFINE _EDITOR = "emacs" (CHAR)
DEFINE O VERSION = "Oracle9i Enterprise Edition Release
9.2.0.1.0 - Production
With the Partitioning, OLAP and Oracle Data Mining options
JServer Release 9.2.0.1.0 - Production" (CHAR)
DEFINE O RELEASE = "902000100" (CHAR)
DEFINE _RC = "0" (CHAR)
DEFINE _USER
                = "scott" (CHAR)
DEFINE PW = "tiger" (CHAR)
```

MySQL/Mariadb do not have, to date some of these interactive terminal based features

For most of the current versions of DB server have added web based functions

One can use phpMyadmin mySQLweb

Bipin C. DESAI

3/

# A short introduction to ER & SQL

**Notes** 



Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

🕩 Bipin C Desai

## **Database Languages**

- ❖ A Database Management System (DBMS) provides two types of languages; they may also be viewed as components of the DBMS language:
  - ◆ Data Definition Language (DDL)
    - ☐ Language (notation) for *defining* and modifying a database schema
    - ☐ It includes syntax for declaring tables, indexes, views, constraints, etc.)
  - ◆ Data Manipulation Language (DML)
    - □ Language for *accessing* and *manipulating* the data (organized/stored according to the appropriate data model)

# **Query Languages**

- \* Theoretical:
  - Relational Algebra, Relational Calculus, Datalog
- \* Commercial: SQL
- First there were two: SEQUEL (Ingres) and SQL(R)
- ❖ SQL developed originally at IBM in 1976
  - ♦ First standard: SQL-86
  - ◆ Second standard: SQL-92
  - ◆ Latest standard: SQL-99, or SQL3,

SQL3 standard has over 1,000 pages of

document

- ❖ SQL is the de-facto standard for RDBMS
- The SQL query language components:
  - ◆ DDL (e.g., create)
  - ◆ DML(e.g., select, insert, update, delete)
- Bipin C Desai

3

# **Simple SQL Queries**

A SQL query has a form:

SELECT . . .

FROM . . .

WHERE ...;

The **SELECT** clause defines the schema of the result

The **FROM** clause gives the source relation(s) of the query

The **WHERE** clause is one or more predicates to 'select" the tuples of interest.

The query result is a relation and it is **unnamed**.

ⓑ Bipin C Desai

# **Example "theoretical" SQL Query**

Relation schema:

Course (Cno, Cname, credits)

Query in natural language (English):

Find all the courses stored in the database

Query in SQL:

**SELECT** \*

FROM Course;

Here the "\*" in "SELECT \*" means all attributes in the relation(s) involved.

Bipin C Desai

5

## **More Examples SQL Query**

\* Relation schema:

Movie (title, year, length, filmType)

Ouery in natural language (English):

Find the titles of all movies stored in the database

Query in SQL:

**SELECT** title

FROM Movie;

Relation schema:

Student (SID, FirstName, LastName, Address, GPA)

Query in natural language (English):

Find the SID of every student whose GPA is greater than 3

Query in SQL:

**SELECT SID** 

**FROM** Student

WHERE GPA > 3;

ⓑ Bipin C Desai

#### The "WHERE" clause

The expressions that may follow WHERE are conditions

Standard comparison operators  $\theta$  includes  $\{=, <>, <, >, <=, >=\}$ 

The values that may be compared include constants and attributes of the relation(s) mentioned in FROM clause

Simple expression

 $\mathbf{A} \, \boldsymbol{\theta} \, \text{Value}$ 

Where A, B are attributes and

AθB

A  $\theta$  B  $\theta$  is a comparison operator We may also apply the usual arithmetic operators, +,-,\*,/, etc. to numeric values before comparing them

$$(year - 1930) * (year - 1930) < 100$$

The result of a comparison is a Boolean value TRUE or FALSE Boolean expressions can be combined by the logical operators **AND**, OR, and NOT

Bipin C Desai

- \* Relation schema: **Movie** (<u>title</u>, <u>year</u>, length, filmType)
- Query: Find the titles of all color movies produced in 1950
- Query in SQL:

**SELECT** title

**FROM** Movie

WHERE filmType = 'color' AND year = 1950;

- Query: Find the titles of color movies that are either made after 1970 or are less than 90 minutes long
- Query in SQL:

**SELECT** title

**FROM** Movie

WHERE (year > 1970 OR length < 90) AND filmType = 'color';

Note the precedence rules, when parentheses are absent:

**AND** takes precedence over **OR**,

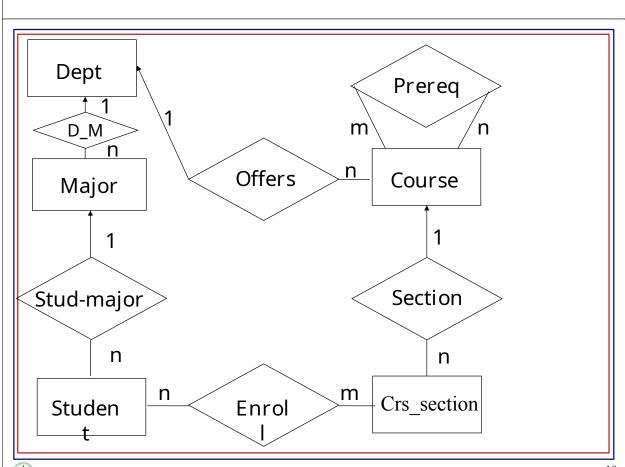
and **NOT** takes precedence over both

Bipin C Desai

# An example of using SQL from E-R to RDBMS

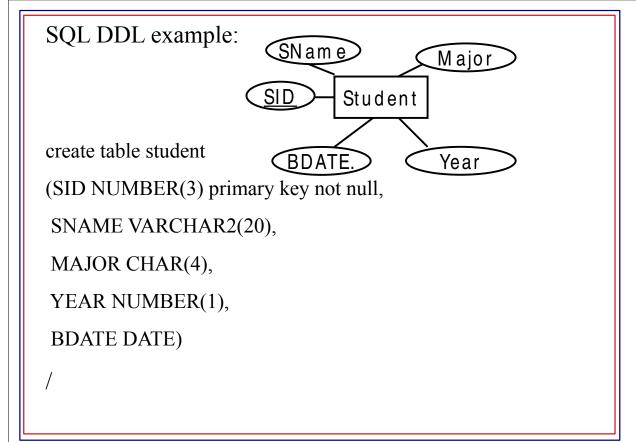
© Bipin C. Desai

Bipin C



Bipin C

١



Bipin C

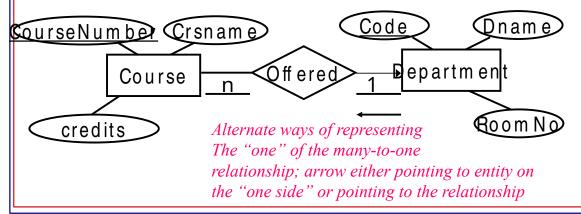
```
insert into student values
(8,'Brenda','COMP','2','13-AUG-77');
(8,'Brenda','COMP',','1977-08-13');
insert into student values
(10,'Mary','ENGL','1','13-MAY-80');
(10,'Mary','ENGL','1','1980-5-13');
insert into student values
(13,'Keily','SENG','4','12-AUG-80');
insert into student values
(14,'Jack','CSAP','1','12-FEB-77');
```

The Bipin C 12

Many to one relationship

A department offers many courses
A given course can be offered by only one department

There is no standard regarding the direction of arrow for the "one" entity.



Bipin C Desai

13

create table course

(COURSE NAME CHAR(20),

COURSE\_NUMBER CHAR(8) primary key NOT NULL,

CREDIT HOURS NUMBER(2),

OFFERING DEPT CHAR(4))

tablespace TUTOR pctfree 5

Note how the relationship w/o an attribute is "merged" with one of the

entity

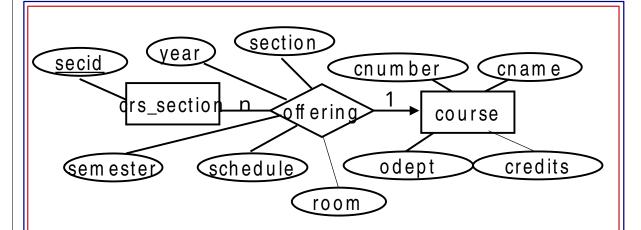
insert into course values
('C++','COMP248',3,'COMP');

insert into course values ('DATA STRUCTURES ','COMP352',3,'COMP');

insert into course values ('OPERATING SYSTEMS','COMP346',4,'COMP');

insert into course values
('DATABASE','COMP353',4,'COMP');

The Bipin C



NB: Here course section (crssection) is really a "weak" entity; However, in most cases it is promoted a "strong" entity by introducing an identifying key attribute section ID (secid)

```
create table crs_section

(SECID NUMBER(6) primary key NOT NULL,

COURSE_NUM CHAR(8),

SECTION CHAR(2),

SEMESTER CHAR(4),

YEAR CHAR(4),

SCHEDULE CHAR(10),

ROOM CHAR(7))

tablespace TUTOR pctfree 2
```

Bipin C

```
insert into crs_section values (85,'COMP352','A','FALL', '1998','TH16001715','H123'); insert into crs_section values (90,'COMP353','B','FALL','1999','MW08451000','H631');
```

® Bipin C

```
create table enrolment

(STUDENT_NUMBER NUMBER(3) not null,

SECTION_ID NUMBER(6) not null,

GRADE CHAR(1),

primary key(student_number, section_id))

tablespace TUTOR pctfree

/

insert into enrolment values(8,85,null);

insert into enrolment values(10,90,null);

insert into enrolment values(8,90,null);

insert into enrolment values(14,90,null);
```

Bipin C

The Bipin C 20

# More examples of using E-R modeling

© Bipin C. Desai

The Bipin C 2

Professors have a SIN, a name, an age, a rank, and a research specialty.

Projects have a project number, a sponsor name (e.g., NSERC),

a starting date, an ending date, and a budget.

Graduate students have a SIN, a name, an age, and a degree program (e.g., M.S. or Ph. D.).

Each project is managed by a professor (principal investigator).

Each project is worked on by one or more professors (co-investigators).

Professors can manage and/or work on multiple projects.

Each project is worked on by one or more graduate students (research assistants).

When a graduate student works on a project, is supervised by a participating professor.

Graduate students can work on multiple projects.

Departments have a department number, a department name, and a main office.

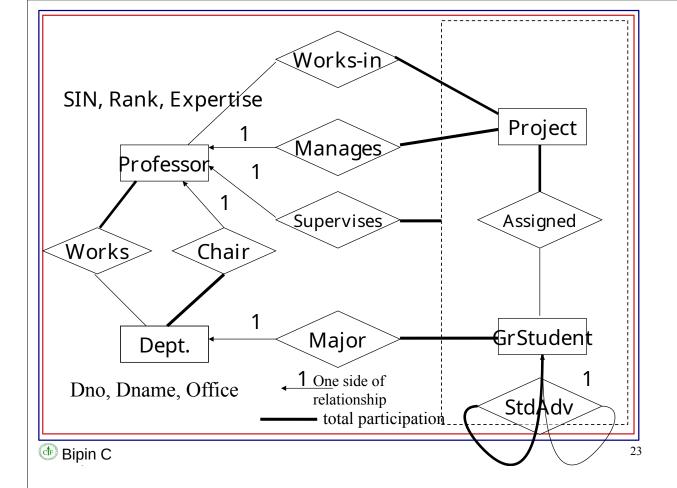
Departments have a chairman who runs the department.

Professors work in one or more departments(%time)

Graduate students have one major department for their degree.

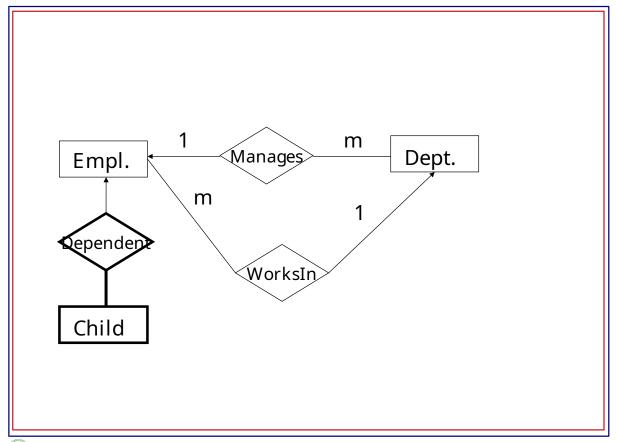
Each graduate student has another, more senior graduate student (student advisor)

d Bipin C 22



In a company database, you need to store information about

employees (SIN,salary and phone),
departments (dno, dname, budget), and
children of employees (with name and age as attributes).
Employees work in departments;
Each employee works in only one department;
A department could have many employees;
Each department is managed by an employee;
Each department has only one manager(an employee);
A manager could manage many departments
A child can only be identified by name
An employee has only one child with a given name
only one parent can declare a child as a dependent



The Bipin C 25

Each musician that records at Notown has an SIN, a name,

an address, and a phone number.

Musicians often share the same address,

no address has more than one phone.

Each instrument that is used in songs recorded at Notown has a name and a musical key

Each album that is recorded on has a title, a copyright date, and an album identifier.

Each song recorded has a title and an author.

Each musician may play several instruments, and

a given instrument may b e played by several musicians.

Each album has a number of songs on it,

but no song may appear on more than one album.

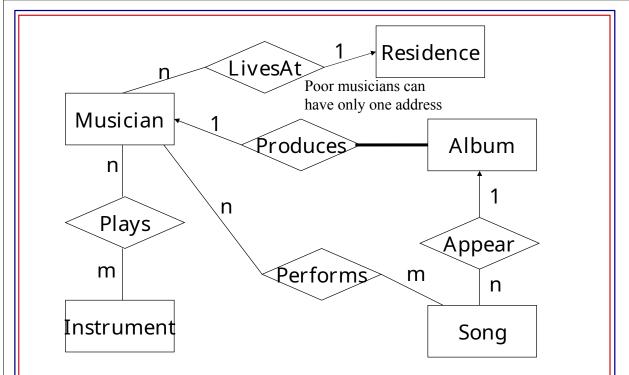
Each song is performed by one or more musicians, and

a musician may perform a number of songs.

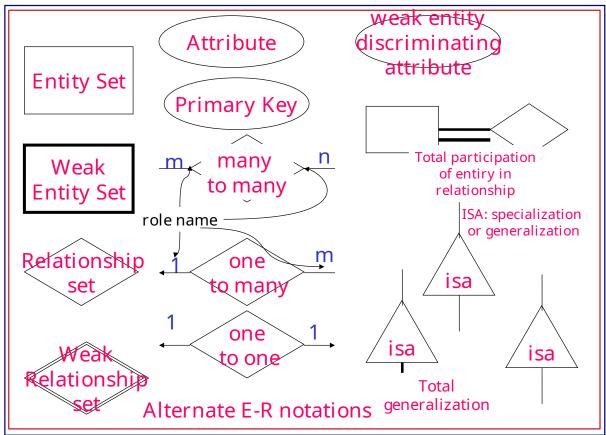
For each album, there is exactly one musician that acts as its producer.

A musician may produce several albums.

d Bipin C 26

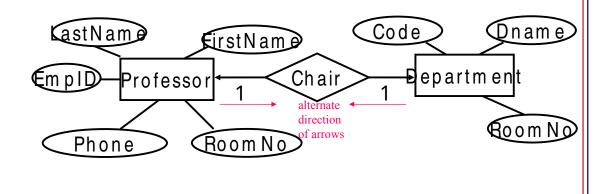


Notes: Since a songs must appear on only one album, Appear is a many to one relationship. Similarly for Produces. Album requires total participation in Produces. Some songs may not be recorded and there may be some instruments that nobody can play!



d Bipin C

- A one-one relationship between **Department** and **its** chair ( a dept. has one chair and a prof. is the chair of at most one dept) is represented by
  - arrows pointing to both Department and Professor or
  - indicated by a line with the number 1 on it.
  - ◆ Sometimes the arrow is in the opposite direction(pointing to the diamond)



Bipin C Desai

20

**ODL** allows only **binary** relationships, i.e., relationships involving two classes.

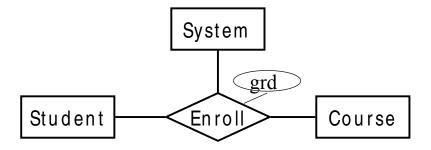
E/R model makes it convenient to define (n-ary) relationships – relationships involving n entity sets

A **n-ary** relationship in an **E/R** diagram is represented by lines from the relationship (diamond) to each of the participating entity sets (rectangles).

Bipin C Desai

30

# N-ary Relationships



Multiplicity of this ternary relationship: n to n to n

Enroll(sid(fk  $\rightarrow$ Student), cno(fk  $\rightarrow$  Course),, sys(fk  $\rightarrow$  System), grd)

What is the problem here?

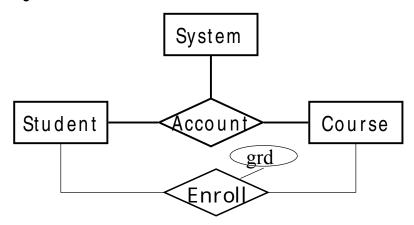
What is the schema for Enroll?

How is the n:n:n relationship mapped to a relation(table)?

Bipin C Desai

31

# N-ary Relationships

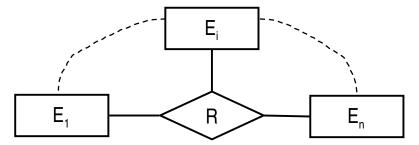


Multiplicity of this ternary relationship: n to n to n

 $\begin{array}{l} \text{Account}(\ \underline{\textbf{sid}}(\text{fk} \to \text{Student}), \ \underline{\textbf{cno}}(\text{fk} \to \text{Course}), \ \underline{\textbf{sys}}(\text{fk} \to \text{System}), \ \underline{\textbf{uid}}) \\ \text{Enroll}(\underline{\textbf{sid}}(\text{fk} \to \text{Student}), \ \underline{\textbf{cno}}(\text{fk} \to \text{Course}), \ \text{grd}) \end{array}$ 

Bipin C Desai

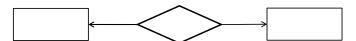
# N-ary Relationships



Multiplicity of this N-ary relationship:  $n_1$ ----  $n_i$  ----  $n_n$  Any of this  $n_i$  could be 1 Any of this could be multiple

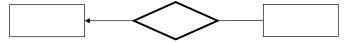
Bipin C Desai

3



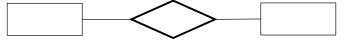
If a binary relationship between two entity sets is 1-to-1,

- the primary key of the relationship is the key of the entity from either of the '1' side, the other side is a foreign key (would be unique)



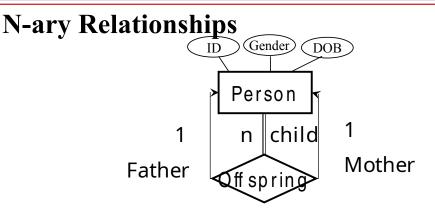
If a binary relationship between two entity sets is 1-to-many,

- the primary key of the relationship is from the 'm' side, the '1' side is theforeign key (would be unique)! This 'convention' may be reversed for convenience -specially if the number of entities on the one side is much smaller!



If a binary relationship between two entity sets is m-to-n,

- the primary key is composite, consisting of the primary key of the entities from each side of the relationship



Multiplicity of this ternary relationship: 1 to 1 to n

Offspring(fid (fatherID  $\rightarrow$  Person(ID)), motherID (fk  $\rightarrow$  Person(ID)) <u>ChildID</u> (fk  $\rightarrow$  Person(ID)))

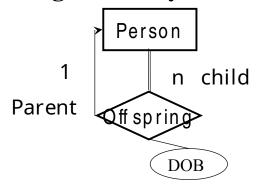
Who here is the father and mother? What is the key?

Bipin C

35

```
create table person(
                                                   create table offspring(
ID number primary key,
                                                   fid number, mid number,
gender char(1),
                                                   cid number primary key,
DOB date);
                                                   foreign key (fid)
                                                           references person(id),
insert into person values(1, 'M', '11-Jan-1900');
insert into person values(2,'F', '11-Jan-1902');
                                                   foreign key (mid)
insert into person values(121,'M', '11-Jan-1925');
                                                            references person(id),
insert into person values(122,'F', '11-Jan-1927'); foreign key (cid)
insert into person values(3,'M', '11-Jan-1901');
                                                            references person(id));
insert into person values(4,'F', '11-Jan-1903');
                                                   insert into offspring values(1,2,121)
insert into person values(341,'M', '11-Jan-1926'); insert into offspring
insert into person values(342,'F', '11-Jan-1928');
                                                            values(1,2,122);
insert into person
                                                   insert into offspring
         values(1213421,'M', '11-Jan-1948');
                                                   values(3,4,341);
insert into person
                                                   insert into offspring
         values(1213422,'F', '11-Jan-1950');
                                                            values(3,4,342);
                                                   insert into offspring
                                                   values(121,342, 1213421);
   Could two tuples exist in offspring with
                                                  insert into offspring
   the same cid???
                                                   values(121, 342, 1213422);
```

### Replacing a ternary relation by a binary relation



What is the schema for Offspring here?

What is a possible inconsistency problem?

Who is the father, mother??

Offspring (<u>IDC</u>, IDF, IDM, DOB)

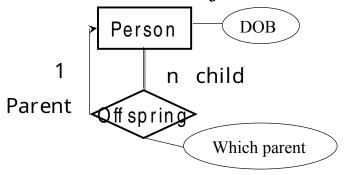
Father(<u>IDC</u>, <u>IDF</u>, DOB) Mother(<u>IDC</u>, <u>IDM</u>, DOB)

- not the same ER
- duplication of DOB
- Composite key

Bipin C

37

# Replacing a ternary relation by a binary relation --- an alternate *non-normal form*



What is the schema for Offspring? Is there a duplication problem? What is the primary key?

TVI UITTI V UITUCU ATTITUUTC			
121	1	Father	
	2	Mother	
122	1	Father	
	2	Mother	

Multivalued attribute

# Roles in Relationships

- It is possible that the same entity set appears two or more times in a relationship
- Suppose, we want to capture the relationship between two courses, one of which is the pre-requisite/follow-on of the other

© Bipin C Desai

Each line to the entity set represents a different role that the entity set plays in the relationship Follow-on

CourseNumber name

Credits

Person

Date

(d) Bipin C Desai

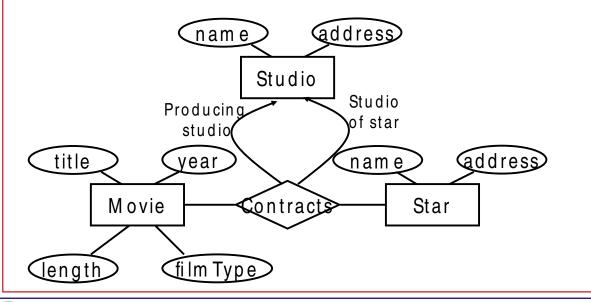
```
create table prereq
(COURSE_Number CHAR(8),
PREREQ CHAR(8),
primary key (course_number, prereq))
tablespace TUTOR pctfree 2
/
```

Bipin C

```
insert into prereq values('COMP353','COMP352'); insert into prereq values('COMP353','COMP346'); insert into prereq values('COMP352','COMP248');
```

The Bipin C 42

Suppose, each star is under contract with a single studio
The studio of the star may enter into a contract with another
studio to allow that star to act in a particular movie



d Bipin C Desai

## Converting n-ary relationship

Any n-ary relationship may be converted into a **collection** of **binary** relationships **without loosing** any information???

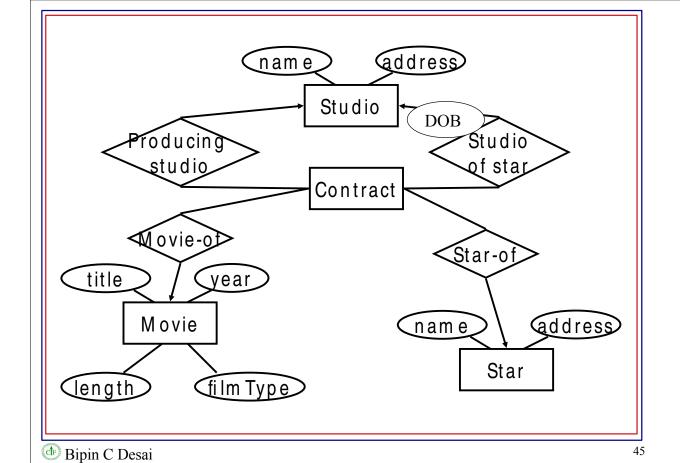
Introduce a **new** entity set – **connecting** *existing* entity set – whose entities might be thought of as tuples of the relationship for the n-ary relationship

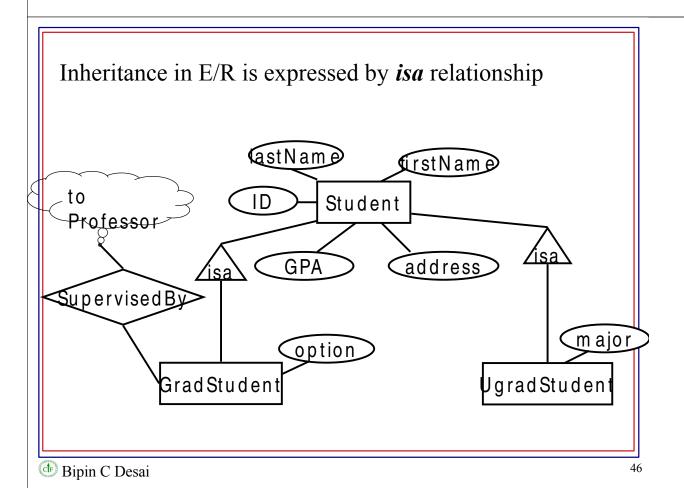
Introduce **many-to-one** relationships from the **connecting** entity set to each of the entity sets participating in the original n-ary relationship

If an entity set plays **more than one** role, then it is the target of one relationship for **each role** 

Usually doesn't convey the same semantics – limitation of modeling

ⓑ Bipin C Desai





- There is a subtle difference between the concept of inheritance in ODL and in the E/R model
- ❖ In **ODL**, an object must be a member of exactly one class
- ❖ In the **E/R** model
  - ♦ We shall view an entity as having "components" belonging to several entity sets that are "part of" a single **isa**-hierarchy
  - ◆ The "components" are connected into a single entity by the **isa** relationships
  - ♦ The entity has whatever attributes any of its components has, and it participates in whatever relationships its components participate in
  - P We need to represent an entity (e.g., CartoonMurderMystery) in the diagram only if it has attributes and/or relationships of its own

Bipin C Desai

4'

# **Constraints**

- There are some important **aspects** of the **real world** that **cannot** be represented using the **ODL** or **E/R** model introduced so far
- The additional information about these aspects often takes the form of **constraints** on the data
- Sometimes modeling this additional information goes beyond the **structural** and **type** constraints imposed by **classes**, **entity sets**, **attributes**, and **relationships**

d Bipin C Desai

**Keys** are (sets of) attributes that **uniquely** identify an object within its class or an entity within its entity set;

#### $K \subseteq R$ . no two entities may agree in all their key values

**Single-value constraints** are requirements that the value of an attribute be unique. In addition to key constraints, other attributes must a have a single-value constraints.

Also in an "one" relationship

**Referential integrity constraints** are requirements that a value referred to by some object must actually exist in the database; *This means, no dangling pointers.* 

**Domain constrains** require that the value of an attribute must be drawn from a specific set of values (called attribute domain), or lies within a specific range

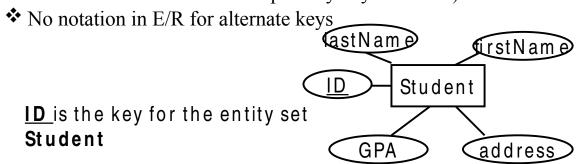
**General constraints** – arbitrary assertions that must hold on the DB

Bipin C Desai

49

# Keys

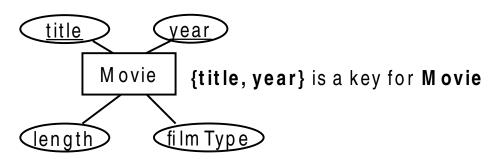
- A *super key* is a set of attributes whose values uniquely identify an entity in the entity set; this set may not be minimal.
- \* A minimal super key is called a (candidate) key.
- An entity may have more than one **key**. One of them is picked as the primary key; others may be called *alternate keys*
- ❖ In **E/R**, we <u>underline</u> the key attribute(s) of an entity set (i.e., those attributes that form the primary key of the set)



d Bipin C Desai

# Example

- \*What should we select as a key for **Movie**?
- **\*** title?
  - ♦ there could be different movies with the same name
- **\*** {title, year}?
  - ♦ there still could be two movies made in the same year, with the same title, but that's very unlikely

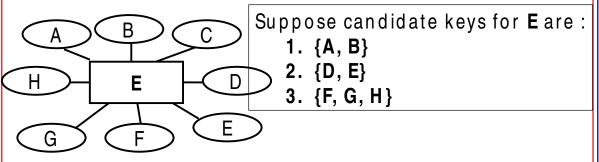


https://en.wikipedia.org/wiki/Lists\_of\_film\_remakes

Bipin C Desai

51





\* Which of the three should "we" pick as the **primary key**?

Criteria to choose a **primary key** when there are more than one candidate:

Total size

Number of attributes

Convenience

A combination of the above

Bipin C Desai 52

### Single Value Constraint

- **❖** In **E**/**R**:
  - attributes are **atomic**
  - lack an arrow  $(\rightarrow)$  can be used to express the multiplicity
  - ◆ What about multi-valued attributes or relationships?

With the E/R model introduced so far, we cannot express the following options regarding the value of a single-valued **attribute**:

- Require that the value of that attribute to be present(not null)
- Or the presence of the value be optional (null allowed)

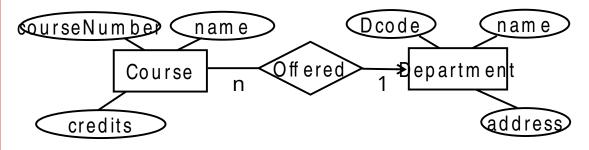
If the choice is not explicit, then we may conclude:

- The value must exist if an attribute is part of the key
- The value is optional, otherwise

h Bipin C Desai

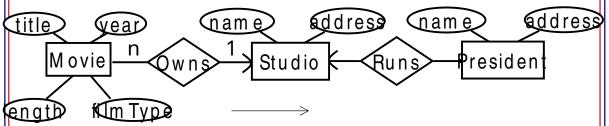
### **Referential Integrity Constraints**

- For relationships:
  - **♦** Single-value + Existence = Referential Integrity Constraint
- \*We extend the **arrow** notation to indicate a reference is mandatory (to support referential integrity)
  - The department that gives a course must always exist in the **Department** entity set



d Bipin C Desai

# **Referential Integrity Constraints**



open arrow to denote ref. intg.

- The studio owning a movie must always be present in the Studios (extent of the Studio entity set)
- \* If a president runs a studio, then that studio exists in the Studios
- \* There could be studios without a president (temporarily)

Bipin C Desai

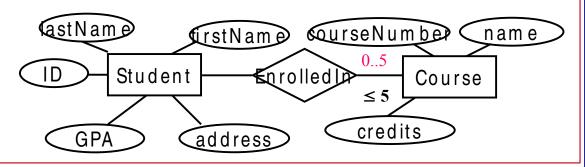
### **Domain constraints**

- \*Domain constraints restrict the values of an attribute to be drawn from a set
  - ◆ In ODL, we give a type to the attributes and hence limit their set of values
  - ◆ ODL does **not** support other restrictions, such as that the value should be within a certain range
  - ◆ E/R, in general, does **not** support imposing domain constraints

h Bipin C Desai

# Relationship degree constraints

- \* Relationship degree constraints restrict the number that an entity/object can participate in a relationship
  - ◆ For example, we can impose a constraint saying that a student cannot be enrolled in more that 5 courses
  - ◆ In ODL, we could use, instead of a set of references, an **array** of size 5
  - ◆ In E/R, we may attach a bounding number to the corresponding link

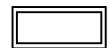


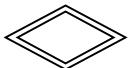
Bipin C Desai

57

# Weak Entity / Relationship Sets

- ❖ A **strong** entity set has a primary key
- \*A weak entity set does not have sufficient attributes to form a primary key. It should be part of a onemany relationship (with no descriptive attributes) with a strong entity set
- \*Discriminator of a weak entity set is a set of attributes that distinguishes among the entities corresponding to a strong entity
- \*Primary key of a weak entity set = primary key of the strong entity + discriminator of the weak entity
- \*Represented in E/R model by

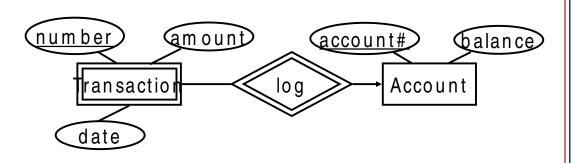




Bipin C Desai

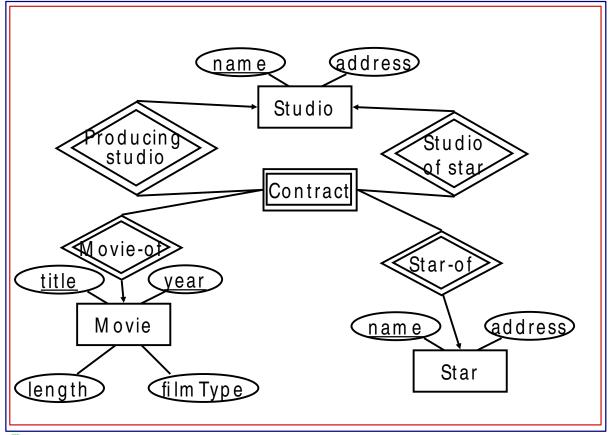
### **Example**

- Log records transactions done by an ATM
- \* Each transaction has a number, date, and an amount
- ❖ Different accounts might have transactions by the same number, on the same date, and for the same amount



Bipin C Desai

59



Bipin C Desai

### **Design Principles**

- Design should
  - ♦ Reflect reality
  - ♦ Avoid redundancy
    - □ Redundant information takes space
    - □ Could cause inconsistency
  - Be as simple as possible
- \*Be careful when choosing between using attributes and using classes or entity sets. Remember that
  - ◆ An attribute is **simpler** to implement than either a class/entity set or a relationship
  - ♦ If something has more information associated with it than just its name, it probably needs to be an entity set or a class

Bipin C Desai

61

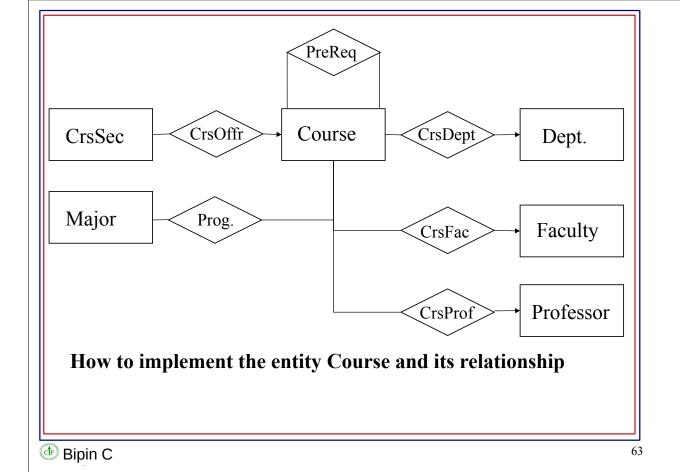
Consider the entity set course in a typical university:

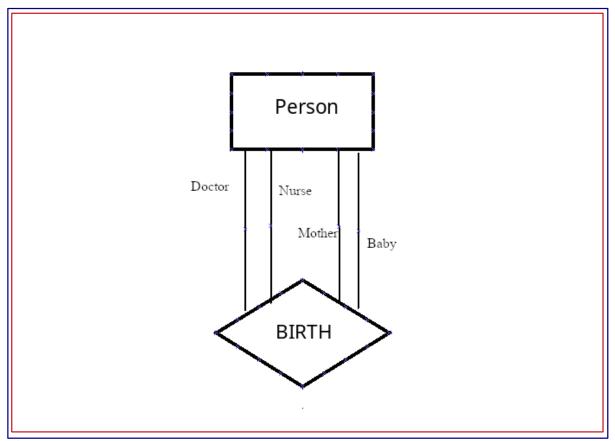
It could be involved in many relationships:
one to many relationships with
the offering department,
the offerings faculty
the professor coordinating the course

A many to one relationship with sections for the course

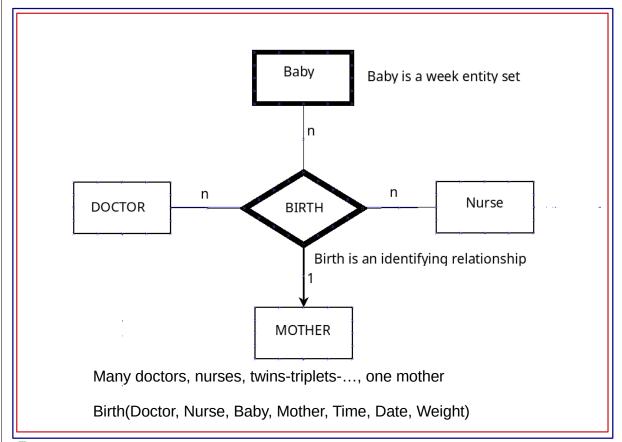
many to many relationships with the major program in which it is required the pre-requisites (follow up) courses

© Bipin C 62



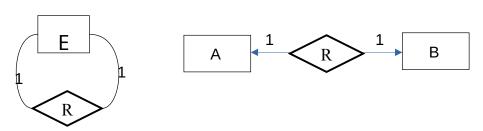


de Bipin C 64



Bipin C

Many doctors, nurses, twins-triplets-..., one mother Birth(Mother, Baby, Time, Date, Weight, Doctor, Nurse) D1 N1 T1 M1 В1 D1 W1 D2 N2 N3 D1 N1 М1 B2 T2 W2 D1 N2 D2 N3 N1 D1 M1 В3 T3 W3 N2 D1 D2 N3 D3



Create table R\_AB(A char(2) primary key, B char(2) unique, C integer (5));

insert into R\_AB values ('A1','B1',11);

Query OK, 1 row affected (0.010 sec)

insert into R\_AB values ('A2','B1',11);

ERROR 1062 (23000): Duplicate entry 'B1' for key 'B'

insert into R\_AB values ('A1','B2',11);

ERROR 1062 (23000): Duplicate entry 'A1' for key 'PRIMARY'

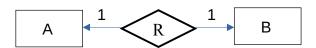
insert into R\_AB values ('A2','B2',22); Query OK, 1 row affected (0.003 sec)

select \* from R\_AB;

+		-+-		+-		+	
	A		В		С		
+		-+-		+-		+	
	A1		В1		11		
	A2		В2		22		
+		-+-		+-		+	
2	rov	vs.	in	set	(0.00	0 (	sec

**Bipin C** 

67

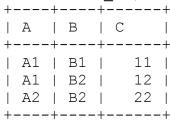


Create table S\_AB(A char(2), B char(2), primary key(A,B),, C integer (5)); insert into S\_AB values ('A1','B1',11),('A1','B2',12), ('A2','B2',22);

Query OK, 3 rows affected (0.009 sec)

Records: 3 Duplicates: 0 Warnings: 0

#### select \* from S AB;



3 rows in set (0.000 sec)

© Bipin C 68

# **Design decision**

Merge the one-to-many relationships
CrsDept, CrsFac, CrsProf
in the schema for Course; all attributes of
these relationships are also included in the
schema for Course

Create a separate relation for each one-to-many relationships

In this case the relation for Course would have a higher arity; but requires one less join to get details for the department, faculty or professor for a given course

Similarly, merge the one-to-many relationship CrsSec in the schema for CrsSec

Create a relation for the many to many relationships Program and PreReq

### **Relational Database**

### Relational Algebra – SQL

#### Notes



Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

Bipin C Desai

#### **Attributes and Domains**

An object or entity is characterized by its properties (attributes or data elements). The set of allowable values for an attribute is the domain of the attribute.

**Domain**. We define a domain, D<sub>i</sub>, as a set of values of the same data type.

Each Attribute is defined on some underlying domain; more than one attribute may share a domain.

#### **Tuples, Relations and Their Schemes**

A relation consists of a homogeneous set of tuples.

Since each tuple in a relation represents an identifiable instance of an entity (object type), duplicate tuples are not allowed.

The number of attributes in the relation gives the **degree** or **arity** of the relation.

The **cardinality** of an instance of a relation, at a point in time, is derived from the count of the tuples in the instance. The cardinality could change over time Relation Representation

#### APPLICANT:

Name	Age	Profession
John Doe	55	Analyst
Mirian Taylor	31	Programmer
Abe Malcolm	28	Receptionist
Adrian Cook	33	Programmer
Liz Smith	33	Manager

Bipin C Desai

1

**Key**. A subset of attributes X of a relation  $R(\mathbf{R})$ ,  $X \in \mathbf{R}$ , with the following time independent properties is called the **key** of the relation:

**Unique Identification**: The values of X uniquely identify a tuple.  $s[X] = t[X] \Rightarrow s = t$ .

**Non-redundancy**: No proper subset of X has the unique identification property.

There may be more than one key in a relation; all such keys are known as **candidate keys**.

One of the candidate keys is chosen as the **primary key**; the others are known as alternate keys.

An attribute that forms part of a candidate key of a relation is called a **prime attribute**.

EMPLOYEE (Emp#, Emp\_Name, Profession)
PRODUCT (Prod#, Prod\_Name, Prod\_Details)
JOB\_FUNCTION (Job#, Title)
ASSIGNMENT (Emp#, Prod#, Job#)

EMPLOYEE		PRODUC	. [		
Emp#	Name	Profession	Prod#	Prod_Name	Prod_Details
101	Jones	A nalyst	HEAP1	HEAP_SORT	ISS module
103	Smith	Programmer	BIN S9	BINARY_SEARCH	ISS/ R module
104	Lalonde	Receptionist	FM6	FILE_MANAGER	ISS/ R-PC subsys
106	Letitia	VP Marketing	B++1	B++_TREE	ISS/ R turbo sys
107	Evan	VP R & D	B++2	B++_TREE	ISS/ R-PC turbo
110	Drew	VP Operation			
112	Smith	Manager	_		

Bipin C Desai

\_\_\_\_

#### JOB\_FUNCTION

Job#	Title
1000	CEO
700	Chief Programmer
800	Manager
600	A nalyst

#### **ASSIGNMENT**

Emp#	Prod#	Job#
107	HEAP1	800
101	HEAP1	600
110	BINS9	800
103	HEAP1	700
101	BINS9	700
110	FM6	800
107	B++1	800

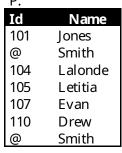
The attributes *Emp*#, *Prod*#, and *Job*# in the relation ASSIGNMENT are known as **foreign** keys.

The Bipin C Desai

#### A null value for an attribute:

- a value that is either not known at the time, or P:
- the value is known but not recorded, or
- no value is applicable for some tuples

<u>P:</u>	
Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith



Emp#	Name	Manager
101	Jones	@
103	Smith	110
104	Lalonde	107
107	Evan	110
110	Drew	112
112	Smith	112

Bipin C Desai

-

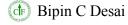
**Integrity rule 1 (Entity Integrity).** If attribute A of relation  $R(\mathbf{R})$  is a component of the primary key of  $R(\mathbf{R})$ , then A cannot accept null values.

**Integrity Rule 2** (**Referential Integrity**). Given two relations R and S. Suppose R refers the relation S via a set of attribute which forms the primary key of S and, hence, this set of attributes forms a foreign key in R. Then, the value of the foreign key in a tuple in R must either be equal to the primary key of a tuple of S or be entirely null.

- All tuples which contain references to the deleted tuple should also be deleted. **cascading** deletion
- A tuple which is referred by other tuples in the database cannot be deleted.
- In the third option, the tuple is deleted, however, the foreign key attributes of all referencing tuples are set to null(otherwise "dangling" pointers!)

d Bipin C Desai

- All tuples which contain references to the deleted tuple should also be deleted. This is **cascading** deletion
- A tuple which is referred by other tuples in the database cannot be deleted.
- In the third option, the tuple is deleted, however, the foreign key attributes of all referencing tuples are set to null (otherwise "dangling" pointers!)



C

### **Query Languages**

- \* The Relational model supports simple, powerful query languages which:
  - ◆ Have formal foundation based on logic.
  - ◆ Allows for implementation which can be optimized.
- \* Allow data access and modification.
- \* These languages are not general purpose programming languages, however, most DBMS vendors have added their own enhancements to improve its functionality.

Bipin C Desai

# Relational Algebra, Calculus

Relational Algebra(RA) and Relational Calculus(RC) are the foundation for implemented languages (e.g. SQL)

- \* RA is operational, and is useful for representing the plan of execution of a query.
- \* RC is declarative allowing users to describe what they want. (i.e. it is non-operational)
  - Understanding RA & RC is vital to the understanding of SQL and query processing!

Your textbook may not cover RC!

Bipin C Desai

Relational algebra is a collection of operations to manipulate relations.

**Basic Operations**: Three of these four basic operations - union, intersection and difference - require that operand relations be union-compatible. (Same number (and order) of attributes on identical (at least compatible) domains

Q:

Id	Name
101	Jones
105	Letitia
107	Evan
110	Drew

Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

Bipin C Desai

UNION ( $\cup$ ) If we assume that P(**P**) and Q(**Q**) are two union-compatible relations, then:

The union of  $P(\mathbf{P})$  and  $Q(\mathbf{Q})$  is the set-theoretic union of  $P(\mathbf{P})$  and  $Q(\mathbf{Q})$ .

The resultant relation,  $R = P \cup Q$ , has tuples drawn from P and Q, such that

$$R = \{t \mid t \in P \lor t \in Q\}$$
 and

$$max(P,Q) \le R \le P + Q$$

Union operation is associative and commutative

$$P \cup Q \cup S = P \cup (Q \cup S) = (P \cup Q) \cup S = (P \cup S) \cup Q$$

Bipin C Desai

13

### Q:

Id	Name
101	Jones
105	Letitia
107	Evan
110	Drew

### $P \cup Q$ :

Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

### **P**:

Ia	iv ame
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

#### $Q \cup P$ :

٦ -	•
Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

Bipin C Desai

14

## DIFFERENCE (-)

The difference operation removes common tuples from the first relation.

$$R = P - Q$$
 such that

$$R = \{ t \mid t \in P \& t \notin Q \} \text{ and } 0 \le R \le P$$

Difference operation is non-associative and non-commutative.

Is 
$$P - Q = Q - P$$
?  
Is  $P - (Q - S) = (P - Q) - S$ ?

Bipin C Desai

1.4

Q	:
	н

110

Id	Name
101	Jones
105	Letitia
107	Evan

Drew

P:

Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

P - Q:

Id	Name
103	Smith
104	Lalonde
112	Smith

Q - P:

Id Name

### INTERSECTION ( $\cap$ )

The intersection operation selects the common tuples from the two relations.

$$R = P \cap Q$$
 where

$$R = \{t \mid t \in P \& t \in Q\} \text{ and } 0 \le R \le \min(P, Q)$$

The intersection operation is really unnecessary as it can be very simply expressed as:

$$P \cap Q = P - (P - Q)$$

$$Q \cap P = Q - (Q - P)$$

Is 
$$P - (P - Q) = Q - (Q - P)$$
?

Bipin C Desai

17

P	٠(
1	•

Id	Name
103	Smith

104 Lalonde

105 Letitia

107 Evan

110 Drew

112 Smith

Q

Id	Name
101	Jones
105	Letitia

105 Ledda 107 Evan

107 Evan

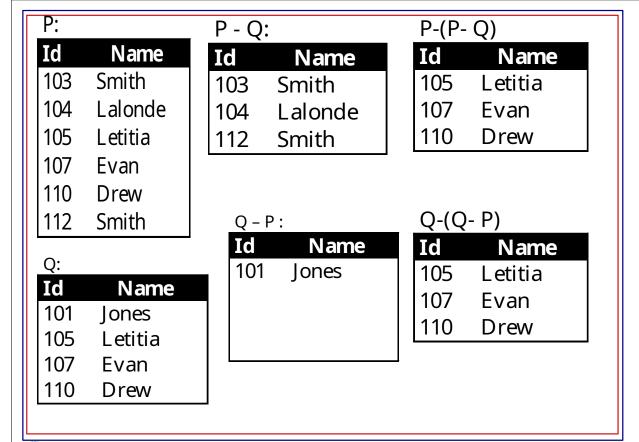
110 Drew

 $P \cap Q$ :

Id	Name
105	Letitia
107	Evan
110	Drew

 $Q \cap P$ :

Id	Name
105	Letitia
107	Evan
110	Drew



Bipin C Desai

19

### RENAMING ( $\rho$ )

The renaming operation  $\rho^*$  is used to rename relations or its attributes. The operation:

$$\rho(R(modattributes), rel_exp)$$

takes a relation expression and the result is named R with some of the attributes, specified in the modattributes, are renamed

The format of modattributes is:

 $modattributes := <oldname \rightarrow newname>|$ 

<position → newname> <,modattributes>

\* $\rho$  rho is the 17<sup>th</sup> letter of the Greek alphabet

de Bipin C Desai

## Employee(Emp#, Ename, Address, Phone, DOB)

 $\rho(Q_{\text{ Emp\#} \rightarrow \text{ ID, Ename} \rightarrow \text{Name}} \Pi_{\text{ Emp\#,Ename}} \text{ EMPLOYEE)}$ 

Q:

<u> </u>	
Id	Name
101	Jones
105	Letitia
107	Evan
110	Drew

Bipin C Desai

\_

### **CARTESIAN PRODUCT (×)**

The extended cartesian or simply the cartesian product of two relations is the concatenation of tuples belonging to the two relations.  $R = P \times Q$ 

The scheme of the result relation is given by:  $\mathbf{R} = \mathbf{P} | |\mathbf{Q}|$ .

The degree of the result relation is given by:

$$|\mathbf{R}| = |\mathbf{P}| + |\mathbf{Q}|.$$

Q:	
Id	Name
101	Jones
105	Letitia
107	Evan
110	Drew
D.	

I	D	•
		•

Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

Id	Name	Id	Name
101	Jones	101	Jones
101	Jones	103	Smith
101	Jones	104	Lalonde
110	Drew	112	Smith

Bipin C Desai

# **PROJECTION** $(\prod)$ $\prod_X R$

It should be noted that the projection operation reduces the arity if the number of attributes in X is less than the arity of the relation. It may also reduce the cardinality of the result relation since duplicate tuples removed.

P:

Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

 $\Pi_{\text{Name}}P$ :

Name
Jones
Smith
Lalonde
Letitia
Evan
Drew

Bipin C Desai

### SELECTION (σ)

The selection operation, yields a "horizontal subset" of a given relation. Any finite number of predicates connected by boolean operators may be specified in the selection operation.

### P:

<u> </u>	
Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia
107	Evan
110	Drew
112	Smith

### **O** Id<106 P:

Id	Name
101	Jones
103	Smith
104	Lalonde
105	Letitia

Bipin C Desai

25

## JOIN (0 0 )

The join operator, as the name suggests, allows the combining of two relations to form a single new relation.

- first compute the cartesian product
- followed by selecting those tuples where the common attribute(s) has(have) the same value(s).

Bipin C Desai

26

Project (Proj#, Pname, Pleader) Assign(P#, E#)

Employee(Emp#, Ename, Address, Phone, DOB)

- **❖** Get Emp# of employees working on Proj# comp353.
- **Get complete details of employee working on comp353.**
- **❖** Get complete detatils of employes working on the Database project.
- ❖ Get complete details of employees working on both comp353 and comp354
- Get Emp# (complete details) of employees working on two projects.
- Get names of employees working in projects where Ma is the project leader.

Bipin C Desai

27

Project (Proj#, Pname, Pleader) Assign(P#, E#)

Employee(Emp#, Ename, Address, Phone, DOB,)

**❖**Get Emp# of employees working on Proj# comp353.

$$\Pi_{E\#} (\sigma_{P\#=comp353}(Assign))$$

**❖**Get complete details of employee working on comp353.

Employee 
$$\triangleright \triangleleft_{\text{Emp\#=E\#}} \Pi_{\text{E\#}} (\sigma_{\text{P\#=comp353}}(\text{Assign}))$$

Project (Proj#, Pname, Pleader) Assign(P#, E#)
Employee(Emp#, Ename, Address, Phone, DOB,)

**❖**Get complete details of employees working on the Database project(s) - a project name.

$$X = \Pi_{E\#} ( (Assign) \rhd \lhd_{Proj\#=P\#} (\sigma_{Pname="Database"}(Project)))$$
  
 $Employee \rhd \lhd_{Emp\#=E\#} X$ 

Combining in one RA expression:

Employee 
$$\triangleright \triangleleft_{\text{Emp\#=E\#}} \Pi_{\text{E\#}}$$
 ( (Assign)  $\triangleright \triangleleft_{\text{Proj\#=P\#}}$  ( $\sigma_{\text{Pname="Database"}}$ (Project)))

Bipin C Desai

29

Project (Proj#, Pname, Pleader) Assign(P#, E#)
Employee(Emp#, Ename, Address, Phone, DOB,)

**❖**Get complete details of employees working on both comp353 and comp354

$$\begin{split} Employee \; \rhd \lhd_{\; Emp\#=E\#} & \Pi_{E\#} \left( \sigma_{P\#=comp353}(Assign) \right) \cap \\ Employee \; \rhd \lhd_{\; Emp\#=E\#} & \Pi_{E\#} \left( \sigma_{P\#=comp354}(Assign) \right) \\ or & Employee \; \rhd \lhd_{\; Emp\#=E\#} \left( \Pi_{E\#} \left( \sigma_{P\#=comp353}(Assign) \right) \cap \\ & \Pi_{E\#} \left( \sigma_{P\#=comp354}(Assign) \right) \end{split}$$

Project (Proj#, Pname, Pleader) Assign(P#, E#)

Employee(Emp#, Ename, Address, Phone, DOB,)

**❖** Get complete details of employees working on (any)two projects

$$\begin{split} &\rho(A1(P1\#,E1\#),Assign)\\ &X=A1{\times}Assign \qquad Y=\Pi_{E\#}\left(\right.\sigma_{P\#{\neq}P1\#\,\wedge\,E\#=E1\#}(X))\\ &Employee\ {\triangleright}{\triangleleft}_{Emp\#=E\#}Y \end{split}$$

### **Combining in one RA expression:**

