CS 235: Introduction to Databases

Svetlozar Nestorov

Administrivia

- Three handouts today:
  - Course info, Schedule, Slides.
- Latest info will always be online!
- Come to class!
- Ask questions!
- Give feedback!
- Have fun!

What is a DBMS?

- A Database Management System manages very large amounts of data and provides:
  - persistent storage,
  - efficient access,
  - concurrent access,
  - secure, atomic access.

Examples

- (Almost) Everything is a database!
  - Banking systems
  - Reservation systems
  - Libraries
  - The Web
- Varying degrees of structure, organization, and efficiency.

Relational Model

- Based on tables, as:

acct#	name	balance
12345	Kate	100000.21
76543	Rick	89.01
23500	Tom	555999.02
34567	Alice	285.48
…	…	…

- Today used in *most* DBMS's.

Relational Database Marketplace

- In 2005, $13.8B (up 8.3% from 2004)
- Market shares (revenue):
  - Oracle 48.6% (48.9%)
  - IBM 22.0% (22.4%)
  - Microsoft 15.0% (13.9%)
  - Teradata 3.2% (3.2%)
  - Sybase 2.9% (3.0%)
  - Other 8.2% (8.5%)

Source: Gartner

IT Experts Use …

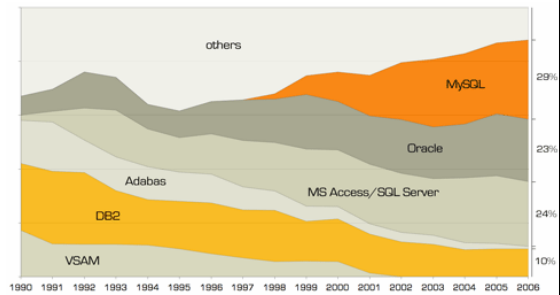

Three Aspects to Studying DBMS's

1. Modeling and design
   - Allows exploration of issues before committing to an implementation.
2. Programming
   - Queries and DB operations like update; connectivity
3. DBMS implementation
   - What's under the hood.

CS235 = (1) + (2)

Entity-Relationship Model

- First step of database design.
  - Represent the real world with diagrams.
- Entity corresponds to an object.
- Entity set corresponds to a class.
  - Set of *similar* objects.
- Attribute = property of entities in entity set.
  - Similar to fields of a struct.

E/R Diagrams

- Entity set → rectangle
- Attribute → oval

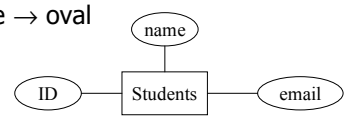

- Other conventions also exist.

Relationships

- Connect two or more entity sets.
- Represented by diamonds.

Relationship Set

- The value of a relationship set is the set of connected entities.
  - Think of the value as a table.
  - One column for each connected entity set.
  - One row for each connection.

Students	Courses
Alice	CS 235
Tom	CS 33510
Rick	CS 235
…	…

Multiway Relationships

- Binary relationships are most common.
- But, sometimes we need a relationship connecting 3 or more entity sets.
- Example: relationship among students, courses, TA's.

Multiway Relationships: Example

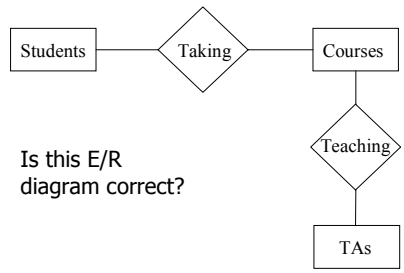

Is this E/R diagram correct?

Multiway Relationships

- Works in CS235, because the TA is a TA of all students. Connection student-TA is *only* via the course.
- But what if students were divided into sections, each headed by a TA?
  - Then, a student in CS235 would be related to only one of the TA's for CS235. Which one?
- Need a 3-way relationship to tell.

