CSC 214: INTRODUCTION TO DATA MANAGEMENT
Introduction to DBMS and Relational Theory

COURSE
   COURSE CODE: CSC 214
   COURSE TITLE: DATA MANAGEMENT I
   CREDIT UNIT: 2
   COURSE STATUS: COMPULSORY

INTRODUCTION TO THE COURSE

COURSE DESCRIPTION
CSC 214 IS A COMPREHENSIVE INTRODUCTION TO THEORIES AND PRACTICES IN DATABASE DESIGN AND MANAGEMENT. DATABASE MANAGEMENT I CONCENTRATES ON PRINCIPLES, DESIGN, IMPLEMENTATION AND APPLICATION OF DATABASE MANAGEMENT SYSTEMS. STUDENTS ARE INTRODUCED TO THE FUNDAMENTAL THEORIES, CONCEPTS AND TECHNIQUES NEEDED TO PROPERLY UNDERSTAND AND IMPLEMENT THE RELATIONAL DATABASE MODEL WHICH IS THE BEDROCK OF TODAY'S MAINSTREAM DATABASE PRODUCTS.

COURSE JUSTIFICATION
AN IN-DEPTH UNDERSTANDING OF THE PRINCIPLES AND APPLICATION OF DATABASE SYSTEMS IS A CRITICAL SUCCESS FACTOR FOR INFORMATION PROFESSIONALS TAKING LEADERSHIP ROLES IN FUTURE INFORMATION SYSTEMS INITIATIVES. THIS COURSE OFFERS STUDENTS THE OPPORTUNITY OF RIGOROUS STUDY OF THE TRADITIONAL PRINCIPLES OF DATABASE DESIGN, IMPLEMENTATION AND USAGE.

COURSE OBJECTIVES
UPON SUCCESSFUL COMPLETION, STUDENTS SHOULD BE ABLE TO:
- DEMONSTRATE GOOD KNOWLEDGE OF BASIC DATABASE CONCEPTS, INCLUDING THE STRUCTURE AND OPERATION OF THE RELATIONAL DATA MODEL.
- UNDERSTAND AND SUCCESSFULLY APPLY LOGICAL DATABASE DESIGN PRINCIPLES, INCLUDING E-R DIAGRAMS AND DATABASE NORMALIZATION.
- ASSESS THE QUALITY AND EASE OF USE OF DATA MODELLING AND DIAGRAMMING TOOLS.
- DESIGN AND IMPLEMENT A SMALL DATABASE PROJECT.
- DESCRIBE AND DISCUSS SELECTED ADVANCED DATABASE TOPICS, SUCH AS DISTRIBUTED DATABASE SYSTEMS AND THE DATA WAREHOUSE.

COURSE CONTENTS
INFORMATION STORAGE & RETRIEVAL, INFORMATION MANAGEMENT APPLICATIONS. INFORMATION CAPTURE AND REPRESENTATION, ANALYSIS & INDEXING, SEARCH, RETRIEVAL, INFORMATION PRIVACY; INTEGRITY, SECURITY; SCALABILITY, EFFICIENCY AND EFFECTIVENESS.
INTRODUCTION TO DATABASE SYSTEM:
COMPONENTS OF DATABASE SYSTEMS DBMS FUNCTIONS, DATABASE ARCHITECTURE AND DATA INDEPENDENCE. DATA MODELLING, ENTITY-RELATIONSHIP MODEL, DATABASE, DESIGN USING ENTITY-RELATIONSHIP AND SEMANTIC OBJECT MODELS, RELATIONAL DATA MODEL, PROCESS OF DATABASE DESIGN.

COURSE REQUIREMENT
THERE ARE NO FORMAL PREREQUISITES FOR THIS COURSE.

METHOD OF GRADING
METHOD OF GRADING

S/N	GRADING	SCORE(%)
1.	**CONTINUOUS ASSESSMENTS** • C.AI	7%

		• C.AII (Mid-Semester Test) • C.AIII	15% 8%
2.	Assignment		
3.	Practical (Laboratory work)/ Case Studies		
4.	Final Examination		70%
5.	Total		100

Course Delivery Strategies:

Lectures are delivered via electronic media (e-learning platform and power point presentations) and other available multimedia resources. Students are also encouraged to work with our programmers and avail themselves of laboratory facilities for practical work. Students are expected to demonstrate their understanding of concepts by completing given tasks in class and submitting assignments as at when due.

Resources used/Reading Material:

Books

- Database Management System (2Ed) by Raghu Ramakrishtan and Johannes Gehrke
- Database Systems: Design, Implementation, and Management (10Ed) by Carlos Coronel, Steven Morris, and Peter Rob (2012). Cengage Learning, Boston. ISBN-13: 978-1-111-96960-8
- Database principles and design (3Ed) by Colin Ritchie (2008). Cengage Learning, London. ISBN-13: 978-1-84480-540-2.
- Database System, the complete book (2Ed) by Hector G. M., Jeffrey D. U., Jennifer W. (2009). Pearson Education Inc. New Jersey. ISBN 0-13-606701-8
- Relational Theory for Computer Professionals by C. J. Date (2013). O'Reilly Media, Inc. Sebastopol. ISBN: 978-1-449-36943-9

Online resources

- Database Management Systems Relational, Object-Relational and Object-Oriented Data Models. Center for Objekt Teknology. Available online: http://www.cit.dk/COT/reports/reports/Case4/05-v1.1/cot-4-05-1.1.pdf
- http://www.help2engg.com/dbms/dbms-languages
- Database Management System by tutorialpoint. Available online https://www.tutorialspoint.com/dbms/dbms_tutorial.pdf

Data: raw representation of unprocessed facts, figures, concepts or instruction. It can exist in any form, usable or not. Data are facts presented without relation to other things. E.g. It is raining

Information: information is data that has been given meaning by way of relational connection. This "meaning" can be useful, but does not have to be. In computer parlance, a relational database makes information from the data stored within it. Information embodies the understanding of some sort. E.g. the temperature dropped to 15 degrees and then it started raining.

DATABASE

A Database is a shared, integrated computer structure that is repository to:
   -End-user data, that is, raw facts of interest to the end user
   -Metadata, or data about data describes of the data characteristics and the set of relationships that link the data found within the db.

A database is a collection of data, typically describing the activities of one or more related organizations. For example, a university database might contain information about the following:

- Entities such as students, faculty, courses, and classrooms.
- Relationships between entities, such as students' enrollment in courses, faculty teaching courses, and the use of rooms for courses.

Proper storage of data in a database will enhance efficient
- Data Management
- Data processing
- Data retrieval

Database System
Refers to an organization of components that define and regulate the collection, storage, management from general management point of view, the DB system is composed of
- Hardware
- Software
- People –system administrators: database systems operations
  - DB administrators: manage the DBMS and ensure the DB is
  - functioning properly
  - DB designers
  - System analysts and programmers design and implement the application programs
  - end user
- Procedures
- Data

DBMS: A database management system (DBMS) is a collection of programs that manages the database structure and controls access to the data stored in the database. In a sense, a database resembles a very well-organized electronic filing cabinet in which powerful software (the DBMS) helps manage the cabinet's contents

ADVANTAGES OF A DBMS

- **Data independence:** This is the technique that allow data to be changed without affecting the applications that process it. We can change the way the database is physically stored and accessed without having to make corresponding changes to the way the database is perceived by the user. Changing the way the database is physically stored and accessed is almost always to improve performance; and the fact that we can make such changes without having to change the way the database looks to the user means that existing application programs, queries, and the like can all still work after the change. Application programs should be as independent as possible from details of data representation and storage. The DBMS can provide an abstract view of the data to insulate application code from such details.
- **Efficient data access:** A DBMS deploys sophisticated techniques to store and retrieve data efficiently.
- **Data integrity control:** the DBMS can enforce integrity constraints on the data. For example, before inserting salary information for an employee, the DBMS can check that the department budget is not exceeded. Also, updating the status for supplier S1 to 200 will rejected, if status values are supposed never to exceed 100
- **Security Control:** the DBMS can enforce access controls that govern what data is visible to different classes of users. Users are only allowed to perform an operation he or she is allowed to carry out on data.
- **Concurrent access and crash recovery:** A DBMS schedules concurrent accesses to the data in such a manner that users can think of the data as being accessed by only one user at a time. Further, the DBMS protects users from the effects of system failures.
- **Reduced application development time:** Clearly, the DBMS supports many important functions that are common to many applications accessing data stored in the DBMS. This, in conjunction with the high-level interface to the data, facilitates quick development of applications. Such applications are also likely to be more robust than applications developed from scratch because

MANY IMPORTANT TASKS ARE HANDLED BY THE DBMS INSTEAD OF BEING IMPLEMENTED BY THE APPLICATION.
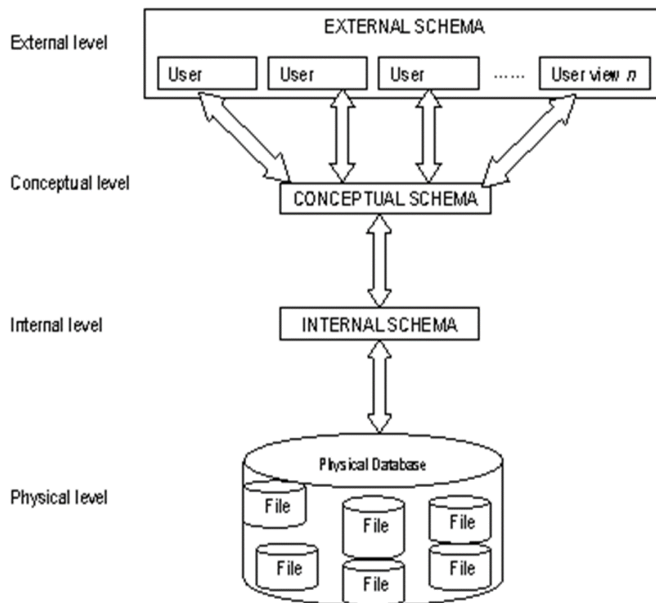
DBMS ARCHITECTURE

THE DBMS PROVIDES USERS WITH AN ABSTRACT VIEW OF THE DATA IN IT I.E. THE SYSTEM HIDES CERTAIN DETAILS OF HOW THE DATA IS STORED AND MAINTAINED FROM USERS. A DBMS CAN BE VIEWED AS DIVIDED INTO LEVELS OF ABSTRACTION. A COMMON ARCHITECTURE GENERALLY USED IS THE ANSI/SPARC (AMERICAN NATIONAL STANDARDS INSTITUTE - STANDARDS PLANNING AND REQUIREMENTS COMMITTEE) MODEL.

THE ANSI/SPARC MODEL ABSTRACTS THE DBMS INTO A 3-TIER ARCHITECTURE AS FOLLOWS:
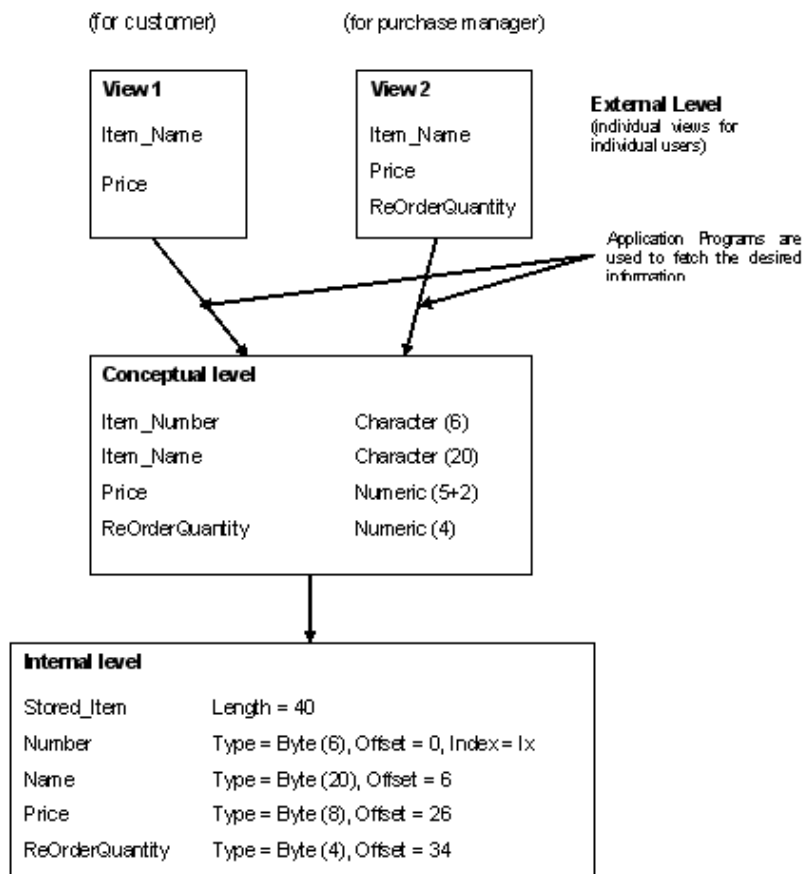EXTERNAL LEVEL
CONCEPTUAL LEVEL
INTERNAL LEVEL

ANSI/SPARC 3-TIER DBMS ARCHITECTURE

I. **EXTERNAL LEVEL:** THE EXTERNAL LEVEL IS THE USER'S VIEW OF THE DATABASE AND CLOSEST TO THE USERS. IT PRESENTS ONLY THE RELEVANT PART OF THE DBMS TO THE USER. E.G. A BANK DATABASE STORES A LOT MORE INFORMATION BUT AN ACCOUNT HOLDER IS ONLY INTERESTED IN HIS/HER ACCOUNT DETAILS SUCH AS THE CURRENT ACCOUNT BALANCE, TRANSACTION HISTORY ETC. AN EXTERNAL SCHEMA DESCRIBES EACH EXTERNAL VIEW. THE EXTERNAL SCHEMA CONSISTS OF THE DEFINITION OF THE LOGICAL RECORDS AND THE RELATIONSHIPS IN THE EXTERNAL VIEW. IN THE EXTERNAL LEVEL, THE DIFFERENT VIEWS MAY HAVE DIFFERENT REPRESENTATIONS OF THE SAME DATA.

II. **CONCEPTUAL LEVEL:** AT THIS LEVEL OF DATABASE ABSTRACTION, ALL THE DATABASE ENTITIES AND RELATIONSHIPS AMONG THEM ARE INCLUDED. CONCEPTUAL LEVEL PROVIDES THE COMMUNITY VIEW OF THE DATABASE AND DESCRIBES WHAT DATA IS STORED IN THE DATABASE AND THE RELATIONSHIPS AMONG THE DATA. IN OTHER WORDS, THE CONCEPTUAL VIEW REPRESENTS THE ENTIRE DATABASE OF AN ORGANIZATION. IT IS A COMPLETE VIEW OF THE DATA REQUIREMENTS OF THE ORGANIZATION THAT IS INDEPENDENT OF ANY STORAGE CONSIDERATION. THE CONCEPTUAL SCHEMA DEFINES CONCEPTUAL VIEW. IT IS ALSO CALLED THE LOGICAL SCHEMA. THERE IS ONLY ONE CONCEPTUAL SCHEMA PER DATABASE. THE FIGURE SHOWS THE CONCEPTUAL VIEW RECORD OF A DATA BASE.

III. **INTERNAL LEVEL OR PHYSICAL LEVEL:** THE LOWEST LEVEL OF ABSTRACTION IS THE INTERNAL LEVEL. IT IS THE ONE CLOSEST TO PHYSICAL STORAGE DEVICE. THIS LEVEL IS ALSO TERMED AS PHYSICAL LEVEL, BECAUSE IT DESCRIBES HOW DATA ARE ACTUALLY STORED ON THE STORAGE MEDIUM SUCH AS HARD DISK, MAGNETIC TAPE ETC. THIS LEVEL INDICATES HOW THE DATA WILL BE STORED IN THE DATABASE AND DESCRIBE THE DATA STRUCTURES, FILE STRUCTURES AND ACCESS METHODS TO BE USED BY THE DATABASE. THE INTERNAL SCHEMA DEFINES THE INTERNAL LEVEL. THE INTERNAL SCHEMA CONTAINS THE DEFINITION OF THE STORED RECORD, THE METHODS OF REPRESENTING THE DATA FIELDS AND ACCESSED METHODS USED. THE FIGURE SHOWS THE INTERNAL VIEW RECORD OF A DATABASE.

(for customer)   (for purchase manager)

View 1
Item_Name

Price

View 2
Item_Name

Price

ReOrderQuantity

External Level
(individual views for
individual users)

Application Programs are
used to fetch the desired
information

Conceptual level

Item_Number	Character (6)
Item_Name	Character (20)
Price	Numeric (5+2)
ReOrderQuantity	Numeric (4)

Internal level

Stored_Item	Length = 40
Number	Type = Byte (6), Offset = 0, Index = Ix
Name	Type = Byte (20), Offset = 6
Price	Type = Byte (8), Offset = 26
ReOrderQuantity	Type = Byte (4), Offset = 34

DBMS LANGUAGES
The workings of a DBMS is controlled by four different languages, they are

- **Data Definition Language (DDL):** Used by the DBA and database designers to specify the conceptual schema of a database. In many DBMSs, the DDL is also used to define internal and external schemas (views). In some DBMSs, separate storage definition language (SDL) and view definition language (VDL) are used to define internal and external schemas. SDL is typically realized via DBMS commands provided to the DBA and database designers. Some examples include:
    - CREATE - to create objects in the database
    - ALTER - alters the structure of the database
    - DROP - delete objects from the database
    - TRUNCATE - remove all records from a table, including all spaces allocated for the records are removed
    - COMMENT - add comments to the data dictionary
    - RENAME - rename an object

- **Data Manipulation Language (DML):** these statements managing data within schema objects. They specify database retrievals and updates. DML commands (data sublanguage) can be embedded in a general-purpose programming language (host language), such as COBOL, C, C++, or Java.
    - A library of functions can also be provided to access the DBMS from a
    - programming language
    - Alternatively, stand-alone DML commands can be applied directly (called a query
    - language).
  Some examples in SQL include:
    - SELECT - Retrieve data from the a database
    - INSERT - Insert data into a table

- ➢ **UPDATE** - Updates existing data within a table
- ➢ **DELETE** - deletes all records from a table, the space for the records remain
- ➢ **MERGE** - **UPSERT** operation (insert or update)
- ➢ **CALL** - Call a PL/SQL or Java subprogram
- ➢ **EXPLAIN PLAN** - explain access path to data
- ➢ **LOCK TABLE** - control concurrency

➢ **Data Control Language (DCL):** used for granting and revoking user access on a database
  - ➢ To grant access to user – **GRANT**
  - ➢ To revoke access from user – **REVOKE**

➢ **Transaction Control (TCL):** Statements are used to manage the changes made by DML statements. It allows statements to be grouped together into logical transactions.

Some examples include:
  - ➢ **COMMIT** - save work done
  - ➢ **SAVEPOINT** - identify a point in a transaction to which you can later roll back
  - ➢ **ROLLBACK** - restore database to original since the last **COMMIT**
  - ➢ **SET TRANSACTION** - Change transaction options like isolation level and what rollback segment to useIn practical data definition language, data manipulation language and data control languages are not separate language; rather they are the parts of a single database language such as **SQL**.

Example

Write the SQL code that will create the table structure for a table named **EMP_1**. This table is a subset of the **EMPLOYEE** table. The basic **EMP_1** table structure is summarized in the following table. **EMP_NUM** is the primary key and **JOB_CODE** is the FK to **JOB**.