```
Employee \rhd \lhd_{Emp\#=E\#} (\Pi_{E\#} (\sigma_{P\#\neq P1\# \land E\#=E1\#} (\rho(A1(P1\#,E1\#),Assign) \times Assign)))
```

Bipin C Desai

3

Project (Proj#, Pname, Pleader) Assign(P#, E#)

Employee(Emp#, Ename, Address, Phone, DOB,)

**❖**Get complete details of employees working in projects where Ma(an employee name) is the project leader.

```
X= (\sigma_{Ename="Ma"}(Employee)))
Y=\Pi_{Proj\#} (Project \bowtie_{Emp\#=Pleader} X)
Z=\Pi_{E\#} (Assign \bowtie_{Proj\#=P\#} Y)
Employee \bowtie_{Emp\#=E\#} Z
```

Employee 
$$\triangleright \triangleleft_{\text{Emp\#=E\#}} (\Pi_{\text{E\#}} (\text{Assign} \triangleright \triangleleft_{\text{Proj\#=P\#}} (\Pi_{\text{Proj\#}} (\text{Project} \triangleright \triangleleft_{\text{Emp\#=Pleader}} (\sigma_{\text{Ename="Ma"}} (\text{Employee})))))))$$

Bipin C Desai

32

## Complete set of RA operations

$$\{\sigma, \Pi, \cup, -, \times\}$$

The above is a complete set of RA operations.

The others could be expressed as a sequence of operations from this set.

$$R \cap S \equiv (R \cup S) - ((R - S) \cup (S-R))$$
  
 $R \bowtie_B S \equiv \sigma_B (R \times S) \text{ etc}$ 

(fr) Bipin C Desai

33

*Definition*: **Theta-join**. The **theta-join** of two relations  $P(\mathbf{P})$  and  $Q(\mathbf{Q})$  is defined as

$$R = P \stackrel{\bowtie}{}_B Q$$
 such that 
$$R = \{t \mid t_1 | | t_2 \wedge t_1 \in P \wedge t_2 \in Q \wedge B\}$$

where B is a selection predicate consisting of terms of the form:  $(t_1[A_i] \theta t_2[B_i])$  for i = 1, 2, ..., n,

where  $\theta_i$  is some comparison operator ( $\theta \in \{=,\neq,<,\leq,>,\geq\}$ ), and  $A_i$  and  $B_i$  are some domain compatible attributes of the relation schemes **P** and **Q** respectively.

$$0 \le |R| \le |P| * |Q|$$
$$|R| = |P| + |Q|$$

Two common and very useful variants of the join are the **equi- join** and the **natural-join**.

If two relations that are to be joined have no domain compatible attributes, then the natural join operation is equivalent to a simple cartesian product.

- -The equi-join and the theta joins are "horizontal subsets" of the cartesian product.
- -The natural join is equivalent to an equi-join with a subsequent projection to eliminate the duplicate attributes.

Bipin C Desai

3.5

SQL is both the data definition and data manipulation language of the relational database systems

**Create table** <relation> (<attribute list>, <integrity constraint list>)

Where the <attribute list> is specified as:

<Attribute list> ::= <attribute name> (<data type>)[not
null][,<attribute list>]

and <integrity constraints list> is specified as:

<Integrity constraint list> ::= <integrity1>|<integrity constraint list>

and <integrity> could be a primary key(a1, a2, ... am), not null or a named constraint.

Clar Bipin C Desai

### create table EMPLOYEE

(Empl\_No integer not null, Name char(25), Skill char(20), Pay\_Rate decimal(10,2) Primary key Empl\_No)

select [distinct] < target list>
from < relation list>
[where < predicate>]

Bipin C Desai

37

## Database Schema

- Employee(Name, Sin, Dept#, MGRSIN)
- ❖ Dept(Dname, Dept#, MgrSin, Bcode)
- Project(Pname, Proj#, Dept#, Lab)
- ❖ Assign(Proj#, EmpSin, Hours)
- EmpDet(Sin, Address, Salary, DOB)
- EmplDepd(Sin, DepName, HowR)

<sup>(d)</sup> Bipin C Desai

# Queries

Employee(Name, Sin, Dept#, MGRSIN)
Dept(Dname, Dept#, MgrSin, Bcode)
Project(Pname, Proj#, Dept#, Lab)

- \*Names of Employees in Dept 101? select Name from Employee where Dept#= 101
  - \* Details of Employee in Dept. 101? select \* from Employee where Dept#= 101

Bipin C Desai

39

```
Employee(Name, Sin, Dept#, MGRSIN)
Dept(Dname, Dept#, MgrSin, Bcode)
Project(Pname, Proj#, Dept#, Lab)
```

\* For all projects in the Software Engg. Lab, find the DName, Manager's name? select Dname, Name from Employee e, Dept d, Project p where Lab = 'Software Engg.' and p.Dept#=d.Dept# and d.MgrSin=e.Sin

\* For all projects in the Software Engg. Lab, find the DName, Manager's address etc.

Ch Bipin C Desai 40

Employee(Name, Sin, Dept#, MGRSIN)
Dept(Dname, Dept#, MgrSin, Bcode)
Project(Pname, Proj#, Dept#, Lab)
EmpDet(Sin, Address, Salary, DOB)

\*For all projects in the Software Engg. Lab, find the DName, Manager's name, address etc.

select Dname, Name, Address, Salary, DOB from Employee e, Dept d, Project p, EmpDet t where Lab = 'Software Engg.' and p.Dept#=d.Dept# and d.MgrSin=e.Sin and e.Sin=t.Sin

Bipin C Desai

41

### **Query Tree**

Consider the relation:

Employee(Emp#, Ename, City, Phone, YOB,)

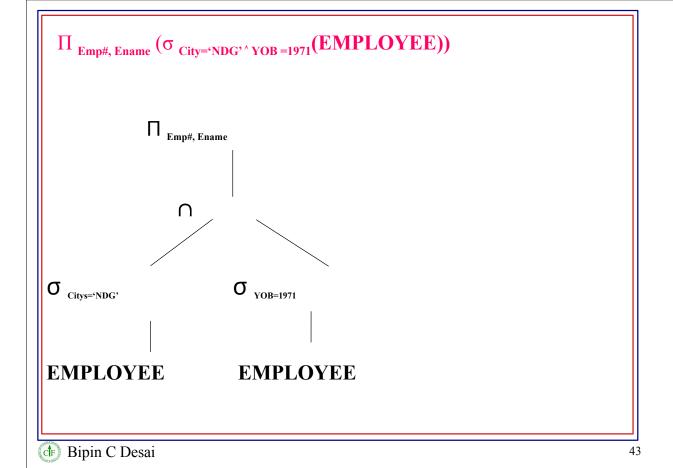
Suppose we want to find Emp# and names of employees who live in NDG(an address) and who were born in 1971(YOB).

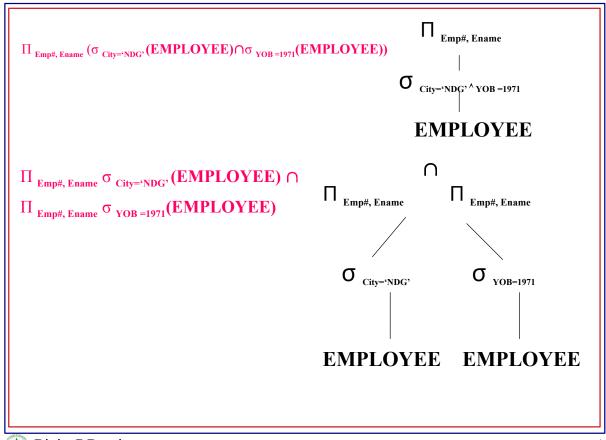
We can express this query in one of the following ways:

```
\begin{array}{l} \prod_{Emp\#, Ename} \left(\sigma_{City=`NDG'^{A}YOB=1971}(EMPLOYEE)\right) \\ \text{or} \\ \prod_{Emp\#, Ename} \left(\sigma_{City=`NDG'}(EMPLOYEE)\cap\sigma_{YOB=1971}(EMPLOYEE)\right) \\ \text{or} \\ \prod_{Emp\#, Ename} \sigma_{City=`NDG'}(EMPLOYEE)\cap \\ \prod_{Emp\#, Ename} \sigma_{YOB=1971}(EMPLOYEE) \end{array}
```

Bipin C Desai

42





Clar Bipin C Desai 44

Project (Proj#, Pname, Pleader) 100 projects
Empl (Emp#, Ename, City, Ph, DOB,) 500 employees
Assign (P#, E#) 1500 assignments

Get complete details of employees working on the DB project(s). Suppose there are 10 DB projects, distribution is uniform.

Av. of 3 projects for each employee; Av. of 15 employees per project Number of assignments to DB project is 150 (10% of assignments).

If no employee works on more than one DB project, then maximum number of tuples in output would be 150.

Two possible ways of expressing this query

$$\Pi_{\text{Emp\#, Ename, City, Ph, DOB}}(\sigma_{\text{E\#=Emp\#}}(\sigma_{\text{P\#=Proj\#}}(\sigma_{\text{Pname="DB"}}(\text{Empl$\times$Assign$\times$Project)$))})$$

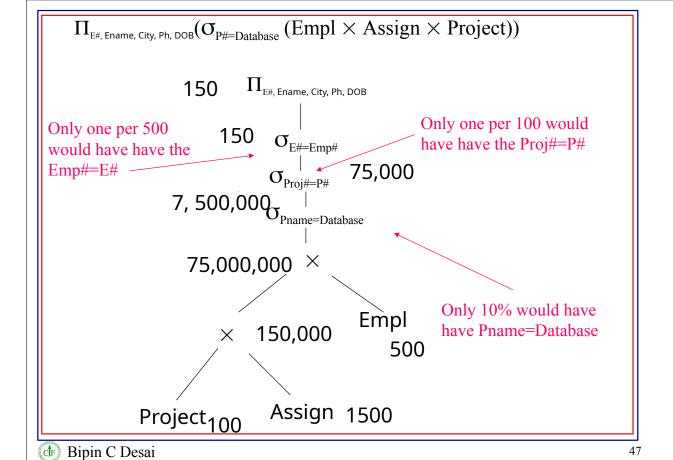
$$Empl \; \rhd \lhd_{\; Emp\#=E\#} \; (\Pi_{E\#} \left( \; (Assign) \; \rhd \lhd_{\; Proj\#=P\#} \; (\sigma_{Pname="DB"}(Project))))$$

45

Pr Pn		PE	Em En	
P1 DB	F	<sup>2</sup> 1 E1	E1 N1	
P2 X	F	<sup>2</sup> 2 E2	E2 N2	
		•	•	
	Pr Pn	ΡE	Em En	
	P1 DB	P1 E1	E1 N1 E2 N2	
		 D2 F2	 E1 N1	
		IZ LZ	E2 N2	
		•••••	•••	
	P2 X	P1 E1	E1 N1 E2 N2	
		•••••	•••	
		P2 E2	E1 N1 E2 N2	
		•••••	•••	

Bipin C Desai

46



Empl  $Arr def = H (T_{E\#} (Assign) 
Arr def = Proj\#=P\# (\sigma_{P\#=Database}(Project))))$ Result 150  $Arr def = H (Assign) 
Arr def = Proj\#=P\# (\sigma_{P\#=Database}(Project))))$   $Arr def = H (Assign) 
Arr def = Proj\#=P\# (\sigma_{P\#=Database}(Project))))$   $Arr def = H (Assign) 
Arr def = Proj\#=P\# (\sigma_{P\#=Database}(Project))))$   $Arr def = H (Assign) 
Arr def = Proj\#=P\# (\sigma_{P\#=Database}(Project))))$  Arr def = H (Assign) 
Arr def = H (Assign)

Of Bipin C Desai 48

# Division (÷)

 $P(\mathbf{P})$ :  $Q(\mathbf{Q})$ : R(R)(result):

A B B A

 $a_1 b_1 b_1 a_1$ 

 $a_1 b_2 b_2 a_5$ 

 $a_2$   $b_1$  The result of dividing P by Q is the relation R

 $a_3$   $b_1$  which has two tuples. For each tuple in R, its

 $a_4$   $b_2$  product with the tuples of Q must be in P. In

our example  $(a_1,b_1)$  and  $(a_1,b_2)$  must both be

 $a_5$   $b_1$  tuples in P; the same is true for  $(a_5,b_1)$  and

 $a_5 b_2 (a_5, b_2)$ .

The Cartesian product of Q and R is a subset

of P.

Bipin C Desai

49

 $P(\mathbf{P}):$   $Q(\mathbf{Q}):$   $R(\mathbf{R})$  is:  $Q(\mathbf{Q}):$   $R(\mathbf{R})$  is:

A B B A B A

 $a_1 b_1$ 

 $\begin{bmatrix} a_1 & b_2 & & & b_1 \\ & b_1 & & a_1 & & b_1 \end{bmatrix}$ 

 $a_2$   $b_1$   $a_2$   $b_2$   $a_3$   $b_4$   $a_2$ 

 $a_3$   $b_1$   $a_3$   $b_3$   $a_3$   $a_3$   $a_3$   $a_3$ 

 $a_4$   $b_2$   $a_5$  R(R) is:

 $a_5$   $b_1$  B A

The division operation is useful where a query involves the phrase "for all objects having all of the specified properties".

 $a_1$   $a_2$ 

 $a_2$ 

 $a_4$ 

a<sub>5</sub>

Project (Proj#, Pname, Pleader) Assign(P#, E#) Employee(Emp#, Ename, Address, Phone, DOB,)

Get complete details of employees working on all Database projects.

Find the Proj# of all Database project as DBPROJNO  $\rho_{1\rightarrow P\#}$  (DBPROJNO,  $\Pi_{Proj\#}(\sigma_{Pname=Database}Project))$ 

Find the specified details for the required employees by dividing **Assign** by DBPROJNO and join the result with **Employee**.

ASSIGN  $\div$  DBPROJNO  $\bowtie$  Emp#=E# **Employee** 

Bipin C Desai

5

Project (Proj#, Pname, Pleader) Assign(P#, E#) Employee(Emp#, Ename, Address, Phone, DOB,)

Get complete details of employees working exactly on all DB projects.

Find the Proj# of all Database project as DBPROJNO  $\rho_{1\rightarrow P\#}$  (DBPROJNO,  $\Pi_{Proj\#}(\sigma_{Pname=Database}Project)$ )

Find those employees who work on all DB projects by dividing

**Assign** by DBPROJNO (some of them work on other projects as well.).

ALLDB = ASSIGN  $\div$  DBPROJNO

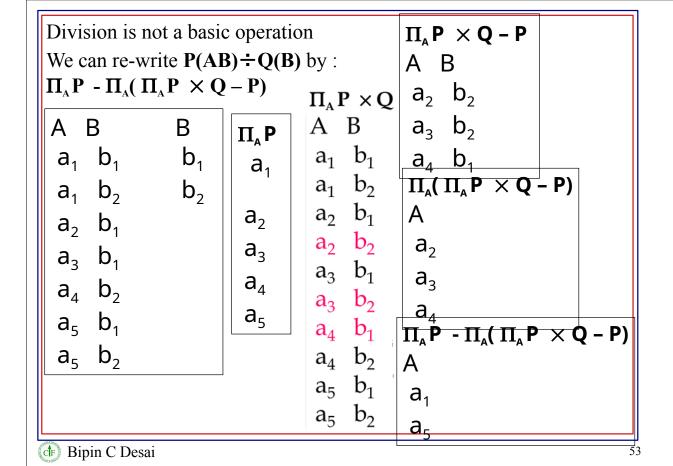
Find those tuples not involving assignments to DB projects

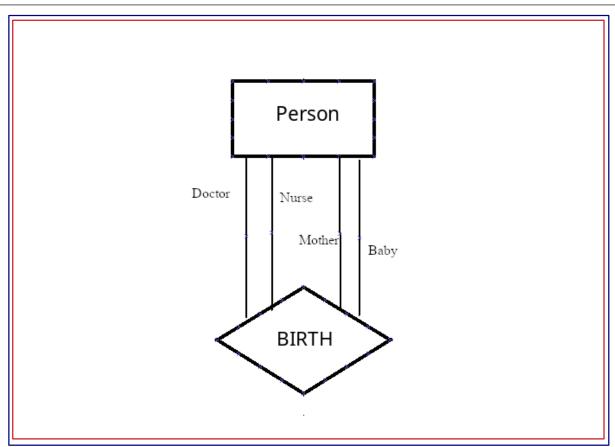
NOTDBONLY = ASSIGN – DBPROJNO  $\times$  ALLDB

Required employees: ONLYDB = ALLDB -  $\Pi_{E\#}$  NOTDBONLY

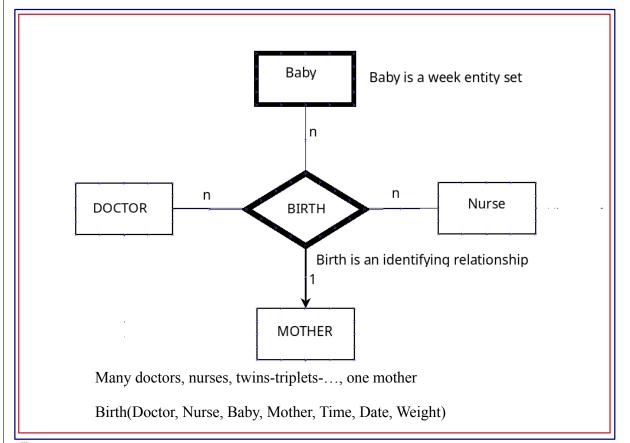
Result is: ONLYDB ▷ Emp#=E# Employee

de Bipin C Desai





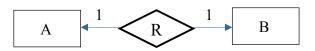
d Bipin C Desai



h Bipin C Desai

Many doctors, nurses, twins-triplets-..., one mother Birth(Mother, Baby, Time, Date, Weight, Doctor, Nurse) D1 N1 M1B1 T1 D1 W1D2N2 N3 D1 N1 M1B2 T2 D1 W2 D2N2 N3 D1 N1 M1 W3 T3 B3 D1 D2N2 D3N3

db Bipin C Desai



Create table R\_AB(A char(2), B char(2), primary key(A,B),, C integer (5));

insert into R\_AB values ('A1','B1',11),('A1','B2',12), ('A2','B2',22);

**NOT 1:1** 

Query OK, 3 rows affected (0.009 sec)

Records: 3 Duplicates: 0 Warnings: 0

#### select \* from R AB;

3 rows in set (0.000 sec)

Bipin C Desai

57



Create table R\_AB(A char(2) primary key, B char(2) unique, C integer (5));

insert into R\_AB values ('A1','B1',11);

Query OK, 1 row affected (0.010 sec)

insert into R\_AB values ('A2','B1',11);

ERROR 1062 (23000): Duplicate entry 'B1' for key 'B'

insert into R\_AB values ('A1','B2',11);

ERROR 1062 (23000): Duplicate entry 'A1' for key 'PRIMARY'

insert into R\_AB values ('A2','B2',22);

Query OK, 1 row affected (0.003 sec)

select \* from R AB;

0	010	-	-	1011			,		
+		-+-		+		+			
	Α		В		С				
+		-+-		+		+			
-	Α1	1	В1	- 1		11			
	A2		В2			22			
+		-+-		+		+			
2	rov	vs.	in	set	(0	000	sec)		

© Bipin C Desai

1:1

## Some special characters used in DB & HTML codes

```
\land ∧ \lor ∨ \exists ∃ \neg \exists ¬ ∃ \forall ∀ \Sigma ∑ \prod ∏ \cup ∪ \cap ∩ \subseteq ⊆ \subset ⊂ \supset ⊃ \supseteq ⊇ \lceil &rcel; \rceil ⌈ \lceil ⌋ \rceil ⌊ \equiv ≡ \neq ≠ \in ∈ \notin ∉ \rightarrow → \sigma σ \pi π \rho ρ \theta θ \phi φ \GammaΓ \times × \div ÷ \leq ≤ \geq ≥ \lceil ∣ \triangleright \triangleleft ⋈
```

see confsys.encs.concordia.ca/CrsMgr/html-symbols.html

Bipin C Desai

- 59

# **Relational Model**



# **Relational Database Design**



To be used in the spirit of copy-forward! https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

🕩 Bipin C Desai

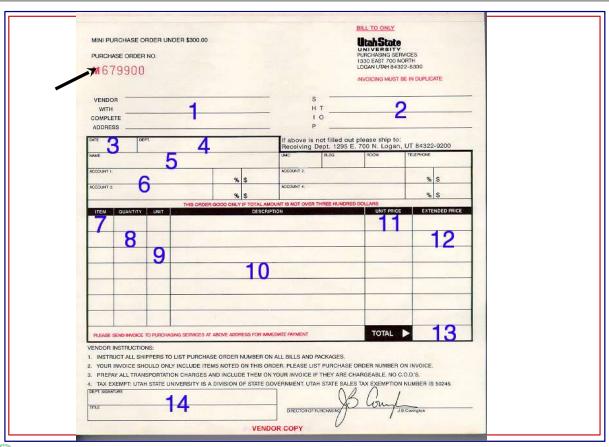
- Relation Scheme and Relational Design
- \* Anomalies in Database:
- \* A Consequence of Bad Design
- Functional Dependency
- Normal Forms

A relation scheme **R** is a plan which gives the attributes involved in one or more relations defined on this relation scheme.

$$\{A 1, A 2, \dots, A_n\},\$$
  
A *i* is defined on domain **D***i*

Relation R on the relation scheme  $\bf R$  is a finite set of mappings or tuples  $\{t_1, t_2, ..., t_p\}$ 

di Bipin C Desai



Bipin C Desai

```
Vendor\hbox{-}\textbf{V}, Address\hbox{-}\textbf{A}, Date\hbox{-}\textbf{D}, Account\hbox{-}\textbf{C}, ShipAddress\hbox{-}\textbf{S} \ ,
```

Partno-**P**, Qty-**Q**, Unit price -**U**, Description-**T**, linetotal -**L**,

•

•

•

Tax- **X**Total - **\$**Signature **G**Date **E** 

PURCHASE\_ORDER (**Pid**, V, A, D, C, S, P, Q, U, T, L, X, \$, G, E)

🕩 Bipin C Desai

\_\_\_\_

### **Anomalies in Database: A Consequence of Bad Design**

Redundancies: the same information is stored more than once.

**Update Anomalies**: The multiple copies may lead to updates which become inconsistent.

**Insertion Anomalies**: Cannot insert some fact unless some other fact is inserted.

**Deletion Anomalies**: Deleting one fact may delete another.

PURCHASE ORDER (**PID**, V, A, D, C, S, P, Q, U, T, L, X, \$, G, E)

What are the redundancies and/or anomalies in PURCHASE\_ORDER?

Bipin C Desai

Functional Dependencies (FDs)

Given attributes X and Y

Y is said to be functionally dependent on X if a given value for each attribute in X, uniquely determines the value of the attributes in Y.

X is called the determinant of the functional dependency (FD) and the FD is denoted as  $X \rightarrow Y$ .

🕩 Bipin C Desai

,

#### STDINF (Name, Course, Phone No, Major, Prof, Grade)

Name	Course	Phone_No	Major Pi	rof Gi	rade
Jones Ng	353 329	237-4539 427-7390	Comp Sci Chemistry	Smith Turner	A B
Jones	328	237-4539	Comp Sci	Clark	В
Martin Dulles	456 293	388-5183 371-6259	Physics Decision Sci	James Cook	A C
Duke	491	823-7293	Mathematics	Lamb	В
Duke	356	823-7293	Mathematics	Bond	in prog
Jones	492	237-4539	Comp Sci	Cross	in prog
Baxter	379	839-0827	English	Broes	С

# The key of **STDINF** is (*Name*, Course)

 ${Name \rightarrow Phone\_No; Name \rightarrow Major; }$  $Name, Course \rightarrow Grade; Course \rightarrow Prof{}$ 

di Bipin C Desai

```
Vendor-V, Address-A, Date-D, Account-C, ShipAddress-S, For each line
Partno-P, Qty-Q, Unit price -U, Description-T, linetotal -L, Total Tax- X

Total for the PO - $
PO approval signature G

Date of PO approval signature E

PURCHASE_ORDER (Pid, V, A, D, C, S, P, Q, U, T, L, X, $, G, E)

FDs in Purchase Order from common business rules

Pid \rightarrowADCSX$GE,

PidP \rightarrow QU,
QU \rightarrow L,
P \rightarrowT
```

Bipin C Desai

9

## Problems due to Redundancy

- \* *Redundancy* creates problems associated with relational schemas:
  - ◆ redundant storage, insert/delete/update anomalies
- ❖ Integrity constraints, e.g., *functional dependencies*, can be used to identify relational schemas with potential anomalies.
- \* Main refinement technique: <u>decomposition</u>
  - ◆ Decompose R(ABCD) with R1(ABC) and R2(BCD).
- Decomposition should be used with care.
  - ◆ Decompose or not to that is the question!

© Bipin C

**Definition**: The decomposition of a relation scheme

$$\mathbf{R} = (A_1, A_2, ...., A_n)$$

is its replacement by a set of relation schemes

$$\{R_1, R_2, ..., R_m\}$$
, such that

$$\mathbf{R_i} \subseteq \mathbf{R}$$
 for  $1 \le i \le m$  and

$$R_1 \cup R_2 \cup ... \cup R_m = R$$
.

Bipin C Desai

1

Inference Axioms: Assume that we have a relation scheme  $R(A_1, A_2, A_3, ..., A_n)$ ; R is a relation on the scheme R and W, X, Y, Z are subsets of R.

F1: Reflexivity:  $X \rightarrow X$ .

**F2:** Augmentation:  $X \rightarrow Y = XZ \rightarrow Y$ , and  $XZ \rightarrow YZ$ .

**F3:** Transitivity:  $X \rightarrow Y$  and  $Y \rightarrow Z = X \rightarrow Z$ .

F4: Additivity:  $X \rightarrow Y$  and  $X \rightarrow Z = X \rightarrow YZ$ .

**F5:** Projectivity:  $X \rightarrow YZ = X \rightarrow Y$  and  $X \rightarrow Z$ .

**F6:** Pseudo-trans:  $X \rightarrow Y$  and  $YZ \rightarrow W = XZ \rightarrow W$ .

R <i>A</i>	В	С	D	Ε
a1	b <sub>1</sub>	c2	d1	<b>e</b> 1
a2	b <u>2</u>	<b>C1</b>	d2	e2
a3	b1	c2	d1	<b>e</b> 3
a3	p3	c3	d3	<b>e</b> 4
a1	b <u>2</u>	<b>C</b> 1	d2	<b>e</b> 5
a4	b4	<b>C4</b>	d4	e6
a3	b <u>2</u>	<b>C</b> 1	d2	<b>e</b> 7
a5	b4	<b>C4</b>	d4	e8

Relation R on the scheme  $\mathbf{R}(A, B, C, D, E)$ 

Illustrates the inference axioms.

FDs:  $B \rightarrow CD$ ,  $B \rightarrow C$ ,  $C \rightarrow D$ ,  $B \rightarrow D$  may hold

Bipin C Desai

13

If **F** is a set of FD's on a relation scheme **R** then  $\mathbf{F}^+$ , the *closure* of **F**, is the smallest set of FD's such that  $\mathbf{F}^+ \supseteq \mathbf{F}$  and no FD can be derived from **F** by using the inference axioms, that are not contained in  $\mathbf{F}^+$ .

If  $\mathbf{R}$  is not specified, then it is assumed to contain all the attributes that appear in  $\mathbf{F}$ .

Suppose R(A, B, C, D):

the FDs that hold on R are:

$$F = \{A \rightarrow B, B \rightarrow C, C \rightarrow D\}$$

Then F<sup>+</sup> also contain:

$$A \rightarrow C$$
,  $A \rightarrow D$ ,  $B \rightarrow D$ 

### Closure of a set of FDs

**Example**: Let  $F = \{W \rightarrow X, X \rightarrow Y, W \rightarrow XY\}$  then  $F^+$  includes the set  $\{W \rightarrow W, X \rightarrow X, Y \rightarrow Y, W \rightarrow X, X \rightarrow Y, W \rightarrow XY, W \rightarrow Y\}$ . The first three FD's follow from axiom F1; the next three FD's are in F, and hence in  $F^+$ . Since  $W \rightarrow XY$  then by axiom F5  $W \rightarrow X$  and  $W \rightarrow Y$ . However,  $F^+$  does not contain a FD, e.g.  $W \rightarrow Z$ , since Z is not contained in the set of attributes that appear in F.

🕩 Bipin C Desai

15

## Functional Dependencies (FDs)

- An FD is a statement about *all* allowable relations on a relational schema.
  - lacktriangle Must be identified based on semantics of application (not from an instance r of a relation on the schema R)
  - Given some allowable instance r of R, we can check if it violates some FD f, but we cannot tell if f holds over R!
- $\star$  K is a candidate key for R means that K  $\to$  R
- \* We require the candidate keys to be minimal
- $If K' \subset K \text{ then } K' \neg \rightarrow R$ 
  - lack However, X  $\rightarrow$  R does not require X to be *minimal(superkey)*!

🕩 Bipin C

## FDs again!

Consider relation

HrEmps (Sin, Name, Grade, Rate, HrWrk):

let us denote it by listing the attributes simply using: SNGRH

- ◆ The schema is a plan for the co-occurrence of the *set* of attributes {S,N,G,R, H}.
- ◆ We may alternately refer to this set of attributes by using just the relation scheme name. (e.g., HrEmps for {SNGRH})
- ◆ In implementation, we usually have only in relation on each relation scheme
- ❖ Some FDs on Hourly\_Emps:
  - Sin is the candidate key:  $S \rightarrow SNGRH$
  - ♦ *Grade* determines Rate (*hourly wages*):  $G \rightarrow R$  (*transitive FD*)

Dipin C 1

# Example

**Emps** 

• Problems due to  $G \rightarrow R$ :

<u>Update anomaly</u>: Can we change R in just the 1st tuple of SNGRH? (e.g., change value of R to 11!)

Insertion anomaly: What if we want to insert an employee and don't Emp know the hourly wage for her grade? Also, we can't insert a rate for a grade unless we have an employee with that grade!

<u>Deletion anomaly</u>: If we delete all employees with grade 35, we lose the information about the wage for this grade!

	S	N	G	R	H
•	123-223-666	Evan	48	10	40
	231-315-368	Lalonde	22	8	30
	131-243-650	Letitia	35	9	30
	434-263-751	Drew	35	9	32
	612-674-134	Ma	48	10	40

	$ \mathbf{S} $	N	G	H
	123-223-666	Evan	48	40
)	231-315-368	Lalonde	22	30
	131-243-650	Letitia	35	30
	434-263-751	Drew	35	32
	612-674-134	Ma	48	40

Wages

G	R
48	10
35	9
22	8

Bipin C

**Definition**: The closure of X under a set of functional dependencies F and written as  $X^+$ , is the set of attributes  $\{A_1, A_2, ..., A_m\}$  such that the FD  $X \rightarrow A_i$  for  $A_i \in X^+$  follows from F by the inference axioms for functional dependencies.

Having found  $X^+$ , we can test if  $F = X \rightarrow Y$  by checking if  $Y \subseteq X^+$ .  $X \rightarrow Y$  is logically implied by F, if and only if  $Y \subseteq X^+$ .

🕩 Bipin C Desai

10

**Lemma**:  $\mathbf{F} = \mathbf{X} \rightarrow \mathbf{Y}$  if and only if  $\mathbf{Y} \subseteq \mathbf{X}^+$ .

**Proof**: Suppose that  $\mathbf{Y} \subseteq \mathbf{X}^+$ . Then by the definition of  $\mathbf{X}^+$ ,  $\mathbf{X} \to A$  can be derived from  $\mathbf{F}$  using the inference rules, for each  $A \in \mathbf{Y}$ .

Now, by the soundness of these rules,

 $\mathbf{F} = \mathbf{X} \rightarrow A \text{ for each } A \in \mathbf{Y}$ 

and by the additivity rule,

$$F = X \rightarrow Y$$
.

Now, suppose that  $\mathbf{F} = \mathbf{X} \to \mathbf{Y}$ . Then by completeness of the inference rules,  $\mathbf{X} \to \mathbf{Y}$  can be derived from  $\mathbf{F}$  using them. By projectivity,  $\mathbf{X} \to A$  can be derived for each  $A \in \mathbf{Y}$ .

This clearly implies that  $Y \subseteq X^+$  by the definition of  $X^+$ .

🕩 Bipin C Desai

# Algorithm to compute X+

Bipin C Desai

21

## Membership Algorithm

```
Input: A set of Fds \mathbf{F}, and the FD \mathbf{X} \to \mathbf{Y}.

Output: Is \mathbf{X} \to \mathbf{Y} \in \mathbf{F^+} or not?

Body:

Compute \mathbf{X^+}

if \mathbf{Y} \subseteq \mathbf{X^+}. then \mathbf{X} \to \mathbf{Y} \in \mathbf{F^+} := true;

else \ \mathbf{X} \to \mathbf{Y} \in \mathbf{F^+} := false;
```

© Bipin C Desai

**Definition**: Given two sets of FD's  $\mathbf{F}$  and  $\mathbf{G}$  over a relation scheme  $\mathbf{R}$ .  $\mathbf{F}$  and  $\mathbf{G}$  are equivalent (i.e.,  $\mathbf{F} \equiv \mathbf{G}$ ) if the closure of  $\mathbf{F}$  is identically equal to the closure of  $\mathbf{G}$  (i.e.,  $\mathbf{F}^+ = \mathbf{G}^+$ ). If  $\mathbf{F}$  and  $\mathbf{G}$  are equivalent then  $\mathbf{F}$  covers  $\mathbf{G}$  and  $\mathbf{G}$  covers  $\mathbf{F}$ .

**Definition**: Given a set of FD's **F**, we say that it is **non-redundant** if no proper subset **F'** of **F** is equivalent to **F**, i.e., no **F'** exists such that **F'\*= F\***.

Bipin C Desai

22

Title: Algorithm: Non-redundant cover

Input: A set of FD's F

Output: A non-redundant cover of F

### **Body**

```
G := F; // initialize G to F
for each FD X → Y in G do
    if X → Y ∈ {F -(X → Y )}<sup>+</sup>

    // i.e.{F-(X → Y)} → X → Y
    then F := {F - (X → Y )};
G := F; // G is the non-redundant cover of F
end;
```

If 
$$\mathbf{F} = \{A \rightarrow BC, CD \rightarrow E, E \rightarrow C, D \rightarrow AEH, ABH \rightarrow BD, DH \rightarrow BC\}$$

then a non-redundant cover for F is:

$$\{A \rightarrow BC, E \rightarrow C, D \rightarrow AEH, ABH \rightarrow BD\}.$$

CD+ under (F-  $CD \rightarrow E$ ) is CDAEHB and it includes E the RHS.

Hence  $CD \rightarrow E$  is redundant.

Bipin C Desai 2

**Definition**: A set of functional dependencies  $\mathbf{F}_{c}$  is a *canonical cover* if every FD in  $\mathbf{F}_{c}$  satisfies the following :

- each FD in  $\mathbf{F}_c$  is simple, (recall that in a simple FD the right hand side has a single attribute i.e., each FD is of the form  $\mathbf{X} \rightarrow A$ );
- for no FD  $X \to A$  with  $Z \subset X$  is  $\{(F_c (X \to A)) \cup (Z \to A)\} \to F_c$ . In other words the left hand side of each FD does not have any extraneous attributes or the *FD's in*  $F_c$  are left reduced;
- no FD  $X \to A$  is redundant i.e.,  $\{ F_c (X \to A) \}$  does not logically imply  $F_c$ .

A canonical cover is sometimes called *minimal*.

Given a set  $\mathbf{F}$  of functional dependencies we can find a canonical set  $\mathbf{F}_c$ ; Obviously  $\mathbf{F}_c$  covers  $\mathbf{F}$ .

<sup>de</sup> Bipin C Desai

**Example**: If  $F = \{A \rightarrow BC, CD \rightarrow E, E \rightarrow C, D \rightarrow AEH, ABH \rightarrow BD, DH \rightarrow BC\}$ 

then a non-redundant cover for **F** is  $\{A \rightarrow BC, E \rightarrow C, D \rightarrow AEH, ABH \rightarrow BD\}$ :

and the canonical cover is  $\{A \rightarrow B, A \rightarrow C, E \rightarrow C, D \rightarrow A, D \rightarrow E, D \rightarrow H, AH \rightarrow D\}.$ 

Bipin C Desai

**Definition**: Given a relation scheme  $\mathbf{R} \{A_1A_2A_3...A_n\}$ , and a set of functional dependencies  $\mathbf{F}$ , a **key**  $\mathbf{K}$  of  $\mathbf{R}$  is a subset of  $\mathbf{R}$  such that the following are satisfied:

- 
$$\mathbf{K} \rightarrow A_1 A_2 A_3 ... A_n$$
 is in  $\mathbf{F}^+$ 

For any  $\mathbf{Y} \subset \mathbf{K}$ ,  $\mathbf{Y} \to A_1 A_2 A_3 ... A_n$  is not in  $\mathbf{F}^*$ 

Given R(A, B, C, D) with F{C $\rightarrow$  D, C $\rightarrow$  A, B $\rightarrow$  C}

Find C+ the closure of C under F.

C<sup>+</sup> is initialized to C

Using C→ D we augment C<sup>+</sup> to CD

Using  $C \rightarrow A$ , we augment  $C^+$  to CDA

No other change is possible; hence closure of C under F is: CDA

Find the candidate key for R.

The closure of B under F is BCDA

Hence B is a candidate key.

🕩 Bipin C Desai

29

Given R(A, B, C, D) with F{D  $\rightarrow$  A, B  $\rightarrow$  C}

Find C<sup>+</sup> the closure of C under F.

C<sup>+</sup> is initialized to C

Since C doesn't appear on the RHS of any FDs in F,

no change is possible; hence closure of C under F is: C

### Find the candidate key for R.

The closure of D under F is DA

The closure of B under F is BC

Since neither of the determinants are possible candidate keys, However,  $BD \rightarrow ABCD$ 

Hence BD is a candidate key.

**Example**: If **R** (*ABCDEH*) and **F** = { $A \rightarrow BC, CD \rightarrow E, E \rightarrow C, D$  $\rightarrow$  AEH, ABH  $\rightarrow$  BD, DH  $\rightarrow$  BC}, Attributes in Left only: Left & Right Right only A, B, C, D, E, H none none Closure of two attributes not involving key already found Closure of one attribute  $A^+ = ABC$  $AB^+ = ABC$  $BH_{+} = BH$  $B^+ = B$  $AC^+ = ABC$ CE+=CE $C_+ = C$  $AE^+ = ABEC$  $CH^+ = CH$  $D^+ = DAEHBC$  $AH^+ = ABCHDE$  $EH^+ = ECH$ 

 $E^{+} = EC$   $H^{+} = H$   $BC^{+} = BEC$   $BE^{+} = BEC$ 

*D* is a candidate key of **R** since  $D \rightarrow ABCDEH$  is in **F+** No other single attribute candidate key.

Also, AH is candidate key since AH+= under **F** is ABCDEH. Since D is only in one RHS, the key must include D or AH No other keys:

Superkeys are DX or AHX where  $X \subseteq \mathbf{R}$ 

Bipin C Desai

31

### Full Functional Dependency

**Definition**: Given a relational scheme **R** and a FD  $X \rightarrow Y$ , then **Y** is *fully functionally dependent* on **X** if there is no **Z**, where **Z** is a proper subset of **X** such that  $Z \rightarrow Y$ .

Thus, the dependency  $X \rightarrow Y$  is left reduced, if there are no extraneous attributes in the left hand side of the dependency.

**Example:** Given **R** (
$$ABCDEH$$
) and **F** = { $A \rightarrow BC, CD \rightarrow E, E \rightarrow C, D \rightarrow AEH, ABH \rightarrow BD, DH \rightarrow BC}$ 

The FD  $ABH \rightarrow BD$  is not left reduced since  $A \rightarrow B$  allows us to eliminate B from the LHS. Also, in  $AH \rightarrow B$  is not left reduced since we have  $A \rightarrow B$  and we can thus eliminate it. So the FD  $ABH \rightarrow BD$  can be replaced by simply  $AH \rightarrow D$ 

**Example**: In the relation scheme **R** (*ABCDEH*) with the FD's,

**F** = { $A \rightarrow BC, CD \rightarrow E, E \rightarrow C, CD \rightarrow AH, ABH \rightarrow BD, DH \rightarrow BC$ }, the dependency  $A \rightarrow BC$  is left reduced and BC is fully functionally dependent on A. However, the functional dependency  $ABH \rightarrow D$ , is not left reduced, the attribute B being extraneous in this dependency.

**Definition**: An 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 **nonprime**.

Bipin C Desai

**Example**: If **R** (ABCDEH) and **F** = { $A \rightarrow BC, CD \rightarrow E, E \rightarrow C, AH \rightarrow D$ }; then AH is the only candidate key of R. The attributes A and H are prime, and the attributes B, C, D, and E are nonprime.

**Definition**: Given a relation scheme **R** with the functional dependencies **F** defined on the attributes of **R**. Let **K** be a candidate key. If **X** is a proper subset of **K**, and if  $\mathbf{F} = \mathbf{X} \to A$ , then, A is said to be *partially dependent* on **K**.

In the relation scheme

**STUDENT\_COURSE\_INFO**(*Name*, *Course*, *Grade*, *Phone\_No*, *Major*, *Course\_Dept*) with the FD's, **F** = {Name → Phone\_NoMajor, Course → Course\_Dept, NameCourse → Grade}.

Here *NameCourse* is a candidate key, *Name* and *Course* are prime attributes. *Grade* is fully functionally dependent on the candidate key.

Phone\_No, Course\_Dept, and Major are partially dependent on the candidate key.

Bipin C Desai

### Transitive Dependency

**Definition**: Given a relation scheme **R** with the functional dependencies **F** defined on the attributes of **R**.

Let **X** and **Y** be subsets of **R** and let *A* be an attribute of **R** such that  $X \not\subset Y$ ,  $A \not\subset XY$ .

If the set of functional dependencies  $\{X \to Y, Y \to A\}$  is implied by F (i.e.,  $F = X \to Y \to A$  and  $F \neg = Y \to X$ ), then A is *transitively dependent* on X.

In the relation scheme

**Employee**(*Emp\_Name*, *Department*, *Manager*) with the function dependencies

F = {Emp\_Name → Department, Department → Manager},
Emp \_Name is the key and Manager is transitively
Dependent on the key since
Emp Name → Department → Manager.

Bipin C Desai

**Content Preserving**: the original relation can be derived from the decomposed relations (lossless join decomposition)

**Dependency Preserving**: the original set of constraints can be derived from the dependencies in the decomposed relations.

**Free of interrelation join constraints**: if there are no dependencies that can only be derived from the join of two or more decomposed relations

**Definition**: An unnormalized relation contains **nonatomic** values.

**Definition**: A relation scheme is said to be in the **first normal form** (**1NF**) if the values in the domain of each attribute of the relation are atomic.

1NF has NO NON-ATOMIC ATTRIBUTES.

**Definition**: A relation scheme **R**<**S**, **F**> is in the **second normal form** (**2NF**) if all non-prime attributes are fully functionally dependent on the relation key(s).

2NF has NO PARTIAL DEPENDENCY

Bipin C Desai

Assignment(Emp#, Name, Dept, Proj#, Hours, Lab) Emp# → NameDept, Proj# → Lab, Emp#Proj# → Hours

Example of an unnormalized (non-normal form) relation

di Bipin C Desai

Smith D1 P1 123 5 L1 123 Smith P2 30 L1 D1 234 Ma D2 P1 20 L1 234 Ma P3 10 L2 D2 234 Ma L3 D2 P4 5 35 345 Russo D1 L1 P1

Bipin C Desai

4

Assignment(Emp#, Name, Dept, Proj#, Hours, Lab) Emp# → NameDept, Proj# → Lab,				
123 Smith D1 123 Smith D1 123 Smith D1 234 Ma D2 234 Ma D2 234 Ma D2 345 Russo D1 23 23	2 L1 3 L2 4 L3 3 P1 5 3 P2 30			

Assignment(Emp#, Name, Dept, Proj#, Hours, Lab) Emp# → NameDept, Proj# → Lab, Dept → Lab Emp#Proj# → Hours

123 Smith D1 P1 35 L1 234 Ma D2 P2 35 L2 345 Russo D1 P3 35 L1

One can't say that there is no anomaly from the contents in the database at a given point in time!

Dipin C Desai

**Definition**: A relation scheme R < S, F > is in the **third normal form(3NF)**if for all nontrivial FD in F + of the form  $X \rightarrow A$ , either X contains a key (i.e., X is a superkey) or A is a prime attribute.

A database scheme is in the third normal form if every relation scheme included in the database scheme is in the third normal form.

BNF HAS NO TRANSITIVE DEPENDENCY

Lossless Join Decomposition

**Definition**: A decomposition of a relation scheme R < S, F > into the relation schemes  $R_i$  ( $1 \le i \le n$ ) is said to be *lossless join decomposition* or simply *lossless* if for every relation r(R) that satisfies the FD's in F, the natural join of the projections of r gives the original relation R: i.e.,

$$\mathbf{r} = \boldsymbol{\Pi}_{\mathbf{R}\mathbf{1}}(\mathbf{r}) \; \trianglerighteq \!\! \triangleleft \; \boldsymbol{\Pi}_{\;\mathbf{R}\mathbf{2}}(\mathbf{r}) \; \trianglerighteq \!\! \triangleleft \; .... \; \trianglerighteq \!\! \triangleleft \; \boldsymbol{\Pi}_{\;\mathbf{R}\mathbf{n}}(\mathbf{r})$$

If  $r \subset \Pi_{R1}(r) \bowtie \Pi_{R2}(r) \bowtie .... \bowtie \Pi_{Rn}(r)$  then the decomposition is called *lossy*.

$$r \subseteq \Pi_{R1}(r) \ \trianglerighteq \lnot \ \Pi_{R2}(r) \ \trianglerighteq \lnot \ .... \ \trianglerighteq \lnot \ \Pi_{Rn}(r) \ \text{is always true}.$$

<sup>di</sup> Bipin C Desai

Assignment(Emp#, Name, Dept, Proj#, Hours, Lab)						
Emp# → NameDept, Proj# → Lab,						
Emp#I	Proj# → Ho	ours				
			P1	L1		
123	Smith	D1	P2	L1		
			P3	L2		
123	Smith	D1	P4	L3		
234	Ma	D2		LO		
234	Ma	D2	123	P1	5	
234	Ma	D2	123	P2	30	
345	Russo	וט	234	P1	20	
			234	P3	10	
			234	P4	5	
			345	P1	35	

