



Relational Algebra 1

Unit 3.3

Extra SQL

<http://www.dcs.napier.ac.uk/~cs171/LJOld/database/>

Relational Algebra



Relational Algebra is :

- the formal description of how a relational database operates
- the mathematics which underpins SQL operations.

Operators in relational algebra are not necessarily the same as SQL operators, even if they have the same name.

Terminology



- **Relation** - a set of tuples
- **Tuple** - a collection of attributes which describe some real world entity.
- **Attribute** - a real world role played by a named domain.
- **Domain** - a set of atomic values.
- **Set** - a mathematical definition for a collection of objects which contains no duplicates.

TABLE

ROW

COLUMN



Operators - Write

INSERT – into a relation. This operator is the **same as SQL**.

DELETE - provides a condition on the attributes of a relation to determine which tuple(s) to remove from the relation. This operator is the **same as SQL**.

MODIFY - changes the values of one or more attributes in one or more tuples of a relation, as identified by a condition operating on the attributes of the relation. This is **equivalent to SQL UPDATE**.



Operators - Retrieval

There are **two groups** of operations:

- Mathematical **set** theory based **operations**:
UNION, INTERSECTION, DIFFERENCE, and CARTESIAN
PRODUCT


- Special **database operations**:

SELECT (*not the same as SQL SELECT*),
PROJECT
JOIN



Relational SELECT

- **SELECT** is used to obtain a **subset of the tuples** of a relation that satisfy a select condition.
- For example, *find all employees born after 1st Jan 1950:*
- **SELECT** `dob > '01/JAN/1950'` (**employee**)






Relational PROJECT

The **PROJECT** operation is used to select a **subset of the attributes** of a relation by specifying the names of the required attributes.

For example, to *get a list of all employees surnames and employee numbers*:

- **PROJECT** `surname, empno` (**employee**)






SELECT and PROJECT

SELECT and PROJECT can be combined together.
For example, to get a *list of employee numbers for employees in department number 1*:

- **PROJECT** *empPno* (**SELECT** *depno = 1* (**employee**))

Mapping this back to SQL gives:

```
SELECT empno
FROM employee
WHERE depno = 1;
```





Set Operations - semantics

Consider two relations R and S:

- **UNION** of R and S

the union of two relations is a relation that includes **all the tuples that are either** in R or in S, **or** in **both** R and S.
Duplicate tuples are eliminated.

- **INTERSECTION** of R and S

the intersection of R and S is a relation that includes **all tuples that are in both** R and S.

