

Relational databases

GridKa School of Computing 2014

mario.lassnig@cern.ch

Agenda

I. The basics

II. Safe use of databases

III. Efficient use of databases

IV. Application development

I

The basics

What's a database, really?

a collection of associated data

consensus

a collection that necessitates a
database management system

(modern interpretation: can't reasonably fit in Excel)

database management system

Codd, 1974 - relational model

Relational Software, 1979 - Oracle

RDBMS

Oracle

Database (Commercial)

MySQL (Commercial & Libre)

Microsoft SQL Server (Commercial)

IBM DB2 (Commercial)

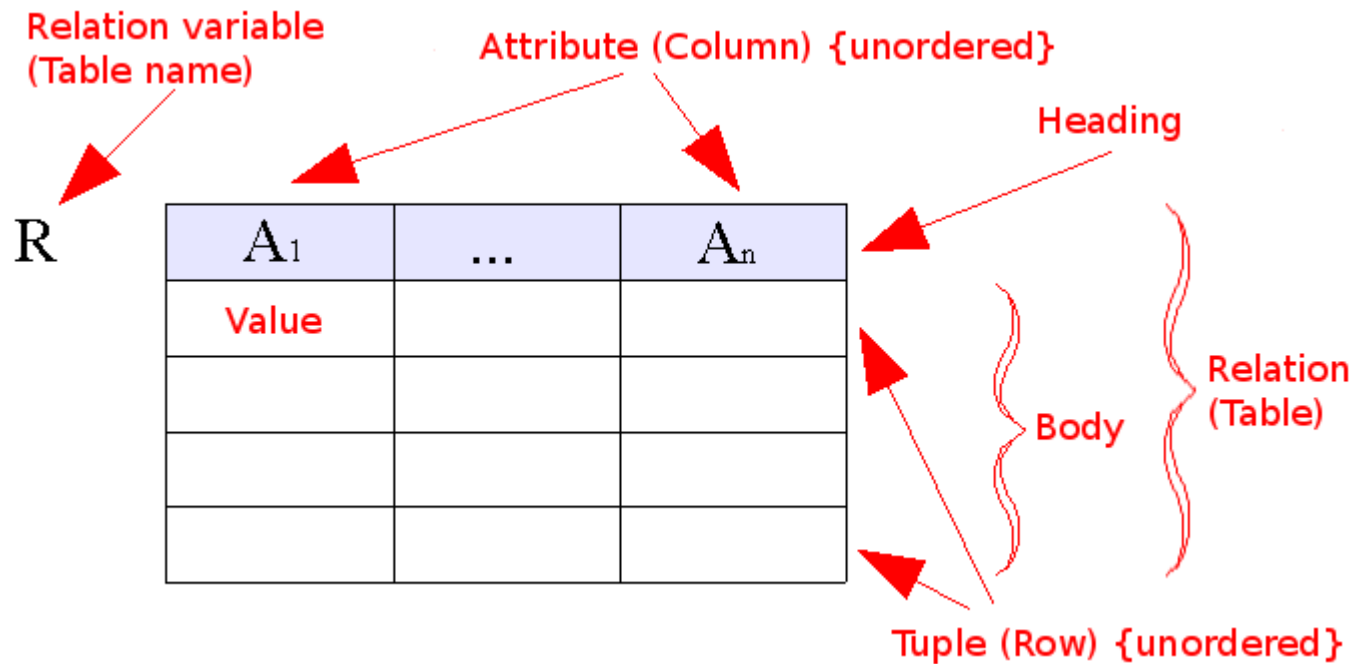
SAP Sybase (Commercial)

PostgreSQL (Libre)

MariaDB (Libre)

SQLite (Libre)

Relational model



Session 1 - Setup

the environment

creating databases

creating users

granting permissions

Session 2 - CRUD

creating tables

inserting rows

querying rows

updating rows

deleting rows

Session 3 - Relations

Joining tables

natural (cross product)

inner (equivalence, difference, ...)

outer (missing values)

self

Views

temporary

materialized (snapshot)

II

Safe use of databases

Session 4 - Constraints

Primary keys

Uniqueness

Enums

Session 5 - Transactions


Dirty read

Non-repeatable read

Phantom read

Deadlock

Session 6 - SQL Injection



III

Efficient usage

Topics

Indexes

Partitioning (range, hash, list, composite)

Execution plans

IV

Application development

Topics

Basic structure

Full application

Relational databases

GridKa School of Computing 2014

mario.lassnig@cern.ch