```
Join the first two relations:
```

When we join the third relation,

Creates extraneous tuples the extraneous tuples are eliminated!

Hence, the decomposition is lossless

```
123 Smith D1 P1 L1
123 Smith D1 P2 L1
123 Smith D1 P3 L2
124 Smith D1 P4 L3
234 Ma
         D2 P1 L1
234 Ma
         D2 P2 L1
         D2 P3 L2
234 Ma
         D2
234 Ma
             P4 L3
345 Russo D1
             P1 L1
345 Russo D1
             P2 L1
345 Russo D1
             P3 L2
345 Russo D1
             P4 L3
```

```
123 P1 5
123 P2 30
234 P1 20
234 P3 10
234 P4 5
345 P1 35
```

Bipin C Desai

47

# Assignment(Emp#, Name, Dept, Proj#, Hours, Lab)

### Emp# $\rightarrow$ NameDept, Proj# $\rightarrow$ Lab, Emp#Proj# $\rightarrow$ Hours

```
123
     Smith
            D1
                 P1
                       5
                         L1
                               123
                                     Smith
                                              D1
                                                   Р1
                                                          5
                                                             L1
     Smith
123
            D1
                 P2
                      30
                          L1
                               123
                                     Smith
                                              D1
                                                   Р1
                                                         20
                                                             L1
234
            D2 P1
     Ma
                      20 L1
                               123 Smith
                                            D1
                                                 Ρ1
                                                       35
                                                            L1
234
            D2
                      10 L2
     Ma
                 Р3
                                     Smith
                                                   Р2
                                                         30
                                                             L1
                               123
                                              D1
234
     Ma
            D2
                 Ρ4
                      5
                         L3
345
                      35
     Russo
            D1
                 Ρ1
                          L1
                               234
                                     Ма
                                              D2
                                                   Ρ1
                                                          5
                                                             L1
                               234
                                              D2
                                                   Р1
                                                         20
                                                             L1
                                     Ma
                               234
                                              D2
                                                   Ρ1
                                                         35
                                                             L1
                                     Ma
                         5 L1
                    P1
123
    Smith D1 L1
                               234
                                             D2
                                                   P2
                                                         30
                                                             L1
                    P2
                        30 L1
                                     Ма
234
           D2
               L1
    Ма
                    P1
                        20 L1
                               234
                                                   Р3
                                                             L2
                                     Ма
                                              D2
                                                         10
234
    Ма
           D2
               L2
                    Р3
                        10 L2
234
    Ma
           D2
               L3
                               234
                                              D2
                                                   Ρ4
                                                          5
                                                             L3
                                     Ma
                    P4
                         5 L3
           D1
345
    Russo
               L1
                    P1
                        35 L1
                               345
                                                   Р1
                                                          5
                                                             L1
                                     Russo
                                             D1
                               345
                                              D1
                                                   Р1
                                                         20
                                                             L1
                                     Russo
  A lossy join decomposition
                               345
                                              D1
                                                   Ρ1
                                                         35
                                                             L1
                                     Russo
```

Bipin C Desai

345

Russo

Ρ2

D1

30

L1

**Definition**: Given a relation scheme R < S, F > where F is the associated set of functional dependencies on the attributes in S. Consider that R is decomposed into the relation schemes  $R_1, R_2, ..., R_n$  with the functional dependencies  $F_1, F_2, ..., F_n$ .

Then this decomposition of **R** is **dependencies preserving**, the closure of **F'** (where  $F' = F_1 \cup F_2 \cup ... \cup F_n$ ) is identical to  $F^+$  (i.e.,  $F'^+ \equiv F^+$ ).

© Bipin C Desai

**Theorem**: A decomposition of relation scheme  $\mathbf{R} < (X, Y, Z)$ ,

- F> into say  $R_1<(X,Y)$ ,  $F_1>$  and  $R_2<(X,Z)$ ,  $F_2>$  is:
- (i) dependency preserving if every functional dependency in
- **R** can logically derived from the functional dependencies of
- $\mathbf{R}_1$  and  $\mathbf{R}_2$  i.e.,  $(\mathbf{F}_1 \cup \mathbf{F}_2)^+ = \mathbf{F}^+$ , and
- (ii) is lossless if the common attributes **X** of  $\mathbf{R}_1$  and  $\mathbf{R}_2$  form
- a superkey of at least one of these i.e.,  $X \rightarrow Y$  or  $X \rightarrow Z$ .

```
for each decomposed relation R, do
 if an attribute A_i is included in \mathbf{R}_i,
  then TABLE_LOSSY(i,j) := \alpha_{Ai}
  else TABLE_LOSSY(i,j) := \beta_{iAi}
change := true
while (change) do
 for each FD X \rightarrow Y in F do
   if rows i and j exist such that the \alpha_r symbol appears in
          each column corresponding to the attributes of X
          then if one of the symbol in the Y column is \alpha_r, the other \beta_r
          then make replace \beta_r with \alpha_r
           else if the symbols are \ensuremath{\mathbb{G}_{pm}} and \ensuremath{\mathbb{G}_{qm}}
                                      then make both of them, say, \mathbb{S}_{nm};
      else change := false
i := 1
If there is a row will all \alpha then the decomposition is lossless
```

Bipin C Desai

**Example:** R(A,B,C,D) with the functional dependencies  $F \{A \rightarrow B, A \rightarrow C, C \rightarrow D\}$ . Consider the dependence preserving decomposition of R into  $R_1(A,B,C)$  and  $R_2(C,D)$ .

 $A \quad B \quad C \quad D \qquad A \quad B \quad C \quad D$ 

**Example**: R(A, B, C, D, E) with the functional dependencies F  $\{AB \rightarrow CD, A \rightarrow E, C \rightarrow D\}$ . Then the decomposition of **R** into  $R_1(A,B,C)$  and  $R_2(B,C,D)$  and  $R_3(C,D,E)$  is lossy.

	F { <i>AB</i> –	CD, A	$\rightarrow$ E, C	$\rightarrow D$ }.	
	A	В	C	D	E
$R_1(A,B,C)$	α	α	α	β	β
$R_2(B,C,D)$	β	α	α	α	β
$R_3(C,D,E)$	β	β	α	α	α
	C-	$\rightarrow D$			
	A	В	C	D	$\mathbf{E}$
$R_1(A,B,C)$	α	α	α	$\beta\alpha$	β
$R_2(B,C,D)$	β	α	α	α	β
$R_3(C,D,E)$	β	β	α	α	$\alpha$

No further changes – no row with all  $\alpha$  Hence , lossy decomposition!

🕩 Bipin C Desai

53

# Algorithm to check if a decomposition is dependency preserving

**Input:** A relation scheme and a set F of FDs: a projection  $(R_1, R_2, ..., R_n)$  of R with the FDs  $(F_1, F_2, ..., F_n)$ .

**Output:** Whether the decomposition is dependency preserving or not.

### Body:

```
F'+\_=\_F+:= true;
F':=\emptyset;
for i:=1 	ext{ to } n 	ext{ } do
F':=F' \cup F_i;
for each 	ext{ FD } X \rightarrow Y \in F 	ext{ } and 	ext{ } while (F'+\_=\_F+) 	ext{ } do
// 	ext{ compute } X'+, 	ext{ the closure of } X 	ext{ } under F'
if Y \not\subset X'+ 	ext{ } then 	ext{ } F'+\_=\_F+:= 	ext{ } false;
```

**Example**: Consider the relation scheme  $\mathbf{R}(A,B,C,D)$  with the Ds  $\mathbf{F} = \{A \to B, A \to C, C \to D\}$ . Here, the decomposition of  $\mathbf{R}$  into  $\mathbf{R}_1 < (A,B,C)$ ,  $\{A \to B, A \to C\}$  and  $\mathbf{R}_2 < (C,D)$ ,  $\{C \to D\} >$  is dependence preserving, since in this case each FD in  $\mathbf{F}$  is included in  $\mathbf{F}'$  (where  $\mathbf{F}' = \mathbf{F}_1 \cup \mathbf{F}_2$ ).

**Example**: Consider the relation Student\_Advisor(*Name*, *Dept*, *Advisor*) with the FDs  $\mathbf{F} = \{N \to D, N \to A, A \to D\}$ . Here, its decomposition into S\_Pr<(N, A),  $\{N \to A\}$ >, and D\_A <(D, A),  $\{A \to D\}$ > is dependence preserving, since  $N \to D$  is implied by  $(N \to A) \cup (A \to D)$ ; in addition the decomposition is lossless.

© Bipin C Desai

**Example**: Consider  $\mathbf{R}(A,B,C,D)$  with the FDs

**F**  $\{A \rightarrow B, A \rightarrow C, C \rightarrow D\}$  and its decomposition into

 $\mathbf{R}_1(A,B,C)$  with the FDs  $\mathbf{F}_1 = \{A \to B, A \to C\}$  and

 $\mathbf{R}_2(C,D)$  with the FDs  $\mathbf{F}_2 = \{C \to D\}$ .

This decomposition is dependence preserving since all the original FD's can be logically derived from  $\mathbf{F_1}$  and  $\mathbf{F_2}$ .

**Example**:  $\mathbf{R}(A,B,C,D)$  with the FDs F  $\{A \to B, A \to C, A \to D\}$  is decomposed into  $\mathbf{R}_1(A,B,D)$  with the FDs  $\mathbf{F}_1 = \{A \to B, A \to D\}$  and  $\mathbf{R}_2(B,C)$  with the FDs  $\mathbf{F}_2 = \{\}$  is not dependence preserving since the FD  $A \to C$  is not implied by any FD's in  $\mathbf{R}_1$  or  $\mathbf{R}_2$ .

**Example** The decomposition of the relation Concentration (Student, Major\_or\_Minor, Dept, Advisor), with the FDs  $\{SM_mF_s \rightarrow A, A \rightarrow F_s\}$  into the relations  $SM_mA$  and  $F_sA$  is not FD preserving since  $\mathbf{F'} = A \rightarrow F_s$  and the FD  $SM_mF_s \rightarrow A$  is not implied by  $\mathbf{F'}$ .

Bipin C Desai

```
Third Normal Form Decomposition Algorithm
Input: A relation Scheme \mathbf{R}, a set of Canonical FDs \mathbf{F}_c, and \mathbf{K} a candidate key of \mathbf{R}.(K must have any attributes \not\in \mathbf{F}_c)
Output: A collection of third normal form relation schemes (\mathbf{R}_1, \mathbf{R}_2, ..., \mathbf{R}_i) which are dependency preserving and lossless. \mathbf{i} := 0
if there is a dependency \mathbf{X} \to \mathbf{Y} in \mathbf{F}_c such that all the attributes that remain in \mathbf{R} are included in it{
i:= i+1; output \mathbf{R} as \mathbf{R}_i { \mathbf{X}, \mathbf{Y}};}
else{ for each FD \mathbf{X} \to A in \mathbf{F}_c {
i:= i+1; form \mathbf{R}_i < {\mathbf{X}, A}, \mathbf{F}_i { \mathbf{X} \to A} > With < (\mathbf{X}, A), {\mathbf{X} \to A} > and < (\mathbf{X}, A), {\mathbf{X} \to A} > with if \mathbf{F}_j for 1 \le j \le i not satisfies \mathbf{K} \subseteq \mathbf{X} {
i:= i+1; form \mathbf{R}_i { \mathbf{K}}
```

**SHIPPING**(Ship, Capacity, Date, Cargo, Value)

Ship → Capacity, ShipDate → Cargo, CargoCapacity → Value

The given set of FD's is in canonical form.

A candidate key of the relation is ShipDate.

Decompose into:

 $\mathbf{R}_1(Ship, Capacity)$  with the FD:  $Ship \rightarrow Capacity$ ,

 $\mathbf{R}_{>}(Ship, Date, Cargo)$  with the FD:  $ShipDate \rightarrow Cargo$ ,

 $\mathbf{R}_3$ (Cargo, Capacity, Value) with the FD: CargoCapacity  $\rightarrow$ Value

Bipin C Desai

59

```
Consider the relation scheme Student_info(Student(S), Major(M), Student_Department(S_d), Advisor(A), Course(C), Course_Department(C_d), Grade(G), Professor(P), Prof_Department(P_d), Room(R), Day(D), Time(T)) with \{S \rightarrow M, S \rightarrow A, M \rightarrow S_d, S \rightarrow S_d, A \rightarrow S_d, C \rightarrow C_d, C \rightarrow P, P \rightarrow P_d, RTD \rightarrow C, RTD \rightarrow P, TPD \rightarrow R, TSD \rightarrow R, TDC \rightarrow R, TPD \rightarrow C, TSD \rightarrow C, SC \rightarrow G\} Redundant FDs \{S \rightarrow Sd, RTD \rightarrow P, TDC \rightarrow R, TPD \rightarrow C, TSD \rightarrow R\} The primary key is TSD.

3NF decomposition is: \langle \mathbf{R}_1(SMA), \{S \rightarrow MA\} \rangle; \langle \mathbf{R}_2(MSd), \{M \rightarrow Sd\} \rangle, \langle \mathbf{R}_3(ASd);, \{A \rightarrow Sd\} \rangle; \langle \mathbf{R}_4(CCdP), \{C \rightarrow CdP\} \rangle; \langle \mathbf{R}_5(PPd), \{P \rightarrow Pd\} \rangle; \langle \mathbf{R}_6(RTDC), \{RTD \rightarrow C\} \rangle; \langle \mathbf{R}_7(TPDR), \{TPD \rightarrow R\} \rangle; \langle \mathbf{R}_8(TSDR), \{TSD \rightarrow R\} \rangle; \langle \mathbf{R}_8(SCG), \{SC \rightarrow G\} \rangle.
```

Name	Student#	Course	Grade		
Jones	23714539	353	А		
Ng	42717390	329	А		
Jones	23714539	328	in prog		
Martin	38815183	456	С		
Dulles	37116259	293	В		
Duke	82317293	491	С		
Duke	82317293	353	in prog		
Jones	23714539	491	С		
Evan	11011978	353	A+		
Baxter	83910827	379	in prog		
	The Grade	Relation			
Suppose the FDs are $Student\# \rightarrow Name \text{ and } Name \rightarrow Student\#?$					
Is it in 3NF? Any redundancies?					

Bipin C Desai

**Definition**: A normalized relation scheme R < S, F > is in the **Boyce Codd normal form** if for every nontrivial FD in F + of the form  $X \rightarrow A$  where  $X \subseteq S$  and  $A \in S$ , X is a superkey of R.

A database scheme is in the BCNF if every relation scheme in the database scheme is in the BCNF.

The relation GRADE is not in the BCNF because of the dependencies *Student#* → *Name* and *Name* → *Student#* are nontrivial and their determinants are not superkeys of GRADE.

# i := 0; S := { R(U) }; all\_BCNF := false; Find a non-redundant cover F' from F while ( ¬all\_BCNF ) { if ∃((X → Y) ∈ F' + Y ⊄ X) (XY ⊆ R<sub>j</sub>) X ¬→R<sub>j</sub>{ i := i+1; <R<sub>i</sub>{X, Y}, X → Y> ∪ S R<sub>j</sub> := R<sub>j</sub> - Y; } else all\_BCNF := true; }

Algorithm: Lossless BCNF Decomposition

Bipin C Desai

```
Example: Let us find a BCNF decomposition of the relation: \mathbf{SHIPPING}(Ship, Capacity, Date, Cargo, Value) \{S \rightarrow Cap, SD \rightarrow Cargo, CargoCap \rightarrow V \} There are no redundant FD'S in the set Since S \rightarrow Cap and since Ship \rightarrow \mathbf{SHIPPING} replace \mathbf{SHIPPING} with: \mathbf{R}_1(S, Cap) and \mathbf{R}_2(S, D, Cargo, V). The decomposition is lossless but not FD preserving: the FD CargoCap \rightarrow V is not implied by \{Ship \rightarrow Cap, SD \rightarrow Cargo\}. A BCNF decomposition which is lossless and FD preserving: \mathbf{R}_1(Cargo, Capacity, Value) with the FD CargoCapacity \rightarrow Value, \mathbf{R}_2(Ship, Capacity) with the FD Ship \rightarrow Capacity \mathbf{R}_3(Ship, Date, Cargo) with the FD ShipDate \rightarrow Cargo
```

Given  $F_1$ ={PersonName  $\rightarrow$  City, Street; PersonName,CompName  $\rightarrow$  Salary; CompName  $\rightarrow$  CompCity; PersonName  $\rightarrow$  MgrName} and Given  $F_2$ ={CompName  $\rightarrow$  CompCity; PersonName, CompName, CompCity  $\rightarrow$  Salary; PersonName  $\rightarrow$  City; PersonName  $\rightarrow$  Street; PersonName, City  $\rightarrow$  MgrName}

Does  $F_1$  cover  $F_2$ ?

Bipin C Desai

6

Given F₁={PersonName → City, Street; PersonName,CompName → Salary; CompName → CompCity; PersonName → MgrName}

Candidate key: PersonNameCompanyName No redundant attributes on the LHS. No redundant FDs

 $R(P_n, C_n, M_n, C, S, C_c, \$)$  can be decomposed, using the 3NF algorithm (FD preserving and losslessly ) into:

 $R_1(P_nCS)$ ,  $R_2(P_nC_n\$)$ ,  $R_3(C_nC_c)$ ,  $R_4(P_nM_n)$ 

```
Given F_1={PersonName → City, Street;
 PersonName,CompName → Salary;
 CompName → CompCity;
 PersonName → MgrName}
 Candidate key: PersonNameCompanyName P<sub>n</sub>C<sub>n</sub>
 No redundant attributes on the LHS.
 No redundant FDs
 R(P_n, C_n, M_n, C, S, C_c, \$) can be decomposed, using the BCNF
 algorithm as follows:
 P_nC_nM_nCSC_c$
Do not use the FD P_nC_n \rightarrow $
R_1(P_nCS) P_nC_nM_nC_e$
WHY???
      R_2(P_nC_n\$) \sim P_nC_nM_nC_c
               R_3(P_nM_n) P_nC_nC_c
                  R_4(C_nC_c) P_nC_n
 P_nC_n is already in LHS of R_2, we can combine it with it(drop it).
Bipin C Desai
```

```
Given F_1={PersonName \rightarrow City, Street;
PersonName,CompName → Salary;
CompName → CompCity;
PersonName → MgrName}
Candidate key: PersonNameCompanyName
No redundant attributes on the LHS.
No redundant FDs
Since P<sub>n</sub>C<sub>n</sub> is the key we need not use it in decompoing it in
the second step(as shown on the previous slide!!)
Hence, R(P_n, C_n, M_n, C, S, C_c, \$) can be decomposed,
alternatively, using the BCNF algorithm as follows:
   P<sub>n</sub>C<sub>n</sub>M<sub>n</sub>CSC<sub>c</sub>$
R_1(P_nCS) P_nC_nM_nC_c$
       R_2(P_nM_n) P_nC_nC_c$
                  R_3(C_nC_c) R_4(P_nC_n\$)
```

Given R = 
$$\{ABCDEGIJK\}$$
,  
 $\{AB \rightarrow CDE, E \rightarrow G, B \rightarrow G, BG \rightarrow AIJ, IJ \rightarrow K\}$   
 $F = \{AB \rightarrow C \qquad BG \rightarrow A \qquad AB \rightarrow D \qquad BG \rightarrow I$ 

$$\begin{array}{ccccc} AB \rightarrow D & BG \rightarrow I \\ AB \rightarrow E & BG \rightarrow J \\ E \rightarrow G & IJ \rightarrow K \\ B \rightarrow G \} \end{array}$$

$$F_c = \{B \rightarrow A, B \rightarrow C, B \rightarrow D, B \rightarrow E, B \rightarrow I, B \rightarrow J, IJ \rightarrow K, E \rightarrow G\}$$

B is a candidate key.

Bipin C Desai

69

### (ABCDEGIJK)

$$F_c = \{B \rightarrow A, B \rightarrow C, B \rightarrow D, B \rightarrow E, B \rightarrow I, B \rightarrow J, E \rightarrow G, IJ \rightarrow K \}$$

B is a candidate key.

3NF: Since there is no single FD which includes all attributes in R, we create a relation for each FD:

R1(AB), R2(BC), R3(BD), R4(BE), R5(BI), R6(BJ), R7(EG), R8(IJK)

Combine the relations with the same LHS:

R1'(ABCDEIJ), R7(EG), R8(IJK)

Why did we not include G in R1'?

(ABCDEGIJK)

$$F_{c} = \{B \rightarrow A, B \rightarrow C, B \rightarrow D, B \rightarrow E,$$
  

$$B \rightarrow I, B \rightarrow J, E \rightarrow G, IJ \rightarrow K \}$$

B is a candidate key.

BCNF decomposition

$$\left(ABCDEGIJK\right) \text{ since } (IJ \to K) \in F'^+ \land K \not\subset IJ) \land (IJK \subseteq R) \land IJ \to R$$

$$R1(IJK),\,IJ \to K \qquad (ABCDEGIJ) \text{ since } (E \to G) \in F'^+ \land G \not\subset E) \land (EG \subseteq R) \land E \dashv \to R$$

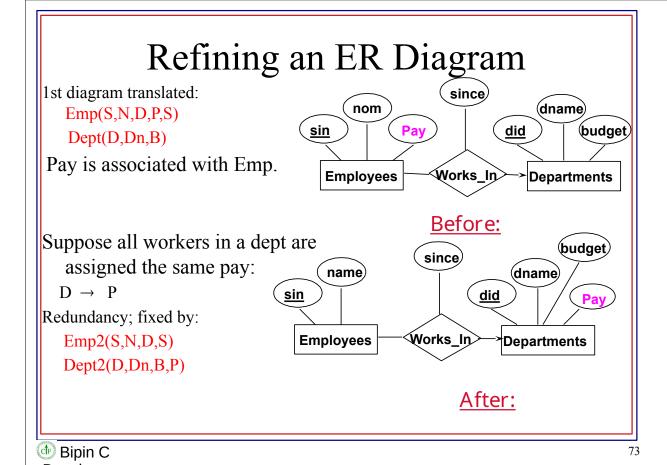
R2(EG), E 
$$\rightarrow$$
 G R3(ABCDEIJ), B  $\rightarrow$  A, B  $\rightarrow$  C,  
B  $\rightarrow$  D, B  $\rightarrow$  E,  
B  $\rightarrow$  I, B  $\rightarrow$  J,

Bipin C Desai

71

### Decompose:

Projects<{Employee, Project, Dept,Part, QtyUsed, HrsWorked}, {Employee,Project → HrsWorked; Project → Dept; Project,Part → QtyUsed }>



### Normal Forms: Conclusions

- -Returning to the issue of schema refinement, the first question to ask is whether any refinement is needed!
- -If a relation is in a certain *normal form* (BCNF, 3NF etc.), it is known that certain kinds of problems are avoided/minimized. This can be used to help us decide whether decomposing the relation will help.
- Role of FDs in detecting redundancy:
- Consider a relation R with 3 attributes, ABC.
  - No FDs hold: There is no redundancy here.
  - However if A → B: Then, several tuples could have the same A value, and if so, they'll all have the same B value! We need refinement!

ⓓ Bipin C 74

### Boyce-Codd Normal Form (BCNF)

- Reln R with FDs F is in BCNF if, for all  $X \rightarrow A$  in
  - $A \in X$  (called a *trivial* FD), or
  - X contains a key for R.
- In other words, R is in BCNF if the only non-trivial FDs that hold over R are key constraints.
  - No dependency in R that can be predicted using FDs alone.
  - If we are shown two tuples that agree upon the X value,
  - we cannot infer the A value in one tuple from the A
  - value in the other.
  - If example relation is in BCNF, the 2 tuples must be identical (since X is a key).

X	Y	A
X	y1	a
X	y2	?

🕩 Bipin C

75

### Third Normal Form (3NF)

- Relation R with FDs F is in 3NF if, for all  $X \to A$  in
  - $A \in X$  (called a *trivial* FD), or
  - X contains a key for R, or
  - A is part of some key for R.
- Minimality of a key is crucial in third condition above!
- If R is in BCNF, obviously it is also in 3NF.
- If R is in 3NF, some redundancy is possible. It is a compromise, used when BCNF not achievable (e.g., no ``good'' decomposition, or performance considerations).
  - Lossless-join, dependency-preserving decomposition of R into a collection of 3NF relations always possible.

Bipin C

### When is R not in 3NF?

- If 3NF violated by  $X \rightarrow A$ , one of the following holds:
  - X is a subset of some key K(Partial Dep)
    - We store (X, A) pairs redundantly.
  - X is not a proper subset of any key.(Trans. Dep)
    - There is a chain of FDs  $K \to X \to A$ , which means that we cannot associate an X value with a K value unless we also associate an A value with an X value.
- But: even if relation is in 3NF, these problems could arise.
  - (Member, Chalet, Date, Card),  $M \rightarrow C$ ,  $C \rightarrow M$  is in 3NF, but for each reservation of member, same (M, C) pair is stored.
- Thus, 3NF is indeed a compromise relative to BCNF.

d Bipin C 7'

### Decomposition of a Relation Scheme

- Suppose that relation R contains attributes A1 ... An. A <u>decomposition</u> of R consists of replacing R by two or more relations such that:
  - Each new relation scheme contains a subset of the attributes of R (and no attributes that do not appear in R), and
  - Every attribute of R appears as an attribute of one of the new relations.
- Intuitively, decomposing R means we will store instances of the relation schemes produced by the decomposition, instead of instances of R.

<sup>d</sup> Bipin C 78

### **Example Decomposition**

- Decompositions should be used only when needed.
  - SNPGRH has FDs  $S \rightarrow SNPGRH$  and  $G \rightarrow R$
  - Second FD causes violation of 3NF; R values repeatedly associated with G values. Easiest way to fix this is to create a relation GR to store these associations, and to remove R from the main schema:
    - i.e., we decompose SNPGRH into SNPGH and GR
- The information to be stored consists of SNPGRH tuples. If we just store the projections of these tuples onto SNPGH and GR, are there any potential problems that we should be aware of?

🕩 Bipin C

70

Given R(A, B, C, D) with F{C $\rightarrow$  D, C $\rightarrow$  A, B $\rightarrow$  C}

If R is not in BCNF, decompose it into a set of BCNF relations that preserve the FDs.

B is the candidate key.

Both  $C \rightarrow D$  and  $C \rightarrow A$  cause BCNF violations.

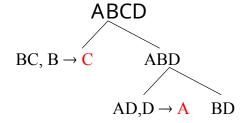
One way to obtain a (lossless) join preserving decomposition is to decompose R into

AC, BC, and CD.

Given R(A, B, C, D) with  $F\{D \rightarrow A, B \rightarrow C\}$ 

Here BD is a candidate key. R is in 1NF but not 2NF due to the partial dependencies.

$$B \rightarrow C$$
 and  $D \rightarrow A$ 



The decomposition: AD, BC, BD

- obtained by first decomposing R into AD, BCD;
- followed by decomposing BCD into BC and BD

is BCNF and lossless and join-preserving.

Bipin C Desai

01

## Review of Relational Design

©Bipin C. DESAI

Example relation:

**EMPLOYEE** 

(EID, Project, Component, EName, Building, Room, TelNo)

Note: Keys are underlined.

What are the FDs?

What is the normal form of the relation?

-Only one phone in each room.

R(I, P, C, N, B, R, T)

FD:  $\{I \rightarrow NTBR, BR \rightarrow T, T \rightarrow BR\}$ 

Key: IPC

All partial dependencies.

R1(BRT), R2(INT), R3(IPC)

Bipin C Desai

83

Example relation:

**EMPLOYEE** 

(EID, Project#, Component, Qty, EName, Building, Room, TelNo, Hours)

Note: Keys are underlined.

What are the FDs? What is the normal form of the relation?

- -Only one phone in each room.
- -There is a m-to-n relationship between projects and components
- -Each employee works a number of hours on a project R(I, P, C, Q, N, B, R, T, H)

Key: IPC

FD:  $\{I \rightarrow NTBR, BR \rightarrow T, T \rightarrow BR, IP \rightarrow H, PC \rightarrow Q\}$ 

R(I, P, C, Q, N, B, R, T, H)

Key: IPC

FD:  $F\{I \rightarrow NTBR, BR \rightarrow T, T \rightarrow BR, IP \rightarrow H, PC \rightarrow Q\}$ 

The corresponding F<sub>c</sub> is

 $\mathsf{F_c}\{\mathsf{I} \to \mathsf{N, I} \to \mathsf{T, BR} \to \mathsf{T, T} \to \mathsf{BR, IP} \to \mathsf{H, PC} \to \mathsf{Q}\}$ 

A 3NF decomposition is: R(INT), R(BRT), R(IPH), R(PCQ), R(IPC),

It is also in BCNF!

🕩 Bipin C Desai

85

R(I, P, C, N, B, R, T)

Key: IPC

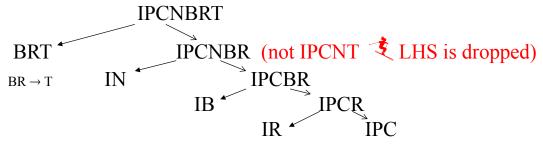
FD:  $\{I \rightarrow N, I \rightarrow T, I \rightarrow B, I \rightarrow R, BR \rightarrow T, T \rightarrow BR\}$ 

3NF decomposition: R1(IN), R2(IT), R3(IB), R4(IR), R5(BRT),

R6(IPC)

Can we combine R1 .... R4?

BCNF decomposition:

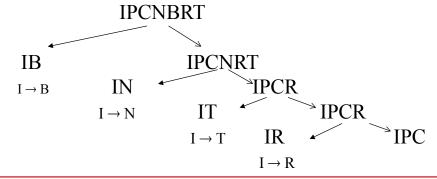


R(I, P, C, N, B, R, T)

Key: IPC

FD:  $\{I \rightarrow N, I \rightarrow T, I \rightarrow B, I \rightarrow R, BR \rightarrow T, T \rightarrow BR\}$ 

Another BCNF Decomposition which is lossless but NOT FD preserving:



🕩 Bipin C Desai

07

### **Example: 1NF but not 2NF**

ORDER(SuplNo, Address, Distance, PartNo, Price)

Assume each supplier is located in only one Address.

What are the FDs?

What are the anomalies?

PO (SuplNo,, PartNo, Price)

Supplier (SuplNo, Address, Distance)

What are the FDs in each and the normal form of each?

Any anomalies?

Bipin C Desai

89

### **Decomposition (into 3NF):**

SUPPLIER\_Address (SuplNo, Address)

Address\_DISTANCE (Address, Distance)

### **Example (3NF but not BCNF):**

Can\_Supply (SuplNo, SuplName, Address, PartNo, Price)

### **Functional Dependencies:**

We assume that SuplNo, Address, PartNo are always unique Thus we have two candidate keys:

(SuplNo, PartNo) and (SuplName, Address, PartNo)

and we have the following dependencies:

(SuplNo, PartNo) → Price

(SuplNo, PartNo) → SuplName, Address

(SuplName, Address, PartNo) → Price

(SuplName, Address, PartNo) → SuplNo

SuplName, Address → SuplNo

 $SuplNo \rightarrow SuplName, Address$ 

🕩 Bipin C Desai

91

### **Decomposition (into BCNF)**

SUPPLIER (SuplNo, SuplName)
SUPPLIER\_PARTS (SuplNo, PartNo, Quantity)

### A relation is in BCNF iff every determinant is a candidate key

BCNF addresses the situations which 3NF does not handle. In many real DB design the relations in 3NF are also in BCNF.

When is a relation in 3NF not in BCNF: it has multiple composite candidate keys, and these candidate keys are non-disjoint (at least one common attribute)

### **Example:**

Can\_Supply (SuplNo, SuplName, Address, PartNo, Price)

Can Supply is an example of a relation in 3NF but not in BCNF

Can\_Supply exhibits the above properties).

di Bipin C Desai

The following relation is in 3NF, and also in BCNF:

SUPPLIERS (SuplNo, Suplname, Address, PostalCode)

We assume that each supplier has a unique Suplname, so that SuplNo and Suplname are both candidate keys.

These candidate keys are *not* composite keys and hence the 3NF is also BCNF(all FDs the LHS is a candidate key)

### **Functional Dependencies:**

SuplNo → Address

SuplNo → PostalCode

SuplNo → SuplName

SuplName → SuplNo

SuplName → Address

SuplName  $\rightarrow$  PostalCode

#### Anomalies even in a BCNF relations:

SUPPLIERS (SuplNo, SuplName, Address, PostalCode)

**INSERT**: We cannot record the Address for a SuplNo without also knowing the SuplName

**DELETE**: If we delete the row for a given SuplName, we lose the information that the SuplNo is associated with a given Address.

**UPDATE**: Since SuplName is a candidate key (unique), there are none.

#### **Decomposition:**

SUPPLIER\_INFO (SuplNo, Address, PostalCode)
SUPPLIER NAME (SuplNo, SuplName)

🕩 Bipin C Desai

95

$$R(ABC) F=\{AB \rightarrow C, C \rightarrow B\}$$

This is in 3NF but not in BCNF.

There is no need (no way) to decompose this relation!

$$R(X,Y,Z)$$
 with  $F=\{XY \rightarrow Z, YZ \rightarrow X, XZ \rightarrow Y\}$ 

The candidate keys are: XY, YZ and XZ.

This relation is in BCNF since the determinant of each FD is a candidate key!

There is no need (no way) to decompose this relation!

### **Relational Calculus**



To be used in the spirit of copy-forward! https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

Bipin C Desai

# Propositional Logic ...

A proposition is a statement that is either true or false (but not both).

In propositional logic, we assume a collection of atomic propositions are given, e.g. p, q, r, s, t, ....

p = "COMP5531 is about databases"

q = "COMP5531 is an important course"

r = "databases is an important course"

 $\neg p =$  "COMP5531 is not about databases"

 $p \land q =$  "COMP5531 is about databases and COMP5531 is an important course"

 $p \land q \rightarrow r =$  "COMP5531 is about databases and COMP5531 is an important course then databases is an important course"

## **Propositional Logic.**

We form compound propositions by using **logical operators**: and, or, not, exclusive or, implication(if-then), biconditional(iff)).

A **tautology** is a compound proposition that always evaluates to **true**. e.g.: p v ¬p

A **contradiction** is a compound proposition that always evaluates to **false**.

A **predicate** is a property or description of subjects in the universe of discourse.

In the previous slide, predicates are italicized: is about databases, is an important course

Bipin C Desai

3

## **Propositional & Predicate Logic - Relational Calculus**

Knowing the two propositions

p = "COMP5531 is about databases"

r = "databases is an important course"

Can we say that COMP5531 is an important course?

Example of a relation as predicates: Assignment (E#, P#, H) expresses the fact that Employee E# is assigned to project P# for H hours

Its value is true if an Employee *E*# is assigned *H* hours to project *P*# *else it is false* 

In the database this is expressed by having a tuple in the table for Assignment

⇔ indicates derivable or follows in both directions

#### **Assertion of Universality**

 $\forall x : P(x) \Leftrightarrow \neg \exists x : \neg P(x)$ 

If everything is (true), there exists nothing that is not (true).

#### **Denial of Existence**

 $\forall x : \neg P(x) \Leftrightarrow \neg \exists x : P(x)$ 

If everything **is not(true)**, there exists nothing that **is(true)**.

#### **Denial of Universality**

 $\neg \forall x : P(x) \Leftrightarrow \exists x : \neg P(x)$ 

If not everything **is(true)**, there exists something that **is not(true)**.

#### **Assertion of Existence**

 $\neg \forall x : \neg P(x) \Leftrightarrow \exists x : P(x)$ 

If not everything is not(true), there exists something that

**is(true)** Bipin C Desai

5

## De Morgan's Laws

It can be shown that the following, called De Morgan's laws are equivalent:

$$P(x) ^ Q(x) \equiv \neg (\neg P(x) ^ {\lor} \neg Q(x))$$

$$P(x) \lor Q(x) \equiv \neg(\neg P(x) \land \neg Q(x))$$

A generalization of De Morgan's Law involving the  $\forall$ ,  $\exists$  quantifiers is obtained as shown in the following.

#### Assertion of Universality & Assertion of Existence

$$\forall x(P(x)) \equiv \neg(\exists x)(\neg P(x))$$
 and  $\exists x(P(x)) \equiv \neg(\forall x)(\neg P(x))$ 

In formal systems, the acceptable sentences (or formulae) are usually called **well-formed formulae** (wff).

In the wff ( $\forall x$ )(P(x) & Q(y)), where  $\forall$  is the universal quantifier (for all), x is **bound** and y is **free**.

Tuple calculus formulae are built from **atoms** of the form:

- $A_1$ .  $x \in R$  where R is a relation and x is a tuple variable.
- $A_2$ .  $x \theta y$  or  $x \theta c$  where  $\theta \in \{=, \neq, <,, >, \geq\}$ , x and y are variables and c is a constant: x, y, c are domain compatible Formulae are built from atoms using the following rules:
- **B**₁. An atom is a formula.
- $\beta_2$ . If f and g are formulae, then are:  $\neg f$ , (f),  $f \lor g$ ,  $f \land g$ ,  $f \to g$
- $B_3$ . If f(x) is a formula, where x is free, then  $\exists x(f(x))$ , and  $\forall x(f(x))$
- are also formulae; however, x is now bound.

the formula  $f \rightarrow g$ , meaning **if** f **then** g, is equivalent to  $\neg f \circ g$ .

Bipin C Desai

## Relational Calculus

**Relational calculus** is a query system wherein queries are expressed as variables and formulae on these variables. Such a formula describes the properties of the required result relation without specifying the method of evaluating it.

Tuple and domain calculi are collectively referred to as relational calculus.

Bipin C Desai

8

A query in tuple relational calculus is expressed as a formula:

 $\{t \mid P(t)\}$ 

This is the formula that finds all tuples t such that the predicate P is true.

A formula may use a constant to specify a particular value, while a variable is used as a place holder for the values in an expression or procedure.

We can also specify logical connectors such as "not" (or negation; denoted by  $\neg$ ), "or" ( $\lor$ ), "and" ( $\land$ ), and "implication" ( $\rightarrow$ ), universal (or **for all**; denoted by  $\forall$  and existential (or **for some**; denoted by  $\exists$ )

Bipin C Desai

.

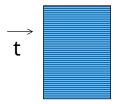
PROJECT (*Project*#, *Project\_Name*, *Chief\_Architect*)
EMPLOYEE (*Emp*#, *EmpName*)
ASSIGNED\_TO (*Project*#, *Emp*#)

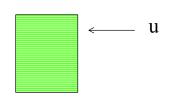
- ◆Obtain the **employee numbers** of employees working on project P1.
- ◆Obtain **employee details** for those employees assigned to project P1
- ◆Get complete details of employees working on **a** Database project.
- ◆Get complete details of employees working on **al**l Database projects.
- List the complete details of employees working on **both** P1 and P2.
- ◆List the complete details of employees working on **either** P1 **or** P2 **or both**.



◆Obtain the employee numbers of employees working on project P1.

```
\{t(Emp\#) \mid \exists u(u \in ASSIGNED\_TO \land u[Project\#]='P1' \land t[Emp\#] = u[Emp\#]) \}
```





Bipin C Desai

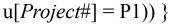
11

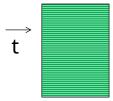
PROJECT (*Project*#, *Project\_Name*, *Chief\_Architect*)
EMPLOYEE (*Emp*#, *EmpName*)
ASSIGNED\_TO (*Project*#, *Emp*#)

◆ Obtain employee details for those employees assigned to the project P1

```
\{t \mid t \in employee \land \exists u(u \in ASSIGNED\_TO \land u(
```

$$u[Emp\#]=t[Emp\#] \land$$







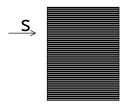
<del>-----</del>

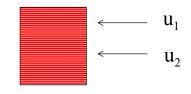
(h) Bipin C Desai

List the complete details of employees working on both P1 and P2.

```
{s | s \in employee \land \(\perp \underline{\pmathbb{I}}\) u_1, u_2 (u_1 \in \text{assigned_to} \( \land \underline{\pmathbb{U}}\) \( \land \text{u}_2 \in \text{assigned_to} \land \( \land \underline{\pmathbb{I}}\) u_1[Emp#] \( \land \underline{\pmathbb{U}}\) \( \land \underline{\pmathbb{I}}\) \( \land \underline{\pmathbb{I}}\)
```

Find s such that s is from employee and there exists tuples u<sub>1</sub>,u<sub>2</sub> both from assigned\_to such that a number of predicates are being satisfied



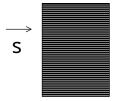


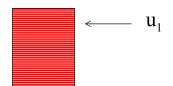
Bipin C Desai

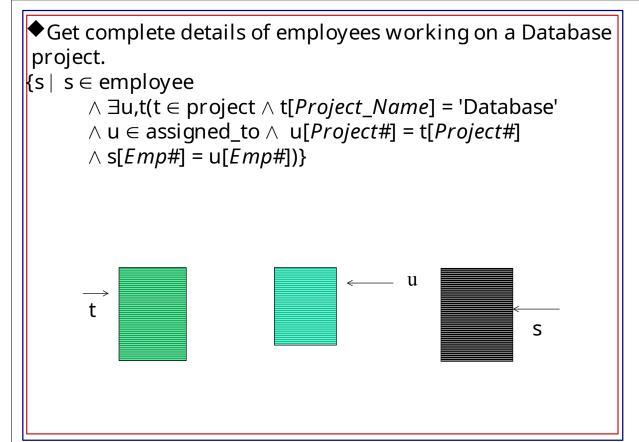
11

List the complete details of employees working on either P1 or P2 or both.