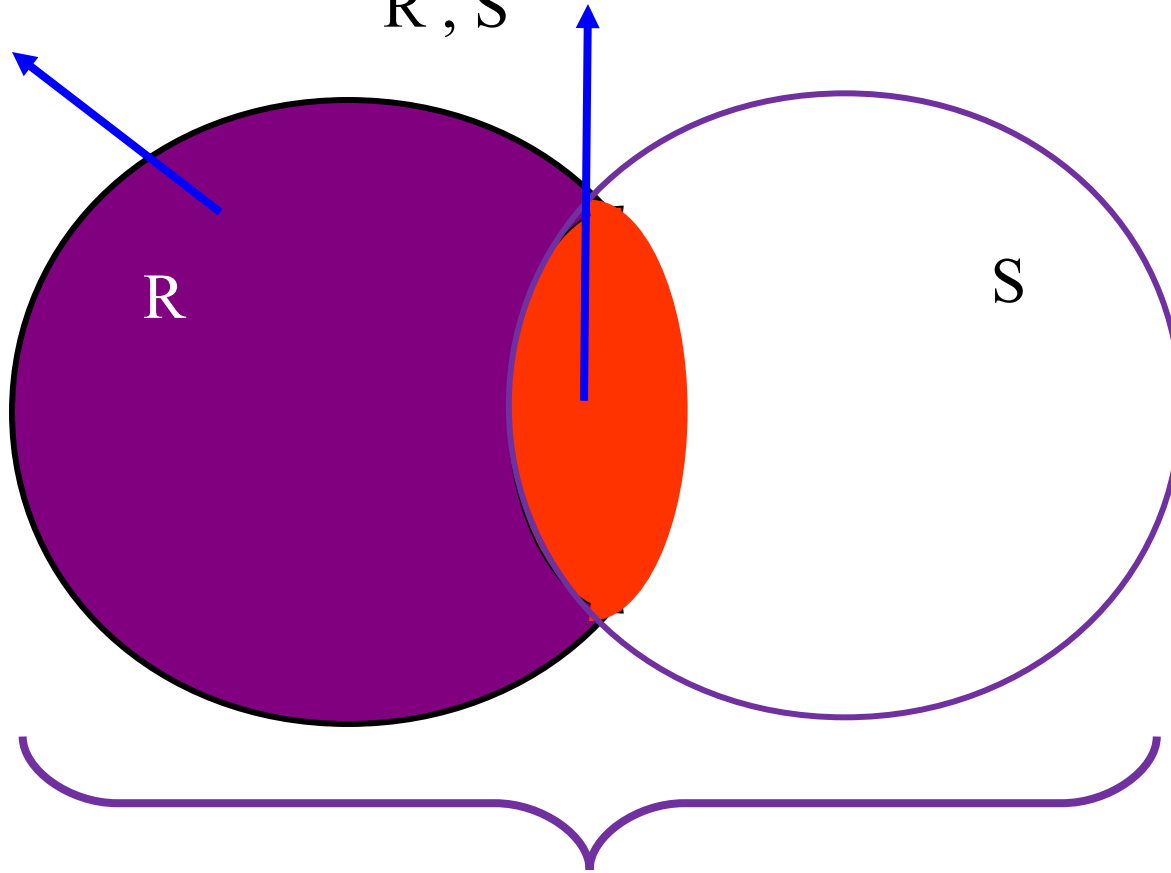
- **DIFFERENCE** of R and S

the difference of R and S is the relation that contains all **tuples that are in R but** that are **not in S.**

