SCC.201
Database Management Systems

2023 - Week 1 – ER diagrams.

Uraz C Turker & Ricki Boswell

What will you learn today?

- General knowledge of Data, Databases, and Database Management Systems.

- Introduction to database design: Entity-Relation Diagrams.

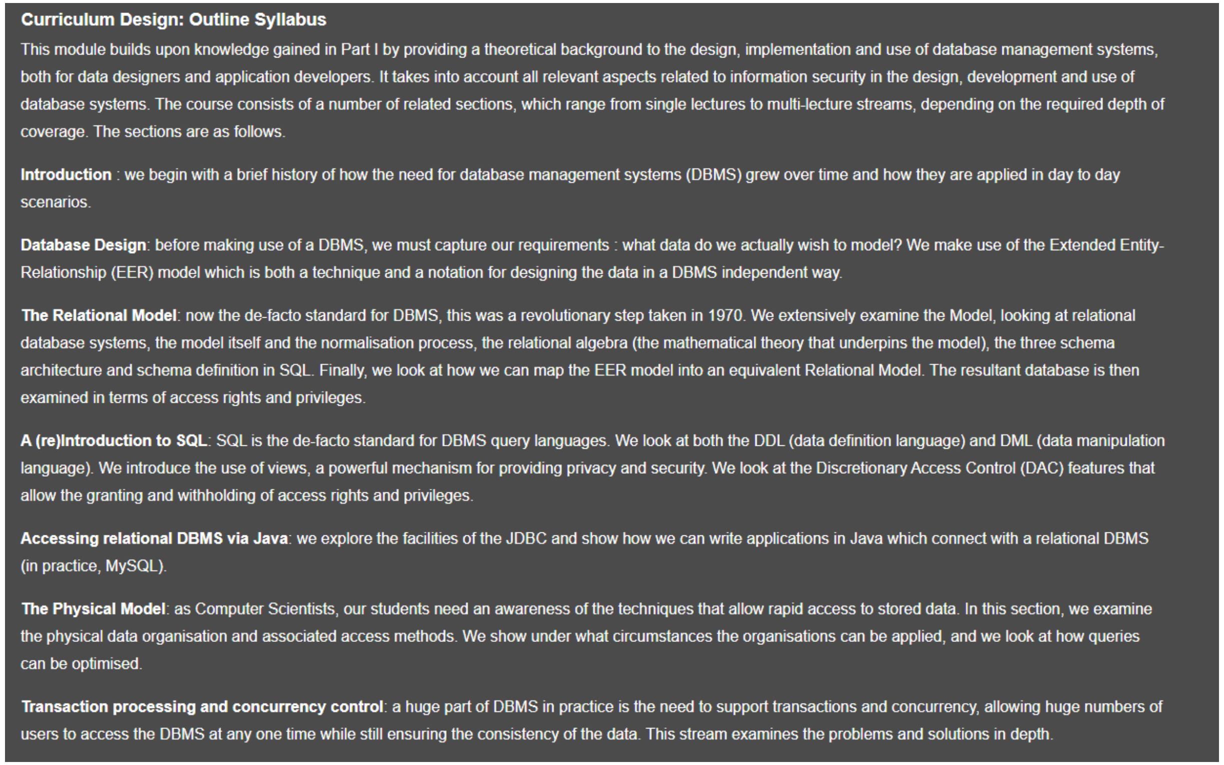

Curriculum Design: Outline Syllabus

This module builds upon knowledge gained in Part I by providing a theoretical background to the design, implementation and use of database management systems, both for data designers and application developers. It takes into account all relevant aspects related to information security in the design, development and use of database systems. The course consists of a number of related sections, which range from single lectures to multi-lecture streams, depending on the required depth of coverage. The sections are as follows.

Introduction : we begin with a brief history of how the need for database management systems (DBMS) grew over time and how they are applied in day to day scenarios.

Database Design: before making use of a DBMS, we must capture our requirements : what data do we actually wish to model? We make use of the Extended Entity-Relationship (EER) model which is both a technique and a notation for designing the data in a DBMS independent way.

The Relational Model: now the de-facto standard for DBMS, this was a revolutionary step taken in 1970. We extensively examine the Model, looking at relational database systems, the model itself and the normalisation process, the relational algebra (the mathematical theory that underpins the model), the three schema architecture and schema definition in SQL. Finally, we look at how we can map the EER model into an equivalent Relational Model. The resultant database is then examined in terms of access rights and privileges.

A (re)Introduction to SQL: SQL is the de-facto standard for DBMS query languages. We look at both the DDL (data definition language) and DML (data manipulation language). We introduce the use of views, a powerful mechanism for providing privacy and security. We look at the Discretionary Access Control (DAC) features that allow the granting and withholding of access rights and privileges.

Accessing relational DBMS via Java: we explore the facilities of the JDBC and show how we can write applications in Java which connect with a relational DBMS (in practice, MySQL).

The Physical Model: as Computer Scientists, our students need an awareness of the techniques that allow rapid access to stored data. In this section, we examine the physical data organisation and associated access methods. We show under what circumstances the organisations can be applied, and we look at how queries can be optimised.

Transaction processing and concurrency control: a huge part of DBMS in practice is the need to support transactions and concurrency, allowing huge numbers of users to access the DBMS at any one time while still ensuring the consistency of the data. This stream examines the problems and solutions in depth.

What is a Data, Database, Database Management System (DBMS)?

What is data?

What is a Data, Database, Database Management System (DBMS)?

Data is anything that can be represented using Binary notation.

We can then encode data in File:

0	1	0	0	1	1	1	0
0	1	0	0	1	1	1	1
0	0	1	1	1	0	1	0

N
O
:

Source of Data

- Data is being created whenever we:
  - Eat,
  - Travel,
  - Chat,
  - Talk over a line on a phone, etc.

A source of data: Games

The aim of having data?

- is to retrieve information.
- To achieve this we need to process this data.
- What is the fundamental step to process data?
- To establish a flexible, understandable, and a common representation for the data.

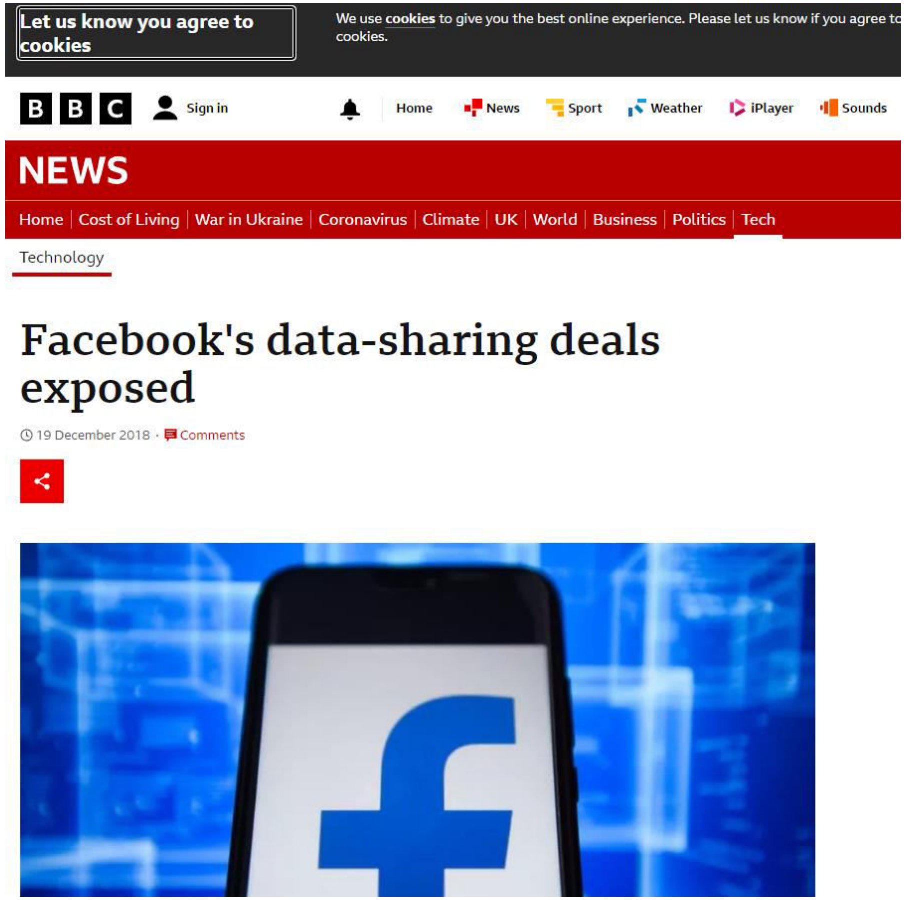

Let us know you agree to cookies

We use cookies to give you the best online experience. Please let us know if you agree t cookies.

BBC
Sign in
Home
News
Sport
Weather
iPlayer
Sounds
NEWS
Home | Cost of Living | War in Ukraine | Coronavirus | Climate | UK | World | Business | Politics | Tech
Technology

Facebook's data-sharing deals exposed

19 December 2018 · Comments

How to represent data?

- Files? File systems?
  - Hard to manage, code, and maintain.
  - A straightforward mechanism to store data.
- Consider the following query:

  "What is the expected monthly net income for the next financial year?"
- What would happen if the data is altered at the same time?
- How to guarantee the uniqueness of the data?

What is a Data, Database, Database Management System (DBMS)?

Database: It represents a massive amount of persistent data according to a logical model

The logical model can be represented graphically by.

UML or Entity Relationship Diagrams.

The logical model allows us to reason about the logical representation of the data to gain the flexibility to establish an excellent data representation power to aid processing.

What is a Data, Database, Database Management System (DBMS)?

DBMS: Provides efficient, reliable, convenient, safe, multi-user friendly storage and access to massive amounts of persistent data.

What is a Data, Database, Database Management System (DBMS)?

DBMS: Provides efficient, reliable, convenient, safe, multi-user friendly storage and access to massive amounts of persistent data.

This is why they are very popular.

DBMS in a Computer Game.

https://www.csgodatabase.com/about/

Key concepts

- Data model
  - Schema, XML, Graph
- Schema vs Data
- Data Definition Language (DDL)
- Data Manipulation Language (DML)

Key people

- DBMS Implementer.
- DB Designer.
- DB Application Developer.
- DB Admin.

The logical model: Entity relationship model.

A graphical representation of the logical view of a Database.

Communication method (a single image worth 1000s of words!)

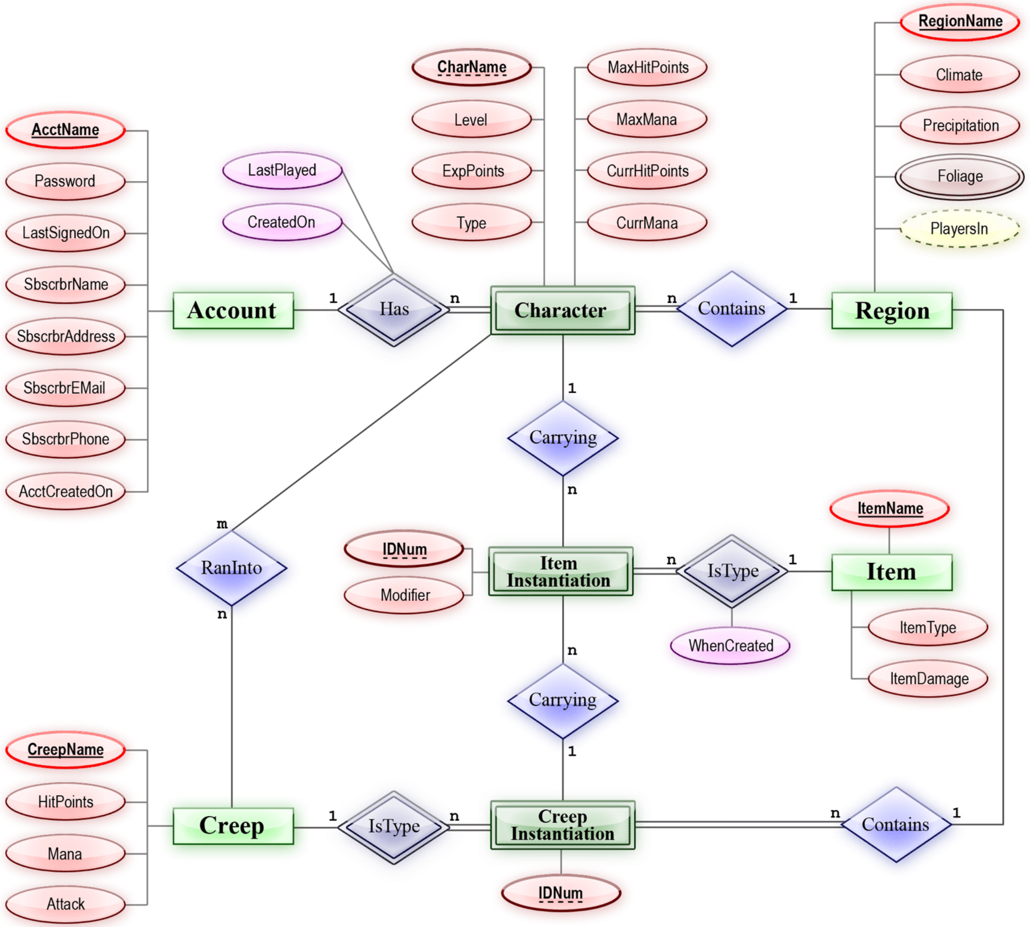

AcctName
Password
LastSignedOn
SbscrbrName
SbscrbrAddress
SbscrbrEMail
SbscrbrPhone
AcctCreatedOn
Account
1
Has
n
LastPlayed
CreatedOn
CharName
Level
ExpPoints
Type
MaxHitPoints
MaxMana
CurrHitPoints
CurrMana
Character
n
Contains
1
Region
RegionName
Climate
Precipitation
Foliage
PlayersIn
1
Carrying
n
m
RanInto
n
IDNum
Modifier
Item Instantiation
n
IsType
1
Item
ItemName
WhenCreated
ItemType
ItemDamage
n
Carrying
1
CreepName
HitPoints
Mana
Attack
Creep
1
IsType
n
Creep Instantiation
IDNum
n
Contains
1

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable attributes.
- An entity must have attributes.
- Pilot -> Name, Surname, Hours_Flight, Height, Weight.
- Car -> Model, Weight, Length, Max_Speed.

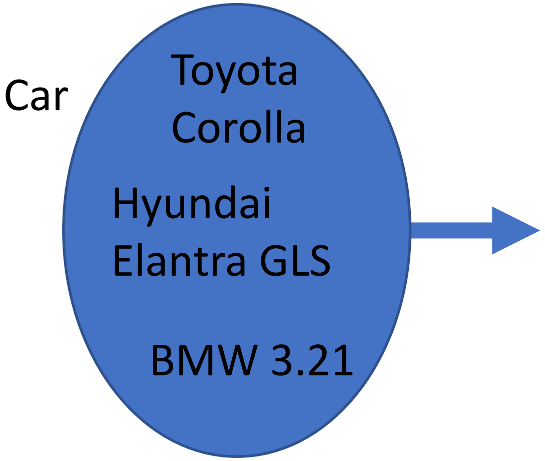

Entity set: A conceptual representation of entities w.r.t the same set of attributes is called an Entity Set.

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.
- Rectangular shape is used for entities.

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.
- Rectangular shape is used for entities.

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.
- Rectangular shape is used for entities.

- Oval shape is used for attributes.

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.
- Rectangular shape is used for entities.

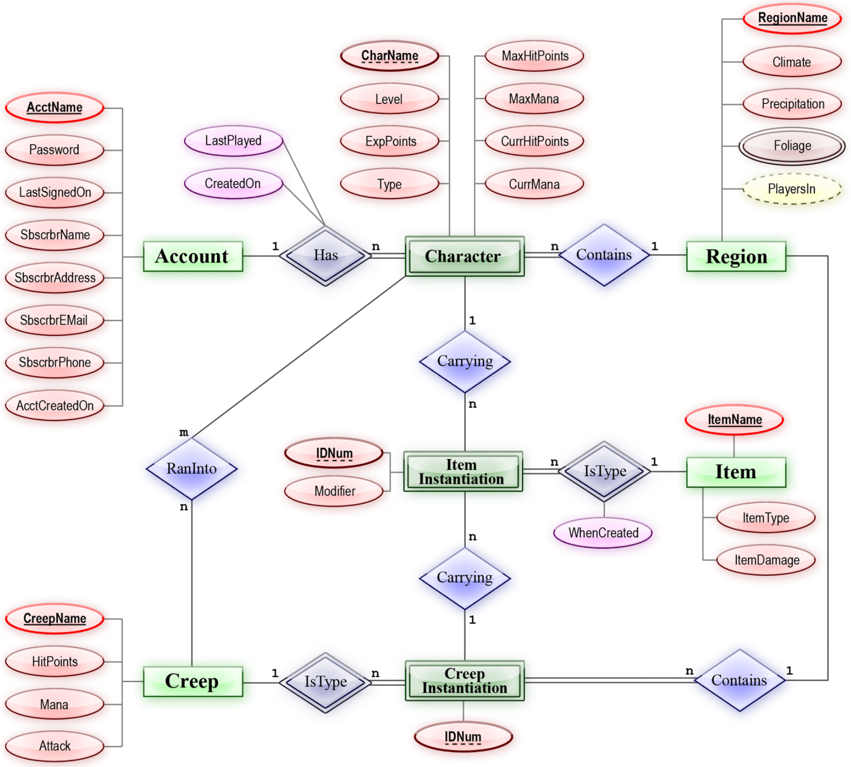

- Oval shape is used for attributes.

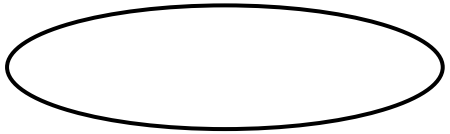

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.

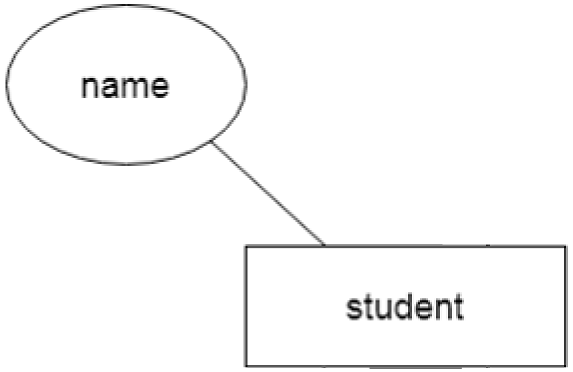

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.

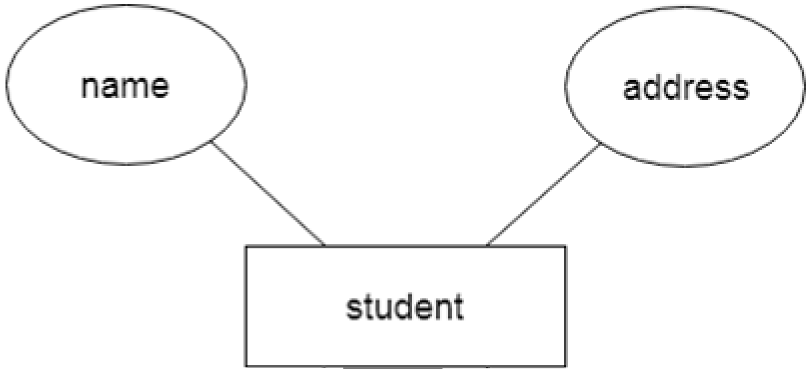

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.

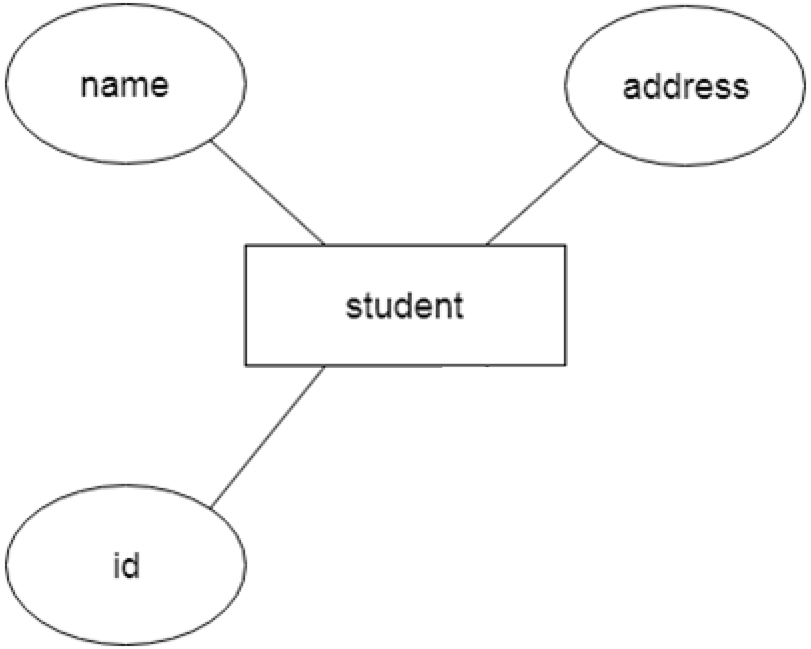

Key concepts of ER: Entity

- **Entity:** An object that has distinguishable properties.
- An entity must have attributes.

