

### The Relational Data Model

The **relational model** (E. Codd, 1970) supports powerful, simple and declarative languages with which operations on data are expressed. We define operations on relations whose results are themselves relations.

### Domain

A **domain** is a set of values.

#### Examples:

- a set of integers is a domain
- set of character strings
- set of character strings of length 20
- set of real numbers
- set {0,1}

### Cartesian Product

The **Cartesian product** (or just **product**) of domains  $D_1, D_2, \dots, D_k$ , written  $D_1 \times D_2 \times \dots \times D_k$ , is the set of all  $k$ -tuples  $(d_1, d_2, \dots, d_k)$  such that  $d_1$  is in  $D_1$ ,  $d_2$  is in  $D_2$ , and so on.

#### Example:

$k=2$   $D_1=\{0,1\}$   $D_2=\{a,b,c\}$   
 $D_1 \times D_2 = \{(0,a), (0,b), (0,c), (1,a), (1,b), (1,c)\}$

### Relation

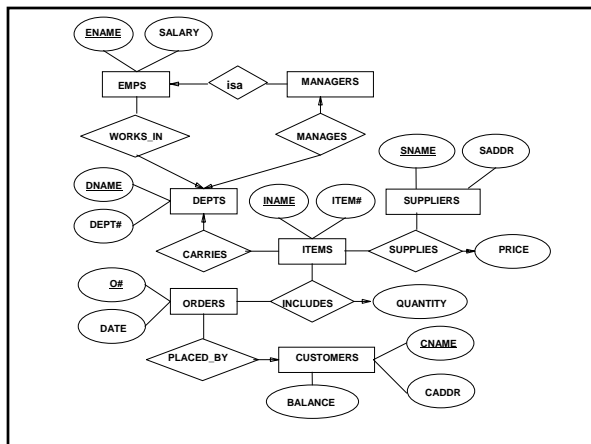
A **relation** is any subset of the Cartesian product of one or more domains. We shall assume that a relation is finite unless we state otherwise.

#### Example:

$\{(0,a), (0,c), (1,b)\}$  – subset of  $D_1 \times D_2$

The members of a relation are called **tuples**. Each relation that is a subset of some product  $D_1 \times D_2 \times \dots \times D_k$  of  $k$  domains is said to have **arity** (or **degree**)  $k$ . A tuple  $(d_1, d_2, \dots, d_k)$  has  $k$  components: the  $i$ th component is  $d_i$ . A tuple with  $k$  components is called a **k-tuple**. It helps to view a relation as a table, where each row is a tuple and each column corresponds to one component. The columns are often given names, called **attributes**. The set of attributes names for a relation is called the **relation scheme**. If we name a relation **REL**, and its relation scheme has attributes  $A_1, A_2, \dots, A_k$ , we often write the relation scheme as **REL**  $(A_1, A_2, \dots, A_k)$ .

The collection of relation schemes used to represent information is called a **(relational) database scheme**, and the current values of the corresponding relations form the **(relational) database**.



The data of an entity-relationship diagram is represented by two sorts of relations:

1. An entity set **E** can be represented by a relation whose relation scheme consists of all the attributes of the entity set. Each tuple of the relation represents one entity in the current instance of **E**.

**Example:** The entity set **CUSTOMERS** is represented by the relation **CUSTOMERS** (CNAME, CADDR, BALANCE)

If  $E$  is an entity set whose entities are identified through a relationship with some other entity set  $F$ , then the relational scheme also has the attributes of  $F$  that are needed for the key of  $E$ .

**Example:** The relation for entity set **MANAGERS** has only one attribute, **ENAME**, which is the key for **MANAGERS**. The value of **ENAME** for a given manager is the name of the employee entity that is this manager.

**MANAGERS (ENAME)**

2. A relationship  $R$  among entities  $E_1, E_2, \dots, E_k$  is represented by a relation whose relational scheme consists of the attributes in the keys for each of  $E_1, E_2, \dots, E_k$ . By renaming attributes if necessary, we make certain that no two entity sets in the list have attributes with the same name, even if they are the same entity set.

The relation scheme for the entity sets is (each entity set has the same name as the relation):

1. EMPs (ENAME, SALARY)
2. MANAGERS (ENAME)
3. DEPTS (DNAME, DEPT#)
4. SUPPLIERS (SNAME, SADDR)
5. ITEMS (INAME, ITEM#)
6. ORDERS (O#, DATE)
7. CUSTOMERS (CNAME, CADDR, BALANCE)

Now let us consider the relationships. We should not create a relation for the **isa** relationship, since it would just consist of the **ENAME** attribute repeated (and renamed in one repetition). And would hold exactly the same information as the **MANAGES** relation; that is, it would list the names of all those employees who are managers. The other six relationships yield the following relation schemes:

8. WORKS\_IN (ENAME, DNAME)
9. MANAGES (ENAME, DNAME)
10. CARRIES (INAME, DNAME)
11. SUPPLIES (SNAME, INAME, PRICE)
12. INCLUDES (O#, INAME, QUANTITY)
13. PLACED\_BY (O#, CNAME)

In each case, the set of attributes is the set of keys for the entity sets connected by the relationship of the same name as the relation.

