By *Marina Barsky*

Entity-Relationship Model

Lecture 1

- A *data model* is a collection of concepts for describing data.

- A *schema* is a description of a particular collection of data, using a given data model.

- A *database instance* is a collection of data compliant with the schema

Databases model the real world

- “Data Model” allows us to translate real world things into structures computers can store

- Many models exist:
  - Network
  - Hierarchical
  - Relational
  - Object-Oriented
  - XML
  - etc.

In **relational** model

- Database = set of named **relations** (or **tables**)
- Each relation has a set of named **attributes** (or **columns**)
- Each **tuple** (or **row**) has a value for each attribute
- Each attribute has a **type** (or **domain**)
- Relations are connected using **keys**

Enrolled

sid	cid	grade
53666	Carnatic101	C
53666	Reggae203	B
53650	Topology112	A
53666	History105	B

Students

sid	name	login	age	gpa
53666	Jones	jones@cs	18	3.4
53688	Smith	smith@eecs	18	3.2
53650	Smith	smith@math	19	3.8

Example: relation *Movies*

This *relation* describes movies:

- Title
- Year when the movie was made
- Movie's length
- Whether the movie was in color

Title	Year	Length	FilmType
Star Wars	1997	124	color
Mighty Ducks	1991	104	color
Wayne's World	1992	95	color
…	…	…	…

Steps in Database Design

- Requirements Analysis
  - user needs; what must database do?
- Conceptual Design
  - high level modeling (often done w/ER diagrams)
- Logical Design
  - translate ER into DBMS data model
- Schema Refinement
  - consistency, normalization
- Physical refinement
  - indexes, disk layout
- Security Design
  - who accesses what, and how

ER Model: entity

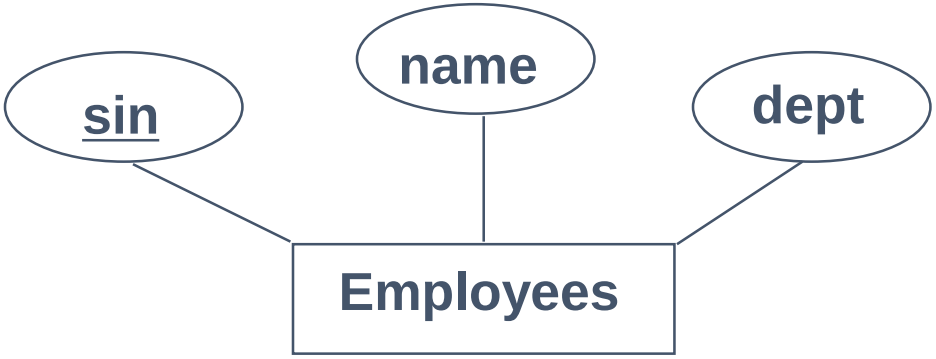

- *Entity:* Real-world object, distinguishable from other objects. An entity is described using a set of *attributes*.
- *Entity Set*: A collection of similar entities. E.g., all employees.
  - All entities in an entity set have the same set of attributes
  - Each entity set has a *key (underlined)*.
  - Each attribute has a *domain*.

ER Model: relationship

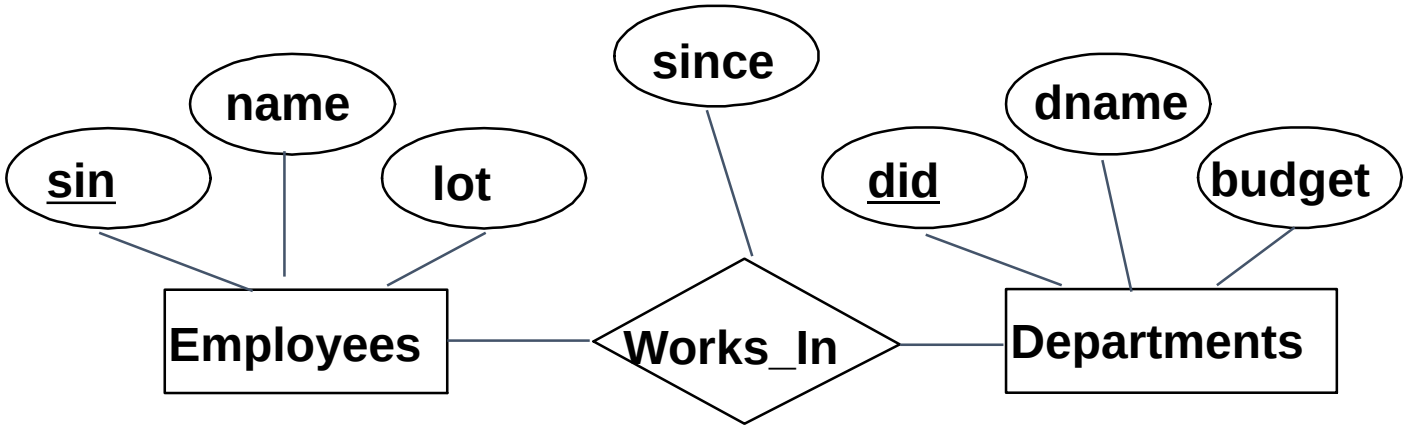

- *Relationship*: Association among two or more entities. E.g., John works in Pharmacy department.
  - relationships can have their own attributes.
- *Relationship Set*: Collection of similar relationships.

Conceptual design with entities and relationships

- What are the *entities* and *relationships* in the enterprise?
- What information about these entities and relationships should we store in the database?
- What *integrity constraints* or *business rules* hold?
- A database `schema' in the ER Model can be represented pictorially (*ER diagrams*).
- Can easily map an ER diagram into a relational schema.

ER Model: entity role

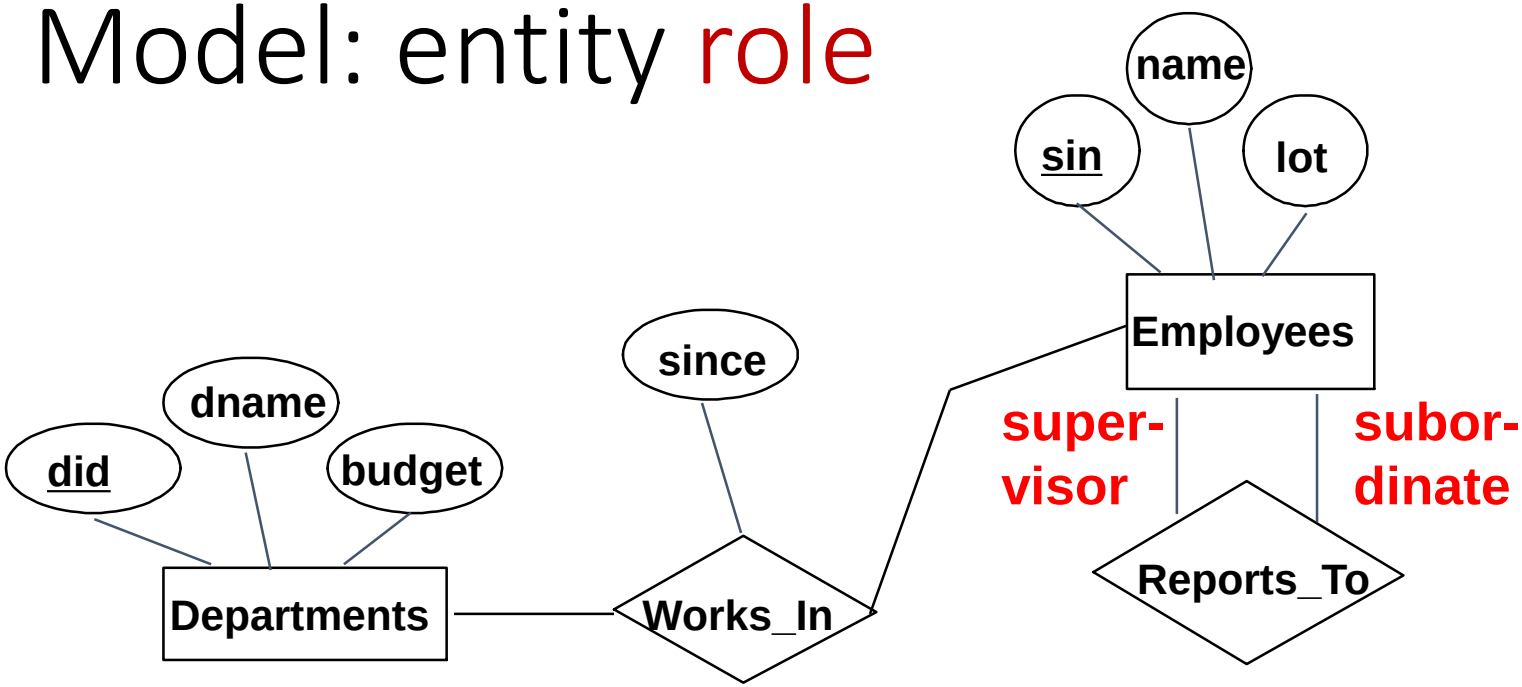

- Same entity set can participate in different relationship sets, or in different "roles" in the same set.