Sets

DIFFERENCE
R, S

INTERSECTION
R, S




UNION
R, S



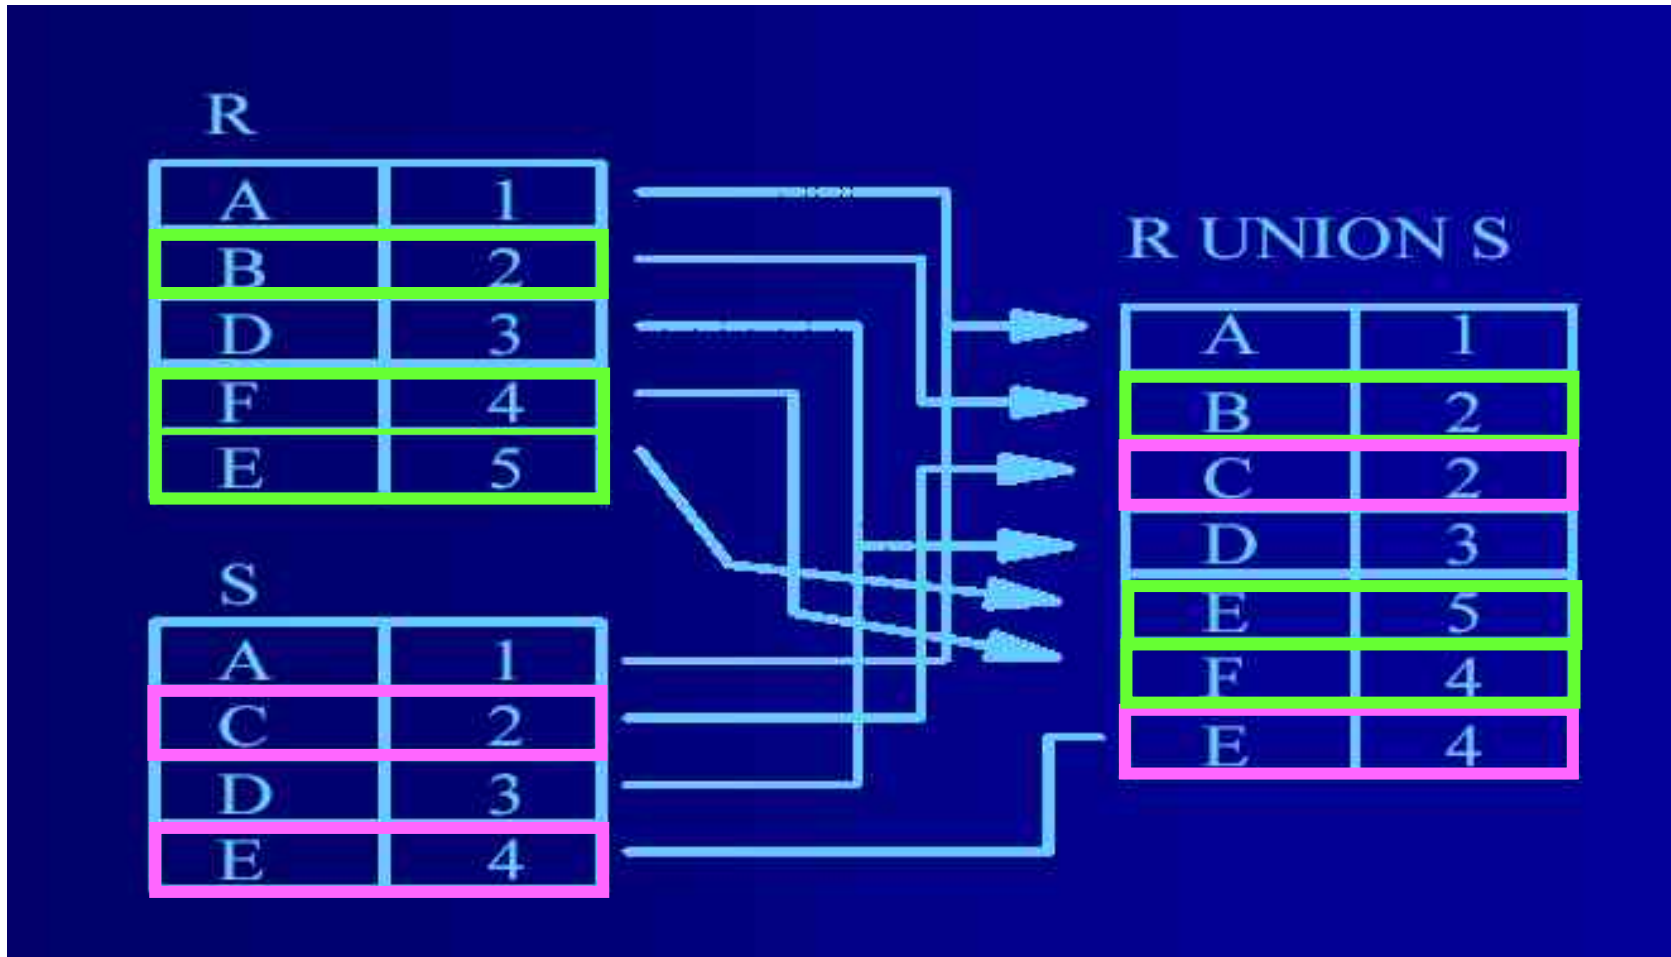
Set Operations - Requirements

For set operations to function correctly the relations R and S must be **union compatible**. Two relations are union compatible if:

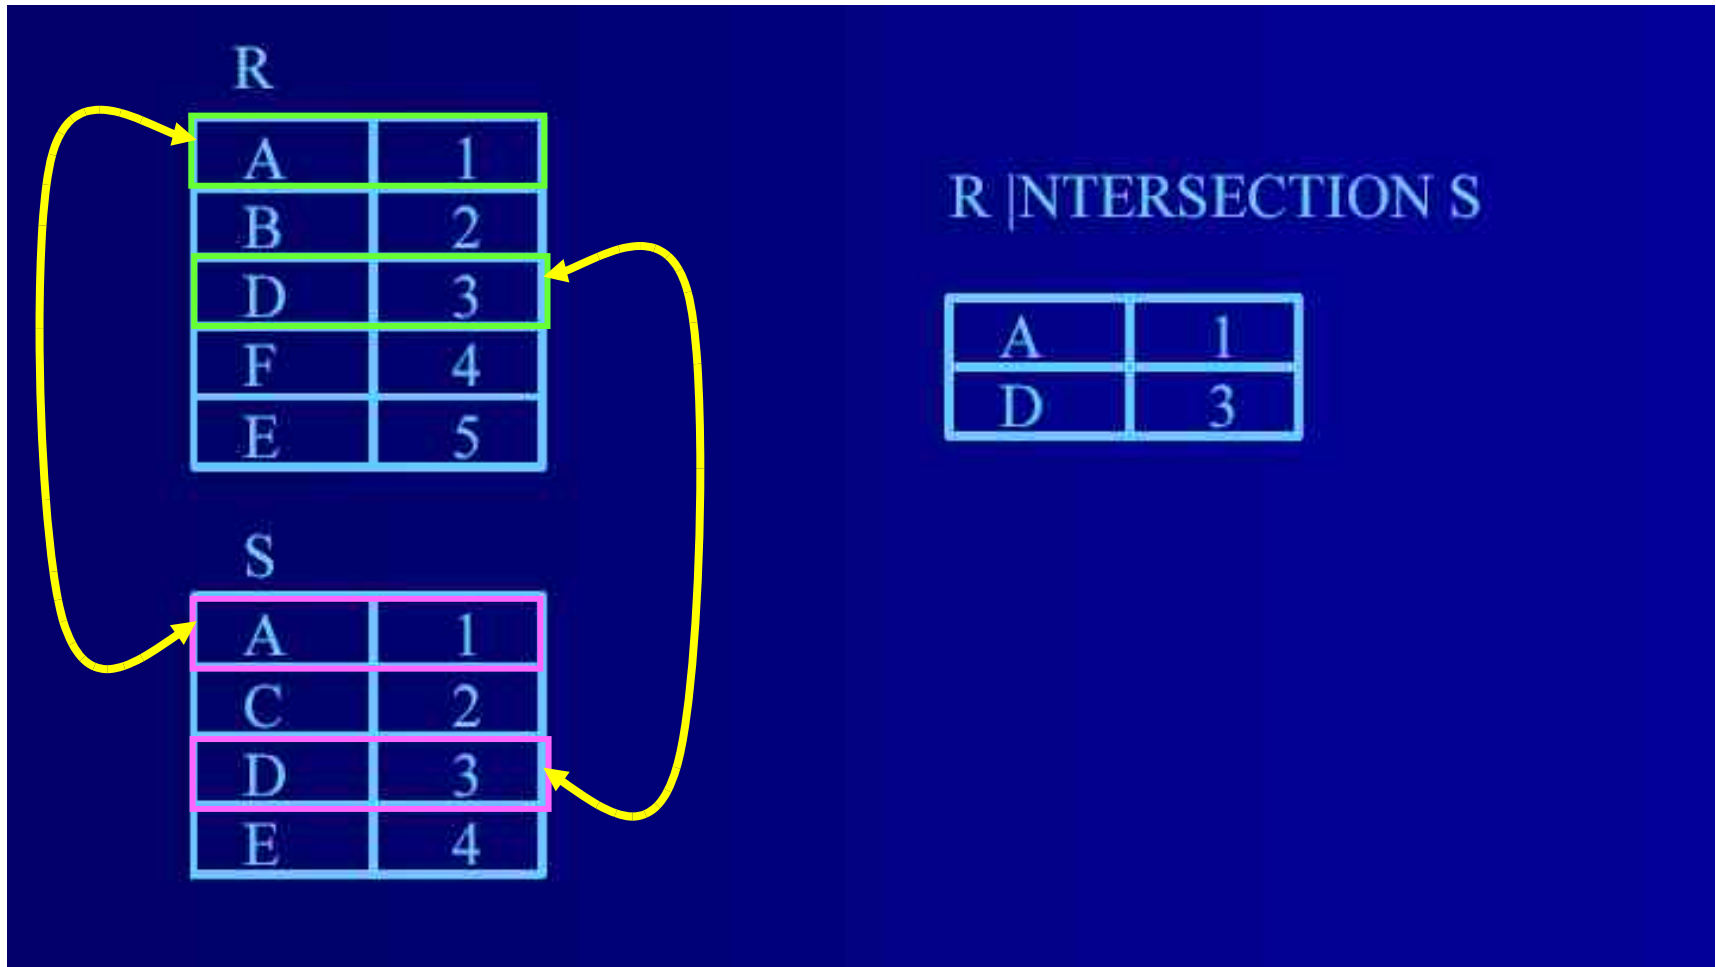
- they have the **same number of attributes**
- the **domain** of each attribute *in column order* **is the same** in both R and S