Hint: Primary Key cannot contain null value

ATTRIBUTE (FIELD) NAME	DATA DECLARATION
EMP_NUM	CHAR(3)
EMP_LNAME	VARCHAR(15)
EMP_FNAME	VARCHAR(15)
EMP_INITIAL	CHAR(1)
EMP_HIREDATE	DATE
JOB_CODE	CHAR(3)

```
CREATE TABLE EMP_1(
EMP_NUM CHAR(6) NOT NULL,
EMP_LNAME VARCHAR(15),
EMP_FNAME VARCHAR(15),
EMP_INITIAL CHAR(1),
EMP_HIREDATE DATE,
JOB_CODE CHAR(3),
PRIMARY KEY (EMP_NUM),
FOREIGN KEY(JOB_CODE) REFERENCES JOB (JOB_CODE)
) ;
```

Having created the table structure in (a), write the SQL code to enter the first two rows for the table **EMP_1** below:

EMP_NUM	EMP_LNAME	EMP_FNAME	EMP_INITIAL	EMP_HIREDATE	JOB_CODE
101	News	John	G	08-Nov-00	502
102	Senior	David	H	12-Jul-89	501
103	Arbough	June	E	01-Dec-96	500
104	Ramoras	Anne	K	15-Nov-87	501
105	Johnson	Alice	K	01-Feb-93	502
106	Smithfield	William		22-Jun-04	500
107	Alonzo	Maria	D	10-Oct-93	500
108	Washington	Ralph	B	22-Aug-91	501
109	Smith	Larry	W	18-Jul-97	501

```
INSERT INTO EMP_1
(EMP_NUM, EMP_LNAME, EMP_FNAME, EMP_INITIAL, EMP_HIREDATE, JOB_CODE)
VALUES
("101", "News", "John", "G", "08-Nov-00", "502"),
("102", "Senior", "David", "H", "12-Jul-89", "500");
```

Assuming the data shown in the EMP_1 table have been entered, write the SQL code that will list all attributes for a job code of 502.

```
SELECT * FROM EMP_1
WHERE JOB_CODE = '502';
```

Write the SQL code that will save the changes made to the EMP_1 table.

```
COMMIT WORK;
```
NB: WORK is optional.

DBMS Data Model

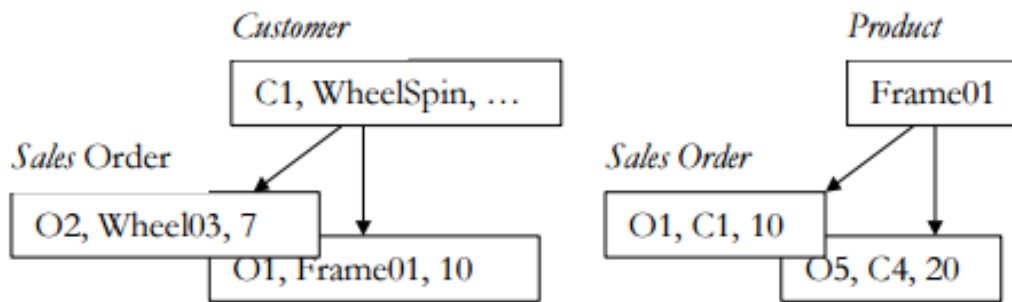A data model is a notation for describing data or information. The description generally consists of three parts:

I. Structure of the data: The data structures used to implement data in the computer are sometimes referred to, in discussions of database systems, as a physical data model.

II. Operations on the data: In database data models, there is usually a limited set of operations that can be performed but dba can describe database operations at a very high level, yet have the database management system implement the operations efficiently.

III. Constraints on the data. Database data models usually have a way to describe limitations on what the data can be. These constraints can range from the simple (e.g., "a day of the week is an integer between 1 and 7" or "a movie has at most one title") to some very complex limitations.

Traditionally, there are four DBMS. These four data models also represent the historical developments of the DBMS:

Hierarchical Database Model

This is the oldest DBMS data model. In this model, information is organized as a collection of inverted trees of records. The record at the root of a tree has zero or more child records; the child records, in turn, serve as parent records for their immediate descendants. This parent-child relationship recursively continues down the tree. The records consist of fields, where each field may contain simple data values (e.g. integer, real, text)., or a pointer to a record. The pointer graph is not allowed to contain cycles. Some combinations of fields may form the key for a record relative to its parent. Only a few hierarchical DBMSs support null values or

VARIABLE-LENGTH FIELDS.

Customer

| C1, WheelSpin, … |

Product

| Frame01 |

Sales Order

| O2, Wheel03, 7 |

| O1, Frame01, 10 |

Sales Order

| O1, C1, 10 |

| O5, C4, 20 |

EXAMPLE OF HIERARCHICAL DATA MODEL

APPLICATIONS CAN NAVIGATE A HIERARCHICAL DATABASE BY STARTING AT A ROOT AND SUCCESSIVELY NAVIGATE DOWNWARD FROM PARENT TO CHILDREN UNTIL THE DESIRED RECORD IS FOUND.

SEARCHING DOWN A HIERARCHICAL TREE IS VERY FAST SINCE THE STORAGE LAYER FOR HIERARCHICAL DATABASES USE CONTIGUOUS STORAGE FOR HIERARCHICAL STRUCTURES. ALL OTHER TYPES OF QUERIES REQUIRE SEQUENTIAL SEARCH TECHNIQUES. A DDL FOR HIERARCHICAL DATA MODEL MUST ALLOW THE DEFINITION OF RECORD TYPES, FIELDS TYPES, POINTERS, AND PARENT-CHILD RELATIONSHIPS. AND THE DML MUST SUPPORT DIRECT NAVIGATION USING THE PARENT-CHILD RELATIONSHIPS AND THROUGH POINTERS.

LIMITATIONS

➢ HIERARCHICAL MODEL ONLY PERMITS ONE TO MANY RELATIONSHIP. THE CONCEPT OF LOGICAL RELATIONSHIP IS OFTEN USED TO CIRCUMVENT THIS LIMITATION. LOGICAL RELATIONSHIP SUPERIMPOSE ANOTHER SET OF CONNECTION BETWEEN DATA ITEMS SEPARATE FROM THE PHYSICAL TREE STRUCTURE. THIS OF COURSE INCREASES ITS COMPLEXITY

➢ OFTEN A NATURAL HIERARCHY DOES NOT EXIST AND IT IS AWKWARD TO IMPOSE A PARENT-CHILD RELATIONSHIP. POINTERS PARTIALLY COMPENSATE FOR THIS WEAKNESS, BUT IT IS STILL DIFFICULT TO SPECIFY SUITABLE HIERARCHICAL SCHEMAS FOR LARGE MODELS AND THIS MEANS PROGRAMS HAVE TO NAVIGATE VERY CLOSE TO THE PHYSICAL DATA STRUCTURE LEVEL, IMPLYING THAT THE HIERARCHICAL DATA MODEL OFFERS ONLY VERY LIMITED DATA INDEPENDENCE.

➢ LACK OF AD HOC QUERY CAPABILITY PLACED BURDEN ON PROGRAMMERS TO GENERATE CODE FOR REPORTS

NETWORK MODEL:

IT REPRESENTS COMPLEX DATA RELATIONSHIPS MORE EFFECTIVELY THAN THE HIERARCHICAL MODEL. THE MAJOR IMPROVEMENT IS THAT THE ONE-TO-MANY LIMITATION WAS REMOVED; THE MODELS STILL VIEWS DATA IN A HIERARCHICAL ONE-TO-MANY STRUCTURE BUT NOW RECORD MAY HAVE MORE THAN ONE PARENT. NETWORK DATA MODELS REPRESENT DATA IN A SYMMETRIC MANNER, UNLIKE THE HIERARCHICAL DATA MODEL (DISTINCTION BETWEEN A PARENT AND A CHILD). INFORMATION IS ORGANIZED AS A COLLECTION OF GRAPHS OF RECORD THAT ARE RELATED WITH POINTERS. MORE FLEXIBLE THAN A HIERARCHICAL DATA MODEL AND STILL PERMITS EFFICIENT NAVIGATION.
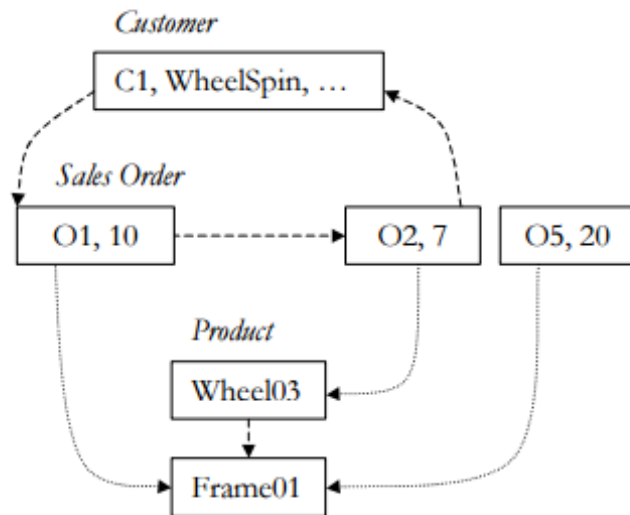
EXAMPLE OF NETWORK DATA MODEL

THE RECORDS CONSIST OF LISTS OF FIELDS (FIXED OR VARIABLE LENGTH WITH MAXIMUM LENGTH), WHERE EACH FIELD CONTAINS A SIMPLE VALUE (FIXED OR VARIABLE SIZE). THE NETWORK DATA MODEL ALSO INTRODUCES THE NOTION OF INDEXES OF FIELDS AND RECORDS, SETS OF POINTERS, AND PHYSICAL PLACEMENT OF RECORDS. A DDL FOR NETWORK DATA MODELS MUST ALLOW THE DEFINITION OF RECORD TYPES, FIELDS TYPES, POINTERS AND INDEXES. AND THE DML MUST ALLOW NAVIGATION THROUGH THE GRAPHS THROUGH THE POINTERS AND INDEXES. PROGRAMS ALSO NAVIGATES CLOSELY TO THE PHYSICAL

STORAGE STRUCTURES, IMPLYING THAT THE NETWORK DATA MODEL ONLY SUPPORTS LIMITED DATA INDEPENDENCE, AND ARE THEREFORE DIFFICULT TO MAINTAIN AS THE DATA MODELS EVOLVE OVER TIME.

CONCEPTS INTRODUCED UNDER THE NETWORK MODEL INCLUDE:

- SCHEMA: CONCEPTUAL DESIGN OF THE ENTIRE DATABASE USUALLY MANAGED BY THE DBA
- SUB-SCHEMA: VIRTUAL VIEW OF PORTIONS OF THE DATABASE VISIBLE TO APPLICATION PROGRAMMERS
- DATA MANAGEMENT LANGUAGE: ENABLES DEFINITION OF AND ACCESS TO THE SCHEMA AND SUB-SCHEMA. IT CONSIST OF DDL TO CONSTRUCT THE SCHEMA AND DML TO DEVELOP PROGRAMS
- DATA DEFINITION LANGUAGE

LIMITATIONS

- ➢ CUMBERSOME
- ➢ LACK OF AD HOC QUERY CAPABILITY PLACED BURDEN ON PROGRAMMERS TO GENERATE CODE FOR REPORTS
- ➢ STRUCTURAL CHANGE IN THE DATABASE COULD PRODUCE HAVOC IN ALL APPLICATION PROGRAMS

THE RELATIONAL DATABASE MODEL

DEVELOPED BY E.F. CODD (IBM) IN 1970, THE RELATIONAL DATA MODEL HAS A MATHEMATICAL FOUNDATION IN RELATIONAL ALGEBRA. THE MODEL IS BASED ON FIRST-ORDER PREDICATE LOGIC AND DEFINES A TABLE AS AN N-ARY RELATION. DATA IS ORGANIZED IN RELATIONS (TWO-DIMENSIONAL TABLES). EACH RELATION CONTAINS A SET OF TUPLES (RECORDS). EACH TUPLE CONTAINS A NUMBER OF FIELDS. A FIELD MAY CONTAIN A SIMPLE VALUE (FIXED OR VARIABLE SIZE) FROM SOME DOMAIN (E.G. INTEGER, REAL, TEXT, ETC.).

ADVANTAGES OF RELATIONAL MODEL

- BUILT-IN MULTILEVEL INTEGRITY: DATA INTEGRITY IS BUILT INTO THE MODEL AT THE FIELD LEVEL TO ENSURE THE ACCURACY OF THE DATA; AT THE TABLE LEVEL TO ENSURE THAT RECORDS ARE NOT DUPLICATED AND TO DETECT MISSING PRIMARY KEY VALUES; AT THE RELATIONSHIP LEVEL TO ENSURE THAT THE RELATIONSHIP BETWEEN A PAIR OF TABLES IS VALID; AND AT THE BUSINESS LEVEL TO ENSURE THAT THE DATA IS ACCURATE IN TERMS OF THE BUSINESS ITSELF. (INTEGRITY IS DISCUSSED IN DETAIL AS THE DESIGN PROCESS UNFOLDS.)

- **Logical and physical data independence from database applications:** Neither changes a user makes to the logical design of the database, nor changes a database software vendor makes to the physical implementation of the database, will adversely affect the applications built upon it.

- **Guaranteed data consistency and accuracy:** Data is consistent and accurate due to the various levels of integrity you can impose within the database. (This will become quite clear as you work through the design process.)

- **Easy data retrieval:** At the user's command, data can be retrieved either from a particular table or from any number of related tables within the database. This enables a user to view information in an almost unlimited number of ways.

One commonly perceived disadvantage of the relational database was that software programs based on it ran very slowly.

Some definitions

RELATION: A relation, as defined by E. F. Codd, is a set of tuples (d1, d2, ..., dn), where each element dj is a member of Dj, a data domain, for each j=1, 2, ..., n. A **data domain** is simply a data type. It specifies a data abstraction: the possible values for the data and the operations available on the data. For example, a String can have zero or more characters in it, and has operations for comparing strings, concatenating string, and creating strings. A relation is a truth predicate. It defines what attributes are involved in the predicate and what the meaning of the predicate is. In relational data model, relations are represented in the table format. This format stores the relation among entities. A table has rows and columns, where rows represent records and columns represent the attributes. E.g.

title	year	length	genre
Gone With the Wind	1939	231	drama
Star Wars	1977	124	sciFi
Wayne's World	1992	95	comedy

TUPLE: A single row of a table, which contains a single record for that relation is called a tuple. A tuple has attribute values which match the required attributes in the relation. The ordering of attribute values is immaterial. Every tuple in the body of a given relation is required to conform to the heading (attribute) of that relation, i.e. it contains exactly one value, of the applicable type, for each attribute, and nothing else besides

ATTRIBUTE: The columns of a relation are named by attributes. Attributes appear at the tops of the columns. Usually, an attribute describes the meaning of entries in the column below. For instance, the column with attribute length holds the length, in minutes, of each movie.

ATTRIBUTE DOMAIN: Every attribute has some predefined value scope, known as attribute domain

ATTRIBUTE VALUE/INSTANCE: An attribute value is the value for an attribute in a particular tuple. An attribute value must come from the domain that the attribute specifies. Most relational DBMS allows NULL attribute values. Each attribute value in a relational model must be atomic i.e. it must be of some elementary type such as integer or string. It is not permitted for a value to be a record structure, set, list, array, or any other type that reasonably can have its values broken into smaller components.

SCHEMAS: The name of a relation and the set of attributes for a relation is called the schema for that relation. The schema is depicted by the relation name followed

by a parenthesized list of its attributes. Thus, the schema for relation Movies above is

      Movies (title , year, length, genre)

In the relational model, a database consists of one or more relations. The set of schemas for the relations of a database is called a relational database schema, or just a database schema.

Data Types: All attributes must have a data type. The following are the primitive data types that are supported by **SQL** (Structured Query Language) systems.

    i.   Character strings of fixed or varying length. The type **CHAR**(n) denotes a fixed-length string of up to n characters. **VARCHAR**(n) also denotes a string of up to n characters. The difference is implementation-dependent; typically **CHAR** implies that short strings are padded to make n characters, while **VARCHAR** implies that an endmarker or string-length is used. Normally, a string is padded by trailing blanks if it becomes the value of a component that is a fixed-length string of greater length. For example, the string 'foo' if it became the value of a component for an attribute of type **CHAR**(5), would assume the value 'foo ' (with two blanks following the second o).

    ii.   Bit strings of fixed or varying length. These strings are analogous to fixed and varying-length character strings, but their values are strings of bits rather than characters. The type **BIT** (n) denotes bit strings of length n, while **BIT VARYING** (n) denotes bit strings of length up to n.

    iii.   The type **BOOLEAN** denotes an attribute whose value is logical. The possible values of such an attribute are **TRUE, FALSE.**

    iv.   The type **INT** or **INTEGER** (these names are synonyms) denotes typical integer values. The type **SHORTINT** also denotes integers, but the number of bits permitted may be less, depending on the implementation (as with the types int and short int in **C**).

    v.   Floating-point numbers can be represented in a variety of ways. We may use the type **FLOAT** or **REAL** (these are synonyms) for typical floating point numbers. A higher precision can be obtained with the type **DOUBLE PRECISION**. We can also specify real numbers with a fixed decimal point. For example, **DECIMAL**(n,d) allows values that consist of n decimal digits, with the decimal point assumed to be d positions from the right. Thus, 0123.45 is a possible value of type **DECIMAL**(6,2). **NUMERIC** is almost a synonym for **DECIMAL**, although there are possible implementation-dependent differences.

    vi.   Dates and times can be represented by the data types **DATE** and **TIME**, respectively. These values are essentially character strings of a special form. We may, in fact, coerce dates and times to string types, and we may do the reverse if the string "makes sense" as a date or time.

A relational database Schema is depicted by stating both the attributes and their datatype:

Movies (
title CHAR(100),
year INT,
length INT,
genre CHAR(10),
studioName CHAR(30),
producer INT
)

Relation instance: A finite set of tuples in the relational database system represents relation instance. Relation instances do not have duplicate tuples.

{
      <Person SSN# = "123-45-6789" Name = "Art Larsson" City = "San Francisco">,
      <Person SSN# = "231-45-6789" Name = "Lino Buchanan" City = "Philadelphia">,
      <Person SSN# = "321-45-6789" Name = "Diego Jablonski" City = "Chicago">
}

It is more common and concise to show a relation value as a table. All ordering within the table is artificial and meaningless.

Design theory for Relational Database

A common problem with schema design involve trying to combine too much into one relation thus leading to redundancy. Thus, improvements to relational schemas pay close attention to eliminating redundancy. The theory of "dependences" is a well-developed theory for relational databases providing guidelines on how to develop good schema and eliminate flaws if any. The first concept we need to consider is Functional Dependency (FD).

Functional Dependency (FD)

FUNCTIONAL DEPENDENCY: the term functional dependence can be defined most easily this way:

Definition:

Let A and B be subsets of the attribute of a relation R. Then the functional dependency (FD)

$A \rightarrow B$

holds in R if and only if, whenever two tuples of R have the same value for A, they also have the same value for B. A and B are the determinant and the dependent, respectively, and the FD overall can be read as "A functionally determines B" or "B is functionally dependent on A," or more simply just as
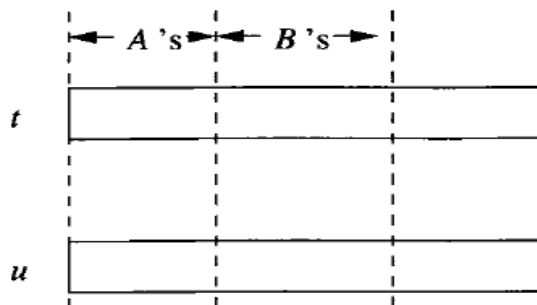
$A \rightarrow B$

If A and B are composite, then we have

$A1, A2, ..., An \rightarrow B1, B2, ..., Bm$

This is also equivalent to

$A1, A2, ..., An \rightarrow B1, A1, A2, ..., An \rightarrow B2, ..., A1, A2, ..., An \rightarrow Bm$

The attribute(s) B is functionally dependent on attributes(s)A, if A determines B.
e.g. STU_PHONE is functionally dependent on STU_NUM.