E/R for movies: elements

Entity – certain movie

The set of all movies constitutes an Entity set

Attributes

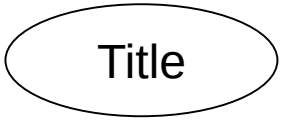

Relationships

Movies: Entity-Relationship diagram

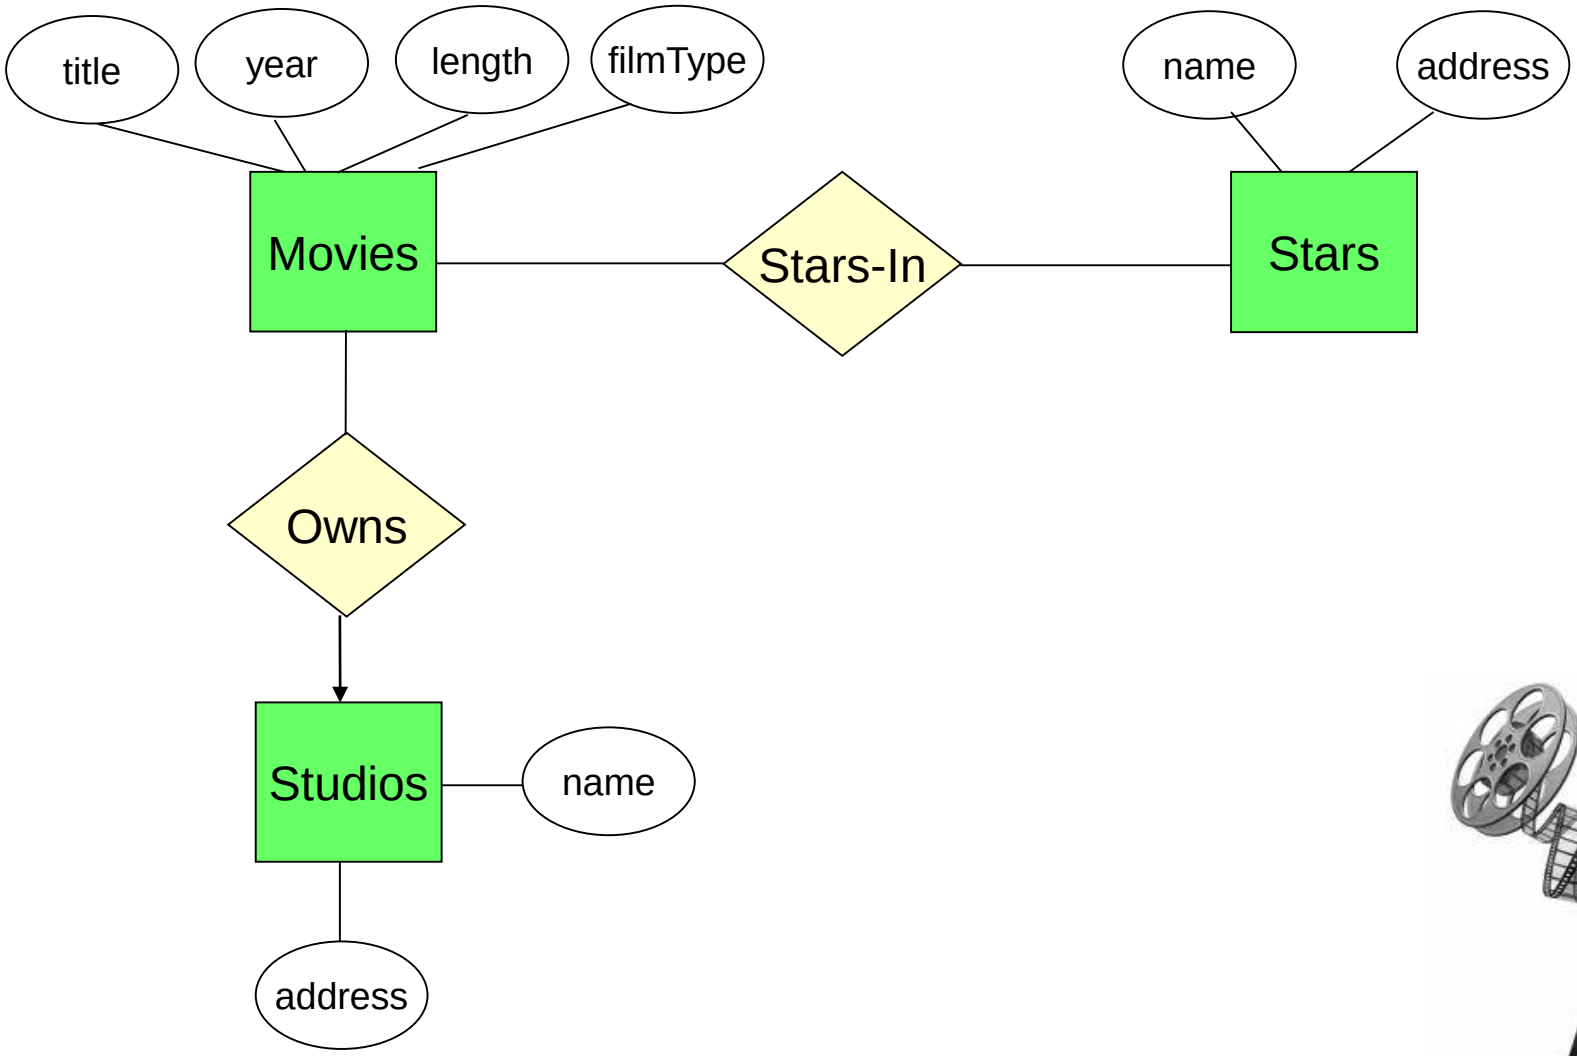

Relationships