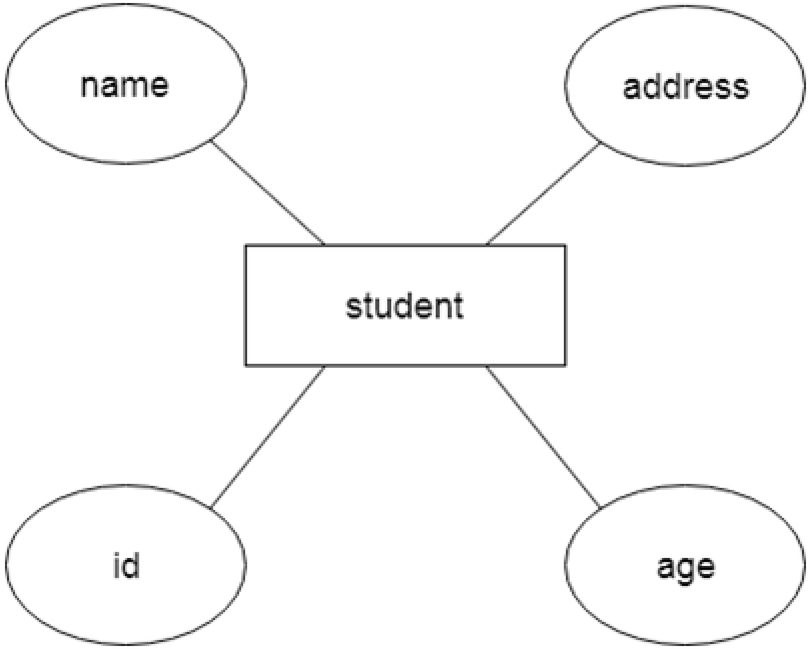

Key concepts of ER: Entity

Car -> Model, Weight, Length, Max_Speed.

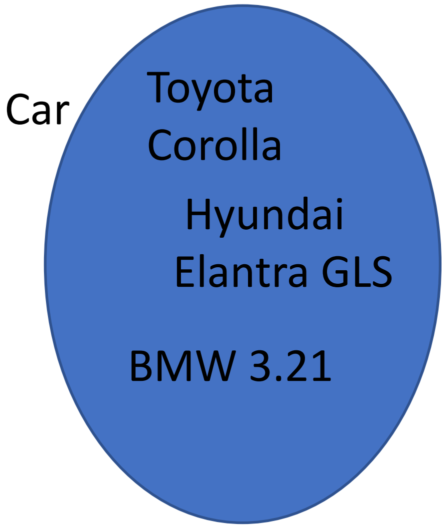

Car

Toyota Corolla

Hyundai Elantra GLS

BMW 3.21

Model	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.16	200
Hyundai E.GLS	1400	3.16	210

Which attribute distinguishes an entity (the whole row) within the entity set Car?

Brand uniquely identifies an entity (the whole row) within this entity set.

Key concepts of ER: Entity

- Car -> Model, Weight, Length, Max_Speed.

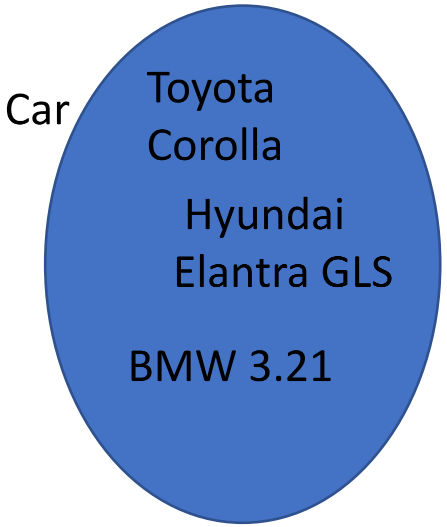

Car

Toyota Corolla

Hyundai Elantra GLS

BMW 3.21

Model	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.16	200
Hyundai E.GLS	1400	3.16	210

Which attribute distinguishes an entity (the whole row) within the entity set Car?

Model uniquely identifies an entity (the whole row) within this entity set.

The attribute that uniquely identifies an entity within this entity set is called a **‘key attribute’**.

Key concepts of ER: Entity

- Car -> Model, Weight, Length, Max_Speed.

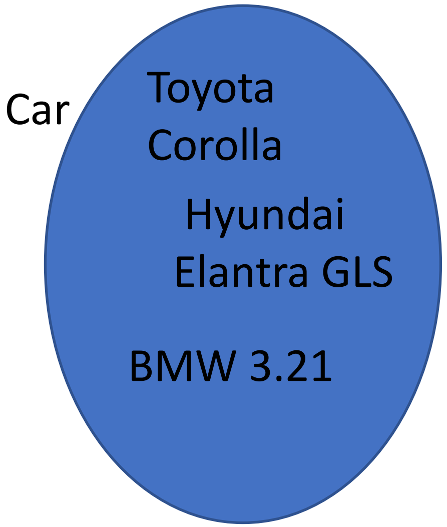

Model	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

In this example, attributes ‘Length’ and ‘Model’ uniquely identify an entity.

The set of all possible key attributes is called ‘**candidate keys**’. {Model,Length}

The DB designer selects one among the candidate keys as **the primary key attribute** for the entity set.

Key concepts of ER: Entity

- Car -> Model, Weight, Length, Max_Speed.

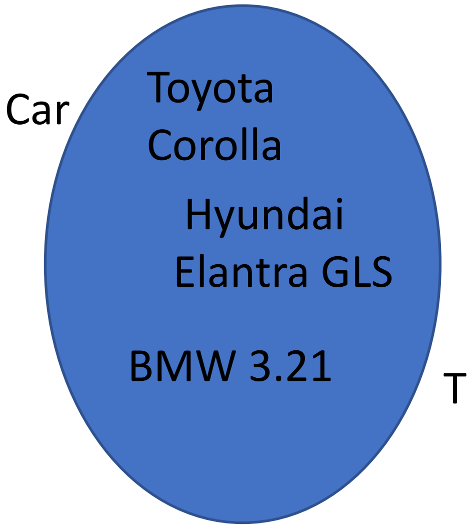

Model	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

The primary key attributes are represented with solid underlines (Brand is selected).

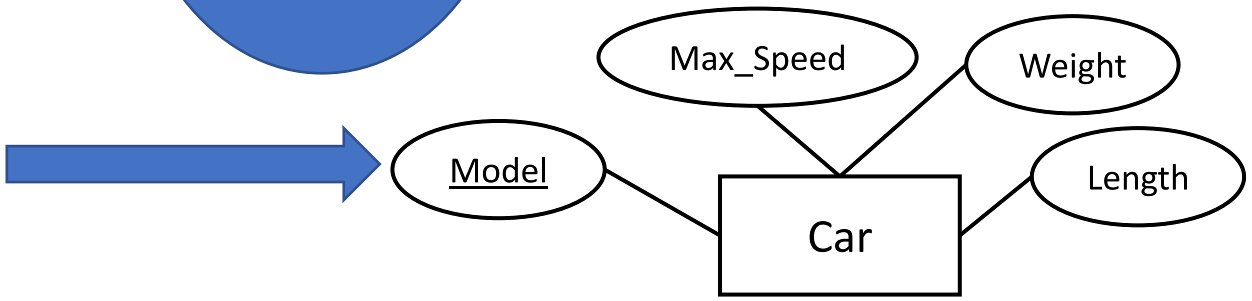

Key concepts of ER: Entity

- Car -> Model, Weight, Length, Max_Speed.

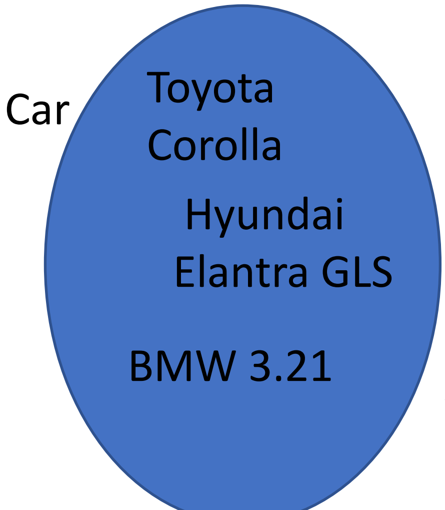

Model	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

Some attributes are computed: Derived attributes.

Derived attributes are given using dashed lines.

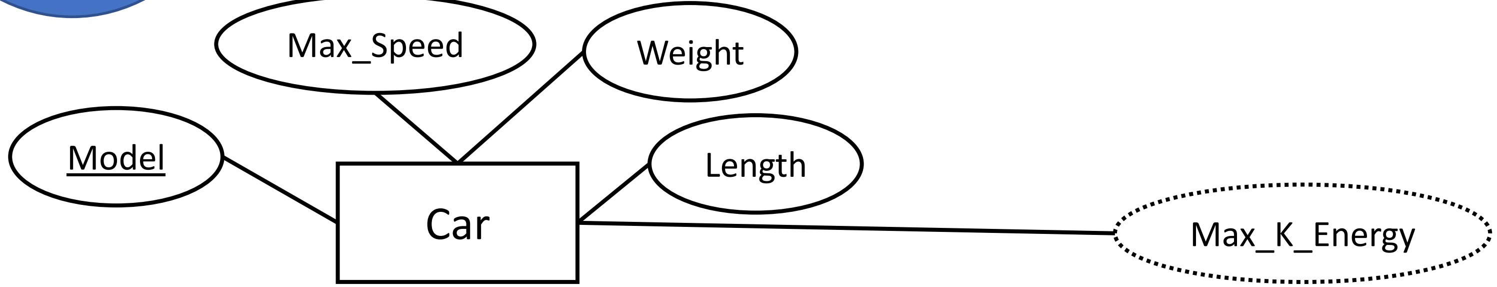

Key concepts of ER: Entity

- Car -> Model, Weight, Length, Max_Speed.

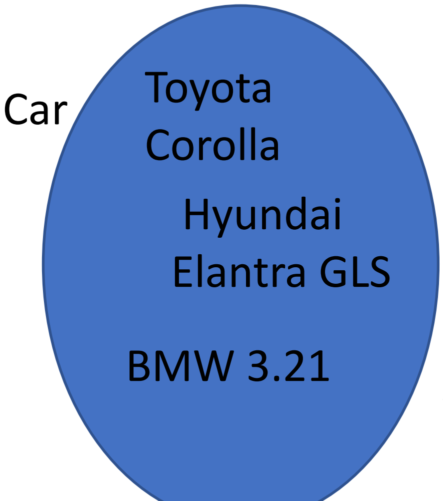

Car

Toyota Corolla

Hyundai Elantra GLS

BMW 3.21

Model	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

Some attributes have multi-values: Multi-valued attributes.

Multi-valued attributes are given using double lines.

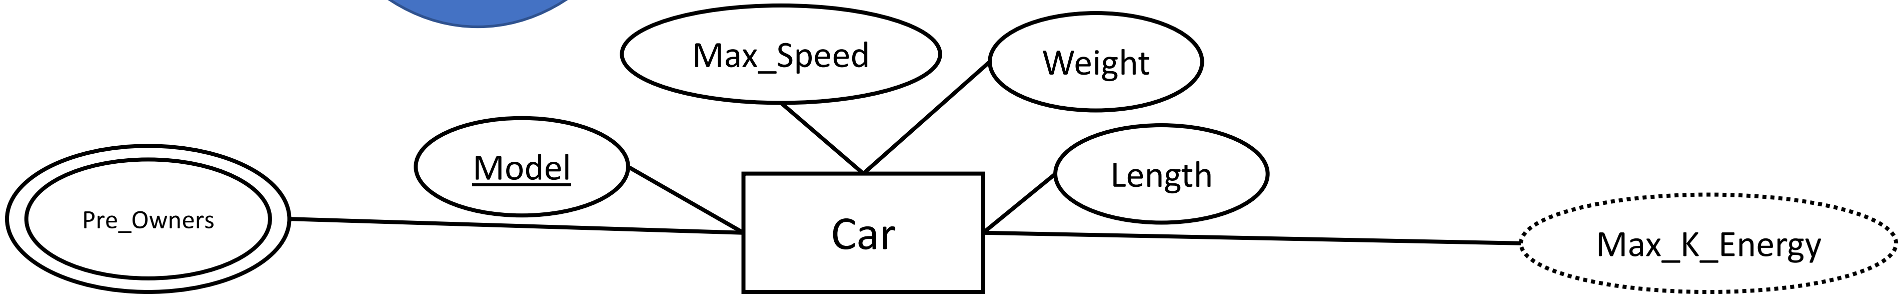

Max_Speed

Weight

Model

Length

Pre_Owners

Car

Max_K_Energy

Check yourself!

- Which are multi-valued attributes?
- Which are Derived Attributes?
- Which are the primary key attributes?

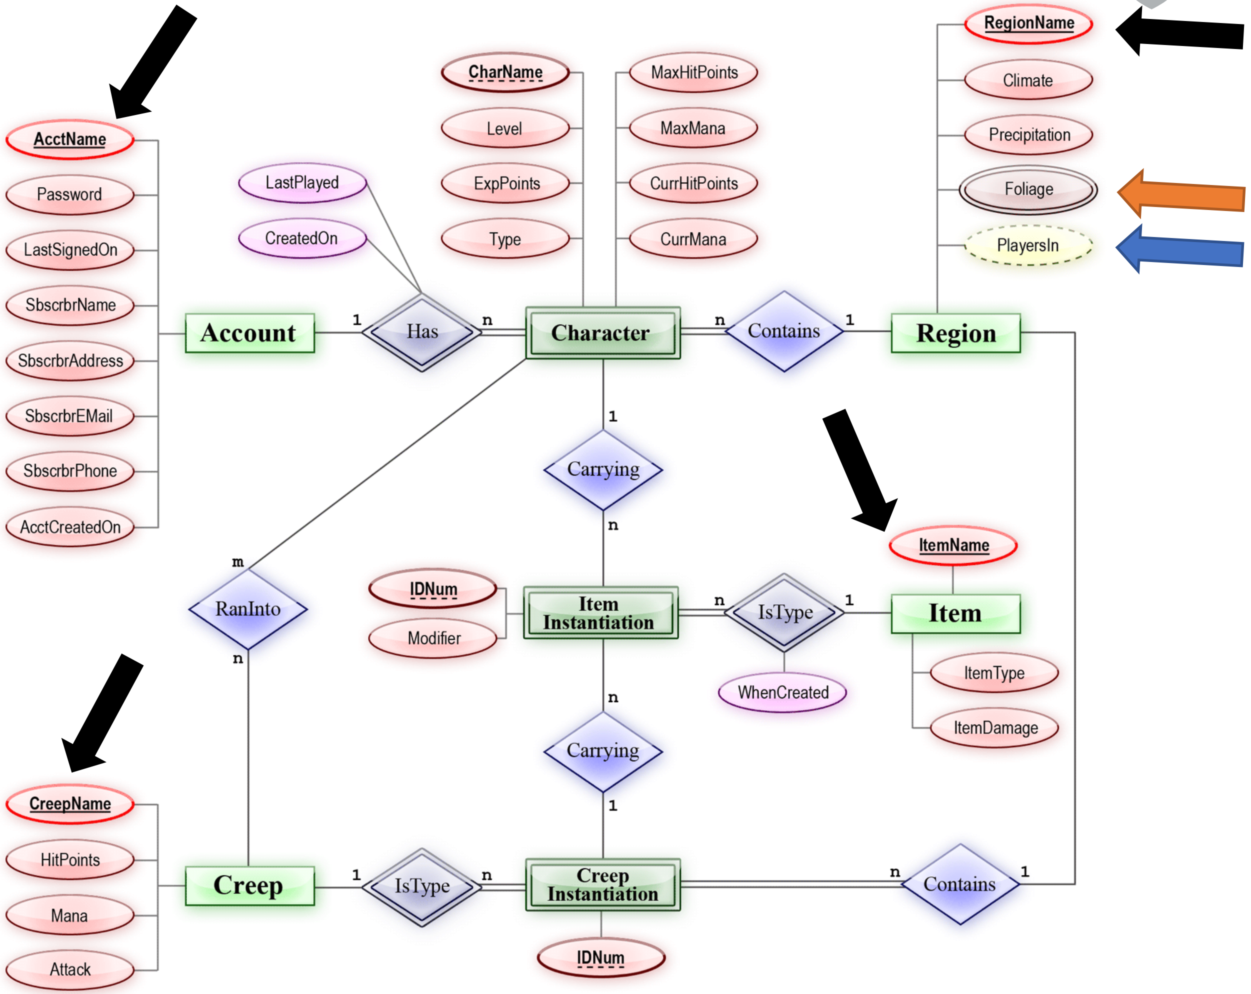

Relations

- In real life, we usually interact with several Entity sets.
  - Cars, Car_Repairers, Drivers, etc.
- To reason queries like:
  - “Tell me the Car repairer for BMW 3.21” we need to be able to associate these entities logically.
- A **<u>Relation</u>** is a logical model for associating two or more **Entities**:
  - *BMW 3.21 can be repaired by Mechanic Alex.*
- A **Relation set** covers all such relations between two or more **Entity sets**.

Key concepts of ER: Relation

A diamond represents a **Relation set.**

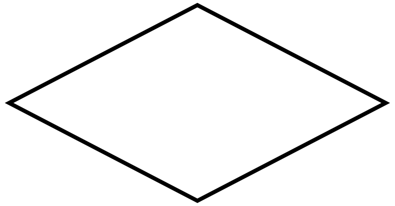

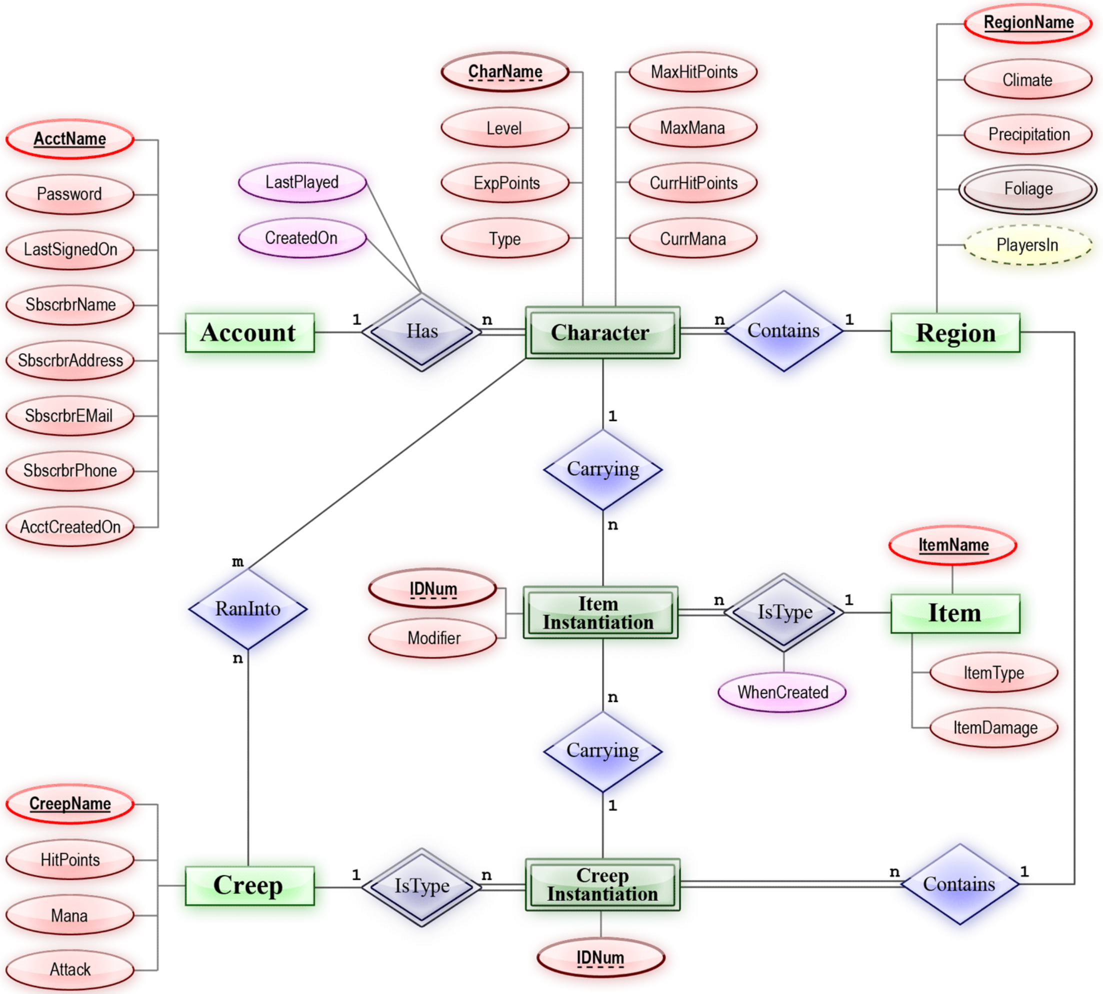

LET’S BEGIN

Recap

- We have seen Database, Database Management System.
- Entity relationship diagram.
    - An **entity** is an object with a set of attributes, and an **Entity set** is a set of entities.
    - The attribute that can distinguish an entity within the Entity set is a **key attribute**.
    - There can be several key attributes (**candidate keys**).