STU_NUM is not functionally dependent on STU_PHONE because the STU_PHONE value 2267 is associated with two STU_NUM values: 324274 and 324291. (This could happen when roommates share a single land line phone number.)

The functional dependence definition can be generalized to cover the case in which the determining attribute values occur more than once in a table.

Functional dependence can then be defined this way:

Attribute B is functionally dependent on A if all of the rows in the table that agree in value for attribute A also agree in value for attribute B.

RELATION KEYS: The key's role is based on a concept known as **determination**. I.e. the statement "A determines B" indicates that if you know the value of attribute A, you can look up (determine) the value of attribute B. E.g.:

AN INVOICE NUMBER IDENTIFIES ALL OF THE INVOICE ATTRIBUTES SUCH AS INVOICE DATE AND THE CUSTOMER NAME.

IF WE KNOW STU_NUM IN A STUDENT TABLE WE CAN LOOK UP (DETERMINE) STUDENT'S LAST NAME, GRADE POINT AVERAGE, PHONE NUMBER, ETC.

STU_NUM	STU_LNAME	STU_FNAME	STU_INIT	STU_DOB	STU_HRS	STU_CLASS	STU_GPA	STU_TRANSFER	DEPT_CODE	STU_PHONE	PROF_NUM
321452	Bowser	William	C	12-Feb-1975	42	So	2.84	No	BIOL	2134	205
324257	Smithson	Anne	K	15-Nov-1981	81	Jr	3.27	Yes	CIS	2256	222
324258	Brewer	Juliette		23-Aug-1969	36	So	2.26	Yes	ACCT	2256	228
324269	Oblonski	Walter	H	16-Sep-1976	66	Jr	3.09	No	CIS	2114	222
324273	Smith	John	D	30-Dec-1958	102	Sr	2.11	Yes	ENGL	2231	199
324274	Katinga	Raphael	P	21-Oct-1979	114	Sr	3.15	No	ACCT	2267	228
324291	Robertson	Gerald	T	08-Apr-1973	120	Sr	3.87	No	EDU	2267	311
324299	Smith	John	B	30-Nov-1986	15	Fr	2.92	No	ACCT	2315	230

TABLE NAME: STUDENT

THE SHORTHAND NOTATION FOR "A DETERMINES B" IS

A → B.

IF A DETERMINES B, C, AND D, WE WRITE

A → B, C, D.

FOR THE STUDENT EXAMPLE WE CAN WRITE:

STU_NUM → STU_LNAME, STU_FNAME, STU_INIT, STU_DOB, STU_TRANSFER

IN CONTRAST, STU_NUM IS NOT DETERMINED BY STU_LNAME BECAUSE IT IS QUITE POSSIBLE FOR SEVERAL STUDENTS TO HAVE THE LAST NAME SMITH.

PROPER UNDERSTANDING OF THE PRINCIPLE OF DETERMINATION IS VITAL TO THE UNDERSTANDING OF A CENTRAL RELATIONAL DATABASE CONCEPT KNOWN AS **FUNCTIONAL DEPENDENCE (FD)**.

DEFINITIONS

KEY ATTRIBUTE(S): WE SAY A SET OF ONE OR MORE ATTRIBUTES {A1, A2, ..., AN} IS A **KEY FOR A RELATION R** IF:

I. THOSE ATTRIBUTES FUNCTIONALLY DETERMINE ALL OTHER ATTRIBUTES OF THE RELATION. THAT IS, IT IS IMPOSSIBLE FOR TWO DISTINCT TUPLES OF R TO AGREE ON ALL OF A1, A2, ..., AN (UNIQUENESS).

II. NO PROPER SUBSET OF {A1, A2, ..., AN} FUNCTIONALLY DETERMINES ALL OTHER ATTRIBUTES OF R; I.E., A KEY MUST BE MINIMAL.

WHEN A KEY CONSISTS OF A SINGLE ATTRIBUTE A, WE OFTEN SAY THAT A (RATHER THAN {A}) IS A KEY. AN ATTRIBUTE THAT IS PART OF A KEY IS CALLED **KEY ATTRIBUTE**.

CONSIDER THE RELATION MOVIES BELOW:

title	year	length	genre	studioName	starName
Star Wars	1977	124	SciFi	Fox	Carrie Fisher
Star Wars	1977	124	SciFi	Fox	Mark Hamill
Star Wars	1977	124	SciFi	Fox	Harrison Ford
Gone With the Wind	1939	231	drama	MGM	Vivien Leigh
Wayne's World	1992	95	comedy	Paramount	Dana Carvey
Wayne's World	1992	95	comedy	Paramount	Mike Meyers

ATTRIBUTES {TITLE, YEAR, STARNAME} FORM A KEY FOR THE RELATION MOVIES BECAUSE IT MEETS THE TWO CONDITIONS:

CONDITION 1:

DO THEY FUNCTIONALLY DETERMINE ALL THE OTHER ATTRIBUTES? YES

CONDITION 2:

DO ANY PROPER SUBSET OF {TITLE, YEAR, STARNAME} FUNCTIONALLY DETERMINES ALL OTHER ATTRIBUTES?

{TITLE, YEAR} DO NOT DETERMINE STARNAME THUS {TITLE, YEAR} IS NOT A KEY.

{YEAR, STARNAME} IS NOT A KEY BECAUSE WE COULD HAVE A STAR IN TWO MOVIES IN THE SAME YEAR; THEREFORE

{YEAR, STARNAME} → TITLE IS NOT AN FD.

{title, starName} is not a key, because two movies with the same title, made in different years, can have a star in common.

Therefore, no proper subset of {title, year, starName} functionally determines all other attributes

Super Key (shortened: super set of keys): An attribute or a combination of attributes that is used to identify the records uniquely is known as Super Key. It is to be noted that some superkeys are not (minimal) keys. Note that every superkey satisfies the first condition of a key: it functionally determines all other attributes of the relation. However, a superkey need not satisfy the second condition: minimality. A table can have many Super Keys. E.g. of Super Key

- ID
- ID, Name
- ID, Address
- ID, Department_ID
- ID, Salary
- Name, Address

Candidate Key: It can be defined as minimal Super Key or irreducible Super Key. In other words an attribute or a combination of attribute that identifies the record uniquely but none of its proper subsets can identify the records uniquely. E.g. of Candidate Key

Code
Name, Address

Primary Key: A Candidate Key that is used by the database designer for unique identification of each row in a table is known as Primary Key. A Primary Key can consist of one or more attributes of a table. E.g. of Primary Key - Database designer can use one of the Candidate Key as a Primary Key.

In this case we have "Code" and "Name, Address" as Candidate Key,
The designer may prefer "Code" as the Primary Key as the other key is the combination of more than one attribute.

Null values should never be part of a primary key, they should also be avoided to the greatest extent possible in other attributes too. A null is no value at all. It does not mean a zero or a space. There are rare cases in which nulls cannot be reasonably avoided when you are working with non-key attributes. For example, one of an EMPLOYEE table's attributes is likely to be the EMP_INITIAL. However, some employees do not have a middle initial. Therefore, some of the EMP_INITIAL values may be null. Null can also exist because of the nature of the relationship between two entities. Conventionally, the existence of nulls in a table is often an indication of poor database design. Nulls, if used improperly, can create problems because they have many different meanings. For example, a null can represent:

An unknown attribute value.
A known, but missing, attribute value.
A "not applicable" condition.

Foreign Key: A foreign key is an attribute or combination of attributes in one base table that points to the candidate key (generally it is the primary key) of another table. The purpose of the foreign key is to ensure referential integrity of the data i.e. only values that are supposed to appear in the database are permitted. E.g.

Consider two table
Employee (EmployeeID, EmployeeName, DOB, DOJ, SSN, DeptID, MgrID) and
DeptTbl (Dept_ID, Dept_Name, Manager_ID, Location_ID)

Dept_ID is the primary key in Table DeptTbl, the DeptID attribute of table Employee (dependent or child table) can be defined as the Foreign Key as it can reference to the Dept_ID attribute of the table DeptTbl (the referenced or parent table), a Foreign Key value must match an existing value in the parent table or be NULL.

COMPOSITE KEY: If we use multiple attributes to create a Primary Key then that Primary Key is called Composite Key (also called a Compound Key or Concatenated Key).

FULL FUNCTIONAL DEPENDENCY (FFD): If the attribute (B) is functionally dependent on a composite key (A) but not on any subset of that composite key, the attribute (B) is fully functionally dependent on (A).

ALTERNATE KEY: Alternate Key can be any of the Candidate Keys except for the Primary Key.

SECONDARY KEY: The attributes that are not even the Super Key but can be still used for identification of records (not unique) are known as Secondary Key.
E.g. of Secondary Key can be Name, Address, Salary, Department_ID etc. as they can identify the records but they might not be unique.

Table name: PRODUCT   Database name: Ch03_SaleCo
Primary key: PROD_CODE
Foreign key: VEND_CODE

PROD_CODE	PROD_DESCRIPT	PROD_PRICE	PROD_ON_HAND	VEND_CODE
001278-AB	Claw hammer	12.95	23	232
123-21UUY	Houselite chain saw, 16-in. bar	189.99	4	235
QER-34256	Sledge hammer, 16-lb. head	18.63	6	231
SRE-657UG	Rat-tail file	2.99	15	232
ZZX/3245Q	Steel tape, 12-ft. length	6.79	8	235

link

VEND_CODE	VEND_CONTACT	VEND_AREACODE	VEND_PHONE
230	Shelly K. Smithson	608	555-1234
231	James Johnson	615	123-4536
232	Annelise Crystall	608	224-2134
233	Candice Wallace	904	342-6567
234	Arthur Jones	615	123-3324
235	Henry Ortozo	615	899-3425

Table name: VENDOR
Primary key: VEND_CODE
Foreign key: none

AN EXAMPLE OF RELATIONAL DB WITH PRIMARY KEY AND FOREIGN KEY

EXERCISE
Suppose R is a relation with attributes A1, A2, ..., An. As a function of n, tell how many superkeys R has, if:
a) The only key is A1.
b) The only keys are A1 and A2
c) The only keys are {A1, A2} and {A3, A4}
d) The only keys are {A1, A2} and {A1, A3}

RULES ABOUT FUNCTIONAL DEPENDENCIES

These rules guide us on how we can infer a functional dependency from other given FD's.
E.g., given that a relation R (A, B, C) satisfies the FD's
A —> B and B —> C,
then we can deduce that R also satisfies the FD
A —> C.

PROOF:
Consider two tuples of R that agree on A

Let the tuples agreeing on attribute A be (a, b1, c1) and (a, b2, c2)

Since R satisfies A → B, and these tuples agree on A, they must also agree on B. That is, b1 = b2

The tuples are now (a, b, c1) and (a, b, c2), where b is both b1 and b2.
Similarly, since R satisfies B → C, and the tuples agree on B, they agree also on C.
Thus, c1= c2; i.e., the tuples do agree on C.

We have proved that any two tuples of R that agree on A also agree on C, and that is the FD
A → C.
This rule is called the transitive rule

The Splitting/Combining Rule
Recall that the FD:
A1, A2, ..., An → B1, B2, ..., Bm
is equivalent to the set of FD's:
A1, A2, ..., An → B1, A1, A2, ..., An → B2, ..., A1, A2, ..., An → Bm

In other words, we may split attributes on the right side so that only one attribute appears on the right of each FD. Likewise, we can replace a collection of FD's having a common left side by a single FD with the same left side and all the right sides combined into one set of attributes. In either event, the new set of FD's is equivalent to the old. The equivalence noted above can be used in two ways.

- We can replace an FD

    A1, A2, ..., An → B1, B2, ..., Bm by a set of FD's
    A1, A2, ..., An → Bi for i = 1, 2, ..., m
    We call this transformation the splitting rule.

- We can replace a set of FD's

    A1, A2, ..., An → Bi for i = 1, 2, ..., m by the single FD
    A1, A2, ..., An → B1, B2, ..., Bm.
    We call this transformation the combining rule.

E.g. the set of FD's:

    title year → length
    title year → genre
    title year → studioName
is equivalent to the single FD:

    title year → length, genre, studioName

The splitting/ combining rule is stated as follows:
Suppose we have two tuples that agree in A1, A2, ..., An. As a single FD, we would assert "then the tuples must agree in all of B1, B2, ..., Bm." As individual FD's, we assert "then the tuples agree in B1, and they agree in B2, and, ..., and they agree in Bm."

Trivial-dependency rule.

Trivial Functional Dependencies: If a functional dependency (FD) $\alpha \rightarrow \beta$ holds in Relation R, then the term trivial is attached to the dependency if it is satisfied by all possible r(R)
i.e. $\alpha \rightarrow \beta$ is trivial if $\beta \subseteq \alpha$ or $\beta \cup \alpha = R$
where β is a subset of α, then it is called a trivial FD.
e.g.
title, year → title
title → title
are both trivial FD

There is an intermediate situation in which some, but not all, of the attributes on the right side of an FD are also on the left. This FD is not trivial.

**Non-trivial:** If an FD X → Y holds, where Y is not a subset of X, then it is called a non-trivial FD.

This can be simplified by removing from the right side of an FD those attributes that appear on the left. That is: The FD
A1, A2, ..., An → B1, B2, ..., Bm is equivalent to
A1, A2, ..., An → C1, C2, ..., Ck
where the C's are all those B's that are not also A's.

Completely non-trivial: If an FD X → Y holds, where x intersect Y = $\emptyset$, it is said to be a completely non-trivial FD.
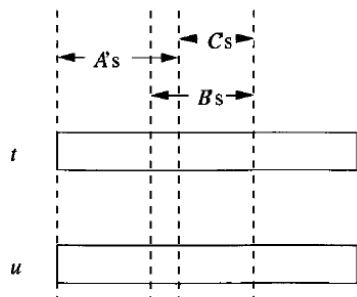
If t and   Then they
u agree   must agree
on the A's on the Bs

So surely
they agree
on the Cs

Trivial dependency rule

Computing the Closure of Attributes

Given a set α = {A1, A2, ..., A n} of attributes of R and a set of functional dependencies FD, we need a way to find all of the attributes of R that are functionally determined by α. This set of attributes is called the closure of α under F and is denoted α^+. Finding α^+ is useful because:

- if α^+ = R, then α is a superkey for R
- With closure we can find all FD's easily
    - To check if X → A
    - Compute X+
    - Check if A $\in$ X
- if we find α^+ for all $\alpha \subseteq$ R, we've computed F$^+$ (except that we'd need to use decomposition to get all of it).

Formal definition of closure:
Suppose α = {A1, A2, ..., An} is a set of attributes and S is a set of FD's. The closure of α under the FD's in S is the set of attributes B such that every relation that satisfies all the FD's in set S also satisfies A1, A2, ..., An → B. That is, A1, A2, ..., An → B follows from the FD's of S.

We denote the closure of a set of attributes A1, A2, ..., An by
{A1, A2, ..., An}$^+$.
Note that A1, A2, ..., An are always in {A1, A2, ..., An}$^+$ because the FD A1, A2, ..., An → Ai is trivial when i is one of 1,2,... , n.
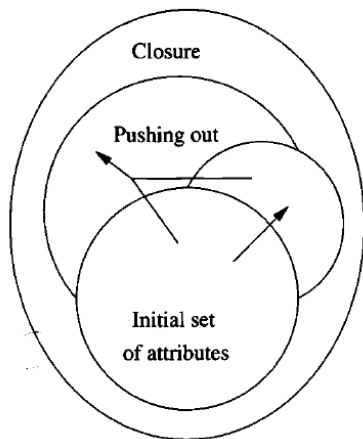
: Computing the closure of a set of attributes

THE FIGURE ABOVE ILLUSTRATES THE CLOSURE PROCESS:

STARTING WITH THE GIVEN SET OF ATTRIBUTES, WE REPEATEDLY EXPAND THE SET BY ADDING THE RIGHT SIDES OF FD'S AS SOON AS WE HAVE INCLUDED THEIR LEFT SIDES. EVENTUALLY, WE CANNOT EXPAND THE SET ANY FURTHER, AND THE RESULTING SET IS THE CLOSURE.

AN ALGORITHM FOR COMPUTING α^+:

```
result := α
repeat
        temp := result
        for each functional dependency  β → γ  in F do
        if  β ⊆ result then
                result := result ∪ γ
until temp = result
```

EXAMPLE:

CONSIDER A RELATION WITH ATTRIBUTES A, B, C, D, E, AND F. SUPPOSE THAT THIS RELATION HAS THE FD'S

AB → C, BC → AD, D → E, AND CF → B.

WHAT IS THE CLOSURE OF {A, B}?

SOLUTION

FIRST, SPLIT BC → AD INTO BC → A AND BC → D.

RESULT = {A, B}.

FOR AB → C

    AB ⊆ RESULT, SO WE HAVE

    RESULT = RESULT ∪ C I.E. RESULT = {A, B, C}.

FOR BC → C AND BC → D

    BC ⊆ RESULT, SO WE HAVE

    RESULT = RESULT ∪ A AND D I.E., RESULT = {A, B, C, D}

FOR D → E

    D ⊆ RESULT, SO WE HAVE

    RESULT = RESULT ∪ E I.E. RESULT = {A, B, C, D, E}

NO MORE CHANGES TO RESULT ARE POSSIBLE, THUS, {A, B}$^+$ = {A, B, C, D, E}.

BY COMPUTING THE CLOSURE OF ANY SET OF ATTRIBUTES, WE CAN TEST WHETHER ANY GIVEN FD A1, A2, ..., AN → B FOLLOWS FROM A SET OF FD'S S.

FIRST COMPUTE {A1, A2, ..., AN}$^+$ USING THE SET OF FD'S S. IF B IS IN {A1, A2, ..., AN}$^+$, THEN A1, A2, ..., AN → B DOES FOLLOW FROM S, AND IF B IS NOT IN {A1, A2, ..., AN}$^+$, THEN THIS FD DOES NOT FOLLOW FROM S.

More generally, A1, A2, ..., An → B1, B2, ..., Bm follows from set of FD's S if and only if all of B1, B2, ..., Bm are in $\{A1, A2, ..., An\}^+$

Example:
Consider the relation and FD's in the example above, Suppose we wish to test whether AB → D follows from these FD's. We compute $\{A, B\}^+$, which is {A, B, C, D, E}. Since D is a member of the closure, we conclude that AB → D does follow.
On the other hand, consider the FD
D → A. To test whether this FD follows from the given FD's, first compute $\{D\}^+$.
$\{D\}^+$ = {D, E}. Since A is not a member of {D, E}, we conclude that D → A does not follow.

Armstrong's Axioms

If F is a set of functional dependencies then the closure of F, denoted as F^+, is the set of all functional dependencies logically implied by F. Armstrong's Axioms are a set of rules, that when applied repeatedly, generates a closure of functional dependencies.

- **Reflexivity / reflexive rule:** If $\{B1, B2, ..., Bm\} \subseteq \{A1, A2, ..., An\}$, then A1, A2, ..., An → B1, B2, ..., Bm. These are what we have called trivial FD's.
- **Augmentation rule:** If A1A2 ... An → B1B2 ... Bm, then A1A2 ... AnC1C2 ... Ck → B1B2, ... BmC1C2 ... Ck for any set of attributes C1, C2, ..., Ck
  Since some of the C 's may also be A's or B's or both, we should eliminate from the left side duplicate attributes and do the same for the right side.
- **Transitivity rule:** If A1, A2, ..., An → B1, B2, ..., Bm and B1, B2, ..., Bm → C1, C2, ..., Ck hold in relation R, then A1, A2, ..., An → C1, C2, ..., Ck also holds in R.