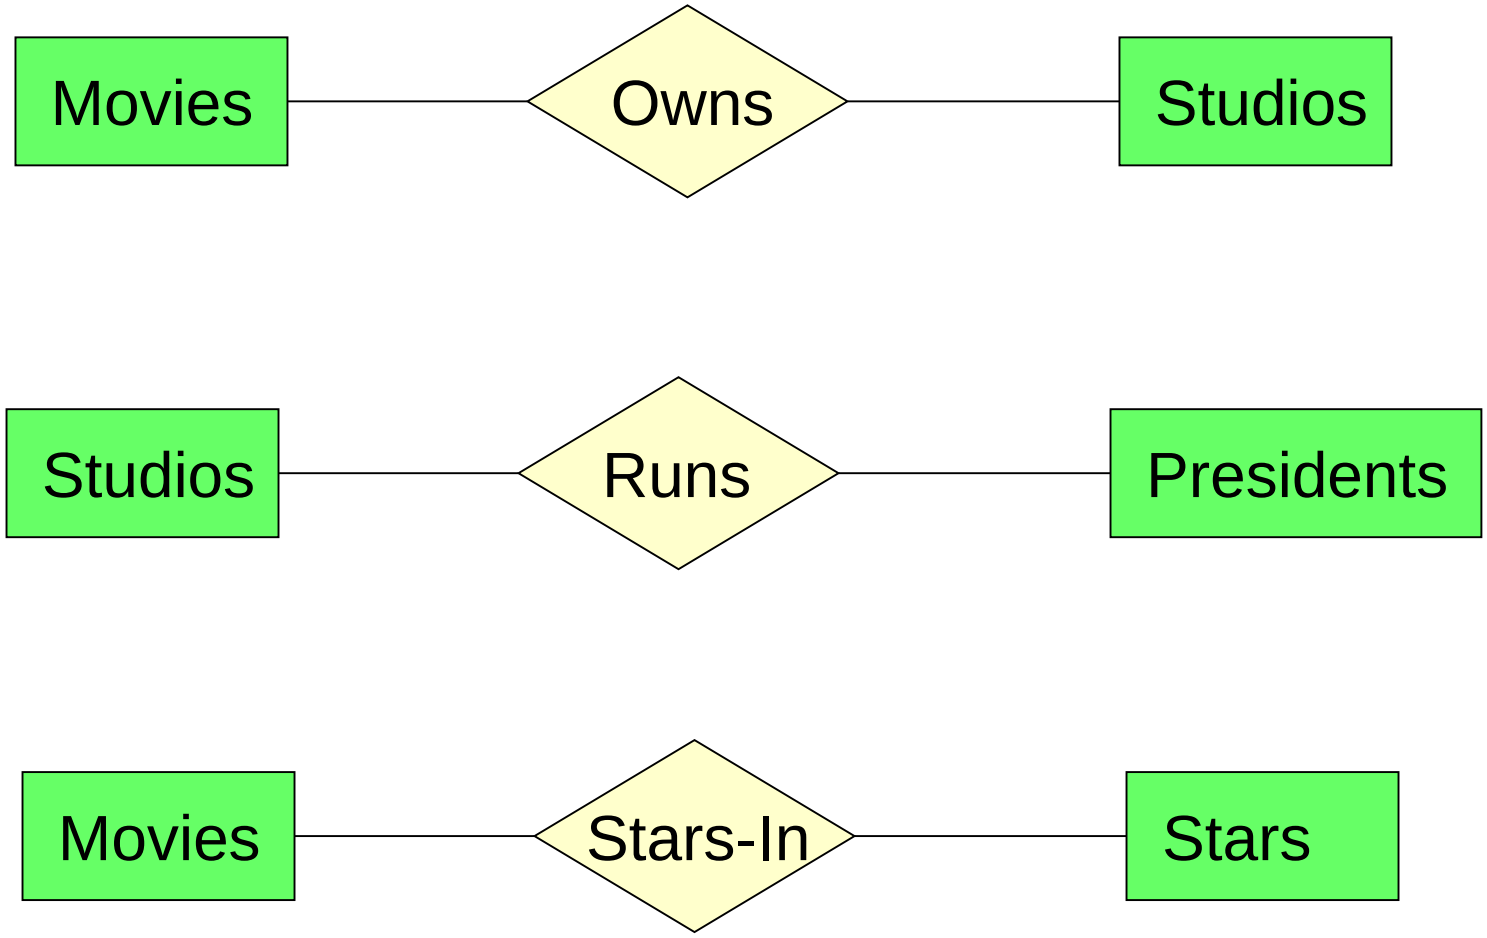

Multiplicity of Relationships

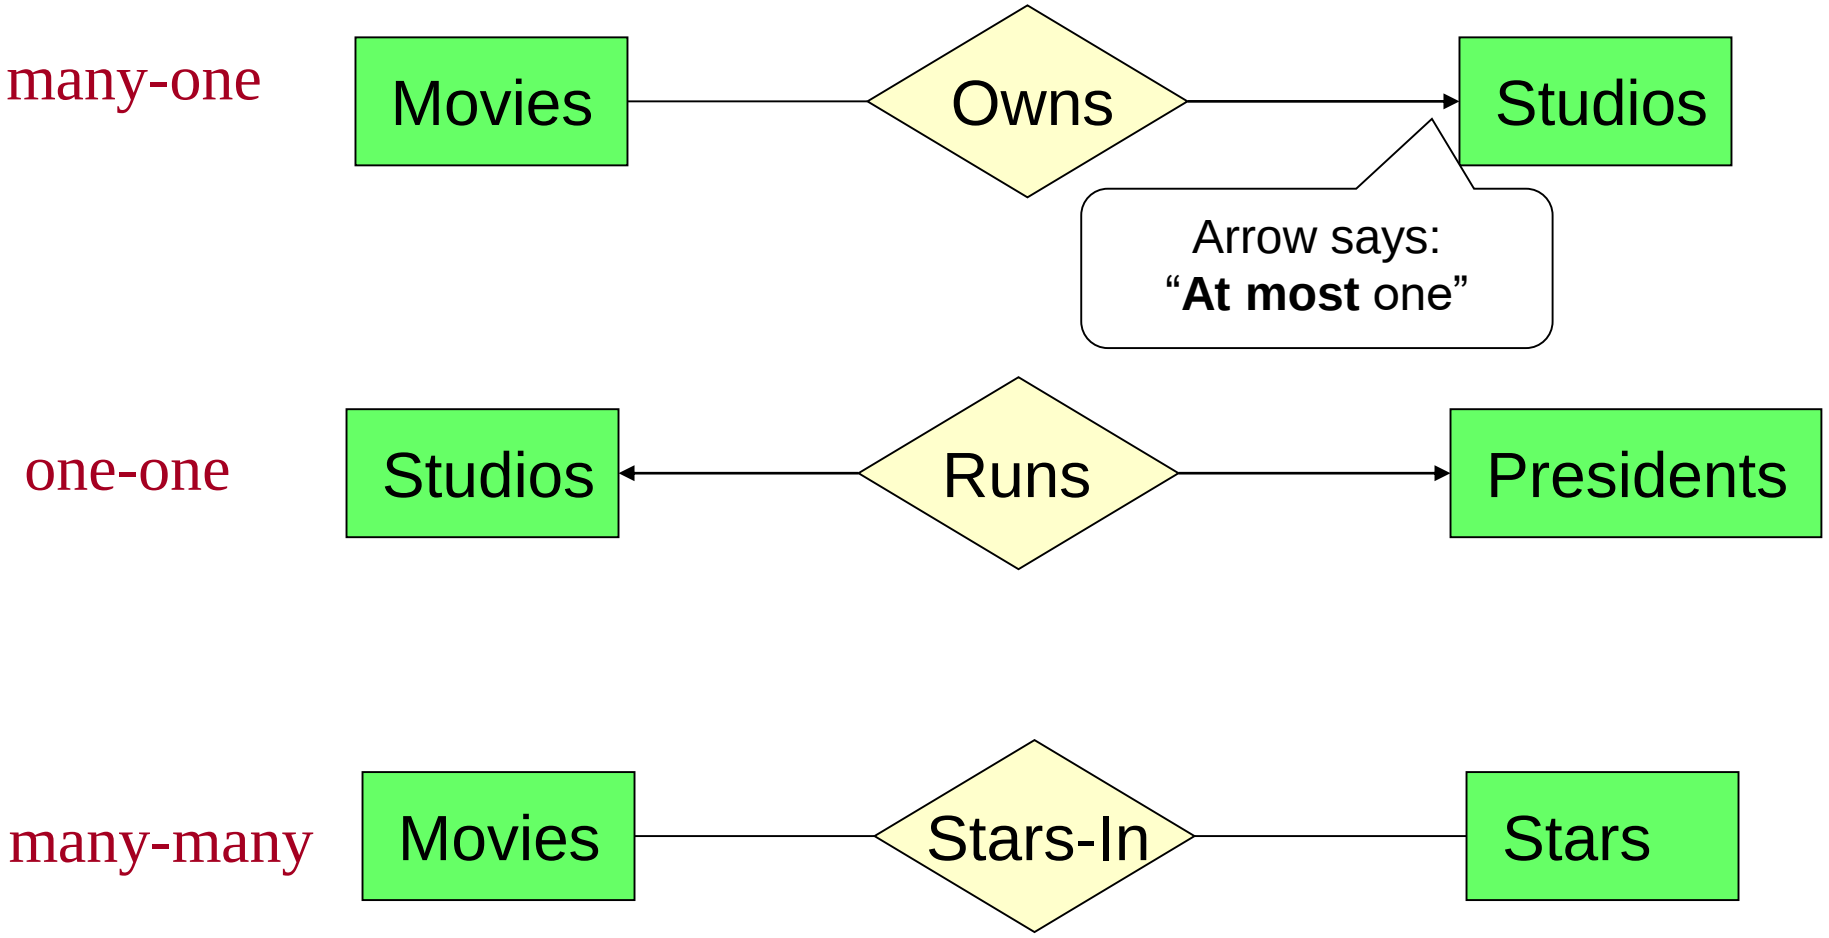

Sometimes binary relationships aren't enough!

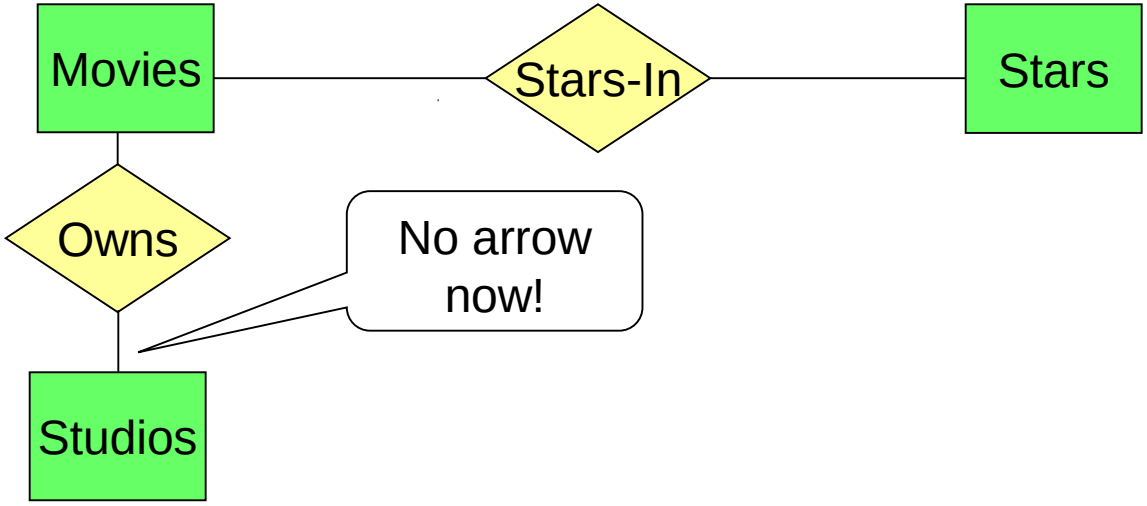

What could go wrong with this design?

Which stars a studio is paying for a given movie?

