Introduction to Relational Databases

Hugh Darwen

hugh@dcs.warwick.ac.uk
www.dcs.warwick.ac.uk/~hugh

CS252.HACD: Fundamentals of Relational Databases
Section 1: Introduction

1

Some Preliminaries

The theory taught in this part of the course was originally devised by Edgar F. Codd in 1969. His seminal paper (1970) was entitled *A Relational Model of Data for Large Shared Data Banks*.

We will use a language called **Tutorial D** for examples and exercises.

We will use *Rel*, an implementation of **Tutorial D,** for our on-line work.

2

What Is a Database?

An *organised*, machine-readable collection of *symbols*, to be *interpreted* as a *true* account of some *enterprise*.

Machine-*updatable*, too …
… so a database is also a collection of *variables*.

A database is typically available to a community of *users*, with possibly varying requirements.

3

"Organised Collection of Symbols"

For example:

StudentId	Name	CourseId
S1	Anne	C1
S1	Anne	C2
S2	Boris	C1
S3	Cindy	C3

The symbols are organised into rows and columns, thus forming a table. One of the rows is different in kind from the others.

4

"To Be Interpreted as a True Account"

For example (from the table just shown):

StudentId	Name	CourseId
S1	Anne	C1

Perhaps those green symbols, organised as they are with respect to the blue ones, are to be understood to mean:

"Student S1, named Anne, is enrolled on course C1."

5

"Collection of Variables"

ENROLMENT

StudentId	Name	CourseId
S1	Anne	C1
S1	Anne	C2
S2	Boris	C1
S3	Cindy	C3
S4	Devinder	C1

ENROLMENT is a *variable*. Perhaps the table we saw earlier was once its *value*. If so, it (the variable) has been *updated* since then (the row for S4 has been added).

6

What Is a Relational Database?

A database whose symbols are organised into a collection of *relations*. Here is a relation, shown in tabular form:

StudentId	Name	CourseId
S1	Anne	C1
S1	Anne	C2
S2	Boris	C1
S3	Cindy	C3
S4	Devinder	C1

Might be the value currently assigned to ENROLMENT, a *relation variable* ("relvar").

7

"Relation" not equal to "Table"

This table is different from the one we have just seen, but it represents the same relation:

Name	StudentId	CourseId
Devinder	S4	C1
Cindy	S3	C3
Anne	S1	C1
Boris	S2	C1
Anne	S1	C2

In other words, the relation represented does not depend on the order in which we place the rows or the columns in the table.

8

Anatomy of a Relation

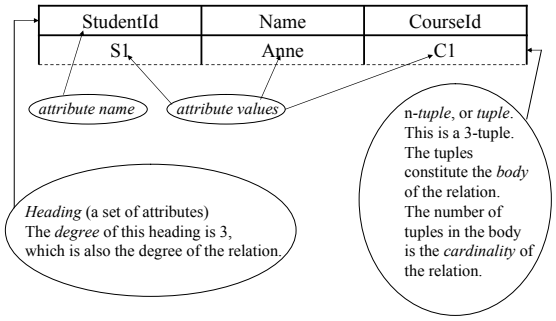

9

What Is a DBMS?

A piece of software for managing databases and providing access to them.

A DBMS responds to *imperatives* ("statements") given by *application programs*, custom-written or general-purpose, executing on behalf of users.

Imperatives are written in the *database language* of the DBMS (e.g., SQL).

Responses include completion codes, messages and results of *queries*.

10

What Does a DBMS Do?

In response to requests given by application programs:

- creates and destroys variables
- takes note of integrity rules (*constraints*)
- takes note of *authorisations* (who is allowed to do what, to what)
- updates variables (honouring constraints and authorisations)
- provides results of *queries*
- and more

Now, how does a *relational* DBMS do these things? …

11

Create and Destroy Variables

Creation (in **Tutorial D**):

```
VAR ENROLMENT BASE RELATION
        { StudentId   SID ,
          Name        CHAR,
          CourseId    CID }
      KEY { StudentId, CourseId } ;
```

Destruction:

```
DROP VAR ENROLMENT ;
```

12

Take Note of Integrity Rules

E.g., can't have more than 20,000 enrolments altogether. In **Tutorial D**:

```
CONSTRAINT MAX_ENROLMENTS
    COUNT ( ENROLMENT ) <= 20000  ;
```

And if a constraint ceases to be applicable:

```
DROP CONSTRAINT MAX_ENROLMENTS ;
```

13

Take Note of Authorisations

E.g. (perhaps – but not in **Tutorial D**):

```
PERMISSION U9_ENROLMENT FOR User9
TO READ ENROLMENT  ;

PERMISSION U8_ENROLMENT FOR User8
TO UPDATE ENROLMENT ;
```

Permissions sometimes need to be withdrawn:

```
DROP PERMISSION U9_ENROLMENT ;
```

14

Updates Variables

E.g.:

```
DELETE ENROLMENT WHERE StudentId =
SID ( 'S4' ) ;

UPDATE ENROLMENT WHERE StudentId =
SID ( 'S1' ) Name := 'Ann' ;

INSERT ENROLMENT
   RELATION {
      TUPLE { StudentId SID ( 'S4' ) ,
              Name 'Devinder' ,
              CourseId  CID ( 'C1' ) } }  ;
```

15

Provides Results of Queries

E.g.: Who is enrolled on course C1?

```
(ENROLMENT WHERE CourseId = CID('C1'))
   { StudentId, Name }
```

The result is another relation! In tabular form:

StudentId	Name
S1	Anne
S2	Boris
S4	Devinder

16

EXERCISE

Consider this table:

A	B	A
1	2	3
4		5
6	7	8
9	9	?
1	2	3

Give three reasons why it cannot be representing a relation.

17

IMPORTANT BIT OF ADMIN

Are you:
(a) an overseas student visiting for just one year, or
(b) doing a degree that is completely outside the CS department? (In which case you should complete an Unusual Option Form, obtainable from your dept secretary)

If so, on exit from this lecture:

Write your name, ITS userid, and course code on the sheet of paper provided.

Otherwise you won't get access to the software needed for the Worksheets!

18