Recap

- The DB designer will **select the primary key** attribute among candidate keys.
- Attributes are atomic unless they are Multi-valued attributes.
- Some attributes are derived by computing the values of other attributes.
- We can relate two or more entities through a **relation**.
- We can relate entity sets through **relationship sets.**

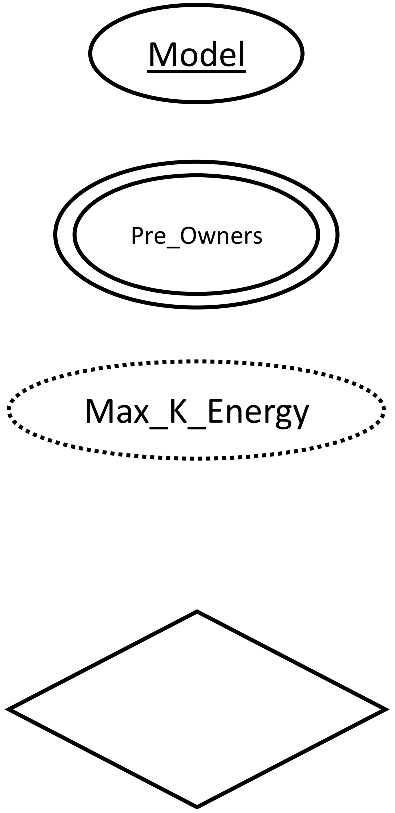

Key concepts of ER: Relation

A relation set ‘Repairs’ relates two Entity Sets.

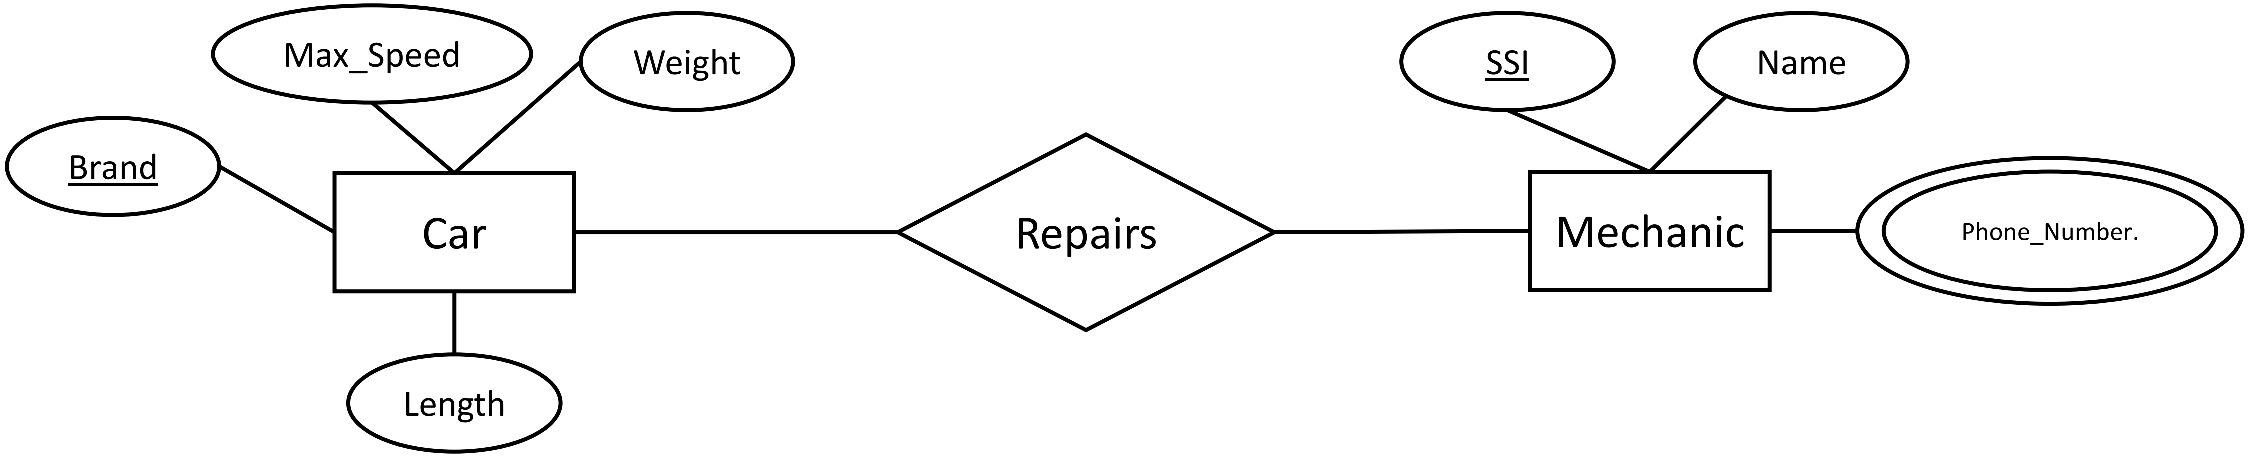

Key concepts of ER: Relation

- Relating two or more entity sets requires the designer to set mapping cardinalities of this relation.

SSI	Name	Phone_Nu
87542702	Tom	75315567
68201937	Uraz	75335521
23139827	Nick	75315544

Name	Brand
Nick	Hyundai.
Tom	Corolla..
Uraz	BMW..

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

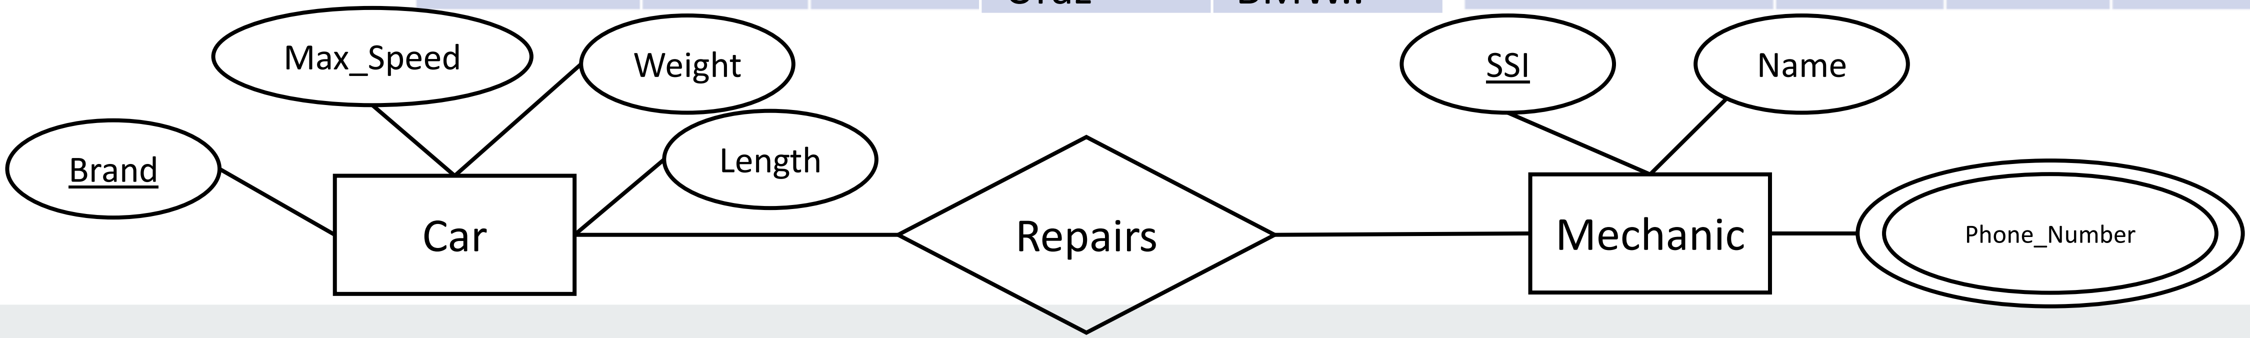

ER Model

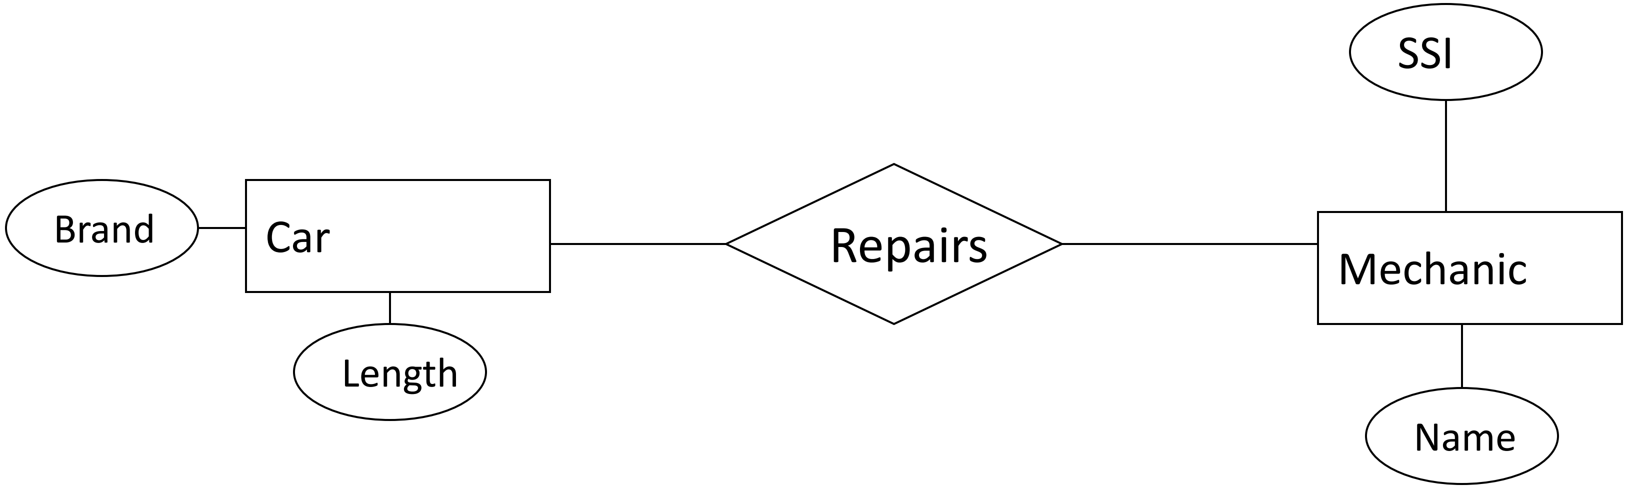

- Relationship sets also have attributes

ER Model

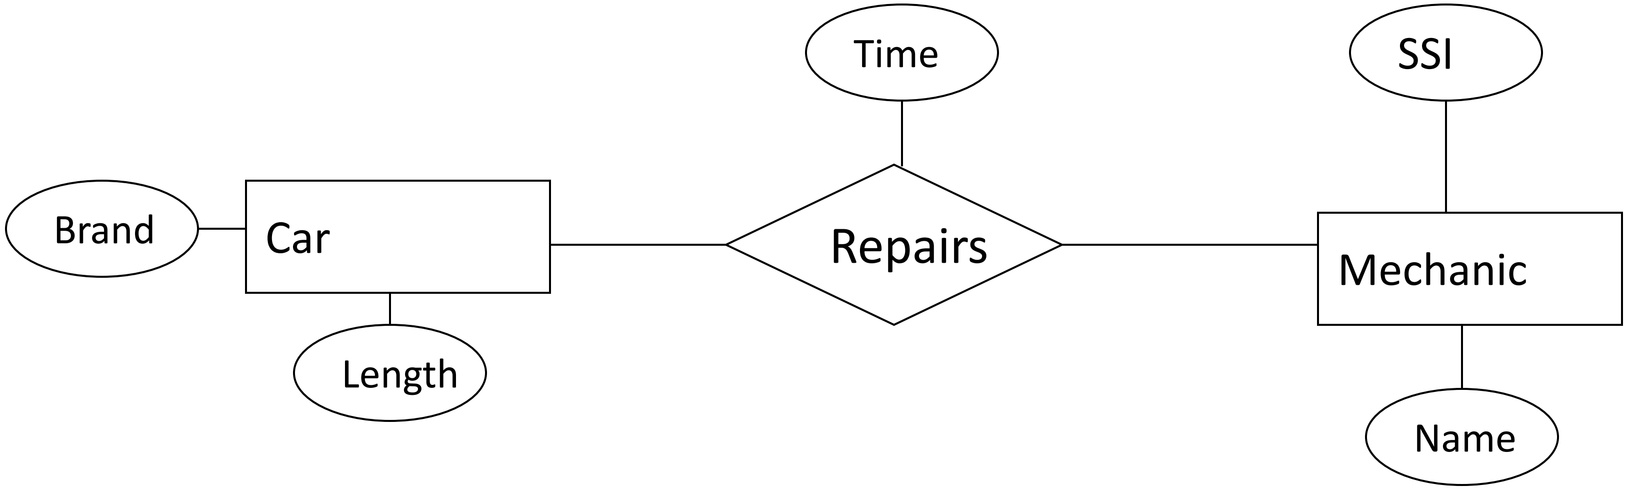

- Relationship sets also have attributes
- We will talk about the key in a relationship set later.

ER Model

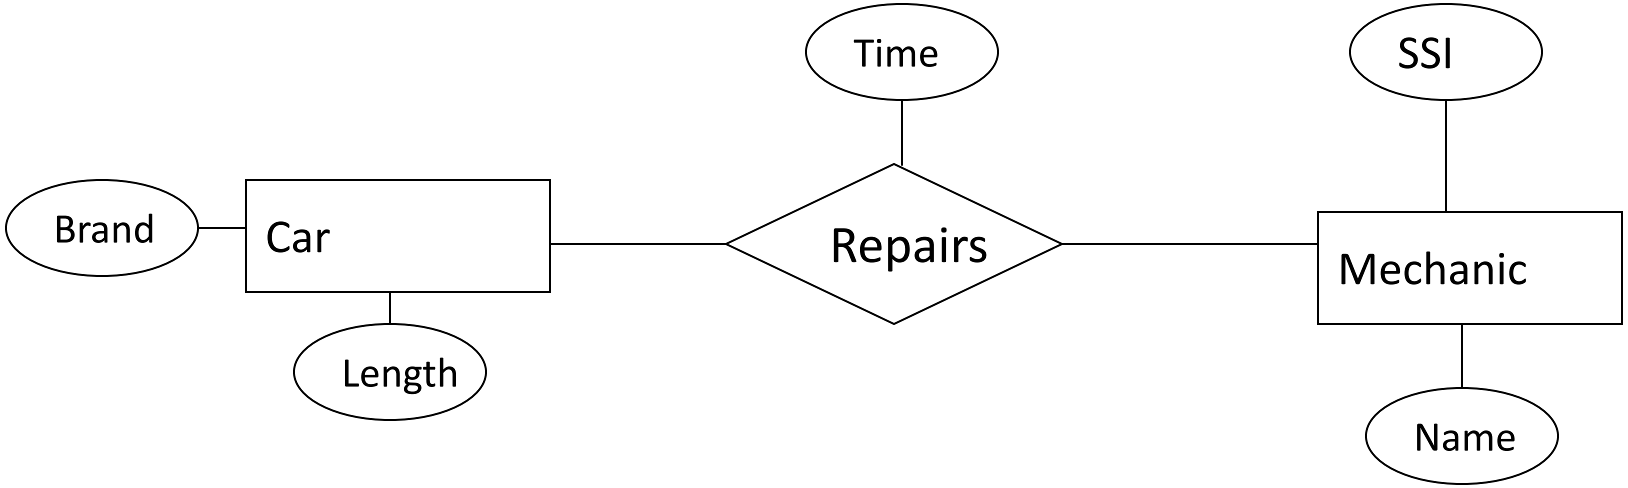

- **Degree of a relationship set** is the number of entity sets that participate in a relationship
- **Binary relationship** sets involve two entity sets

ER Model

- Ternary relationship sets involve three entity sets

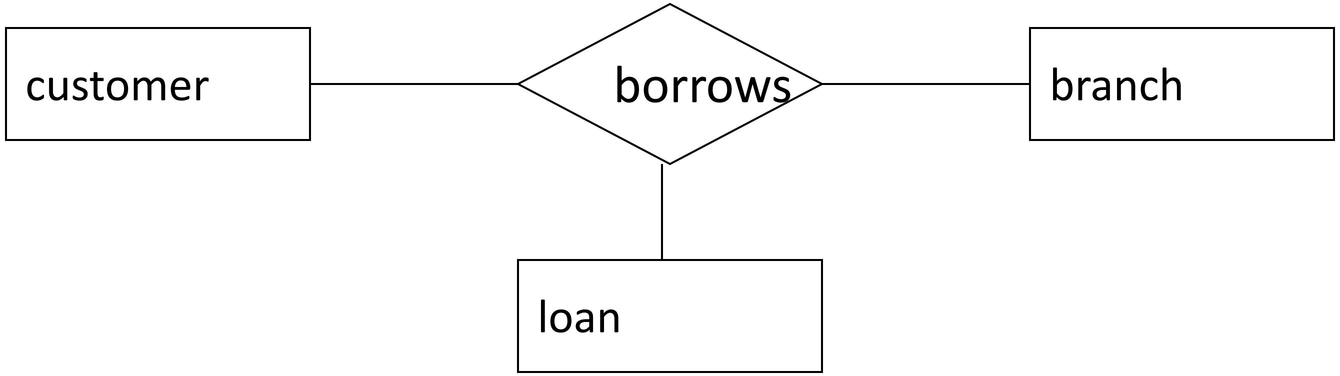

Ternary relationship sets.

