

Relational Model

CSE462 Database Concepts

Demian Lessa

Department of Computer Science and Engineering
State University of New York, Buffalo

January 21–24, 2011

Data Model

- A **data model** is a set of concepts used to describe the structure of the data and the constraints it should obey. It also provides operations for data retrieval and modification.
- In database design, we typically use three different data models.
 - **Conceptual**: concepts are closer to the way users perceive them.
 - **Logical**: falls between the other two, balancing user views with some representation details.
 - **Physical**: representation details such as data organization on disk, available access methods, etc.

Data Model

Why do different data models?

- Requirements Analysis
 - Determines data, applications, critical operations, etc.
- Conceptual Design
 - High-level description of data and constraints (e.g., ERM).
- Logical Design
 - Conversion of the conceptual design into a database schema.
- Schema Refinement
 - Redundancy elimination through a process called normalization.
- Physical Design
 - Considers workloads, indexes, clustering/partitioning, etc.
- Application and Security Design...

Overview

Relational Model (1970)

- Originally proposed by Ted Codd.
- Separates physical implementation from conceptual view.
- Models data **independently** from its intended or actual use.
 - Describes data both **minimally** and **mathematically**.
 - A relation describes an association between data items— **tuples** with **attributes**.
 - Uses standard mathematical (logical) operations over the data— **relational algebra** or **relational calculus**.

Definitions

The relational model represents data as two-dimensional **relations**.

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

Table: The `Movies` relation.

- Each **row** in the `Movies` relation represents a movie and each **column** a movie property. Every column header is an **attribute**.
- A **relation schema** consists of a relation name and a *set of attributes*.
Notation: `Movies(title, year, length, genre)`.
- A **database schema** is the set of all relation schemas in the database.

Definitions (cont.)

- Every relation attribute is associated with a **domain**, an elementary type such as `int` or `string`. Domains may optionally be included in relation schemas:
`Movies(title: string, year: int, length: int, genre: string)`
- Rows of a relation are called **tuples**. A tuple has one component for each relation attribute. Every tuple component must have a value that belongs to the corresponding column's domain or is `NULL` (`NULL` is not a value!). The **arity** of a tuple is the number of components in the tuple.
Notation: ('Gone With the Wind', 1939, 231, 'drama').

Required

- Read sections 2.1 and 2.2 of chapter #2.
- Review the movies database schema of section 2.2.8.

Exercise 2.2.1

The relations below constitute part of a banking database.

acctNo	type	balance	firstName	lastName	idNo	account
12345	savings	12000	Robbie	Banks	901-222	12345
23456	checking	1000	Lena	Hand	805-333	12345
34567	savings	25	Lena	Hand	805-333	23456

Table: The Accounts relation.

Table: The Customers relation.

Specify:

- The attributes of each relation.
- The tuples of each relation.
- The components of the first tuple of each relation.
- The relation schema for each relation.
- The database schema.
- A suitable domain for each attribute.
- Another equivalent way to present each relation.

Exercise 2.2.3:

Considering orders of tuples and attributes, how many different ways are there to represent a relation instance if the instance has:

- Three attributes and three tuples?
- Four attributes and five tuples?
- n attributes and m tuples?

Classwork #1

Web Page	Day	Hits
index.html	2011-01-21	18
schedule.html	2011-01-21	12
syllabus.html	2011-01-21	11
index.html	2011-01-22	18
schedule.html	2011-01-22	9
syllabus.html	2011-01-22	6

Web Statistics: Snapshot of our course's web site statistics.

- Specify a schema for `WebStats`.
 - Include attribute names, their domains, and a minimal key.
- Can ("index.html", 2011-01-22, 15) be inserted into `WebStats`?
 - Justify your answer based on your answer above.

SQL

- Structured Query Language (SQL)** is a standardized language used to specify and manipulate relational databases. It consists of a data definition language (DDL) and a data manipulation language (DML). The current standard is SQL:2008.
- SQL defines three kinds of relation– stored relations (tables), computed relations (views), and temporary tables.