  If some of the C 's are among the A's, we may eliminate them from the right side by the trivial-dependencies rule

  To test whether A1, A2, ..., An → C1, C2, ..., Ck holds,
  we need to compute the closure
  $\{A1, A2, ..., An\}^+$ with respect to the two given FD's.
  The FD A1, A2, ..., An → B1, B2, ..., Bm tells us that all of B1, B2, ..., B m are in $\{A1, A2, ..., A n\}^+$.
  Then, we can use the FD B1, B2, ..., Bm → C1, C2, ..., Ck to add C1, C2, ..., Ck to $\{A1, A2, ..., An\}^+$.
  Since all the C's are in $\{A1, A2, ..., An\}^+$ we conclude that
  A1, A2, ..., An → C1, C2, ..., Ck holds for any relation that satisfies both
  A1, A2, ..., An → B1, B2, ..., Bm and B1, B2, ..., Bm → C1, C2, ..., Ck.

Additional rules:

- **Union:** If X → Y and X → Z, then X → Y Z
- **Pseudotransitivity:** If X → Y and W Y → Z, then W X → Z
- **Composition:** If X → Y and Z → W, then XZ → Y W

Transitive dependence: an attribute Y is said to be transitively dependent on attribute X if Y is functionally dependent on another attribute Z which is functionally dependent on X.

Closure of FD's set
Given Relation R and a set of FD's F that holds in R:
The closure of F in R (denoted F^+) is the set of all FD's F in R that are logically implied by F i.e. s the set of all regular FDs that can be derived from F

```
algorithm (F)
 /* F is a set of FDs */
F⁺ = ∅
for each possible attribute set X
      Compute the closure X⁺ of X on F
      for each attribute A ∈ X⁺
            add to F⁺ the FD: X → A
return F⁺
```

EXAMPLE:

ASSUME THERE ARE 4 ATTRIBUTES A, B, C, D AND THAT $F = \{A \rightarrow B, B \rightarrow C\}$. TO COMPUTE F+ WE FIRST GET:

$A^+ = AB^+ = AC^+ = ABC^+ = \{A, B, C\}$
$B^+ = BC^+ = \{B, C\}$
$C^+ = \{C\}$
$D^+ = \{D\}$
$AD^+ = \{A, D\}$
$BC^+ = \{B, C\}$
$BD^+ = BCD^+ = \{B, C, D\}$
$ABD^+ = ABCD^+ = \{A, B, C, D\}$
$ACD^+ = \{A, C, D\}$

EXERCISE

CONSIDER A RELATION WITH SCHEMA R (A, B, C, D) AND FD'S $AB \rightarrow C$, $D \rightarrow D$ AND $D \rightarrow A$.

  I.   WHAT ARE ALL THE NONTRIVIAL FD'S THAT FOLLOW FROM THE GIVEN FD'S? YOU SHOULD RESTRICT YOURSELF TO FD'S WITH SINGLE ATTRIBUTES ON THE RIGHT SIDE.
  II.   WHAT ARE ALL THE KEYS OF R?
  III. WHAT ARE ALL THE SUPERKEYS FOR R THAT ARE NOT KEYS?

RELATIONAL SET OPERATORS

THE RELATIONAL DATA MODEL IS DESIGNED IN SUCH WAYS THAT DATA MAY BE PROCESSED WITH MATHEMATICAL OPERATIONS. DATA IN RELATIONAL TABLES ARE OF LITTLE USE UNLESS THEY ARE MANIPULATED TO YIELD MEANINGFUL INFORMATION. RELATIONAL ALGEBRA FORMS THE THEORETICAL BASIS FOR MANIPULATING TABLE CONTENT USING EIGHT OPERATORS; FOUR RELATIONAL OPERATORS AND FOUR SET OPERATORS. RELATIONAL OPERATORS TAKE ONE OR TWO RELATIONS AS INPUTS AND RETURN RELATIONS AS THE RESULT WHILE SET OPERATORS TAKE ONE OR TWO SETS AS INPUTS AND RETURN SETS AS THE RESULT.

 FOUR RELATIONAL OPERATIONS
   ➢ PROJECT
   ➢ SELECT
   ➢ JOIN
   ➢ DIVISION
FOUR SET OPERATIONS
   ➢ UNION
   ➢ DIFFERENCE
   ➢ INTERSECTION
   ➢ CARTESIAN PRODUCT

VERY FEW DBMSS ARE CAPABLE OF SUPPORTING ALL EIGHT RELATIONAL OPERATORS. TO BE CONSIDERED MINIMALLY RELATIONAL, THE DBMS MUST SUPPORT THE KEY RELATIONAL OPERATORS SELECT, PROJECT, AND JOIN.

  1. SELECT, ALSO KNOWN AS RESTRICT, YIELDS VALUES FOR ALL THE ROWS FOUND IN A TABLE THAT SATISFY A GIVEN CONDITION. SELECT YIELDS A HORIZONTAL SUBSET OF A TABLE.

2. **PROJECT** yields all values for selected attributes. **PROJECT** yields a vertical subset of a table

3. **UNION**: combines all rows from two or more tables, excluding duplicate rows. In order to be used in a **UNION**, the tables must be **UNION** compatible, that is:
    - ➤ The relations must all have the same number of attributes.
    - ➤ Corresponding columns must all have identical data types and lengths.
    When these criteria are met, the tables are said to be union compatible.

4. **INTERSECT**: yields only the rows that appear in both tables. As with **UNION**, the tables must be union-compatible to yield valid results.

5. **DIFFERENCE**: yields all rows in one table that are not found in the other table. As with the **UNION**, the tables must be **UNION**-compatible to yield valid results.

| STU_FNAME | STU_LNAME | | DIFFERENCE | EMP_FNAME | EMP_LNAME | | yields | | STU_FNAME | STU_LNAME |

STU_FNAME	STU_LNAME
George	Jones
Jane	Smith
Peter	Robinson
Franklin	Johnson
Martin	Lopez

DIFFERENCE

EMP_FNAME	EMP_LNAME
Franklin	Lopez
William	Turner
Franklin	Johnson
Susan	Rogers

yields →

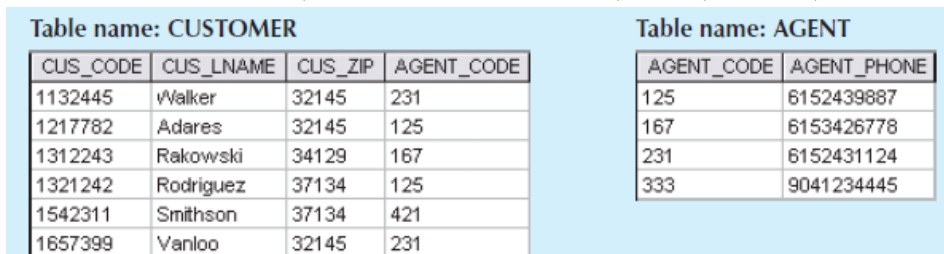| STU_FNAME	STU_LNAME
George	Jones
Jane	Smith
Peter	Robinson
Martin	Lopez
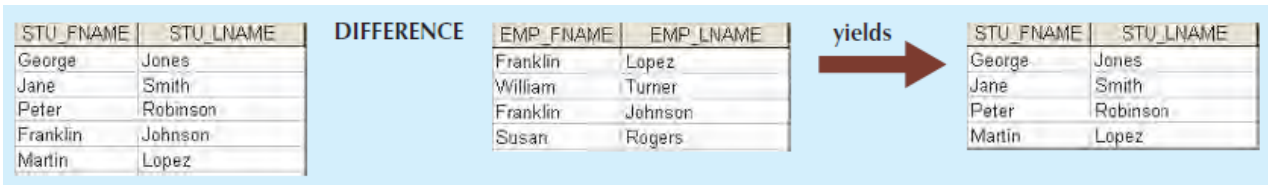
6. **PRODUCT**: YIELDS ALL POSSIBLE PAIRS OF ROWS FROM TWO TABLES- ALSO KNOWN AS CARTESIAN PRODUCT. THEREFORE, IF ONE TABLE HAS SIX ROWS AND THE OTHER TABLE HAS THREE, THE **PRODUCT** YIELDS A LIST COMPOSED OF 6 X 3= 18 ROWS.

P_CODE	P_DESCRIPT	PRICE
123456	Flashlight	5.26
123457	Lamp	25.15
123458	Box Fan	10.99
213345	9v battery	1.92
254467	100W bulb	1.47
311452	Powerdrill	34.99

PRODUCT

STORE	AISLE	SHELF
23	W	5
24	K	9
25	Z	6

yields →

P_CODE	P_DESCRIPT	PRICE	STORE	AISLE	SHELF
123456	Flashlight	5.26	23	W	5
123456	Flashlight	5.26	24	K	9
123456	Flashlight	5.26	25	Z	6
123457	Lamp	25.15	23	W	5
123457	Lamp	25.15	24	K	9
123457	Lamp	25.15	25	Z	6
123458	Box Fan	10.99	23	W	5
123458	Box Fan	10.99	24	K	9
123458	Box Fan	10.99	25	Z	6
213345	9v battery	1.92	23	W	5
213345	9v battery	1.92	24	K	9
213345	9v battery	1.92	25	Z	6
311452	Powerdrill	34.99	23	W	5
311452	Powerdrill	34.99	24	K	9
311452	Powerdrill	34.99	25	Z	6
254467	100W bulb	1.47	23	W	5
254467	100W bulb	1.47	24	K	9
254467	100W bulb	1.47	25	Z	6

7. **JOIN**: JOINS TWO TABLES TOGETHER USING A SHARED KEY USUALLY EITHER THE PRIMARY KEY OR FOREIGN KEY. JOIN ALLOWS THE USE OF INDEPENDENT TABLES LINKED BY COMMON ATTRIBUTES. JOIN IS A FUNDAMENTAL CONCEPT IN RELATIONAL DATABASE. A JOIN CAN EITHER BE INNER JOIN OR OUTER JOIN. AN INNER JOIN IS A JOIN THAT ONLY RETURNS MATCHED RECORDS FROM THE TABLES THAT ARE BEING JOINED E.G. *NATURAL JOIN, EQUIJOIN, THETA JOIN*. IN AN OUTER JOIN, THE MATCHED PAIRS WOULD BE RETAINED, AND ANY UNMATCHED VALUES IN THE OTHER TABLE WOULD BE LEFT NULL. WE LOOK AT TYPES OF JOIN BELOW:

- **NATURAL JOIN (INNER JOIN)**: A NATURAL JOIN LINKS TABLES BY SELECTING ONLY THE ROWS WITH COMMON VALUES IN THEIR COMMON ATTRIBUTE(S). A NATURAL JOIN IS THE RESULT OF A THREE-STAGE PROCESS:
  A. **PRODUCT** OF THE TABLES IS CREATED
  B. **SELECT** IS PERFORMED ON THE OUTPUT OF STEP A) TO YIELD ONLY THE ROWS WHOSE VALUES ARE EQUAL.
  C. A **PROJECT** IS PERFORMED ON THE RESULTS OF STEP B TO YIELD A SINGLE COPY OF EACH ATTRIBUTE, THEREBY ELIMINATING DUPLICATE COLUMNS. THE FINAL OUTCOME OF A NATURAL JOIN YIELDS A TABLE THAT DOES NOT INCLUDE UNMATCHED PAIRS AND PROVIDES ONLY COPIES OF THE MATCHES.

Table name: CUSTOMER

CUS_CODE	CUS_LNAME	CUS_ZIP	AGENT_CODE
1132445	Walker	32145	231
1217782	Adares	32145	125
1312243	Rakowski	34129	167
1321242	Rodriguez	37134	125
1542311	Smithson	37134	421
1657399	Vanloo	32145	231

Table name: AGENT

AGENT_CODE	AGENT_PHONE
125	6152439887
167	6153426778
231	6152431124
333	9041234445

THE TWO TABLES TO BE USED FOR JOIN. IN THE FOLLOWING EXAMPLE,

```
SELECT*
FROM Customer
NATURAL JOIN Agent
```

NOTE A FEW CRUCIAL FEATURES OF THE NATURAL JOIN OPERATION:

➢ IF NO MATCH IS MADE BETWEEN THE TABLE ROWS, THE NEW TABLE DOES NOT INCLUDE THE UNMATCHED ROW. IN THAT CASE, NEITHER AGENT_CODE 421 NOR THE CUSTOMER WHOSE LAST NAME IS SMITHSON IS INCLUDED. SMITHSON'S AGENT_CODE 421 DOES NOT MATCH ANY ENTRY IN THE AGENT TABLE.

➤ THE COLUMN ON WHICH THE JOIN WAS MADE—THAT IS, AGENT_CODE—OCCURS ONLY ONCE IN THE NEW TABLE.

➤ IF THE SAME AGENT_CODE WERE TO OCCUR SEVERAL TIMES IN THE AGENT TABLE, A CUSTOMER WOULD BE LISTED FOR EACH MATCH. FOR EXAMPLE, IF THE AGENT_CODE 167 WERE TO OCCUR THREE TIMES IN THE AGENT TABLE, THE CUSTOMER NAMED RAKOWSKI, WHO IS ASSOCIATED WITH AGENT_CODE 167, WOULD OCCUR THREE TIMES IN THE RESULTING TABLE. (A GOOD AGENT TABLE CANNOT, OF COURSE, YIELD SUCH A RESULT BECAUSE IT WOULD CONTAIN UNIQUE PRIMARY KEY VALUES.)

CUS_CODE	CUS_LNAME	CUS_ZIP	CUSTOMER.AGENT_CODE	AGENT.AGENT_CODE	AGENT_PHONE
1132445	Walker	32145	231	125	6152439887
1132445	Walker	32145	231	167	6153426778
1132445	Walker	32145	231	231	6152431124
1132445	Walker	32145	231	333	9041234445
1217782	Adares	32145	125	125	6152439887
1217782	Adares	32145	125	167	6153426778
1217782	Adares	32145	125	231	6152431124
1217782	Adares	32145	125	333	9041234445
1312243	Rakowski	34129	167	125	6152439887
1312243	Rakowski	34129	167	167	6153426778
1312243	Rakowski	34129	167	231	6152431124
1312243	Rakowski	34129	167	333	9041234445
1321242	Rodriguez	37134	125	125	6152439887
1321242	Rodriguez	37134	125	167	6153426778
1321242	Rodriguez	37134	125	231	6152431124
1321242	Rodriguez	37134	125	333	9041234445
1542311	Smithson	37134	421	125	6152439887
1542311	Smithson	37134	421	167	6153426778
1542311	Smithson	37134	421	231	6152431124
1542311	Smithson	37134	421	333	9041234445
1657399	Vanloo	32145	231	125	6152439887
1657399	Vanloo	32145	231	167	6153426778
1657399	Vanloo	32145	231	231	6152431124
1657399	Vanloo	32145	231	333	9041234445

STEP 1: CARTESIAN PRODUCT OF THE 2 TABLES

CUS_CODE	CUS_LNAME	CUS_ZIP	CUSTOMER.AGENT_CODE	AGENT.AGENT_CODE	AGENT_PHONE
1217782	Adares	32145	125	125	6152439887
1321242	Rodriguez	37134	125	125	6152439887
1312243	Rakowski	34129	167	167	6153426778
1132445	Walker	32145	231	231	6152431124
1657399	Vanloo	32145	231	231	6152431124

STEP 2: SELECT YIELD ONLY THE ROWS FOR WHICH THE AGENT_CODE VALUES ARE EQUAL. THE COMMON COLUMNS ARE REFERRED TO AS THE JOIN COLUMNS

CUS_CODE	CUS_LNAME	CUS_ZIP	AGENT_CODE	AGENT_PHONE
1217782	Adares	32145	125	6152439887
1321242	Rodriguez	37134	125	6152439887
1312243	Rakowski	34129	167	6153426778
1132445	Walker	32145	231	6152431124
1657399	Vanloo	32145	231	6152431124

STEP 3: PROJECT ELIMINATES DUPLICATE COLUMNS TO YIELD ONLY AGENT_CODE

▪ EQUIJOIN, LINKS TABLES ON THE BASIS OF AN EQUALITY CONDITION THAT COMPARES SPECIFIED COLUMNS OF EACH TABLE. THE OUTCOME OF THE EQUIJOIN DOES NOT ELIMINATE DUPLICATE COLUMNS, AND THE CONDITION OR CRITERION USED TO JOIN THE TABLES MUST BE EXPLICITLY DEFINED. THE EQUIJOIN TAKES ITS NAME FROM THE EQUALITY COMPARISON OPERATOR (=) USED IN THE CONDITION. E.G.

```
SELECT*
FROM Customer
JOIN Agent on (AGENT_CODE = CUSTOMER.AGENT_CODE)
```

▪ THETA JOIN: IF ANY OTHER COMPARISON OPERATOR SUCH AS (<, >, ...) IS USED, THE JOIN IS CALLED A THETA JOIN.

```
SELECT*
FROM Customer
JOIN Agent on (AGENT_CODE > CUSTOMER.AGENT_CODE)
```

- **Outer Join:** In an outer join, the matched pairs would be retained, and any unmatched values in the other table would be left null. It is an easy mistake to think that an outer join is the opposite of an inner join. However, it is more accurate to think of an outer join as an "inner join plus." The outer join still returns all of the matched records that the inner join returns, plus it returns the unmatched records from one of the tables. The **SQL OUTER JOIN** operator (+) is used only on one side of the join condition only. The subtypes of **OUTER JOIN** are:
  - ➢ Left outer join or left join
  - ➢ Right outer join or right join
  - ➢ Full outer join

Syntax
```
Select *
FROM table1, table2
WHERE conditions [+];
```
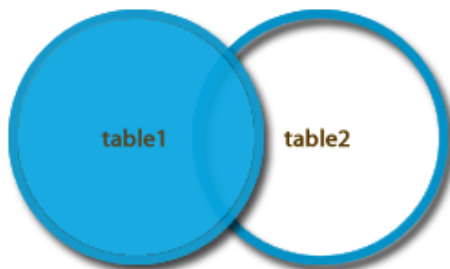
- The **LEFT JOIN** (specified with the keywords **LEFT JOIN** and **ON**) joins two tables and fetches all matching rows of two tables for which the SQL-expression is true, plus rows from the first table that do not match any row in the second table.

Left Join: Syntax
```
SELECT *
FROM table1
LEFT [OUTER] JOIN table2
ON table1.column_name=table2.column_name;
```


Pictorial representation of Left join

E.g.
```
SELECT *
FROM CUSTOMER
LEFT OUTER JOIN AGENT
ON CUSTOMER.AGENT_CODE = AGENT_CODE
```

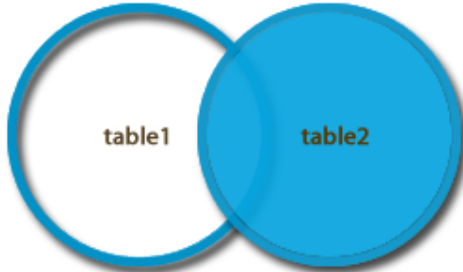| CUS_CODE	CUS_LNAME	CUS_ZIP	AGENT_CODE	AGENT_PHONE
1217782	Adares	32145	125	6152439887
1321242	Rodriguez	37134	125	6152439887
1312243	Rakowski	34129	167	6153426778
1132445	Walker	32145	231	6152431124
1667399	Vanloo	32145	231	6152431124
1542311	Smithson	37134	421	

Left outer join or Left Join

- The **RIGHT JOIN**, joins two tables and fetches rows based on a condition, which are matching in both the tables (before and after the **JOIN** clause

mentioned in the syntax below), and the unmatched rows will also be available from the table written after the JOIN clause (mentioned in the syntax below).