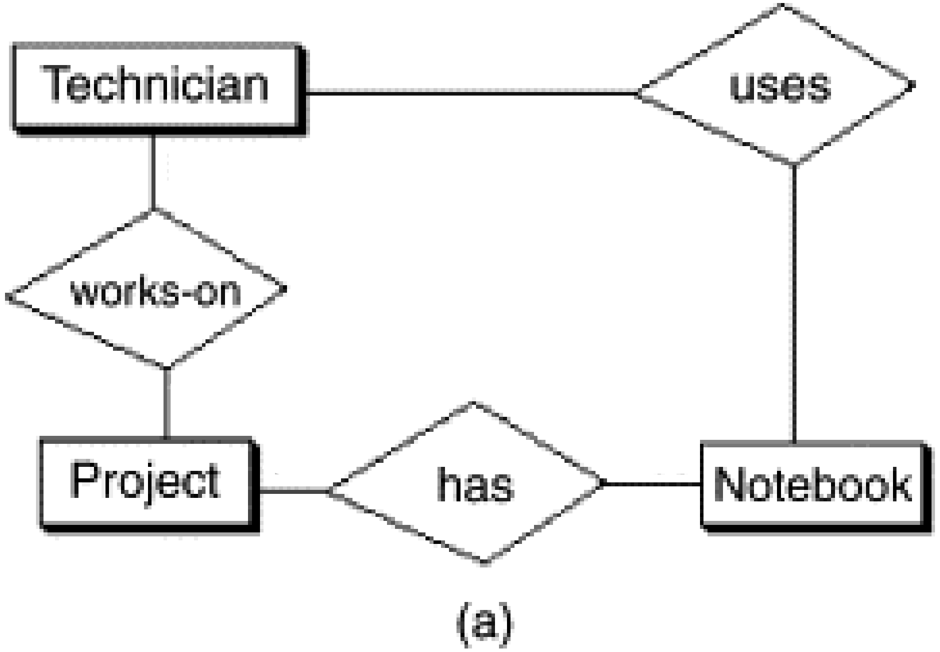

ER Model

- We may have relationships among the entities that belong to the same entity set
- each entity has a role in such a relationship

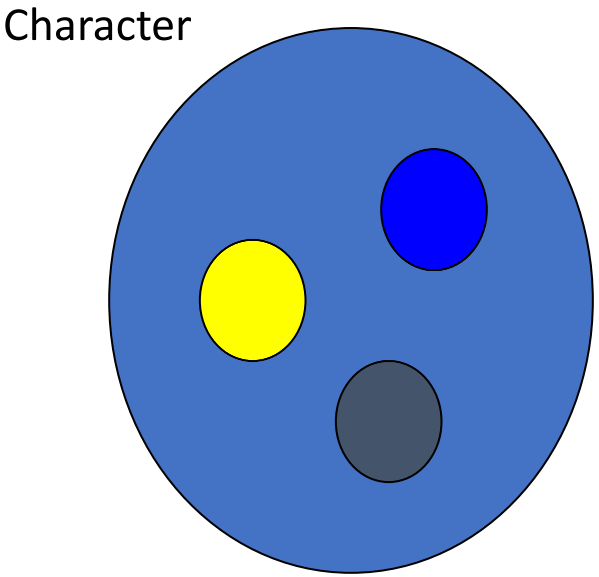

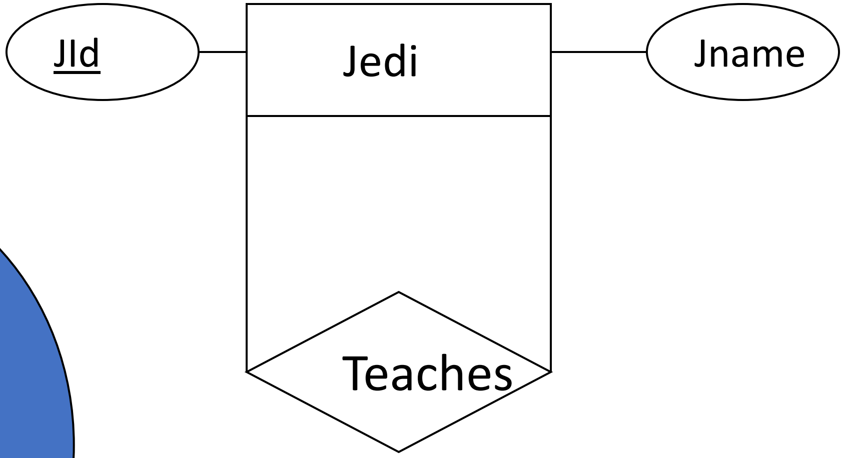

ER Model

- We may have relationships among the entities that belong to the same entity set
- **each entity has a role in such a relationship.**

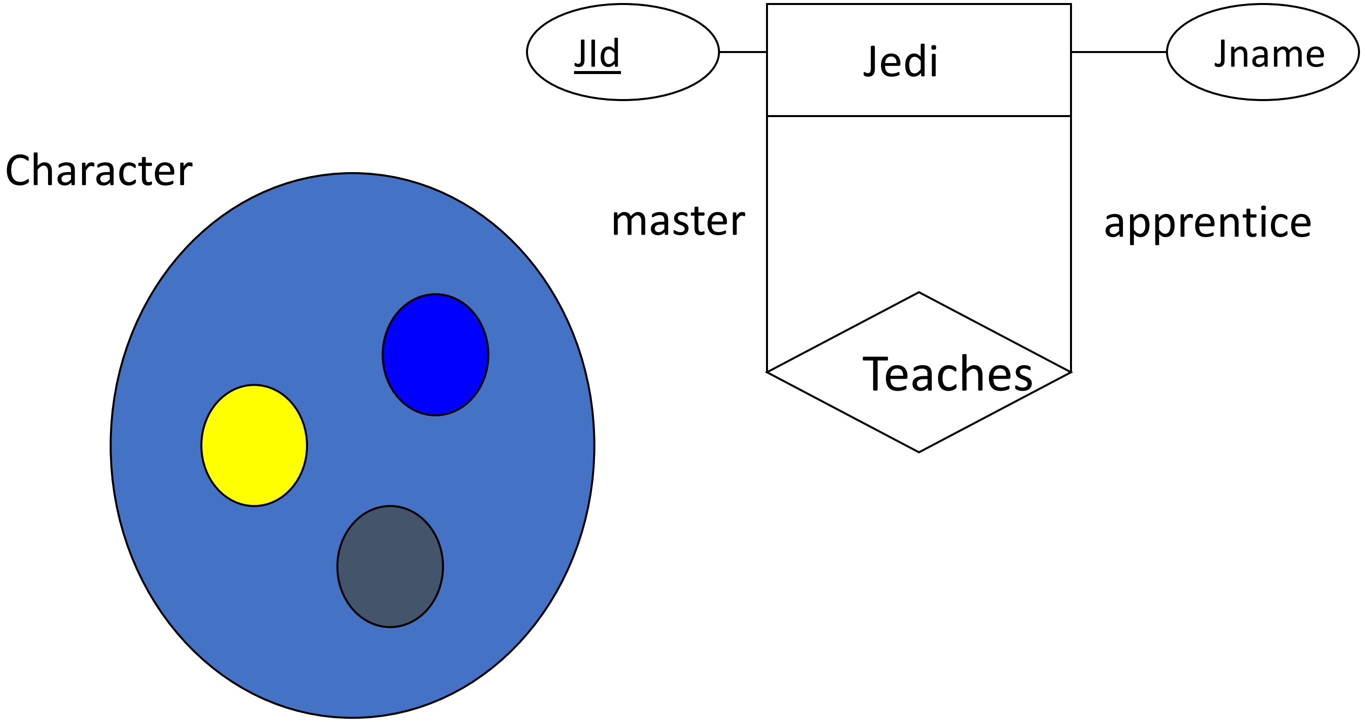

ER Model

- We may have relationships among the entities that belong to the same entity set
- each entity has a role in such a relationship.
- **What is the degree of the following relationship set (2 or 1)?**

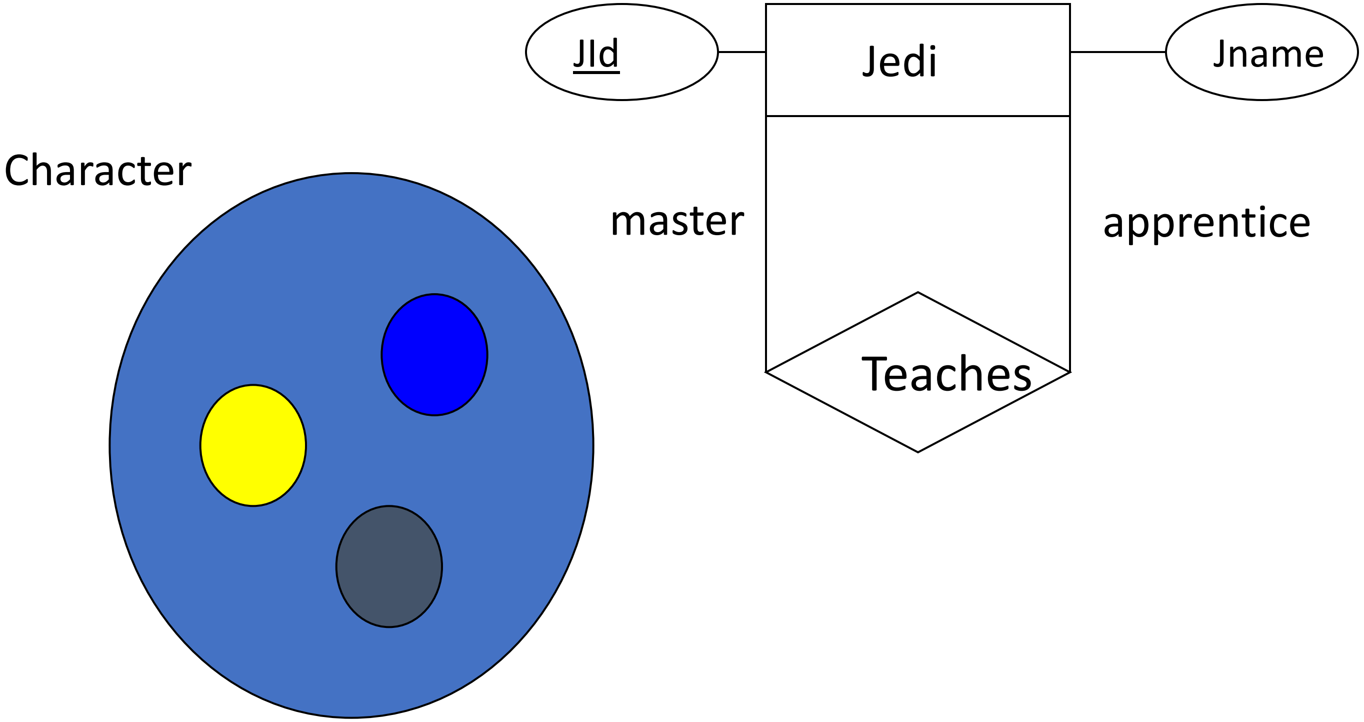

ER Model

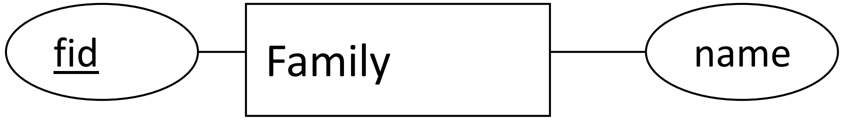

ER Model

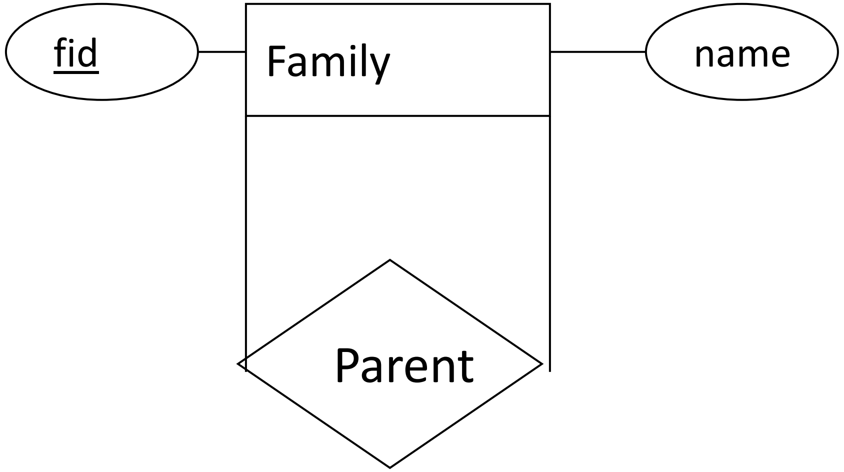

ER Model

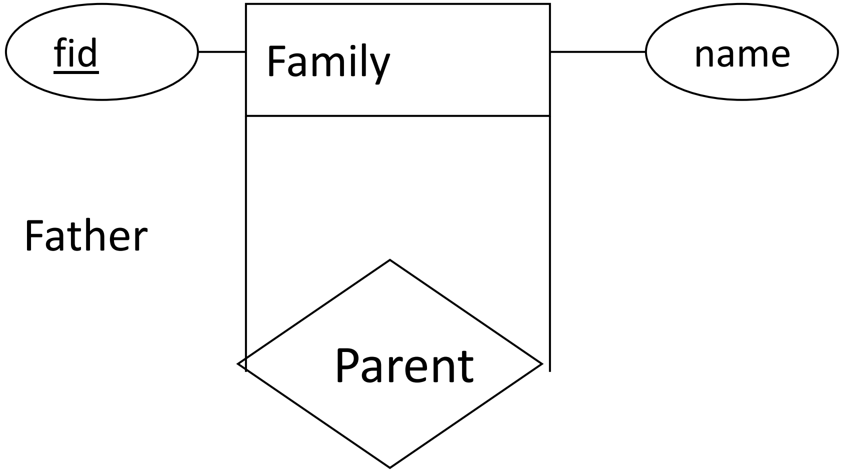

ER Model

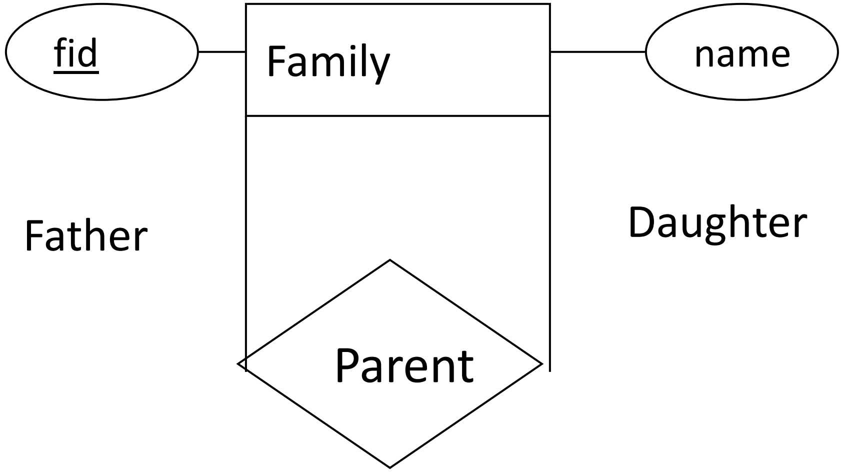

ER Model

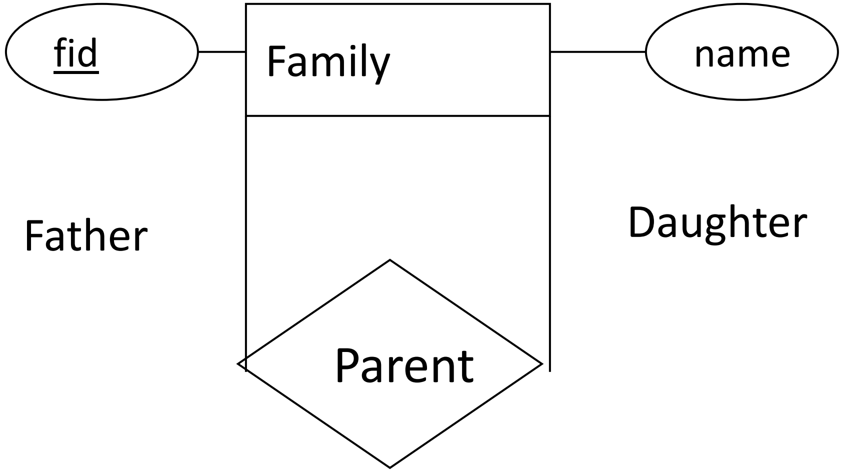

Key concepts of ER: Cardinalities.

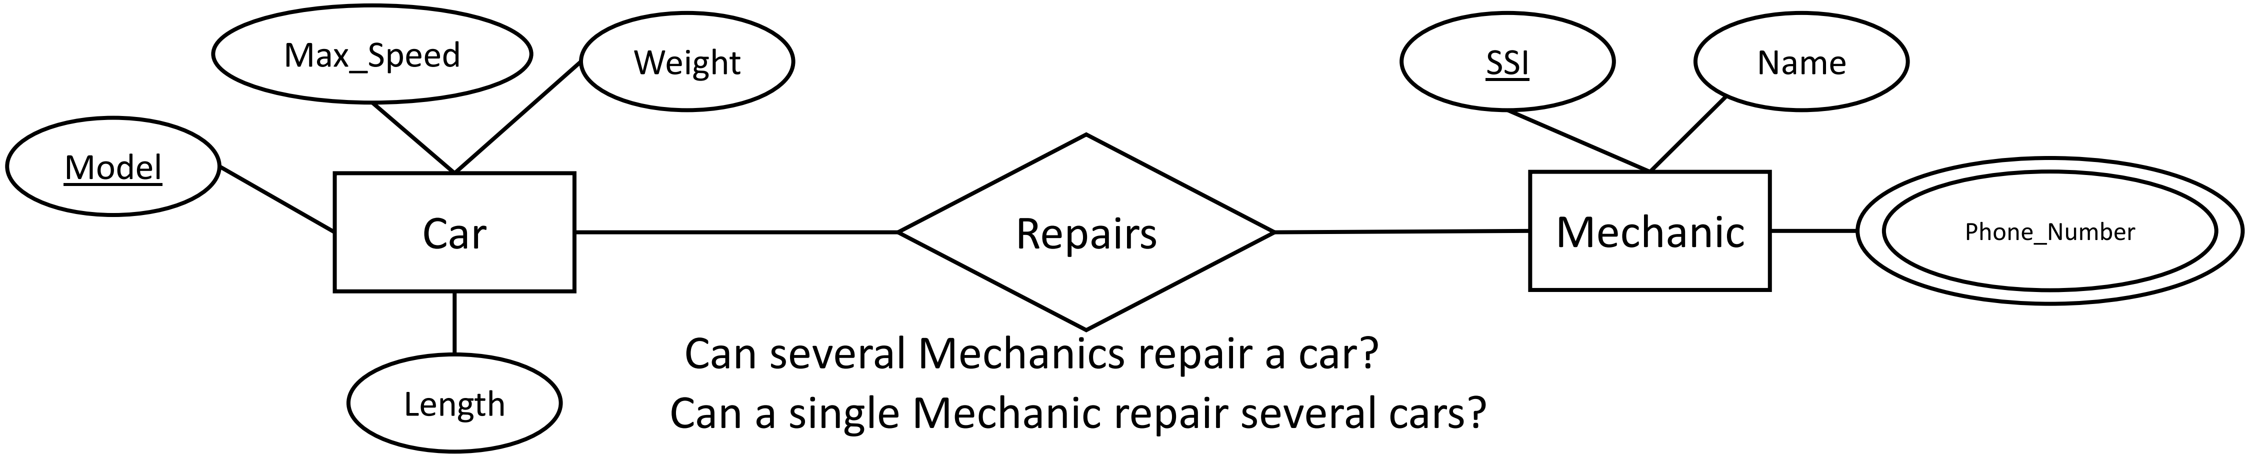

Can several Mechanics repair a car?

Can a single Mechanic repair several cars?

Yes/No -> Designer will decide.

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

SSI	Name	Phone_Number
87542702	Tom	75315567, 75315264
68201937	Uraz	75335521, 75334567
23139827	Nick	75315544, 75315237

Mapping cardinalities

- Relating two or more entity sets requires the designer to set mapping cardinalities of this relation.
- One-to-One relationship (ex: marriage relationship set between husbands and wives)

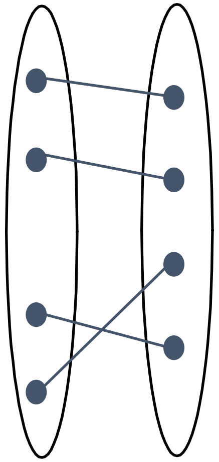

1-to-1

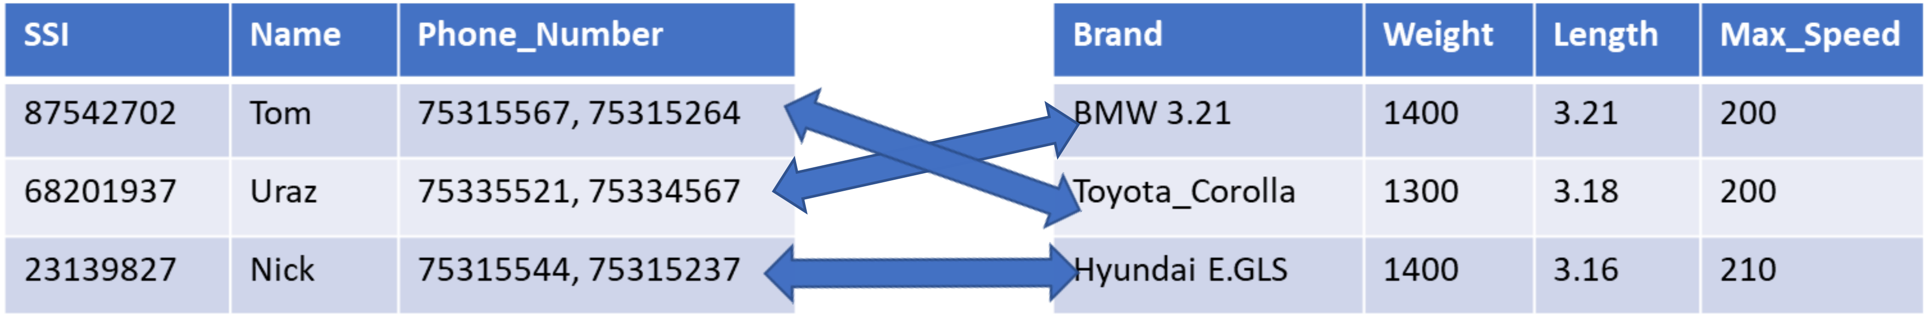

SSI	Name	Phone_Number
87542702	Tom	75315567, 75315264
68201937	Uraz	75335521, 75334567
23139827	Nick	75315544, 75315237

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

Mapping cardinalities

- Relating two or more entity sets requires the designer to set mapping cardinalities of this relation.
- One-to-One relationship (ex: marriage relationship set between husbands and wives)

Name	Brand
Tom	Toyota..
Nick	Hyundai.
Uraz	BMW..