3-Way Relationship

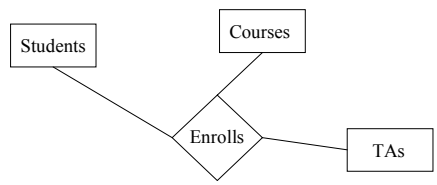

Students	Courses	TAs
Alice	CS 235	Tom
Rick	CS 235	Kate
Elaine	CS 235	Tom
…	…	…

Beers-Bars-Drinkers Example

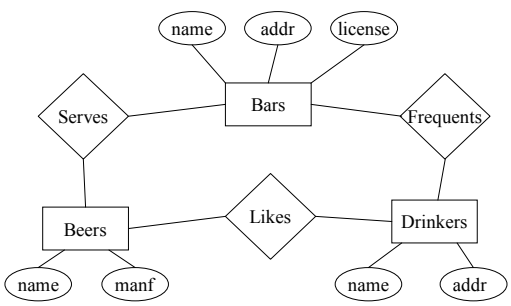

Multiplicity of Relationships

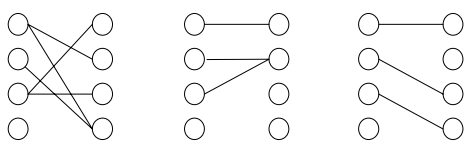

Many-many Many-one One-one

Representation of Many-One

- E/R: arrow pointing to "one."
  - Rounded arrow = "exactly one."
- Other conventions also exist.

Drinkers Have Favorite Beers

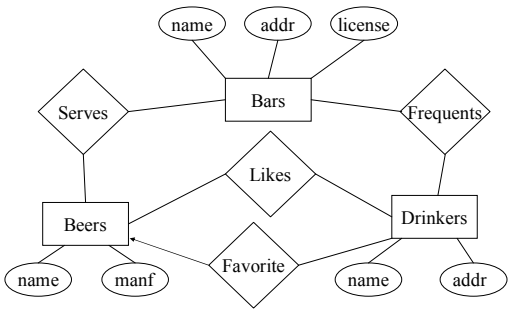

One-One Relationships

- Put arrows in both directions.

- Design issues:
  - Is the rounded arrow justified?
  - Here, manufacturer is an E.S.; in earlier diagrams it is an attribute. Which is right?

Attributes on Relationships

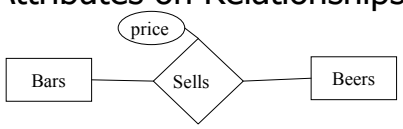

- Shorthand for 3-way relationship:

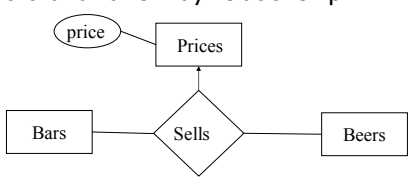

Attributes on Relationships

- A true 3-way relationship.
  - Price depends jointly on beer and bar.
- Arrow convention for multiway relationships: "all other E.S. determine one of these."
  - Not sufficiently general to express any possibility.
  - However, if price, say, depended only on the beer, then we could use two 2-way relationships: price-beer and beer-bar.
  - Better solution?

Converting Multiway to 2-Way

- Necessary in some object-oriented models.
- Create a new connecting E.S. to represent rows of a relationship set.
  - E.g., (Jimmy's, Bud, $3.50) for the *Sells* relationship.
- Many-one relationships from the connecting E.S. to the others.

Converting Multiway to 2-Way

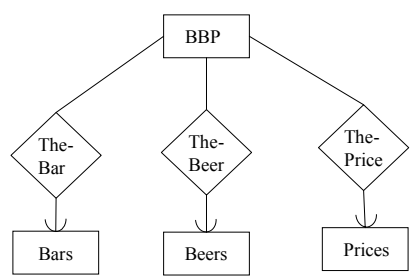

Roles

- Sometimes an E.S. participates more than once in a relationship.
- Label edges with roles to distinguish.

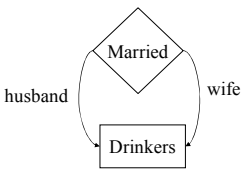

husband	wife
Tom	Kate
Johnny	Vanessa
…	…

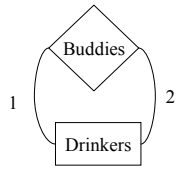

Buddy1	Buddy2
Alice	Kate
Kate	Alice
Rick	Tom
Tom	Rick
…	…

- Notice Buddies is symmetric, Married not.
  - Cannot specify symmetric in E/R.
- Should we replace *husband* and *wife* by one relationship *spouse*?