Syntax

```
SELECT *
FROM table1
RIGHT [OUTER] JOIN table2
ON table1.column_name=table2.column_name;
```


Pictorial representation of Right Join

E.g.

```
SELECT *
FROM CUSTOMER
RIGHT OUTER JOIN AGENT
ON CUSTOMER.AGENT_CODE = AGENT_CODE
```

| CUS_CODE	CUS_LNAME	CUS_ZIP	AGENT_CODE	AGENT_PHONE
1217782	Adares	32145	125	6152439887
1321242	Rodriguez	37134	125	6152439887
1312243	Rakowski	34129	167	6153426778
1132445	Walker	32145	231	6152431124
1657399	Vanloo	32145	231	6152431124
			333	9041234445

Right Join

- Full outer join: the FULL OUTER JOIN combines the results of both left and right outer joins and returns all (matched or unmatched) rows from the tables on both sides of the join clause.

Syntax

```
SELECT *
FROM table1
FULL OUTER JOIN table2
ON table1.column_name=table2.column_name;
```

- ON table1.column_name=table2.column_name;More specifically, if an outer join is produced for tables CUSTOMER and AGENT, two scenarios are possible

8. The DIVIDE operation uses one single-column table (e.g., column "A") as the divisor and one 2-column table (i.e., columns "A" and "B") as the dividend. The tables must have a common column (e.g., column "A"). The output of the DIVIDE operation is a single column with the values of column "A" from the dividend table rows where the value of the common column (i.e., column "A") in both tables matches.
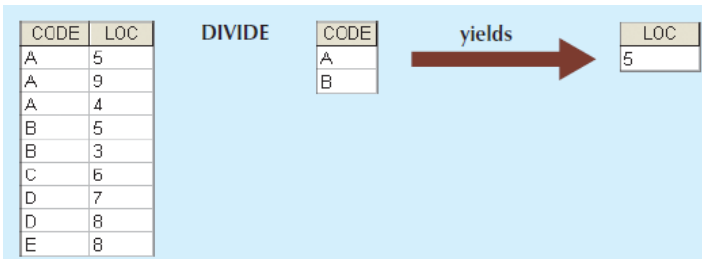
Sales Table

Export dest. code	Export dest. name	Date
12	The Kingdom of Minanmi	3/5
12	The Kingdom of Minanmi	3/10
23	Alpha Empire	3/5
25	The Kingdom of Ritol	3/21
30	The Kingdom of Sazanna	3/25

Export Destination Table

Export dest. code	Export dest. name
12	The Kingdom of Minanmi
23	Alpha Empire

Date
3/5

The result of the division will be the date on which there are exports to both The Kingdom of Minanmi and the Alpha Empire.

DIVIDE OPERATION

Relationships within the Relational Database

Relationships are classified as: one-to-one (1:1), one-to-many (1:M), and many-to-many (M:N or M:M). In developing a good database designs, we must focus on the following points:

- ➤ The 1:M relationship is the relational modeling ideal. Therefore, this relationship type should be the norm in any relational database design.
- ➤ The 1:1 relationship should be rare in any relational database design.
- ➤ M:N relationships cannot be implemented as such in the relational model. We will later consider how any M:N relationship can be changed into two 1:M relationships.

The 1:M Relationship
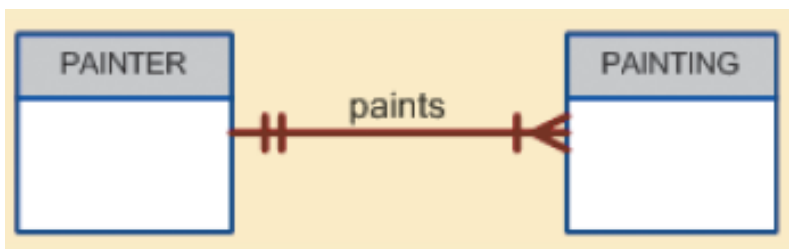
The 1:M relationship between PAINTER and PAINTING

Table name: PAINTER
Primary key: PAINTER_NUM
Foreign key: none

Database

PAINTER_NUM	PAINTER_LNAME	PAINTER_FNAME	PAINTER_INITIAL
123	Ross	Georgette	P
126	Itero	Julio	G

Table name: PAINTING
Primary key: PAINTING_NUM
Foreign key: PAINTER_NUM

PAINTING_NUM	PAINTING_TITLE	PAINTER_NUM
1338	Dawn Thunder	123
1339	Vanilla Roses To Nowhere	123
1340	Tired Flounders	126
1341	Hasty Exit	123
1342	Plastic Paradise	126

THE IMPLEMENTED 1:M RELATIONSHIP BETWEEN PAINTER AND PAINTING

THE ONE-TO-MANY (1:M) RELATIONSHIP IS EASILY IMPLEMENTED IN THE RELATIONAL MODEL BY PUTTING THE PRIMARY KEY OF THE 1 SIDE IN THE TABLE OF THE MANY SIDE AS A FOREIGN KEY.
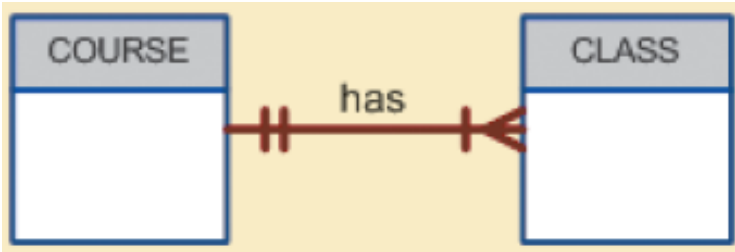
THE 1:M RELATIONSHIP BETWEEN COURSE AND CLASS

Table name: COURSE
Primary key: CRS_CODE
Foreign key: none

Database name: Ch03_TinyCollege

CRS_CODE	DEPT_CODE	CRS_DESCRIPTION	CRS_CREDIT
ACCT-211	ACCT	Accounting I	3
ACCT-212	ACCT	Accounting II	3
CIS-220	CIS	Intro. to Microcomputing	3
CIS-420	CIS	Database Design and Implementation	4
QM-261	CIS	Intro. to Statistics	3
QM-362	CIS	Statistical Applications	4

Table name: CLASS
Primary key: CLASS_CODE
Foreign key: CRS_CODE

CLASS_CODE	CRS_CODE	CLASS_SECTION	CLASS_TIME	CLASS_ROOM	PROF_NUM
10012	ACCT-211	1	MWF 8:00-8:50 a.m.	BUS311	105
10013	ACCT-211	2	MWF 9:00-9:50 a.m.	BUS200	105
10014	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
10015	ACCT-212	1	MWF 10:00-10:50 a.m.	BUS311	301
10016	ACCT-212	2	Th 6:00-8:40 p.m.	BUS252	301
10017	CIS-220	1	MWF 9:00-9:50 a.m.	KLR209	228
10018	CIS-220	2	MWF 9:00-9:50 a.m.	KLR211	114
10019	CIS-220	3	MWF 10:00-10:50 a.m.	KLR209	228
10020	CIS-420	1	W 6:00-8:40 p.m.	KLR209	162
10021	QM-261	1	MWF 8:00-8:50 a.m.	KLR200	114
10022	QM-261	2	TTh 1:00-2:15 p.m.	KLR200	114
10023	QM-362	1	MWF 11:00-11:50 a.m.	KLR200	162
10024	QM-362	2	TTh 2:30-3:45 p.m.	KLR200	162

THE IMPLEMENTED 1:M RELATIONSHIP BETWEEN COURSE AND CLASS

THE 1:1 RELATIONSHIP: AS THE 1:1 LABEL IMPLIES, IN THIS RELATIONSHIP, ONE ENTITY CAN BE RELATED TO ONLY ONE OTHER ENTITY, AND VICE VERSA. FOR EXAMPLE, ONE DEPARTMENT CHAIR—

A PROFESSOR—CAN CHAIR ONLY ONE DEPARTMENT, AND ONE DEPARTMENT CAN HAVE ONLY ONE DEPARTMENT CHAIR.

THE ENTITIES **PROFESSOR** AND **DEPARTMENT** THUS EXHIBIT A 1:1 RELATIONSHIP.
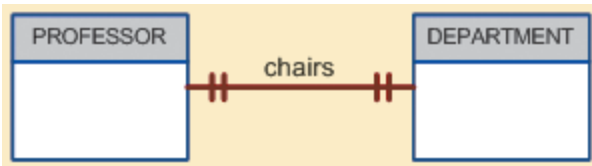
THE 1:M RELATIONSHIP BETWEEN **PROFRSSOR** AND **DEPARTMENT**

IF WE THE EXAMINE THE **PROFESSOR** AND **DEPARTMENT** TABLES, WE NOTE SOME IMPORTANT FEATURES:

- EACH PROFESSOR IS A COLLEGE EMPLOYEE; THUS, THE PROFESSOR IDENTIFICATION IS THROUGH THE **EMP_NUM**. (HOWEVER, NOTE THAT NOT ALL EMPLOYEES ARE PROFESSORS—THERE'S ANOTHER OPTIONAL RELATIONSHIP.)

- THE 1:1 **PROFESSOR** CHAIRS **DEPARTMENT** RELATIONSHIP IS IMPLEMENTED BY HAVING THE **EMP_NUM** AS FOREIGN KEY IN THE **DEPARTMENT** TABLE. NOTE THAT THE 1:1 RELATIONSHIP IS TREATED AS A SPECIAL CASE OF THE 1:M RELATIONSHIP IN WHICH THE "MANY" SIDE IS RESTRICTED TO A SINGLE OCCURRENCE. IN THIS CASE, **DEPARTMENT** CONTAINS THE **EMP_NUM** AS A FOREIGN KEY TO INDICATE THAT IT IS THE DEPARTMENT THAT HAS A CHAIR.

- ALSO, NOTE THAT THE **PROFESSOR** TABLE CONTAINS THE **DEPT_CODE** FOREIGN KEY TO IMPLEMENT THE 1:M **DEPARTMENT** EMPLOYS **PROFESSOR** RELATIONSHIP. THIS IS A GOOD EXAMPLE OF HOW TWO ENTITIES CAN PARTICIPATE IN TWO (OR EVEN MORE) RELATIONSHIPS SIMULTANEOUSLY. THE PRECEDING "**PROFESSOR** CHAIRS **DEPARTMENT**" EXAMPLE ILLUSTRATES A PROPER 1:1 RELATIONSHIP. IN FACT, THE USE OF A 1:1 RELATIONSHIP ENSURES THAT TWO ENTITY SETS ARE NOT PLACED IN THE SAME TABLE WHEN THEY SHOULD NOT BE. HOWEVER, THE EXISTENCE OF A 1:1 RELATIONSHIP SOMETIMES MEANS THAT THE ENTITY COMPONENTS WERE NOT DEFINED PROPERLY. IT COULD INDICATE THAT THE TWO ENTITIES ACTUALLY BELONG IN THE SAME TABLE! AS RARE AS 1:1 RELATIONSHIPS SHOULD BE, CERTAIN CONDITIONS ABSOLUTELY REQUIRE THEIR USE. ONE SUCH CONDITION IS THE CONCEPT CALLED GENERALIZATION HIERARCHY, WHICH IS A POWERFUL TOOL FOR IMPROVING DATABASE DESIGNS UNDER SPECIFIC CONDITIONS TO AVOID A PROLIFERATION OF NULLS. ONE OF THE CHARACTERISTICS OF GENERALIZATION HIERARCHIES IS THAT THEY ARE IMPLEMENTED AS 1:1 RELATIONSHIPS.

TABLE NAME: **PROFESSOR**
PRIMARY KEY: **EMP_NUM**
FOREIGN KEY: **DEPT_CODE**

EMP_NUM	DEPT_CODE	PROF_OFFICE	PROF_EXTENSION	PROF_HIGH_DEGREE
103	HST	DRE 156	6783	Ph.D.
104	ENG	DRE 102	5561	MA
105	ACCT	KLR 229D	8665	Ph.D.
106	MKT/MGT	KLR 126	3899	Ph.D.
110	BIOL	AAK 160	3412	Ph.D.
114	ACCT	KLR 211	4436	Ph.D.
155	MATH	AAK 201	4440	Ph.D.
160	ENG	DRE 102	2248	Ph.D.
162	CS	KLR 203E	2359	Ph.D.
191	MKT/MGT	KLR 409B	4016	DBA
195	PSYCH	AAK 297	3550	Ph.D.
209	CS	KLR 333	3421	Ph.D.
228	CS	KLR 300	3000	Ph.D.
297	MATH	AAK 194	1145	Ph.D.
299	ECON/FIN	KLR 284	2851	Ph.D.
301	ACCT	KLR 244	4683	Ph.D.
335	ENG	DRE 208	2000	Ph.D.
342	SOC	BBG 208	5514	Ph.D.
387	BIOL	AAK 230	8665	Ph.D.
401	HST	DRE 156	6783	MA
425	ECON/FIN	KLR 284	2851	MBA
435	ART	BBG 185	2278	Ph.D.

The 1:M DEPARTMENT employs PROFESSOR relationship is implemented through the placement of the DEPT_CODE foreign key in the PROFESSOR table.

The 1:1 PROFESSOR chairs DEPARTMENT relationship is implemented through the placement of the EMP_NUM foreign key in the DEPARTMENT table.

Table name: DEPARTMENT
Primary key: DEPT_CODE
Foreign key: EMP_NUM

DEPT_CODE	DEPT_NAME	SCHOOL_CODE	EMP_NUM	DEPT_ADDRESS	DEPT_EXTENSION
ACCT	Accounting	BUS	114	KLR 211, Box 52	3119
ART	Fine Arts	A&SCI	435	BBG 185, Box 128	2278
BIOL	Biology	A&SCI	387	AAK 230, Box 415	4117
CS	Computer Info. Systems	BUS	209	KLR 333, Box 56	3245
ECON/FIN	Economics/Finance	BUS	299	KLR 284, Box 63	3126
ENG	English	A&SCI	160	DRE 102, Box 223	1004
HST	History	A&SCI	103	DRE 156, Box 284	1867
MATH	Mathematics	A&SCI	297	AAK 194, Box 422	4234
MKT/MGT	Marketing/Management	BUS	106	KLR 126, Box 55	3342
PSYCH	Psychology	A&SCI	195	AAK 297, Box 438	4110
SOC	Sociology	A&SCI	342	BBG 208, Box 132	2008

The M:N Relationship: A many-to-many (M:N) relationship is not supported directly in the relational environment. However, M:N relationships can be implemented by creating a new entity in 1:M relationships with the original entities.

To explore the many-to-many (M:N) relationship, consider a rather typical college environment in which each STUDENT can take many CLASSes, and each CLASS can contain many STUDENTs. The ER model for this M:N relationship is below:
The ERM's M:N relationship between STUDENT and CLASS
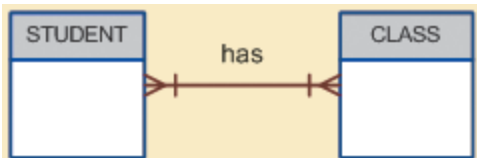
Note the features of the ERM above:
- Each CLASS can have many STUDENTs, and each STUDENT can take many CLASSes.
- There can be many rows in the CLASS table for any given row in the STUDENT table, and there can be many rows in the STUDENT table for any given row in the CLASS table.

To examine the M:N relationship more closely, imagine a small college with two students, each of whom takes three classes. The table below shows the enrollment data for the two students.

SAMPLE STUDENT ENROLLMENT DATA

STUDENT'S LAST NAME	SELECTED CLASSES
Bowser	Accounting 1, ACCT-211, code 10014 Intro to Microcomputing, CIS-220, code 10018 Intro to Statistics, QM-261, code 10021
Smithson	Accounting 1, ACCT-211, code 10014 Intro to Microcomputing, CIS-220, code 10018 Intro to Statistics, QM-261, code 10021

TABLE NAME: STUDENT
PRIMARY KEY: STU_NUM
FOREIGN KEY: NONE

STU_NUM	STU_LNAME	CLASS_CODE
321452	Bowser	10014
321452	Bowser	10018
321452	Bowser	10021
324257	Smithson	10014
324257	Smithson	10018
324257	Smithson	10021

Table name: CLASS
Primary key: CLASS_CODE
Foreign key: STU_NUM

CLASS_CODE	STU_NUM	CRS_CODE	CLASS_SECTION	CLASS_TIME	CLASS_ROOM	PROF_NUM
10014	321452	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
10014	324257	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
10018	321452	CIS-220	2	MWF 9:00-9:50 a.m.	KLR211	114
10018	324257	CIS-220	2	MWF 9:00-9:50 a.m.	KLR211	114
10021	321452	QM-261	1	MWF 8:00-8:50 a.m.	KLR200	114
10021	324257	QM-261	1	MWF 8:00-8:50 a.m.	KLR200	114

GIVEN THE DATA RELATIONSHIP AND THE SAMPLE DATA IN THE TABLE ABOVE, IT CAN BE WRONGLY ASSUMED THAT M:N RELATIONSHIP CAN BE IMPLEMENTED BY SIMPLY ADDING A FOREIGN KEY IN THE MANY SIDE OF THE RELATIONSHIP THAT POINTS TO THE PRIMARY KEY OF THE RELATED TABLE. THIS NOT CORRECT

- THE TABLES WILL CREATE MANY REDUNDANCIES. FOR EXAMPLE, NOTE THAT THE STU_NUM VALUES OCCUR MANY TIMES IN THE STUDENT TABLE. IN A REAL-WORLD SITUATION, ADDITIONAL STUDENT ATTRIBUTES SUCH AS ADDRESS, CLASSIFICATION, MAJOR, AND HOME PHONE WOULD ALSO BE CONTAINED IN THE STUDENT TABLE, AND EACH OF THOSE ATTRIBUTE VALUES WOULD BE REPEATED IN EACH OF THE RECORDS SHOWN HERE. SIMILARLY, THE CLASS TABLE CONTAINS MANY DUPLICATIONS: EACH STUDENT TAKING THE CLASS GENERATES A CLASS RECORD. THE PROBLEM WOULD BE EVEN WORSE IF THE CLASS TABLE INCLUDED SUCH ATTRIBUTES AS CREDIT HOURS AND COURSE DESCRIPTION.
- GIVEN THE STRUCTURE AND CONTENTS OF THE TWO TABLES, THE RELATIONAL OPERATIONS BECOME VERY COMPLEX AND ARE LIKELY TO LEAD TO SYSTEM EFFICIENCY ERRORS AND OUTPUT ERRORS.

THE PROBLEMS INHERENT IN THE MANY-TO-MANY (M:N) RELATIONSHIP CAN EASILY BE AVOIDED BY CREATING A
COMPOSITE ENTITY (ALSO REFERRED TO AS A BRIDGE ENTITY OR AN ASSOCIATIVE ENTITY). BECAUSE SUCH A TABLE IS USED TO LINK THE TABLES THAT WERE ORIGINALLY RELATED IN AN M:N RELATIONSHIP, THE COMPOSITE ENTITY STRUCTURE INCLUDES—AS FOREIGN KEYS—AT LEAST THE PRIMARY KEYS OF THE TABLES THAT ARE TO BE LINKED. THE DATABASE DESIGNER CAN THEN DEFINE THE COMPOSITE TABLE'S PRIMARY KEY EITHER BY: USING THE COMBINATION OF THOSE FOREIGN KEYS OR CREATE A NEW PRIMARY KEY. IN THE EXAMPLE ABOVE, WE CAN CREATE THE COMPOSITE ENROLL TABLE CLASS AND STUDENT. IN THIS EXAMPLE, THE ENROLL TABLE'S PRIMARY KEY IS THE COMBINATION OF ITS FOREIGN KEYS CLASS_CODE AND STU_NUM. BUT THE DESIGNER COULD HAVE DECIDED TO CREATE A SINGLE-ATTRIBUTE NEW PRIMARY KEY SUCH AS ENROLL_LINE, USING A DIFFERENT LINE VALUE TO IDENTIFY EACH ENROLL TABLE ROW UNIQUELY. (MICROSOFT ACCESS USERS MIGHT USE THE AUTONUMBER DATA TYPE TO GENERATE SUCH LINE VALUES AUTOMATICALLY).