```
 \{s \mid s \in employee \land \exists u_{1,} u_{2}(u_{1} \in assigned\_to \land \\ u_{2} \in assigned\_to \land ((s[Emp\#] = u_{1}[Emp\#] \land \\ u_{1}[Project\#] = P1') \\ \lor (s[Emp\#] = u_{2}[Emp\#] \land u_{2}[Project\#] = P2'))) \}
```







14

The universe of discourse for a particular branch of mathematics is a set that contains everything of interest for that subject.

If P and Q are formulas, then "if P then Q" or "P implies Q" is written  $P\Rightarrow Q$ , using the conditional symbol,  $\Rightarrow$ .

Bi-conditional, written ⇔, corresponds to the phrase "if and only if" or "iff" for short.

The denial or negation of of  $P \Rightarrow Q$  can be expressed as:

$$\neg (P\Rightarrow Q) \Leftrightarrow \neg (\neg P \lor Q) \text{ or } \neg (P\Rightarrow Q) \Leftrightarrow \neg \neg P) \land (\neg Q)$$
  
 $\neg (P\Rightarrow Q) \Leftrightarrow P \land \neg Q$ 

```
Now we use: \forall x \neg P(x) \Leftrightarrow \neg \exists x P(x) and re-write the above predicate as: s \in \text{employee} \land \neg \exists t (\neg (P(t) \rightarrow \exists u Q(u,t)))
Substituting f \rightarrow g by its equivalent form \neg f \lor g:

We get: s \in \text{employee} \land \neg \exists t (\neg (\neg P(t) \lor \exists u Q(u,t)))

Now move the negation in and stop it just after the \lor s \in \text{employee} \land \neg \exists t (\neg (\neg P(t) \lor \exists u Q(u,t)))

s \in \text{employee} \land \neg \exists t (P(t) \land \neg \exists u Q(u,t))

\{s \mid s \in \text{employee} \land \neg \exists t (t \in \text{project} \land t [\text{Project\_Name}] = \text{'Database'} \land \neg \exists u (u \in \text{assigned\_to} \land u [\text{Project\#}] = t [\text{Project\#}] \land s [\text{Emp\#}] = u [\text{Emp\#}])\}
```

di Bipin C Desai

(b) Bipin C Desai

```
mysql> select * from Assign;
                 mysql> select * from Project;
     1 10000
                         353 | Database |
                                            10000
                         354 | Database |
                                            10000
       10002 |
                         534 | OS
                                            10005
                         574 | VOIP
                                            10005
                    4 rows in set (0.00 sec)
     | 10004
          ----+mysql> select * from Employee;
  10001 | Piere
                 | Laval
                                  5144555445 | 1956-10-12
  10002 | Nathalie | Brossard
                               | 5147454555 | 1976-04-01
 10003 | Mary | Dorval | 5145544455 | 1965-10-21
  10004 | Sabrina | St. Laurent | 5144445555 | 1987-01-31 |
                  | Montreal | 5144454555 | 1964-02-29
```

db Bipin C Desai

```
Get details of employees working on alI Database projects.  \{s \mid s \in \text{employee} \\ \land \forall \ t(\ (t \in \text{project} \land t[\textit{Project\_Name}] = '\text{Database'}) \\ \rightarrow (\exists u(u \in \text{assigned\_to} \land u[\textit{Project\#}] = t[\textit{Project\#}] \\ \land \ s[\textit{Emp\#}] = u[\textit{Emp\#}]\ )\ )\ ) \}  replacing ( f ) \rightarrow ( g ) by its equivalent form ( \negf ) ^{\vee} ( g ) :  \{s \mid s \in \text{employee} \\ \land \ \forall t(\ (t \not\in \text{project} \lor t[\textit{Project\_Name}] \neq '\text{Database'}\ ) \\ \lor (\exists u(u \in \text{assigned\_to} \land u[\textit{Project\#}] = t[\textit{Project\#}] \\ \land \ s[\textit{Emp\#}] = u[\textit{Emp\#}]\ )\ )\ )
```

The Bipin C Desai 24

```
insert Employee values(10006, 'John', 'Ndg', 5144455555, '1988-04-01');
insert Assign values (353, 10006),(354, 10006),(534, 10006),(574,10006);
mysql> select * from Employee s
where not exists (select * from Project t
             where not exists(select * from Assign u
                   where u.Pno = t.ProjNo and
                    s.EmpNo = u.Eno));
mysql> select * from Employee e
where not exists
   (select *
   from Project
   where ProjNo NOT IN
         (select distinct a.Pno
         from Assign a
         where a.Eno = e.EmpNo) );
| EmpNo | Ename | Address | Phone | DOB
+----+----+----+
| 10006 | John | Ndg | 5144455555 | 1988-04-01 |
```

```
mysql> update Employee set DOB = '1988-01-01' where
EmpNo=10006;
mysgl> update Employee set DOB = DATE ADD(DOB.
INTERVAL 365 DAY) where EmpNo=10006;
mysql> select * from Employee where EmpNo=10006;
+----+
| EmpNo | Ename | Address | Phone | DOB
+----+
| 10006 | John | Ndg | 5144455555 | 1988-12-31 | +----+
1 row in set (0.00 sec)
mysql>update Employee set DOB = DATE_SUB(DOB,
INTERVAL 365 DAY) where EmpNo=10006;
mysql> select * from Employee where EmpNo=10006;
+----+
| EmpNo | Ename | Address | Phone | DOB
+----+
| 10006 | John | Ndg | 5144455555 | 1988-01-01 | +----+
1 row in set (0.02 sec)
```

di Bipin C Desai

```
Example: Get the employee numbers of employees, other than employee 107, who work on at least all those projects that employee 107 works on"  \{ t[Emp\#] \mid t \in assigned\_to \land \\ \forall u_1(u_1 \in assigned\_to \land u_1[Emp\#] = 107 \\ \rightarrow \exists u_2(u_2 \in assigned\_to \land u_2[Emp\#] \neq 107 \\ \land u_1[Project\#] = u_2[Project\#] \land t[Emp\#] = u_2[Emp\#])) \}  Alternately we can write this query without the logical implication by substituting its equivalent form \neg f \lor g:  \{ t[Emp\#] \mid t \in assigned\_to \land \\ \forall u_1(u_1 \not\in assigned\_to \land u_1[Emp\#] \neq 107 \\ \lor \exists u_2(u_2 \in assigned\_to \land u_2[Emp\#] \neq 107 \\ \land u_1[Project\#] = u_2[Project\#] \land t[Emp\#] = u_2[Emp\#])) \}
```

To avoid procedural operation, such as projection, in a calculus query, we could define t to be on the relation scheme (*Emp#*) and rewrite this query expression as:

```
{ t(Emp#) |
∀u<sub>1</sub>(u<sub>1</sub> ∉ assigned_to ∨ u<sub>1</sub>[Emp#] ≠ 107
∨ ∃u<sub>2</sub>(u<sub>2</sub> ∈ assigned_to ∧ u<sub>2</sub>[Emp#] ≠ 107
∧ u<sub>1</sub>[Project#] = u<sub>2</sub>[Project#]∧ t[Emp#] = u<sub>2</sub>[Emp#]))}
```

**Example**: Get employee numbers of employees who do not work on project P2.

```
{ t[Emp#] | t ∈ assigned_to ∧
¬∃u(u ∈ assigned_to ∧ u[Project#] = P2
∧ t[Emp#] = u[Emp#])}
```

Alternatively, we can express this query in the following equivalent form:

```
{ t[Emp\#] | t \in assigned\_to \land \forall u(u \notin assigned\_to \lor t[Emp\#] ≠ u[Emp\#] \lor u[Project\#] ≠ P2)}
```

Bipin C Desai

**Example:** Compile a list of employee numbers of employees who work on all projects.

```
{ t[Emp\#] | t ∈ assigned_to \land

\forallp(p ∈ PROJECT \rightarrow \existsu( u ∈ assigned_to

\land p[Project\#] = u[Project\#]

\land t[Emp\#] = u[Emp\#]))}
```

The above can be rewritten as:

```
{ t[Emp\#] \mid t \in assigned\_to \land \forall p(p \notin PROJECT \lor \exists u(u \in assigned\_to \land p[Project\#] = u[Project\#] \land t[Emp\#] = u[Emp\#]))}
```

**Example:** Get employee numbers of employees, not including employee 107, who work on at least one project that employee 107 works on.

```
{ t[Emp\#] | t ∈ assigned_to \land

∃s,u (s ∈ assigned_to \land u ∈ assigned_to

\land s[Project\#] = u[Project\#]

\land s[Emp\#] = 107

\land t[Emp\#] = u[Emp\#])}
```

Bipin C Desai

21

Consider the division operation on the two relations,  $P(\mathbf{P})$  and  $Q(\mathbf{Q})$ , where  $\mathbf{Q} \subseteq \mathbf{P}$ :

$$\begin{split} R &= P \div Q \\ R &= \{t \mid t \in P[\textbf{P} \cdot \textbf{Q}] \ \land \ \forall s (s \in Q \land (t \cap s \in P)) \} \\ R &= \{t \mid t \in P[\textbf{P} \cdot \textbf{Q}] \ \land \ \forall s (s \in Q \rightarrow \exists u (u \in P \land u[\textbf{Q}] = s \land u[\textbf{P} \cdot \textbf{Q}] = t[\textbf{P} \cdot \textbf{Q}])) \} \\ R &= P \div Q = \Pi_{\textbf{P} \cdot \textbf{Q}}(P) - \Pi_{\textbf{P} \cdot \textbf{Q}}(\Pi_{\textbf{P} \cdot \textbf{Q}}(P) \times Q) - P) \end{split}$$

db Bipin C Desai

33

### A domain calculus expression is of the form

$$\{X \mid f(X)\}$$

where f is a formula on X, and X represents a set of domain variables.

$$A_1$$
.  $X \in \mathbf{R}$ 

$$A_2$$
.  $x \theta y \text{ or } x \theta c$ 

where  $\theta$  is one of the comparison operators x and y are domain compatible variables, and c is a domain compatible constant.

 $B_1$ . An atom is a formula.

 $B_2$ . If f and g are formulae, then  $\neg f$ , (f),  $f \land g$ ,  $f \land g$ ,  $f \rightarrow g$  are also formulae.

 $B_3$ . If f(X) is a formula where X is free, then  $\exists X(f(X))$ , and  $\forall X(f(X))$  are also formulae

db Bipin C Desai 34

PROJECT(**Project**#,Project\_Name,Chief\_Architect)
EMPLOYEE(**Emp**#, EmpName)
ASSIGNED\_TO (**Project**#, **Emp**#)

Get employee numbers for employees working on project number P1

 $\{e \mid \exists p \ (\langle e, p \rangle \in assigned\_to \land p = P1) \}$ 

In this can, we can drop the quantifier and simplify

the query as:

 $\{e \mid \langle e, p \rangle \in assigned\_to \land p = P1\}$ 

Bipin C Desai

25

Get employee details such that the employee is assigned to the project P1

$$\{\langle e_1, m \rangle \mid \exists e_2 \ (\langle p, e_2 \rangle \in \text{assigned\_to} \\ \land \langle e_1, m \rangle \in \text{employee}) \land p = P1 \land e_1 = e_2\}$$

Compile the details of employees working on a Database project.

{e,m 
$$\mid \exists p_1, e_1, p_2, n_2 (< p_1, e_1 > \in assigned\_to$$
  
  $\land < e, m > \in employee$   
  $\land < p_2, n_2, c_2 > \in project$   
  $\land e_1 = e \land p_1 = p_2 \land n_2 = Database)}$ 

de Bipin C Desai

Compile the details of employees working on both P1 and P2.

(b) Bipin C Desai

Obtain the employee numbers of employees, other than the employee 107, who work on at least all those projects that employee 107 works on.

{e |  
$$\in$$
 assigned\_to  $\forall$  p<sub>1</sub>,e<sub>1</sub>(  
 1,e<sub>1</sub>>  $\in$  assigned\_to  $\land$  e<sub>1</sub> =107  
  $\rightarrow$  ( $\exists$ p<sub>2</sub>,e<sub>2</sub>( 2,e<sub>2</sub>>  $\in$  assigned\_to  
  $\land$  e<sub>2</sub>  $\neq$  107  $\land$  p<sub>1</sub> = p<sub>2</sub> $\land$  e = e<sub>2</sub>))}

An equivalent form of this query

{e |  ∈ assigned\_to ∧  

$$\forall p_1,e_1(  
 $\lor (\exists p_2,e_2(  
 $\land e_2 \neq 107 \land p_1 = p_2 \land e = e_2))$ }$$$

di Bipin C Desai

Get employee numbers of employees who do not work on the P2 project.

{e | 
$$\exists p \ ( \in assigned\_to$$
  
  $\land \forall p_1,e_1 \ ( \notin assigned\_to$   
  $\lor p_1 \neq P2 \lor e_1 \neq e))}$ 

What are the employee numbers of employees who work on all projects?"

{e | 
$$\exists p \ ( < p,e > \in assigned\_to$$
  
  $\land \forall p_1(< p_1,n_1,c_1 > \in project$   
  $\rightarrow < p_1,e > \in assigned\_to ))}$ 

Bipin C Desai

Acquire the employee numbers of employees, other than employee 107, who work on at least one project that employee 107 works on.

{e | 
$$\exists$$
 p,p<sub>1</sub>,e<sub>1</sub>,p<sub>2</sub>,e<sub>2</sub>(  $\in$  assigned\_to  
  $\land$  1,e<sub>1</sub>>  $\in$  assigned\_to  
  $\land$  2,e<sub>2</sub>>  $\in$  assigned\_to  
  $\land$  e<sub>2</sub>  $\neq$  107  $\land$  p<sub>1</sub> = p<sub>2</sub>  $\land$  e<sub>1</sub> = 107  $\land$  e = e<sub>2</sub>)}

The following is covered in predicate logic discussions: Here Q(x) is any predicate of variable x and  $\equiv$  means Logically equivalent

Negating a proposition such as  $\forall x Q(x)$  requires negating the predicate and changing the quantifier from the Universal to the existential

Thus we can replace a negated universal quantifier with a predicate Q(x) as follows to its logically equivalent form

$$\neg [\forall x Q(x)] \equiv \exists x [\neg Q(x)]$$
 or  $\exists x [\neg Q(x)] \equiv \neg [\forall x Q(x)]$ 

Similarly, we can negate a proposition with an existential quantifier as follows:

$$\neg [\exists x Q(x)] \equiv \forall x [\neg Q(x)]$$
 or  $\forall x [\neg Q(x)] \equiv \neg [\exists x Q(x)]$ 

Bipin C Desai

1

Using the last formula from the previous slide :

$$\neg[\exists x Q(x)] \equiv \forall x[\neg Q(x)]$$

If we reverse the sides and move the negation inside:

$$\forall x(\neg Q(x)) \equiv \neg(\exists x)(\neg Q(x))$$

Now if we substitute, in the above:

$$P(x)$$
 for  $\neg Q(x)$  and  $\neg P(x)$  for  $\neg \neg Q(x)$  i.e.,  $\neg P(x)$  for  $Q(x)$ 

We get:

$$\forall x(P(x)) \equiv \neg[(\exists x)(\neg Q(x))]$$
 or

$$\forall x(P(x)) \equiv \neg[(\exists x) (\neg \neg P(x))] \text{ or }$$

$$\forall x(P(x)) \equiv \neg[(\exists x) (P(x))]$$

It is the last form that we have used!!

```
Get details of employees working on all Database projects.
\{s \mid s \in employee \land
    \neg\exists t((t \in project \land
            \wedge t[Project_Name] = 'Database'
                 \land \neg \exists u (u \in assigned\_to \land \exists u (u \in assigne
         u[Project\#] = t[Project\#] \land s[Emp\#] = u[Emp\#]))
select distinct a.empno
                 from
                                                                                   assigned_to a, project p
                 where
                                                                                          not exists
                                                                 (select *
                                                                    from project p
                                                                     where p.projname = 'database' and
                                                                                           not exists
                                                                                               (select *
                                                                                                  from assigned_to a1
                                                                                                  where a1.projno = p.projno and
                                                                                                                          a1.empno = a.empno));
```

Get details of employees working on **al**l Database projects and **only** on database projects.

```
\{s \mid s \in employee \land \neg \exists t((t \in project \land \land t[Project\_Name] = 'Database' \land \neg \exists u(u \in assigned\_to \land u[Project\#] = t[Project\#] \land s[Emp\#] = u[Emp\#]))) \land (\neg \exists t1(t1 \in project \land t1[Project\_Name] \neq 'Database' \land \exists u1(u1 \in assigned\_to \land u1[Project\#] = t1[Project\#] \land s[Emp\#] = u1[Emp\#])))\}
```

```
select distinct a.empno
          assigned_to a, project p
  where
           not exists
        (select *
        from project p
        where p.projname = 'database' and
           not exists
           (select *
            from assigned_to a1
            where a1.projno = p.projno and
               a1.empno = a.empno)
      and not exists (select *
               from project p1
               where p1.projname <> 'database' and
           exists (select *
            from assigned_to a2
            where a2.projno = p1.projno and
               a2.empno = a.empno));
```

# Evolving database systems

MariaDB [mysql]> show engines;

Engine	Support	·
MRG_MyISAM MyISAM MEMORY	YES   YES	Collection of identical MyISAM tables   MyISAM storage engine   Hash based, stored in memory,   useful for temporary tables
CSV Aria InnoDB  SEQUENCE PERFORMANCE_SCHEMA	YES DEFAULT YES	CSV storage engine   Crash-safe tables with MyISAM heritage   Percona-XtraDB, Supports transactions, row-level   locking, foreign keys and encryption for tables   Generated tables filled with sequential values   Performance Schema

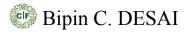
Bipin C Desai

47

Engine	Support	Comment			
 Aria	YES	Crash-safe tables with MyISAM heritage. Used for internal temporary tables and privilege table:			
MRG_MyISAM	YES	Collection of identical MyISAM tables			
MEMORY	YES	Hash based, stored in memory, useful for temporary tables			
BLACKHOLE	YES	/dev/null storage engine (anything you write to it disappears)			
MyISAM	YES	Non-transactional engine with good performance and small data footprint			
CSV	YES	Stores tables as CSV files			
ARCHIVE	YES	gzip-compresses tables for a low storage footprint			
FEDERATED	YES	Allows one to access tables on other MariaDB servers, supports transactions and more			
PERFORMANCE_SCHEMA	YES	Performance Schema			
InnoDB	DEFAULT	Supports transactions, row-level locking, foreign keys and encryption for tables			
SEQUENCE	YES	Generated tables filled with sequential values			
rows in set (0.015	sec)				

di Bipin C Desai

## SQL - II



Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

Bipin C Desai

This set of slides is extensive with many examples using Oracle and MariaDB/MySQL.

These example should be tried out for better understanding.

We will not go through them in details: this is left as exercises!

Here are the main SQL topics – their syntax etc.- DBMS dependent! Data Types

SQL statements

Group/having

**Functions** 

Views, Null and operations on null

Joins – cross. inner, natural, outer,

Implementation of joins

Representing relationships, key constraints

Complex relationships

Constraints, triggers (row/statement, before/after);

mutating triggers

Aggregation

#### **DATA TYPES**

-The prompt for MariaDB/MySQL is followed by database name

Bipin C Desai

3

Oracle also uses **varchar2**(n); it's truly varying length **varchar2** is not yet supported in MySQL/Mariadb!

Times:SQL2 form is TIME 'hh:mm:ss[.ss...]'

Dates: SQL2 format is DATE 'yyyy-mm-dd' (m =0 or 1) supported in MySQL/Mariadb
Oracle's default dates format is 'dd-mon-yy'
Example: create table BDays(name char(25), d DATE);
insert into BDays values ("Martha Smith", '18-nov-1962');
Oracle function to\_date converts a date value into default format, e.g.,
insert into BDays values("Martha Smith",
to date('1962-nov-18', 'yyyy-mm-dd'));

SQL is case insensitive

However, the case is significant for *strings* and could be for names of tables and columns

Two strings  $s_1$  and  $s_2$  are equal if

- they have the same sequence of characters and

- the same case Note: some DBMS are

The strings are compared alphabetically case sensitive for the

'fodder' < 'folder' names of tables and

'bat' < 'batman' columns: all DBMS are

string LIKE pattern

Ordinary characters in *pattern* match only ordinary characters in *string* 

case sensitive for strings

The special character % in *pattern*, matches any sequence of zero or more characters in *string* 

The special character \_ in *pattern*, matches any **one** character in *string* 

Bipin C

\_\_\_

#### Find all students with "de" in their name

select Name, Dept

from students

where Name like '%de" %';

Note: an apostrophe in a string represented by two apostrophes "without any intervening spaces

Expressing special characters \_ and % in a string by an using a preceding escape character.

**SQL** allows any character to be used as an escape character with the *escape* keyword

string LIKE 'x\_%x%' ESCAPE 'x'

Here  $\mathbf{x}$  is the escape character

The sequence  $x_a$  and  $x_b$  is taken to be a single a and b

This pattern matches any string that begins with \_ and ends with %

We can apply the 3 most common set operations union, intersection, except(difference) to relations **R** and **S**, provided the relations are *compatible*.

If two SQL queries produce compatible relations as their result, then we may combine these queries using: **union**, **intersection**, **except** 

The SQL implementation of **union**, **intersection**, and **except** operation normally eliminate duplicates; the modifier ALL is used to keep the duplicates:

```
R \text{ UNION ALL } S, |t \in R| = n, |t \in S| = m, |t \in R \cup S| \le n + m

R \text{ INTERSECT ALL } S, |t \in R| = n, |t \in S| = m, |t \in R \cap S| = min(n,m)

R \text{ EXCEPT ALL } S, |t \in R| = n, |t \in S| = m, |t \in R - S| = max(0,n-m)
```

NOTE: in many of the SQL examples, to save space on slides, we are not including constraints such a primary key which are usually evident

Bipin C

7

#### Relation schemas:

```
Faculty (Name, Dept, Position, salary, gender)
Student (Name, Dept, Major, sex)
```

#### Query:

Give the names and the departments of students and professors.

```
(SELECT Name, Dept FROM Faculty)

UNION

(SELECT Name, Dept FROM Student);
```

Bipin C

8

```
create table Faculty ( Name varchar2(30),
                                              create table Student (
                                               Name varchar2(28),
Dept varchar2(20), Position varchar2(20),
                                               Dept varchar2(20),
salary dec(9,2), gender char(1));
                                               Major varchar2(20),
                                                sex char(1));
insert into faculty values ('Smith', 'CS', 'Prof', 81000.00, 'F');
insert into faculty values ('Brown', 'CS', 'Assoc Prof', 75000.00, 'M');
insert into student values ('Brown', 'CS', 'Info', 'F');
                                    (SELECT Name, Dept FROM Faculty)
(SELECT Name, Dept FROM Faculty)
                                            UNION ALL
        UNION
                                    (SELECT Name, Dept FROM Student);
(SELECT Name, Dept FROM Student);
                                                       DEPT
                                    NAME
NAME
               DEPT
                                    Smith
                                                       CS
                CS
Brown
                                    Brown
                                                        CS
Smith
                CS
                                                        CS
                                    Brown
```

Bipin C

g

```
Relation schemas:
   TA (Name, stipend, course)
   Student (Name, Dept, Major, sex)
Query: Find names of female TAs who are majoring in
  the department of Computer Science
(SELECT name
                       Find names of TAs who are not majoring
FROM TA)
                       in the department of Computer Science.
         INTERSECT
(SELECT name
                       (SELECT name
FROM Student
WHERE Dept="Computer Science"

FROM TA)
                                         EXCEPT
         and sex = "F");
                           (SELECT Name
                           FROM Student
                           WHERE Dept = "Computer Science");
```

Bipin C

```
A tuple in SQL is represented by a parenthesized list of scalar values;
  (Smith, 'CompSci') or (Smith, Student.Dept)
If a tuple t has the same number of components as a table (relation) R,
  then it makes sense to compare t and R
t \text{ IN } R -- this is true iff t \in R
t \Leftrightarrow ANY R -- true if t is neither greater nor less than any tuple in R
Relation schemas: Student (Name, Dept, Major, sex),
TA (Name, stipend, course), GRADE(Name, course, gr)
Find students who got an A in the course where they are TAs
   select t.Name
                                               We are conveniently
  from TA t
                                               ignoring time!
  where (t.name, t.course) in
                                  (select Name, course
                                   from grade
                                   where gr = 'A');
```

Bipin C

11

#### INSERT, DELETE, ALTER

The form of a insert statement:

insert into relation[list of attributes] values(list of values); insert into relation *select statement*;

The form of a delete statement:

delete from relation where <condition>;

Delete every tuple in the relation satisfying the condition

The form of an update statement:

update relation set <new-value assignments> where <condition>;

Adding Columns

alter table relation add <column declaration>;

Removing Columns

Add & Remove with care!
alter table relation DROP <column name>;

select * from EMPL;									
	Eid	Name		title		salary		emailsuffix	
+-	+-		+-		+-		+-	+	
	33	John		SrProgr		120000		coldmail.org	
	34	Jenny		SrProgr		110000		coldmail.org	
	35	Anne		WebDev		90000		gonemail.com	
	36	Mary		WebDev		85000		comemail.com	
	37	Freddy		Progr		75000		netmail.com	
	38	Johnny		Progr		80000		netmail.com	
	39	Art		Progr		75000		netmail.com	
	40	Albert		Progr		70000		netmail.com	
	41	Susan		WebProgr		90000		gonemail.com	
	42	Paul		WebProgr		85000		gonemail.com	
	43	Edward		DBProgr		75000		coldmail.org	
	44	Kim		WebDev		110000		coldmail.org	
	45	Roger		DBA		150000		comemail.com	
	46	Danny		NetAdmin		100000		sizzlingmail.com	
	47	Mike		Mkt Mgr		120000		gonemail.com	
+ =	48	MaryAnne 	  -	Mkt Mg 	 ·+-	90000	 -+-	speedymail.com	

13

# Group by and Having

```
The Group by clause is to group the data by the value(s) of one (or more) column(s)
```

The predicate for the GROUP BY clause is HAVING

```
CREATE TABLE EMPL (
Eid int unsigned not null auto_increment primary key,
Name varchar(20),
title varchar(30),
salary int,
emailsuffix varchar(60)
);
```

ALTER TABLE EMPL AUTO\_INCREMENT = 100;

(b) Bipin C Desai

15

```
select title, emailsuffix
from EMPL
GROUP BY title, emailsuffix;
  title
            | emailsuffix
  DBA
          | comemail.com
| DBProgr | coldmail.org
| Mkt Mg | speedymail.com
| Mkt Mgr | gonemail.com
| NetAdmin | sizzlingmail.com |
| Progr | netmail.com
| SrProgr | coldmail.org
| WebDev | coldmail.org
| WebDev | comemail.com
| WebDev | gonemail.com
| WebProgr | gonemail.com
11 rows in set (0.00 sec)
```

17

```
select title, emailsuffix from EMPL GROUP BY title, emailsuffix having count(title)>2;
```

```
+----+
| title | emailsuffix |
+----+
| Progr | netmail.com |
+----+
```

select title, emailsuffix from EMPL GROUP BY title, emailsuffix having count(title)>=2 ORDER BY title;

Bipin C Desai

19

select title, salary from EMPL where emailsuffix like '%org' GROUP BY title, emailsuffix having count(title)>=2 ORDER BY salary;

# Merge rows

Using the select statement to merge multiple rows into 1 row: MySQL: the **group concat** notation".

mysql> select C, group\_concat(B) as Bs

- -> from R
- -> group by C;

```
C | Bs |
 ----+
    12 | 10,11 | | a1 | 10 | 12 | 14 | 9 | | a2 | 11 | 12 | 17 | 0 | | a3 | 9 | 14 |
   14 | 9 |
    17 I 8
+----+
3 rows in set (0.03 sec)
```

```
mysql> select * from R;
 +---+
 | A | B | C
| a4 | 8 | 17 |
```

Bipin C Desai

# Merge rows

Using the select statement to merge multiple rows into 1 row:

MySQL: the **group concat** notation".

mysql> select C, group\_concat(distinct B separator ';' ) as Bs

-> from R

```
-> group by C;
                      mysql> select * from R;
                     | A | B | C
 C | Bs
 ----+
                     | a1 | 10 | 12 |
   12 | 10;11 |
                             11 | 12 |
                     | a2 |
   14 | 9 |
                             9 | 14 |
                      | a3 |
   17 | 8
                             8 | 17 |
                      | a4 |
                     | a5 | 10 |
3 rows in set (0.03 sec) two 10+ are not repeated ---+
```

```
Merge rows: Oracle
SELECT C,
listagg(B, ', ') WITHIN GROUP (ORDER BY B) AS Bs
FROM R
GROUP BY C;
                           select * from R;
C
     Bs
   10, 11
12
14
                          | a1 | 10 | 12 |
     9
17
                            a2 | 11 |
                                            12
3 rows returned in 0.07 seconds | a3 | 9 |
                                           14 |
                                           17 |
                           | a4 |
```

23

```
Update part of text in a column
```

Handy to update part of an existing text column in a table!

select message from account\_email where message like '%confsys%'; 11 rows in set (0.001 sec)

```
update account_email set message =replace (message ,'confsys.encs', 'ideas.encs'); Query OK, 11 rows affected (0.010 sec)
```

select message from account\_email where message like '%confsys%'; Empty set (0.001 sec)

#### Revert:

update account email

-> set message =replace (message ,'ideas.encs','confsys.encs'); Query OK, 11 rows affected (0.009 sec)

#### **Functions**

Most database systems have a multitude of functions:

- Comparison Functions and Operators
- Logical Operators
- Control Flow Functions
- String Functions
- Mathematical Functions
- Date and Time Functions
- Encryption and Compression Functions
- Bit Functions
- Full-Text Search Functions
- Cast functions and Operators
- Information Functions
- XML Functions

Bipin C Desai

25

# Regular expression

REGEXP is used to give a pattern scheme for a string comparison of the pattern with a string using the syntax:

expr REGEXP pat

If there is a match the REGEXP function returns 1, else 0.

If either expression or pattern is NULL, the function returns NULL.

Some meta-characters: ^, \* , . , [ ] , ( ) , {m,n}

- Match the beginning of a string.
- \$ Match the end of a string.
- . Match any character

Suggestion: Look up manual/tutorials and try examples.

```
Date format for MySQL is YYYY-MM-DD - the SQL2 default
To set Oracle's default date format to YYYY-MM-DD
Internally Oracle stores both date and time as a single value
$conn = OCILogon($my Ora id,$My-Ora PW,$My Ora db)
//Set Oracle's date format to YYYY-MM-DD
$stmt = OCIParse($conn,"ALTER SESSION SET
          NLS DATE FORMAT='YYYY-MM-DD'");
OCIExecute($stmt,OCI DEFAULT);
create table bdate( Name char(25), bday date);
insert into bdate values('Jane', '20-Jan-83');
select * from bdate:
NAME
                 BDAY
Jane
                 20-JAN-83
MariaDB>insert into bdate values('Jane', '1983-01-20');
```

```
ALTER SESSION SET NLS DATE FORMAT='YYYY-MM-DD';
SQL> select * from bdate;
NAME
                      BDAY
                     1983-01-20
Jane
SELECT Name, TO CHAR(bday, 'YYYY/MM/DD') AS Birthday
FROM bdate;
NAME
                 BIRTHDAY
          1983/01/20
Jane
In Oracle:
The functions TO CHAR or TO DATE return part of the
date/time.
TRUNC will return the first day of the period. ROUND will
round up at mid year/mid month (July 1 or 16th day)
```

de Bipin C Desai

CREATE TABLE supplies (supname char(14), part# number(4), price number(7,2));

CREATE TABLE project (proj# number(4), projname char(14), projmgr number(4));

CREATE TABLE usedin (proj# number(4), part# number(4), qty number(3)); CREATE TABLE empls (emp# number(4), empname char(14), address char(14));

Bipin C Desai

2

SQL> select *	from empls;		SQL> select * f	rom assi	igned_to;
EMP#	EMPNAME	ADDRESS	PROJ#	EMP#	HOURS
120	Hardrock	Outremont	353		-
II	Eliza		753		
			353		-
II.		Laval	353 451		40 40
127	Jim	Montreal	753		40
129	Sun	Brossard	353		
131	Moon	Beaconsfield	451		10
			321	120	40
II .	Dr. Dolittle		326	142	40
141	Knowit	Montreal	326		
142	Softee	NDG	451	135	1
143	Dr. Knowall	Montreal			
SQL> sele	ct * from proj	ect;	EMP# EMPNAME	1	ADDRESS
PROJ	# PROJNAME	PROJMGR	135 Dr. D	olittle	Laval
35	3 database	135			
45	1 database	141			
32	1 software	120			
32	6 hardware	142			
75	3 database	135			

USEDIN COMP321 COMP321	1 2	5 2	SUPPLIES SUPNAME	PART# 	PRICE
COMP321 COMP326 COMP326 COMP326 COMP353 COMP353 COMP451 COMP451 COMP753 COMP753 COMP753 COMP753 COMP753 COMP753	3 4 5 6 1 8 9 7 1 2 3 4	3 1 3 4 5 1 2 3 4 3 6 4	PROVIBEC PROVIBEC PROVIBEC PROVIBEC SUPORIO SUPBEC NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY NDG-SUPPLY MDG-SUPPLY MDG-SUPPLY SUPBEC MANIBEC MDG-SUPPLY	1 2 3 4 1 2 1 2 1 2 3 4 7 1 1	710.2 815.3 325 795.99 695.99 799.98 699.99 799.99 324.99 795.98 754 699.98 727.99 699.99
SUPORIO MANIPART SUPBEC NDG-SUPPLY		Toronto Winnipeg Laval NDG	MDG-SUPPLY MDG-SUPPLY MDG-SUPPLY	2 3 4	799.99 324.99 795.98

31

```
PROJECT(Project\#, Project\_Name, Chief\_Architect)
EMPLOYEE (Emp\#, EmpName)
ASSIGNED_TO (Project\#, Emp\#)

• Get details of employees working on all Database projects.

{s \mid s \in employee

\land \forall t(t \in project \land t[Project\_Name] = 'Database'

\rightarrow \exists u(u \in assigned\_to \land u[Project\#] = t[Project\#]

\land s[Emp\#] = u[Emp\#])}

replacing f \rightarrow g by its equivalent form \neg f \lor g:

{s \mid s \in employee

\land \forall t(t \not\in project \lor t[Project\_Name] \neq 'Database'

\lor \exists u(u \in assigned\_to \land u[Project\#] = t[Project\#]

\land s[Emp\#] = u[Emp\#])}
```

(h) Bipin C Desai

```
Using \forall x(P(x)) \equiv \neg(\exists x)(\neg P(x)) Assertion of Universality \{s \mid s \in \text{employee} \\ \land \neg(\exists t)(\neg (t \notin \text{project} \lor t[Project\_Name] \neq '\text{Database'} \\ \lor \exists u(u \in \text{assigned\_to} \land u[Project\#] = t[Project\#] \\ \land s[Emp\#] = u[Emp\#]))\}
\{s \mid s \in \text{employee} \\ \land \neg(\exists t) (t \in \text{project} \land t[Project\_Name] = '\text{Database'} \\ \neg(^{\lor} \exists u(u \in \text{assigned\_to} \land u[Project\#] = t[Project\#] \\ \land s[Emp\#] = u[Emp\#]))\}
\{s \mid s \in \text{employee} \\ \land \neg(\exists t) (t \in \text{project} \land t[Project\_Name] = '\text{Database'} \\ \land \neg \exists u(u \in \text{assigned\_to} \land u[Project\#] = t[Project\#] \\ \land s[Emp\#] = u[Emp\#]))\}
```

### ALTERNATE SCHEME: Using set operation for answering "ALL" type queries select \* from empls s Set of all projects where not exists( (select Proj# with name = 'database' from project t where t.ProjName = 'database') minus Set of projects with (select u.Proj# name='database' from assigned to u, project t1 assigned to where u.Proj# = t1.Proj#employee s and s.EMP# = u.Emp# and t1.ProjName = 'database'));

(d) Bipin C Desai

### **USE OF VIEW**

A view is a materialized(virtual) table that can be used in any SQL query

Create a view (project numbers) of the projects managed by an employee with the name 'Dr. Dolittle'

Optional renaming of attributes used in view definitions

## **Using VIEW**

Find suppliers who can supply all parts used in a project managed by Dr. Dolittle, and the corresponding project number(s)

```
 \{S \mid s \in \text{supplies } ^{ } d \in \text{do\_project} \\  \  ^{ } S[Supname] = s [Supname] \\  \  ^{ } S[Project\#] = d [Projectnumb] \\  \  ^{ } \neg (\exists u) (u \in \text{used\_in } ^{ } u[Project\#] = d[Projectnumb] \\  \  ^{ } \neg \exists t (t \in \text{supplies } ^{ } u[Part\#] = t[Part\#] \\  \  ^{ } t[Supname] = s[Supname] )) \}
```

(d) Bipin C Desai

## A view can appear where a relation name is allowed Find suppliers who can supply all parts used in a project managed by Dr. Dolittle

There exists a supplier s and select unique s.supname, d.proj# project d such that there are no from supplies s, do\_project d parts used in d that is not where exists Is this predicate required?? supplied by this supplier s

```
(select *
                                             PROJ#
                              SUPNAME
from supplies s1
where s1.supname 

supname 

MDG-SUPPLY
                                             753
and not exists
                                             753
                             NDG-SUPPLY
   (select *
                             PROVIBEC
                                             753
   from usedin u
   where u.proj# = d.proj#
     and not exists
            (select *
            from supplies s2
            where s2.supname = s1.supname
                         and s2.part#=u.part#)))
```

A view can appear where a relation name is allowed Find suppliers who can supply all parts used in a project managed by Dr. Dolittle

```
select unique s.supname, d.proj# There exists a supplier s and
                                     project d such that there are no
from supplies s, do_project d
                                     parts used in d that is not
where not exists
                                     supplied by this supplier s
 (select *
                                                        PROJ#
  from usedin u
  where u.proj# = d.proj#
                                       MDG-SUPPLY
  and not exists
                                                        753
                                       NDG-SUPPLY
    (select *
                                                        753
                                       PROVIBEC
      from supplies s2
      where s2.supname = s.supname
      and s2.part#=u.part#))
```

Bipin C Desai

#### An Alternative

Using the view do\_project and set difference operations to write the SQL query for:

Find suppliers who can supply all parts used in a project managed by Dr. Dolittle

	SUPNAME	PROJ#
select unique s.supname, d.proj# from supplies s, do_project d where not exists (select u.part#	MDG-SUPPLY NDG-SUPPLY PROVIBEC	753 753 753
from usedin u where u.proj# = d.proj#	Parts used in one of Dolittle's project	
minus (select s2.part# from supplies s2 where s2.supname =	Parts suppli = s.supname))	ed by s

d Bipin C Desai

39

# Find suppliers who can supply all parts used in a project managed by Dr. Dolittle

```
select unique s.supname, d.proj#

from supplies s, do_project d

where not exists

(select s2.part#

from supplies s2

where s2.supname = s.supname

minus

(select u.part#

from usedin u

where u.proj# = d.proj# )))

Parts used in one of

Dolittle's project
```

What is wrong with this query??

If a supplier supplies all parts used in the project but also other parts than this supplier would not be included.

Bipin C Desai

41

How to update a view?

Translate modification of the view to the corresponding modification on the base tables used in the view definition—be able to identify the base relation(s) and attribute(s)

Should we allow updates on views?

Yes, however it depends - some problems can arise

Some simple views can be updated

Known as **updatable views (easy if primary keys are part of view)** 

Many views cannot be updated

This is due to the so called view-update anomaly

insert into do\_project values('Proj1'); ⇒(Proj1, null, null)

Note: null should be allowed for the base attributes

Would the insertion cause the insertion of (Proj1, null, emp# of Dr. Do..)?

SQL provides a formal definition of when modifications to a view are permitted

- it is permitted if the view is defined by selecting some attributes from one relation **R**, which could be an "updatable" view itself
- the view definition uses **SELECT** (but not **SELECT DISTINCT**)
- the WHERE clause does not involve R in a sub query
- the list in the **SELECT** clause includes "enough" attributes that for every tuple inserted into the view, the tuple inserted into the base relation will "yield" the inserted tuple of the view
- the **NOT NULL** constraints on the base relation will not be violated

Bipin C

4

SQL allows user defined data types - **domains** 

We can define a domain as follows:

create domain <name> as <type description>default value;

To create a domain with default value:

create domain Projnumbers as number(4) default 9999;

To change the default for a domain:

alter domain Projnumbers set default 0;

To delete a domain definition:

drop domain Projnumbers;

# Arithmetic operations on NULLs

Result of an arithmetic operator, when at least one of the operands has a value of NULL, is NULL

if x have the value NULL, then x+3 is also NULL

However, NULL is not a constant

NULL + 3 is *illegal* 

Some basic arithmetic rules are not applicable.

Suppose x is a numeric value

$$x * 0 = 0$$
, but if x is NULL then  $x * 0$  is NULL

$$x - x = 0$$
, but if x is NULL then  $x - x$  is NULL

© Bipin C 45

## In 3-Valued Logic we may assume that:

**TRUE** = 1, **FALSE** = 
$$0$$
, **UNKNOWN** =  $1/2$ 

$$x$$
 **AND**  $y = \min(x, y)$ ,  $x$  **OR**  $y = \max(x, y)$ , **NOT**  $x = 1-x$ 

X	Y	XAND Y	X OR Y	NOT X
TRUE	TRUE	TRUE	TRUE	FALSE
TRUE	UNKNOWN	UNKNOWN	TRUE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE
UNKNOWN	TRUE	UNKNOWN	TRUE	UNKNOWN
UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
UNKNOWN	FALSE	FALSE	UNKNOWN	UNKNOWN
FALSE	TRUE	FALSE	TRUE	TRUE
FALSE	UNKNOWN	FALSE	UNKNOWN	TRUE
FALSE	FALSE	FALSE	FALSE	TRUE

```
Null value and logical operations

Two value logic: x OR (NOT x) = 0 OR 1| 1 OR 0 = 1 = TRUE

For 3-valued logic:
x OR (NOT x) = max(1/2,(1-1/2)) = 1/2 = UNKNOWN ≠ (TRUE)

Note: We can't treat NULL as a constant:
grade ( Name, course, gr)

Consider query:
select *
from grade
WHERE gr <= "c" or gr > "c";

Here the result is expected to be the grade relation.
```

If null values are allowed for gr, then the above query returns only tuples of **grade** where the value of gr is not NULL.

Bipin C

```
In SQL2, there are other forms for expressing \times and \bowtie
Cartesian Product of Employee and Position
Employee (Empl No, Name, Skill, Pay rate) Position (Posting No, Skill)
  select * from Employee, Position;
  select * from Employee CROSS JOIN position;
Theta join of Employee, Position
                                                  select *
   select * from Employee JOIN Position ON
                                                   from Employees
                     Position. Skill = Employee. Skill; JOIN Position
                                                   USING (skill);
   select distinct * from Employee JOIN Position ON
                     Position. Skill = Employee. Skill;
To remove
duplicates
Natural join of Employee, Position
   select * from Employee NATURAL JOIN Position
```

## **OUTER JOIN** -- computes the join relations preserving dangling tuples by padding them with NULLs

A tuple in R is dangling if it doesn't join with any tuple in S; similarly a tuple in S is dangling if it doesn't join with any tuple in R

**FULL OUTER JOIN:** It pads dangling tuples of **R** and **S** preserving them **LEFT OUTER JOIN:** It pads dangling tuples of **R** only; tuples of **R** preserved **RIGHT OUTER JOIN:** It pads dangling tuples of *S* only; tuples of *S* preserved

	A	В
R:	1	2
	2	3
	В	C
- 1	_	

6

R FULL OUTER JOIN S: R LEFT OUTER JOIN S: R RIGHT OUTER JOIN S:

В

2

 $\mathbf{C}$ 5

6

A	В	C	A	В	C	A	
1	2	5 6	1	2	5	1	-
1	2	6	1	2	6	1 1	
2	3		2	3		<b>i</b>	
	7		Ē				

In SQL

R |NATURAL| [LEFT | RIGHT | FULL] OUTER JOIN S [ON .

Bipin C

S:

SQL> select * from part;	SQL> select * f	rom supplies;	
	SUPNAME	PART#	PRICE
PART# DESCR	PROVIBEC	1	710.2
	PROVIBEC	2	815.3
1 part1	PROVIBEC	3	325
2 part2	PROVIBEC	4	795.99
3 part3	SUPORIO	1	695.99
-	SUPBEC	2	799.98
4 part4	NDG-SUPPLY	1	699.99
5 part5	NDG-SUPPLY	2	799.99
6 part6	NDG-SUPPLY	3	324.99
<u>-</u>	NDG-SUPPLY	4	795.98
7 part7	NDG-SUPPLY	7	754
8 part8	SUPBEC	1	699.98
9 part9	MANIBEC	1	727.99
-	MDG-SUPPLY	1	699.99
	MDG-SUPPLY	2	799.99
	MDG-SUPPLY	3	324.99
	MDG-SUPPLY	4	795.98

```
select *
from part p right outer join supplies s
on (p.part# = s.part#);
    PART# DESCR SUPNAME PART# PRICE
        1 part1 MDG-SUPPLY
1 part1 MANIBEC
1 part1 SUPBEC
                                               699.99
                                          1
                                                727.99
        1 part1 1 part1
                                           1
                                                699.98
                  SUPORIO PROVITE
                                          1
                   NDG-SUPPLY
                                                699.99
                                          1
        1 part1
                                                695.99
        1 part1
                   PROVIBEC
                                          1
                                                 710.2
        2 part2
                                          2
                                                 799.99
                   MDG-SUPPLY
                                          2
                                                799.99
        2 part2
                   NDG-SUPPLY
                                           2
        2 part2
                                                799.98
                    SUPBEC
                                          2
        2 part2
                   PROVIBEC
                                                 815.3
        3 part3
                                          3
                  MDG-SUPPLY
                                                324.99
                  NDG-SUPPLY
                                          3
        3 part3
                                                324.99
                                          3
        3 part3
                   PROVIBEC
                                                   325
                                          4
                                                795.98
        4 part4
                   MDG-SUPPLY
                                                795.98
                   NDG-SUPPLY
                                          4
        4 part4
                                          4
                                                795.99
        4 part4
                   PROVIBEC
        7 part7 NDG-SUPPLY
                                                   754
17 rows selected.
```

SQL> select \* from part p full outer join supplies s on (p.part# = s.part#); PART# DESCR SUPNAME PART# 1 part1 PROVIBEC 710.2 2 part2 PROVIBEC 815.3 PROVIBEC 325 3 part3 3 795.99 4 part4 PROVIBEC 1 part1 SUPORIO 695.99 2 799.98 2 part2 SUPBEC 1 part1 NDG-SUPPLY 1 699.99 2 2 part2 NDG-SUPPLY 799.99 3 324.99 3 part3 NDG-SUPPLY 795.98 4 4 part4 NDG-SUPPLY 7 7 part7 754 NDG-SUPPLY 1 699.98 1 part1 SUPBEC MANIBEC
MDG-SUPPLY
MDG-SUPPLY
MDG-SUPPLY 1 727.99 1 part1 1 2 3 4 699.99 1 part1 2 part2 799.99 3 part3 324.99 795.98 4 part4 MDG-SUPPLY 5 part5 8 part8 6 part6 9 part9

Bipin C Desai

53

# Joins in SQL