**Example:** **SUPPLIES** connects **SUPPLIERS**, **ITEMS**, and **PRICE**, which have keys **SNAME**, **INAME**, and **PRICE**, respectively, and it is these three attributes we see in the scheme for **YVCB**.

**SUPPLIES (SNAME, INAME, PRICE)**

The two relations **MANAGES** and **WORKS\_IN** have the same set of attributes, but of course their meaning are different. The tuple  $(e, d)$  in **MANAGES** means that  $e$  manages department  $d$ , while the same tuple in **WORKS\_IN** means that  $e$  is an employee in department  $d$ .

### Keys of Relations

A set  $S$  of attributes of a relation  $R$  is a **key** if

1. No instance of  $R$  that represents a possible state of the world can have two tuples that agree in all the attributes of  $S$ , yet are not the same tuple, and
2. No proper subset of  $S$  has property (1).

**Example:**

In the relation **SUPPLIES**, **SNAME** and **INAME** together form a key. If there are two tuples  $(s, i, p_1)$  and  $(s, i, p_2)$  in **SUPPLIES**, then supplier  $s$  would apparently sell item  $i$  both at price  $p_1$  and at price  $p_2$ , a situation that means our data is faulty. This observation justifies condition (1).

To check (2) we have to consider the proper subset, that is **SNAME** alone and **INAME** alone. Neither should satisfy condition (1). For example, it is quite possible that we find the two tuples

(Acme, Brie, 3.50)  
 (Acme, Perrier, 1.25)

in **SUPPLIES** at the same time, and although they agree on **SNAME**, they are not the same tuple. Similarly, we might find

(Acme, Brie, 3.50)  
 (Ajax, Brie, 3.95)

showing that **INAME** alone does not satisfy condition (1).

A relation may have more than one key.  
**Example:** Consider **DEPTS (DNAME, DEPT#)**. We do not give two departments the same name, and we do not give two departments the same number, so we may declare that **DNAME** is a key and **DEPT#** is a different key. But it is useful to select one unique key.

**Primary key** is a unique key selected from among several choices, all of which are called **candidate keys**.

The rules are:

1. If a relation comes from an entity set, a set of attributes is a key for that relation if it is a key for the entity set.
2. If a relation comes from a many-many relationship, then the key for the relation is normally the set of all the attributes.
3. If a relation comes from a one-to-one relationship between entity set **E** and **F**, then the key for **E** and the key for **F** are both keys for the relation. That relations, like entity sets, can have more than one set of attributes that is a candidate key.
4. If a relation comes from a relationship that is many-one from  $E_1, E_2, \dots, E_{k-1}$  to  $E_k$ , then the set of attributes that is the union of the keys for  $E_1, E_2, \dots, E_{k-1}$  is normally a key for the relation.

**Examples:**  
 In **SUPPLIES** the lone key consists of two attributes, since the relationship is many-one from **SUPPLIERS** and **ITEMS** to **PRICE** (4), and the first two entity sets have keys **SNAME** and **INAME**, respectively. Thus, **{SNAME, INAME}** forms a key for relation **SUPPLIES**.

The relation **DEPTS** has two candidate keys, each consisting of one attribute. **DNAME** is a key for entity set **DEPTS**, but we might well decide that **DEPT#** also should be a key, since the **YVCB** probably does not intend to give two departments the same number.

Relations with Common Keys  
 When two relations have a candidate key in common, we can combine the attributes and receive a relation whose set of attributes is the union of the two sets.

**Example:**  
 Relations **DEPTS** and **MANAGES** each have **DNAME** as a candidate key; in one case it is the primary key and in the other not. We may thus replace **DEPTS** and **MANAGES** by one relation

**DEPTS (DNAME, DEPT#, MGR)**

The new relation has the same name **DEPTS**. The attributes **DNAME** and **DEPT#** are the same as the attributes of the same name in the old **DEPTS** relation, while **MGR** is intended to be the attribute **ENAME** from **MANAGES**.

Dangling Tuples  
 Tuples that need to share a value with a tuple in another relation, but find no such value, are called **dangling tuples**. To avoid this problem we add to the database scheme information about **existence constraints**: if a tuple **v** appears in attribute **A** of some tuple in relation **R**, then **v** must also appear in attribute **B** of some tuple in relation **S**. We can store **null values** in certain fields. This null value may appear if the field is not a primary key.

