CS411
Database Systems

02: The Entity-Relationship Model

Kazuhiro Minami

Steps in building a DB application

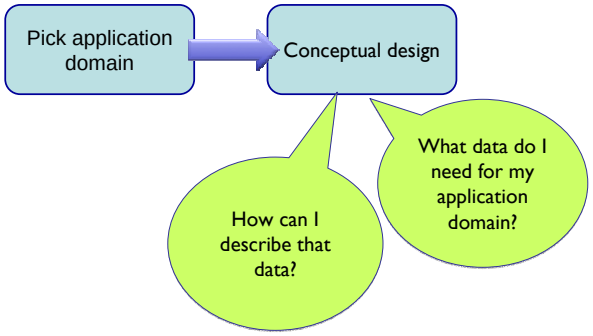

Steps in building a DB application

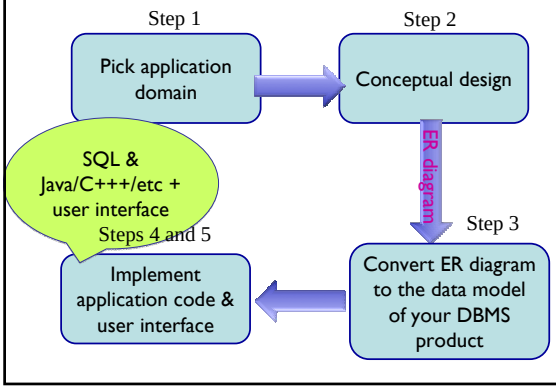

Entity Relationship (ER) Model by Peter Chen

- Born in Taiwan
- Ph.D from Harvard University in 1973
- Professor at Louisiana State University

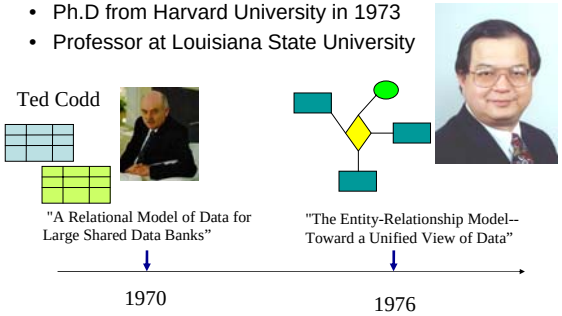

The ER model is very simple

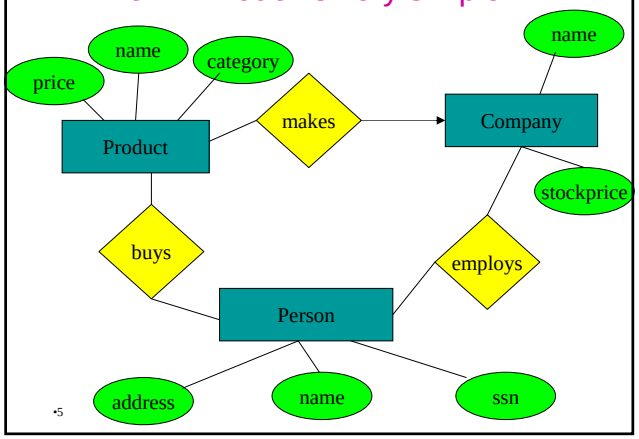

Entity

- real-world object distinguishable from other objects
- described by its *attributes*

Attribute

- Has an atomic domain: string, integers, date, …

Entity set: all have the same set of attributes

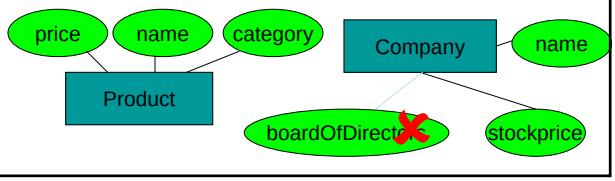

Relationships

If A, B are sets, then a relation R is a subset of A x B.

A = {1, 2, 3} B = {a, b, c, d}

R = {(1,a), (1,c), (3,b)}

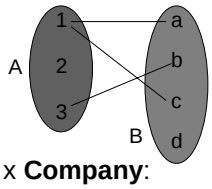

makes is a subset of **Product** x **Company**:

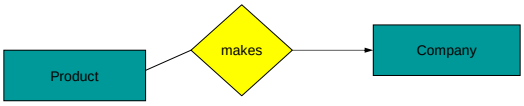

We can show the cardinality of a relationship

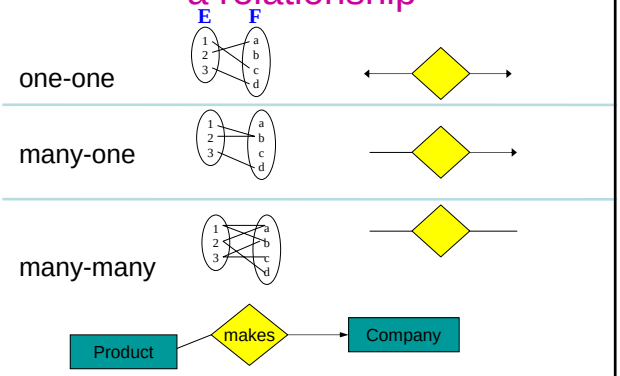

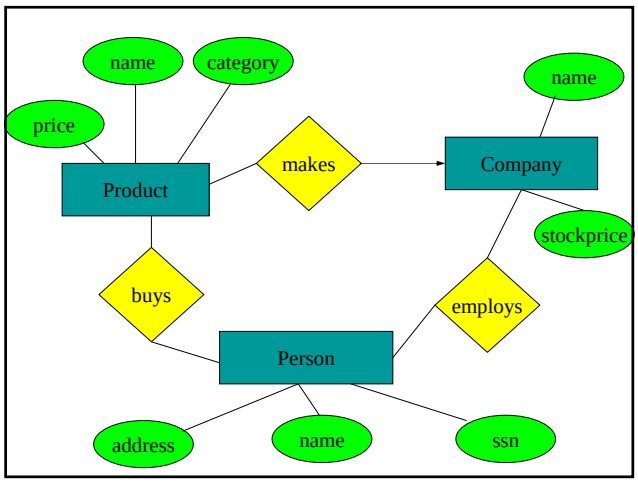

Exercise

Design a database for a bank, including information about customers and their accounts. Information about a customer includes their name, address, phone, and Social Security number. Accounts have numbers, types (e.g., saving, checking) and balances. Also record the customer(s) who own an account.

1) Draw the E/R diagram for this database.
2) Change your diagram so an account can have only one customer.
3) Further change your diagram so a customer can have only one account
4) Change your original diagram in (1) so that a customer have a set of addresses.

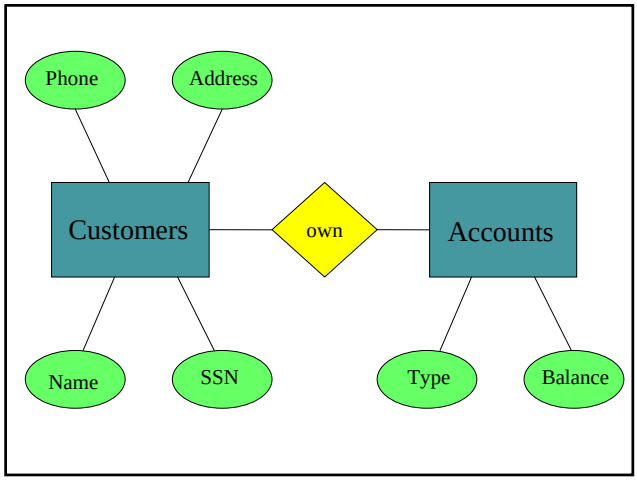

How do we model a multiway relationship?

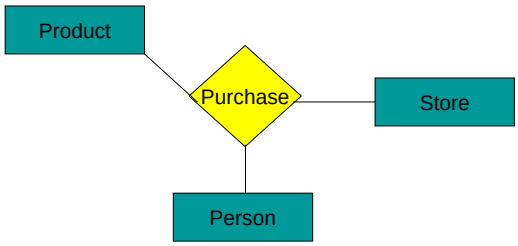

Can still model as a mathematical set (how?)

What do arrows mean in n-way relationships?

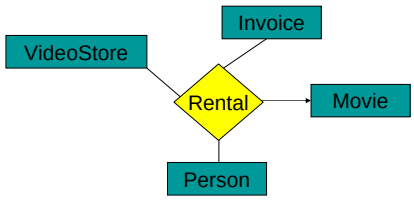

If I know the store, person, and invoice, then there is only one possible movie.

"VideoStore, Invoice, and Person ***determines*** Movie"

What if there are several arrows?

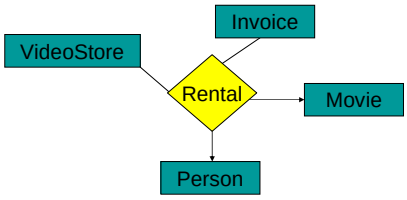

store, person, invoice determines movie;
store, invoice, movie determines person

How do I say "invoice determines store"?

No good way; best approximation:

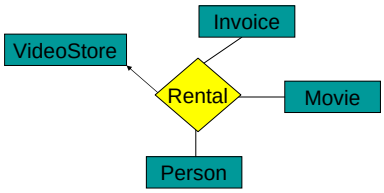

Why is this incomplete?

What if we need an entity set twice in one relationship?