Solution: Three-way relationship

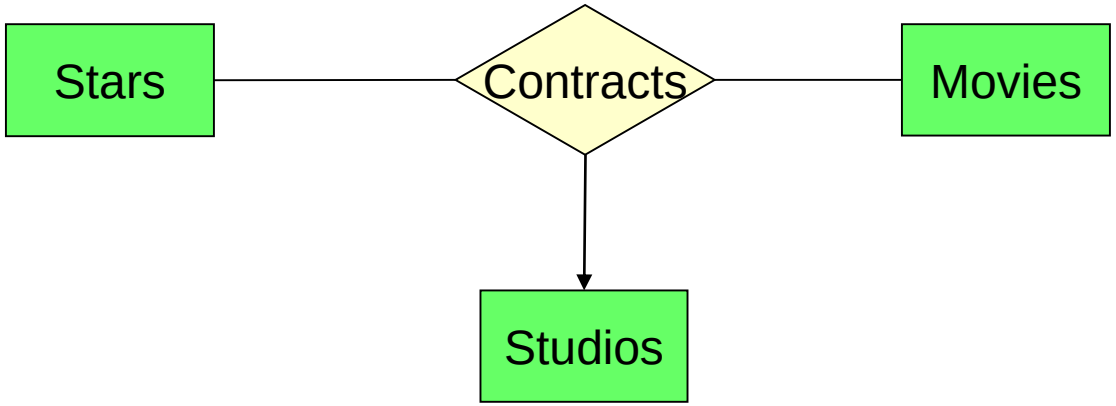

Example: Attributes on Relationships

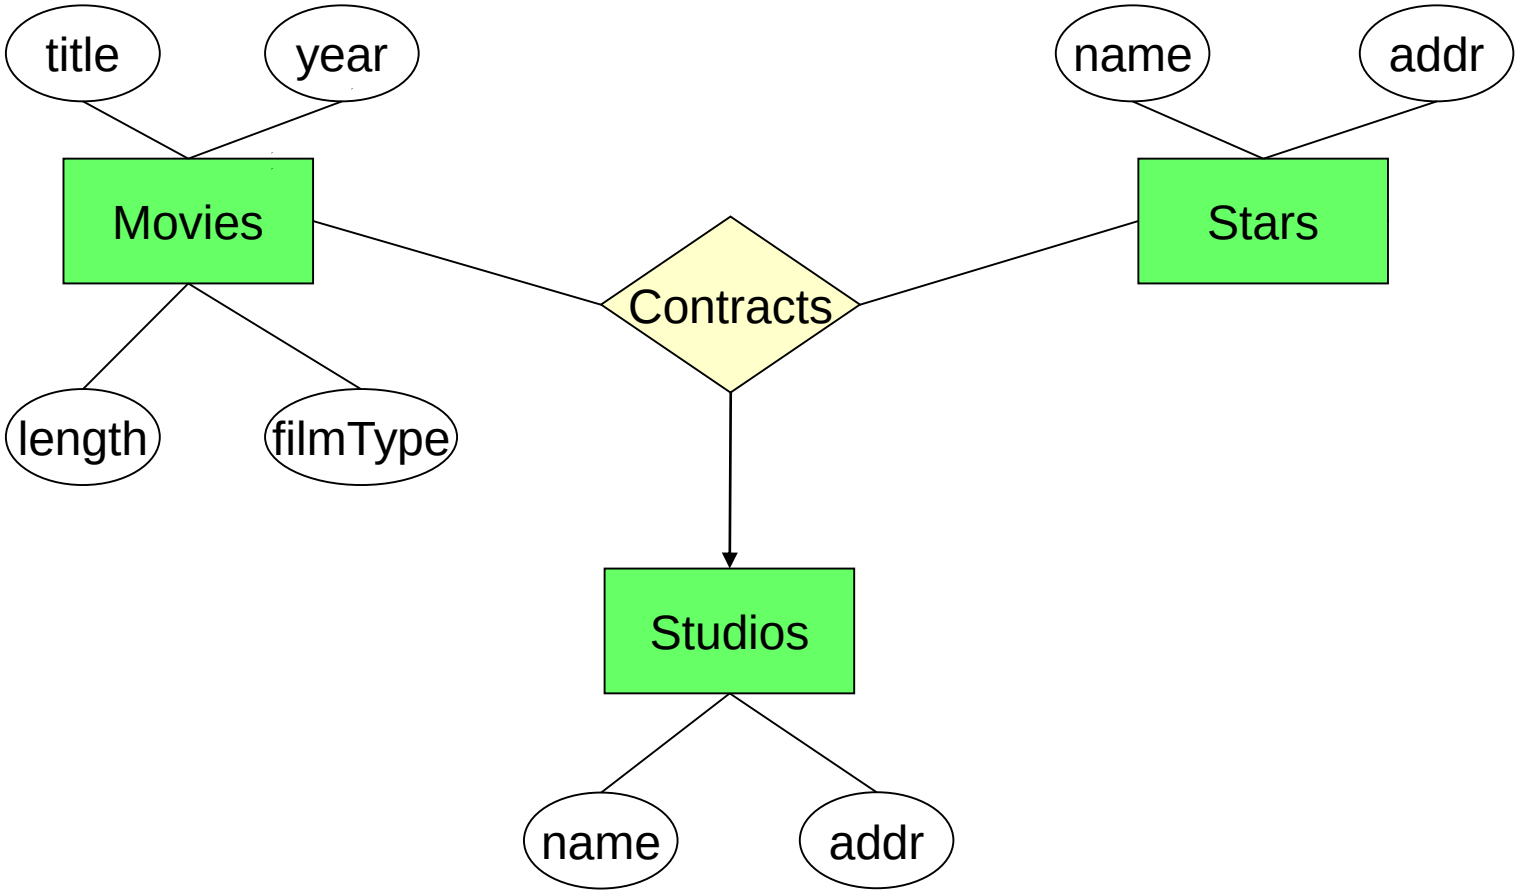

Example: Attributes on Relationships

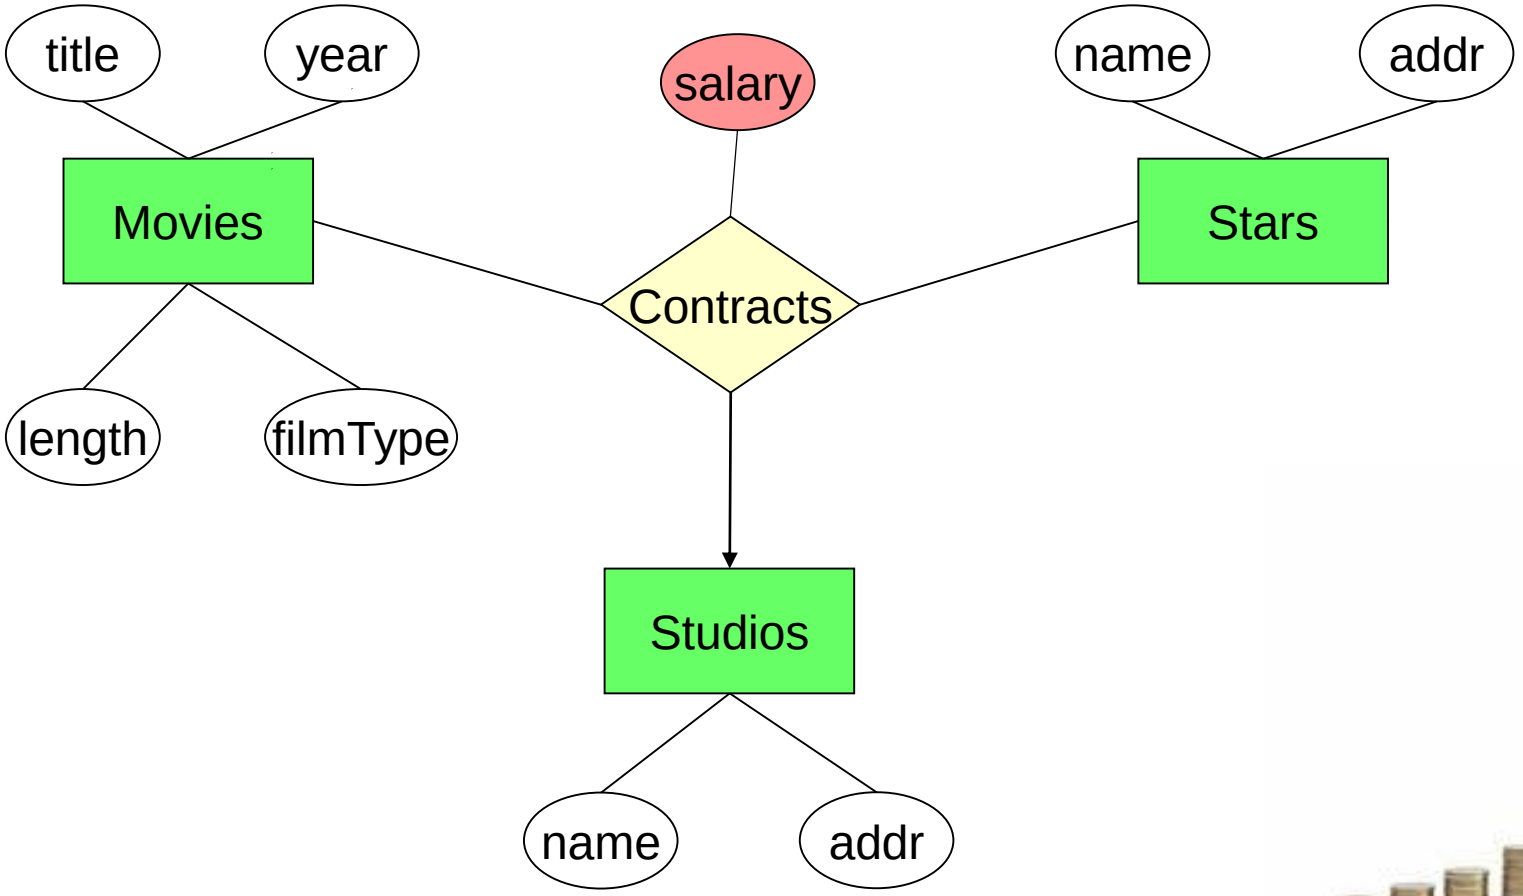

Example: Roles in a relationship

- An entity set can appear two or more times in a relationship.
- Each line to the entity set represents a different role.

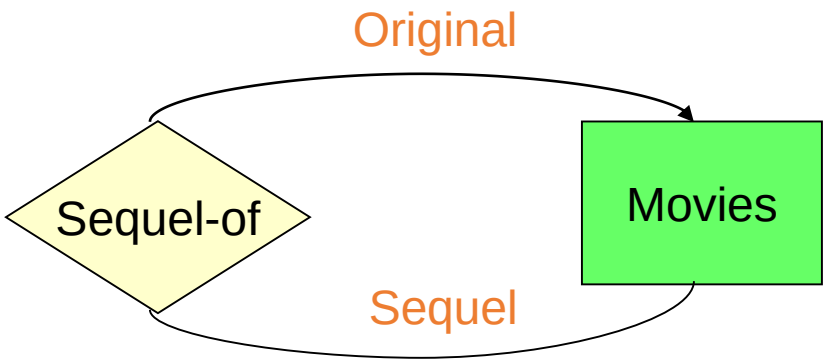

- A movie may have many sequels, but for each sequel there is at most one original movie.

Another Example

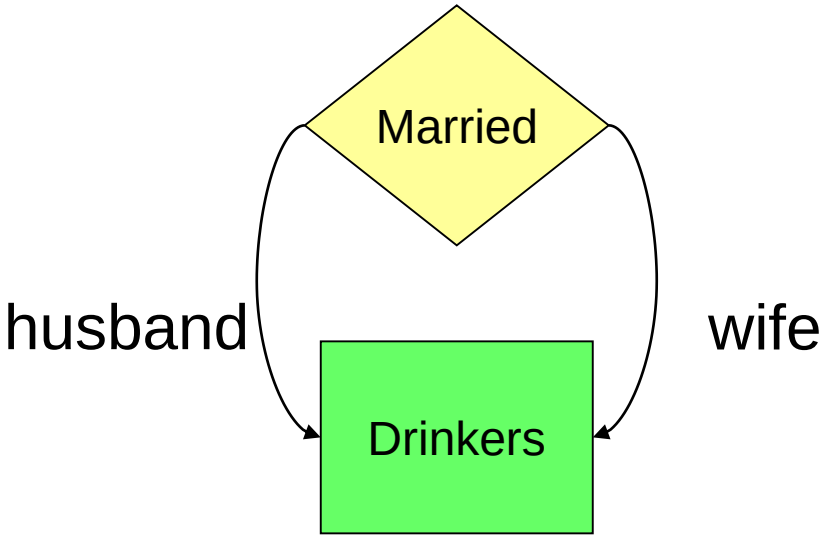

Husband	Wife
Bob	Ann
Joe	Sue
…	…

Another role

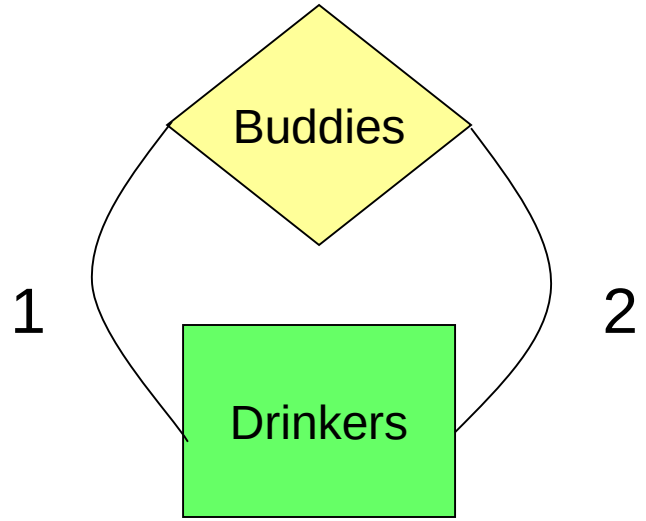

Buddy1	Buddy2
Bob	Ann
Joe	Sue
Ann	Bob
Joe	Moe
…	…

Design example 1. “Bars-Beer-Drinkers” (BBD)

- Bars sell some beers.
- Drinkers like some beers.
- Drinkers frequent some bars.