Table name: STUDENT
Primary key: STU_NUM
Foreign key: none

STU_NUM	STU_LNAME
321452	Bowser
324257	Smithson

Table name: ENROLL
Primary key: CLASS_CODE + STU_NUM
Foreign key: CLASS_CODE, STU_NUM

CLASS_CODE	STU_NUM	ENROLL_GRADE
10014	321452	C
10014	324257	B
10018	321452	A
10018	324257	B
10021	321452	C
10021	324257	C

Table name: CLASS
Primary key: CLASS_CODE
Foreign key: CRS_CODE

CLASS_CODE	CRS_CODE	CLASS_SECTION	CLASS_TIME	CLASS_ROOM	PROF_NUM
10014	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
10018	CIS-220	2	MWF 9:00-9:50 a.m.	KLR211	114
10021	QM-261	1	MWF 8:00-8:50 a.m.	KLR200	114

Because the ENROLL table links two tables, STUDENT and CLASS, it is also called a linking table.
In other words, a linking table is the implementation of a composite entity.

The ENROLL table yields the required M:N to 1:M conversion. Observe that the composite entity represented by the ENROLL table must contain at least the primary keys of the CLASS and STUDENT tables (CLASS_CODE and STU_NUM, respectively) for which it serves as a connector. Also note that the STUDENT and CLASS tables now contain only one row per entity. The ENROLL table contains multiple occurrences of the foreign key values, but those controlled redundancies are incapable of producing anomalies as long as referential integrity is enforced. Additional attributes may be assigned as needed. In this case, ENROLL_GRADE is selected to satisfy a reporting requirement. Also note that the ENROLL table's primary key consists of the two attributes CLASS_CODE and STU_NUM because both the class code and the student number are needed to define a particular student's grade. Naturally, the conversion is reflected in the ERM, too. The revised relationship is shown below:

Changing the M:N relationship to two 1:M relationships
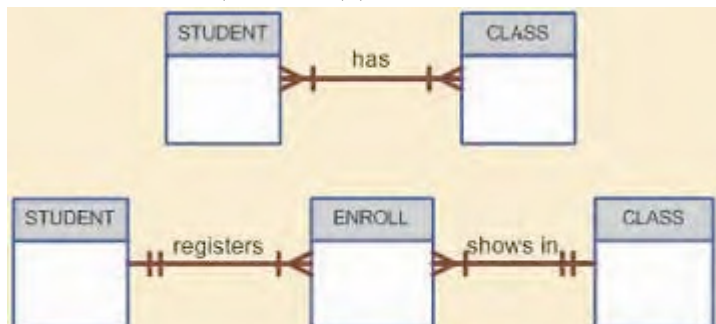
note that the composite entity named ENROLL represents the linking table between STUDENT and CLASS. We can increase the amount of available information even as

WE CONTROL THE DATABASE'S REDUNDANCIES. BELOW IS THE EXPANDED ERM, INCLUDING THE 1:M RELATIONSHIP BETWEEN **COURSE** AND **CLASS**. NOTE THAT THE MODEL IS ABLE TO HANDLE MULTIPLE SECTIONS OF A **CLASS** WHILE CONTROLLING REDUNDANCIES BY MAKING SURE THAT ALL OF THE **COURSE** DATA COMMON TO EACH **CLASS** ARE KEPT IN THE **COURSE** TABLE.

EXPANDED ENTITY RELATIONSHIP MODEL
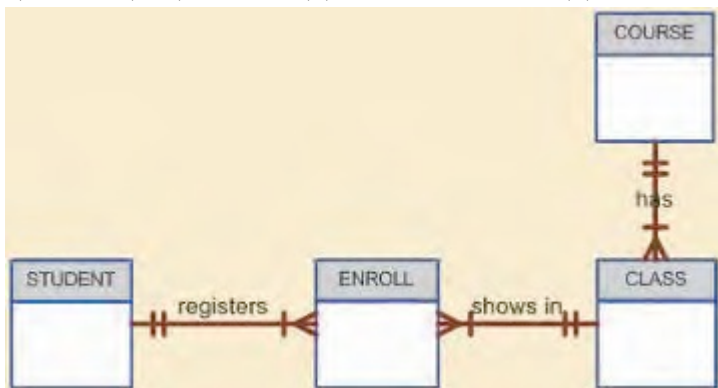
THE RELATIONSHIP DIAGRAM THAT CORRESPONDS TO THE ERM SHOWN ABOVE IS AS BELOW:
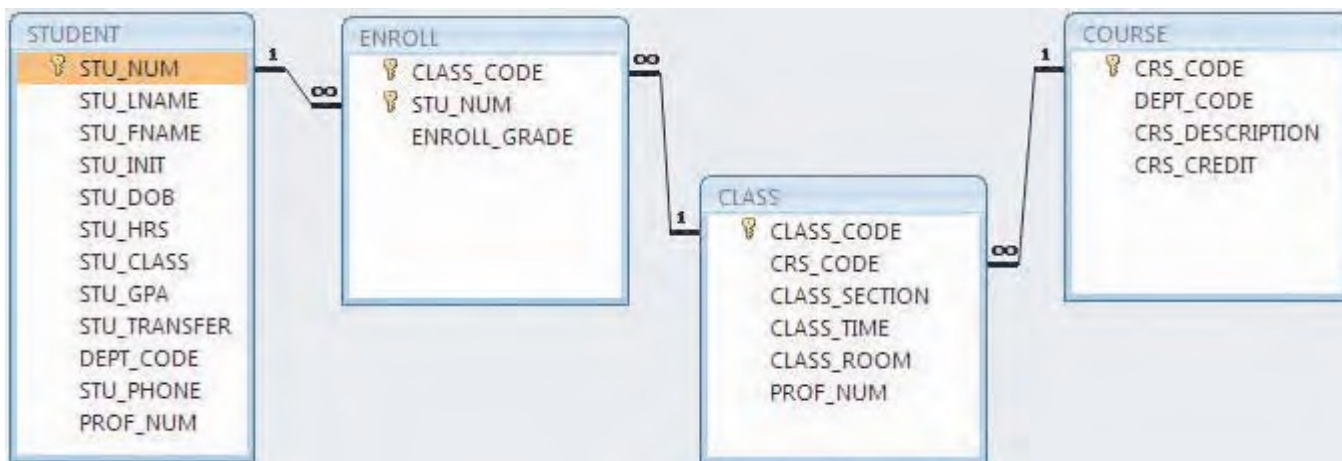
CODD'S RELATIONAL DATABASE RULES

IN 1985, DR. E. F. CODD PUBLISHED A LIST OF 12 RULES TO DEFINE A RELATIONAL DATABASE SYSTEM. THE REASON DR. CODD PUBLISHED THE LIST WAS HIS CONCERN THAT MANY VENDORS WERE MARKETING PRODUCTS AS "RELATIONAL" EVEN THOUGH THOSE PRODUCTS DID NOT MEET MINIMUM RELATIONAL STANDARDS. DR. CODD'S LIST, SERVES AS A FRAME OF REFERENCE FOR WHAT A TRULY RELATIONAL DATABASE SHOULD BE. NOTE THAT EVEN THE DOMINANT DATABASE VENDORS DO NOT FULLY SUPPORT ALL 12 RULES.

DR. CODD'S 12 RELATIONAL DATABASE RULES

RULE	RULE NAME	DESCRIPTION
1	Information	All information in a relational database must be logically represented as column values in rows within tables.
2	Guaranteed Access	Every value in a table is guaranteed to be accessible through a combination of table name, primary key value, and column name.
3	Systematic Treatment of Nulls	Nulls must be represented and treated in a systematic way, independent of data type.
4	Dynamic Online Catalog Based on the Relational Model	The metadata must be stored and managed as ordinary data, that is, in tables within the database. Such data must be available to authorized users using the standard database relational language.
5	Comprehensive Data Sublanguage	The relational database may support many languages. However, it must support one well-defined, declarative language with support for data definition, view definition, data manipulation (interactive and by program), integrity constraints, authorization, and transaction management (begin, commit, and rollback).
6	View Updating	Any view that is theoretically updatable must be updatable through the system.
7	High-Level Insert, Update, and Delete	The database must support set-level inserts, updates, and deletes.
8	Physical Data Independence	Application programs and ad hoc facilities are logically unaffected when physical access methods or storage structures are changed.
9	Logical Data Independence	Application programs and ad hoc facilities are logically unaffected when changes are made to the table structures that preserve the original table values (changing order of columns or inserting columns).
10	Integrity Independence	All relational integrity constraints must be definable in the relational language and stored in the system catalog, not at the application level.
11	Distribution Independence	The end users and application programs are unaware and unaffected by the data location (distributed vs. local databases).
12	Nonsubversion	If the system supports low-level access to the data, there must not be a way to bypass the integrity rules of the database.
	Rule Zero	All preceding rules are based on the notion that in order for a database to be considered relational, it must use its relational facilities exclusively to manage the database.

## THE ENTITY RELATIONSHIP MODEL (ERM)

PETER CHEN FIRST INTRODUCED THE ER DATA MODEL IN 1976; IT WAS THE GRAPHICAL REPRESENTATION OF ENTITIES AND THEIR RELATIONSHIPS IN A DATABASE STRUCTURE THAT QUICKLY BECAME POPULAR BECAUSE IT COMPLEMENTED THE RELATIONAL DATA MODEL CONCEPTS. THE RELATIONAL DATA MODEL AND ERM COMBINED TO PROVIDE THE FOUNDATION FOR TIGHTLY STRUCTURED DATABASE DESIGN. ER MODELS ARE NORMALLY REPRESENTED IN AN ENTITY RELATIONSHIP DIAGRAM (ERD), WHICH USES GRAPHICAL REPRESENTATIONS TO MODEL DATABASE COMPONENTS. THE ERD REPRESENTS THE CONCEPTUAL DATABASE AS VIEWED BY THE END USER. ERDS DEPICT THE DATABASE'S MAIN COMPONENTS: ENTITIES, ATTRIBUTES, AND RELATIONSHIPS. BECAUSE AN ENTITY REPRESENTS A REAL-WORLD OBJECT, THE WORDS ENTITY AND OBJECT ARE OFTEN USED INTERCHANGEABLY. THE NOTATIONS USED WITH ERDS ARE THE ORIGINAL CHEN NOTATION AND THE NEWER CROW'S FOOT AND UML NOTATIONS. SOME CONCEPTUAL DATABASE MODELING CONCEPTS CAN BE EXPRESSED ONLY USING THE CHEN NOTATION. BECAUSE OF ITS IMPLEMENTATION EMPHASIS, THE CROW'S FOOT NOTATION CAN REPRESENT ONLY WHAT COULD BE IMPLEMENTED. IN SUMMARY:

- THE CHEN NOTATION FAVORS CONCEPTUAL MODELING.
- THE CROW'S FOOT NOTATION FAVORS A MORE IMPLEMENTATION-ORIENTED APPROACH.
- THE UML NOTATION CAN BE USED FOR BOTH CONCEPTUAL AND IMPLEMENTATION MODELING.

THE ER MODEL IS BASED ON THE FOLLOWING COMPONENTS:

- **ENTITY:** An entity is anything about which data are to be collected and stored. An entity is represented in the ERD by a rectangle, also known as an entity box. The name of the entity, a noun, is written in the center of the rectangle. The entity name is generally written in capital letters and is written in the singular form: PAINTER rather than PAINTERS, and EMPLOYEE rather than EMPLOYEES. Usually, when applying the ERD to the relational model, an entity is mapped to a relational table. Each row in the relational table is known as an entity instance or entity occurrence in the ER model. Each entity is described by a set of attributes that describes particular characteristics of the entity. For example, the entity EMPLOYEE will have attributes such as a Social Security number, a last name, and a first name. A collection of like entities is known as an entity set. The word entity in the ERM corresponds to a table—not to a row—in the relational environment. The ERM refers to a table row as an entity instance or entity occurrence.

- **ATTRIBUTES:** Attributes are characteristics of entities. For example, the STUDENT entity includes, among many others, the attributes STU_LNAME, STU_FNAME, and STU_INITIAL. In the original Chen notation, attributes are represented by ovals and are connected to the entity rectangle with a line. Each oval contains the name of the attribute it represents. In the Crow's Foot notation, the attributes are written in the attribute box below the entity rectangle. Because the Chen representation is rather space-consuming, software vendors have adopted the Crow's Foot attribute display.

ATTRIBUTES OF THE **STUDENT** ENTITY: CHEN AND CROW'S FOOT
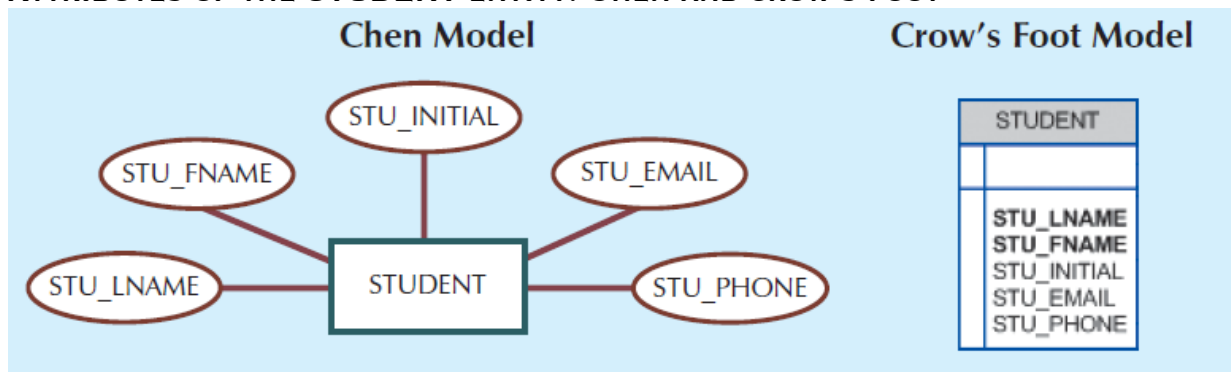
REQUIRED AND OPTIONAL ATTRIBUTES: A required attribute is an attribute that must have a value; in other words, it cannot be left empty. As shown above there are two boldfaced attributes in the Crow's Foot notation. This indicates that a data entry will be required. In this example, STU_LNAME and STU_FNAME require data entries because of the assumption that all students have a last name and a first name. But students might not have a middle name, and perhaps they do not (yet) have a phone number and an e-mail address. Therefore, those attributes are not presented in boldface in the entity box. An optional attribute is an attribute that does not require a value; therefore, it can be left empty.

ATTRIBUTE DOMAINS: Attributes have a domain. A domain is the set of possible values for a given attribute. For example, the domain for the grade point average (GPA) attribute is written (0,4) because the lowest possible GPA value is 0 and the highest possible value is 4. The domain for the gender attribute consists of only two possibilities: M or F (or some other equivalent code). The domain for a company's date of hire attribute consists of all dates that fit in a range (for example, company startup date to current date). Attributes may share a domain. For instance, a student address and a professor address share the same domain of all possible addresses. In fact, the data dictionary may let a newly declared attribute inherit the characteristics of an existing attribute if the same attribute name is used. For example, the PROFESSOR and STUDENT entities may each have an attribute named ADDRESS and could therefore share a domain.

IDENTIFIERS (PRIMARY KEYS): The ERM uses identifiers, that is, one or more attributes that uniquely identify each entity instance. In the relational model,

SUCH IDENTIFIERS ARE MAPPED TO PRIMARY KEYS (PKS) IN TABLES. IDENTIFIERS ARE UNDERLINED IN THE ERD. KEY ATTRIBUTES
ARE ALSO UNDERLINED IN A FREQUENTLY USED TABLE STRUCTURE SHORTHAND NOTATION USING THE FORMAT:
TABLE NAME (<u>KEY_ATTRIBUTE 1</u>, ATTRIBUTE 2, ATTRIBUTE 3, . . . ATTRIBUTE K)

FOR EXAMPLE, A **CAR** ENTITY MAY BE REPRESENTED BY:
   CAR (CAR_VIN, MOD_CODE, CAR_YEAR, CAR_COLOR)
(EACH CAR IS IDENTIFIED BY A UNIQUE VEHICLE IDENTIFICATION NUMBER, OR **CAR_VIN**.)
COMPOSITE IDENTIFIERS: IDEALLY, AN ENTITY IDENTIFIER IS COMPOSED OF ONLY A SINGLE ATTRIBUTE. HOWEVER, IT IS POSSIBLE TO USE A COMPOSITE IDENTIFIER, THAT IS, A PRIMARY KEY COMPOSED OF MORE THAN ONE ATTRIBUTE. E.G. **CLASS** ENTITY OF **CRS_CODE** AND **CLASS_SECTION** INSTEAD OF USING **CLASS_CODE**. EITHER APPROACH UNIQUELY IDENTIFIES EACH ENTITY INSTANCE.

COMPOSITE AND SIMPLE ATTRIBUTES: ATTRIBUTES ARE CLASSIFIED AS SIMPLE OR COMPOSITE. A COMPOSITE ATTRIBUTE, NOT TO BE CONFUSED WITH A COMPOSITE KEY, IS AN ATTRIBUTE THAT CAN BE FURTHER SUBDIVIDED TO YIELD ADDITIONAL ATTRIBUTES. FOR EXAMPLE, THE ATTRIBUTE **ADDRESS** CAN BE SUBDIVIDED INTO STREET, CITY, STATE, AND ZIP CODE. SIMILARLY, THE ATTRIBUTE **PHONE_NUMBER** CAN BE SUBDIVIDED INTO AREA CODE AND EXCHANGE NUMBER. A SIMPLE ATTRIBUTE IS AN ATTRIBUTE THAT CANNOT BE SUBDIVIDED. FOR EXAMPLE, AGE, SEX AND MARITAL STATUS WOULD BE CLASSIFIED AS SIMPLE ATTRIBUTES. TO FACILITATE DETAILED QUERIES, IT IS WISE TO CHANGE COMPOSITE ATTRIBUTES INTO A SERIES OF SIMPLE ATTRIBUTES.

SINGLE-VALUED ATTRIBUTES: A SINGLE-VALUED ATTRIBUTE IS AN ATTRIBUTE THAT CAN HAVE ONLY A SINGLE VALUE. FOR EXAMPLE, A PERSON CAN HAVE ONLY ONE SOCIAL SECURITY NUMBER, AND A MANUFACTURED PART CAN HAVE ONLY ONE SERIAL NUMBER. KEEP IN MIND THAT A SINGLE-VALUED ATTRIBUTE IS NOT NECESSARILY A SIMPLE ATTRIBUTE. FOR INSTANCE, A PART'S SERIAL NUMBER, SUCH AS **SE-08-02-189935**, IS SINGLE-VALUED, BUT IT IS A COMPOSITE ATTRIBUTE BECAUSE IT CAN BE SUBDIVIDED INTO THE REGION IN WHICH THE PART WAS PRODUCED (**SE**), THE PLANT WITHIN THAT REGION (**08**), THE SHIFT WITHIN THE PLANT (**02**), AND THE PART NUMBER (**189935**).