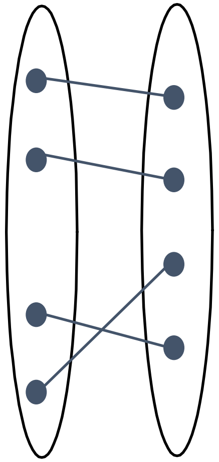

1-to-1

SSI	Name	Phone_Number
87542702	Tom	75315567, 75315264
68201937	Uraz	75335521, 75334567
23139827	Nick	75315544, 75315237

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

Mapping cardinalities

- One-to-One (ex: marriage relationship set between husbands and wives)
- One-to-Many (ex?)

Name	Brand
Uraz	BMW
Uraz	Hyundai
Nick	Toyota

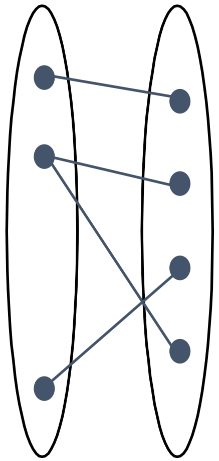

1-to Many

SSI	Name	Phone_Number
87542702	Tom	75315567, 75315264
68201937	Uraz	75335521, 75334567
23139827	Nick	75315544, 75315237

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

Mapping cardinalities

- One-to-One (ex: marriage relationship set between husbands and wives)
- One-to-Many (ex?)
- Many-to-One (ex?)

Name	Brand
Uraz	BMW
Nick	BMW
Tom	Toyota

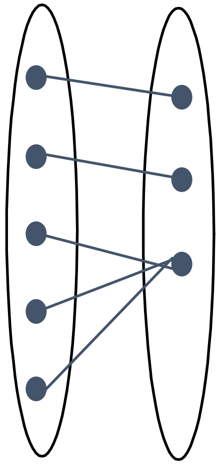

Many-to-1

SSI	Name	Phone_Number
87542702	Tom	75315567, 75315264
68201937	Uraz	75335521, 75334567
23139827	Nick	75315544, 75315237

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

Mapping cardinalities

- One-to-One (ex: marriage relationship set between husbands and wives)
- One-to-Many (ex?)
- Many-to-One (ex?)
- Many-to-Many (ex?)

Name	Brand
Uraz	BMW
Uraz	Hyundai
Nick	Toyota
Nick	BMW

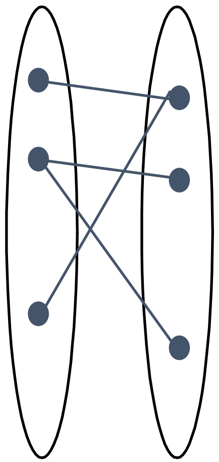

Many-to-Many

SSI	Name	Phone_Number
87542702	Tom	75315567, 75315264
68201937	Uraz	75335521, 75334567
23139827	Nick	75315544, 75315237

Brand	Weight	Length	Max_Speed
BMW 3.21	1400	3.21	200
Toyota_Corolla	1300	3.18	200
Hyundai E.GLS	1400	3.16	210

- Consider the works_in relationship
- If an employee can work in a department and a department can have multiple employees
- What type of relationship is that?

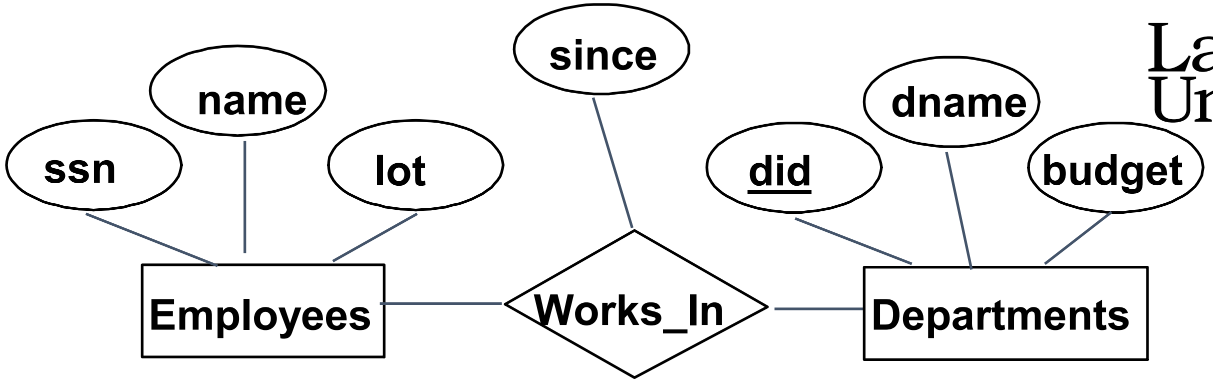

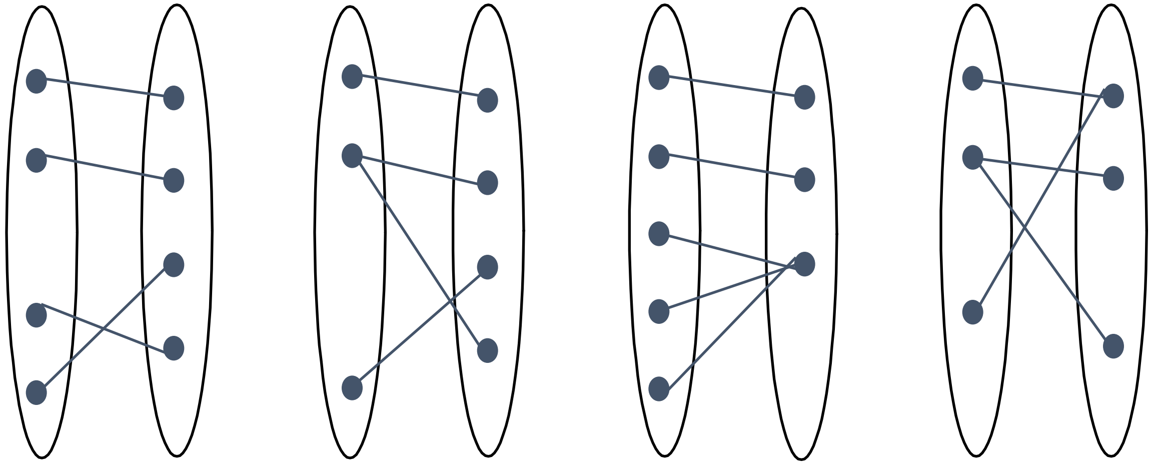

1-to-1 1-to Many Many-to-1 Many-to-Many

- Consider the works_in relationship
- If an employee can work in several departments and a department can have multiple employees
- What type of relationship is that?

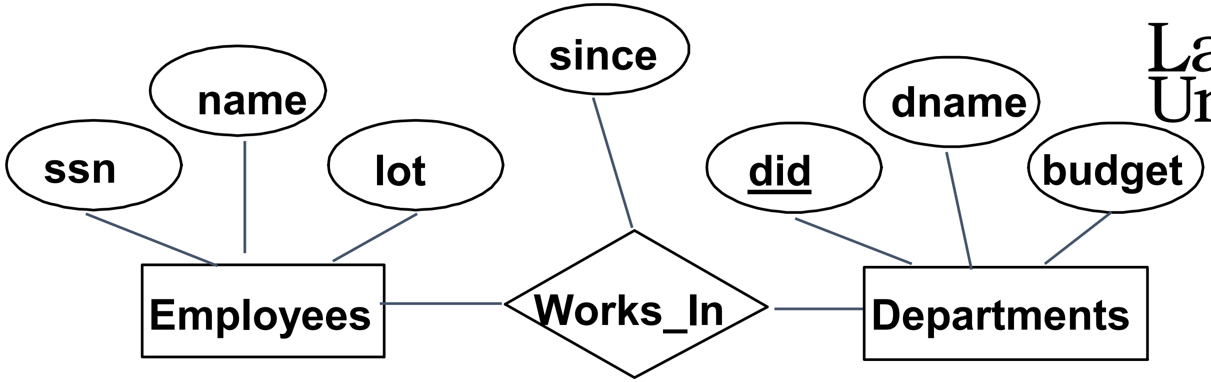

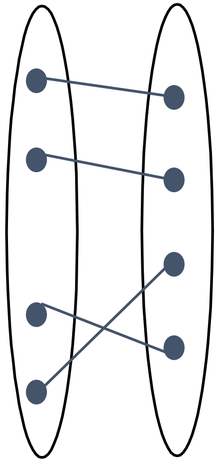
1-to-1

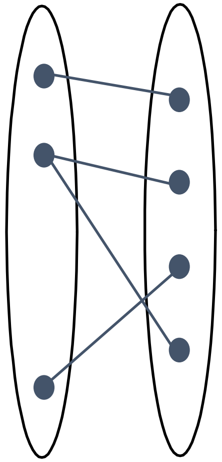
1-to Many

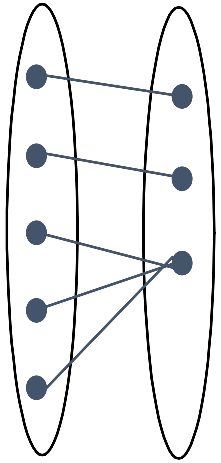
Many-to-1

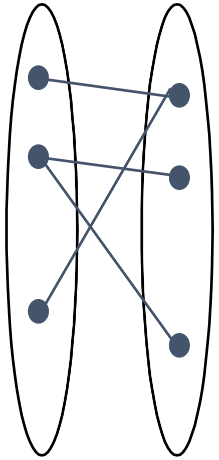
Many-to-Many

How do we encode cardinality into ER diagrams?

- We use Chen’s notation
- 1:1 is for one-to-one
- 1:N is for one to many
- N:1 is for many to one
- N:M is for many to many.

1:1

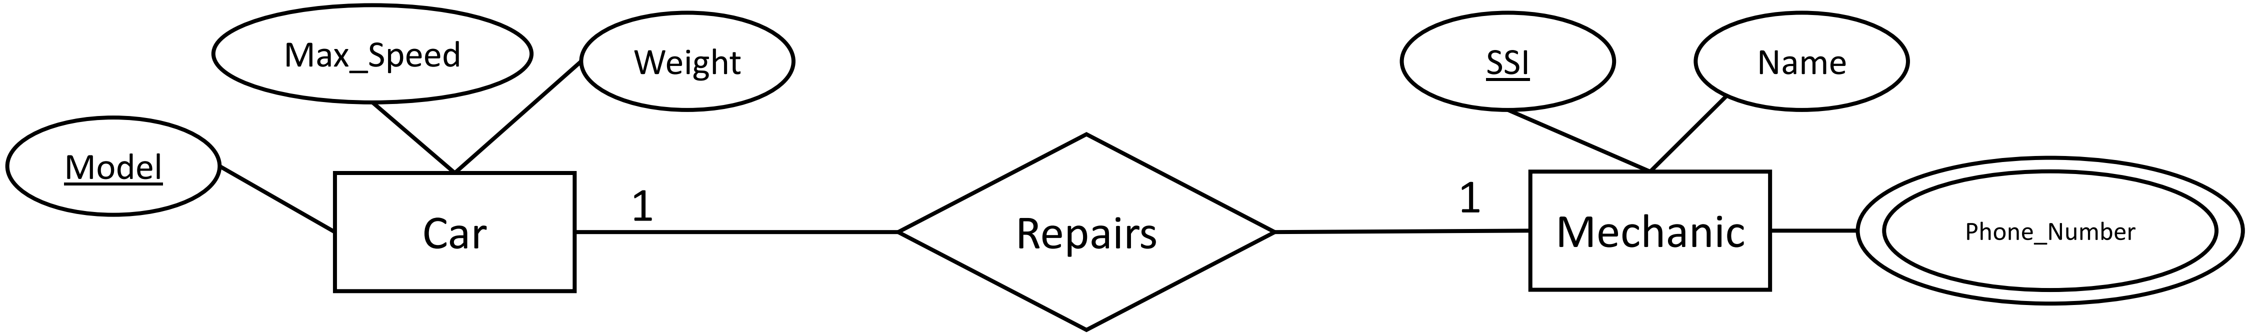
Max_Speed
Weight
Model
Car
1
Repairs
1
Mechanic
SSI
Name
Phone_Number

1:N One to many

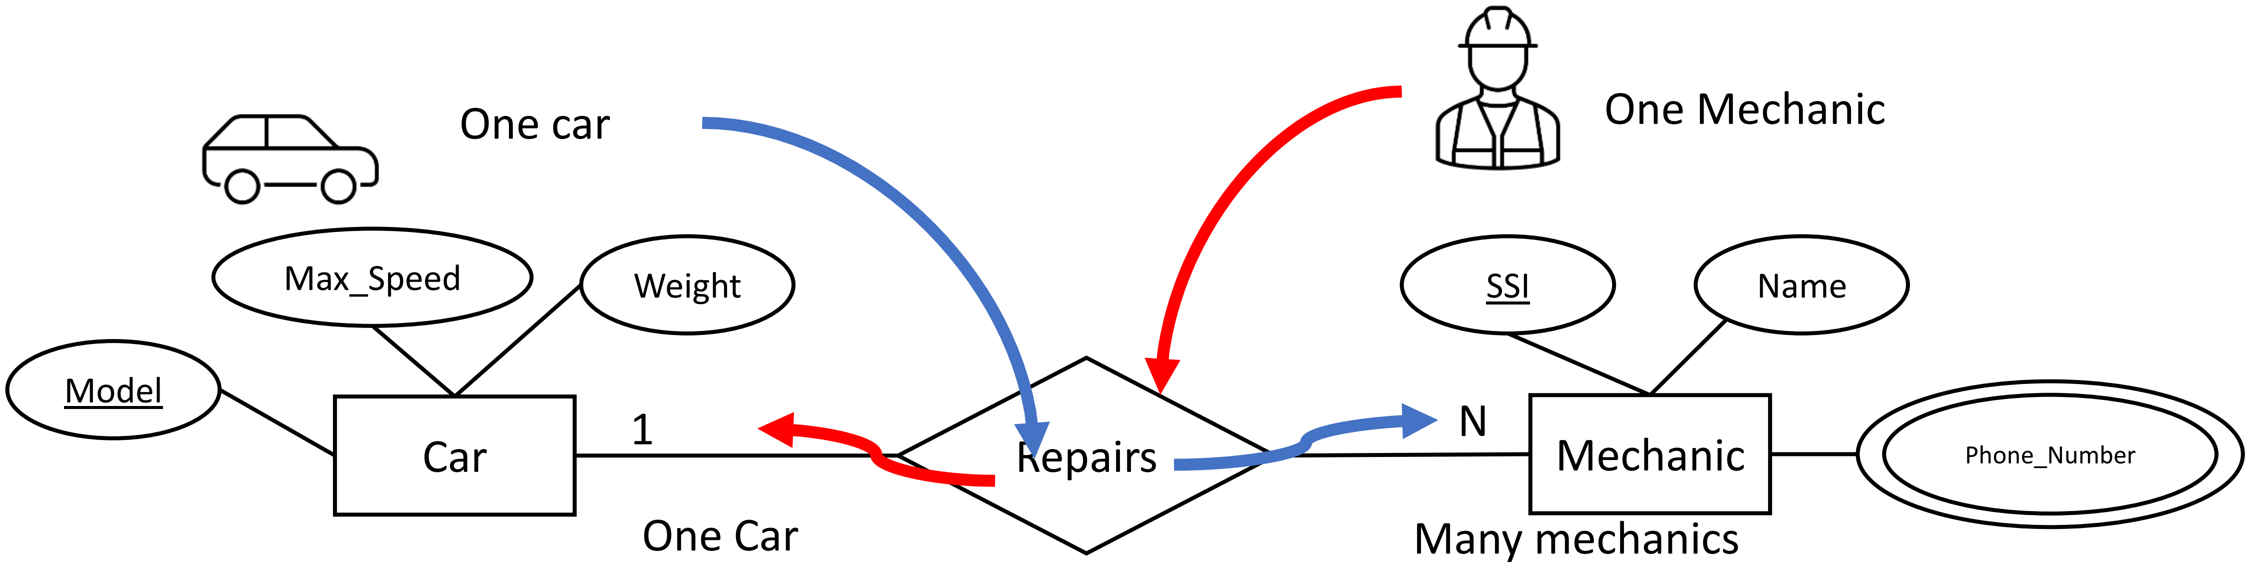

One car can be repaired by many mechanics.
A mechanic can repair one car.

N:1 Many to one

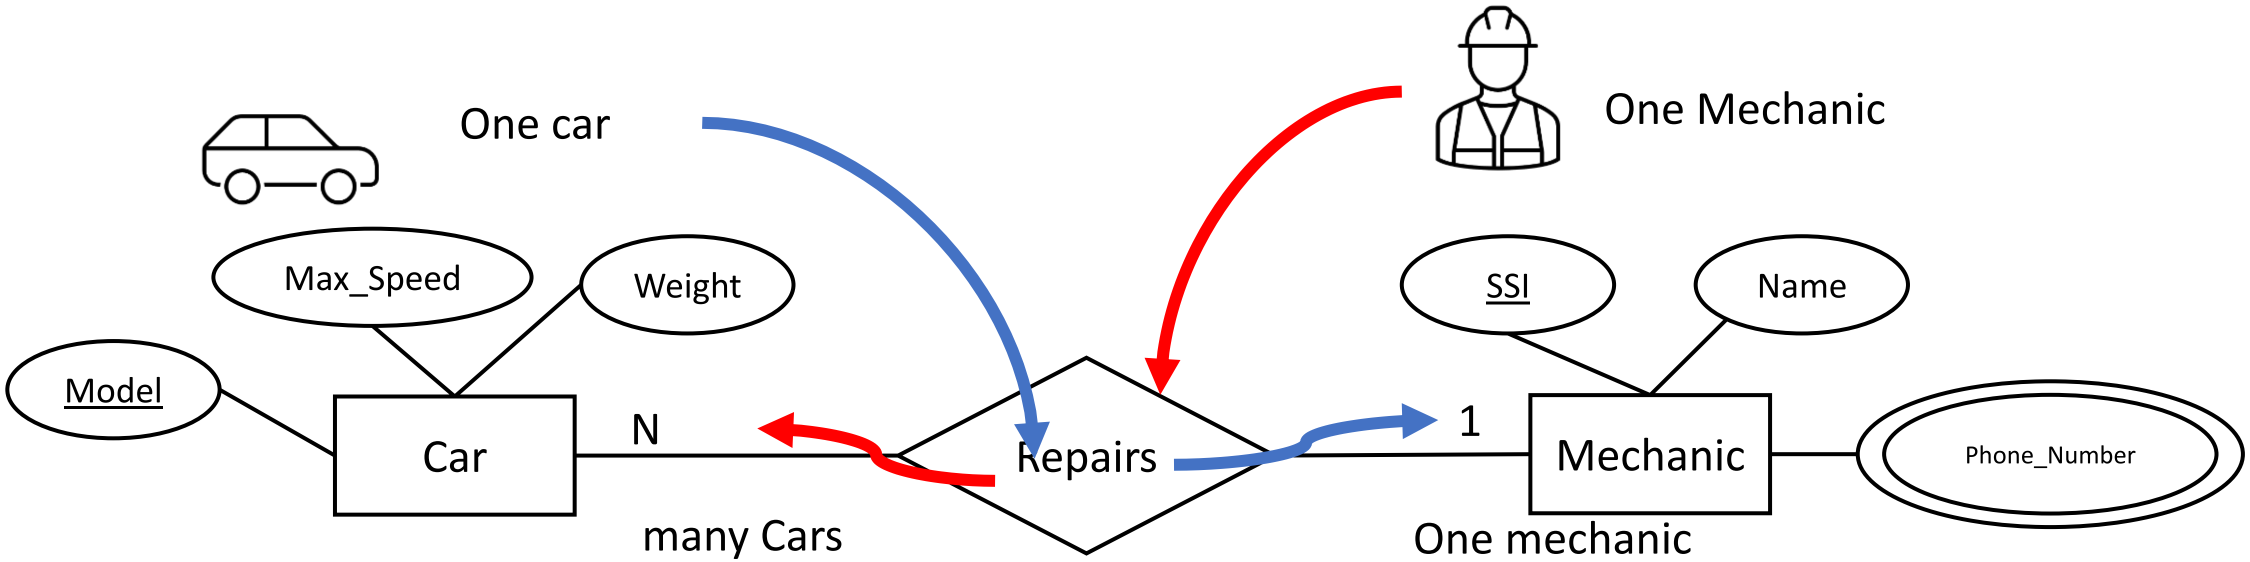

One car can be repaired by a mechanic.
A mechanic can repair many cars.