What would be the E/R diagram?

BBD ER diagram

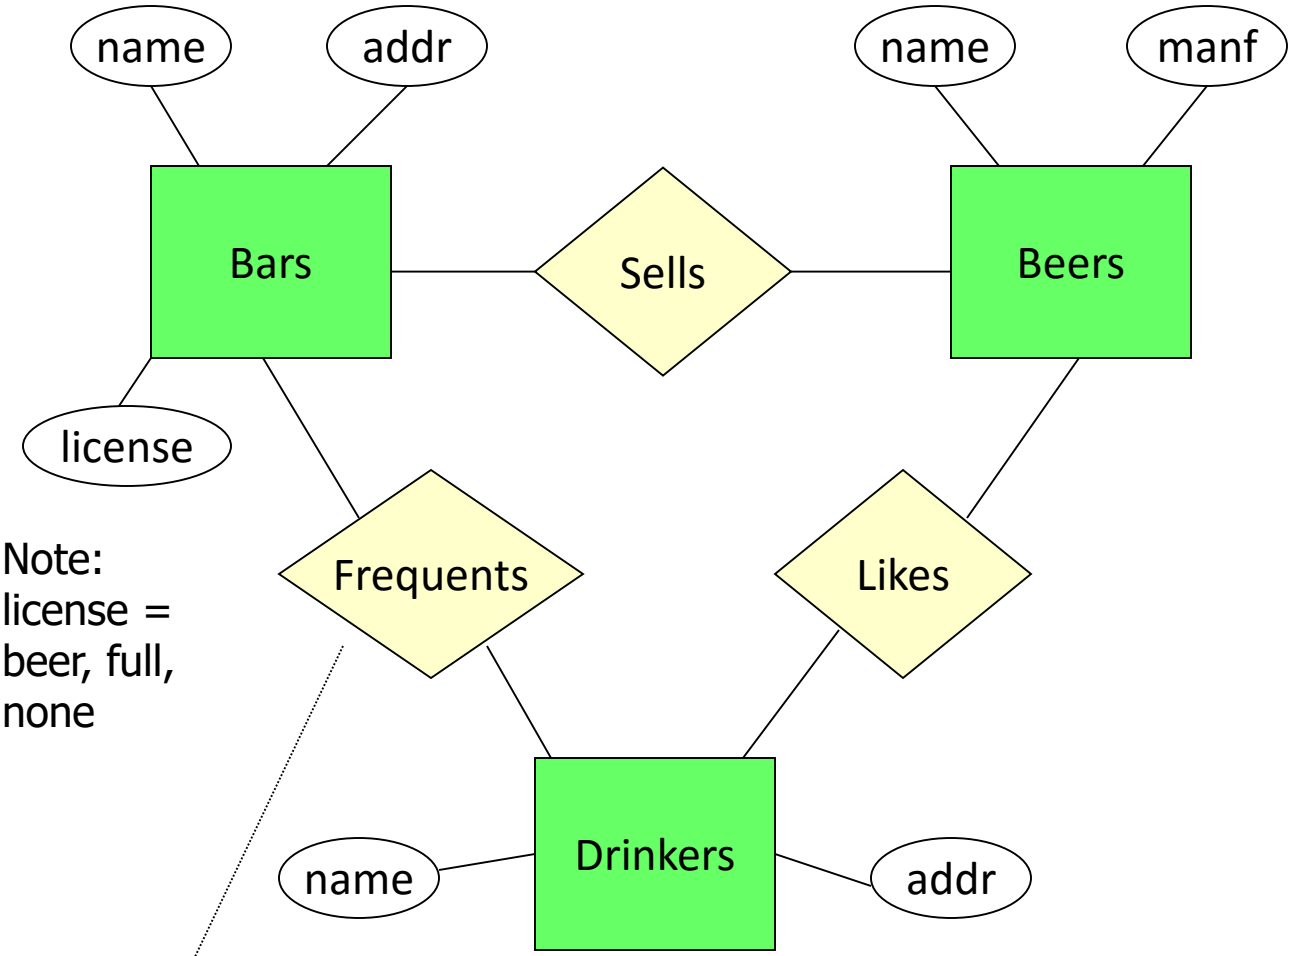

Note:
license =
beer, full,
none

Why we need it?

Bars sell some beers.

Drinkers like some beers.

Drinkers frequent some bars.

BBD Multiway Relationship

- Suppose that drinkers prefer drink certain beers at certain bars (many bear types at multiple bars).

- How do we reflect their preferences in ER diagram?

BBD Multi-way Relationship

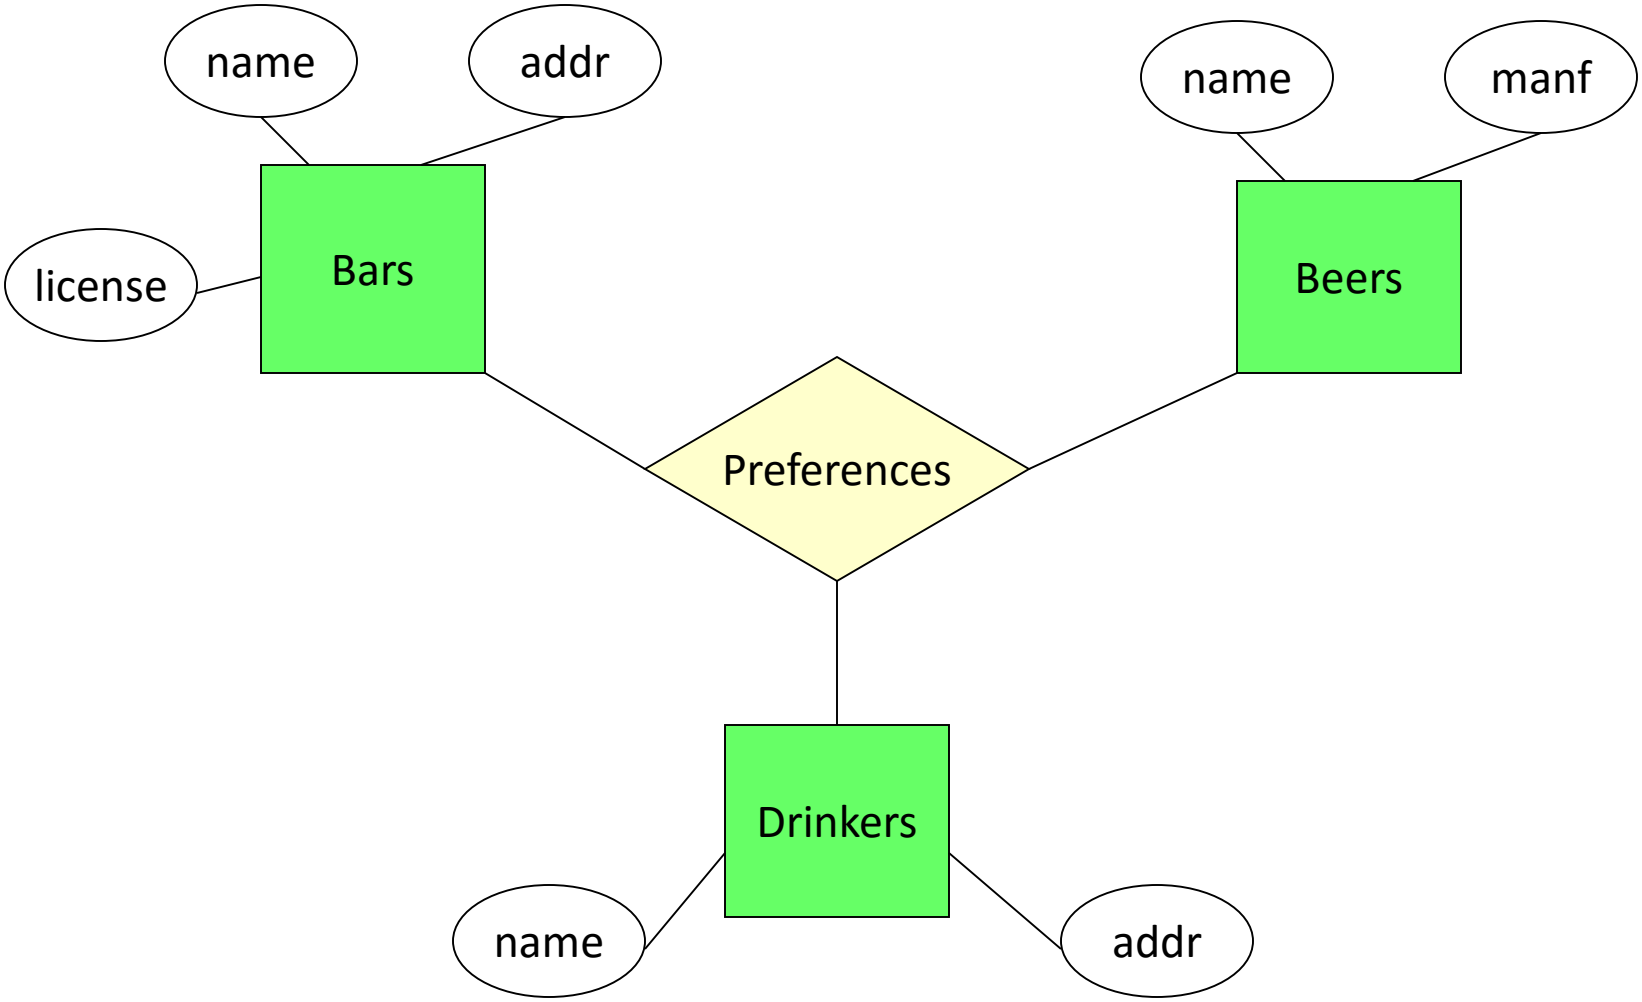

A typical relationship table for ternary (three-way) relationships

Bar	Drinker	Beer
Joe's Bar	Ann	Miller
Sue's Bar	Ann	Bud
Sue's Bar	Ann	Pete's Ale
Joe's Bar	Bob	Bud
Joe's Bar	Bob	Miller
Joe's Bar	Cal	Miller
Sue's Bar	Cal	Bud Lite