```
mysql> select * from R; mysql> select * from S;
+---+
                  +----+
                  | A | B | C |
+---+
                 | 10 |
| a1 | 10 |
          12 |
                         12 | d1 |
| a2 |
     11 |
                  | 11 |
                         12 | d2 |
          12 |
| a3 |
                     6 |
                         14 | d3 |
                      9 |
| a4 |
     8 |
          17 I
                         12 | d4 |
+---+
                  +----+
```

In MySQL(up to version 5.7 at least), JOIN, CROSS JOIN, and INNER JOIN are syntactic equivalents: i.e., they can be used interchangeably.

In standard SQL, they are not equivalent. INNER JOIN is used with an ON clause,

```
select * from Employee JOIN Position ON
```

Position. Skill = Employee. Skill;

select R.a, T.e from R inner join S on R.b = S.b inner join T on S.c = T.c +---+---+

MariaDB [test]> select \* from R join S on R.B=S.B;

CROSS JOIN is used as follows: select \* from R cross join S;

```
| A | B | C | B | C | |
| a1| 10| 12| 10| 12| D1|
| a2| 11| 12| 11| 12| D2|
       9 | 15 |
3 rows in set (0.028 sec)
```

Bipin C Desai

55

+		-+-		-+-		-+-				+	-+
ALL THE FOLLOWING	A	-	В	1	С	1	В		С	D	1
ARE EQUIVALENT IN +		-+-		-+-		-+-		<b>-</b> -		+	-+
MySQL	a1		10		12		10		12	d1	.
Note: result schema R    S	a2		11		12		10		12	d1	
select *	a3		9		14		10		12	d1	
from R <b>JOIN</b> S;	a4		8		17		10		12	d1	
l	a1		10		12		11		12	d2	
select *	a2		11		12		11		12	d2	
from R, S;	a3		9		14		11		12	d2	
l	a4		8		17		11		12	d2	
select *	a1		10		12		6		14	d3	
from R CROSS JOIN S;	a2		11		12		6		14	d3	
I	a3		9		14		6		14	d3	
I	a4		8		17		6		14	d3	
select *	a1		10		12		9		12	d4	
from R <b>INNER</b> JOIN S;	a2		11		12		9		12	d4	
	a3		9		14		9		12	d4	
4*4 rows in result	a4		8		17		9		12	d4	
+		-+-		-+-		-+-				+	<del>-</del> +

Equijoin: join predicate containing an equality operator.

- combines rows that have same values for the specified columns.

If two tables in a join query have no join predicate the DBMS returns a **Cartesian product**.

#### **Outer Join**

An outer join extends the result of a simple join.

An **outer join** returns all rows that satisfy the join condition and those rows from one table for which no rows from the other satisfy the join condition.

Such rows are not returned by a simple join.

Bipin C Desai

5′

```
Joins: Equi-Joins
```

```
select * from R JOIN S on R.B=S.B;
```

	•			В				•
				10				
a2		11	12	11		12	d2	
a3		9	14	9		12	d4	

select \* from R JOIN S on R.B=S.B and R.C=S.C;

```
| A | B | C | B | C | D |
+----+-----+-----+-----+-----+
| a1 | 10 | 12 | 10 | 12 | d1 |
| a2 | 11 | 12 | 11 | 12 | d2 |
```

```
select * from R left outer join S on R.B=S.B;
                        1 C
ΙA
     | B
           I C
                 ΙB
                              l D
               12 |
        10 |
                     10 |
                            12 | d1
| a1 |
1 a2 1
               12 |
                     11 |
                            12 | d2
        11 |
               14 |
                      9 |
                            12 | d4
1 a3 1
         9 |
| a4 |
         8 |
               17 | NULL | NULL | NULL |
select * from R left outer join S
on R.B=S.B and R.C=S.C;
        I C
                        I C
I A I B
               ΙB
+---+
               12 |
                     10 |
                            12 | d1
| a1 |
        10 |
1 a2 1
        11 |
               12 |
                     11 |
                            12 | d2
| a3 |
        9 |
               14 | NULL | NULL | NULL |
        8 |
| a4 |
              17 | NULL | NULL | NULL |
```

59

```
select *
from R right outer join S
on R.B=S.B and R.C=S.C;
+----+
       | C
l A
     ΙB
             | B
                  | C
                         l D
----+
        10 |
             12 |
                  10 |
                       12 | d1
a1
             12 |
                       12 | d2
1 a2
        11
                  11 |
                       14 | d3
| NULL | NULL | NULL |
                 6 I
| NULL | NULL | NULL |
                 9 |
                       12 | d4
```

```
FULL outer join does not exist in MySQL;
Simulated by:
select *
from R left outer join S on R.B=S.B
UNION
select *
from R right outer join S on R.B=S.B;
+----+
       l B
             I C
                    l B
                 12 |
                        10 |
                              12 | d1
 a1
          10 |
1 a2
          11
                 12 I
                        11
                              12
                                   d2
l a3
           9
                 14 |
                         9
                              12
                                   d4
la4
                 17 | NULL | NULL | NULL |
 NULL | NULL | NULL |
                         6
                              14
                                 1 d3
```

```
select *
                            FULL outer join does
from R left outer join S
                            not exist in MySQL;
on R.B=S.B and R.C=S.C
                            Simulated by:
UNION
select *
from R right outer join S
on R.B=S.B and R.C=S.C;
  Α
        l B
                С
                        В
                               I C
                    12 |
 a1
            10
                           10
                                   12
                                         d1
l a2
                    12
                                   12
                                         d2
            11
                           11
 a3
             9 |
                    14 |
                         NULL
                                 NULL |
                                        NULL
 a4
             8 |
                    17 |
                         NULL |
                                 NULL |
                                        NULL |
 NULL |
         NULL | NULL |
                            6
                                   14 |
                                         d3
  NULL |
         NULL | NULL |
                            9
                                   12 |
                                        d4
```

```
Full outer Join: Oracle
                                     The Full outer join
                                     simulation,
  select *
                                     as in MySQL
  from Rr
                                     causes problems
  full outer join
                                     with column
  Ss
                                     headings: the
  on r.B=s.B and r.C=s.C
                                     ones used
        В
              C
                    В
                          C
  Α
                                D
                                     would be from
        10
                          12
              12
                                d1
  a1
                    10
                                     the system catalog!
              12
  a2
        11
                    11
                          12
                                d2
                                d3
                    6
                          14
                                d4
                    9
                           12
        9
  a3
              14
  a4
        8
              17
  6 rows returned in 0.01 seconds
```

Bipin C 63

```
select *
                        ВС
                                    C
                    Α
                                В
                                        D
from R cross join S
                        10
                            12
                                10
                                    12
                                         d1
                    a1
                            12
                                11
                                    12
                                         d2
                        10
                    a1
select *
                        10
                            12
                                6
                                    14
                                        d3
                    a1
from R, S
                        10 12
                                    12
                                        d4
                    a1
                                9
                                    12
                        11
                            12
                                10
                                         d1
                    a2
                            12
                                11
                                         d2
                        11
                                    12
                    a2
                                             Oracle: Cross Join
                        11
                            12
                    a2
                                6
                                    14
                                        d3
                                             - explicit and implicit
                    a2
                        11
                           12
                                    12
                                        d4
                                9
                                              However, Oracle some
                           14 10
                                    12
                    аЗ
                        9
                                        d1
                                              releases get
                                    12 d2
                    a3
                           14
                               11
                        9
                                              chocked up by:
                                   14 d3
                    аЗ
                        9
                           14
                               6
                                   12
                        9
                           14
                               9
                                        d4
                    аЗ
                                             select *
                           17
                               10 12 d1
                        8
                    a4
                                             from R JOIN S;
                           17
                        8
                               11
                                    12
                                        d2
                    a4
                           17
                                        d3
                    а4
                        8
                                6
                                   14
                                             select *
                           17
                                   12
                    a4
                        8
                                9
                                       d4
                                             from R inner join S
```

Bipin C

64

# Natural Join

+-		+-			+-			+
	Α		В			С		
+-		+-			-+-			+
	a1			10	1	1	2	1
	a2			11		1	2	
	a3			9		1	4	
	a4			8		1	7	

+.			+-			-+-		-+
	В		İ	С		İ	D	
+-			+-			-+-		+
		10			12		d1	
		11			12		d2	
		6			14		d3	
		9			12		d4	

select \* from R natural join S

Α В C D 10 12 d1 a1 11 12 a2 d2

2 rows returned in 0.03 seconds

Bipin C Desai

# Self Join

										S	
Ī	A	Ī	В		С	I	Ī	В	1	C	
										12	
İ	a2	İ	11	İ	12	ĺ	İ	11	İ	12	
	a3		9		14			6		14	
-	a4	-	8		17			9		12	
- 1		- 1		- 1			1		- 1		

MariaDB [test]> select r1.A, r2.A, r1.B, r1.C \_\_\_\_\_ To avoid duplicate

- -> from R r1
- -> inner join R r2 on r1.C=r2.C
- -> where r1.A < r2.A
- -> order by r1.A, r2.A

| a1 | a2 | 10 | 12 | +---+

# **Join Operation Execution**

## Hash joins

- In a *hash* join, a DBMS does a full-scan of one of the tables in the join operation to build a main-memory hash table. Then it searches for a matching value in the (hash table) for the other table.

Hash joins need more main memory but it could execute faster for certain types of join, the hash join will execute faster than a nested loop join.

Bipin C Desai

67

# **Join Operation Execution**

## Nested loops join

- The *nested loops* join is one of the original join pans and it is the most common method. In a *nested loops* join, we have two tables: one is the left operand table and the other the right hand table.

The an index for the attribute of left table is accessed to get the row IDs of the rows with the attribute value.

The matching rows of the second table are then probed in a nested loop and matching rows of the two tables are joined using an index range scan.

# Join Operation Execution

Some queries will perform faster with NESTED LOOPS joins, some with HASH joins, while others favor sort-merge joins.

It is difficult to predict what join technique will be fastest *a priori*, so many tuning a database is to test the often use joins and record the statistics to guide the productions operations

Bipin C Desai

6

#### Constraints

Primary keys declarations

Foreign key constraints is a referential integrity constraints

If a supplier supplies part# 4, then we must have **part# 4** in the table for part

Primary key constraint is declared within the DDL SQL command CREATE TABLE using the keywords PRIMARY KEY or UNIQUE

Many DBMSs treat them as synonyms: A table may have one primary key but any number of "unique" declarations

```
SOL> create table x
                                    NOTE: References
 (a number (4) primary key);
                                    must be a primary
Table created.
                                    key or an attribute
SQL> create table y
                                    with an unique
  (s number (4),
                                    attribute
   b number (4) references x(a));
Table created.
SQL> insert into y values(1, 2);
insert into y values (1, 2)
ERROR at line 1:
SQL > insert into x values (2);
1 row created.
SQL> insert into y values(1, 2);
1 row created.
```

```
SQL> create table x(
 2 a number (4) primary key,
 3 b number(4));
                      This is considered as a 'foreign key'
Table created.
                                    Note: The table x must
SQL> create table y()
                                    be created before we can
 2 c number (4) primary key,
                                    create table y
 3 d number (4) references x(a));
                                 SQL> create table z(
Table created.
                                  2 e number (4) references x(a),
SQL> create table z(
                                  3 f number (4) references y(c));
 2 e number (4) references x(b),
                                 Table created
 3 f number (4) references y(c));
ERROR:no matching unique or primary key for this column-list
create table w(
e number (4) references x(a),
f number (4) references y(d));
ERROR no matching unique or primary key for this column-list
```

```
Possible situations violating foreign key constraints:
    Insert:
         SQL> insert into y values(1, 2);
         insert into y values(1, 2)
         ERROR at line 1:ORA-02292: integrity constraint (SCOTT.SYS C002908)
           violated - child record found No tuple in x with primary key 2
                                        Some tuple in y is referencing the
    Update:
                                        current key value of a tuple x
         SQL > update x set a=3;
         ERROR at line 1:
         ORA-02292: integrity constraint (SCOTT.SYS C002908) violated - child record
         found
                                        No tuple in x with primary key 2
         SQL> update y set b=4;
         ERROR at line 1: ORA-02291: integrity constraint (SCOTT.SYS C002908)
           violated - parent key not found
                                         Some tuple in y is referencing the
    Delete:
                                         current key value of a tuple x
         SQL > delete from x where a=2;
         delete from x where a=2
         ERROR at line 1: ORA-02292: integrity constraint (SCOTT.SYS C002908)
           violated - child record found
```

Bipin C

/3

```
Insert with null values is OK!

SQL> insert into y values(1,null);
1 row created.

SQL> insert into y values(2,null);
1 row created.

SQL> select * from y;

C D

1
2
```

There are *three* policies choices for situations violating foreign key constraints

The **reject** policy (*default*)

The system will reject any such violating request and a run-time error will be generated. The database state will *not* change.

In case of update or delete request:

The **cascade** policy: changes to the referenced attributes are "mimicked" at the foreign key (e.g. y.b)

The **set-NULL** policy: set the referencing attribute to NULL (e.g., y.b)

Options/policies may be chosen for deletes and updates, *independently* 

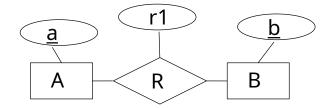
ON DELETE {CASCADE | SET NULL }]

[ ON UPDATE] { CASCADE | SET NULL }]

The policy to be used is a design decision and must conform to the business rules of the underlying application.

Bipin C

75



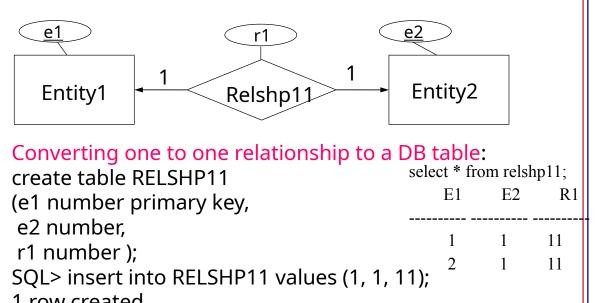
## Attributes of R

Type of rel-ship	From A	From B	attrib
m-m	$a_{pf}$	$b_{pf}$	r
m-1	$a_{pf}$	$b_{\rm f}$	r
1-1	$a_{pf}$	$b_{fu}$	r

 $x_{pf}$  attribute is prime and foreign key  $x_{f}$  attribute is foreign key  $x_{fu}$  attribute is foreign key and unique

For m-1 multiplicity either of the entities could be "A" -2 cases

In the case of 1-1 multiplicity either of the entities could be "A" - 2 choices



1 row created.

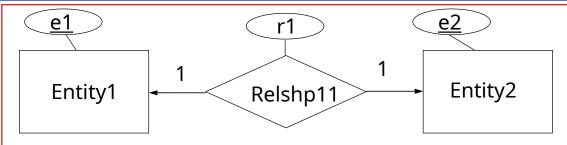
SQL>insert into RELSHP11 values (1, 2, 11);

ERROR at line 1: ORA-00001: unique constraint violated

SQL> insert into RELSHP11 values (2, 1, 11);

1 row created. ← Not a 1-to-1 relationship!

Bipin C Desai



## Converting one to one relationship to a DB table:

create table RELSHP11 (e1 number primary key, e2 number unique not null,

e2 value for two different e1 values (not 1to-1). Without null, we can leave out some values for entity e2. ( not a 1-to-1 relationship!)

Without unique, we could insert the same

r1 number ); SQL> insert into RELSHP11 values (1, 1, 11);

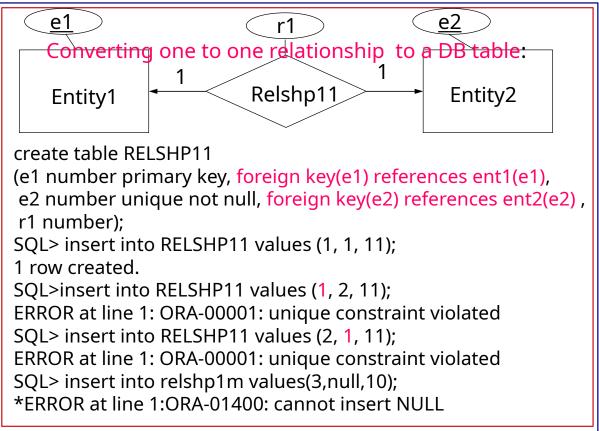
1 row created.

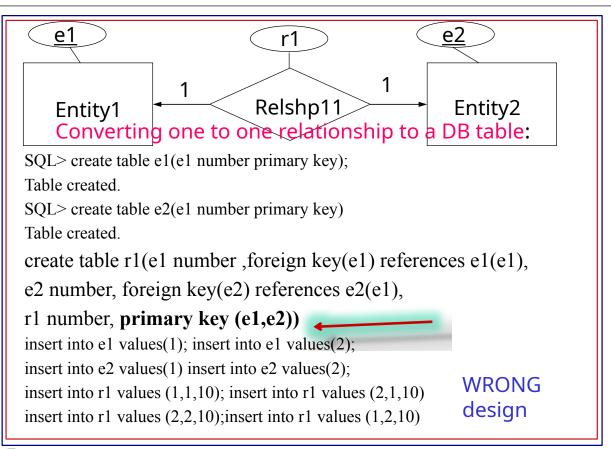
SQL>insert into RELSHP11 values (1, 2, 11);

ERROR at line 1: ORA-00001: unique constraint violated

SQL> insert into RELSHP11 values (2, 1, 11);

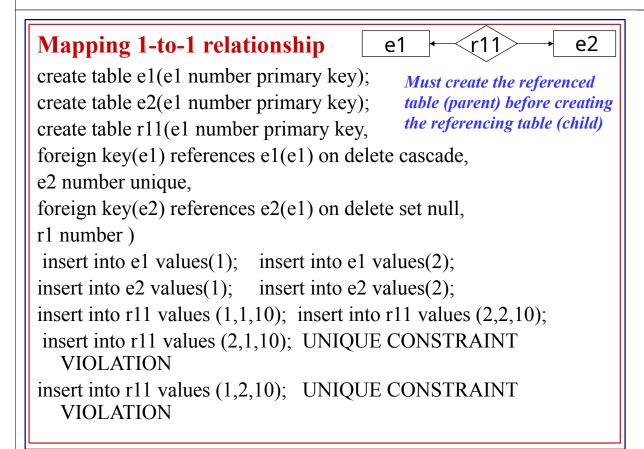
ERROR at line 1: ORA-00001: unique constraint violated





```
create table r1new(e1 number unique, foreign key(e1) references e1(e1),
e2 number unique, foreign key(e2) references e2(e1),
                                               Added the unique attribute
r1 number,
primary key (e1,e2))
                                               With the unique attribute is
insert into r1new values(1,1,10);
                                                the primary key redundant?
1 row created
insert into r1new values(2,1,10)
                           Trying to insert another relationship involving e2.e1=1
ERROR at line 1:
ORA-00001: unique constraint (SCOTT.SYS C004018) violated
insert into r1new values(3,3,10)
                            Referential integrity enforced
ERROR at line 1:
ORA-02291: integrity constraint (SCOTT.SYS C004020) violated - parent key
   not found
drop table e1
                        Since the value(s) in e1 are being referenced
ERROR at line 1:
ORA-02449: unique/primary keys in table referenced by foreign keys
```

© Bipin C 81



```
Mapping many-to-1 relationship
                                                       RAB
                                                                    В
 SQL> create table A(
                                  SQL> create table B(
 al number(3) primary key,
                                  b1 number(4) primary key,
 a2 number (3)); Table created
                                  b2 number (3)); Table created
 SQL> create table RAB(
                                Must create the referenced table (parent)
  r1 number(3),
                                before creating the referencing table (child)
  r2 number(4) primary key,
  constraint fk 1 foreign key (r1) references A(a1),
  constraint fk 2 foreign key (r2) references B(b1)); Table created.
 SQL> insert into RAB values (null,null);
 *ERROR ORA-01400: cannot insert NULL into ("SCOTT"."RAB"."R2")
 SQL> insert into RAB values (null,1);
 *ERROR ORA-02291: integrity constraint (SCOTT.FK 2) violated - parent key not found
 SQL> insert into A values(1,1); 1 row created
 SQL> insert into RAB values (null,1);
 *ERROR ORA-02291: integrity constraint (SCOTT.FK 2) violated - parent key not found
 SQL> insert into B values(11,11);1 row created.
```

SQL> insert into RAB values (1,12);

\*ERROR ORA-02291: integrity constraint (SCOTT.FK\_2) violated - parent key not found SQL> insert into RAB values (2,11);

\*ERROR ORA-02291: integrity constraint (SCOTT.FK\_1) violated - parent key not found SQL> insert into RAB values (1,11); 1 row created.

SQL> delete A where a1=1;

\*ERROR ORA-02292: integrity constraint (SCOTT.FK\_1) violated - child record found SQL> delete B where b1=11;

\*ERROR ORA-02292: integrity constraint (SCOTT.FK\_2) violated - child record found SQL> insert into A values (2, 2); 1 row created.

SQL> insert into RAB values (2,11);

\*ERROR ORA-00001: unique constraint (SCOTT.SYS\_C004197) violated SQL> insert into B values (12, 12); 1 row created.

SQL> insert into RAB values (1, 12); 1 row created.

```
CREATE TABLE Supplier(
SID numeric(10) not null,
SName varchar2(50) not null,
Contact varchar2(50),
CONSTRAINT s_pk PRIMARY KEY (SID, SName));

The "ON CASCADE DELETE" in the foreign key constrain in Parts causes all tuples with the matching SID, SName
```

PNo numeric(10) not null, SNo numeric(10) not null, SName varchar2(50) not null, CONSTRAINT p\_fk

FOREIGN KEY (SNo, SName)

REFERENCES Supplier(SID, SName) ON DELETE CASCADE);

values in Parts to be deleted when a

record in Supplier is deleted

© Bipin C Desai

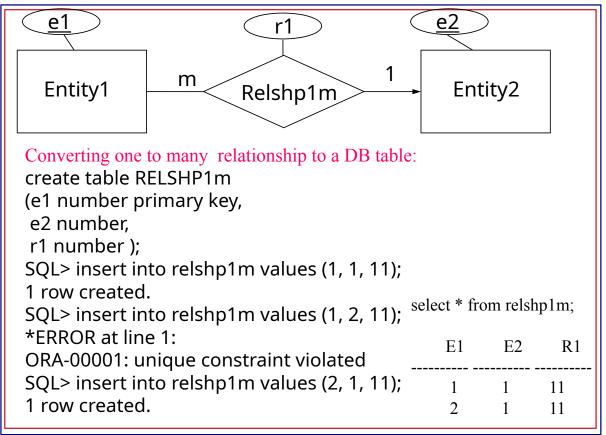
```
Added the on delete clauses
create table e1(e1 number primary key);
create table e2(e1 number primary key);
                                        Changed the composite primary key
create table r11(e1 number primary key,
foreign key(e1) references e1(e1) on delete cascade, SQL> select * from e1;
e2 number unique,
                                                         SOL> select * from e2
foreign key(e2) references e2(e1) on delete set null,
                                                              E1
rl number)
insert into e1 values(1); insert into e1 values(2); 2
insert into e2 values(1);
                         insert into e2 values(2);
insert into r11 values (1,1,10); insert into r11 values (2,2,10);
insert into r11 values (2,1,10); UNIQUE CONSTRAINT VIOLATION
insert into r11 values (1,2,10); UNIQUE CONSTRAINT VIOLATION;
                         SQL> select * from r11;
SQL> delete from e2;
                                   E2
2 rows deleted.
     e2 columns set to null
                                         10
SQL> delete from e1;
                                         10
2 rows deleted.
                                                  SQL> select * from r11;
                                                  no rows selected
            Delete of rows in e1 cascades to r11
```

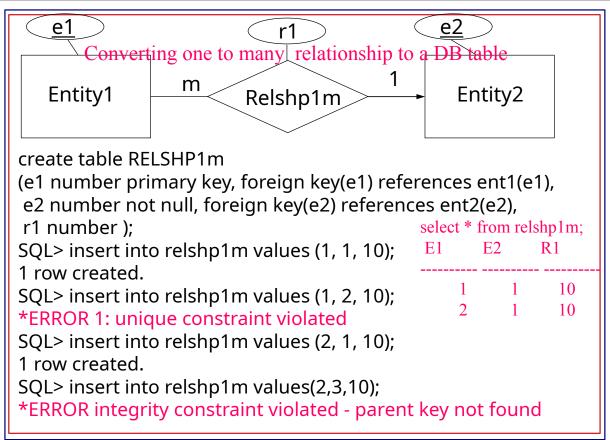
```
create table e1(e1 number primary key);
create table e2(e1 number primary key);
create table r11(e1 number primary key,
foreign key(e1) references e1(e1) on delete cascade,
e2 number unique,
foreign key(e2) references e2(e1) on delete set null,
r1 number )
                        SQL> delete from e2;
SQL> delete from e1;
                        y rows deleted
x rows deleted
SQL> select * from r11;
no rows selected
SQL> insert into r11 values (3,null,null);
ERROR at line 1:
ORA-02291: integrity constraint (SCOTT.SYS_C004071)
```

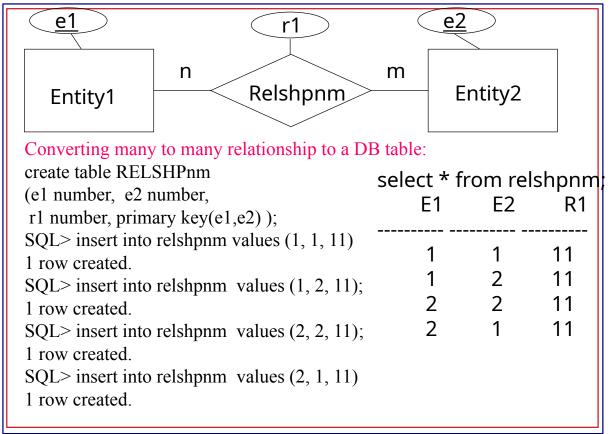
violated - parent key not found

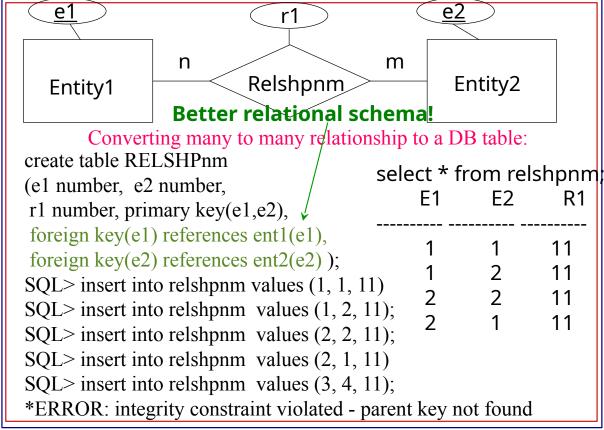
Can't insert a null value in r11 of the foreign key constraint

Bipin C Desai









91

# Alternate Candidate keys — How to implement?

Candidate keys of our friends:

Phone number create table friends(

or name varchar2(30) unique,

(name, address) address varchar(30) unique,

phone decimal(16) primary key);

create table friends1( name varchar2(30),

address varchar(30),

phone decimal(16) primary key,

unique (name,address));

create table friends2(

name varchar2(30), address varchar(30),

phone decimal(16) unique,

primary key (name, address));

# **Alternate Candidate keys**

```
create table friends (
name varchar2(30) unique,
address varchar(30) unique,
phone decimal(16) primary key);

SQL> insert into friends values('smith','montreal',1234);
1 row created.

SQL> insert into friends values('smith','laval',1235);
ERROR: unique constraint (SCOTT.SYS_C003729) violated
SQL> insert into friends values('smith','montreal',1235);
ERROR: unique constraint (SCOTT.SYS_C003729) violated
SQL> insert into friends values('brown','laval',1235);
1 row created.

Only 1 Smith, 1 Montreal, etc.!
```

Bipin C

93

# Alternate Candidate keys

```
create table friends1(
name varchar2(30),
address varchar(30),
phone decimal(16) primary key,
unique (name,address));

SQL> insert into friends1 values('smith','montreal',1234);
1 row created.
SQL> insert into friends1 values('smith','laval',1236);
1 row created.
SQL> insert into friends1 values('smith','laval',1237)
* ERROR at line 1:
ERROR: unique constraint (SCOTT.SYS_C003727) violated
```

# **Alternate Candidate keys**

create table friends2(
name varchar2(30), address varchar(30),
phone decimal(16) unique, primary key(name,address));
SQL> insert into friends2 values('smith','laval',1235);
1 row created.

SQL> insert into friends2 values('smith','montreal',1235); ERROR: unique constraint (SCOTT.SYS\_C003724) violated

SQL> insert into friends2 values('brown','laval',1235);

ERROR: unique constraint (SCOTT.SYS C003725) violated

SQL> insert into friends2 values('smith','laval',1236);

ERROR: unique constraint (SCOTT.SYS\_C003724) violated

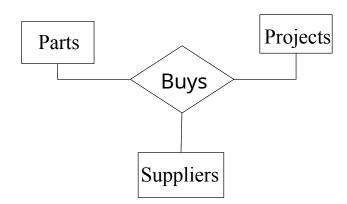
SQL> insert into friends2 values('smith','montreal',1234); 1 row created. SQL> select \* from friends2 /

NAME	ADDRESS	PHONE
smith	montreal	1234
smith	laval	1235

Bipin C

9

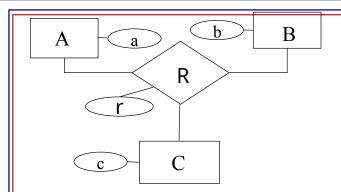
# Ternary Relationships and multiplicity



The above is an example of a three way relationship The multiplicity could be m or 1 for any of the entity sets involved in the relationship

Ignoring the permutation of the entities we need to consider Four cases: m-m-m or m-m-1 or m-1-1 or 1-1-1

Considering permutations there are 8 cases 2\*2\*2 or 1+3+3+1



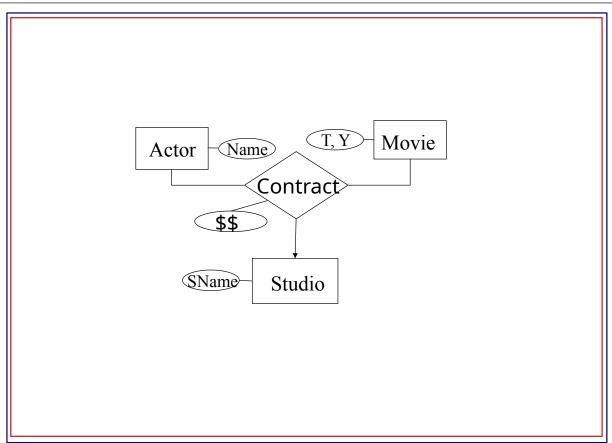
Type	From A	From B	From C	attrib
m-m-m	$a_{pf}$	$b_{pf}$	c <sub>pf</sub>	r
m-m-1	$a_{pf}$	$b_{pf}$	$c_{\rm f}$	r
m-1-1	$a_{pf}$	$b_{\rm f}$	$c_{\rm f}$	r
1-1-1	$a_{pf}$	$b_{fu}$	$c_{\mathrm{fu}}$	r

 $x_{pf}$  attribute is prime and foreign key  $x_{f}$  attribute is foreign key  $x_{fu}$  attribute is foreign key and unique

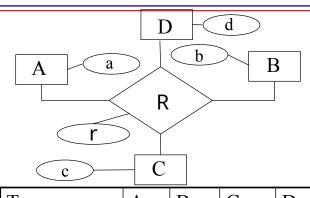
In the case of 1-1-1 multiplicity any of the entities could be used for A

Bipin C Desai

9'



The Bipin C Desai



 $x_{pf}$  attribute is prime and foreign key  $x_f$  attribute is foreign key  $x_{fu}$  attribute is foreign key and unique

Type	A	В	C	D	attrib
m-m-m-m	$a_{pf}$	$b_{pf}$	$c_{pf}$	$d_{pf}$	r
m-m-m-1	$a_{pf}$	$b_{pf}$	$c_{pf}$	$d_{\rm f}$	r
m-m-1-1	$a_{pf}$	$b_{pf}$	$c_{\rm f}$	$d_{\rm f}$	r
m-1-1-1	$a_{pf}$	$b_{\rm f}$	$c_{\rm f}$	$d_{\rm f}$	r
1-1-1-1	$a_{pf}$	$b_{fu}$	c <sub>fu</sub>	$d_{\text{fu}}$	r

In the case of 1-1-1 multiplicity any of the entities could be used for A

Total # permutations would be 2\*2\*2\*2= 1+4+6+4+1=16

Bipin C Desai

99

Given the following relations, find the CS courses that Brenda can take.(Note: she cannot take a course already passed and must have all pre-requisites)

Student(Sno Name)
Dept(Dno, Dname)
Course(Cno, Dno, Cname)
Enroll(Sno,Cno,Grade)
Prereg(Cno,Pcno)

CSCrs =  $\Pi_{Cno} \sigma_{Dname=CS}$  DEPT  $\bowtie$  Course BrendaPassed =  $\Pi_{Cno}$  (( $\sigma_{Grade \neq F}$ Enroll)  $\bowtie$  ( $\sigma_{Name=Brenda}$ Student)) BrendaSatisfiedPre = CSCrs X BrendaPassed BrendaCantTake =  $\Pi_{Cno}$  (Prereq - BrendaSatisfiedPre) BrendaCanTake = (CSCrs - BrendaPassed) - BrendaCantTake

(b) Bipin C Desai

```
select c.cno
from course c, dept d, student s
where d.dept='Computer Science' and s.sname='Brenda' and
c.dno = d.dno and not exists(
                                  Has not already taken and
select e.cno
                                  passed the course
from enroll e
where e.cno=c.cno
s.sno=e.sno and e.grade <>'F') and
not exists(select p1.cno
     from preq p1
     where p1.cno = c.cno and
                                 Has all the pre-req.
      not exists(select e1.cno
          from enroll e1, student s1
           where e1.sno = s1.sno and s1.sname='Brenda' and
            e1.sno = s1.sno and e1.grade <>'F' and
            e1.cno = p1.pcno));
```

# Alternate SQL

```
select c.cno
from course c, dept d, student s
where d.dept='Computer Science' and s.sname='Brenda' and
c.dno = d.dno and c.cno not in(
                                    The course c.cno has not
select e.cno
                                    already been taken and
from enroll e
                                     passed
where s.sno=e.sno and e.grade <>'F') and
not exists(select p1.cno
     from preq p1
     where p1.cno = c.cno and
                                   Has all the pre-req.
     not exists(select e1.cno
            from enroll e1
            where e1.sno = s.sno and
            e1.sno = s.sno and e1.grade <>'F' and
            e1.cno = p1.pcno));
```

d Bipin C Desai

We may associate the NOT NULL constraint with an attribute for a table

#### Two consequences:

- 1. We can't insert a tuple into the table without giving value for the attribute defined with the NOT NULL constraint.
- 2. We can't use the "set-null" policy to fix foreign-key violations for such attributes

Bipin C

```
SQL> create table z( d number (4)
check (d >999));
Table created.
SQL> insert into z values(1);
insert into z values(1)
ERROR at line 1: ORA-02290: check constraint
(SCOTT.SYS_C002909) violated
```

#### Difference between a check and a foreign-key constraint.

The check is done only when a tuple is inserted or updated. A foreign key constraint checks for any update, deletes

Bipin C

103

```
*Example:

CREATE TABLE Star(

name CHAR(30) PRIMARY KEY,

address VARCHAR(255),

gender CHAR(1),

birthdate DATE,

CHECK (gender = 'F' OR name NOT LIKE 'Ms.%')
);

This constraints says that if a star is male (M), then his name must not begin with 'Ms.' (¬condition); Here we
```

used (gender='F'  $\mathbf{OR}$  ¬condition) for (M $\rightarrow$ not Ms).

Bipin C

106

```
CREATE TABLE person(
       name CHAR(30) PRIMARY KEY,
       address VARCHAR(255),
       gender CHAR(1),
        dob DATE,
        CHECK (gender = 'F' OR name NOT LIKE 'Ms.%') );
       SQL> insert into person (name, gender)
          values ('Ms. John Smith', 'M');
       ERROR at line 1:
       ORA-02290: check constraint (SCOTT.SYS C002916)
        violated
       SQL> insert into person (name, gender)
       values ('Ms. George Sands', 'F');
       1 row created.
       SQL> alter table person add constraint Person Adr unique (gender);
       Table altered
       SQL> alter table person add income number (12,2);
       Table altered.
```

Bipin C Desai

10

Assertions, or general constraints, are boolean-valued SQL expressions that must always be true

Sometimes we need a constraint that involves relation as a whole or part of the database schema

Assertions are checked when a mentioned relation changes Assertion in SQL not supported by most DBMS:

```
CREATE ASSERTION PoorPerson CHECK
(NOT EXIST (SELECT *
FROM person
WHERE income > 10000
)
);
```

Bipin C 108

# Constraints and triggers

SQL provides a number of features to express *integrity constraints* (primary, foreign key) as part of the database schema.

Constraints, in essence, provide database designers with more control over the database content

An *active* element is an statement that we write once, store in the database, and "program" to execute when an event occurs. This event is considered as a **trigger** 

An event may be an insertion of a tuple into a predefined table or a specified change in the database that causes a specified (boolean-valued) condition to become true

It is also possible to implement many of the constraints and triggers in scripts such as PHP, JSP etc. However, this is left to the application programmer and has to be included in each application!

Bipin C

100

## **Triggers**

Procedures that are *implicitly* executed when an INSERT, UPDATE, or DELETE statement is issued against an associated table.

Simple procedure is *explicitly* executed by a user, application, or a trigger.

Triggers (one or more) are implicitly executed by Oracle when a triggering INSERT, UPDATE, or DELETE statement is issued, regardless of how it is issued(user or application).

A trigger can restrict DML operations against a table (time of day/week etc.)

A statement in a trigger body could causes another trigger to be fired! Such the triggers are said to be *cascading triggers*.

(h) Bipin C Desai

# Why Triggers?

• A required referential integrity rule cannot be enforced using: the integrity constraints such as:

NOT NULL, UNIQUE key, PRIMARY KEY, FOREIGN KEY, CHECK, update CASCADE, update and delete SET NULL, update and delete SET DEFAULT

A trigger has three parts:

- a triggering *event* (some statement)
  - a trigger *condition*
  - a trigger action
- Enforce referential integrity when child and parent tables are on different nodes of a distributed database
- Enforce complex business rules not definable using integrity constraints

Bipin C Desai

11

## **TRIGGERS**

Triggers are often called event-condition-action rules

- An *Event* is any specified changes in the DB, due to insertions, deletions or updates
- A **Condition** is a predicate or test to determine if the specified trigger is applicable
- An Action is one or more SQL statements

Triggers are not supported in SQL2

Differ from checks or SQL2 assertions in that:

Event is programmable, rather than **implied** by the kind of check

Condition is not available in checks

Action could be any sequence of database operations

Triggers are essential for Active Database Management Systems (ADBMS)

Bipin C

Triggers are compiled by storing the procedure in a text file and compiling it with:

@filename

If we update an entire table with an SQL statement

A row-level trigger will be executed once for each tuple

A **statement-level trigger** will be executed only once for the entire update In a statement-level trigger:

We can not refer to old and new tuples

Instead, we can/should refer to

The set of old tuples – OLD TABLE
The set of new tuples – NEW TABLE

Bipin C

113

**Relation scheme:** Employee(name, empId, salary, dept, supervisorId)

**Constraint:** No employee gets a salary more than his/her supervisor.

**CREATE TRIGGER** Inform supervisor

**BEFORE INSERT OR UPDATE OF** salary, supervisorId **ON** Employee

**NEW ROW AS** new

FOR EACH ROW

WHEN (new.salary > (SELECT salary

FROM Employee

WHERE empId=new.supervisorId))

Begin

**ROLLBACK**;

Inform Supervisor(new.supervisorId, new.empId);

End;

Bipin C

114

## **Row Triggers (before and after)**

A *row trigger* is fired once for **each** row affected by, say, an UPDATE statement.

Row triggers are used when the trigger action depends on data provided by the triggering statement or rows that are affected.

## **Statement Triggers (before and after)**

A statement trigger is fired **once**, regardless of the number of rows in the table that the triggering statement affects (even if no rows are affected)

If a DELETE statement deletes several rows from a table, a statement-level DELETE trigger is fired only once, regardless of the number of rows are deleted from the table.

Useful when the trigger action does not depend on the data provided by the triggering statement or the rows that are affected.

Bipin C Desai

114

**Relation Scheme: Movie**(title, year, length, filmType, studioName, producerC#)

CREATE VIEW ParamountMovie AS SELECT title, year FROM Movie WHERE studioName = 'Paramount';

The following *trigger* replaces an insertion on the view (ParamountMovie) with an insertion on its underlying base table (Movie)

CREATE TRIGGER ParamountInsert This is WRONG for a view!

INSTEAD OF INSERT ON ParamountMovie

REFERENCING NEW ROW AS NewRow

FOR EACH ROW

**INSERT INTO** Movie(title, year, studioName)

**VALUES** (NewRow.title, NewRow.year, 'Paramount');

Bipin C

116

```
CREATE or REPLACE TRIGGER ParamountInsert
INSTEAD OF INSERT ON ParamountMovie
FOR EACH ROW
BEGIN
INSERT INTO Movie(title, year, studioName)
VALUES (:new.title, :new.year, 'Paramount');
                                        To-date Mysql/
end ParamountInsert;
                                        Maraiadb doesn't
                                         have 'instead of
Trigger created.
SQL> show errors trigger ParamountInsertinsert' option
No errors.
insert into ParamountMovie values ('Movie2016', '2016');
1 row created.
SQL> select * from movie;
TITLE
                         YEAR STUDIONAME
Movie2016
                         2016 Paramount
```

h Bipin C Desai

```
SQL> desc student;
                Null?
                         Type
 SID
                NOT NULL NUMBER (7)
 SNAME
                         VARCHAR2 (20)
MAJOR
                         CHAR (4)
 YEAR
                         NUMBER (1)
BDATE
                         DATE
create view cstdnt as
select sid as id, sname as name, bdate as dob
from student
where major='COMP';
View created.
SQL> select * from cstdnt;
        ID NAME
                                DOB
         8 Brenda
                                 13-AUG-89
```

CREATE OR REPLACE TRIGGER CStudentInsert----

INSTEAD OF INSERT ON cstdnt

REFERENCING NEW AS NewRow

FOR EACH ROW

INSERT INTO student(sid, sname, major, year, bdate)

VALUES (NewRow.id, NewRow.name, 'COMP', 1, NewRow.dob)

Warning: Trigger created with compilation errors.

SQL> SHOW ERRORS TRIGGER CStudentInsert; Errors for TRIGGER CSTUDENTINSERT:

#### LINE/COL ERROR

1/7 PL/SQL: SQL Statement ignored

2/51 PL/SQL: ORA-00984: column not allowed here

d Bipin C Desai

SQL> CREATE OR REPLACE TRIGGER CStudentInsert

instead of INSERT ON cstdnt

FOR EACH ROW

INSERT INTO student(sid, sname, major, year, bdate)

VALUES (:new.id, :new.name, 'COMP', 1, :new.dob)

Trigger created.

SQL> select \* from student;

SID	SNAME	MAJO	YEAR	BDATE
8	Brenda	COMP	2	13-AUG-89
10	Dupont	ENGL	1	13-MAY-80
13	Kelly	SENG	4	12-AUG-80
14	Jack	CSAP	1	12-FEB-77

Bipin C Desai

119

SQL> insert into cstdnt values(7,'Drew', '13-Sep-81'); 1 row created.

SQL> select \* from student;

6 rows in set (0.01 sec)

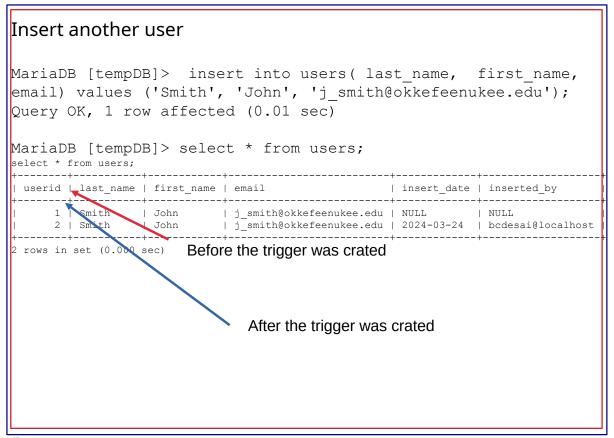
SID	SNAME	MAJO	YEAR	BDATE
8	Brenda	COMP	2	13-AUG-89
10	Dupont	ENGL	1	13-MAY-80
13	Kelly	SENG	4	12-AUG-80
14	Jack	CSAP	1	12-FEB-77
7	Drew	COMP	1	13-SEP-81

(b) Bipin C Desai

```
Other triggers:
DELIMITER |
                                           -before delete
CREATE TRIGGER users before insert
                                           -before update
BEFORE INSERT
                                           -after insert
 ON users FOR EACH ROW
                                           -after update
BEGIN
                                           -after delete
 DECLARE whoinserted varchar(50);
 -- Find username of person performing inserting a new user
 SELECT USER() INTO whoinserted;
 -- Update create date field to current system date
 SET NEW.insert_date = SYSDATE();
 -- Update created by field to the username of the person
      performing the INSERT
 SET NEW.inserted by = whoinserted;
END; |
DELIMITER;
                            show triggers;
```

h Bipin C Desai

```
MariaDB [test] > insert into users( last name, first name, email)
     -> values ('Smith', 'John', 'j smith@okkefeenukee.edu');
Query OK, 1 row affected (0.024 sec)
Now create trigger
MariaDB [test]> DELIMITER |
MariaDB [test] > CREATE TRIGGER users before insert
    -> BEFORE INSERT
    -> ON users FOR EACH ROW
    -> BEGIN
    -> DECLARE whoinserted varchar(50);
    -> -- Find username of person performing inserting a new user
    -> SELECT USER() INTO whoinserted;
    -> -- Update create_date field to current system date
    -> SET NEW.insert date = SYSDATE();
    -> -- Update created by field to the username of the person
        -- performing the INSERT
    -> SET NEW.inserted by = whoinserted;
    -> END; |
Query OK, 0 rows affected (0.011 sec)
MariaDB [test]> DELIMITER ;
```



Bipin C Desai

125

#### **HOW ARE TRIGGERS EXECUTED**

- SQL statement is issued
- Execute any BEFORE statement-level triggers
- For each row affected by the triggering SQL statement
  - Execute any BEFORE row-level triggers
  - Lock and change row, and perform integrity constraint checking The lock is not released until the transaction is committed
  - Execute any AFTER row-level triggers
- Execute any AFTER statement-level triggers

```
Example with triggers etc.
create table University(
Name CHAR(20) PRIMARY KEY,
City CHAR(20));
                                Name of the constraint
create table Engineer(
EID NUMBER(4),
                                Candidate key
SIN NUMBER(9),
Name char(20),
                                    Check constraint
AlmaMater CHAR(20),
HireAge number(2) CHECK (HireAge BETWEEN 25 AND 65).
CONSTRAINT Engineer PK PRÍMARY KEY(eid),
CONSTRAINT Engineer CK UNIQUE(SIN),
FOREIGN KEY (AlmaMater) REFERENCES University(NAME)
);
```

ⓑ Bipin C Desai

```
create table Project(
ProjNo NUMBER(4) primary key,
EID NUMBER(4),
FOREIGN KEY (EID) REFERENCES Engineer(EID)
);

create table Assigned(
ID NUMBER(4),
pno NUMBER(4),
CONSTRAINT Assign_FK1 FOREIGN KEY(ID)
REFERENCES Engineer(EID),
CONSTRAINT Assign_FK2 FOREIGN KEY(pno)
REFERENCES Project(ProjNo)
);
```

# insert into University values('ConU', 'Montreal'); insert into University values('UdeM', 'Montreal'); insert into Engineer values(11, 123456789, 'Smith', 'ConU', 35); insert into Engineer values(12, 234567891, 'Shah', 'UdeM', 73) \* ERROR at line 1: ORA-02290: check constraint (SCOTT.SYS\_C003385) violated insert into Engineer values(12, 234567891, 'Shah', 'UdeM', 33); insert into project values(1, 11); insert into project values(2, 11); insert into project values(3, 11); insert into Assigned values(11, 1); insert into Assigned values(12, 1);

(h) Bipin C Desai

```
SQL> Create or Replace package ProjEngg
as
EnggEid number(4); Package is a collections of end; Procedures and functions

Package created. Prigger of type Row level

SQL> Create or Replace Trigger WhichEngg
Before Insert on Project for each row begin

ProjEngg.EnggEid := :new.EID; end;

/ Trigger created.
```

# 

Bipin C Desai

```
SQL> insert into project values(4, 11);
1 row created.
SQL> select * from project;
                                  SQL> select * from project;
  PROJNO
                                     PROJNO
                                                        EID
                          11
                                               1
                                                             11
                          11
                                                             11
                          11
                                                             11
                                                             11
SQL> insert into project values(5, 11);
ERROR at line 1:
ORA-20001: **** Too many projects for this engineer! ****
ORA-06512: at "SCOTT NUMBEROFPROJS", line 7
ORA-04088: error during execution of trigger
              'SCOTT.NUMBEROFPROJS'
```

## **Mutating trigger**

A trigger that attempts to modify the same table that initiated the trigger is called a mutating trigger.