MULTIVALUED ATTRIBUTES: MULTIVALUED ATTRIBUTES ARE ATTRIBUTES THAT CAN HAVE MANY VALUES. FOR INSTANCE, A PERSON MAY HAVE SEVERAL COLLEGE DEGREES, AND A HOUSEHOLD MAY HAVE SEVERAL DIFFERENT PHONES, EACH WITH ITS OWN NUMBER. SIMILARLY, A CAR'S COLOR MAY BE SUBDIVIDED INTO MANY COLORS (THAT IS, COLORS FOR THE ROOF, BODY, AND TRIM). IN THE CHEN ERM, THE MULTIVALUED ATTRIBUTES ARE SHOWN BY A DOUBLE LINE CONNECTING THE ATTRIBUTE TO THE ENTITY. THE CROW'S FOOT NOTATION DOES NOT IDENTIFY MULTIVALUED ATTRIBUTES.
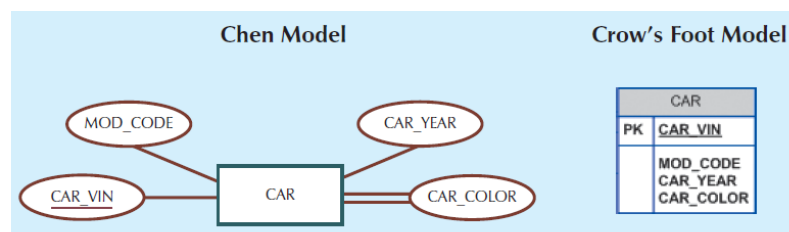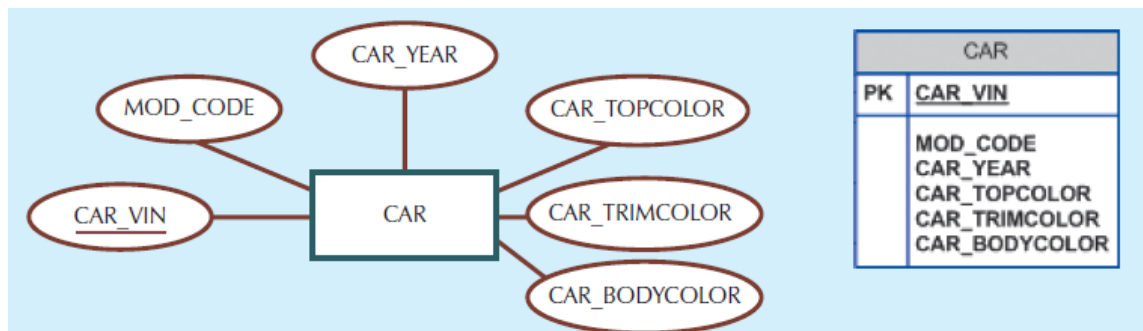

A MULTIVALUED ATTRIBUTE IN AN ENTITY

THE ERD ABOVE CONTAINS ALL OF THE COMPONENTS INTRODUCED THUS FAR. NOTE THAT **CAR_VIN** IS THE PRIMARY KEY, AND **CAR_COLOR** IS A MULTIVALUED ATTRIBUTE OF THE **CAR** ENTITY.

IMPLEMENTING MULTIVALUED ATTRIBUTES
ALTHOUGH THE CONCEPTUAL MODEL CAN HANDLE M:N RELATIONSHIPS AND MULTIVALUED ATTRIBUTES, IT POOR PRACTICE TO IMPLEMENT THEM IN THE RDBMS. IN THE RELATIONAL TABLE, EACH COLUMN/ROW INTERSECTION REPRESENTS A SINGLE DATA VALUE. THE DESIGNER MUST DECIDE ON ONE OF TWO POSSIBLE COURSES OF ACTION TO HANDLE MULTIVALUED ATTRIBUTES:

I. SPLIT THE MULTIVALUED ATTRIBUTE TO CREATE SEVERAL NEW ATTRIBUTES. FOR EXAMPLE, THE CAR ENTITY'S ATTRIBUTE CAR_COLOR CAN BE SPLIT TO CREATE THE NEW ATTRIBUTES CAR_TOPCOLOR, CAR_BODYCOLOR, AND CAR_TRIMCOLOR, WHICH ARE THEN ASSIGNED TO THE CAR ENTITY. ALTHOUGH THIS SOLUTION SEEMS TO WORK, ITS ADOPTION CAN LEAD TO MAJOR STRUCTURAL PROBLEMS IN THE TABLE. FOR EXAMPLE, IF ADDITIONAL COLOR COMPONENTS—SUCH AS A LOGO COLOR—ARE ADDED FOR SOME CARS, THE TABLE STRUCTURE MUST BE MODIFIED TO ACCOMMODATE THE NEW COLOR SECTION. IN THAT CASE, CARS THAT DO NOT HAVE SUCH COLOR SECTIONS GENERATE NULLS FOR THE NONEXISTING COMPONENTS, OR THEIR COLOR ENTRIES FOR THOSE SECTIONS ARE ENTERED AS N/A TO INDICATE "NOT APPLICABLE." ALSO CONSIDER THE EMPLOYEE ENTITY CONTAINING EMPLOYEE DEGREES AND CERTIFICATIONS. IF SOME EMPLOYEES HAVE 10 DEGREES AND CERTIFICATIONS WHILE MOST HAVE FEWER OR NONE, THE NUMBER OF DEGREE/CERTIFICATION ATTRIBUTES WOULD NUMBER 10, AND MOST OF THOSE ATTRIBUTE VALUES WOULD BE NULL FOR MOST OF THE EMPLOYEES.) IN SHORT, WHILE SOLUTION 1 IS PRACTICABLE, IT IS NOT AN ACCEPTABLE SOLUTION.


SPLITTING THE MULTIVALUED ATTRIBUTE INTO NEW ATTRIBUTES

II. CREATE A NEW ENTITY COMPOSED OF THE ORIGINAL MULTIVALUED ATTRIBUTE'S COMPONENTS. THIS NEW ENTITY ALLOWS THE DESIGNER TO DEFINE COLOR FOR DIFFERENT SECTIONS OF THE CAR. (SEE TABLE BELOW).

COMPONENTS OF THE MULTIVALUED ATTRIBUTE

SECTION	COLOR
Top	White
Body	Blue
Trim	Gold
Interior	Blue

ANOTHER BENEFIT WE CAN DERIVE FROM THIS APPROACH IS THAT WE ARE NOW ABLE TO ASSIGN AS MANY COLORS AS NECESSARY WITHOUT HAVING TO CHANGE THE TABLE STRUCTURE.
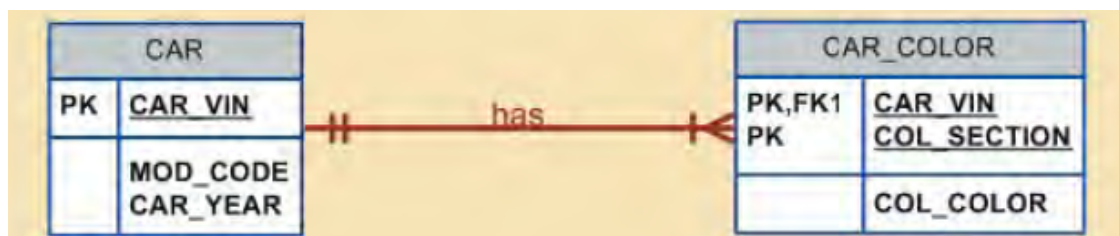
A NEW ENTITY SET COMPOSED OF A MULTIVALUED ATTRIBUTE'S COMPONENTS

NOTE THAT THE ERM SHOWN IN FIGURE ABOVE REFLECTS THE COMPONENTS LISTED IN PREVIOUS TABLE. THIS IS THE PREFERRED WAY TO DEAL WITH MULTIVALUED ATTRIBUTES. CREATING A NEW ENTITY IN A 1:M RELATIONSHIP WITH THE ORIGINAL ENTITY YIELDS SEVERAL BENEFITS: IT'S A MORE FLEXIBLE, EXPANDABLE SOLUTION, AND IT IS COMPATIBLE WITH THE RELATIONAL MODEL!

DERIVED ATTRIBUTES: A DERIVED ATTRIBUTE IS AN ATTRIBUTE WHOSE VALUE IS CALCULATED (DERIVED) FROM OTHER ATTRIBUTES. THE DERIVED ATTRIBUTE NEED NOT BE PHYSICALLY STORED WITHIN THE DATABASE; INSTEAD, IT CAN BE DERIVED BY USING AN ALGORITHM. FOR EXAMPLE, AN EMPLOYEE'S AGE, EMP_AGE, MAY BE FOUND BY COMPUTING THE INTEGER VALUE OF THE DIFFERENCE BETWEEN THE CURRENT DATE AND THE EMP_DOB. IN MICROSOFT ACCESS, WE USE:

```
INT((DATE() - EMP_DOB)/365)
```

IN MICROSOFT **SQL** SERVER, WE USE

```
SELECT DATEDIFF("YEAR", EMP_DOB, GETDATE()),
```

WHERE **DATEDIFF** IS A FUNCTION THAT COMPUTES THE DIFFERENCE BETWEEN DATES. THE FIRST PARAMETER INDICATES THE MEASUREMENT, IN THIS CASE, YEARS. IN ORACLE, WE USE **SYSDATE** INSTEAD OF **DATE()**.

A DERIVED ATTRIBUTE IS INDICATED IN THE CHEN NOTATION BY A DASHED LINE CONNECTING THE ATTRIBUTE AND THE ENTITY. THE CROW'S FOOT NOTATION DOES NOT HAVE A METHOD FOR DISTINGUISHING THE DERIVED ATTRIBUTE FROM OTHER ATTRIBUTES.
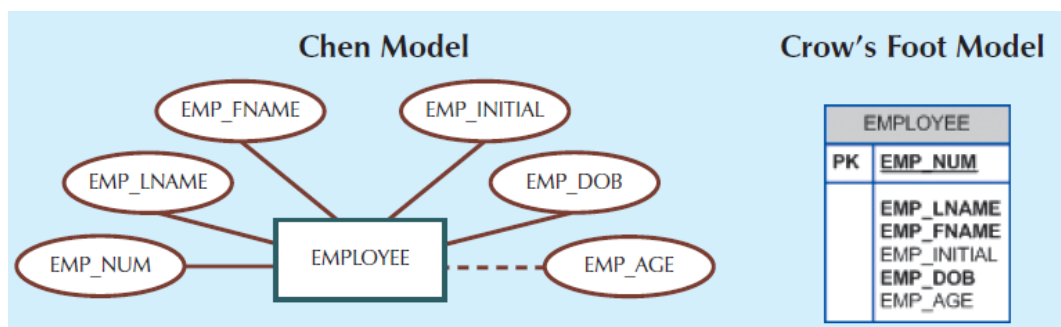
DEPICTION OF A DERIVED ATTRIBUTE

DERIVED ATTRIBUTES ARE SOMETIMES REFERRED TO AS COMPUTED ATTRIBUTES. A DERIVED ATTRIBUTE COMPUTATION CAN BE AS SIMPLE AS ADDING TWO ATTRIBUTE VALUES LOCATED ON THE SAME ROW, OR IT CAN BE THE RESULT OF AGGREGATING THE SUM OF VALUES LOCATED ON MANY TABLE ROWS (FROM THE SAME TABLE OR FROM A DIFFERENT TABLE). THE DECISION TO STORE DERIVED ATTRIBUTES IN DATABASE TABLES DEPENDS ON THE PROCESSING REQUIREMENTS AND THE CONSTRAINTS PLACED ON A PARTICULAR APPLICATION. THE DESIGNER SHOULD BE ABLE TO BALANCE THE DESIGN IN ACCORDANCE WITH SUCH CONSTRAINTS.

TABLE BELOW SHOWS THE ADVANTAGES AND DISADVANTAGES OF STORING (OR NOT STORING) DERIVED ATTRIBUTES IN THE DATABASE.

ADVANTAGES AND DISADVANTAGES OF STORING (OR NOT STORING) DERIVED ATTRIBUTES IN THE DATABASE.

	DERIVED ATTRIBUTE	
	STORED	NOT STORED
Advantage	Saves CPU processing cycles Saves data access time Data value is readily available Can be used to keep track of historical data	Saves storage space Computation always yields current value
Disadvantage	Requires constant maintenance to ensure derived value is current, especially if any values used in the calculation change	Uses CPU processing cycles Increases data access time Adds coding complexity to queries

- **RELATIONSHIPS.** RELATIONSHIPS DESCRIBE ASSOCIATIONS AMONG DATA. MOST RELATIONSHIPS DESCRIBE ASSOCIATIONS BETWEEN TWO ENTITIES. THE THREE TYPES OF RELATIONSHIPS AMONG DATA INCLUDE:
  - ONE-TO-MANY (1:M)
  - MANY-TO-MANY (M:N)
  - AND ONE-TO-ONE (1:1).

THE ER MODEL USES THE TERM CONNECTIVITY TO LABEL THE RELATIONSHIP TYPES. THE NAME OF THE RELATIONSHIP IS USUALLY AN ACTIVE OR PASSIVE VERB. FOR EXAMPLE, A **PAINTER** PAINTS MANY **PAINTINGS**; AN **EMPLOYEE** LEARNS MANY **SKILLS**; AN **EMPLOYEE** MANAGES A **STORE**. ILLUSTRATED BELOW ARE THE DIFFERENT TYPES OF

relationships using two ER notations: the original Chen notation and the more current Crow's Foot notation.
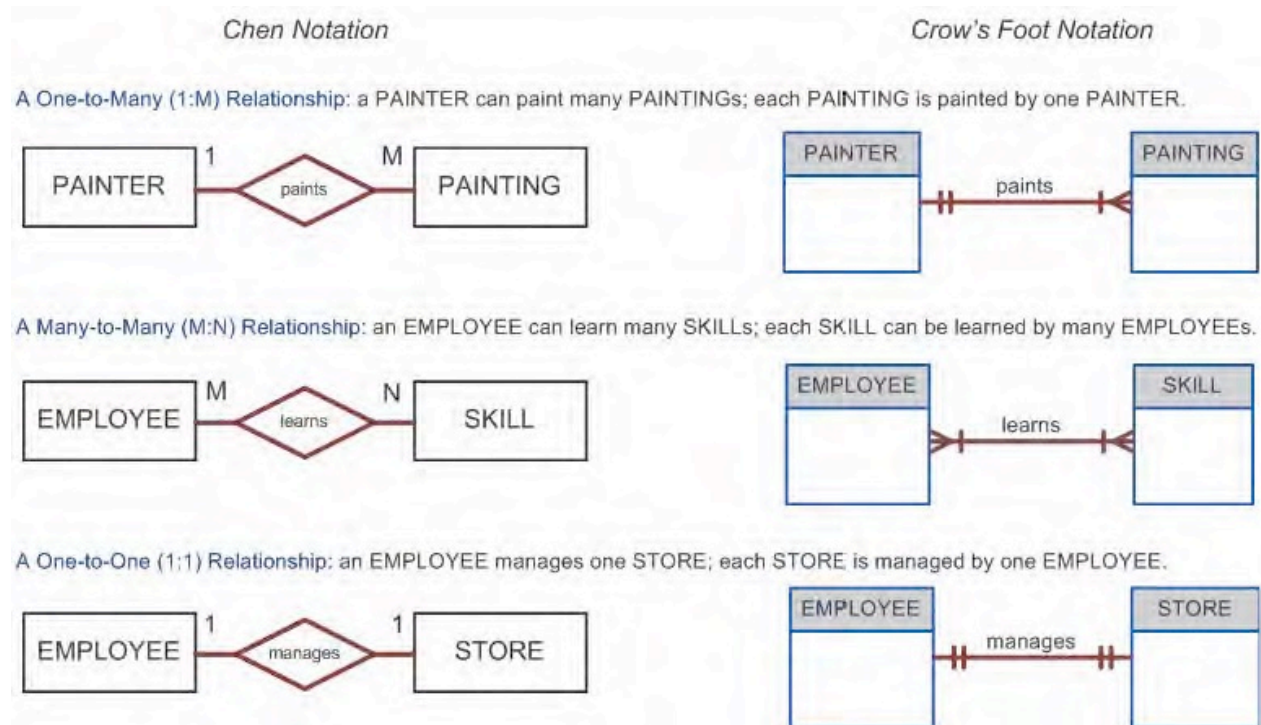
The left side of the ER diagram shows the Chen notation, based on Peter Chen's landmark paper. In this notation, the connectivities are written next to each entity box. Relationships are represented by a diamond connected to the related entities through a relationship line. The relationship name is written inside the diamond. The right side illustrates the Crow's Foot notation. The name "Crow's Foot" is derived from the three-pronged symbol used to represent the "many" side of the relationship. In the basic Crow's Foot ERD represented above, the connectivities are represented by symbols. For example, the "1" is represented by a short line segment, and the "M" is represented by the three-pronged "crow's foot." The relationship name is written above the relationship line. In Figure above, entities and relationships are shown in a horizontal format, but they may also be oriented vertically. The entity location and the order in which the entities are presented are immaterial; just remember to read a 1:M relationship from the "1" side to the "M" side.

Connectivity and Cardinality

As stated above, the term connectivity is used to describe the relationship classification. Cardinality expresses the minimum and maximum number of entity occurrences associated with one occurrence of the related entity. In the ERD, cardinality is indicated by placing the appropriate numbers beside the entities, using the format (x,y). The first value represents the minimum number of associated entities, while the second value represents the maximum number of associated entities. Many database designers who use Crow's Foot modeling notation do not depict the specific cardinalities on the ER diagram itself because the specific limits described by the cardinalities cannot be implemented directly through the database design. Correspondingly, some Crow's Foot ER modeling tools do not print the numeric cardinality range in the diagram; instead, you can add it as text if you want to have it shown.
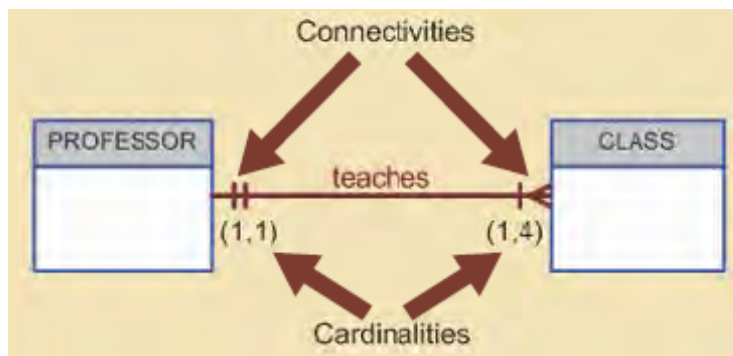
CONNECTIVITY AND CARDINALITY

KNOWING THE MINIMUM AND MAXIMUM NUMBER OF ENTITY OCCURRENCES IS VERY USEFUL AT THE APPLICATION SOFTWARE LEVEL. A COLLEGE MIGHT WANT TO ENSURE THAT A CLASS IS NOT TAUGHT UNLESS IT HAS AT LEAST 10 STUDENTS ENROLLED. SIMILARLY, IF THE CLASSROOM CAN HOLD ONLY 30 STUDENTS, THE APPLICATION SOFTWARE SHOULD USE THAT CARDINALITY TO LIMIT ENROLLMENT IN THE CLASS. HOWEVER, KEEP IN MIND THAT THE DBMS CANNOT HANDLE THE IMPLEMENTATION OF THE CARDINALITIES AT THE TABLE LEVEL—THAT CAPABILITY IS PROVIDED BY THE APPLICATION SOFTWARE OR BY TRIGGERS.