Multiple relationships between two entity sets

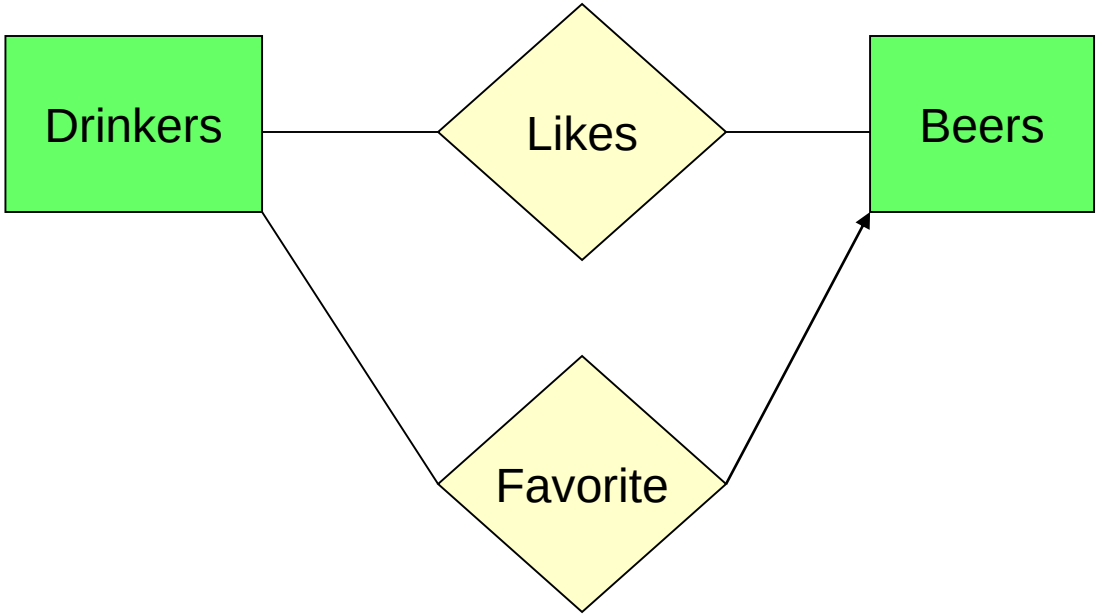

“Exactly one” Multiplicity

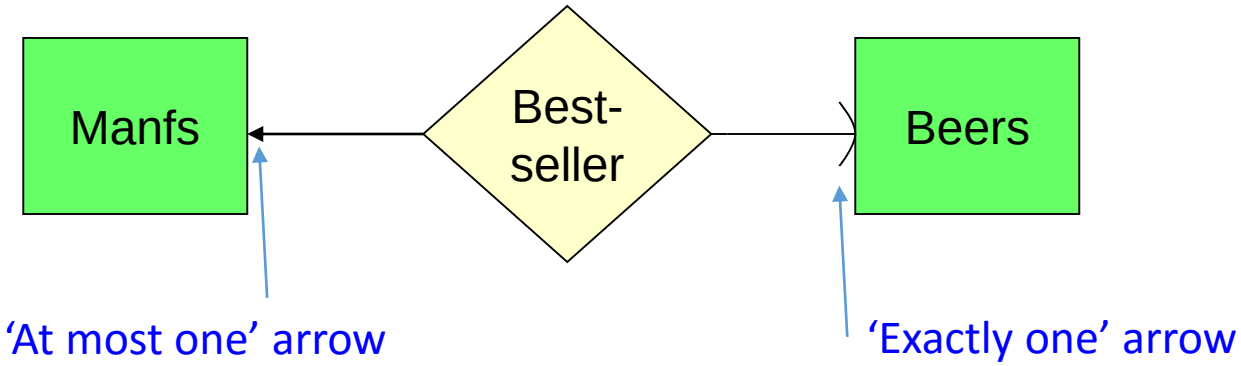

- Some beers are not the best-seller of any manufacturer, so a **rounded arrow** to *Manfs* would be inappropriate.
- But a manufacturer has to have a best-seller (in our model)

Design choices

- Should a concept be modeled as an entity or an attribute?
- Should a concept be modeled as an entity or a relationship?
- Identifying relationships: binary or ternary?

Entity vs. Attribute

- Should *address* be an attribute of Employees or an entity (related to Employees)?
- Depends upon how we want to use address information, and the semantics of the data:
    - If we have several addresses per employee, *address* must be an entity (since attributes cannot be set-valued).
    - If the structure (city, street, etc.) is important, *address* must be modeled as an entity (since attribute values are atomic).
    - If the lifetime of the address differs from the entity, *address* must be modeled as an entity (since attributes are deleted with their entity).

Exercise 1. Bank database

- Let us design a database for a bank, including information about customers and their accounts.

Information about a customer includes their name, address, phone, and SIN number. Accounts have numbers, types (e.g., savings, checking) and balances. We also need to record the customer(s) who own an account. Draw the E/R diagram for this database.

Exercise 1 solution

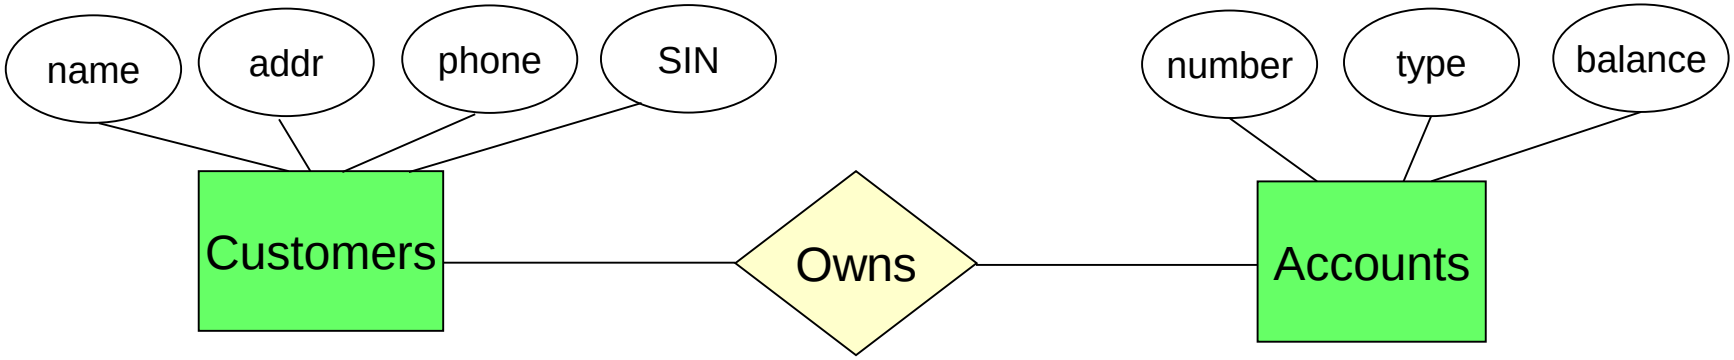

Exercise 1A. Bank database

- Modify your solution as follows:
  a) **Change your diagram so an account can belong to only one customer.**

Exercise 1A solution

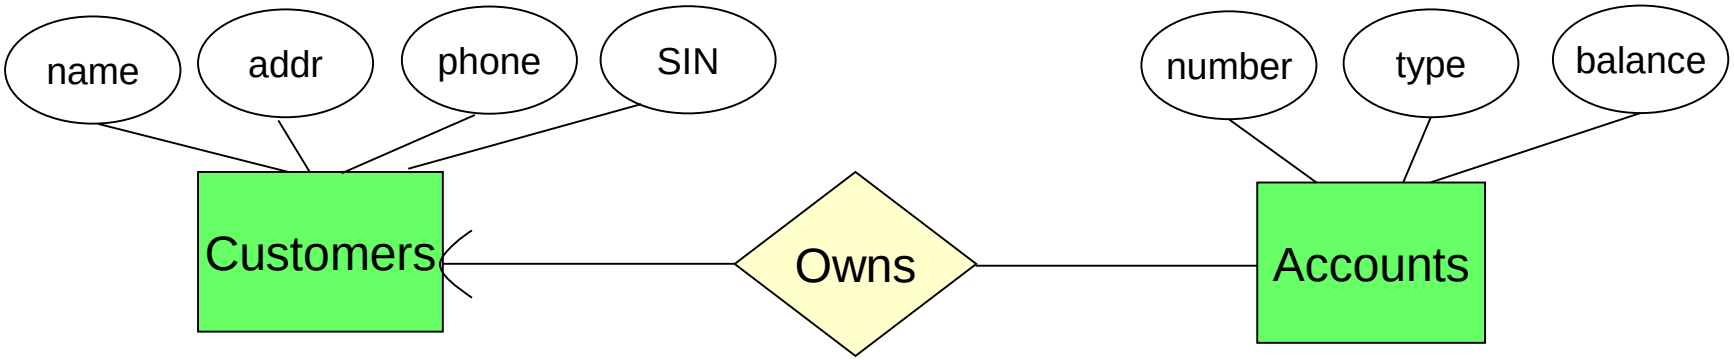

Exercise 1B. Bank database

- Modify your solution as follows:

  a) Change your diagram so an account can have only one customer.

  b) Change your diagram so that a customer can have a **set** of addresses (which are street-city-province triples) and a **set** of phones.