```
CREATE OR REPLACE TRIGGER person st
AFTER INSERT ON PERSON
REFERENCING NEW AS newRow
FOR EACH ROW
DECLARE STATUS CHAR (4);
BEGIN
STATUS := 'Poor';
IF (:newRow.income > 20000) THEN
STATUS := 'Med'; END IF;
IF (:newRow.income > 60000) THEN
STATUS := 'Rich'; END IF;
INSERT INTO person VALUES (:newRow.name, :newRow.dob,
           :newRow.income, STATUS);
END person st;
run
                Trigger created.
```

(b) Bipin C Desai

```
SQL> insert into person (name, dob,income)
values('Jones', '10-jun-68', 61000.00);
insert into person (name, dob,income)
values('Jones', '10-jun-68', 61000.00)

*

ERROR at line 1:
ORA-04091: table SCOTT.PERSON is mutating,
trigger/function may not see it
ORA-06512: at "SCOTT.PERSON_ST", line 11
ORA-04088: error during execution of trigger
'SCOTT.PERSON_ST'
Sql> drop trigger person_st;

Trigger dropped.
```

```
CREATE or REPLACE TRIGGER
Getting around mutation!
                           person1 st
create table person1 (
                           AFTER INSERT ON PERSON1
name char(25) primary key,
                           REFERENCING NEW AS newRow
dob date,
                           FOR EACH ROW
income number(12,2))
                           DECLARE STATUS CHAR(4);
                           BEGIN
                           STATUS := 'Poor';
drop table person cascade
                           IF (:newRow.income > 20000) THEN
constraints
                           STATUS := 'Med'; END IF;
                           IF (:newRow.income > 60000)THEN
                           STATUS := 'Rich'; END IF;
create table person (
                           INSERT INTO person VALUES
name char(25) primary key,
                           (:newRow.name, :newRow.dob,
dob date,
                           :newRow.income, STATUS);
income number (12,2),
                           END person st;
status char (6))
                           run
```

Bipin C Desai

```
SQL> insert into person1 values('Smith', '31-may-70', 21000.00);
SQL> insert into person1 values('John', '3-Apr-68', 11000.00);
SQL> insert into person1 values('Wang', '31-may-70', 60001.00);
SQL> select * from person1;
NAME
                      DOB
                                      INCOME
                           31-MAY-70
                                            21000
Smith
            03-APR-68
John
                                     11000
                    31-MAY-70 60001
Wana
SQL> select * from person;
NAME
                      DOB
                                  INCOME STATUS
Smith
                           31-MAY-70
                                            21000 Med
John
                    03-APR-68 11000 Poor
                    31-MAY-70 60001 Rich
Wang
Now input of dates needs the use of to_date function in Oracle!
insert into person1 values ('Jones', to date('1976/02/29','yyyy/mm/dd'), 29000.00);
```

```
SQL> select * from person;
                        DOB
NAME
                                      INCOME STATUS
Smith
                        31-MAY-70 21000 Med
                        03-APR-68 11000 Poor
John
Wang
                        31-MAY-70
                                     60001 Rich
SQL> drop table person cascade constraints;
NOTE: Data in person1 is still not deleted
SQL> select * from person1;
NAME
                        DOB
                                      INCOME
Smith
                        31-MAY-70
                                      21000
John
                        03-APR-68 11000
                        31-MAY-70 60001
Wang
```

di Bipin C Desai

```
MariaDB [tempDB]> create table person (
name char(25) primary key,
income decimal(12,2),
status char (6));
                    Note: the MySQL trigger syntax is different
CREATE OR REPLACE TRIGGER person st
AFTER INSERT ON person
FOR EACH ROW
BEGIN
DECLARE STATUS CHAR(4);
set STATUS = 'Poor';
IF (new.income > 20000) THEN
set STATUS = 'Med'; END IF;
IF (new.income > 60000)THEN
set STATUS = 'Rich'; END IF;
INSERT INTO person VALUES(new.name, new.income, STATUS);
END
```

MariaDB [tempDB]>insert into person (name, income) values ('Smith', 10000.0);

ERROR 1442 (HY000): Can't update table 'person' in stored function/trigger because it is already used by statement which invoked this stored function/trigger.

Mutating Trigger in mariadb/mysql

de Bipin C Desai

```
Another way to get around mutation
```

```
SQL> create view personv as select * from person;
CREATE OR REPLACE TRIGGER personv st
INSTEAD OF INSERT ON PERSONV
                                   SQL> select * from person;
FOR EACH ROW
DECLARE STATUS CHAR (4);
                                  NAME DOB
                                                   INCOME STATUS
BEGIN
STATUS := 'Poor';
                                  Smith 31-MAY-70 21000 Med
IF (:new.income > 20000) THEN John 03-APR-68 11000 Poor STATUS := 'Med'; END IF; Wang 31-MAY-70 60001 Rich
STATUS := 'Med'; END IF;
IF (:new.income > 60000) THEN Black 13-MAY-45 120000 Rich
STATUS := 'Rich'; END IF;
INSERT INTO person VALUES(:new.name, :new.dob,
             :new.income, STATUS);
END person st;
run
Trigger created.
insert into personv values('Black', '13-may-45', 120000, 'Poor');
1 row created.
```

## SQL> select TRIGGER\_NAME from user\_triggers;

## TRIGGER\_NAME

-----

CSTUDENTINSERT

**ECTRIG** 

NUMBEROFPROJS

PERSON1\_ST

PERSONV\_ST

WHICHENGG

6 rows selected.

insert into personv (name, dob,income) values('Jones', '10-jun-68', 61000); 1 row created.

SQL>	select * from	person;
NAME	DOB	INCOME STATUS
Smith	31-MAY-70	21000 Med
John	03-APR-68	11000 Poor
Wang	31-MAY-70	60001 Rich
Black	13-MAY-45	120000 Rich
Jones	10-JUN-68	61000 Rich

Bipin C Desai

14

# **Another Mutating Trigger**

```
create table EmpName (
   eid number not null,
                             If an employee is deleted, his city
   Name varchar2(30),
                             is also deleted!
   primary key(eid));
create table EmpCity (
   eid number,
   City varchar2(15),
   foreign key (eid) references EmpName(eid) on
  delete cascade);
insert into EmpName values(1,'Smith');
insert into EmpName values(2,'Lee');
insert into EmpCity values(1,'Montreal');
insert into EmpCity values(2,'Laval');
commit;
```

```
create or replace trigger ECTrig
  after delete on EmpCity
  for each row
  declare
    n integer;
begin
select count(*) into n from EmpName;
dbms output.put line ('There are ' || n || 'rows in EmpName');
dbms output.put line('after cascade delete of EmpCity');
dbms output.new line;
                          set serveroutput on; To enable dbms output
end;
                          delete from EmpName where eid = 1;
                          ERROR at line 1:
run
                          ORA-04091: table SCOTT.EMPNAME Is mutating,
Trigger created.
                          trigger/function may not see it
                          ORA-06512: at "SCOTT.ECTRIG", line 4
                          ORA-04088: error during execution of trigger '
                          SCOTT ECTRIG'
```

d Bipin C Desai

```
Solution: Use statement trigger instead of row trigger
create or replace trigger ECTrig
  after delete on EmpCity
            n integer;
  declare
  begin
   select count(*) into n from EmpName;
   dbms output.put ('There are ' || n || ' rows in EmpName');
   dbms output.put line ('after cascade delete of EmpCity');
   dbms output.new line;
                                        There are 1 rows in EmpName
  end;
                                          after cascade delete of EmpCity
                                        1 row deleted.
                                       SQL> select * from EmpName;
run
                                       EID NAME
 set serveroutput on; To enable dbms_output
                                         2
                                             Lee
                                       SQL> select * from EmpCity;
 delete from EmpName where eid = 1;
                                       EID CITY
 1 row deleted.
                                            Laval
```

The mutating trigger error occurs due to the protocol used in Oracle to manage a **read consistent view** of data. (data read is of the same generation)

The error is occurs when a row-level trigger, while executing, accesses the table on which it is based.

The table is said to be mutating.

Mutation will not occur if a single record is inserted in the table (using VALUES clause).

If bulk insertion is done or data is inserted from another table mutation will occur.

The mutating error is not only encountered during queries, but also for insert, updates and deletes present in the trigger.

It is reported that newer release of the Oracle DBs (9i+) reduces the impact of the mutating triggers -but triggers still mutates.

Bipin C Desai

# Another example of Mutation

create table T (A number, B varchar2(10));

```
SQL> create or replace trigger Ttrg
                                            PLS INTEGER
 2 before insert or update or delete
3 on T
                                            PLS INTEGER instead of INTEGER or
                                            NUMBER for an efficient numeric
                                            datatype.
 4 for each row
 5 declare
                                            magnitude range for this datatype is –
                                            2147483647 through 2147483647.
    i pls integer;
 7 begin
                                            require less storage than INTEGER or
     select count(1)
                                            NUMBER values.
 9
     into i
                                            operations use faster machine arithmetic,
 10
     from T;
     dbms output.put line('Trigger success');
 11
     exception
12
      when no data found then
13
                                               count(1) and count(*)
      dbms output.put line('Error');
14
                                                returns the number
15 end;
                                                of rows
16 /
Trigger created.
```

```
SQL> insert into T values(1, 'ABD');
1 row created.
SQL> update T set A=2; Bulk Update
update T
ERROR at line 1:
ORA-04091: table SCOTT.T is mutating, trigger/function may not see it
DRA-06512: at "SCOTT.TTRG", line 4
DRA-04088: error during execution of trigger 'SCOTT.TTRG'
SQL> create table T1 (A number primary key, B varchar2(10));
SQL> insert into T1 values (1, 'ABC');
SQL> insert into T1 values (1, 'ABC');
SOL> insert into T select * from T1; Bulk Insert
insert into T select * from T1
                                T and T1 have the same schema
ERROR at line 1:
DRA-04091: table SCOTT.T is mutating, trigger/function may not see it
DRA-06512: at "SCOTT.TTRG", line 4
ORA-04088: error during execution of trigger 'SCOTT.TTRG'
```

Dipin C Desai

```
Sln: Statement level trigger
create or replace trigger Ttrg
before insert or update or delete on T
Declare i pls_integer;
begin
 select count(1) into I from T;
 dbms_output.put_line('Trigger success');
exception
 when no data found then
 dbms output.put line('Error');
end:
SQL> insert into T select * from T1;
2 rows created.
SQL> select * from T;
        1 ABD
        1 ABC
        2 BCD
```

```
Final example of Mutation
create table T1 (A number primary key, B varchar2(10));
create table T2 (A number, B varchar2(10),
foreign key (A) references T1 on delete cascade);
create or replace trigger T1trg
before insert or update or delete on T1
for each row
Declare i pls integer;
begin
select 1
into i
from T2
where A = :new.A;
dbms output.put line('Trigger success');
exception
when no data found then
dbms output.put line('Error: no data');
end:
```

(b) Bipin C Desai

```
SQL> insert into T1 values (1, 'ABC');
1 row created.
SQL> select * from T1;

SQL> select * from t2;
no rows selected

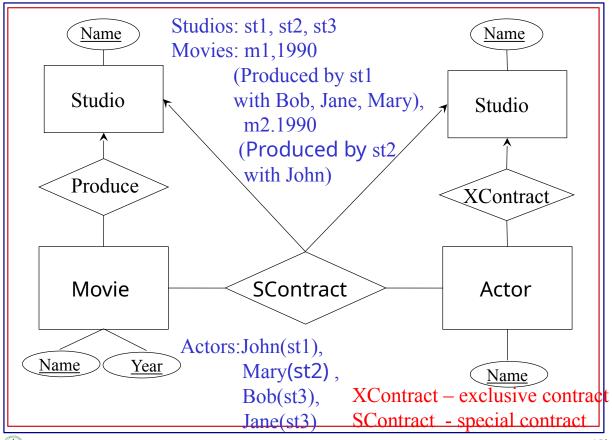
SQL> delete from t1;
delete from t1

*

ERROR at line 1:
ORA-04091: table SCOTT.T2 is mutating, trigger/function may not see it
ORA-06512: at "SCOTT.T1TRG", line 4
ORA-04088: error during execution of trigger 'SCOTT.T1TRG'
```

SQL> insert into T1 values (2, 'BCD'); 1 row created. SQL> insert into T2 values(1, 'XYZ'); 1 row created. SQL> insert into T2 values (2, 'WXY'); 1 row created. SQL> delete from T1 where A= 1; delete from T1 where A= 1 ERROR at line 1: ORA-04091: table SCOTT.T2 is mutating, trigger/function may not see it ORA-06512: at "SCOTT.T1TRG", line 4 ORA-04088: error during execution of trigger 'SCOTT.T1TRG'

Bipin C Desai



```
create table studio(
                                       create table actor(
name varchar2(12) primary key);
                                       name varchar2(12) primary key);
insert into studio values('st1');
                                       insert into actor values('John');
insert into studio values('st2');
                                       insert into actor values('Mary');
insert into studio values('st3');
                                       insert into actor values('Bob');
                                       insert into actor values('Jane');
create table movie(
name varchar2(12), year dec(4),
                                            insert into xcontract
primary key(name, year));
                                                    values('John', 'st1');
insert into movie values('m1',1990);
                                            insert into xcontract
insert into movie values('m2',1990);
                                                    values('Mary', 'st2');
                                            insert into xcontract
create table xcontract(
                                                    values('Bob', 'st3');
aname varchar2(12) primary key,
                                            insert into xcontract
 astudio varchar2(12) not null,
                                                    values('Jane', 'st3')
foreign key (aname) references actor(name),
foreign key (astudio) references studio(name));
```

ⓑ Bipin C Desai

```
create table produce(
sname varchar2(12) not null unique,
mname varchar2(12), myear dec(4),
primary key (mname,myear),
-- must give the above together as foreign key
foreign key (mname, myear) references movie(name, year),
foreign key (sname) references studio(name));
insert into produce values('st1', 'm1',1990);
insert into produce values('st2', 'm2',1990);
```

Ispin C Desai

```
create table scontract( -- special contract
mname varchar2(12), myear dec(4),
aname varchar2(12), pstudio varchar2(12),
astudio varchar2(12),
primary key(mname, myear, aname),
-- This is equivalent to a m-m-1-1 multiplicity
foreign key (mname, myear) references movie(name, year),
foreign key (aname) references actor(name),
foreign key (pstudio) references studio(name),
foreign key (astudio) references studio(name));
insert into scontract values('m1',1990,'Mary', 'st1', 'st2');
insert into scontract values('m1',1990,'Bob',
                                             'st1', 'st3');
insert into scontract values('m1',1990,'Jane', 'st1', 'st3');
insert into scontract values('m1',1990,'Mary', 'st2', 'st1');
* ERROR at line 1:
ORA-00001: unique constraint (SCOTT.SYS C004729) violated
```

Bipin C Desai

155

insert into scontract values('m2',1990,'John', 'st2', 'st1');

-- There is no check in consistency John has exclusive contract with st1

insert into scontract values('m2',1990,'Bob', 'st2', 'st3');

-- These is no check in consistency movie m2 in 1990 is made by st2 Bob is under contact to st3

insert into scontract values('m2',1990,'Jane', 'st2', 'st3');

-- These is no check in consistency movie m2 in 1990 is made by st2 Jane is under contact to st3

SQL> select * MNAME		eract; ANAME	PSTUDIO	ASTUDIO
m1	1990	Mary	st1	st2
m1	1990	Bob	st1	st3
m1	1990	Jane	st1	st3
m2	1990	John	st2	st1
m2	1990	Bob	st2	st3
m2	1990	Jane	st2	st3

```
SQL> select * from xcontract;
            ASTUDIO
Jane
                    Does not maintain the consistency
          st1
st2
John
                     John is under contract to st1 not st3
Mary
                     Jane is under contract to st3 not st1
Bob
           st3
SQL> delete from scontract;
6 rows deleted.
insert into scontract values('m1',1990,'Jane','st1','st1');
insert into scontract values('m2',1990,'John','st2','st3');
SQL> select * from scontract;
            MYEAR ANAME
                                            ASTUDIO
MNAME
                                PSTUDIO
              1990 Jane
m1
                                st1
                                            st1
m2
              1990 John
                           st2
                                            st3
```

db Bipin C Desai

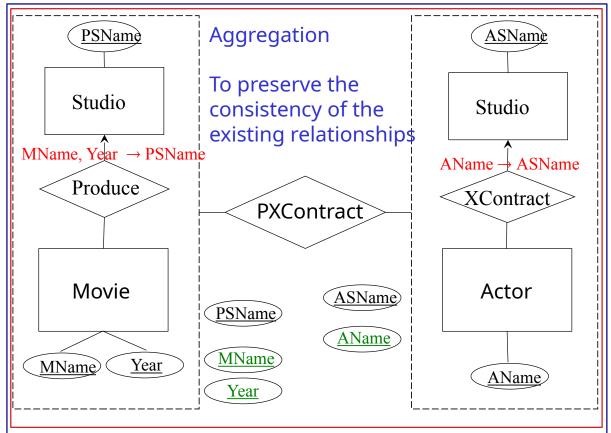
```
create table scontract1(
mname varchar2(12),
myear dec(4),
aname varchar2(12),
pstudio varchar2(12),
astudio varchar2(12),
primary key(mname, myear, aname, pstudio, astudio),
-- This is equivalent to a 1 to 1 to 1 multiplicity
foreign key (mname, myear) references movie(name, year),
foreign key (aname) references actor(name),
foreign key (pstudio) references studio(name),
foreign key (astudio) references studio(name));
```

insert into scontract1 values('m1',1990,'Mary', 'st1', 'st2'); insert into scontract1 values('m1',1990,'Bob', 'st1', 'st3'); insert into scontract1 values('m1',1990,'Jane', 'st1', 'st3'); insert into scontract1 values('m2',1990,'John', 'st2', 'st1'); insert into scontract1 values('m2',1990,'Bob', 'st2', 'st3'); insert into scontract1 values('m2',1990,'Jane', 'st2', 'st3');

insert into scontract1 values('m1',1990,'Mary', 'st3', 'st1');

- -- Allows the same movie to be made by different studio
- -- Allows the same actor to be under contract to >1 studio!

d Bipin C Desai



```
drop table pxcontract;
create table pxcontract(
mname varchar2(12),
year dec(4),
pstudio varchar2(12),
aname varchar2(12),
astudio varchar2(12),
primary key(pstudio, mname, year),
foreign key (mname, year) references produce(mname, myear),
foreign key (aname) references xcontract(aname));
SQL> insert into pxcontract values('m1',1990, 'st1', 'Mary', 'st2');
1 row created.
SQL> insert into pxcontract values('m1',1990, 'st1', 'Bob', 'st3');
* ERROR at line 1:
ORA-00001: unique constraint (SCOTT.SYS_C004747) violated
```

(b) Bipin C Desai

### Use triggers to maintain consistency

```
drop table pxcontract;
create table pxcontract(
mname varchar2(12),
year dec(4),
aname varchar2(12),
primary key( mname, year, aname),
foreign key (mname,year)
    references produce(mname,myear),
foreign key (aname) references xcontract(aname));
```

```
USE TRIGGER for consistency
```

create view vcontract as select \* from scontract;

CREATE OR REPLACE TRIGGER SP\_Trig

Instead of INSERT ON vcontract

FOR EACH ROW

Declare pstudio varchar2(12); astudio varchar2(12);

begin

select p.sname into pstudio from produce p

where :new.mname=p.mname

And :new.myear=p.myear;

select x.astudio into astudio from xcontract x

where :new.aname=x.aname;

**INSERT INTO scontract values** 

(:new.mname,:new.myear, :new.aname, pstudio, astudio);

END SP Trig;

•

Bipin C Desai

163

insert into vcontract values('m1',1990,'Mary', 'st3', 'st1'); 1 row created.

SQL> select \* from scontract;

 MNAME
 MYEAR ANAME
 PSTUDIO
 ASTUDIO

 m1
 1990 Mary
 st1
 st2

insert into vcontract values('m1',1990,'Mary', 'st1', 'st2');

#### \* ERROR at line 1:

ORA-00001: unique constraint (SCOTT.SYS\_C004729) violated

ORA-06512: at "SCOTT.SP\_TRIG", line 8

ORA-04088: error during execution of trigger 'SCOTT.SP\_TRIG'

insert into vcontract values('m1',1990,'Bob', 'st1', 'st3'); SQL> select \* from scontract;

MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
m1	1990	Mary	st1	st2
m1	1990	Bob	st1	st3

insert into vcontract values('m1',1990,'Jane', 'st1', 'st3'); 1 row created.

SQL> select \* from scontract;

MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
m1	1990	Bob	st1	st2
m1	1990		st1	st3
m1	1990		st1	st3

Bipin C Desai

insert into vcontract values('m2',1990,'John', 'st2', 'st1');

SQL> select *	from sconti	act;
MNAME	MYEAR	ANAI

MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
m1	1990	Mary	st1	st2
m1	1990	Bob	st1	st3
m1	1990	Jane	st1	st3
m2	1990	John	st2	st1

insert into vcontract values('m2',1990,'Bob', 'st2', 'st3'); SQL> select \* from scontract;

MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
m1		Mary	st1	st2
m1 m1	1990 1990	Bob Jane	st1 st1	st3 st3
m2		John	st2	st1
m2	1990	Вор	st2	st3

insert into vcontract values('m2',1990,'Jane', 'st2', 'st3');

# SQL> select \* from scontract;

MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
m1	1990	Mary	st1	st2
m1	1990	Bob	st1	st3
m1	1990	Jane	st1	st3
m2	1990	John	st2	st1
m2	1990	Bob	st2	st3
m2	1990	Jane	st2	st3

SQL> insert into vcontract values('m1',1990,'Mary', 'st3', 'st1') \*ERROR at line 1:

ORA-00001: unique constraint (SCOTT.SYS C004729) violated

ORA-06512: at "SCOTT.SP\_TRIG", line 8

ORA-04088: error during execution of trigger 'SCOTT.SP TRIG'

Bipin C Desai

167

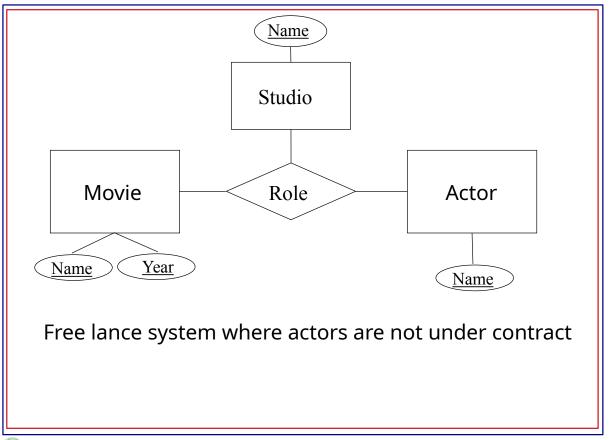
l	MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
l					
l	m1	1990	Mary	st1	st2
l	m1	1990	Bob	st1	st3
l	m1 m1 m2 m2 m2	1990	Jane	st1	st3
l	m2	1990	John	st2	st1
l	m2	1990	Bob	st2	st3
l	m2	1990	Jane	st2	st3
ı					

### 6 rows selected.

# SQL> select \* from vcontract;

MNAME	MYEAR	ANAME	PSTUDIO	ASTUDIO
m1	1990	Mary	st1	st2
m1	1990	Bob	st1	st3
m1	1990	Jane	st1	st3
m1 m1 m1 m2 m2 m2	1990	John	st2	st1
m2	1990	Bob	st2	st3
m2	1990	Jane	st2	st3
6 rows selected.				

(d) Bipin C Desai



Bipin C Desai

```
DATES:
```

# How to specify the beginning weekday of the week

select to\_char(trunc(sysdate,'DAY'),'fmDay"," Month DD"," YYYY')

AS First\_week\_day from dual;

FIRST WEEK DAY

-----

Sunday, November 17, 2002

### If we want Monday to be the beginning of the week:

SQL> alter session set nls\_territory=FRANCE;

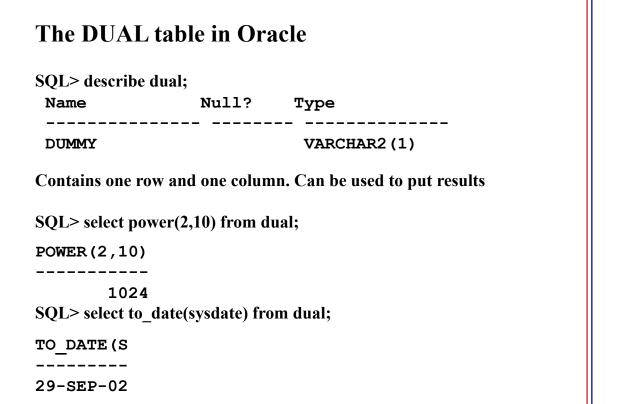
Session altered.

SQL> select to\_char(trunc(sysdate,'DAY'),'fmDay"," Month DD", "YYYY') AS First week day from dual;

FIRST\_WEEK\_DAY

-----

Monday, November 18, 2002



ⓑ Bipin C Desai

**TRUNCate function and dates** truncate to first day of week select to char(trunc(sysdate,'DAY')) as FirstDayofWeek from dual; FIRSTDAYO select to char(trunc(sysdate)) from dual; TO\_CHĀR(T 14-MAR-04 16-MAR-04 Date of query (a Tuesday) select to char(trunc(sysdate,'DAY'),'fmDay') as FirstDay from dual; FIRSTDAY Format fully after truncating to first day of week, month, year Sunday select to char(trunc(sysdate,'MONTH'),'fmMonth') as Month from dual MONTH select to char(trunc(sysdate,'MONTH')) from dual; TO\_CHĀR(T March 01-MAR-04 select to\_char(trunc(sysdate),'fmYear') as year from dual;

VEAD select to\_char(trunc(sysdate,'YEAR')) from dual; YEAR TO\_CHĀR(T

Bipin C Desai

Two Thousand Four

173

01-JAN-04

WEEKDAY
-----SUNDAY

To find the first business day of the week for a particular date:

SQL> select to\_char(trunc((select DOB from person where Name='Smith'),'DAY'), 'fmDay') from dual;

Monday

To find the first business day of the week two days from a particular date:

SQL> select to\_char(trunc((select DOB from person where Name='Smith'),'DAY') + 2, 'fmDay') from dual;

Wednesday

```
To find the day of the week for a specified date:

SQL> select to_char(to_date('18-Nov-02'), 'Day') As Weekday
from dual;

WEEKDAY
------
Monday

Find the first business day after the birthday of Smith:

select to_char(trunc((select DOB from person
where Name='Smith'),
'DAY') +2, 'fmDay') As
Smith_Bday2busWeek from dual;

SmithQ

Finding where a date is(half, quarter)

select TO_NUMBER(TO_CHAR((select DOB from person
where Name='Smith'), 'Q')) as SmithQ from dual;
```

d Bipin C Desai

### **DECODE** function

```
DECODE(expression, if1, then1, if2, then2, ..., ifn,thenn, else)
```

```
Create table duedate(
                                    assign# number (2) primary key,
Create tables assignment(
                                    duedate date):
sid number(7),
                                    Insert into duedate
assign# number(2),
                                       values(1, '17-jan-2004');
submitdate date
                                    insert into duedate
primary key (sid, assign#))
                                       values(2, '14-feb-2004');
insert into assignment values (123, 1, '17-jan-2004');
insert into assignment values (124, 1, '18-jan-2004');
insert into assignment values (125, 1, '19-jan-2004');
insert into assignment values (123, 2, '17-feb-2004');
insert into assignment values (124, 2, '18-feb-2004');
insert into assignment values (125, 2, '19-feb-2004');
```

```
select a.sid, a.assigno,
 decode(trunc(a.submitdate - d.duedate), 0, null, 1, 'one day',
 2, 'two days', 3, 'three days', 4, 'four days', 'Too late') as Late days
from assignment a, duedate d
where a.assigno = d.assigno
order by sid, assigno;
           ASSIGNO LATE DAYS
 SID
   123
                 1
   123
                2
                       three days
   124
                       one day
                1
                2
                       four days
   124
   125
                1
                       two days
                2
   125
                       Too late
```

Dipin C Desai

```
To calculate the number of business days between two days: store the
following code in a file say: numbusdays.sql
define frdate = '&1'
define todate = '&2' FROM_DATE TO_DATE BUSINESS_DAY
set verify off
                  20-Nov-02
                                  24-Dec-02
                                                  25
select
 '&frdate' From Date, '&todate' To Date,
  1 + to_date('&todate') - to_date('&frdate') - How many weeends?
                                              a saturday?
  ((TRUNC(to_date('&todate'),'D') -
                                                 a sunday?
         TRUNC(to_date('&frdate'),'D'))/7)*2
  + DECODE(to char(to date('&todate'),'D'),7,-1,0)
  + DECODE(to char(to date('&frdate'),'D'),1,-1,0) Business Days
from dual
                   DECODE (exp, if, then, else, ..)
Then one can interactively call it:
SQL> @numbusdays 20-Nov-02 24-Dec-02
```

```
SQL> create table interval (startdate char(10), enddate char(10));
insert into interval values('1998.04.11','1998.09.30');
insert into interval values('1998.04.15','1998.10.01');
insert into interval values('1998.05.11','1998.06.17');
insert into interval values('1998.06.14','1998.10.12');
```

SQL> SELECT STARTDATE, ENDDATE FROM INTERVAL WHERE TO DATE('1998.04.17','YYYYY.MM.DD') BETWEEN TO DATE(STARTDATE, 'YYYY.MM.DD') AND TO DATE(ENDDATE, 'YYYY.MM.DD');

STARTDATE ENDDATE

1998.04.11 1998.09.30

1998.04.15 1998.10.01

Bipin C Desai

### If the interval is stored as dates:

SQL> create table intervaldate (startdate date, enddate date);

SQL> insert into intervaldate

select TO DATE(startdate, 'YYYY.MM.DD'),

TO DATE(enddate, 'YYYY.MM.DD')

from interval:

# SQL> select \* from intervaldate; SQL> select startdate, enddate+1 from intervaldate;

STARTDATE	ENDDATE	STARTDATE	
		JIANIDAIE	ENDDAIETI
15-APR-98	30-SEP-98 01-OCT-98		01-OCT-98 02-OCT-98
	17-JUN-98 12-OCT-98	11-MAY-98	18-JUN-98 13-OCT-98

### **SQL> SELECT startdate**, enddate FROM intervaldate WHERE TO DATE('1998.07.03','YYYYY.MM.DD') **BETWEEN** startdate AND enddate:

STARTDATE ENDDATE

11-APR-98 30-SEP-98

15-APR-98 01-OCT-98

14-JUN-98 12-OCT-98

SQL> select TO CHAR(startdate, 'YYYY-MM-DD:HH:MI:SS') as starttime from intervaldate;

STARTTIME

1998-04-11:12:00:00

1998-04-15:12:00:00

1998-05-11:12:00:00

1998-06-14:12:00:00

Bipin C Desai

# SQL> select TO CHAR(startdate+

8/24 + 13/1440 + 12/86400'YYYY-MM-DD:HH:MI:SS') as NewStartTime

from intervaldate;

NEWSTARTTIME

1998-04-11:08:13:12

1998-04-15:08:13:12

1998-05-11:08:13:12

1998-06-14:08:13:12 ENDOFMONTH

Last day of month:

**SQL>select LAST\_DAY(enddate)** 30-SEP-98

as EndofMonth 31-OCT-98 from intervaldate; 30-JUN-98

31-OCT-98

### SQL> select Name, Trunc(MONTHS\_BETWEEN(Sysdate, DOB)/12) as Age from person;

NAME	AGE
Smith	32
John	34
Wang	32

Age in five years?

select Name,

Trunc(MONTHS\_BETWEEN(ADD\_MONTHS(Sysdate,60), DOB)/12) as Age from person;

NAME	AGE
Smith	37
John	39
Wang	37

Bipin C Desai

183

```
create or replace function time between (start tm in date, end tm in date,
 hours_only varchar2 default 'N') return varchar2 as
-- If "hours only" is null or "N", the return will be a string formatted like:
-- 2 days, 3 hrs, 5 mins, 10 secs
-- If "hours_only" is not "N", then the return is a value in hours, like 102.325
ret_val varchar2(80);
start_sec number;
                                                   if upper(hours_only) = 'N' then
                                                       if full_sec > 3599 then
 full_sec number;
                                                         hours := floor(full_sec / 3600);
balance
          number;
                                                         balance := mod(full sec, 3600);
 minutes
                                                         full sec := balance;
hours
          number;
days
                                                        if hours > 1 then
                                                           ret val := ret val || to char(hours) ||' hrs, ';
 function get_sec (time_in in date) return number as
                                                        else
begin
                                                          ret val := ret val || to char(hours) ||' hr, ';
   return to_number(to_char(time_in,'SSSSS'));
                                                        end if;
end;
                                                      end if;
                                                      if full_sec > 59 then
                                                       minutes := floor(full_sec / 60);
 start sec := get sec(start tm);
                                                         balance := mod(full sec, 60);
 end_sec := get_sec(end_tm);
-- check if end time is in the same day as start time full_sec := balance; if to_char(start_tm,'YYMMDD') = if minutes > 1 then
       to_char(end_tm,'YYMMDD') then
                                                           ret_val := ret_val||to_char(minutes)||' mins, ';
   full_sec := end_sec - start_sec;
                                                        else
   days := 0;
                                                          ret_val := ret_val||to_char(minutes)||' min, ';
                                                        end if;
  days := trunc(end_tm - start_tm);
                                                      end if;
  if days > 0 then
                                                       ret_val := ret_val||to_char(full_sec)||' secs';
    ret_val := to_char(days)||' days, ';
   end if;
                                                     -- Calculate the time difference in hours,
   if end_sec > start_sec then
                                                      -- to three decimal places
    full_sec := end_sec - start_sec;
                                                      ret_val := to_char((24 * days) + round((full_sec / 3600),3));
    full_sec := 86400 - start_sec + end_sec;
                                                    end if;
  end if;
                                                    return ret val;
 end if;
                                                   end;
                                                   grant execute on time_between to public;
                                                   create public synonym time_between for time_between;
```

select Name, time\_between(dob, SysDate, 'N') AS Age from person where name='Smith';

NAME AGE
-----Smith 11866 days, 10 hrs, 38 mins, 14 secs

The Bipin C Desai

# Oracle Editing SQL Buffer

Command append txt change /old/new/ change /txt clear buffer del get file input input txt list	abbrev. a text c /old/new/ c /txt cl buff  i iI txt l	Oper. on crnt. line/all lines adds text at the end of a line change old to new in a line delete text from a line delete all lines in the buffer delete a line load file into buffer add one or more lines add text as a line list all lines of buffer
<b>list</b> n	l n (n)	list line n and make it current
list *	*	list crnt. Line
<b>list</b> last	l last	list last line
<b>list</b> m n	l m n	list lines m – n
<b>save</b> file	sav file	save buffer to file

```
SQL> desc user_catalog;
Name Null? Type

TABLE_NAME NOT NULL VARCHAR2 (30)
TABLE_TYPE VARCHAR2 (11)

SQL> select * from cat;
TABLE_NAME TABLE_TYPE user_catalog

ASSIGNED_TO TABLE
BONUS TABLE
COURSE TABLE
CRS_SECTION TABLE
DEPT TABLE
DEPT TABLE
DEPTMAJOR TABLE
DISTANCE TABLE
DO PROJECT VIEW
DO PROJ_SUP VIEW
DUMMY TABLE
EMAIL_INFO TABLE
```

IBipin C Desai

```
SQL> desc assigned_to;
Name Null? Type
                            NUMBER (4)
 PROJ#
 EMP#
                               NUMBER (4)
 • • • • •
SQL> select table name from user tables;
TABLE NAME
ASSIGNED TO
BONUS
COURSE
SQL> select TABLESPACE NAME from user tables;
TABLESPACE NAME
SYSTEM
SYSTEM
TUTOR
TUTOR ......
```

SQL> desc user_tables; Name	Null? Type
TABLE_NAME TABLESPACE_NAME CLUSTER_NAME IOT_NAME etc. SQL> desc user tab colum	NOT NULL VARCHAR2 (30)  VARCHAR2 (30)  VARCHAR2 (30)  VARCHAR2 (30)
Name	Null? Type
TABLE_NAME COLUMN_NAME DATA_TYPE DATA_TYPE_MOD DATA_TYPE_OWNER etc.	NOT NULL VARCHAR2 (30) NOT NULL VARCHAR2 (30) VARCHAR2 (106) VARCHAR2 (3) VARCHAR2 (30)

(d) Bipin C Desai

```
SQL> desc user_views;
                       Null?
                                       Type
 Name
 VIEW NAME
                   NOT NULL
                                   VARCHAR2 (30)
 TEXT—LENGTH
                                   NUMBER
 TEXT
                                    LONG
 TYPE TEXT LENGTH
                                   NUMBER
                                   VARCHAR2 (4000)
 TYPE TEXT
 OID TEXT LENGTH
                                   NUMBER
                                   VARCHAR2 (4000)
VARCHAR2 (30)
 OID TEXT
VIEW TYPE OWNER
 VIEW TYPE
                                   VARCHAR2 (30)
 SUPERVIEW NAME
SQL> select view_name from user_views;
VIEW NAME
DO PROJECT
DO PROJ SUP
TE\overline{M}P1
```

SQL> desc user_triggers;						
Name	Null?	Type				
TRIGGER_NAME		VARCHAR2(30)				
TRIGGER_TYPE		VARCHAR2 (16)				
TRIGGERING_EVENT		VARCHAR2 (227)				
TABLE OWNER		VARCHAR2(30)				
BASE_OBJECT_TYPE		VARCHAR2(16)				
TABLE NAME		VARCHAR2(30)				
COLUMN NAME		VARCHAR2 (4000)				
REFERENCING_NAMES		VARCHAR2 (128)				
WHEN CLAUSE		VARCHAR2 (4000)				
STATUS		VARCHAR2(8)				
DESCRIPTION		VARCHAR2 (4000)				
ACTION TYPE		VARCHAR2 (11)				
TRIGGER_BODY		LONG				

(h) Bipin C Desai

```
SQL> select TRIGGER_NAME from user_triggers;

TRIGGER_NAME

EMP_SAL_RAISE
PERSON1_ST

select TRIGGER_NAME, TRIGGER_TYPE, TRIGGERING_EVENT, TABLE_OWNER from user_triggers
where TRIGGER_NAME='PERSON1_ST';

TRIGGER_NAME TRIGGER_TYPE TRIGGERING_EVENT TABLE_OWNER

PERSON1_ST AFTER EACH ROW INSERT SCOTT
```

```
SQL> SET PAGESIZE 66
SQL> COLUMN object type FORMAT A20
SQL> COLUMN object name FORMAT A30
SQL> COLUMN status FORMAT A10
SQL> BREAK ON object type SKIP 1
SQL> SELECT object type, object name, status
 FROM user objects
 WHERE object type IN ('PACKAGE', 'PACKAGE BODY',
           'FUNCTION', 'PROCEDURE',
           'TYPE','TYPE BODY',
                'TRIGGER');
OBJECT_TYPE OBJECT_NAME STATUS
                                    VALID
FUNCTION
              BUSINESS DAYS
                 PERSON1 ST
TRIGGER
                                     VALID
```

h Bipin C Desai

```
select text
from user source
where name='WORKING DAYS';
TEXT
FUNCTION working days(date1 IN DATE, date2 IN DATE)
RETURN NUMBER IS workdays NUMBER;
BEGIN
workdays := TRUNC(date2) - TRUNC(date1) + 1
    - ((TRUNC(to date(date2,'D'))-TRUNC(to date(date1),'D'))/7)*2;
 IF TO CHAR(date2,'D') = '7' THEN
  workdays := workdays - 1;
 END IF:
 IF TO CHAR(date1,'D') = '1' THEN
  workdays := workdays - 1;
 END IF;
 RETURN(workdays);
end;
```

```
select text
from user source
where name=PERSONV ST';
TEXT
TRIGGER personv st
INSTEAD OF INSERT ON PERSONV
FOR EACH ROW
DECLARE STATUS CHAR(4);
BEGIN
STATUS := 'Poor';
IF (:new.income > 20000) THEN
STATUS := 'Med'; END IF;
IF (:new.income > 60000)THEN
STATUS := 'Rich'; END IF;
INSERT INTO person VALUES(:new.name, :new.dob, :new.income, STATUS);
END person st;
```

ⓑ Bipin C Desai

```
select object name
       from user_procedures;
                             select text
       OBJECT_NAME
                             from user_source
       WORKING_DAYS
                             where name='WORKING DAYS';
TEXT
FUNCTION working_days(date1 IN DATE, date2 IN DATE)
RETURN NUMBER IS workdays NUMBER;
BEGIN
workdays := TRUNC(date2) - TRUNC(date1) + 1
((TRUNC(to_date(date2,'D'))-TRUNC(to_date(date1),'D'))/7)*2;
 IF TO CHAR(date2,'D') = '7' THEN
 workdays := workdays - 1;
 END IF;
IF TO_CHAR(date1,'D') = '1' THEN
 workdays := workdays - 1;
 END IF;
 RETURN(workdays);
end;
```

# Relation Algebra, Bags and Constraints

### **Notes**



To be used in the spirit of copy-forward! Pl. see: https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

🕩 Bipin C Desai

### **BAGS**

In a set, there are no duplicates.

A set (collection of similar objects) having multiple occurrences of one or more members is called a "bag".

Implementation of relational model allow a relation(table) to have duplicates.

This is specially so for intermediate results(a convenience) and if no primary keys are defined for a table.

Thus in a relation that is a bag, duplicate tuples are allowed(though not required – so a bag may have no duplicates at a given point in time.)

A stored table is not a bag; since constraints enforcements require that duplicates are not stored (each row has an UNIQUE row identifier)

	A	В	С		В	C	
Instance	a1	<b>b</b> 1	c1	П <sub>вс</sub> R without	b1	c1	
R which is not a bag:	a1	<b>b2</b>	c2	eliminating the	b2	c2	
not a bag.	a3	b1	c1	duplicates tuples	b1	c1	
	a4	<b>b2</b>	c2	is a bag!	b2	c2	
	a5	<b>b5</b>	<b>c5</b>		1	-	
			•	,	b5	c5	

**Faster projection operations:** no need to check duplicates for each tuple in the output relation

**Correct computation** compute the *average* of values under attribute B in the projection of R.

**Faster bag Unions:** Computing ( $\mathbb{R} \cup_{\text{Bag}} \mathbb{S}$ ) is cheaper than ( $\mathbb{R} \cup_{\text{Set}} \mathbb{S}$ ). Given R has n tuples and S has m tuples, then the costs of evaluating these queries would be  $\mathbb{O}(n+m)$  and  $\mathbb{O}(n_*m)$ , respectively.

$$R \cup_{Bag} S \equiv R \cup_{B} S$$
  $R \cup_{Set} S \equiv R \cup_{S} S \equiv R \cup S$ 

Bipin C Desai

3

 $R \cup_B S$  (**bag union** of R and S): the bag of tuples that are in R, in S, or in both. If a tuple t appears n times in R, and m times in s, then t appears n+m times in bag  $R \cup_B S$ 

$$R \cup_B S = \{ t:k \mid t:n \in R \land t:m \in S \land k = n+m \}$$

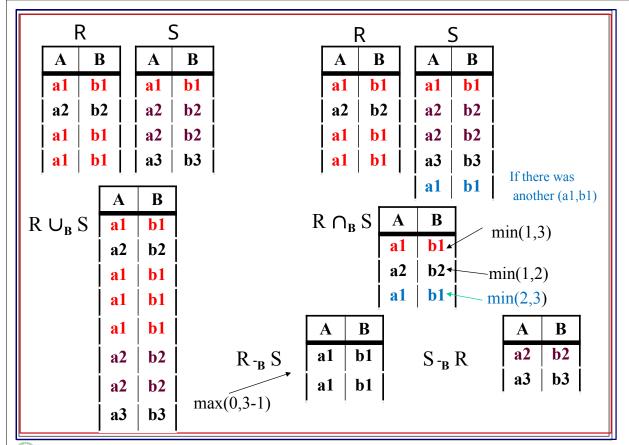
 $R \cap_B S(bag intersection of R and S):$  the bag of tuples that appear in both R and S. If a tuple t appears n times in R, and m times in S, then the number of occurrences of t in bag  $R \cap_B s$  is min(n,m)

$$R \cap_B S = \{ t:k \mid t:n \in R \land t:m \in S \land k = min(n,m) \}$$

 $R -_B S$  (bag difference of R and S): is defined as follows:

$$R -_{B} S = \{ t:k \mid t:n \in R \wedge t:m \in S \wedge k = max(0, n-m) \}$$

$$S -_B R = \{t:k \mid t:n \in R \land t:m \in S \land k = max(0, m-n) \}$$



Bipin C Desai

5

### **BAG PROJECTION**

Let  $\mathbf{R}$  be a relation scheme, and  $\mathbf{R}$  be a bag over  $\mathbf{R}$ .

The **projection** operator is used to produce, from R, a new bag that has only some of **R**'s columns

If elimination of one or more attributes during the projection causes the same tuple to be created from several tuples, these **multiple tuples are not eliminated** from the result of a bag-projection

A	В	C
a1	b1	c1
a1	<b>b2</b>	c2
a3	<b>b</b> 1	c1
a4	<b>b2</b>	c2
a5	<b>b</b> 5	<b>c</b> 5

В	C
b1	c1
b2	c2
b1	c1
b2	c2
b5	c5

Bipin C Desai

6

### BAG SELECTION σ<sub>C</sub>R

The tuples in the output relation are those that satisfy the predicate C, which involves attributes of R

Duplicates are eliminated in a set selection but **not so** from the result of a bag-selection

**Note**: The selection operation  $\sigma$  in RA is not the same as the **SELECT** clause in SQL which is the projection part of the DML component of SQL

A	В	C	lι	A	В	C
a1	b1	c1	1	a1	b1	c1
a1	<b>b</b> 1	c2	l	a1	b1	c2
а3	<b>b</b> 1	c1	İ	a3	b1	c1
a4	<b>b2</b>	c2	i !			
a5	b5	<b>c</b> 5	İ	σ	B=b1	(K):
R - not a bag						

			-			
A	В	C	1			
a1	b1	c1		A	В	C
	b1	c2		a1	b1	c1
a1	!			a1	<b>b</b> 1	c2
a1	<b>b1</b>	c1				: :
a1	<b>b2</b>	c2		a1	b1	c1
a5	!	c5		$\sigma_{l}$	B=b1	S):
<b>S</b> - a bag				1	A bag	

🕩 Bipin C Desai

7

# **Cartesian Product of Bags**

Given R and S, then  $R \times_B S$  is the bag of tuples formed by concatenating pairs of tuples, the first of which comes from R and the second from S.

$$\mathbf{R} \times_{\mathbf{B}} \mathbf{S} = \{ t_{I} \cdot t_{2} | t_{I} \in \mathbf{R} \land t_{2} \in \mathbf{S} \}$$

As in the set cartesion product, each tuple of R is paired with each tuple of S: however, in bag product each tuple is used regardless of whether it is a duplicate or not.

Hence, if a tuple  $t_1$  appears m times in a relation R, and a tuple  $t_2$  appears n times in relation S, then tuple  $t_1 \cdot t_2$  appears  $m \cdot n$  times in the bag-product R  $\times_B$  S

# Joins of Bags $\bowtie$ B(predicate)

The join of bags  $R \bowtie_{B(predicate)} S$  is computed in the same way as the join of sets; however, duplicates are not eliminated!

A   a1   a1	B b1 b1 b1	bi bi	2 c2 3 c3	A   B   B   C     a1   b1   b1   c2     a1   b1   b3   c3     R   T
A	R.B	S.B	C	$\begin{array}{ c c c c } \hline A & B & C \\ \hline \end{array}$
a1	b1	<b>b2</b>	c2	a1 b1 c2
a1	<b>b</b> 1	<b>b3</b>	c3	a1 b1 c2
a1	<b>b1</b>	<b>b3</b>	<b>c3</b>	!!!!
a1	<b>b1</b>	<b>b2</b>	<b>c2</b>	j
<b>a1</b>	<b>b1</b>	<b>b3</b>	c3	$R\bowtie_{_{B(R.B=S.B)}} T$
a1	<b>b1</b>	<b>b3</b>	<b>c3</b>	
ı	R×	$\mathcal{L}_{\mathrm{B}}$ S		•

Bipin C Desai

(

### **Constraints on Relations**

Relational algebra offers a convenient way to express a wide variety of constraints, such as referential integrity and FD's.

There are **two ways** to express constraints in RA

If r is a relational algebraic expression, then  $r = \Phi$  is a constraint that says "the value of r must be empty"

If r and s are relational algebraic expressions, then  $r \subseteq s$  is a constraint that says "every tuple in the result of r is in the result of s" (even when r and s are bags)

A RA constraint may be expressed in more than one way; i.e.

 $r \subseteq s$  could be written as  $r - s = \Phi$ 

If  $\mathbf{r} \subseteq \mathbf{s}$ ,  $\therefore$  no tuple in  $\mathbf{r}$  that is not in  $\mathbf{s}$ , and hence  $\mathbf{r} - \mathbf{s} = \mathbf{\Phi}$ The constraint  $\mathbf{r} = \mathbf{\Phi}$  could be rewritten as:  $\mathbf{r} \subseteq \mathbf{\Phi}$ 

# Referential integrity

If we have a value v in a tuple t of a relation R, then v must also appears as a component of some tuple s of relation S

**Example:** if we have a tuple (s,c,g) in relation **Enrol**(sid,crsno,grade), then there must be a student with sid = s and a course with crsno = c such that s has taken/taking c.

The values s and c in **Enrol** are "referring" to some values outside this relation, and these values **must** exist in the **Student** and **Course** relations **Course**(crsno, name, credits)

 $\pi_{crsno}$  **Enrol**  $\subseteq \pi_{crsno}$  **Course** or equivalently

 $\pi_{\mathit{crsno}}$  Enrol -  $\pi_{\mathit{crsno}}$  Course =  $\Phi$ 

de Bipin C Desai

# **Functional Dependency**

**Definition**: If two tuples of a relation R agree on the attributes X, then they must also agree on the attributes Y.

**Student**(sid, name, dob, gender),  $sid \rightarrow name$ 

To express the FD:  $sid \rightarrow name$  in RA, construct **pairs of Student** tuples, using **Cartesian product**, and see if there is a violation of this FD, using selection with sids equal but names not.

To assert the constraint, we equate the result must be null.

 $\rho(S, \rho(S1, Student) \times \rho(S2, Student))$ 

 $\sigma_{S1.sid=S2.sid \ ^S1.name \ \neq S2.name} S = \Phi$ 

### **Domain Constraints**

**Empl** (*Empl*#, name,dob, gender, salary)

• To express the domain constraint:

The only valid values for the attribute **gender** are 'F' and 'M'

```
\sigma_{\text{gender} \neq \text{'F'} \text{ AND gender} \neq \text{'M'}}, (Empl) = \Phi
```

• To express the following constraint?

Maximum salary of every employee is \$30,000

$$\sigma_{salary>30000}(Empl) = \Phi$$

These are examples of domain constraints

Bipin C Desai

13

# Bags in DBMS

```
create table dept(
dcode number(3),
dname varchar2(30),
location varchar2(30))

/

insert into dept values (100, 'CS', 'EV300');
insert into dept values (100, 'CS', 'EV300');
insert into dept values (100, 'CS', 'EV300');
/
```

SQL> select * from dept; DCODE DNAME			LOCATION			
100	CS		EV300			
100	CS		EV300			
100	CS		EV300			
100	CS		EV300			
Without a primary key ORACLE/MySQL/MariaDB ALLOWS DUPLICATES!						
SQL> desc de	ept;					
Name		Null?	Type			
DCODE			NUMBER(3)			
DNAME			VARCHAR2(30)			
LOCATION			VARCHAR2(30)			

15

Bipin C Desai

SQL> alter table dept add(constraint pk\_const primary key(dcode)); ERROR at line 1: ORA-02437: cannot validate (SCOTT.PK\_CONST) primary key violated Remove duplicate records delete from dept SQL> select \* from dept; where rowid in ( select rowid DCODE DNAME LOCATION from dept 100 CS EV300 minus select max(rowid) from dept d group by d.dcode); 3 rows deleted.