N:N many to many

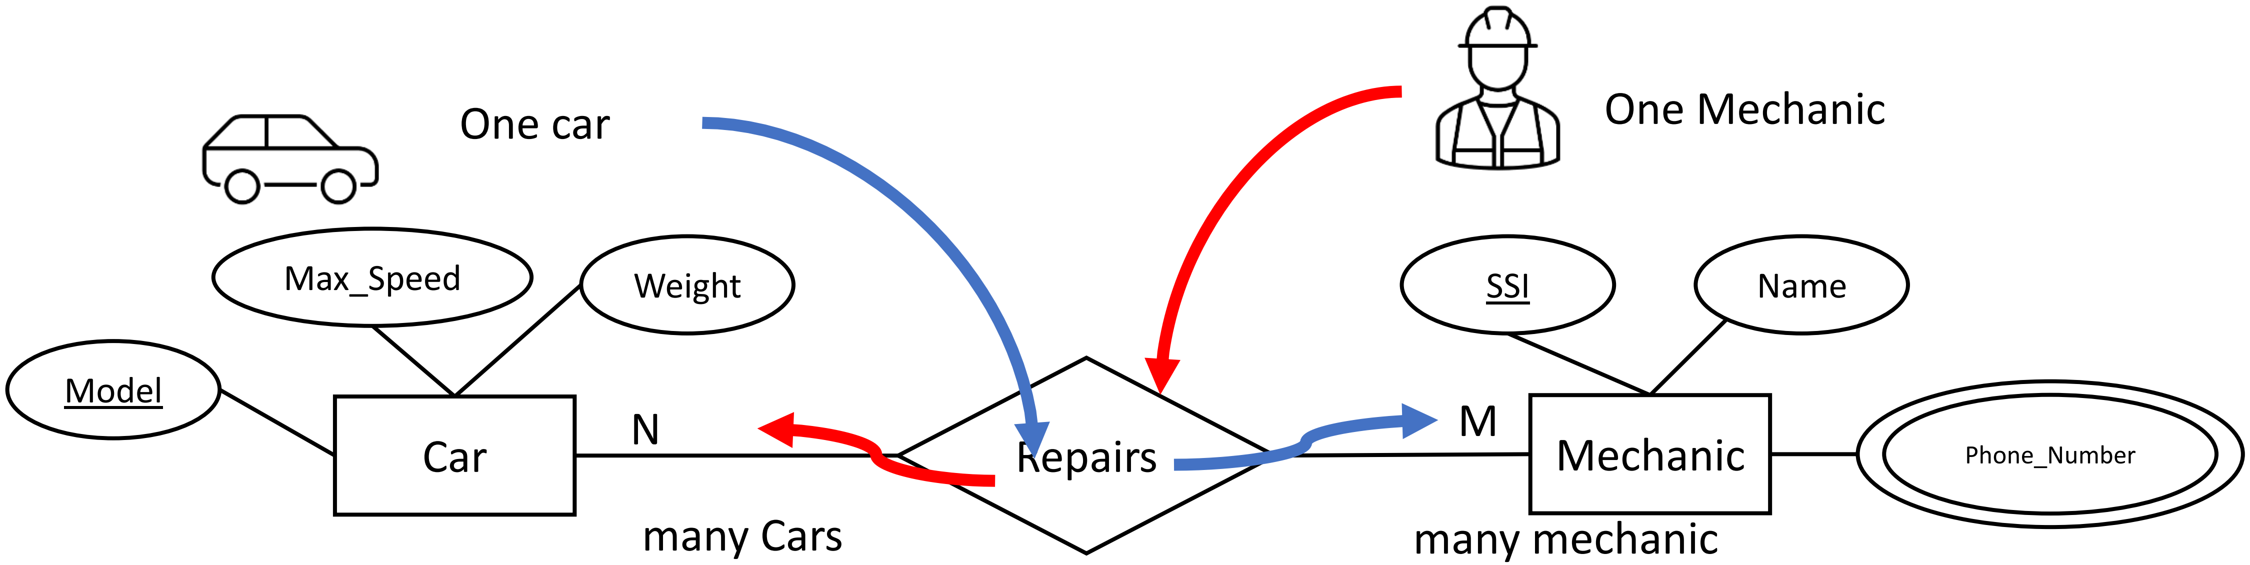

One car can be repaired by many mechanics.
A mechanic can repair many cars.

Participation constraints

- Can there be a mechanic that can not repair a car?
- If not, we need to state there is a **Total Participation.**

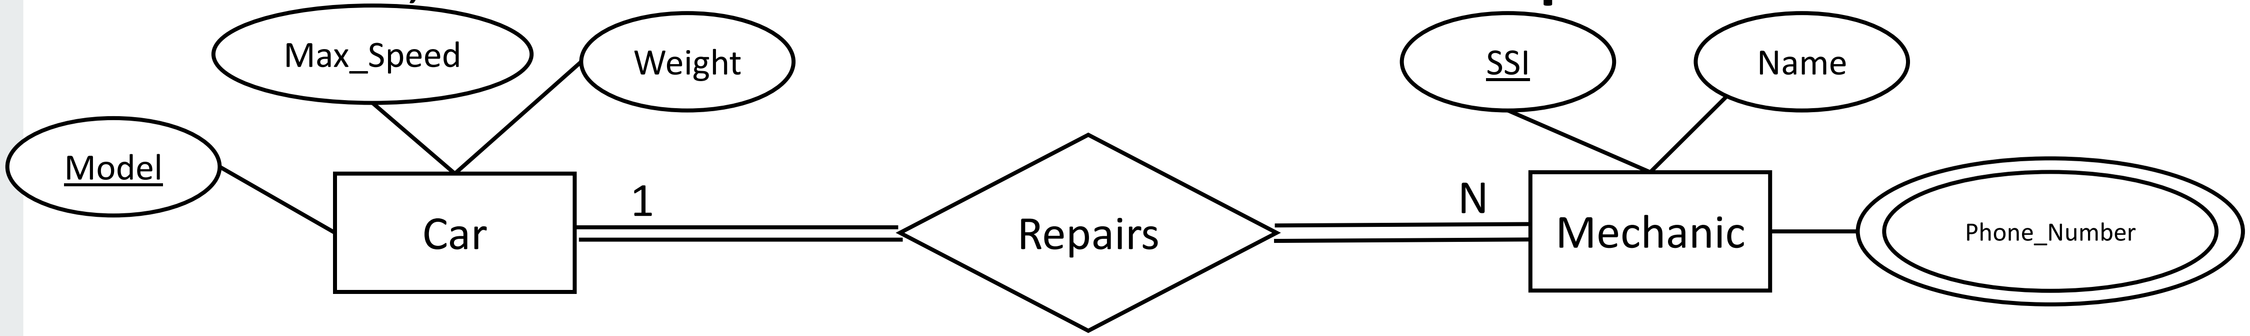

A **double line** identifies total participation.

- If so, then it is **Partial Participation**

A **single line** identifies partial participation.

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Manages).

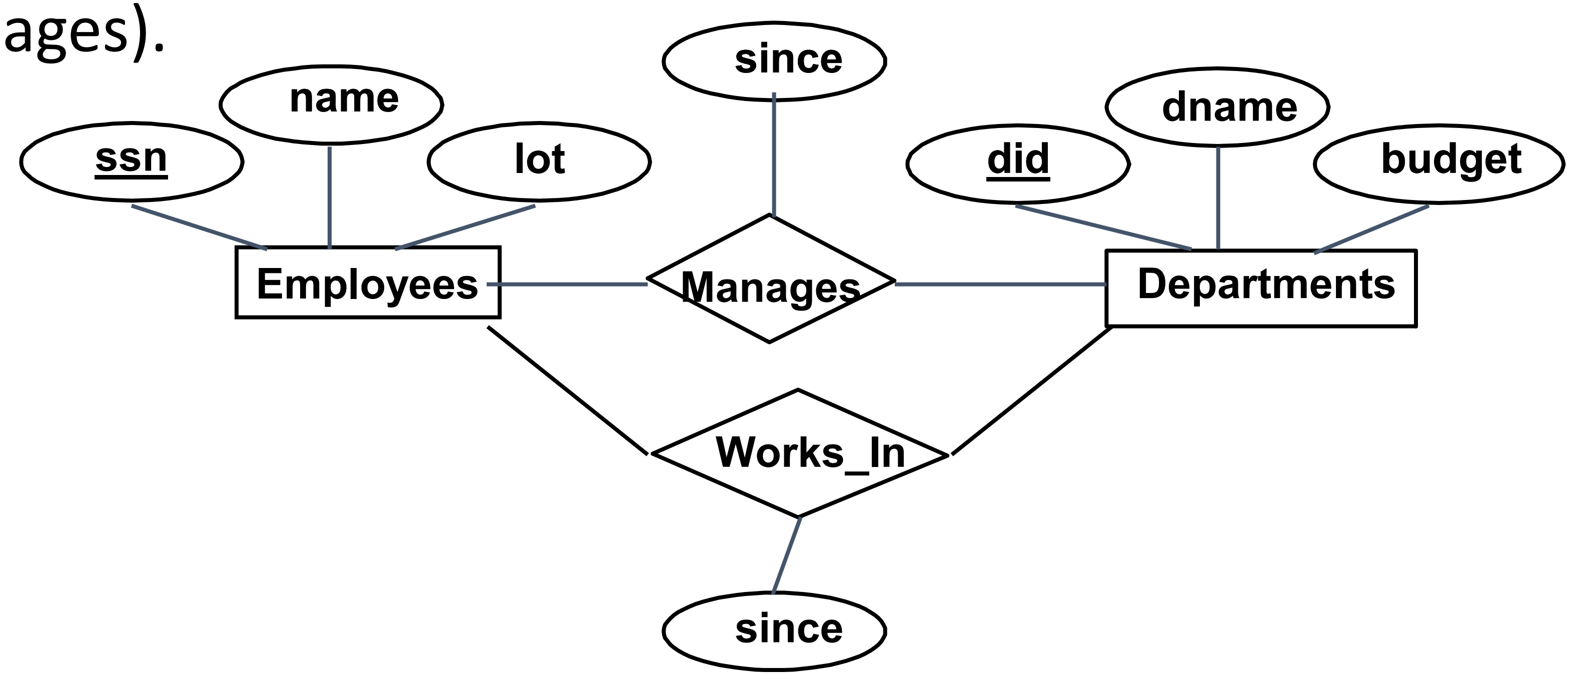

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Manages).

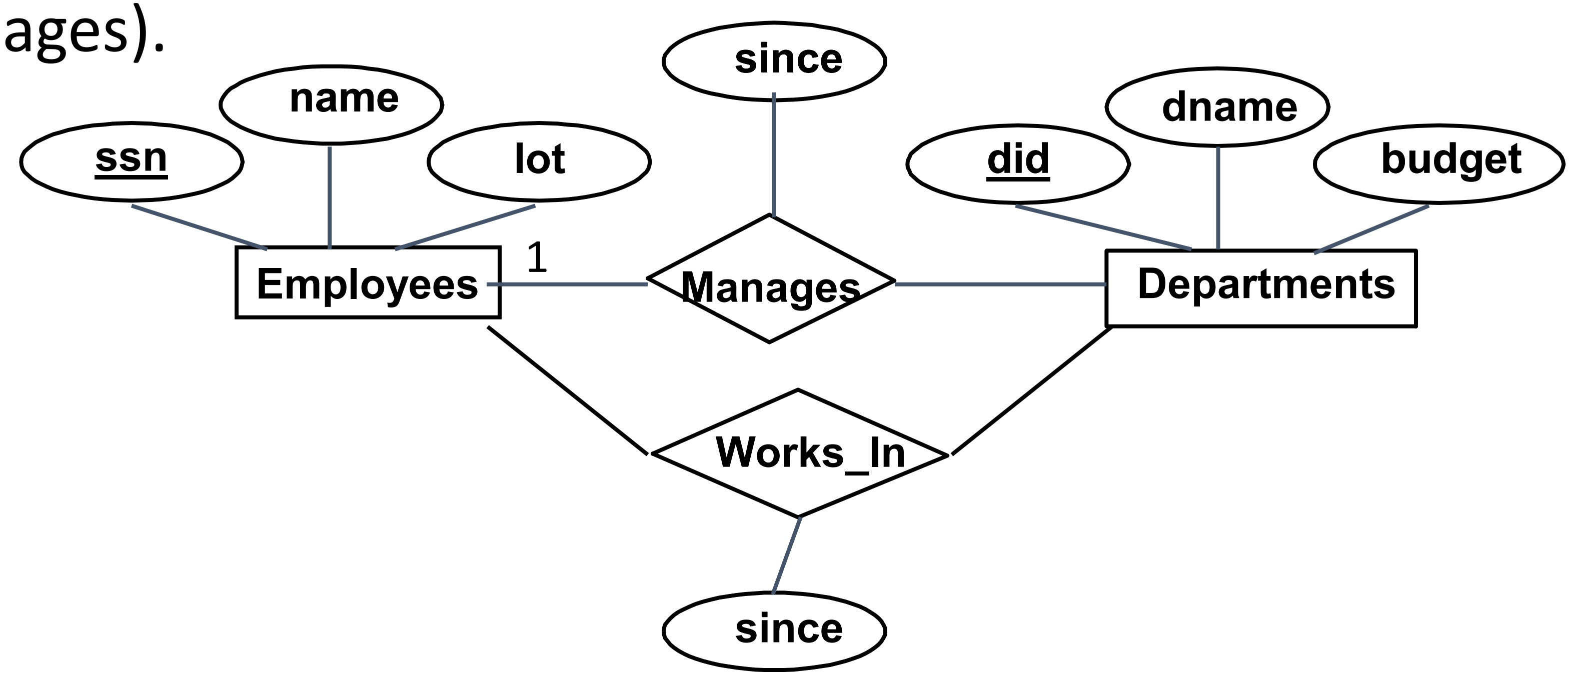

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Manages).

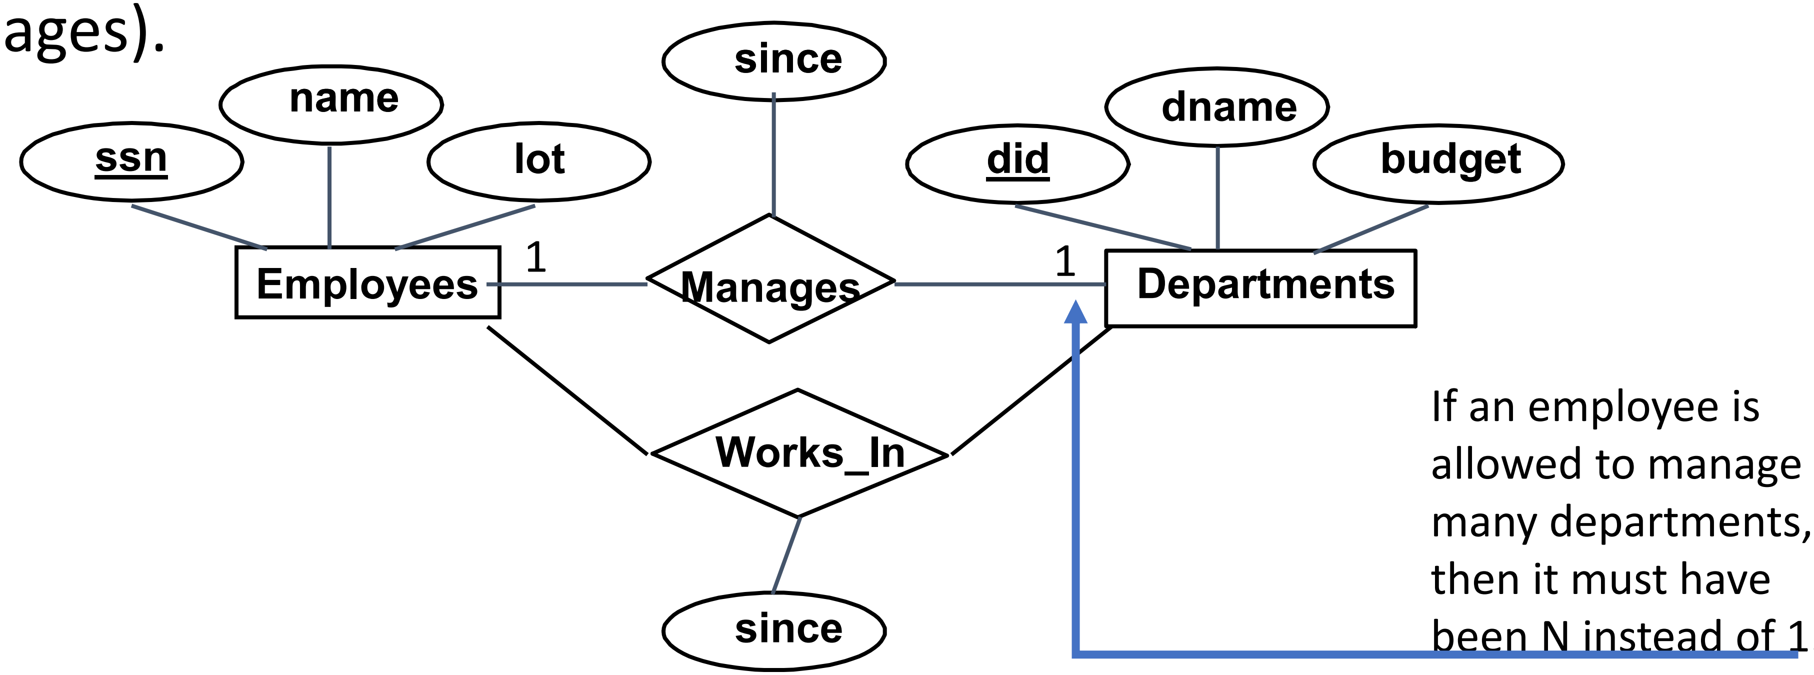

If an employee is allowed to manage many departments, then it must have been N instead of 1.

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Works_In).

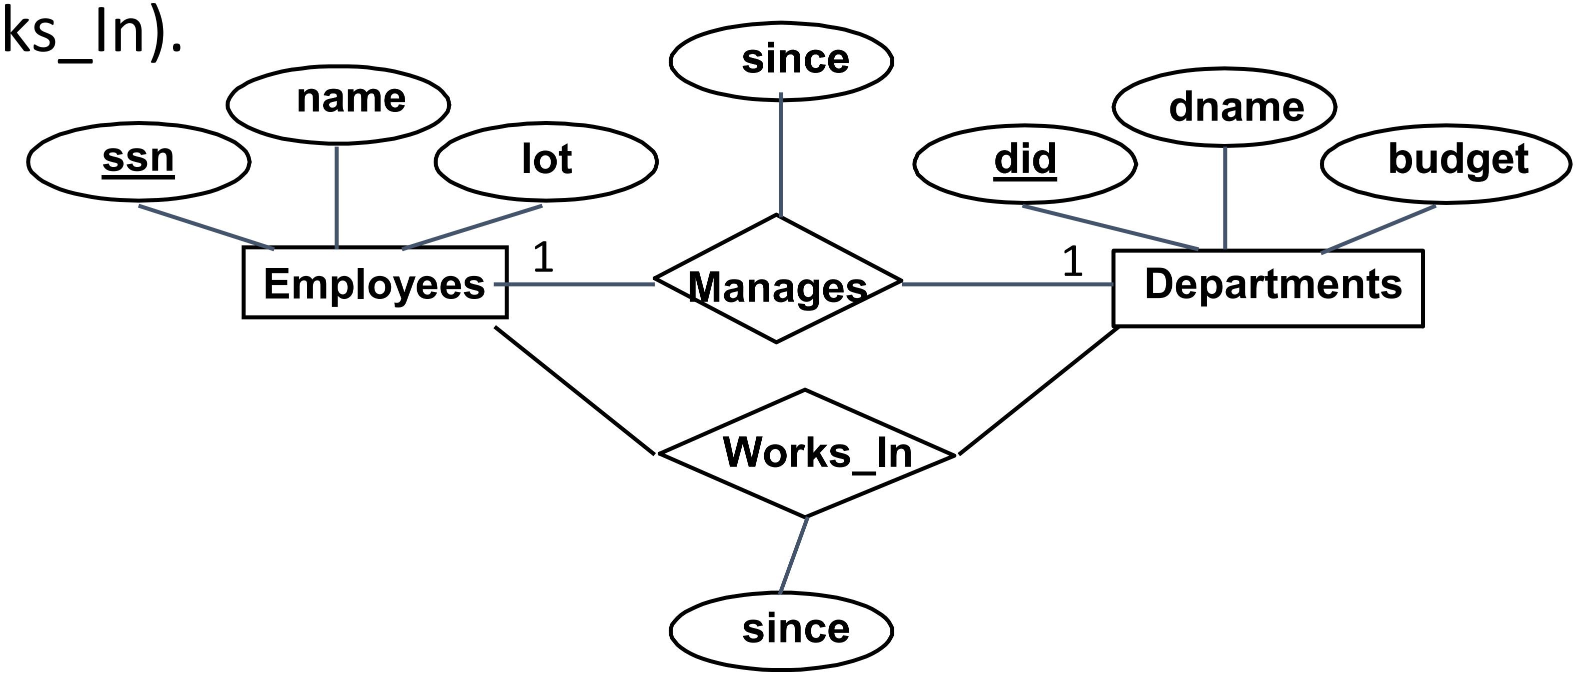

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Works_In).