  Remember that we do not allow attributes to have non-atomic types, such as sets, in the E/R model.

Exercise 1B solution

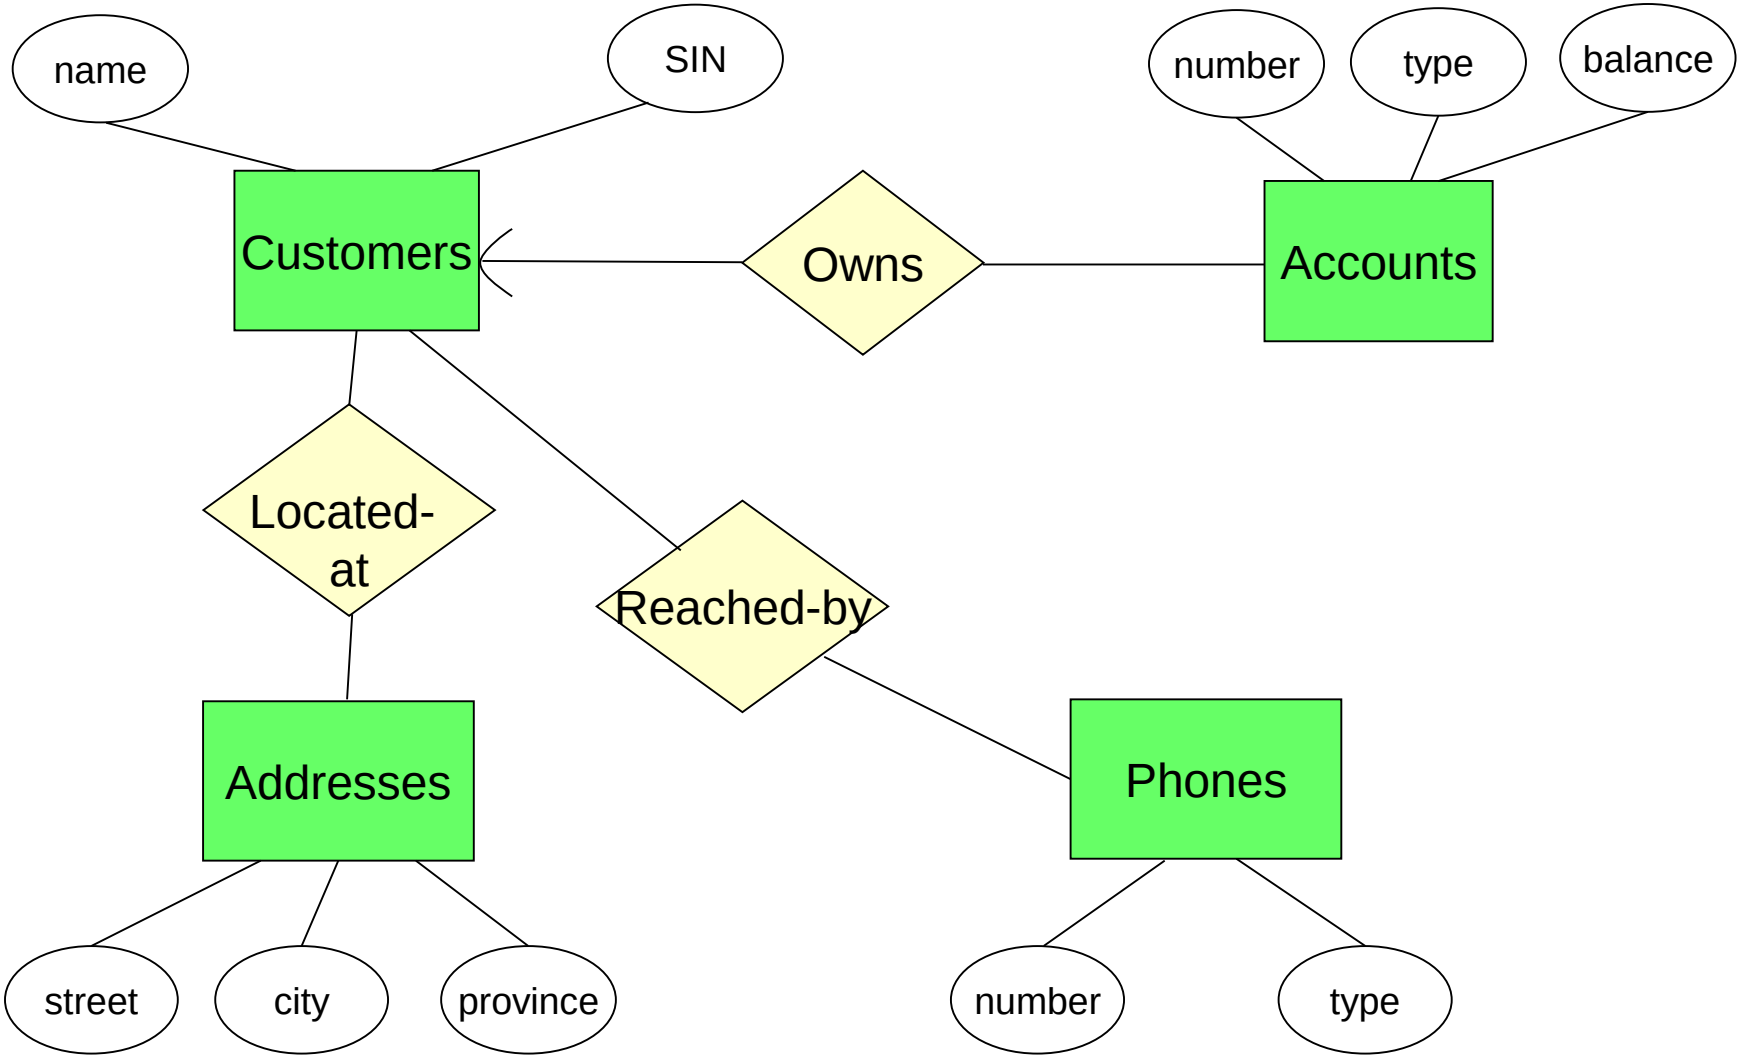

Exercise 1C. Bank database

- Modify your solution as follows:

  a) Change your diagram so an account can have only one customer.

  b) Change your diagram so that a customer can have a set of addresses (which are street-city-province triples) and a set of phones.

  c) Further modify your diagram so that customers can have a set of addresses, and **at each address there is a set of phones**.