SQL: Creating Tables

The `CREATE TABLE` command creates a table by specifying its schema and optional constraints. Simplified syntax:

```
CREATE TABLE tableName (  
    attr1 type1 [column_constraint [...]],  
    ...  
    attrN typeN [column_constraint [...]]  
    [, table_constraint]  
    [, ...]  
);
```

Constraints may be specified either as part of an attribute declaration (column constraint) or provided after all attribute declarations (table constraint). Certain constraints must be specified as column constraints (`DEFAULT`, `NOT NULL`) while others as table constraints (multi-column constraints).

SQL: Data Types

SQL data types (not extensive).

- `BIT (n)`, `BIT VARYING (n)`
 - Bit strings of fixed or varying length.
- `BOOLEAN`
 - Logical values, with three truth values: `TRUE`, `FALSE`, `UNKNOWN`.
- `CHAR (n)`, `VARCHAR (n)`
 - Character strings of fixed or varying length.
- `DATE`, `TIME`, `TIMESTAMP`.
 - Temporal values consisting of date, time, or date-and-time:
- Numbers
 - `INT` (also `INTEGER`), `SMALLINT`: integers.
 - `FLOAT` (also `REAL`): single-precision real numbers.
 - `DECIMAL (n, d)`: higher precision real numbers, where `n` is the number of decimal digits and `d` is the number of significant digits to the right of the decimal point.

SQL: Constraints

Specifying `NOT NULL`, `DEFAULT`, and `CHECK` constraints using the `CREATE TABLE` command:

```
CREATE TABLE tableName (  
  attr1 type1 [[NOT] NULL] [DEFAULT val1] [CHECK(expr1)],  
  ...  
  attrN typeN [[NOT] NULL] [DEFAULT valN] [CHECK(exprN)]  
  [, [CONSTRAINT chk_name] CHECK(expr)]  
  [, ...]  
);
```

Constraints:

- `NOT NULL`: tuples must have a value for that attribute at all times.
- `DEFAULT`: the value a tuple component takes if no value is supplied at insertion. If no default value is specified, `NULL` is used.
- `CHECK`: the boolean expression must evaluate to `TRUE` or `UNKNOWN` for all tuples at all times. The expression for a column constraint may only reference that column, but multiple columns for a table constraint.

SQL: Constraints (cont.)

Specifying `PRIMARY KEY` and `UNIQUE` constraints using the `CREATE TABLE` command:

```
CREATE TABLE tableName (  
  attr1 type1 [PRIMARY KEY] [UNIQUE],  
  ...  
  attrN typeN [PRIMARY KEY] [UNIQUE]  
  [, [CONSTRAINT pk_name] PRIMARY KEY(attr_list)]  
  [, [CONSTRAINT uc_name] UNIQUE(attr_list)]  
  [, ...]  
);
```

Keys may be declared as `PRIMARY KEY` or `UNIQUE`.

- No two tuples in a relation instance may agree on their key attribute values, unless one of those is `NULL`.
- None of the attributes in a `PRIMARY KEY` may be assigned `NULL`.
- A table may have at most one `PRIMARY KEY` but multiple `UNIQUE` keys.
- Multi-attribute keys must be declared as table constraints.

Example #2

The State relation: US state names and abbreviations.

```
CREATE TABLE State (  
  state CHAR(2) PRIMARY KEY,  
  name VARCHAR(30) UNIQUE  
);
```

The City relation: US cities and their associated states and populations.

```
CREATE TABLE City (  
  cid INT PRIMARY KEY,  
  name VARCHAR(100),  
  state CHAR(2),  
  population NUMERIC CHECK(population > 1000),  
  CONSTRAINT ucNameState UNIQUE(name,state), -- named  
  FOREIGN KEY(state) REFERENCES State(state) -- unnamed  
);
```

SQL: Modifying Schemas

- The `DROP TABLE` command removes a table from the database, including all its tuples:

```
DROP TABLE tableName;
```

- The `ALTER TABLE` command allows attributes to be added (1) or dropped (2) from a table:

```
1 ALTER TABLE tableName ADD attr dataType;  
2 ALTER TABLE tableName DROP attr; -- by name
```

- The `ALTER TABLE` command also allows constraints to be added (1) or dropped (2) to a table. For example:

```
1 ALTER TABLE Movies ADD PRIMARY KEY (title, year);  
2 ALTER TABLE Movies DROP CONSTRAINT pkMovies; -- by name
```

Required

- Read section 2.3 of chapter #2.
- Answer exercises 2.3.1 and 2.3.2.