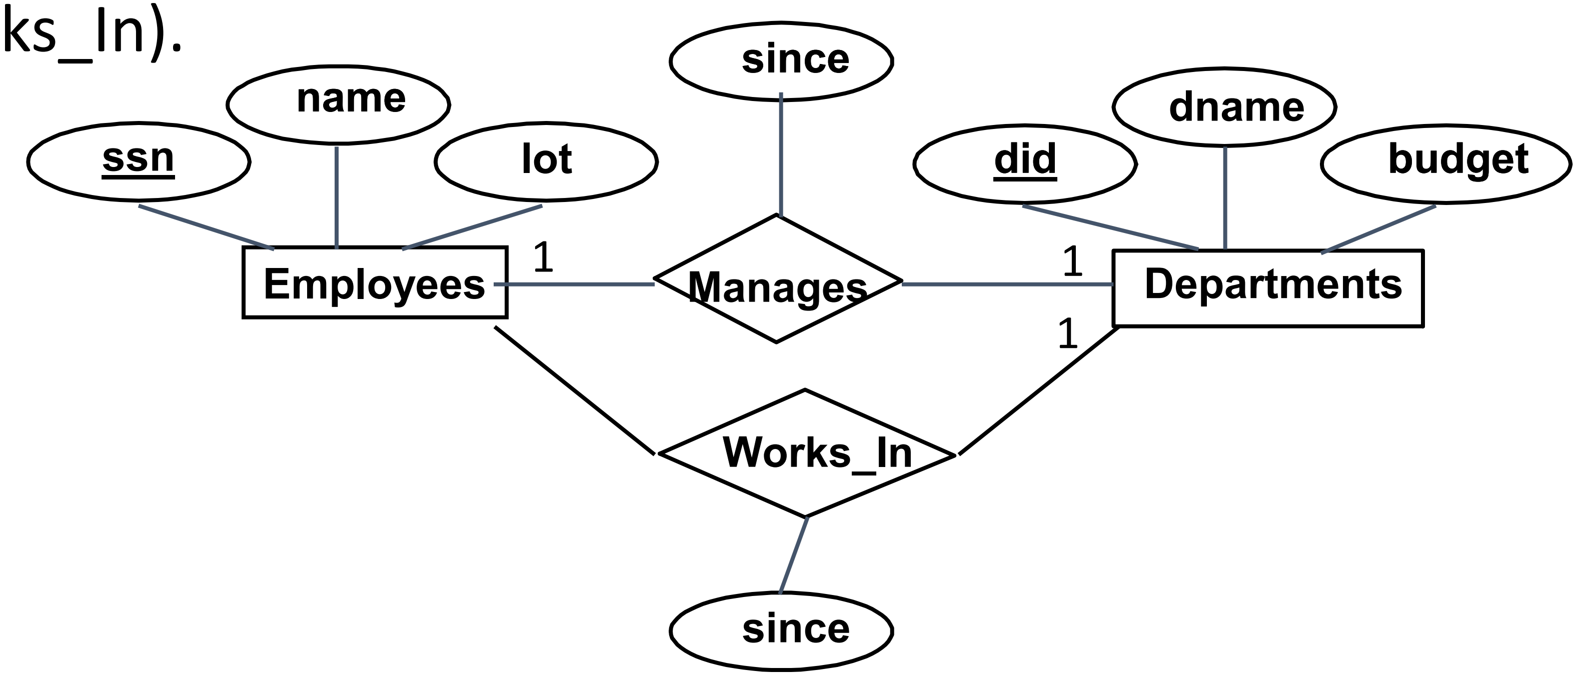

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Works_In).

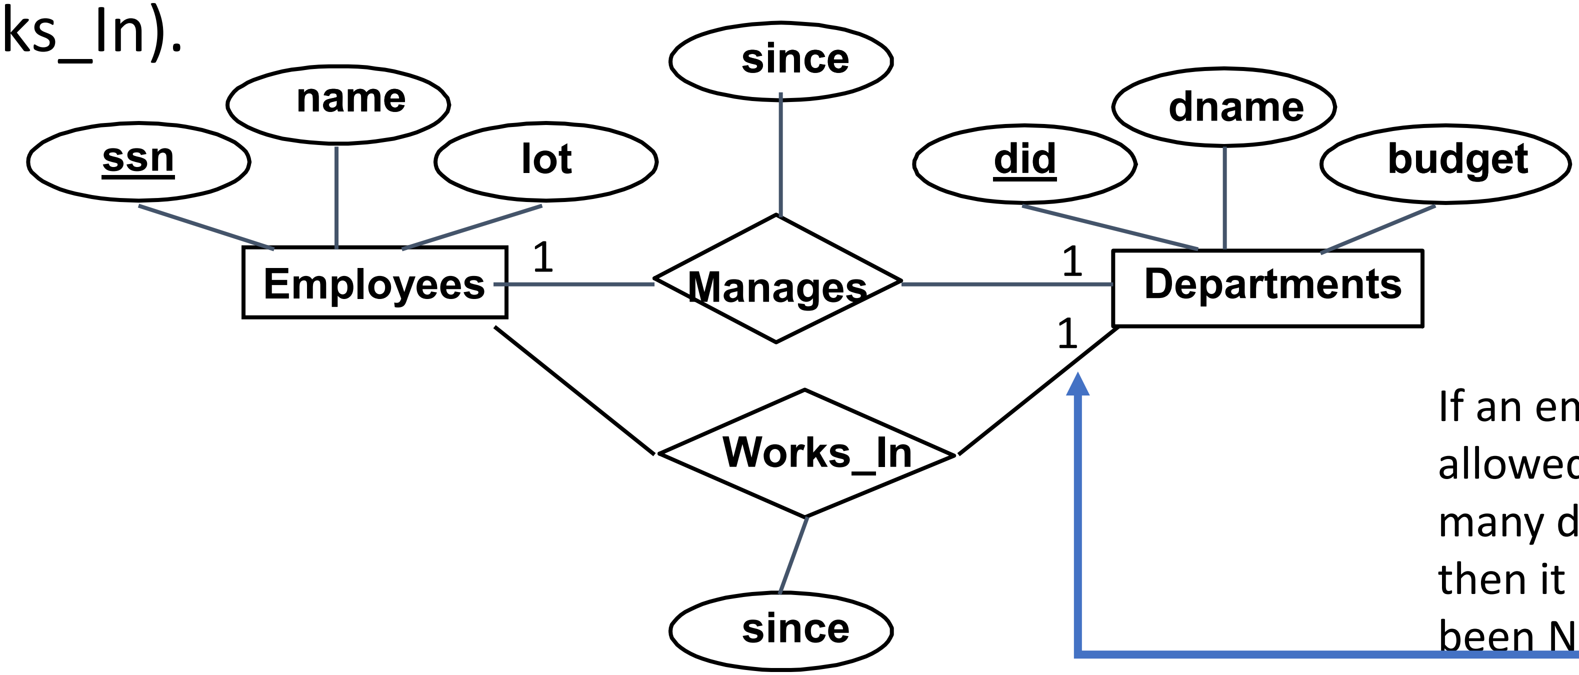

If an employee is allowed to work in many departments, then it must have been N instead of 1.

Let’s fill this.

- Consider the **cardinality** between Employees and Departments (Works_In).

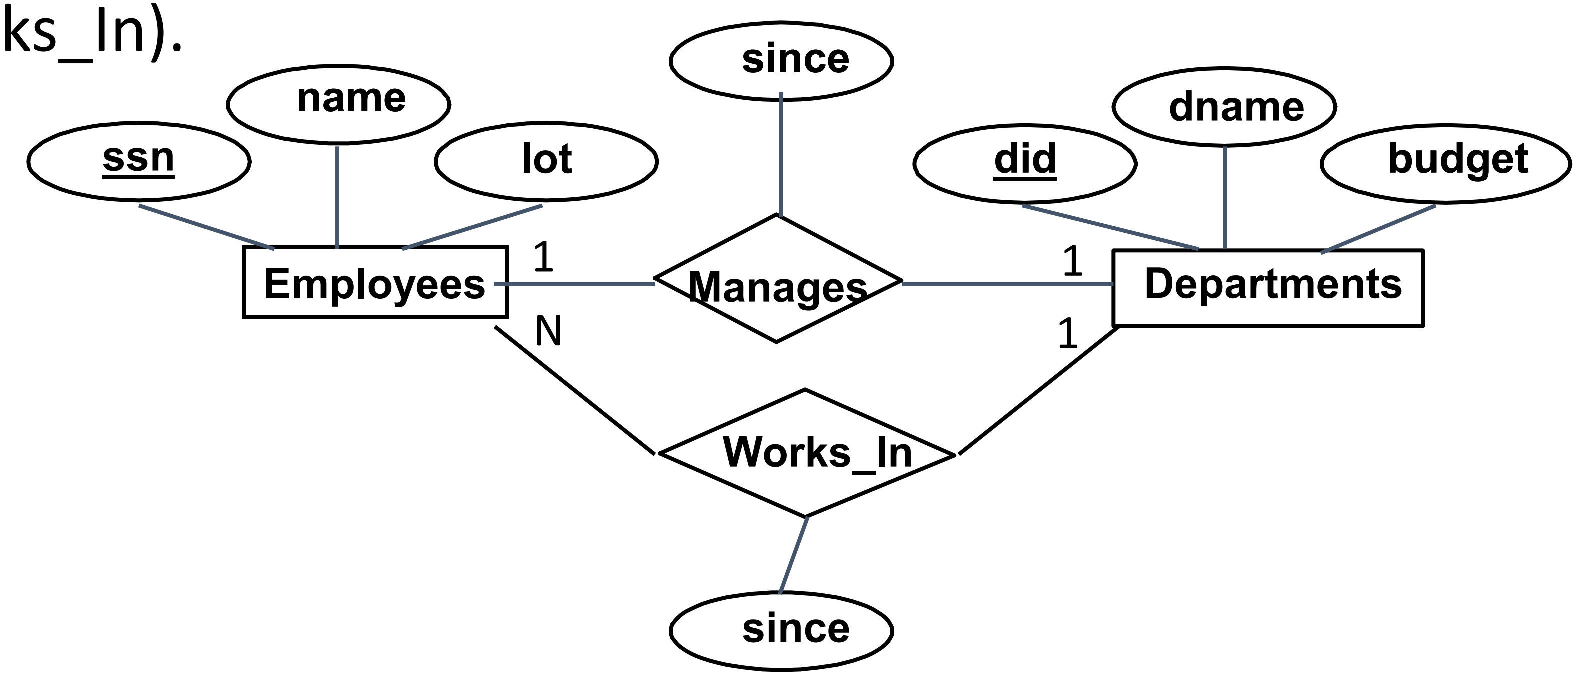

Let’s fill this.

- Consider the **participation constraints** between Employees and Departments (Manages).

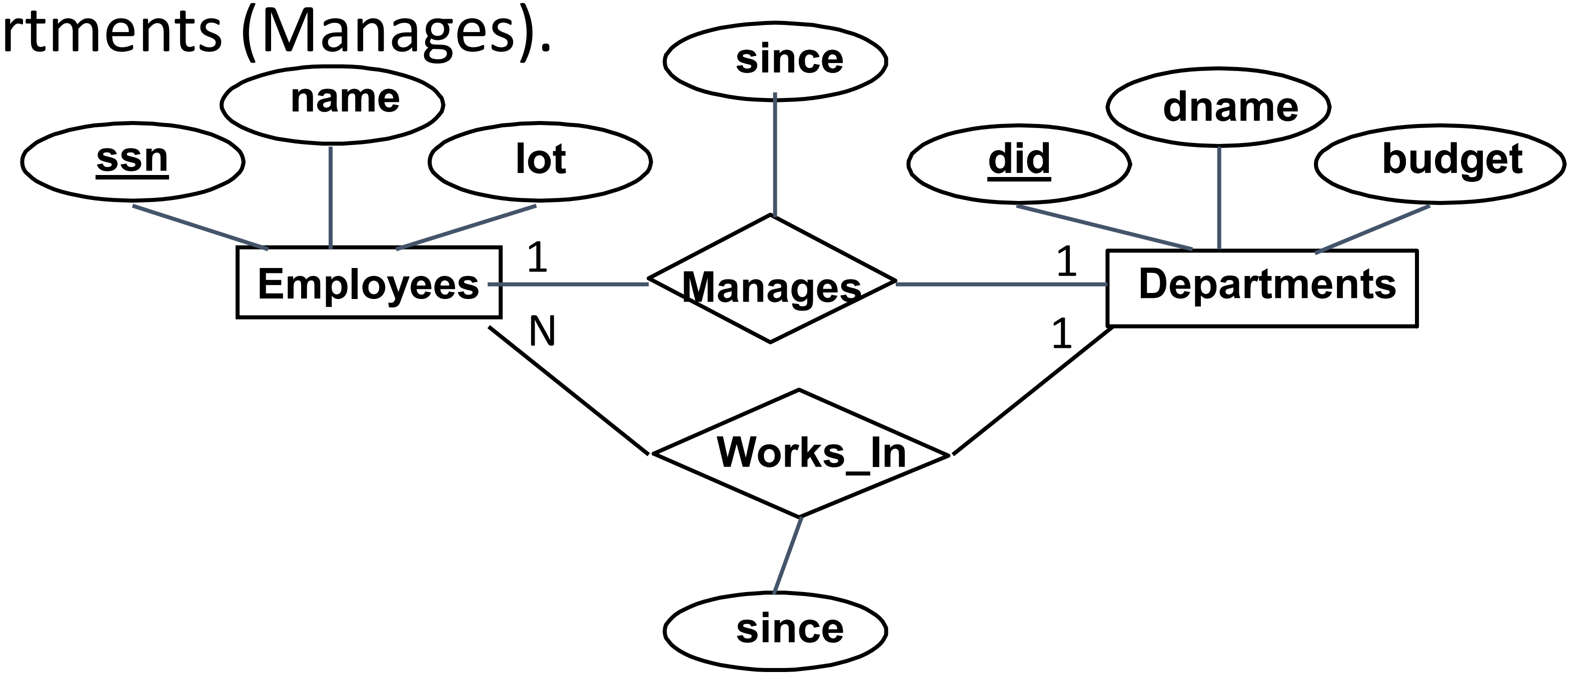

Let’s fill this.

- Consider the **participation constraints** between Employees and Departments (Manages).

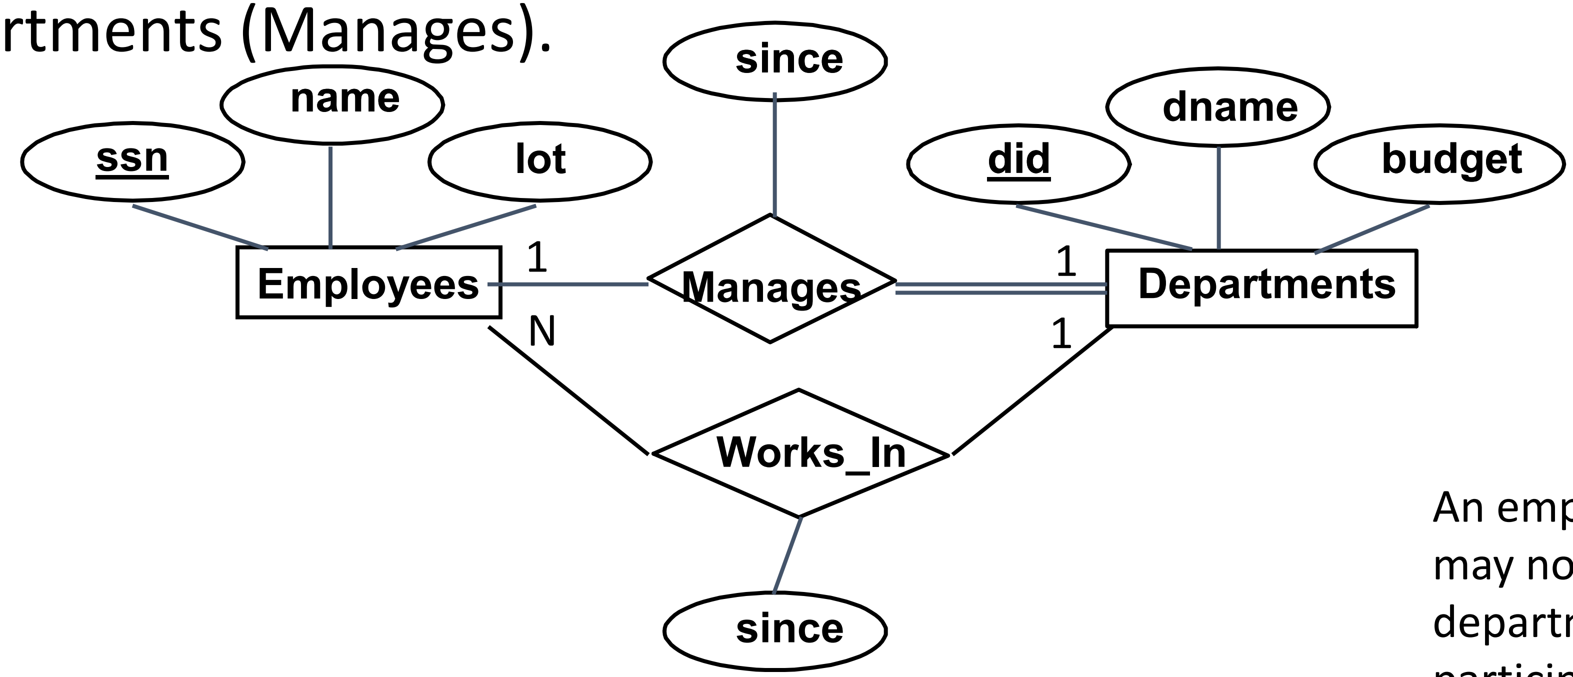

A department must have a manager (total participation).

An employee may or may not manage a department (partial participation).

Let’s fill this.

- Consider the **participation constraints** between Employees and Departments (Works_In).

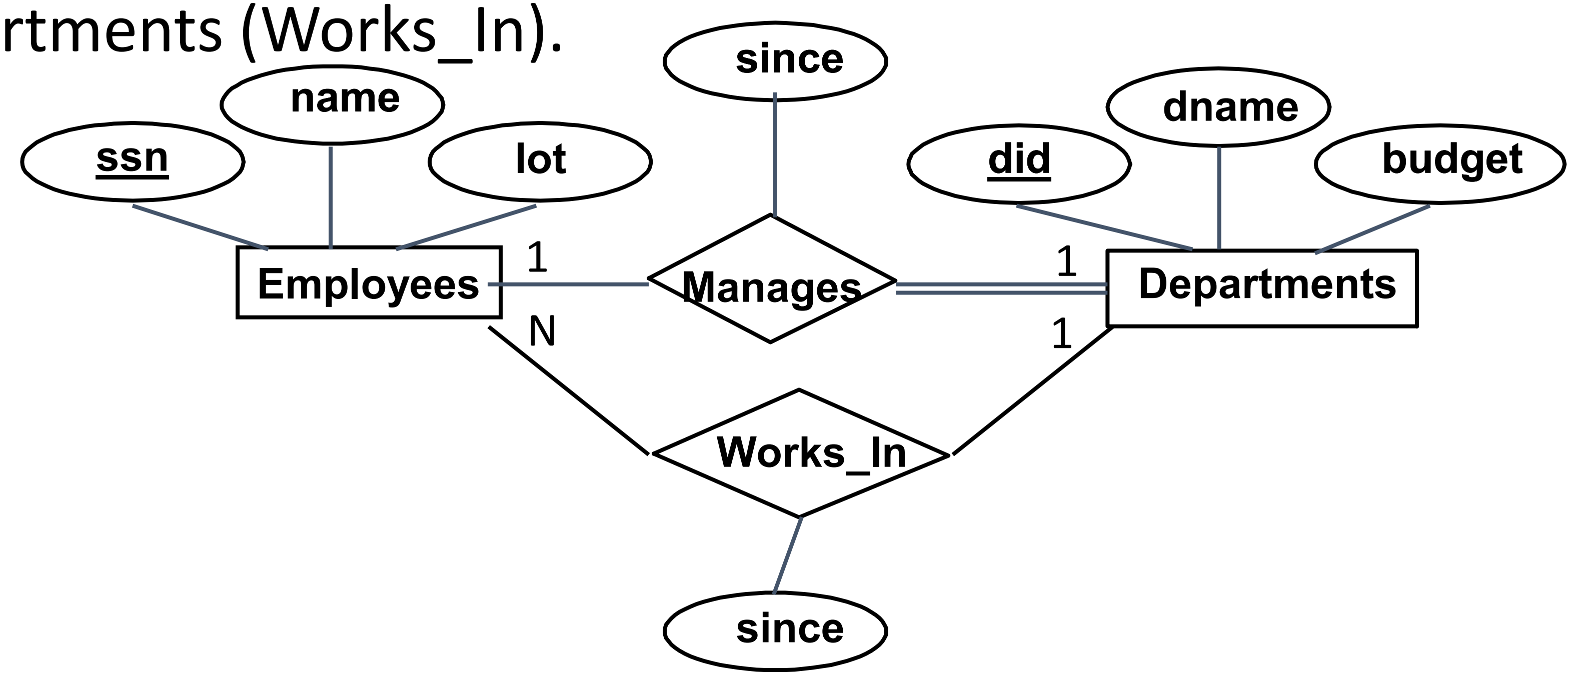

Let’s fill this.