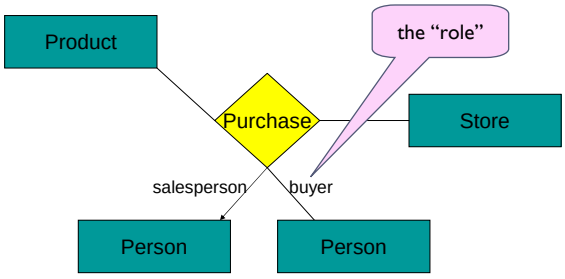

What if we need an entity set twice in one relationship?

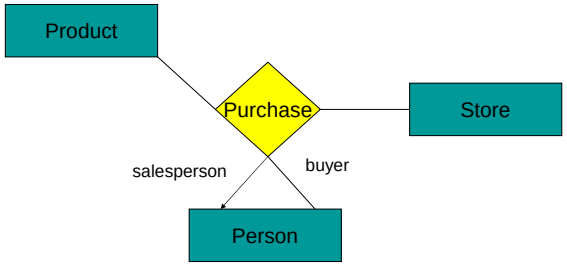

Some versions of the ER model allow attributes on relationships

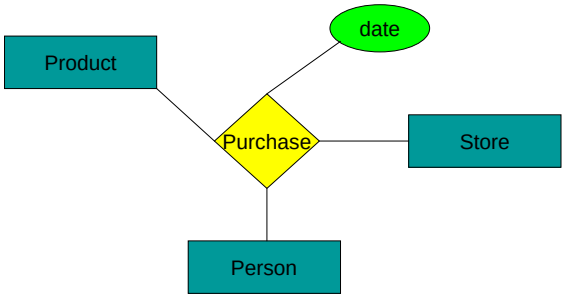

You can "upgrade" a relationship to be an entity set

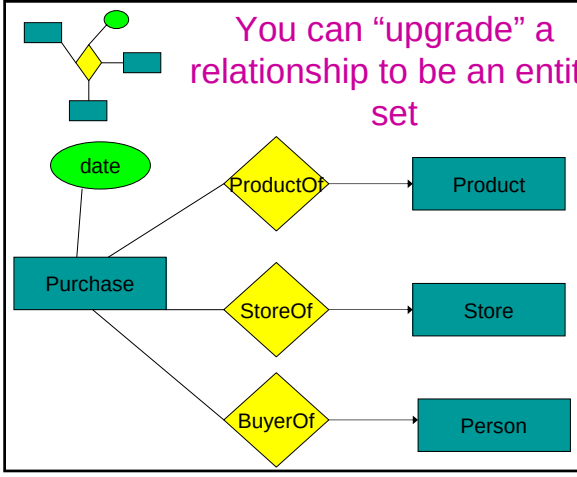

Constraint = assertion about the DB that must always be true

Key: social security number uniquely identifies a person.

Single-value constraint: a person can have only one father.

Referential integrity: if a person works for a company, the company must also be in the DB.

Domain constraint: peoples' ages are between 0 and 150.

General constraint: all others (at most 100 students in this course)

Constraints are very important

- Help us to come up with efficient storage, query processing, etc.
- Help us keep garbage out of the DB
  - Garbage in, garbage out!

Examples:
Erbana, IL 61801
Brittany Speers

Underline the key for each entity set

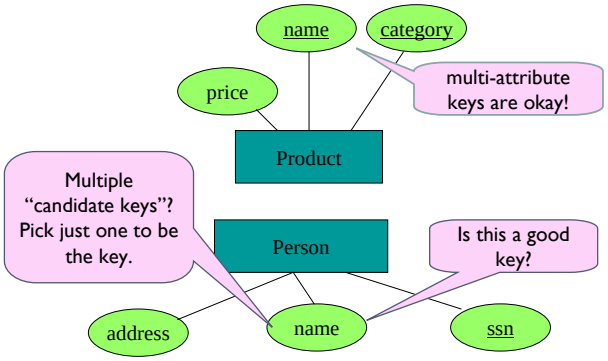

Referential Integrity Constraints

• The reference integrity constraint on relationships explicitly requires a reference to exist
• The DB equivalent of a dangling pointer

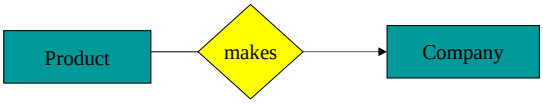

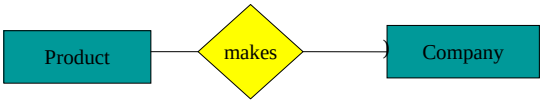

Degree Constraints

- Constraints on degree of a relationship

Sometimes your entity might not seem to have a key

Weak entity set: some or all of its key attributes come from other classes to which it is related.