Exercise 1C solution

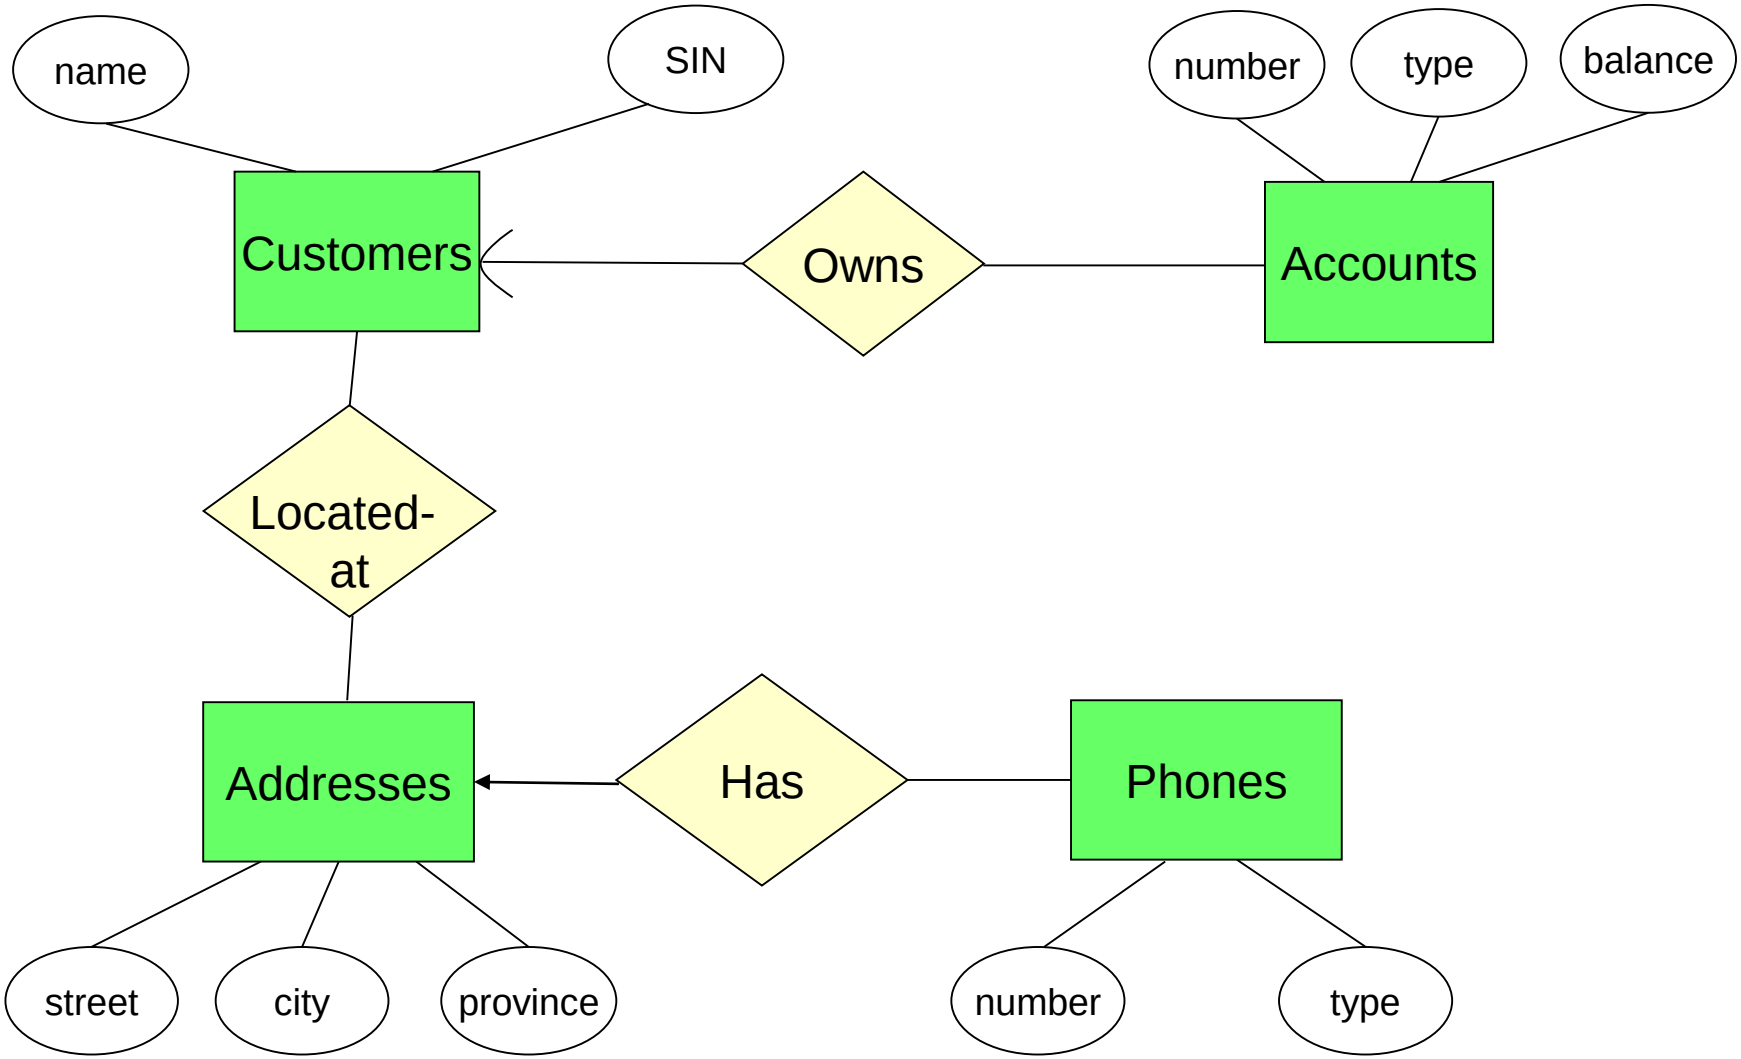

Exercise 2. Teams and fans database

- Give an E/R diagram for a database recording information about teams, players, and their fans, including:
  - For each team, its name, its players, its team captain (one of its players), and the colors of its uniform.
  - For each player, his/her name.
  - For each fan, his/her name, favorite teams, favorite players, and favorite color.

Exercise 2 solution (Variant I)

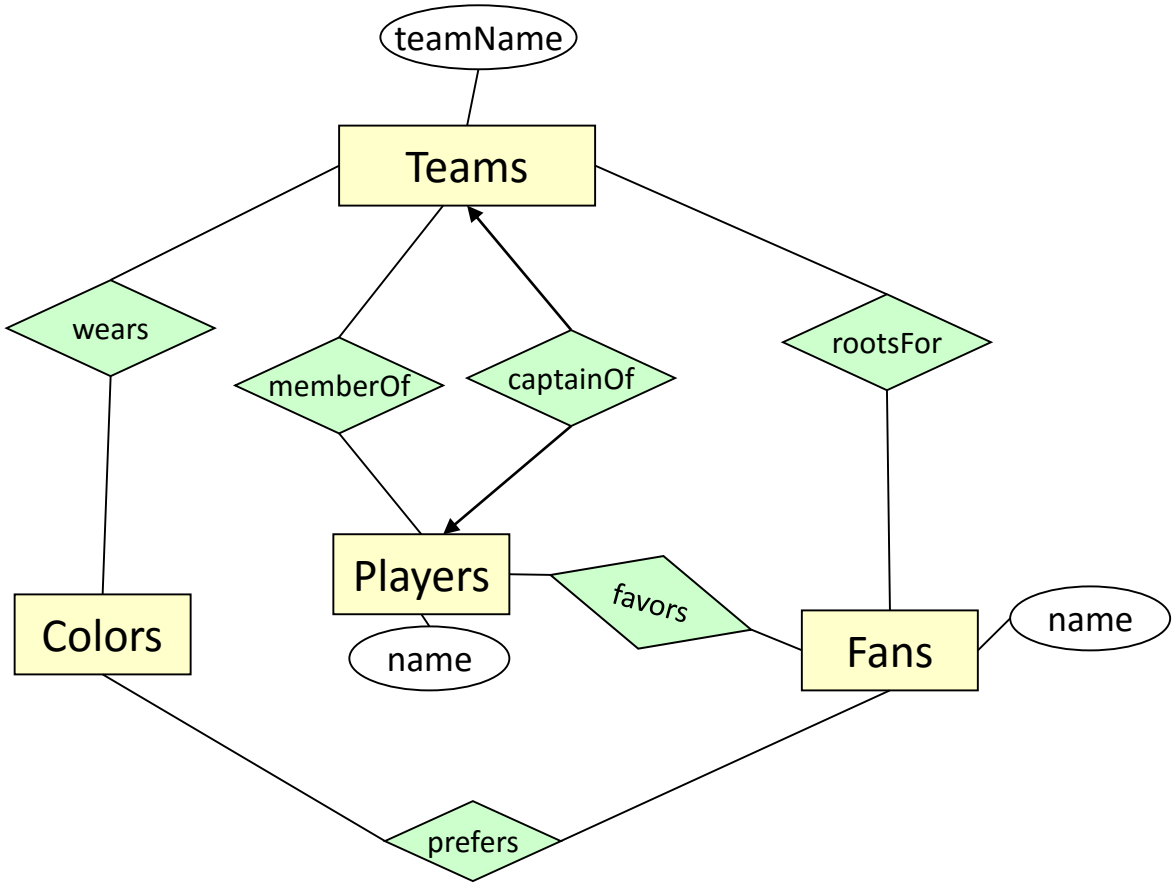

Exercise 2 solution (Variant II)

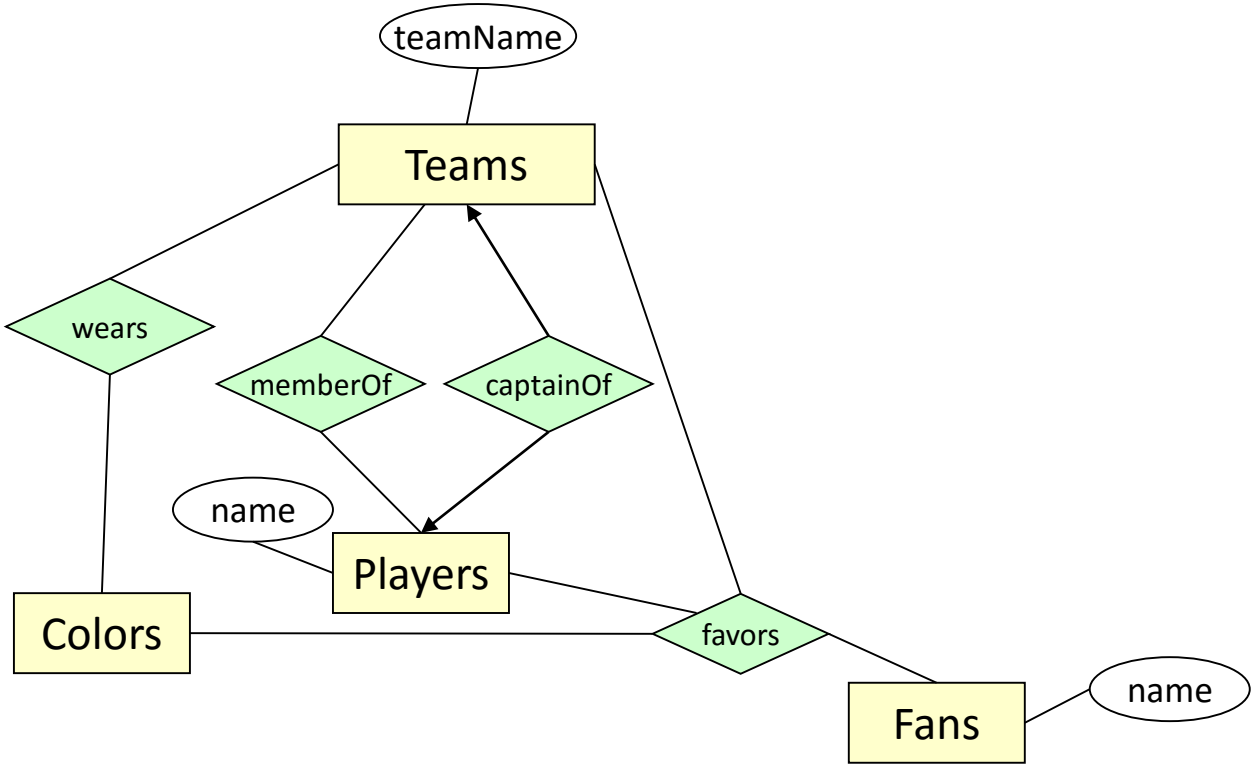

Exercise 2A. Teams and fans database

- Modification A:
  - Suppose we wish to add to the schema a relationship "Led-by" among two players and a team. The intention is that this relationship set consists of triples

    (player1, player2, team)

    such that player 1 played on the team at a time when some

    other player 2 was the team captain.

  Draw the modification to the E/R diagram.

Exercise 2A solution