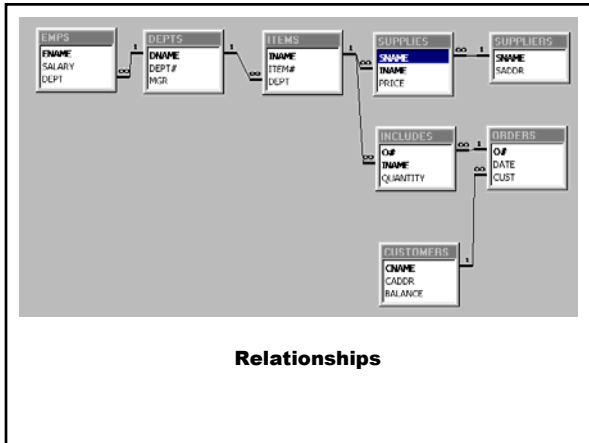
**Example:**  
 If the **YVCB** has a Wine department, whose number is 16, but that temporarily has no manager, we will represent this data with null value for the **MGR** attribute. If we add a manager Truffle of the Gourmet department, whose number is not yet assigned, we could represent this with null value for the **DEPT#** attribute.

**Example:**

The new combined simplified list of relations:

- |  |       |
|--|-------|
| 1. EMPs ( <u>ENAME</u> , SALARY, DEPT)             | 1,8   |
| 2. DEPTS ( <u>DNAME</u> , DEPT#, MGR)              | 2,3,9 |
| 3. SUPPLIERS ( <u>SNAME</u> , SADDR)               | 4     |
| 4. ITEMS ( <u>INAME</u> , ITEM#, DEPT)             | 5,10  |
| 5. ORDERS ( <u>O#</u> , DATE, CUST)                | 6,13  |
| 6. CUSTOMERS ( <u>CNAME</u> , CADDR, BALANCE)      | 7     |
| 7. SUPPLIES ( <u>SNAME</u> , INAME, PRICE)         | 11    |
| 8. INCLUDES ( <u>O#</u> , <u>INAME</u> , QUANTITY) | 12    |

We can see that the new relation **DEPTS** combines **MANAGERS**, **DEPTS**, and **MANAGES**. **DEPTS** and **MANAGES** shared the common candidate key **DNAME**. However, **MANAGERS**, with key, **ENAME**, does not share a common candidate key with these. But **MANAGES** is a one-to-one relationship between **ENAME** and **DNAME**. Hence, these two attributes are in a sense equivalent, we may regard **MANAGES** as if its attributes were **DNAME** rather than **ENAME**.



**Operations in the Relational Data Model**

A relation is a set of **k**-tuples for some **k**, called the **arity** of the relation. The operands of relational algebra are either constant relations or variables denoting relations of a fixed arity.

**1. Union.** The union of relations **R** and **S**, denoted **RUS**, is the set of tuples that are in **R** or **S** or both. **R** and **S** are relations of the same arity.

**Example:**

<b>A B C</b>	<b>D E F</b>	<b>a b c</b>
a b c	b g a	d a f
d a f	d a f	c b d
c b d		b g a
<b>R</b>	<b>S</b>	<b>RUS</b>

**2. Set difference.** The difference of relations **R** and **S**, denoted **R-S**, is the set of tuples in **R** but not in **S**. **R** and **S** have the same arity.

**Example:**

<b>A B C</b>	<b>D E F</b>	<b>a b c</b>
a b c	b g a	c b d
d a f	d a f	
c b d		
<b>R</b>	<b>S</b>	<b>R-S</b>

**3. Cartesian product.** Let **R** and **S** be relations of arity **k<sub>1</sub>** and **k<sub>2</sub>**, respectively. Then **RxS**, the product of **R** and **S**, is the set of all possible **(k<sub>1</sub>+k<sub>2</sub>)**-tuples whose first **k<sub>1</sub>** components form a tuple in **R** and whose last **k<sub>2</sub>** components form a tuple in **S**.

**Example:**

<b>A B C</b>	<b>D E F</b>	<b>A B C D E F</b>
a b c	b g a	a b c b g a
d a f	d a f	a b c d a f
c b d		d a f b g a
<b>R</b>	<b>S</b>	d a f d a f
		c b d b g a
		c b d d a f
		<b>RxS</b>



Relational Algebra as a Query Language

**Example:**

Consider the relation **SUPPLIES**. Which suppliers supply Brie?

$\Pi_{\text{SNAME}}(\sigma_{\text{INAME}=\text{"Brie"}}(\text{SUPPLIES}))$

The result will be a list of all suppliers of Brie.

**Example:**

What items supplier 'Acme' sells for less than \$5, and the prices of each?

$\Pi_{\text{I\&Aacute;NAME,PRICE}}(\sigma_{\text{SNAME}=\text{"Acme"} \wedge \text{PRICE} < 5}(\text{SUPPLIES}))$