EXISTENCE DEPENDENCE: AN ENTITY IS SAID TO BE EXISTENCE-DEPENDENT IF IT CAN EXIST IN THE DATABASE ONLY WHEN IT IS ASSOCIATED WITH ANOTHER RELATED ENTITY OCCURRENCE. IN IMPLEMENTATION TERMS, AN ENTITY IS EXISTENCE-DEPENDENT IF IT HAS A MANDATORY FOREIGN KEY—THAT IS, A FOREIGN KEY ATTRIBUTE THAT CANNOT BE NULL. FOR EXAMPLE, IF AN EMPLOYEE WANTS TO CLAIM ONE OR MORE DEPENDENTS FOR TAX-WITHHOLDING PURPOSES, THE RELATIONSHIP "EMPLOYEE CLAIMS DEPENDENT" WOULD BE APPROPRIATE. IN THAT CASE, THE DEPENDENT ENTITY IS CLEARLY EXISTENCE-DEPENDENT ON THE EMPLOYEE ENTITY BECAUSE IT IS IMPOSSIBLE FOR THE DEPENDENT TO EXIST APART FROM THE EMPLOYEE IN THE DATABASE. IF AN ENTITY CAN EXIST APART FROM ALL OF ITS RELATED ENTITIES (IT IS EXISTENCE-INDEPENDENT), THEN THAT ENTITY IS REFERRED TO AS A STRONG ENTITY OR REGULAR ENTITY.

RELATIONSHIP STRENGTH: THE CONCEPT OF RELATIONSHIP STRENGTH IS BASED ON HOW THE PRIMARY KEY OF A RELATED ENTITY IS DEFINED. TO IMPLEMENT A RELATIONSHIP, THE PRIMARY KEY OF ONE ENTITY APPEARS AS A FOREIGN KEY IN THE RELATED ENTITY. FOR EXAMPLE, THE 1:M RELATIONSHIP BETWEEN VENDOR AND PRODUCT IS IMPLEMENTED BY USING THE VEND_CODE PRIMARY KEY IN VENDOR AS A FOREIGN KEY IN PRODUCT. THERE ARE TIMES WHEN THE FOREIGN KEY ALSO IS A PRIMARY KEY COMPONENT IN THE RELATED ENTITY. RELATIONSHIP STRENGTH DECISIONS AFFECT PRIMARY KEY ARRANGEMENT IN DATABASE DESIGN.

WEAK (NON-IDENTIFYING) RELATIONSHIPS: A WEAK RELATIONSHIP, ALSO KNOWN AS A NON-IDENTIFYING RELATIONSHIP, EXISTS IF THE PK OF THE RELATED ENTITY DOES NOT CONTAIN A PK COMPONENT OF THE PARENT ENTITY. BY DEFAULT, RELATIONSHIPS ARE ESTABLISHED BY HAVING THE PK OF THE PARENT ENTITY APPEAR AS AN FK ON THE RELATED ENTITY. FOR EXAMPLE, SUPPOSE THAT THE COURSE AND CLASS ENTITIES ARE DEFINED AS:

```
COURSE(CRS_CODE, DEPT_CODE, CRS_DESCRIPTION, CRS_CREDIT)
CLASS(CLASS_CODE,  CRS_CODE,  CLASS_SECTION,  CLASS_TIME,
ROOM_CODE, PROF_NUM)
```

IN THIS CASE, A WEAK RELATIONSHIP EXISTS BETWEEN COURSE AND CLASS BECAUSE THE CLASS_CODE IS THE CLASS ENTITY'S PK, WHILE THE CRS_CODE IN CLASS IS ONLY AN FK. IN THIS EXAMPLE, THE CLASS PK DID NOT INHERIT THE PK COMPONENT FROM THE COURSE ENTITY.

TABLE NAME: COURSE

CRS_CODE	DEPT_CODE	CRS_DESCRIPTION	CRS_CREDIT
ACCT-211	ACCT	Accounting I	3
ACCT-212	ACCT	Accounting II	3
CIS-220	CIS	Intro. to Microcomputing	3
CIS-420	CIS	Database Design and Implementation	4
MATH-243	MATH	Mathematics for Managers	3
QM-261	CIS	Intro. to Statistics	3
QM-362	CIS	Statistical Applications	4

Table name: CLASS

CLASS_CODE	CRS_CODE	CLASS_SECTION	CLASS_TIME	ROOM_CODE	PROF_NUM
10012	ACCT-211	1	MWF 8:00-8:50 a.m.	BUS311	105
10013	ACCT-211	2	MWF 9:00-9:50 a.m.	BUS200	105
10014	ACCT-211	3	TTh 2:30-3:45 p.m.	BUS252	342
10015	ACCT-212	1	MWF 10:00-10:50 a.m.	BUS311	301
10016	ACCT-212	2	Th 6:00-8:40 p.m.	BUS252	301
10017	CIS-220	1	MWF 9:00-9:50 a.m.	KLR209	228
10018	CIS-220	2	MWF 9:00-9:50 a.m.	KLR211	114
10019	CIS-220	3	MWF 10:00-10:50 a.m.	KLR209	228
10020	CIS-420	1	W 6:00-8:40 p.m.	KLR209	162
10021	QM-261	1	MWF 8:00-8:50 a.m.	KLR200	114
10022	QM-261	2	TTh 1:00-2:15 p.m.	KLR200	114
10023	QM-362	1	MWF 11:00-11:50 a.m.	KLR200	162
10024	QM-362	2	TTh 2:30-3:45 p.m.	KLR200	162
10025	MATH-243	1	Th 6:00-8:40 p.m.	DRE155	325
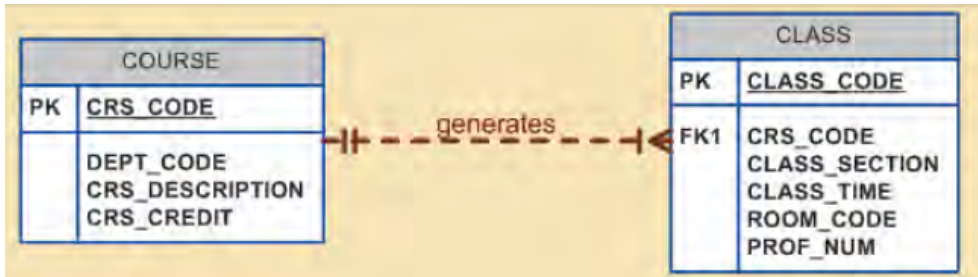
CROW'S FOOT NOTATION DEPICTS A STRONG RELATIONSHIP

STRONG (IDENTIFYING) RELATIONSHIPS: A STRONG RELATIONSHIP, ALSO KNOWN AS AN IDENTIFYING RELATIONSHIP, EXISTS WHEN THE PK OF THE RELATED ENTITY CONTAINS A PK COMPONENT OF THE PARENT ENTITY. FOR EXAMPLE, THE DEFINITIONS OF THE COURSE AND CLASS ENTITIES COURSE(CRS_CODE, DEPT_CODE, CRS_DESCRIPTION, CRS_CREDIT)
CLASS(CRS_CODE, CLASS_SECTION , CLASS_TIME, ROOM_CODE, PROF_NUM)
INDICATE THAT A STRONG RELATIONSHIP EXISTS BETWEEN COURSE AND CLASS, BECAUSE THE CLASS ENTITY'S COMPOSITE PK IS COMPOSED OF CRS_CODE + CLASS_SECTION. (NOTE THAT THE CRS_CODE IN CLASS IS ALSO THE FK TO THE COURSE ENTITY.) THE CROW'S FOOT NOTATION DEPICTS THE STRONG (IDENTIFYING) RELATIONSHIP WITH A SOLID LINE BETWEEN THE ENTITIES. WHETHER THE RELATIONSHIP BETWEEN COURSE AND CLASS IS STRONG OR WEAK DEPENDS ON HOW THE CLASS ENTITY'S PRIMARY KEY IS DEFINED. KEEP IN MIND THAT THE ORDER IN WHICH THE TABLES ARE CREATED AND LOADED IS VERY IMPORTANT. FOR EXAMPLE, IN THE "COURSE GENERATES CLASS" RELATIONSHIP, THE COURSE TABLE MUST BE CREATED BEFORE THE CLASS TABLE. AFTER ALL, IT WOULD NOT BE ACCEPTABLE TO HAVE THE CLASS TABLE'S FOREIGN KEY REFERENCE A COURSE TABLE THAT DOES NOT YET EXIST.
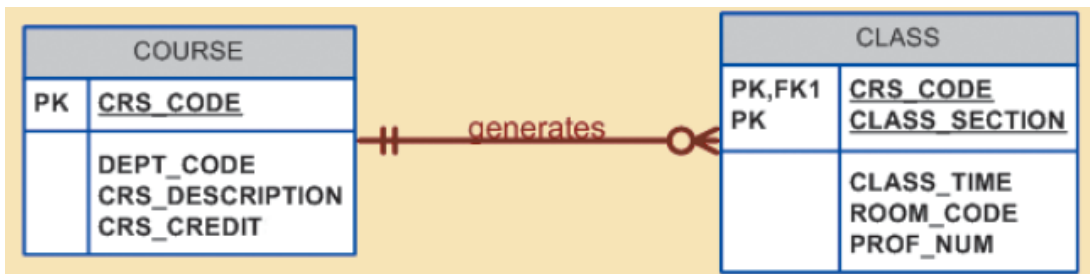

CROW'S FOOT NOTATION DEPICTS A STRONG RELATIONSHIP

WEAK ENTITIES: A WEAK ENTITY IS ONE THAT MEETS TWO CONDITIONS:
- THE ENTITY IS EXISTENCE-DEPENDENT; THAT IS, IT CANNOT EXIST WITHOUT THE ENTITY WITH WHICH IT HAS A RELATIONSHIP.
- THE ENTITY HAS A PRIMARY KEY THAT IS PARTIALLY OR TOTALLY DERIVED FROM THE PARENT ENTITY IN THE RELATIONSHIP.

FOR EXAMPLE, A COMPANY INSURANCE POLICY INSURES AN EMPLOYEE AND HIS/HER DEPENDENTS. FOR THE PURPOSE OF DESCRIBING AN INSURANCE POLICY, AN

EMPLOYEE MIGHT OR MIGHT NOT HAVE A DEPENDENT, BUT THE DEPENDENT MUST BE ASSOCIATED WITH AN EMPLOYEE. MOREOVER, THE DEPENDENT CANNOT EXIST WITHOUT THE EMPLOYEE; THAT IS, A PERSON CANNOT GET INSURANCE COVERAGE AS A DEPENDENT UNLESS S(HE) HAPPENS TO BE A DEPENDENT OF AN EMPLOYEE. DEPENDENT IS THE WEAK ENTITY IN THE RELATIONSHIP "EMPLOYEE HAS DEPENDENT."
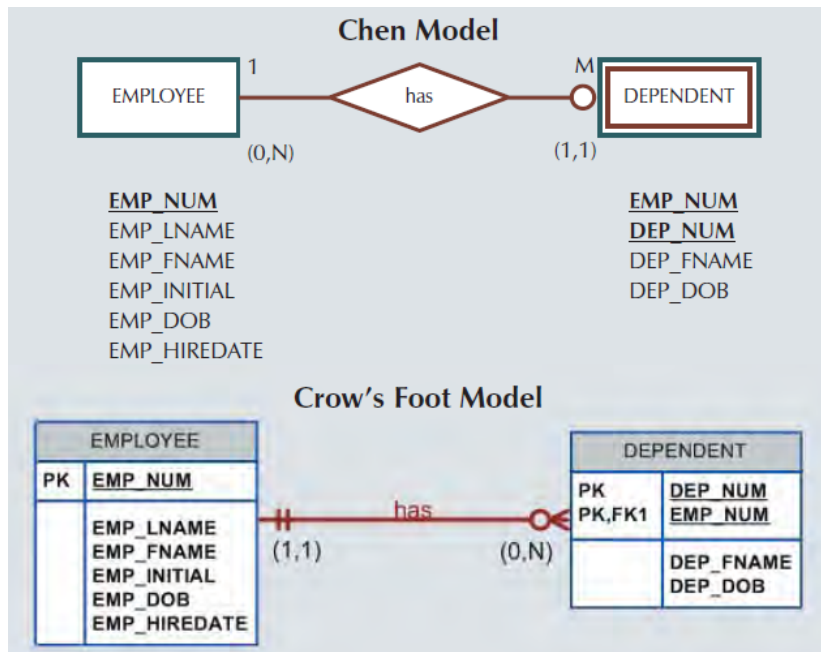
NOTE THAT THE CHEN NOTATION ABOVE IDENTIFIES THE WEAK ENTITY BY USING A DOUBLE-WALLED ENTITY RECTANGLE. THE CROW'S FOOT NOTATION GENERATED BY VISIO PROFESSIONAL USES THE RELATIONSHIP LINE AND THE PK/FK DESIGNATION TO INDICATE WHETHER THE RELATED ENTITY IS WEAK.

A STRONG (IDENTIFYING) RELATIONSHIP INDICATES THAT THE RELATED ENTITY IS WEAK. SUCH A RELATIONSHIP MEANS THAT BOTH CONDITIONS FOR THE WEAK ENTITY DEFINITION HAVE BEEN MET—THE RELATED ENTITY IS EXISTENCE-DEPENDENT, AND THE PK OF THE RELATED ENTITY CONTAINS A PK COMPONENT OF THE PARENT ENTITY. REMEMBER THAT THE WEAK ENTITY INHERITS PART OF ITS PRIMARY KEY FROM ITS STRONG COUNTERPART. FOR EXAMPLE, AT LEAST PART OF THE DEPENDENT ENTITY'S KEY SHOWN IN FIGURE ABOVE WAS INHERITED FROM THE EMPLOYEE ENTITY:

EMPLOYEE (EMP_NUM, EMP_LNAME, EMP_FNAME, EMP_INITIAL, EMP_DOB, EMP_HIREDATE)
DEPENDENT (EMP_NUM, DEP_NUM, DEP_FNAME, DEP_DOB)

CROW'S FOOT SYMBOL	CARDINALITY	COMMENT
O<	(0,N)	Zero or many. Many side is optional.
I<	(1,N)	One or many. Many side is mandatory.
II	(1,1)	One and only one. 1 side is mandatory.
O	(0,1)	Zero or one. 1 side is optional.

CROWFOOT SYMBOLS

- RELATIONSHIP DEGREE: A RELATIONSHIP DEGREE INDICATES THE NUMBER OF ENTITIES OR PARTICIPANTS ASSOCIATED WITH A RELATIONSHIP. A UNARY RELATIONSHIP EXISTS WHEN AN ASSOCIATION IS MAINTAINED WITHIN A SINGLE ENTITY. A BINARY RELATIONSHIP EXISTS WHEN TWO ENTITIES ARE ASSOCIATED. A TERNARY RELATIONSHIP EXISTS WHEN THREE ENTITIES ARE ASSOCIATED. ALTHOUGH HIGHER DEGREES EXIST, THEY ARE RARE AND ARE NOT SPECIFICALLY NAMED. (FOR EXAMPLE, AN ASSOCIATION OF FOUR ENTITIES IS DESCRIBED SIMPLY AS A FOUR-DEGREE RELATIONSHIP.)
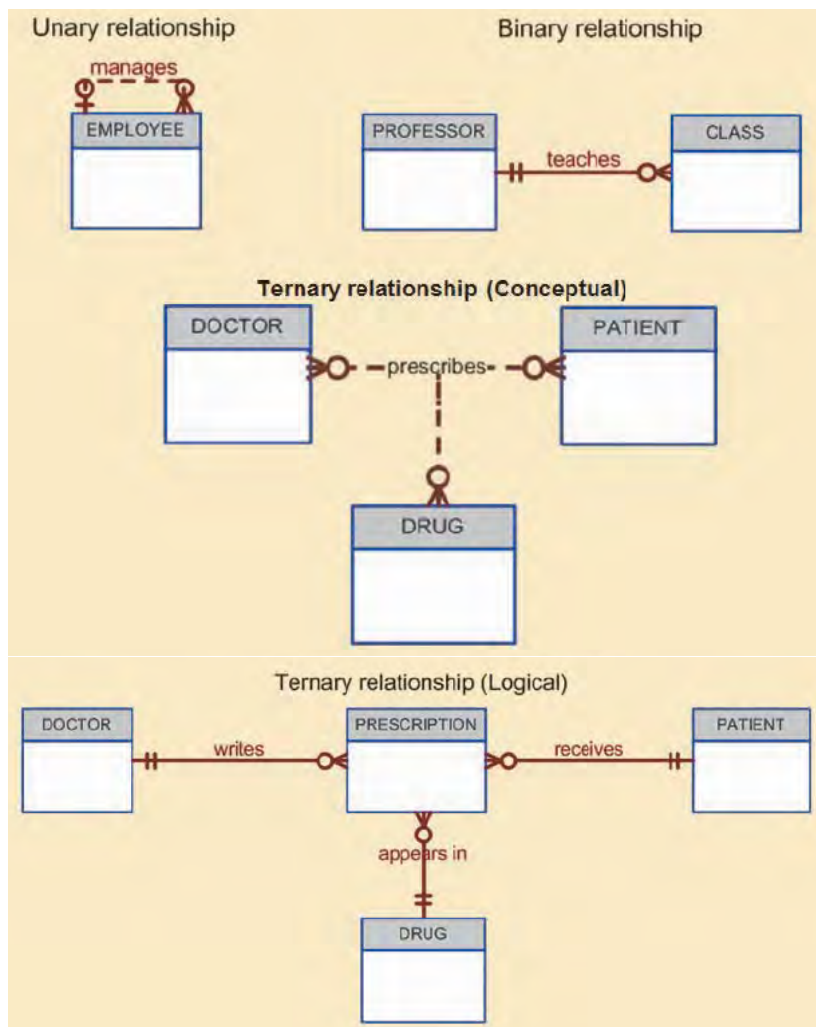
Three types of relationship degree

- o Unary Relationships: In the case of the unary relationship shown above, an employee within the EMPLOYEE entity is the manager for one or more employees within that entity. In this case, the existence of the "manages" relationship means that EMPLOYEE requires another EMPLOYEE to be the manager—that is, EMPLOYEE has a relationship with itself. Such a relationship is known as a recursive relationship.
- o Binary Relationships A binary relationship exists when two entities are associated in a relationship. Binary relationships are most common. In fact, to simplify the conceptual design, whenever possible, most higher-order (ternary and higher) relationships are decomposed into appropriate equivalent binary relationships.
- o Ternary and Higher-Degree Relationships: Although most relationships are binary, the use of ternary and higher-order relationships does allow the designer some latitude regarding the semantics of a problem. A ternary relationship implies an association among three different entities.

Example

Mr Brandon's the owner of SPEED CAFÉ has been having problems with the management of his Café. Having learnt that you are a DB designer, he believes he has finally found a solution. He has asked you to automate the management of his Café. Since this will involve a database backend, you are saddled with the task of showing him a good database model based on the following business rules:

- The café has several employees each having a unique identification number, names and dates of birth.
- An employee is either a "Technical Officer" or "Casual Employee", but not both. A Technical officer has access to one or more computing facilities in

the Café and therefore has login usernames and passwords. Technical officers have varying salary rates based on their ranks. Casual Employee however, do not have access to computing facilities and their salaries are wages (i.e. based on the number of hours worked).

- All Computing facilities in the Café have names (e.g. computer, cable, printer etc.) and date of purchase (remember names are not unique, so you will have to choose a surrogate key).
- Access to Internet facilities in the Café (either by a staff or customer) is through a ticket. Each ticket has a unique ticket number, duration, time of production, period (number of days) of validity and price in Naira.

Draw an implementation oriented ER diagram for SPEED CAFÉ database indicating necessary connectivities, cardinalities and participation constraints (relationship strengths). You can state the necessary assumptions made.