```
SQL> alter table dept add(constraint pk_const primary key(dcode) );

Table altered.

SQL> desc dept;

Name Null? Type

DCODE NOT NULL NUMBER(3)

DNAME VARCHAR2(30)

LOCATION VARCHAR2(30)
```

Bipin C Desai

```
How about MySql:
create table dept(
                              No primary key
dcode numeric(3),
dname varchar(30),
location varchar(30));
insert into dept values (100, 'CS', 'EV300');
insert into dept values (100, 'CS', 'EV300');
insert into dept values (100, 'CS', 'EV300');
mysql> select * from dept;
| dcode | dname | location |
     100 | CS | EV300
    100 | CS
                  | EV300
    100 | CS | EV300
MYSQL ALLOWS DUPLICATES
```

# **Eliminate Duplicates rows**

ⓑ Bipin C Desai

Past & Future(the past re-dressed)!

**Notes** 

Bipin C. DESAI



To be used in the spirit of copy-forward! https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

🕩 Bipin C Desai

Early methods of managing data: did not use SQL!

Definition of data was locked in application programs Data in compact form on external media punched cards (10x80), punched paper tape, magnetic tapes





Computer Museum

## **Hierarchical Data Model (HDM)**

The HDM introduced in 1960s by IBM in IMS -it is based on the parent-child model.

Parent record type, children record types

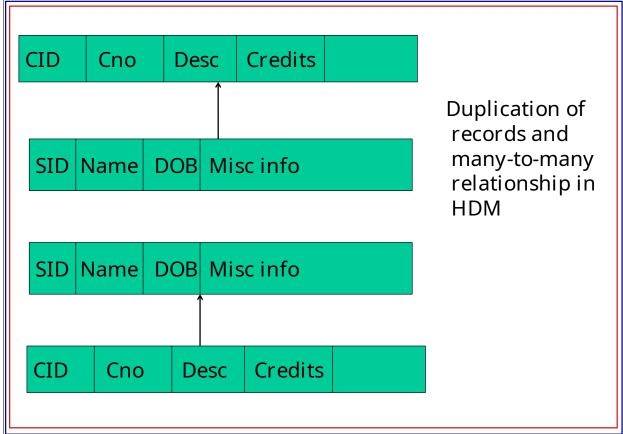
These record types are organized in the form of a rooted tree (hence, no cycles).

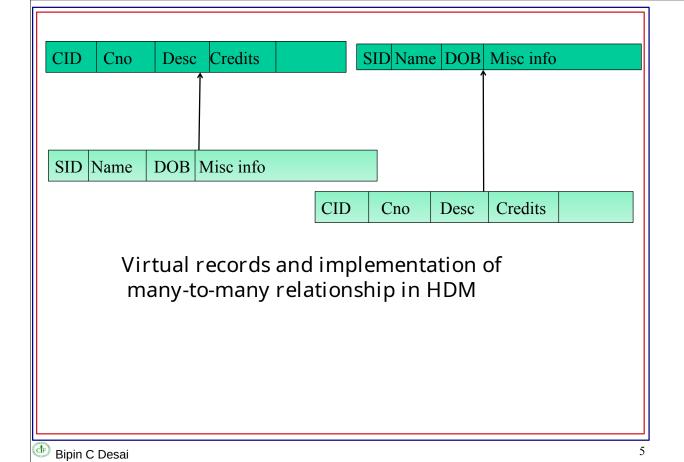
Only one-to-many or one-to-one relationship can be represented Many-to-many relationship requires duplication



Did not use SQL - but not a NoSQL DB!!

Bipin C Desai





HDM databases has features for 'fixed pattern" usage: Pros:

Since the structure is fixed data access faster HDM the relationships is fixed - easy for data integrity HDM represent many data in normally used

– all banking transactions belong to a given account Queries are predictable and uses the hierarchy

### Cons:

Duplication for many-to-many relationships; Difficult to change schema -

IMS currently is in version 15+ and is available for z/OS offers SQL for IMS!

Bipin C Desai

6

### **Network Data model (DBTG)**

In the NDM is based on the set with an owner and a member record types concept: the DBTG set

Each DBTG set can have any number of *set occurrences* (a number of instances of linked records). many-to-many links are not allowed, each set occurrence has precisely one owner, and has zero or more member records.

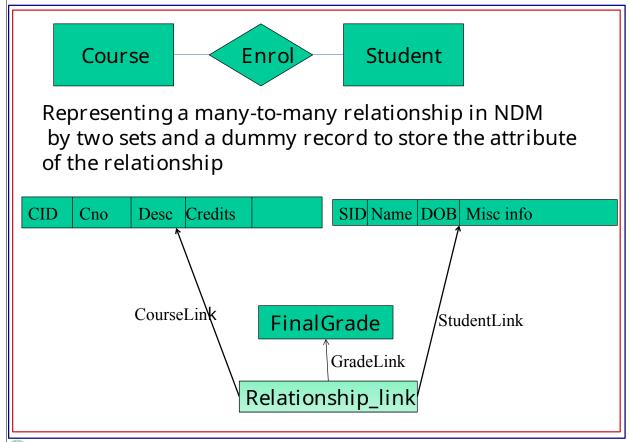
A member record in a set can participate in only one occurrence of a given set at any point.

However, any record can participate simultaneously in several set occurrences of *different* DBTG sets.

Did not use SQL - but not a NoSQL DB!!

Bipin C Desai

7



# **Object Oriented World & ODL**

- ❖ In an object oriented design, the "part of the world" we want to model is thought of as being composed of objects
- Everything is an object and "similar" objects are instances of objects called class
  - ♦ People, Employees, Bank accounts, Students, Course, Airline flights
- ❖ A class simply represents a grouping of similar objects
- \* Every **object** is an **instance** of a **class** and has a unique object identification (OID)
- All objects that are instances of the same class have the same properties and behavior(interaction of objects)
- **ODL** (Object Definition Language) is a proposed standard language for specifying structure of databases
- ❖ ODL is an extension of IDL (Interface Description Language), a component of CORBA (Common Object Request Broker Architecture)

Bipin C Desai

Ç

### **Class Declarations**

- ❖ A declaration of a **class** in ODL, consists of:
  - ♦ The keyword **class**
  - ◆ The **name** of the class
  - ◆ A bracketed { ...} list of **properties** of the class

### **Properties of ODL classes**

- \* ODL classes can have three kinds of properties:
  - **♦** Attributes
    - □ properties whose types are built from **primitive/basic types** such as integers, strings,...
  - **♦** Relationships
    - □ properties whose type is either a **reference** to an object

(x-one) or a **collection** of such references (x-many), where x could be one or many.

- **♦** Methods
  - □ **functions** that may be applied to objects of the class

Bipin C Desai

1

### **Attributes in ODL**

- \* Attributes are the **simplest kinds** of properties
- An attribute describes some aspect of an object by associating, with the object, a value of some simple type
- For example, attributes of a Student object
  - ♦ Student ID
  - ♦ Name
  - **♦** Address
  - ♦ E-mail

Bipin C Desai

12

# Keys in ODL In ODL, we declare keys using the keyword key If a key has more than one attribute, we surround them by (...) □ Example: (two attributes forming a key) class Movie (extent Movies key (title, year)) { attribute string title; ... }; If a class has more than one key, we may list them all, separated by commas □ Example: (A class with two keys) class Employee (extent Employees key empID, SIN) {...};

Bipin C Desai

13

# **Single-Value Constraints in ODL**

- Often, we should *enforce* properties in the database saying that there is at most one value playing a particular role
  - ♦ For example:
    - □ that a movie object has a **unique** title, year, length, and film type
    - □ that a movie is owned by a **unique** studio

# **Single-Value Constraints**

- ❖ In **ODL**:
  - ♦ An attributes is not of collection type (Set, Bag, Array, ... are example of **collection types**.)
  - ♦ A relationship is either a class type or (a single use of) a collection type constructor applied to a class type.
- \* Recall that in **E/R**:
  - attributes are atomic
  - an arrow (→) or a value on the connecting line can be used to express the type of relationship(multiplicity)
  - ♦ How about multi-valued? attributes (No) but relationships (Yes)

Bipin C Desai

14

# **Type system**

A type system consists of

- **♦** Basic types
- ◆ Type constructors: recursive rules whereby complex types are built from simpler ones
- **Atomic types**

Integer Float

Char Character String

Boolean Enumeration

Enumeration is a **list of names** declared to be **synonyms for integers** 

- Class types
  - **♦** Movie

Bipin C Desai

16

### Type constructors in ODL Note: **❖** Set Set, Bag, Array, List and Set <integer> ♦ Set <Movie> Dictionary are called Bag collection types Bag <integer> Bag <Movie> Collection type cannot be Array applied repeatedly Array <integer, 10> ♦ Array <Movie, 3> (nested) **Structure** E.g., it is illegal to ◆ Struct Address {string street, string city} write **!** List Set<Array<integer> List <integer> ♦ List <Student> Dictionary <keyType, valueType> ◆ Dictionary<Student, string>

Bipin C Desai 17

```
class Movie {
    attribute string title;
    attribute integer year;
    attribute enum Film {Colour, BlackAndWhite} filmType; };

class Star {
    attribute string name;
    attribute Struct Address {
        string street,
        Array <char, 10> city
        } homeAddress;
        attribute Struct Address officeAddress; };
```

### **More Examples**

```
class Student {
  attribute string ID;
  attribute string lastName;
  attribute string firstName;
  attribute integer dob;
  attribute string program;
  attribute Struct Address {
    string street,
    string city
    } homeAddress;
};
```

```
class Course {
  attribute string courseNumber;
    attribute string courseName;
    attribute integer NoOfCredits;
  attribute string department;
  };
```

Bipin C Desai

10

# **Expressing Relationships in ODL**

- \* How are **Movies** and **Stars** related?
- \* Movies have actors/actresses(Stars), and Stars have roles in Movies!
- \* Every movie has a star (or stars)
- ❖ In ODL the interaction of classes is expressed by a construct called "relationship"!
- To take into account the fact that a relationship could involve more than one instance of an object from the related class it is expressed as a Set
- Note: In ODL relationship(s) is(are) stored in an object as "OID pointer(s)"; such relationship(s) is(are) not attribute(s)!

# Relationship in ODL: an Example

```
* starOf is a relationship between Movie and Star
class Movie {
   attribute string title;
   attribute integer year;
   attribute integer length;
   attribute enum Film {Colour, B&W} filmType;
   relationship Star starOf;
   };
```

Bipin C Desai

2

- \* How are **Movies** and **Stars** related?
- Not only every movie has a star
- \* But also every star has acted in some movie
- \* To fix this in the **Star** class, we should add the line:

### relationship Movie starredIn;

```
class Star {
    attribute string name;
    attribute Struct Address {
        string street,
        string city
        } address;
    relationship Movie starredIn;
    };
```

#### **Inverse Relationships**

- \* We are omitting a very important aspect of the relationship between movies and stars
- ❖ We need a way to ensure that if a star S is *starOf* movie M, then movie M is *starredIn* for star S
- ❖ In ODL that is done by defining **inverse** of a relationship for each class.

#### StarredIn

Movies Stars

StarOf

Bipin C Desai

2

```
class Movie {
  attribute string title;
  attribute integer year;
  attribute integer length;
  attribute enum Film {colour, B&W} filmType;
  relationship Star starOf
                                    class Star {
          inverse Star::starredIn;
                                       attribute string name;
};
                                       attribute Struct Address {
                                           string street,
What is the problem here!
                                           string city
- how many actors in a movie?
                                       } address;
- how many movies credits for
                                       relationship Movie starredIn
       an actor?
                                         inverse Movie::starOf;
                                    };
```

## Relationships in ODL

- Our model is not quite complete: it is missing an important point!
- ❖ A movie typically has several actors and each actor is featured in many movies.
- To fix this, we need to express the relationship as a set:

#### relationship Set<Star> stars;

```
class Movie {
                                                      class Star {
  attribute string title. Why is this not a set?
                                                         attribute string name;
                                                         attribute Struct Address {
  attribute integer year;
                                                              string street,
  attribute integer length;
  attribute enum Film {colour, B&W} filmType;
                                                              string city
                                                         } address;
          relationship Set<Star starOf
                                                         relationship Set<Movie> starredIn
                    inverse Star::starredIn;
                                                           inverse Movie::starOf;
       The inverse relationship only specifies the name of the relationship in Star
       the set is in Star not in Movie
                                                           What about attributes of a relationship?
```

Bipin C Desai

25

```
    Suppose we introduce another class, Studio, representing companies that produce movies
    How are Movies and Studios related? Every Studio owns several Movies
```

```
class Studio {
    attribute string name;
    attribute string address;
    relationship Set<Movie> owns inverse Movie::ownedBy;
    };
```

\* What about inverse? Every Movie is owned by some Studio

```
class Movie {
    attribute string title;
    attribute integer year;
    attribute integer length;
    attribute enum Film {color, blackAndWhite} filmType;
    relationship Set<Star> starOf inverse Star::starredIn;
    relationship Studio ownedBy inverse Studio::owns;
};
```

## **Multiplicity of relationships**

- \* In general, when we have a pair of inverse relationships, there are four cases:
  - ◆ The relationship is unique in both directions (one case)
  - ◆ The relationship is unique in just one direction (two cases)
  - ◆ The relationship is not unique in any direction (one case)
  - ♦ The *multiplicity* thus refers to the one of these relationships; also denoted as 1-1 (one-one), 1-M (one-many), M-1 (many-one), and M-N (many-many).

Bipin C Desai

27

## Multiplicity of relationships: many-many

A many-many relationship from a class C to a class D is one in which, for each C there is a set of Ds associated with C; in the inverse relationship, a set of Cs is associated with each D

Example: each student can take many courses and each course can be taken by more than one student

```
class Student {
    ...
    relationship Set<Course> takes inverse Course::takenBy;
    };

class Course {
    ...
    relationship Set<Student> takenBy inverse Student:: takes;
    };
```

#### Multiplicity of relationships: many-one

\* A many-one relationship from class C to a class D, is one where for each C there is a at most one D, but no such a constraint in the reverse direction (similarly for one-many

Example, many employees may work in the same department, but each employee works only in one department

🕩 Bipin C Desai

20

## Multiplicity of relationships: one-one

A one-one relationship from class C to class D is one that for each C there is a at most one D, and conversely, for each D there is at most one C

Example: each department has at most one employee as its manager and each employee can manage at most one department

```
class Employee {
    ...
    relationship Department ManagerOf
        inverse Department::manager;
    };
class Department {
    ...
    relationship Employee manager
        inverse Professor:: ManagerOf;
    };
```

#### **Inheritance in Object Oriented System and Subclasses**

- Objects can be organized into a hierarchical inheritance structure
- A child class (or subclass) will inherit properties form a parent class (or all superclasses) higher in the hierarchy.
- \* Often, a class contains objects that have **special properties** not associated with all members of the class
- \* If so, we find it useful to organize the class into *subclasses*, each subclass having its **own special** attributes and/or relationships



Bipin C Desai

3

#### **Subclasses in ODL**

❖ We define a class *C* to be a subclass of another class *D* by following the name *C* in its declaration with a keyword extends and the name *D* 

```
class Cartoon extends Movie {
    relationship Set<Star> voices;
};
```

A subclass *inherits* all the properties of its superclasses

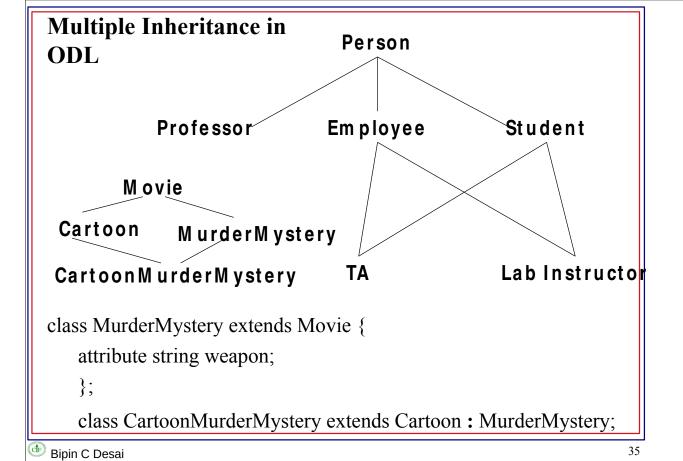
So, each cartoon object has *title*, *year*, *length*, *filmType*, and inherits relationships *stars* and *ownedBy* from Movie, in addition to its own relationship *voices*.

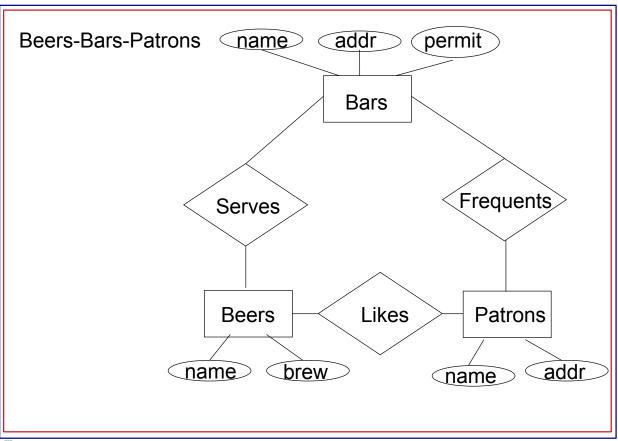
```
Person
               class Person {
                       attribute string lastName;
                       attribute string firstName;
                       attribute integer age;
                       attribute Struct Address {
                      string street,
                      string city
                               } homeAddress;
                                                    Professor
 Student
                                   class Professor extends Person {
class Student extends Person {
                                          attribute string EmpID;
  attribute string ID;
                                          attribute set<string> interest;
  attribute string program;
  };
```

33

#### **Inheritance in ODL**

- ❖ A class may have **more than one** subclass.
- \*A class may have more than one class from which it inherits properties; those classes are its superclasses
- Subclasses may themselves have subclasses, yielding a **hierarchy** of classes where each class inherits the properties of its ancestors.



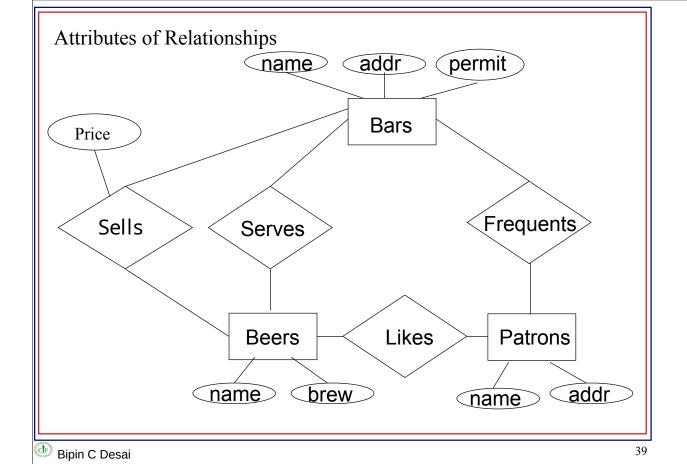


```
class Bars {
class Beers {
                                   attribute string name;
attribute string name;
                                   attribute Struct Addr
attribute string manf;
                                   {string street, string city, string PC}
relationship Set<Bars> servedAt
                                   address:
inverse Bars::serves;
                                   attribute Enum SAQ (full, beer,
                                   BYOB, none } PermitType;
relationship Set<Patrons> fans
                                   relationship Set<Patrons> customers
inverse Patrons::likes;
                                   inverse Patrons::frequents;
                                   relationship Set<Beers> serves
                                   inverse Beers::servedAt;
 Name is given to structure
 & enumeration type for
 possible reuse
```

de Bipin C Desai 37

```
class Patrons {
  attribute string name;
  attribute Struct Bars::Addr
  address;
  relationship Set<Beers> likes
  inverse Beers::fans;
  relationship Set<Bars> frequents
  inverse Bars::customers;
}

Reuse – qualify the name with the class
  for disambiguation
```



The attribute of a relationship converted into a three-way relationship!

Price

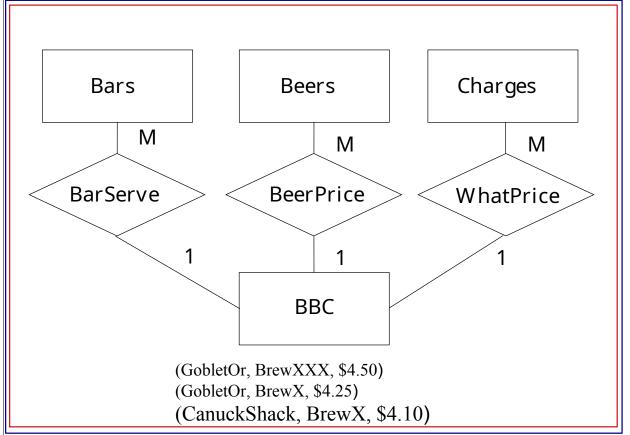
Charges

Bars

Sells

Beers

If price depended only on the beer, then we could use two binary relationships: charge-beer and beer-bar.



```
class Beers {
   attribute string name;
   attribute string manf;
   relationship Set<Bars> servedAt inverse Bars::serves;
   relationship Set<Patrons> fans inverse Patrons::likes;
   relationship Set <BBC> BeerCharges inverse BBC::BeerPrice }
```

43

#### From the internet to the Web

Early 80s: Archie, Veronica; internet file sharing finding systems Late 80s early 90s

HTTP/HTML Tim Berners-Lee, Robert Caillau

Text based browser, Lynx, start of Netscape

HTTP request-response protocol between a client and a server

HTTP session is a sequence of requests-responses

HTML - very simple text markup language

included features for simple formatting and display

HTML based on ideas existing in the late 1980s including:

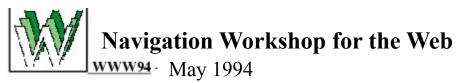
TeX/LaTex, Troff,

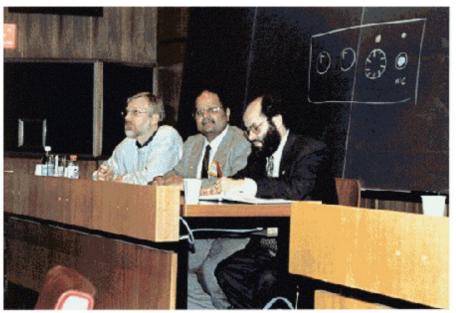
SGML and

the early word processing software

(WordStar, WordPerfect, Word)

May 1994 CERN – Geneva: The first World Wide Web conference





SGML and HTML are mark up languages for information(textual)

HTML was too simple and not extensible

Hence XML E**X**tendible **M**arkup

SGML was extensible but too complex.

**L**anguage

None of these languages do anything: none of these languages are *Turing complete* 

They provide a way to present information which is wrapped in tags. Note the evolution of data/metadata: **metadata** in program; **metadata** in schema; **metadata** with the data

The tags are commonly accepted by a community/group who want to exchange information.

SGML and XML specify the content and structure of a document in a way that allows particular presentations to be generated as needed

```
<!doctype linuxdoc system>
```

- <!-- This a sample SGML file. Comments can appear anywhere It can go over a number of lines. -->
- <article>
- <!-- Article type document -->
- <title>Sample SGML Document
- <!-- Always give a Title. Should be descriptive -->
- <author>Bipin C. DESAI
- <date>March 2000
- <!-- Note the tag minimization the end tags are assumed by the occurrence of a new tag -->
- <abstract>

This document is a sample document using the simple Linuxdoc-SGML DTD: used to write all documents for Linux. There are other DTDs and you can create your very own DTD. However, you have to create all the scripts for its translation to other formats.

</abstract>

Bipin C Desai

47

### A sample DTD saved as notes1.dtd

- <!ELEMENT xslNotes O (title,author,para+)>
- <!ELEMENT title O CDATA>
- <!ELEMENT author O CDATA> O optional end tag
  CDATA character data
- <!ELEMENT para O CDATA>

#### **Use of the sample DTD**

- <?xml version="1.0"?>
- <!DOCTYPE xslNotes SYSTEM "notes1.dtd" >
- <xslNotes>
- <title>XSL Notes</title>
- <author>Bipin C. Desai</author>
- <para> This is paragraph 1.
- <para> This is paragraph 2.
- </xslNotes>

48

Why separate content and structure from presentation and behavior

Once coded, the information can be reused in many formats

Device/Media-independent publishing

One-on-one marketing

Intelligent downstream document processing

Large-scale information management.

XML (Extensible Markup Language): A subset of SGML (ISO 8879) designed for easy implementation

Bipin C Desai

Information in XML form has to be rendered using appropriate formatting mechanism

XML document contains the syntax,

tags are used to provide "keys"

content within the tags represent the "value"

Tags have no predefined meaning but is agreed to by parties involved in the exchange of information

XML by itself conveys only content and structure, not presentation or behavior

XML data is stored in plain text format independent of software/hardware. makes it easy to share data

XML Simplifies Data Interchange

XML applications are designed and adapted to read xml data.

XML Simplifies Platform Changes

New platforms are designed/built so that they can use existing and new XML data

Since all new appliances implement XML features, XML data can be used with diverse devices

> ----- is fresher because more people eat it, more people eat it because ------ is fresher!

Bipin C Desai

5

XQuery is to XML what SQL is to relational databases.

XQuery was designed to query XML data.

XQuery for XML is what SQL for databases

XQuery is built on XPath expressions

XQuery is supported by all major databases

Path Expressions (no joins!)

XQuery uses path expressions to navigate through elements in an XML document.

XPath is used to address (select) parts of documents using

#### path expressions

A path expression is a sequence of document step/tags separated by "/"

Each step operates on the set of instances produced by the previous step Selection predicates may follow any step in a path, in []

#### **XML Schema**

The XML schema defines:
what are the components(elements) in a
corresponding document
Order of these elements
Number of occurrences
Element's contents – could it be empty or it is required and
its contents
Data types, default values

Bipin C Desai 53

Title	Authors array	Publication (Name, Vol, Date, pages)	Meeting	Subject set
Report of the Priorities Workshop	[Caillau, Desai]	(Computer Networks and ISDN Systems; Vol. 27-2,;November 1994; pp. 334-336)	WWW-I	{Web, searching}
Three Paradoxes of Big data	[Richards, King]	(Stanford, CA,;;Sept, 2013,;pp102-1050)	Making End Meet	{Big Data, security, privacy]

A not too correct XML schema to express this type of information is given in the next slide.

Exercise: complete the xml schema and create the xml doc for the above data!

```
Bipin C Desai 55
```

```
<pns:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
<pns:element name="papers" type="Publications" />
<pns:element name="tittle">
  <pns:complexType>
    <pns:sequence>
       <pns:element name="usual_title" type="pns:string"/>
       <pns:element name="alt_title" type="pns:string"/>
    <pns:complexType name="Publications">
  <pns:sequence>
   <pns:element ref="title" minOccurs="0" maxOccurs="unbounded"/>
   <pns:element ref="authors" minOccurs="0" maxOccurs="unbounded"/>
   <pns:element ref="publication" minOccurs="0" maxOccurs="1"/>
   <pns:element ref="meeting" minOccurs="0" maxOccurs="1"/>
   <pns:element ref="subjects" minOccurs="0" maxOccurs="unbounded"/>
 </pns:sequence>
</pns:schema>
```

Predicates (where clause)

XQuery uses predicates to limit the extracted data from XML documents

Storing XML data BLOB Decompose and save as tables

Bipin C Desai

5/

# Storage Devices, Files, and Indexing



To be used in the spirit of copy-forward! https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

Bipin C Desai

#### **Storage Device Selection Criteria**

Capacity vs. cost (What will \$100 buy? How much for 1 Megabytes?) Cost per megabytes of storage has taken a plunge Alas, the need for it has bounded as well.

Permanence

**Portability** 

Relative cost

Performance (Latency, transfer /access rate)

Record size - buffer size, file size.

Accessing method - random/direct or sequential

Data transfer rate

Seek time - time to move read/write head:

average, minimum, maximum

Latency - rotational delay (rpm)

Memory Hierarchy					
Speed	Technology	Application			
1-10's nsec	$I^2L$	fast cache			
	nmos	high speed MM			
	bipolar	buffer			
100's nsec	nmos	main memory			
	core				
100's μsec	CCD	fast back up			
	bubbles				
1-10's msec	floppy disk	• • • • • • • • • • • • • • • • • • •			
fixed head disk					
moving head disk					
10's msec	magnetic tape	security/back up			
100s of ms	optical memor	y large mass			
	tape library	memory,			
	system	archives			

(d) Bipin C Desai

## **Data on External Storage**

<u>Disks:</u> Can retrieve random page at fixed cost

But reading several consecutive pages is much cheaper than reading them in random order

<u>Tapes:</u> Can only read pages in sequence

Cheaper than disks; used for archival storage – extinct??

<u>File organization:</u> Method of arranging a file of records on external storage.

Record id (RID) is sufficient to physically locate record Indexes are data structures that allow us to find the record ids of records with given values in index search key fields

<u>Architecture:</u> Buffer manager stages pages from external storage to main memory buffer pool. File and index layers make calls to the buffer manager.

Bipin C

## Store the database in Main Memory!

Costs too much. \$100 will buy 32GB of DRAM or 1TB SSD today.

Main memory is volatile. We want data to be saved between runs. (Obviously!)

Typical storage hierarchy:

Cache

Main memory (RAM) for currently used data.

Disk for the main database (secondary storage).

Tapes for archiving older versions of the data (tertiary storage).

Bipin C

E

#### **EXTERNAL STORAGE MEDIUMS**

Read/Write Write once read many times(WORM)

(used for archives)

Magnetic Tape Disks/tapes

Disk Robotic storage media

RAID(Redundant CD-Rom.

array of

inexpensive

disks

READ: transfer data to main memory.

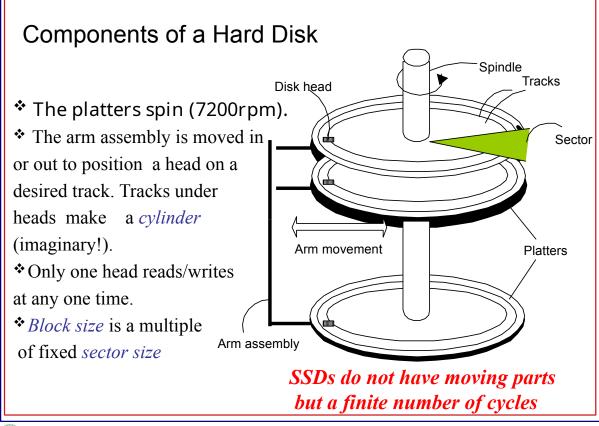
WRITE: transfer data to external device.

READ/WRITE are much slower than main-memory operations!

#### Disks

- Secondary storage device of choice:
   HDD are being replaced by SSD.
- Main advantage over tapes: <u>random access</u> vs. sequential.
- Data is stored and retrieved in units called *disk blocks* or *pages*.
- Unlike RAM, time to retrieve a disk page varies depending upon location on disk.
  - Therefore, relative placement of pages on disk has major impact on DBMS performance!

(b) Bipin C



Bipin C

8

#### Seek Time

Seek Time =  $c_1 + c_2 *$  (number of cylinders to be traversed). Here  $c_1$  and  $c_2$  are constants for a given model of disk drive. Average Seek Time = time to move over 1/3 cylinders. Seek time can be reduced by:

- distributing a file over a number of disk units and
- limiting the range of cylinders on any disk unit.

#### Rotational Latency

The delay between the completion of the seek and the actual transfer of data

actual transfer of data.	<u>RPM</u>	<u> Latency</u>
	10000	3 msec
For a disk rotating at r (RPM)	7200	4.1 msec
$t_1 = \underline{60 * 1000}$ milliseconds	6000	5 msec
2* r	3000	10 msec
	2400	12.5 msec

Bipin C Desai

## Accessing a Disk Page

- Time to a read or a write a disk block:
  - seek time (moving arms to position disk head on track)
  - *rotational delay* (waiting for block to rotate under head)
  - *transfer time* (actually moving data to/from disk surface)
- Seek time and rotational delay dominate.
  - Seek time varies from about 1 to 20msec
  - Rotational delay varies from 0 to 10msec
  - Transfer rate is about 1msec per 4KB page
- Key to lower I/O cost: reduce seek/rotation delays! Hardware vs. software solutions?

SSD obsoletes these!

Bipin C

```
Response time = seek time + latency time + transfer time (5-20 msec) (3-5 msec).

Transfer time = size of transfer/rate of transfer.

Size of transfer corresponds to the data of interest (excluding format information, etc.)

Sequential Read of a number of blocks.

Transfer time = avg. seek time + latency time + (block transfer time) * number of blocks + (min. seek time + latency) * number of cylinders

Problems: Disk scheduling in multi-process environment Approximation: Transfer time = t<sub>efb</sub> * # of blocks,