Union example



Intersection example



Difference example

R

A	1
B	2
D	3
F	4
E	5

S

A	1
C	2
D	3
E	4

R DIFFERENCE S

B	2
F	4
E	5

S DIFFERENCE R

C	2
E	4




Cartesian Product

The Cartesian Product is also an operator which works on two sets. It is sometimes called the CROSS PRODUCT (X) or CROSS JOIN.

It combines the tuples of one relation with all the tuples of the other relation, i.e., every tuple in the first relation is combined with every tuple in the second relation.

Cartesian Product - example



JOIN Operator



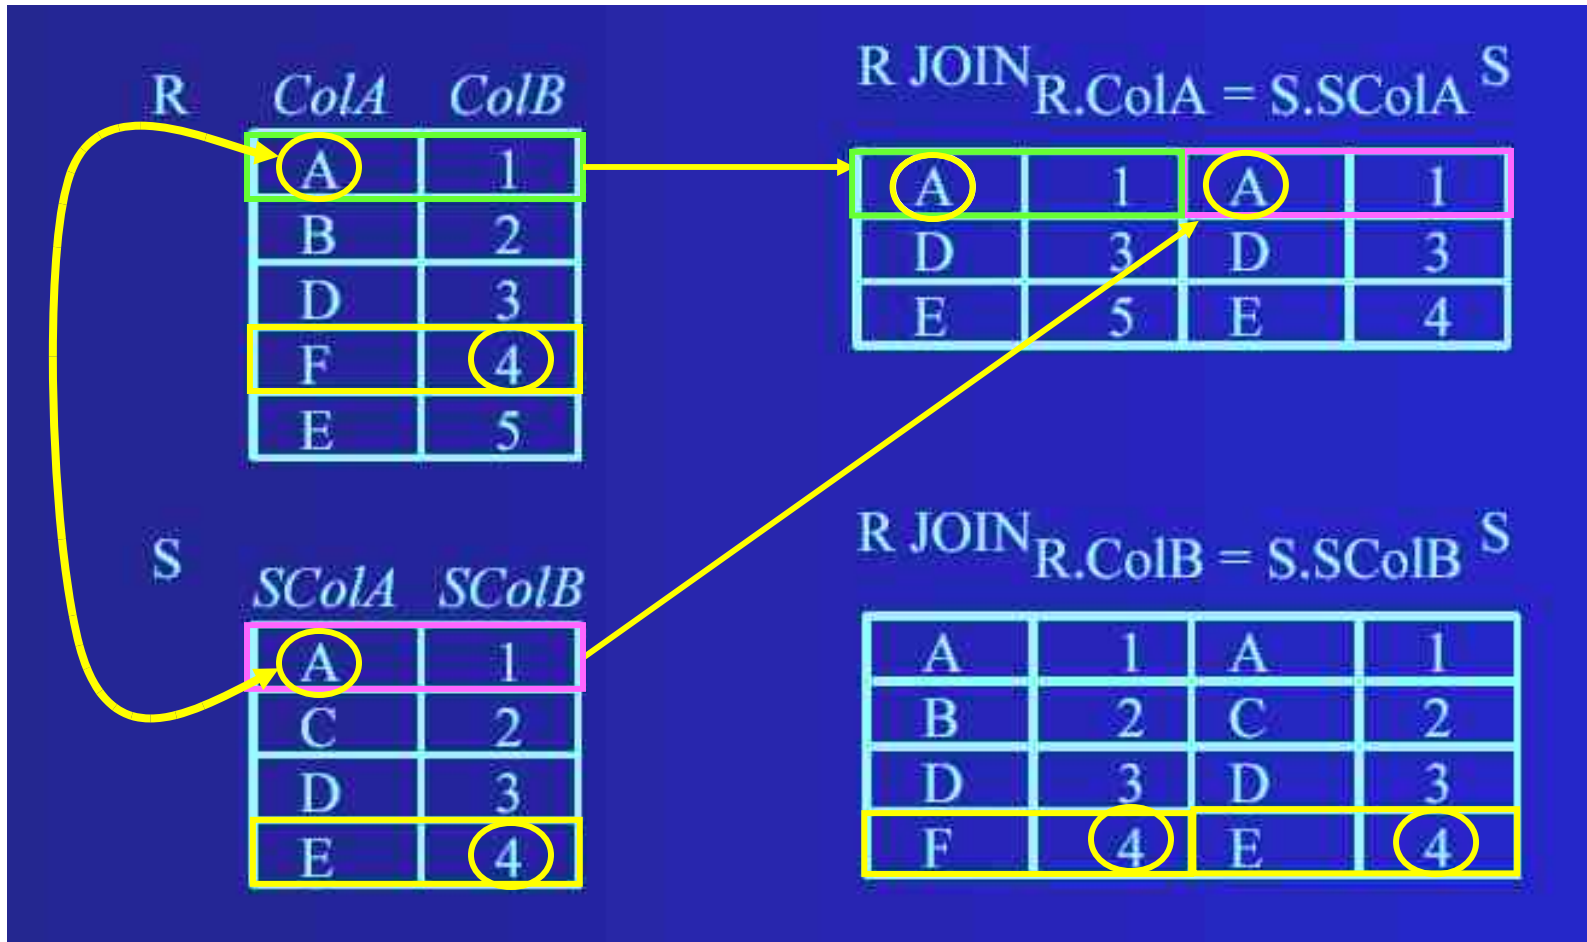
JOIN is used to combine related tuples from two relations:

- In its simplest form the JOIN operator is just the cross product of the two relations.
- As the join becomes more complex, tuples are removed within the cross product to make the result of the join more meaningful. (*conditions / constraints / rules*)
- JOIN allows you to evaluate a *join condition* between the attributes of the relations on which the join is undertaken.

The notation used is:

R JOIN *join condition* **S**
Dr Gordon Russell

JOIN examples



Natural Join

Frequently the JOIN involves an **equality test**, and thus is often described as an equi-join. Such joins result in two attributes in the resulting relation having exactly the same value (redundancy).

A ‘natural join’ will remove the duplicate attribute(s).

- In some systems a natural join will require that the attributes have the same name to identify the attribute(s) to be used in the join. This may require a renaming mechanism.
- If you do use natural joins make sure that the relations do not have two (additional) attributes with the same name by accident E.G. PERSON.Name and COURSE.Name.

OUTER JOINS




Notice that much of the data is lost when applying a join to two relations. In some cases this lost data might hold useful information. **An outer join retains the information that would have been lost** from the tables, replacing missing data with nulls.

As in SQL, there are three forms of the outer join, depending on which data is to be kept.

- **LEFT** OUTER JOIN - keep data from the left-hand table
- **RIGHT** OUTER JOIN - keep data from the right-hand table
- **FULL** OUTER JOIN - keep data from both tables

OUTER JOINS – (Left & Right)



OUTER JOIN – Full