Here, t<sub>efb</sub> is the effective formatted block transfer time.
```

to account for the format information and the ignored seek and latency time.

 $t_{efb} \approx$  1.10 \*  $t_{b}$  , where  $t_{b}$  is the block transfer time

Block transfer time = block size/ rate of transfer

Bipin C Desai

11

#### Random Read of a # of blocks

Transfer time = number of blocks \* (seek + latency + t<sub>b</sub>)

# **Sequential Read from a number of contiguous cylinders**Transfer time = seek time +latency time + t<sub>efb</sub> \* # of block + (min seek time + latency time) \* (# of cylinders -1)

File Organisation the storage required for the file organization sequential the time required to read a random record

indexed sequential the time required to read the next record

direct access the time required to add a record other method the time required to update a record

> the time required to read all records the time required to reorganize a file

> > 13

Choice

- external storage device available simple

- use of the file - type of queries x = y- number of keys range

- mode of retrieval - seq. random Boolean x=y

- mode of update batch - economy of storage on-line

- frequency of use of a file - growth potential of a file

- methods available in the development environment

Bipin C Desai

#### **Updates**:

- insert in sequence
  - at end, at first available location
- delete compress first available location flag as deleted
- modify selected record space for update record size with respect to size of original record
- modify all records

#### **Primary Key Retrieval**

Four (three) possible choices -

- serial file no order (pile)
- sequential ordered wrt primary key
- indexed sequence
- direct access

#### Serial Files (PILE)

Access a random record Access to Next Record Inserting Record Deleting a Record Modifying a Record Reorganisation Single Disk Drive

#### **Sequential File**

Access a random record
Access to Next Record
Inserting Record
Deleting a Record
Modifying a Record
Reorganisation
Single Disk Drive
Two or more Disk Drives

#### **Access to Next Record**

Two or more Disk Drives

Probability of record in same block =  $1 - 1/b_f$ Probability of record not same block =  $1/b_f$ Expected time to get next record. =  $0 * (1 - 1/b_f) + 1 * (1/b_f)*(t_s + t_l + t_b)$ =  $1/b_f*(t_{efb})$ 

Bipin C Desai

15

#### **Modify-in-place or Delete a Record**

- Find it in T<sub>f</sub> (Time to find random record)
- Max. time to modify or mark it as deleted, and wait  $2T_i$  block txf time
- Rewrite it in time = block txf time Total time =  $T_f + 2T_f$

#### **Sector Addressable Disks**

- fixed length arcs of a track track is divided into an integral number of sectors.
- amount of data is fixed by O.S. or by the hardware.
  - simplifies allocation of storage space
  - simplifies address calculations
  - simplifies synchronisation of I/O & computation in sequential processing.

The division of a track into **sectors**:

- -may be implemented completely by **hardware** or
- by **software** controlled formatting operation.

**Block** is a fixed number of bytes that is moved as a unit between storage devices and the main memory. Made up a number of disk sectors.

Bipin C Desai

17

## Arranging Pages on Disk

- 'Next' block concept:
  - blocks on same track, followed by
  - blocks on same cylinder, followed by
  - blocks on adjacent cylinder
- Blocks in a file should be arranged sequentially on disk (by 'next'), to minimize seek and rotational delay.
- For a sequential scan, *pre-fetching* several pages at a time
- "De-fragmentation" to increase access

Bipin C

18

#### **RAID**

- Disk Array: Arrangement of several "inexpensive" disks that gives abstraction of a single, large disk.
- Goals: Increase performance and reliability.
- Two main techniques:
  - Data striping: Data is partitioned; size of a partition is called the striping unit. Partitions are distributed over several disks.
  - Redundancy: More disks => more failures. Redundant information allows reconstruction of data if a disk fails.
- Level 0: No redundancy
- Level 1: Mirrored (two identical copies)
  - Each disk has a mirror image (check disk)
  - Parallel reads, a write involves two disks.
  - Maximum transfer rate = transfer rate of one disk

Bipin C

## Level 0+1: Striping and Mirroring

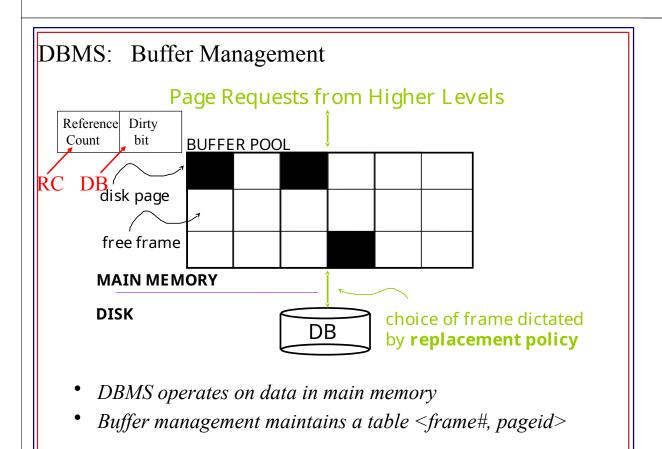
- Parallel reads, a write involves two disks.
- Maximum transfer rate = aggregate bandwidth
- Level 3: Bit-Interleaved Parity
  - Striping Unit: One bit. One check disk.
  - Each read and write request involves all disks; disk array can process one request at a time.
- Level 4: Block-Interleaved Parity
  - Striping Unit: One disk block. One check disk.
  - Parallel reads possible for small requests, large requests can utilize full bandwidth
  - Writes involve modified block and check disk
- Level 5: Block-Interleaved Distributed Parity
  - Similar to RAID Level 4, but parity blocks are distributed over all disks

Bipin C

## Disk Space Management

- Lowest layer of DBMS software manages space on disk.
- Higher levels call upon this layer to:
  - allocate/de-allocate a page
  - read/write a page
- Request for a *sequence* of pages must be satisfied by allocating the pages sequentially on disk! Higher levels don't need to know how this is done, or how free space is managed.

Bipin C 21



Bipin C

22

# When a Page is Requested ...

- If requested page is not in pool:
  - Choose a frame for replacement (LIFO, FIFO, LRU(RC), modified (DB), etc.)
  - If frame is dirty (changed since read into buffer), write it to disk(replacement frame scheme looks for non-dirty frame
  - Read requested page into chosen frame
- *Increment the reference count (RC) of* the page and return its address.

If requests can be predicted (e.g., sequential scans) pages can be <u>pre-fetched</u>

Bipin C

23

## More on Buffer Management

- When a frame is released by an application, the RC is decremented and if the frame is changed, the dirty bit for the frame is set.
- A frame in the buffer may be requested many times, concurrently(reads not update/write)
  - a RC is used to indicate the number of concurrent use of a frame. A frame is a candidate for replacement iff RC = 0.
  - Priority if dirty bit is not set(not modified)
- Concurrency control and recovery may entail additional I/O when a frame is chosen for replacement.

Bipin C

## **Buffer Replacement Policy**

- Frame is chosen for replacement by a *replacement policy*:
  - Least-recently-used (LRU), Clock, MRU etc.
- Policy can have big impact on # of I/O's; depends on the *access pattern*.
- <u>Sequential flooding</u>: Nasty situation caused by LRU + repeated sequential scans.
  - # buffer frames < # pages in file means each page request could cause an I/O.

Bipin C

25

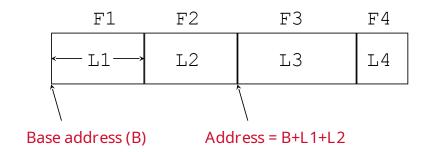
26

# DBMS vs. OS File System

- Differences in different level of support in different OS: portability issues
- Some limitations, e.g., files can't span disks.
- Buffer management in DBMS requires ability to:
  - Manage RC and DB of frames in buffer pool, force a page to disk (important for implementing concurrency control and recovery),
  - adjust *replacement policy*, and pre-fetch pages based on access patterns in typical DB operations.

(b) Bipin C

## Data Record Formats: Fixed Length

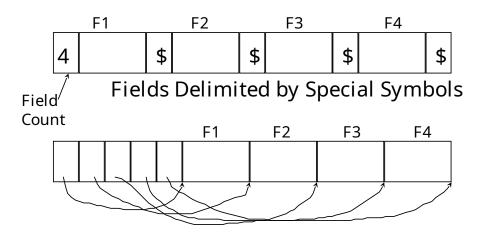


- Information about field types same for all records in a file; stored in *system catalogs*.
- Finding *i'th* field does not require scan of record.

Bipin C 27

# Data Record Formats: Variable Length

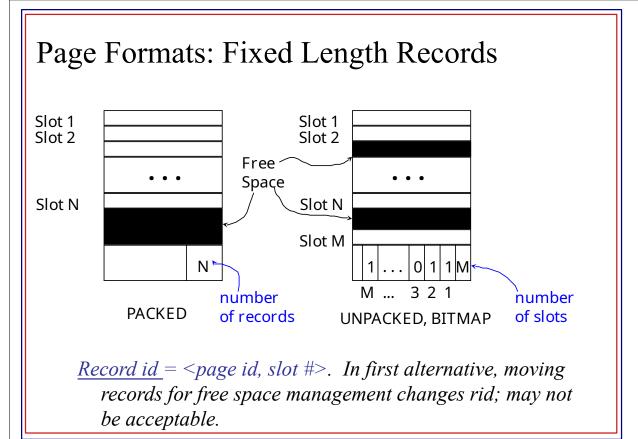
• Two alternative formats (# fields is fixed):



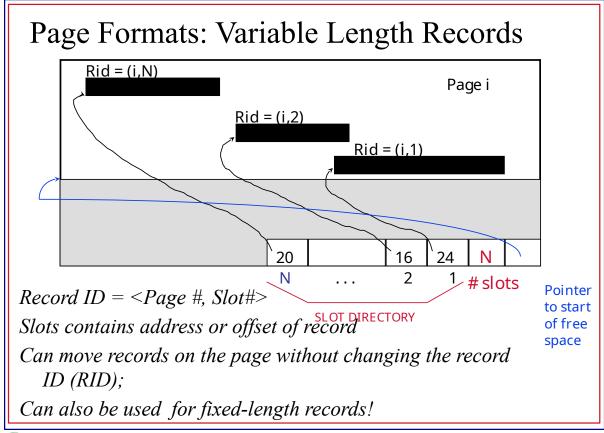
Array of Field Offsets

Second offers direct access to i'th field, efficient storage
 of nulls (special don't know value); small directory overhead

Bipin C



Bipin C 29



Bipin C

## Files of Records

- Page or block is OK when doing I/O, but higher levels of DBMS operate on records, and files of records.
- <u>FILE</u>: A collection of pages, each containing a collection of records. Must support:
  - insert/delete/modify record
  - read a particular record (specified using *record id*)
  - scan all records (possibly with some conditions on the records to be retrieved)



31

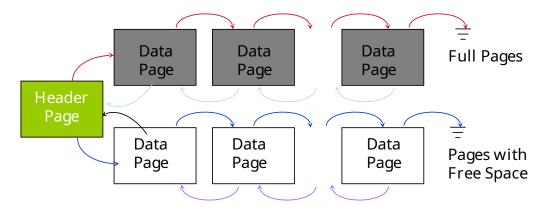
# Unordered (Heap) Files

- Simplest file structure contains records in no particular order.
- As file grows and shrinks, disk pages are allocated and deallocated.
- To support record level operations, we must:
  - keep track of the *pages* in a file
  - keep track of *free space* on pages
  - keep track of the *records* on a page
- There are many alternatives for keeping track of these details.

Bipin C

32

# Heap File Implemented as a List

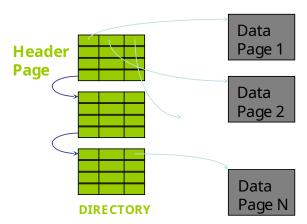


- The Heap file name and its header page address must be stored in a catalog.
- Each page contains 2 `pointers' (forward, reverse) plus data.

Bipin C

33

# Heap File Using a Page Directory



- The entry for a page can include the amount of free space on the page.
- The directory is a collection of pages; for example implemented as a linked list
- Much smaller than linked list of all HF pages!

Bipin C

# System Catalogs

- Catalogs are stored as tables.
- For each table:
  - name, file name, file structure (e.g., Heap file)
  - attribute name and type, for each attribute
  - index name, for each index
  - integrity constraints
- For each view:
  - view name and definition
- For each index:
  - structure (e.g., B+ tree) and search key fields
- Plus statistics, authorization, buffer pool size, etc.

Dipin C 35

## Alternative File Organizations

**Heap files**: Suitable when typical access requires access to all records in a file.

**Sorted Files**: Suitable in cases where the records must be retrieved in some order wrt a "key", or access to a records in a `range' of key values is needed.

Hashed Files: Suitable when random access to records with a given key value is required.

B: The number of blocks (pages) for datab<sub>f</sub>: Blocking factor(# records per block)

t<sub>efb</sub>: Effective time to read or write block

	Heap File	Sorted File	Hashed
			File
Scan all	<b>B</b> t <sub>efb</sub>	<b>B</b> t <sub>efb</sub>	<b>1.25</b> Bt <sub>efb</sub>
recs			
Equality	0.5 Bt <sub>efb</sub>	t <sub>efb</sub> log <sub>2</sub> B	t <sub>efb</sub>
Search			
Range	<b>B</b> t <sub>efb</sub>	t <sub>efb</sub> (log <sub>2</sub> B + #	<b>1.25</b> Bt <sub>efb</sub>
Search		of blocks with	
		matches)	
Insert	<b>2</b> t <sub>efb</sub>	Search + Bt <sub>efb</sub>	<b>2</b> t <sub>efb</sub>
Delete	Search + t.g.	Search + Bt <sub>efb</sub>	<b>2</b> t <sub>efb</sub>
Delete	Jean Ciri Lefb	Jean Cir i Diefb	<b>L</b> efb

Hash 1.25: since pages are only 80% full for avoiding overflows

d Bipin C Desai 37

**INDEX** 

An index is created to speed up access to the records in a file with a given value for a **search key fields**.

Any subset of the fields of a record can be used as search key for an index on the relation.

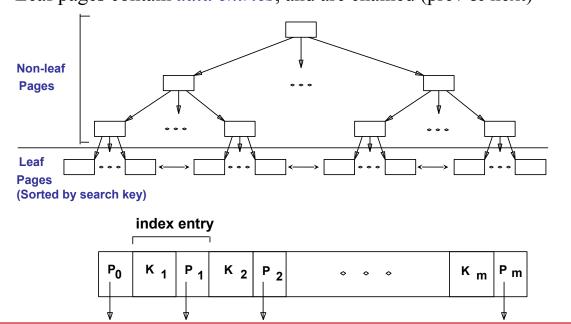
Search key may not be the same as primary key

An index contains a collection of data entries, and supports efficient retrieval of all records with a given search key value **K**.

Bipin C Desai

#### B+ Tree Indexes

- \* Internal nodes (pages) have *index entries*; only used for navigation:
- \* Leaf pages contain *data entries*, and are chained (prev & next)



# 

- Find 28\*? 29\*? All > 15\* and < 30\*
- Insert/delete: Find data entry in leaf, then change it. Need to adjust parent sometimes.
  - And change sometimes bubbles up the tree

Bipin C

41

# Hash-Based Indexes

- Good for equality selections.
- Index is a collection of *buckets*.
  - Bucket = primary page plus zero or more overflow pages.
  - Buckets contain data entries.
- Hashing function h: h(r) = bucket in which (data entry for) record r belongs. h looks at the search key fields of r.
  - No need for "index entries" in this scheme.

## Alternatives for contents of an Index

In an index entry k\* we can store:

Alternative 1: The actual data record with key value k, or

Alternative 2:  $\langle \mathbf{k}, \text{ rid of data record with search key value } \mathbf{k} \rangle$ , or

Alternative 3 < k, list of rids of data records with search key k > k

Choice of alternative for data entries is orthogonal to the indexing technique used to locate data entries with a given key value **k**.

- Examples of indexing techniques: B+ trees, hash-based structures
- Typically, index contains auxiliary information that directs searches to the desired data entries

Bipin C

4

# Alternatives for Data Entries

- Alternative 1:
  - If this is used, index structure is a file organization for data records (instead of a Heap file or sorted file).
  - At most one index on a given collection of data records can use Alternative 1. (Otherwise, data records are duplicated, leading to redundant storage and potential inconsistency.)
  - If data records are very large, # of pages containing data entries is high. Implies size of auxiliary information in the index is also large.

Bipin C

## Alternatives for Data Entries

- Alternatives 2 and 3:
  - Data entries typically much smaller than data records.
     Better than Alternative 1 with large data records, especially if search keys are small. (Portion of index structure used to direct search, which depends on size of data entries, is much smaller than Alternative 1.)
  - Alternative 3 more compact than Alternative 2, but leads to variable sized data entries even if search keys are of fixed length.

Bipin C

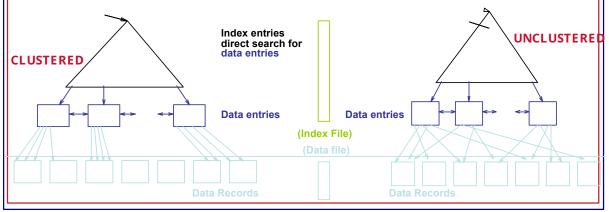
# **Index Classification**

- *Primary* vs. *secondary*: If search key contains primary key, then it is called primary index.
  - *Unique* index: Search key contains a candidate key.
- *Clustered* vs. *un-clustered*: If the order of the data records is the same as, or 'close to', the order of the data entries, then the index is called a clustered index: else un-clustered.
  - Alternative 1 implies clustered; in practice, clustered also implies Alternative 1 (since sorted files are rare).
  - A file can be clustered on at most one search key.
  - Cost of retrieving data records through index varies greatly based on whether index is clustered or not!

## Clustered vs. Unclustered Index

Suppose that Alternative (2) is used for data entries, and that the data records are stored in a Heap file.

- To build clustered index, first sort the Heap file (with some free space on each page for future inserts).
- Overflow pages may be needed for inserts. (Thus, order of data recs is 'close to', but not identical to, the sort order.)



Bipin C

47

# Cost Model for Our Analysis

We ignore CPU costs, for simplicity:

- B: The number of data pages
- R: Number of records per page
- D: (Average) time to read or write disk page
- Measuring number of page I/O's ignores gains of prefetching a sequence of pages; thus, even I/O cost is only approximated.
- Average-case analysis; based on several simplistic assumptions.

These are only estimates to show the overall trends!

Bipin C

# Comparing File Organizations

- Heap files (random order; insert at eof)
- Sorted files, sorted on <age, sal>
- Clustered B+ tree file, Alternative (1), search key <age, sal>
- Heap file with un-clustered B + tree index on search key
   <age, sal>
- Heap file with unclustered hash index on search key < age, sal>

Bipin C

40

# Operations to Compare

Scan: Fetch all records from disk Equality search Range selection Insert a record Delete a record

Bipin C

# Assumptions

- Heap Files: Equality selection on key; exactly one match.
- Sorted Files: Files compacted after deletions.
- Indexes: Alt (2), (3): data entry size = 10% size of record
  - Hash: No overflow buckets.
    - 80% page occupancy => File size = 1.25 data size
  - Tree: 67% occupancy (this is typical).
    - Implies file size = 1.5 data size
- Scans: Leaf levels of a tree-index are chained.
  - Index data-entries plus actual file scanned for unclustered indexes.
- Range searches: We use tree indexes to restrict the set of data records fetched, but ignore hash indexes.

Bipin C

51

# Cost of Operations

- **B:** The number of data pages
- **R:** Number of records per page
- **D:** (Average) time to read or write disk page

	(a) Scan	(b) Equality	(c ) Range	(d) Insert	(e) Delete
(1) H eap	BD	0.5BD	BD	2D	Search +D
(2) Sorted	BD	Dlog 2B	D(log 2 B + # pgs with match recs)	Search + BD	Search +BD
(3) Clustered	1.5BD	Dlog F 1.5B	D(log F 1.5B + # pgs w. match recs)	Search + D	Search +D
(4) Unclust. Tree index	BD(R+0.15)	D(1 + log F 0.15B)	D(log F 0.15B + # pgs w. match recs)	Search + 2D	Search + 2D
(5) Unclust. Hash index	BD(R+0.125)	2D	BD	Search + 2D	Search + 2D

These are estimates using many simplifying assumption:

# Understanding the Workload

For each query in the workload:

- Which relations does it access?
- Which attributes are retrieved?
- Which attributes are involved in selection/join conditions? How selective are these conditions likely to be?

For each update in the workload:

- Which attributes are involved in selection/join conditions? How selective are these conditions likely to be?
- The type of update (INSERT/DELETE/UPDATE), and the attributes that are affected.



# Choice of Indexes

- What indexes should we create?
  - Which relations should have indexes? What field(s) should be the search key? Should we build several indexes?
- For each index, what kind of an index should it be?
  - Clustered? Hash/tree?
- One approach: Consider the most important queries in turn. Consider the best plan using the current indexes, and see if a better plan is possible with an additional index. If so, create it.
  - Obviously, this implies that we must understand how a DBMS evaluates queries and creates query evaluation plans!
  - For now, we discuss simple 1-table queries.
- Before creating an index, must also consider the impact on updates in the workload!

Bipin C 54

# Index Selection Guidelines

- Attributes in WHERE clause are candidates for index keys.
  - Exact match condition suggests hash index.
  - Range query suggests tree index.
    - Clustering is especially useful for range queries; can also help on equality queries if there are many duplicates.
- Multi-attribute search keys should be considered when a WHERE clause contains several conditions.
  - Order of attributes is important for range queries.
  - Such indexes can sometimes enable index-only strategies for important queries.
    - For index-only strategies, clustering is not important!
- Try to choose indexes that benefit as many queries as possible. Since only one index can be clustered per relation, choose it based on important queries that would benefit the most from clustering.

Bipin C

55

# Examples of Clustered Indexes

B+ tree index on E.age can be used to get qualifying tuples.

How selective is the condition?

Is the index clustered?

Consider the GROUP BY query.

If many tuples have E.age > 10, using *E.age* index and sorting the retrieved tuples may be costly.

Clustered *E.dno* index may be better!

Equality queries and duplicates:

Clustering on *E.hobby* helps!

SELECT E.dno FROM Emp E WHERE E.hobby=Stamps

FROM Emp E WHERE E.age>40

SELECT E.dno

SELECT E.dno, COUNT (\*) FROM Emp E WHERE E.age>10 GROUP BY E.dno

Bipin C

# Indexes with Composite Search Keys

Composite Search Keys: Search on a combination of fields.

Equality query: Every field value is indexes using lexicographic order.

equal to a constant value. E.g. wrt <sal,age> index:

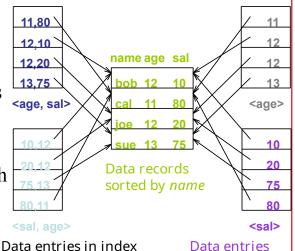
age=20 and sal =75

Range query: Some field value is not a constant. e.g.:age =20; or age=20 and sal > 10

Data entries in index sorted by search key to support range queries.

Lexicographic order, or

Spatial order.



Data entries in index sorted by <sal,age>

Data entries sorted by <sal>

Bipin C

57

# Composite Search Keys

- To retrieve Emp records with age=30 AND sal=4000, an index on <age,sal> would be better than an index on age or an index on sal.
  - Choice of index key orthogonal to clustering etc.
- If condition is: 20<age<30 AND 3000<sal<5000:
  - Clustered tree index on <age,sal> or <sal,age> is best.
- If condition is: *age*=30 AND 3000<*sal*<5000:
  - Clustered <age, sal> index much better than <sal, age> index!
- Composite indexes are larger, updated more often.

# **Index-Only Plans**

A number of queries can be answered without retrieving any tuples from one or more of the tables involved if a suitable index is available.

<E.dno>

SELECT E.dno, COUNT(\*)
FROM Emp E
GROUP BY E.dno

<E.dno,E.sal> Tree index!

SELECT E.dno, MIN(E.sal) FROM Emp E GROUP BY E.dno

<E. age, E.sal > SELECT AVG(E.sal)
or FROM Emp E
<E.sal, E.age > WHERE E.age = 25 AND
Tree index! E.sal BETWEEN 3000 AND 5000

Bipin C

59

# Index-Only Plans (Contd.)

- Index-only plans are possible if the key is <dno,age> or we have a tree index with key <age,dno>
  - Which is better?
  - What if we consider the second query?

SELECT E.dno, COUNT (\*)
FROM Emp E
WHERE E.age=30
GROUP BY E.dno

SELECT E.dno, COUNT (\*)
FROM Emp E
WHERE E.age>30
GROUP BY E.dno

# Index-Only Plans (Contd.)

 Index-only plans can also be found for queries involving more than one table;  $\langle F.dno \rangle$ 

SELECT D.mgr FROM Dept D, Emp E WHERE D.dno=E.dno

<E.dno,E.eid>

SELECT D.mgr, E.eid FROM Dept D, Emp E WHERE D.dno=E.dno

Bipin C

61

# Summary

- Many alternative file organizations exist, each appropriate in some situation.
- If selection queries are frequent, sorting the file or building an *index* is important.
  - Hash-based indexes only good for equality search.
  - Sorted files and tree-based indexes best for range search; also good for equality search. (Files rarely kept sorted in practice; B+ tree index is better.)
- Index is a collection of data entries plus a way to quickly find entries with given key values.

#### Summary (Contd.)

- Data entries can be actual data records, <key, rid> pairs, or <key, rid-list> pairs.
  - Choice orthogonal to *indexing technique* used to locate data entries with a given key value.
- Can have several indexes on a given file of data records, each with a different search key.
- Indexes can be classified as clustered vs. unclustered, primary vs. secondary, and dense vs. sparse. Differences have important consequences for utility/performance.



# Summary (Contd.)

- Understanding the nature of the *workload* for the application, and the performance goals, is essential to developing a good design.
  - What are the important queries and updates? What attributes/relations are involved?
- Indexes must be chosen to speed up important queries (and perhaps some updates!).
  - Index maintenance overhead on updates to key fields.
  - Choose indexes that can help many queries, if possible.
  - Build indexes to support index-only strategies.
  - Clustering is an important decision; only one index on a given relation can be clustered!
  - Order of fields in composite index key can be important.

# Database Index & Performance Optimization

Bipin C. DESAI

© Bipin C Desai 65

Field	Туре	Null	Key	Default	Extra
userid	int(10) unsigned	NO	PRI	NULL	auto_increment
username	varchar(45)	NO	UNI		
password	varchar(45)	NO			
salutation	varchar(45)	NO			
lastname	varchar(64)	NO	İ		
middle name	varchar(30)	NO			
firstname	varchar(64)	NO	İ		
organization	varchar(120)	NO i	į		
department	varchar(255)	NO i	İ		
address	varchar(255)	NO	İ		
city	varchar(70)	NO i	į		
province	varchar(70)	NO	İ		
country	varchar(70)	NO	İ		
postcode	varchar(10)	NO	İ		
email	varchar(70)	NO	İ		
fax	varchar(70)	NO	İ		
phone	varchar(70)	NO i	į		
status	varchar(45)	NO	İ		
register date	datetime	NO	İ	0000-00-00 00:00:00	
last login date	datetime	į NO į	į	0000-00-00 00:00:00	
last conference	int(10)	YES	į	NULL	
receīve email	varchar(20)	NO	İ	NULL	

db Bipin C Desai 66

#### Create INDEX

MariaDB [test]> create index cntr indx on member(country);

Query OK, 0 rows affected (0.061 sec)

Records: 0 Duplicates: 0 Warnings: 0

Creating an index on multiple columns

MariaDB/MySQL allows composite(multi-column) index (up to 16 columns)

Usually 2 or 3 columns are sufficient

CREATE INDEX index name ON TableName (Col1, COL2, COl3);

Bipin C Desai 67

## **Drop INDEX**

Drop index syntax

ALTER TABLE table\_name DROP INDEX index\_name;

Rename index syntax

ALTER TABLE table\_name RENAME INDEX index\_name TO new\_index\_name;

Show indexes syntax

SHOW INDEX FROM tableName;

# EXPLAIN One can use EXPLAIN to see how the DB executes a DML statement DML statements are: SELECT, DELETE, INSERT, REPLACE, and UPDATE statements. EXPLAIN gives execution plan information from the built-in DB optimizer MariaDB [test]> explain select \* from member limit 5; | id | select\_type | table | type | possible\_keys | key | key\_len | ref | rows | Extra | | 1 | SIMPLE | member | ALL | NULL | NULL | NULL | 2625 | |

MariaDB [test]> explain select \* from member limit 1000;

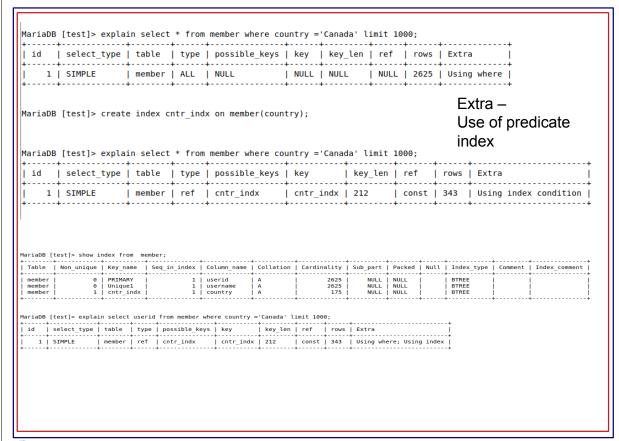
| id | select\_type | table | type | possible\_keys | key | key\_len | ref | rows | Extra |
| 1 | SIMPLE | member | ALL | NULL | NULL | NULL | NULL | 2625 | |

The null for the "possible\_keys" and "key" above are both NULL

+----+

The null for the "possible\_keys" and "key" above are both NULL This indicates that the DB does not have an index it can use The DB will access 2625 rows to generate the result

Bipin C Desai



db Bipin C Desai 70

# SQL III - Relational Object Features

#### Notes



To be used in the spirit of copy-forward! https://users.encs.concordia.ca/~bcdesai/CopyForward.pdf

🕩 Bipin C Desai

- Handling complex data and OO Concept
- Using structured data types and inheritance in SQL
- Object Identity (OID) and reference types in SQL
- Implementing Object features in relational DBMS
- Persistent Programming Languages
- Include object orientation and constructs to deal with added data types in RDBMS.
- Add complex types, including non-atomic values such as nested relations.
- Extend modeling features while retaining the declarative access to data
- Preserve compatibility with SQL

This is what we started with!

123 Smith D1 P1 5 L1 P2 30 L1 234 P1 20 L1 Ma D2 P3 10 L2 P4 5 L3 P1 35 345 Russo D1

Example of a non-normal form (NNF) relation

Bipin C Desai

3

123	Smith	D1	P1	5	L1
			P2	30	L1
234	Ма	D2	P1	20	L1
			P3	10	L2
			P4	5	L3
345	Russo	D1	P1	35	L1

A non-normal form relation  $\rightarrow$  Nested relation

- Abandon atomic attribute requirement by conceptually allowing nested relations — relations within relations
- Maintain the mathematical foundation of relational model
- Allow NNF i. e, non-normal form

# Complex data

Example of a relation with complex data including: array, composite data, sets

Title	Authors array	Publication (Name, Vol, Date, pages)	Meeting	Subject set
Report of the Priorities Workshop	[Caillau, Desai]	(Computer Networks and ISDN Systems; Vol. 27-2,;November 1994; pp. 334-336)	WWW-I	{Web, searching}
Three Paradoxes of Big data	[Richards, King]	(Stanford, CA,;;Sept, 2013,;pp102-1050)	Making End Meet	{Big Data, security, privacy]

Bipin C Desai

# Decomposition of complex data into 4NF decomposition!

A 4NF relation does not have any multivalued dependency of the form  $X \to \to \mathrm{Y}$ 

Report of the Priorities Workshop  $\rightarrow \rightarrow \{Caillau, Desai\}$ 

Title	Authors
Report of the Priorities Workshop	Caillau,
Report of the Priorities Workshop	Desai
Three Paradoxes of Big data	Richards
Three Paradoxes of Big data	King

Bipin C Desai

# Decomposition of complex data into 4NF decomposition (contd.)!

Title	Name	Vol,	Date	pages
Report of the Priorities Workshop	Computer Networks and ISDN Systems	Vol. 27-2	Nov. 1994	pp. 334-336
Three Paradoxes of Big data	Stanford, CA		Sept, 2013,	pp102-1050

Bipin C Desai

# Decomposition of complex data into 4NF decomposition (contd.)!

Title	Meeting
Report of the Priorities Workshop	WWW-I
Three Paradoxes of Big data	Making End Meet

Title	Subject
Report of the Priorities Workshop	Web
Report of the Priorities Workshop	searching
Three Paradoxes of Big data	Big Data
Three Paradoxes of Big data	security
Three Paradoxes of Big data	privacy

8

# **Postgresql**

Ingres was one of first relational 'open source" relational Database that was developed in the early 1970s at UoC, Berkley It gave rise to, among others, SysBase, Microsoft SQL server etc. Ingres was followed by Postgres and Postgresql

It is available for various versions of Linuses and other OS.

To install in

- fedora core: dnf -y install postgres\*
- Debian, Ubuntu: apt install postgresql postgresql-contrib

Bipin C Desai

PostgreSQL is sometimes called an "object-relational database" because it supports table inheritance.

Most of the ORDBMS features were removed from the Postgres and became PostgreSQL.

Once PostgresSQL is installed the database is initialized using: initdb -D /path-to/postgres/data

One can now start the database server using:

pg\_ctl -D /path-to/postgres/data -l logfile start

Once the server is running, a database could be created using: createdb testdb

# Postgresql shell

The shell can be started using the command: psql

To exit from psql use: \q or CTRL-D

To get a list of commands use: \?

To use a particular database use: \c nameofDB;

To list all the databases use; \l;

To stop Postgresql server use the command

pg\_ctl -D /path-to/postgres/data -l logfile stop

Bipin C Desai

11

# Postgres objects

```
CREATE TABLE publication (
title text primary key,
authors text[],
meeting text,
publication text[][],
topics text []
);
```

Bipin C Desai

#### postgres=# \d+ publication;

```
Table "public.publication"
         | Type | Modifiers | Storage | Stats target | Description
-----
          | text | not null | extended |
         | text[] |
                     | extended |
authors
meeting
         | text |
                         | extended |
publication | text[] |
                          | extended |
topics
       | text[] |
                          | extended |
Indexes:
   "publication pkey" PRIMARY KEY, btree (title)
Has OIDs: no
```

Bipin C Desai

```
insert into publication values(
'Report of the Priorities Workshop',

ARRAY['Caillau', 'Desai'],
'WWW-I',

ARRAY[['event', 'Computer Networks and ISDN Systems'],
['volume','27-2'], ['year', 'November 1994'], ['pages','pp. 334-336']],

ARRAY['Web', 'searching']);

postgres=# select * from publication;

Report of the Priorities Workshop | {Caillau, Desai} | www-1 |
{{event, "Computer Networks and ISDN Systems"}, {volume, 27-2},
{year, "November 1994"}, {pages, "pp. 334-336"}} | {Web, searching}
(1 row)
```

Bipin C Desai

SQL 1999 extended to support complex types:

Collection and large object types

Nested relations are an example of collection types

Structured types: arbitrary hierarchies and composite attributes Inheritance

Object orientation: object identifiers and references

SQL 1999 is yet to be fully implemented in most DBMS (2014)

Examples of OODBMS are:

ObjectStore, Objectdatabase++, Objectivity/DB, etc.

Some RDBMS have introduced some object features

OODBMS feature including using object oriented language to manipulate database objects along with the others of RDBMS

(ACID, Query language, Recovery)

```
Oracle: Creating type (class)and a nested table¹

CREATE OR REPLACE TYPE person_typ AS OBJECT (
idno NUMBER,
name VARCHAR2(30),
phone VARCHAR2(20),
MAP MEMBER FUNCTION get_idno RETURN NUMBER,
MEMBER PROCEDURE display_details
(SELF IN OUT NOCOPY person_typ ));

/
The SELF parameter denotes the
object instance currently invoking
the method. NOCOPY allows passing
the argument by reference (i.e., not
copying the argument to the method)
```

Bipin C Desai

17

#### **Member Methods for Comparing Objects**

An object type, with multiple attributes of various data types, has no predefined axis of comparison.

Methods should be specified to compare & order object type variables. The option is to define an map method or an order method for comparing objects, but not both.

#### **Map Methods**

Map methods return values that can be used for comparing and sorting.

Return values can be any built-in data types(except LOBs and BFILEs)

#### **Order Methods**

An order method directly compares values for two particular objects.

Bipin C Desai

```
SQL> desc person_typ;
```

Name Null? Type

\_\_\_\_\_

IDNO NUMBER

NAME VARCHAR2 (30) PHONE VARCHAR2 (20)

#### **METHOD**

-----

MAP MEMBER FUNCTION GET\_IDNO RETURNS NUMBER MEMBER PROCEDURE DISPLAY DETAILS

Bipin C Desai

```
CREATE OR REPLACE TYPE BODY person_typ AS
MAP MEMBER FUNCTION get idno RETURN NUMBER IS
BEGIN
RETURN idno;
END;
MEMBER PROCEDURE display details
   ( SELF IN OUT NOCOPY person typ ) IS
BEGIN -- use the put line procedure of the DBMS OUTPUT
        -- package to display details
DBMS OUTPUT.put line(TO CHAR(idno)|| ' - '|| name|| ' - ' || phone);
END;
                      CREATE OR REPLACE TYPE
END;
                      people typ AS TABLE OF
                      person typ; -- nested table type
Type body created.
```

Bipin C Desai

Creating an Instance of a VARRAY or Nested Table

To create an instance of a collection type by calling the constructor method of the type.

The constructor method is the name of the type.

The elements of the collection is a comma-delimited list of arguments to the method, for example.

person\_typ(1, 'John Smith', '1-650-555-0135')

Bipin C Desai 21

- Create a table that contains an instance of the nested table type people\_typ, named people\_column,
-use the constructor method in a SQL statement to insert values into people typ.

Example: Using the Constructor Method to Insert Values into a Nested Tab

```
CREATE TABLE people_tab (
group_no NUMBER,
people_column people_typ ) -- an instance of nested table
NESTED TABLE people_column STORE AS people_column_nt
/
```

Table created.

Bipin C Desai

```
INSERT INTO people_tab VALUES ( 100, people_typ( person_typ(1, 'John Smith', '1-650-555-0135'), person_typ(2, 'Diane Smith', NULL)));
1 row created.

Create a department_persons Table Using the DEFAULT Clause

CREATE TABLE department_persons (
dept_no NUMBER PRIMARY KEY,
dept_name CHAR(20),
dept_mgr person_typ DEFAULT person_typ(10/John Doe',NULL),
dept_emps people_typ DEFAULT people_typ())

NESTED TABLE dept_emps STORE AS dept_emps_tab;
Table created.
```

Bipin C Desai 23

```
SQL> desc department_persons;
Name
              Null?
                          Type
 DEPT_NO
                NOT NULL NUMBER
 DEPT_NAME
                            CHAR (20)
 DEPT_MGR
                            PERSON_TYP
 DEPT_EMPS
                            PEOPLE_TYP
INSERT INTO department persons VALUES
(101, 'Physical Sciences', person typ(65,'Vrinda Mills', '1-650-555-0125'),
people typ( person typ(1, 'John Smith', '1-650-555-0135'),
person_typ(2, 'Diane Smith', NULL) ) );
INSERT INTO department persons VALUES
(104, 'Life Sciences', person typ(70, 'James Hall', '1-415-555-0101'),
people typ()) -- an empty people typ table
-- Note that people typ() is a literal invocation of the constructor
-- method for an empty people typ nested table.
```

```
select * from department_persons;

DEPT_NO DEPT_NAME

DEPT_MGR(IDNO, NAME, PHONE)

101 Physical Sciences

PERSON_TYP(65, 'Vrinda Mills', '1-650-555-0125')

PEOPLE_TYP(PERSON_TYP(1, 'John Smith', '1-650-555-0135'),

PERSON_TYP(2, 'Diane Smith', NULL))

104 Life Sciences

PERSON_TYP(70, 'James Hall', '1-415-555-0101')

DEPT_NO DEPT_NAME

DEPT_MGR(IDNO, NAME, PHONE)

DEPT_EMPS(IDNO, NAME, PHONE)

DEPT_EMPS(IDNO, NAME, PHONE)

PEOPLE_TYP()
```

Bipin C Desai

25

# Nesting Results of Collection Queries

SELECT d.dept\_emps Example shows the query retrieving the rested collection of employees from the department persons table

The column dept\_emps is a nested table collection of person\_typ type.

The dept\_emps collection column appears in the SELECT list as an Ordinary scalar column.

Querying a collection column in the SELECT list this way nests the elements of the collection in the result row that the collection is associated with. DEPT EMPS(IDNO, NAME, PHONE)

PEOPLE\_TYP(PERSON\_TYP(1, 'John Smith', '1-650-555-0135'), PERSON\_TYP(2, 'Diane Smith', NULL))
PEOPLE\_TYP()

# **Unnesting Results of Collection Queries**

Not all tools or applications can deal with results in a nested format.

To view collection data using tools that require a conventional format, one must un-nest, the collection attribute of a row into one or more relational rows using a TABLE expression TABLE expressions enable you to query a collection in the FROM clause as a table.

In effect, you join the nested table with the row that contains the nested table.

Bipin C Desai

27

TABLE expressions can be used to query any collection value expression, including transient values such as variables and parameters.

Example Un-nesting Results of Collection Queries

SELECT e.\*

FROM department\_persons d, TABLE(d.dept\_emps) e;

IDNO NAME	PHONE
1 John Smith	1-650-555-0135
2 Diane Smith	

# Creating and Populating Simple Nested Tables

```
CREATE TABLE students (
graduation DATE,
math_majors people_typ, -- nested tables (empty)
chem_majors people_typ,
physics_majors people_typ)
NESTED TABLE math_majors STORE AS math_majors_nt
-- storage tables
```

NESTED TABLE chem\_majors STORE AS chem\_majors\_nt NESTED TABLE physics\_majors STORE AS physics\_majors\_nt;

Table created.

Bipin C Desai

29

```
SQL> desc students;
               Null? Type
Name
 GRADUATION
                        DATE
                      PEOPLE TYP
MATH_MAJORS
CHEM MAJORS
                        PEOPLE TYP
PHYSICS MAJORS
                              PEOPLE TYP
The NESTED TABLE..STORE AS clause specifies storage
names for nested tables.
Elements of a nested table are actually stored in a
separate storage table.
Storage names -used to create an index on a nested table.
CREATE INDEX math idno idx ON math majors nt(idno);
CREATE INDEX chem idno idx ON chem majors nt(idno);
CREATE INDEX physics idno idx ON physics majors nt(idno);
```

Bipin C Desai

```
UPDATE students

SET math_majors =

people_typ (person_typ(12, 'Bob Jones', '650-555-0130'),

person_typ(31, 'Sarah Chen', '415-555-0120'),

person_typ(45, 'Chris Woods', '415-555-0124')),

chem_majors =

people_typ (person_typ(51, 'Joe Lane', '650-555-0140'),

person_typ(31, 'Sarah Chen', '415-555-0120'),

person_typ(52, 'Kim Patel', '650-555-0135')),

physics_majors =

people_typ (person_typ(12, 'Bob Jones', '650-555-0130'),

person_typ(45, 'Chris Woods', '415-555-0124'))

WHERE graduation = '01-JUN-03';
```

Bipin C Desai

Bipin C Desai

```
select owner, object_name, object_type
from ALL_OBJECTS
where object_type = 'TYPE'and owner='BCDESAI';
select owner, object_name, object_type
from ALL_OBJECTS
where object_type = 'TABLE' and owner='BCDESAI';
```

• Following are some slides from

Database System Concepts, 6<sup>th</sup> Ed.

©Silberschatz, Korth and Sudarshan corrected for Oracle by BCD

(b) Bipin C Desai 35

# Structured Types and Inheritance in SQL

**Structured types** (a.k.a. **user-defined types**) can be declared & used in SQL

```
create type Name as object

(firstname varchar(20),

lastname varchar(20))

final

create type Address as object

(street varchar(20),

city varchar(20),

zipcode varchar(20))

not final
```

 Note: final and not final indicate whether subtypes can be created

SQL> desc Name; Null? Type Name VARCHAR2(20) FIRSTNAME **LASTNAME** VARCHAR2(20) SQL> desc address; address is NOT FINAL Null? Type Name VARCHAR2(20) STREET VARCHAR2(20) CITY VARCHAR2(20) ZIPCODE

Bipin C Desai

37

# Structured Types and Inheritance in SQL

• Structured types can be used to create tables with composite attributes

create table person (
name Name,
address Address,
dateOfBirth date)

• Dot notation used to reference components: name.firstname

SQL> desc person;

Name Null? Type

-----

NAME NAME

ADDRESS ADDRESS

DATEOFBIRTH DATE

🕩 Bipin C Desai

# Structured Types (cont.)

User-defined row types

```
create type PersonTypeas object (<br/>Warning: Type created with compilation errors.<br/>SQL> show errors<br/>Errors for TYPE PERSONTYPE:<br/>LINE/COL ERROR<br/>10/0 PL/SQL: Compilation unit analysis terminated<br/>2/6 PLS-00320: the declaration of the type of this expression is
```

Once a type is created, we can create one or more tables using our (user-defined) type

create table customer of Customer Person Type

Bipin C Desai

## Structured Types (cont.)

Alternative method which uses unnamed row types.

```
create table person_r(
name row(firstname varchar(20),
lastname varchar(20)),
address row(street varchar(20),
city varchar(20),
zipcode varchar(20)),
dateOfBirth date)
```

row is not supported in oracle!!!

```
Methods
In ORDMS, we can add a method declaration with a structured type.

method ageOnDate (onDate date)

returns interval year

Method body is given separately.

create instance method ageOnDate (onDate date)

returns interval year

for CustomerType
```

begin

return onDate - self.dateOfBirth;

end

We can now find the age of each customer:

select name.lastname, ageOnDate (current\_date)

from customer

Bipin C Desai

41

#### **Constructor Functions**

```
Constructor functions are used to create values of structured types create function Name(firstname varchar(20), lastname varchar(20)) returns Name begin set self.firstname = firstname; set self.lastname = lastname;
```

To create a value of type Name, we use

new Name('John', 'Smith')

Normally used in insert statements

insert into Person values

```
(new Name('John', 'Smith),
new Address('20 Main St', 'New York', '11001'),
date '1960-8-22');
```

🕒 Bipin C Desai

end

```
Type Inheritance
Suppose that we have the following type definition for people:

create type Person

(name varchar(20),

address varchar(20))

Using inheritance to define the student and teacher types

create type Student

under Person
```

under Person
(degree varchar(20),
department varchar(20))
create type Teacher
under Person
(salary integer,
department varchar(20))

Subtypes can redefine methods by using **overriding method** in place of **method** in the method declaration

Bipin C Desai

43

# Multiple Type Inheritance

SQL:1999 and SQL:2003 do not support multiple inheritance If our type system supports multiple inheritance, we can define a type for teaching assistant as follows:

```
create type Teaching Assistant under Student, Teacher
```

To avoid a conflict between the two occurrences of *department* we can rename them

```
create type Teaching Assistant
under
Student with (department as student_dept),
Teacher with (department as teacher_dept)
Each value must have a most-specific type
```

®திirio. ூ. இத்து

#### Table Inheritance

Tables created from subtypes can further be specified as **subtables** E.g. **create table** *people* **of** *Person*;

create table students of Student under people; create table teachers of Teacher under people;

Tuples added to a subtable are automatically visible to queries on the super-table

E.g. query on *people* also sees *students* and *teachers*.

Similarly updates/deletes on *people* also result in updates/deletes on subtables

To override this behaviour, use "only people" in query

Conceptually, multiple inheritance is possible with tables

e.g. teaching\_assistants under students and teachers

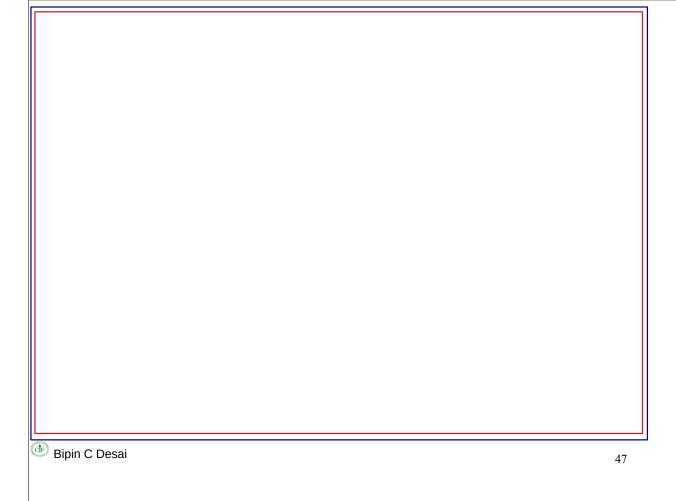
*Not supported in SQL currently:* So we cannot create a person (tuple in *people*) who is both a student and a teacher

Bipin C Desai

45

# JSON Objects

```
SGML ⇒ XML ⇒ JSON
JSON Java Script Object Notation
```



# NOSQL

# Brewer's (CAP) Theorem

There are three core systemic requirements that exist in a special inter-relationship when it comes to designing and deploying applications in a distributed environment

The three requirements are:

Consistency,

Availability and

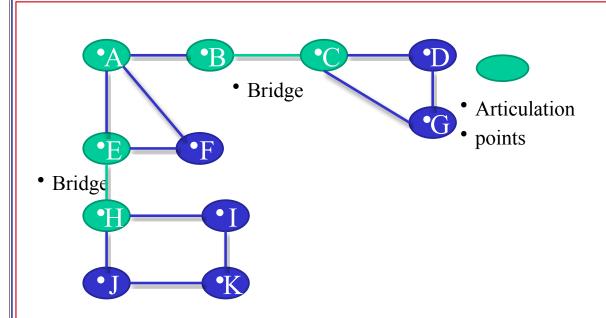
#### **Partition Tolerance**

Compare these with the ACID property that is the traditional requirement

**Atomicity** – all or nothing

**Consistency** - the data goes from one consistent state to another **Isolation** – a transaction is guaranteed to run as if it was the only one **Durability** – any changes made by a transaction are persistent

Bipin C Desai



A distributed system with articulation points (cut vertex) removing which disconnects the graph and bridges(connects subgraphs)

🕩 Bipin C Desai

# Brewer's (CAP) Theorem

Consistency: A constraint of distributed systems that multiple values for the same piece of data are not allowed. Atomicity guarantees that all changes made by a transaction are made or there would be no changes

Availability: Availability means that a service is available. Sites must not to go down at busy periods just because they are busy.

Partition Tolerance: A partition happens when, say, a bridge fails or an articulation node goes down

This causes the network to be partitioned.

Temporary partitions are a possible and critical systems should be tolerant to such events

Bipin C Desai

Dealing with CAP – only two could be guaranteed!

#### **Drop Partition Tolerance**

Run on one system or have bullet proof distributed system (not possible)

#### **Drop Availability Tolerance**

Economically and political downside.

#### **Drop Consistency Tolerance:**

This is the obvious choice in most cases.

Easy to deal with – the masses will not know!!!

Bipin C Desai

# NoSQL

New database applications and new databases – non-relational

Abandon the ACID property – substitute performance, scalability etc. Group most often required data items together

 abandon normalization and the relational approach and hence eliminate joins

Cluster friendly- allow use of multitude of cheap servers

- distributed and partitioned
- No fixed schema (not really!)

Bipin C Desai

53

# NoSQL

Category of "model" and some implementations

**Column family**: BigTable(Google), **Cassandra**, Druid, Hadoop/HBase Unique keys point to multiple columns.

The columns are arranged by column family.

**Document**: Apache CouchDB, Couchbase, MongoDB

Lotus Notes and are similar to key-value stores for semi-structured data The semi-structured documents are stored in JSON like formats.

Key-value: Dynamo(Amazon), FoundationDB, MemcacheDB, Redis

A unique key with pointers to items of data: to implement.

inefficient when accessing small portion of data

Graph: Allegro, Neo4J, InfiniteGraph, OrientDB

A graph theory based model is used

See: http://nosql-database.org/ for a list of NoSQL databases

Bipin C Desai

# Hadoop

Hadoop is a software approach to implements massively parallel computing. http://hadoop.apache.org/

Hadoop modules:

**Hadoop Common**: The common utilities that support the other Hadoop modules.

## Hadoop Distributed File System (HDFS<sup>TM</sup>):

A distributed file system that provides high-throughput access to application data.

**Hadoop YARN**: A framework for job scheduling and cluster resource management.

**Hadoop MapReduce**: A YARN-based system for parallel processing of large data sets.

These modules provide feature that allow data to be spread across thousands of servers with little reduction in performance

Bipin C Desai

- Semi join ⋉ ⋈
- A technique used to support join when a relational database is
- distributed over a number of nodes.
- Suppose table R is on node r and S is on node s and the common
- attribute of R and S is C.
- In semi-join of  $R \ltimes S$ , we proceed as follows
- we send  $\prod_{C} R$  from node r to node s
- at node we do a join of  $(\prod_C R \bowtie \prod_C S)$
- send the result to node r where we do
- $R \bowtie (\prod_C R \bowtie \prod_C S)$
- MapReduce, apply semi join like concept and distribute
- the computation over the nodes
- Answer to the processing needs of large amount of data







• Helen of Troy (1898)

• Paintings by Evelyn de Morgan

Bipin C Desai

57

# Apache Cassandra

Cassandra is a NoSQL Column family implementation

Some of the strong points of Cassandra are:

Highly scalable and highly available with no single point of failure

NoSQL column family implementation

Very high write throughput and good read throughput

SQL-like query language (since 0.8) and support search through secondary indexes

Tunable consistency and support for replication

Flexible schema

```
/usr/bin/cqlsh

Connected to Test Cluster at localhost:9160.

[cqlsh 4.1.1 | Cassandra 2.0.10 | CQL spec 3.1.1 | Thrift protocol 19.39.0]

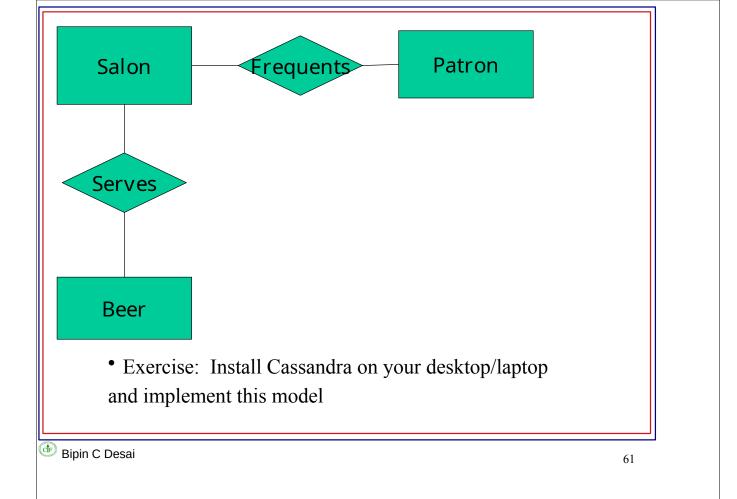
Use HELP for help.

cqlsh> CREATE KEYSPACE testkeyspc WITH REPLICATION = { 'class' : 'SimpleStrategy', 'replication_factor' : 1 };

cqlsh> use testkeyspc;

Keyspace is a "database"
```

Bipin C Desai 59



# Notes Files and Databases



BytePress