- Consider the **participation constraints** between Employees and Departments (Works_In).

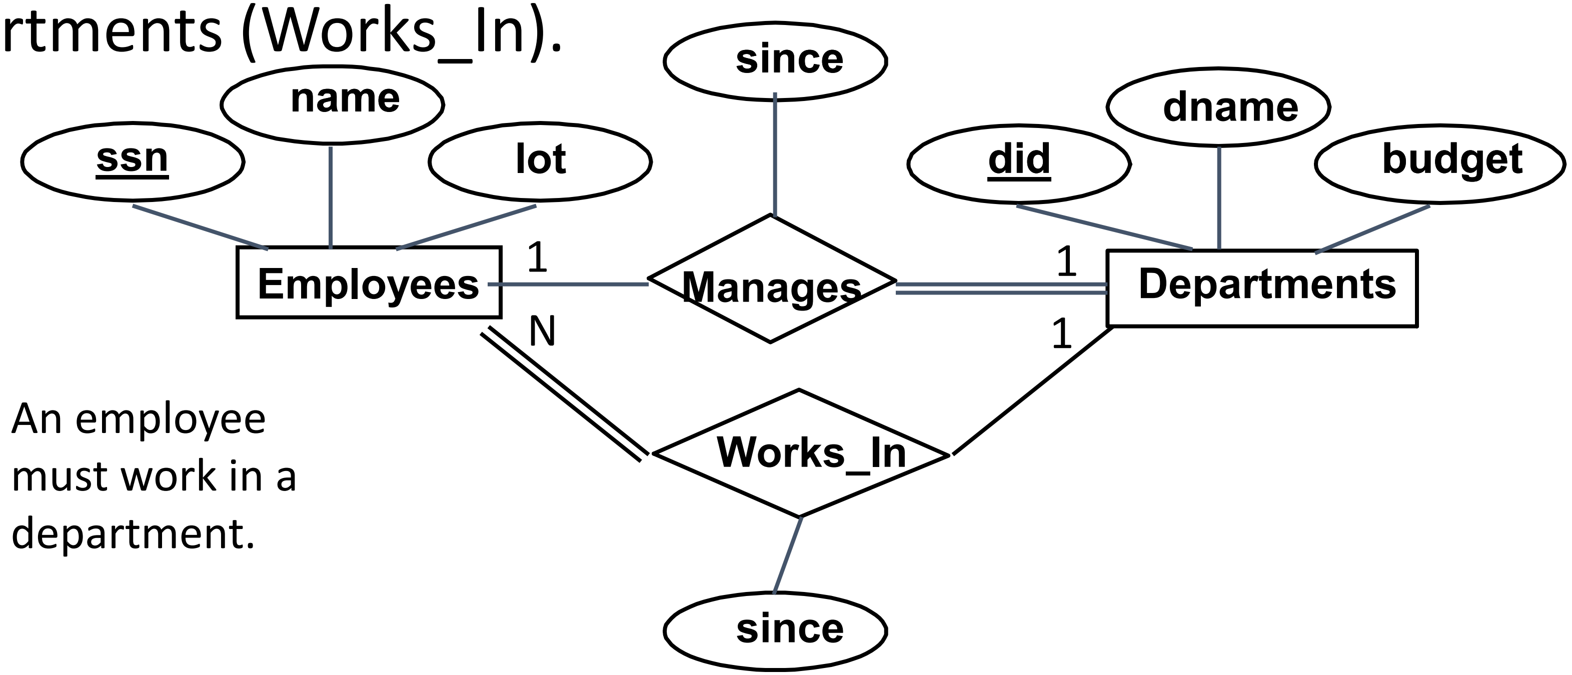

An employee must work in a department.

Let’s fill this.

- Consider the **participation constraints** between Employees and Departments (Works_In).

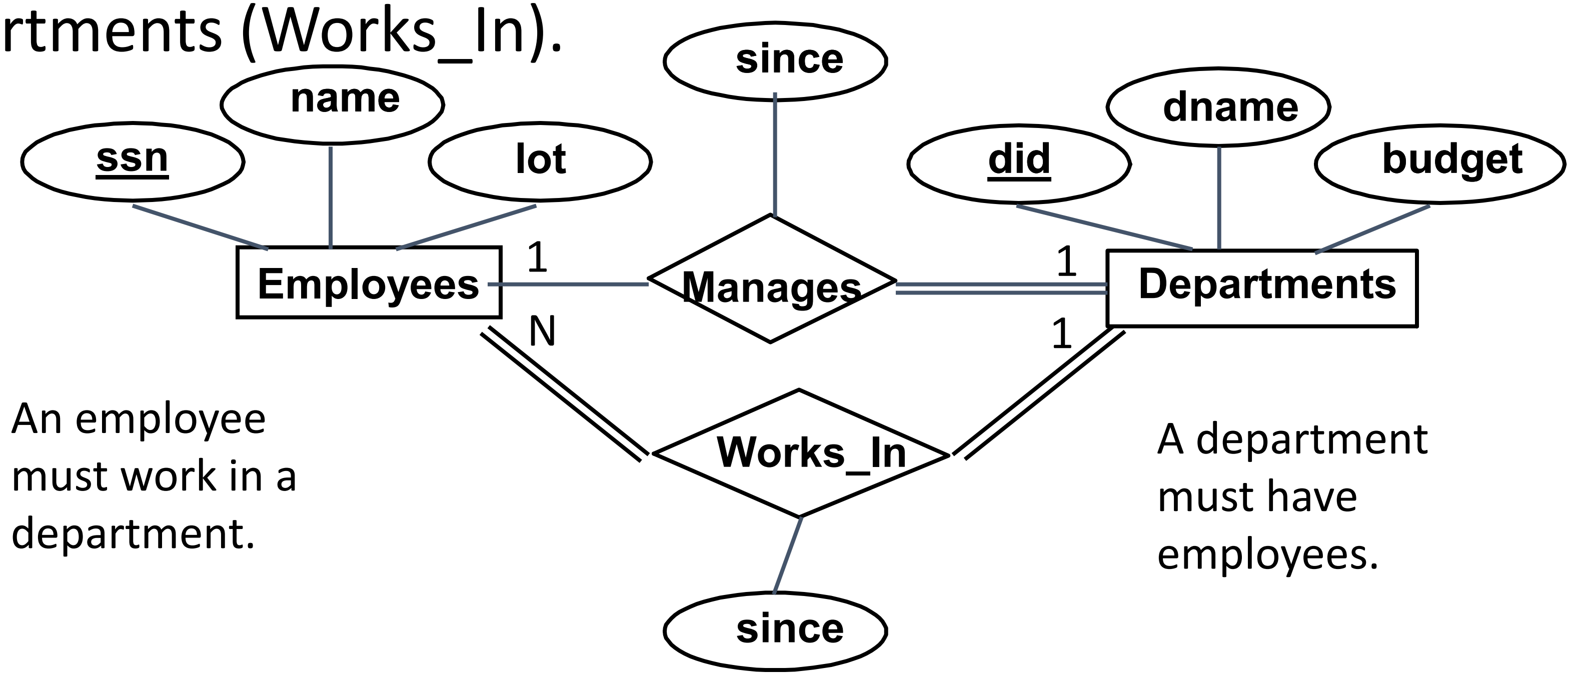

An employee must work in a department.

A department must have employees.

Weak entity set and weak relation sets.

- Not all entity sets have private keys.
- Consider the following ER diagram where a user can Subs. at most once on a given date.

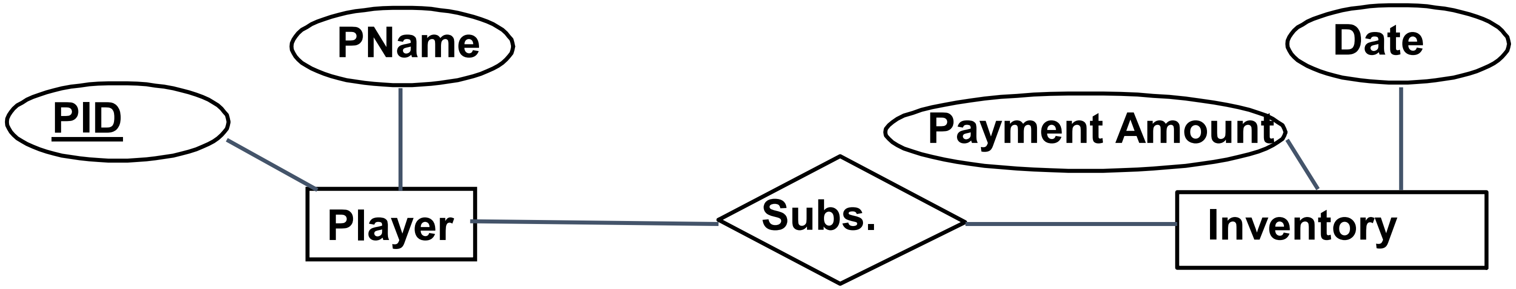

PName	PID
Alex	1
Mark	2
Claudia	3
Summer	4

Payment Amount	Date
1	01/04/1982
12	01/04/1982
1	01/04/1958
3	10/02/1994

Weak entity and weak relation sets.

- Not all entity sets have private keys.
- Consider the following ER diagram where a user can Subs. at most once on a given date.

Week Entity (Inventory) must relate to identifying relation (Subs) with **Total Participation**.

PID.Date (. is a concatenation op.) is a **compound key** for Inventory Weak Entity Set.

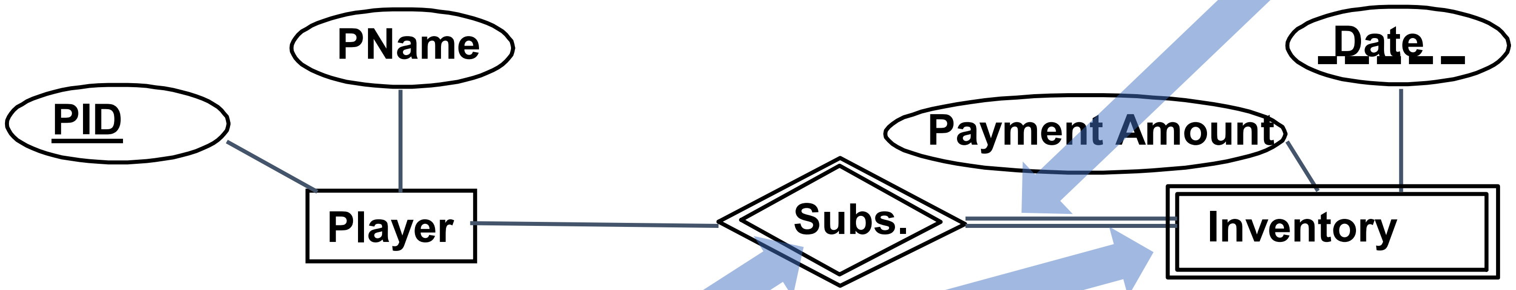

The weak entity attribute used in the compound key is denoted with a dashed underline.

PName	PID
Alex	1
Mark	2
Claudia	3
Summer	4

Week Entity (Inventory) and identifying relation (Subs) are drawn with double lines.

Payment Amount	Date
1	01/04/1982
12	01/04/1982
2	01/04/1958
3	10/02/1994

Massively multiplayer online role-playing game

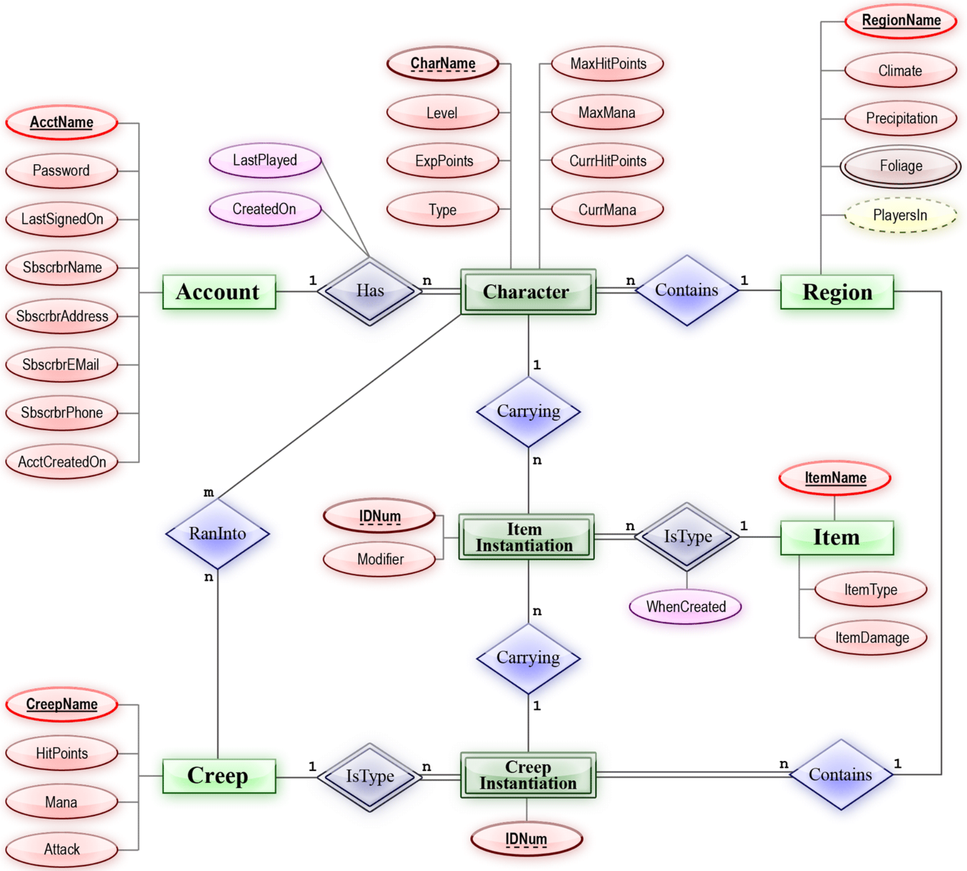

Extended ER

ISA (`is a') Hierarchies

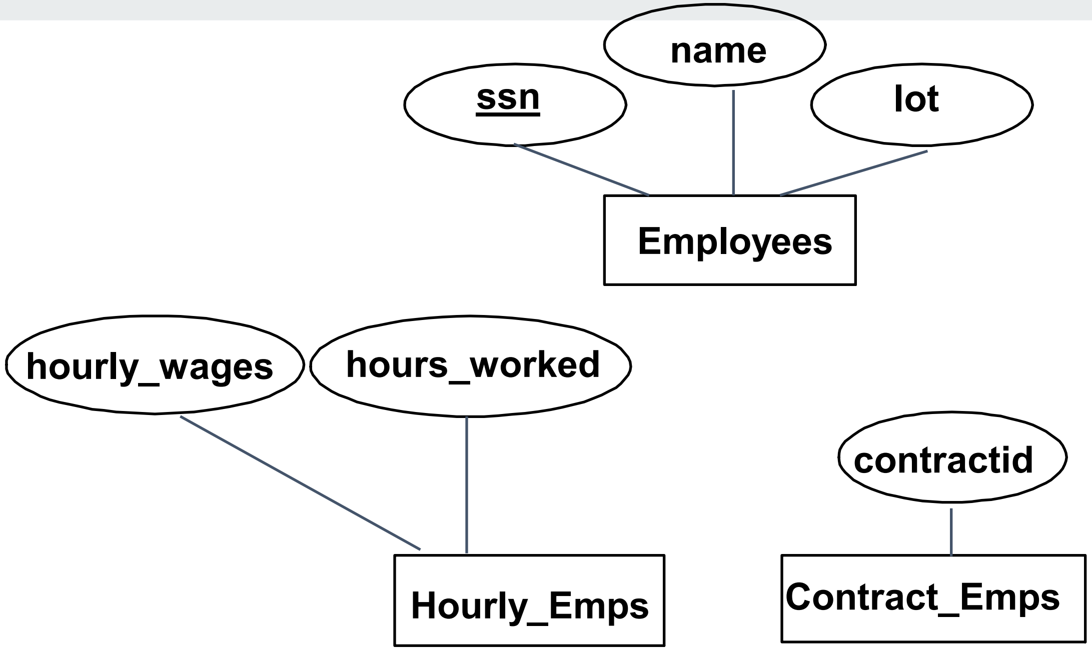

Extended ER

ISA (`is a’) Hierarchies

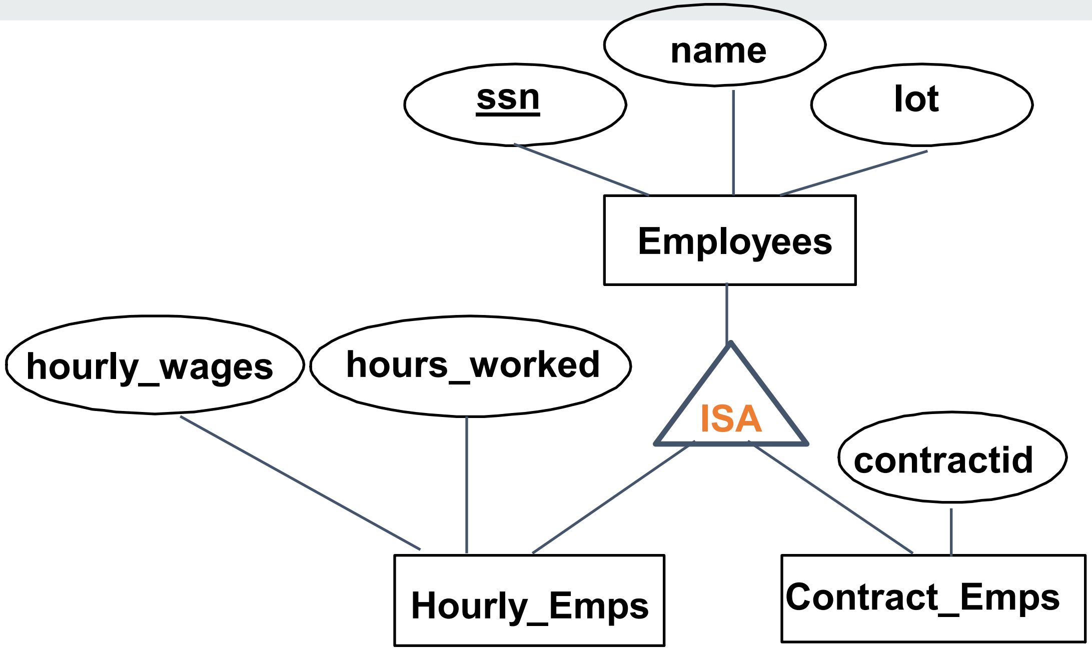

Extended ER
ISA (`is a’) Hierarchies

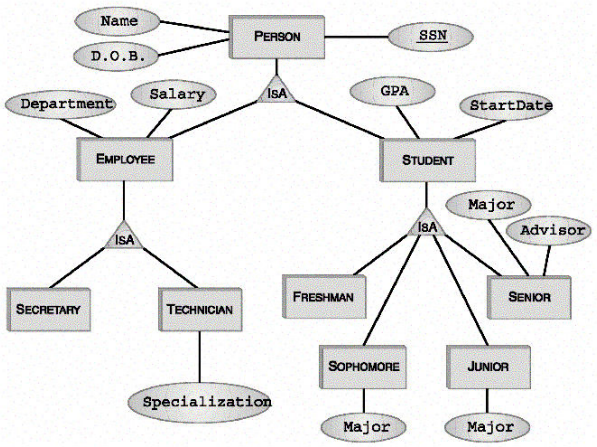

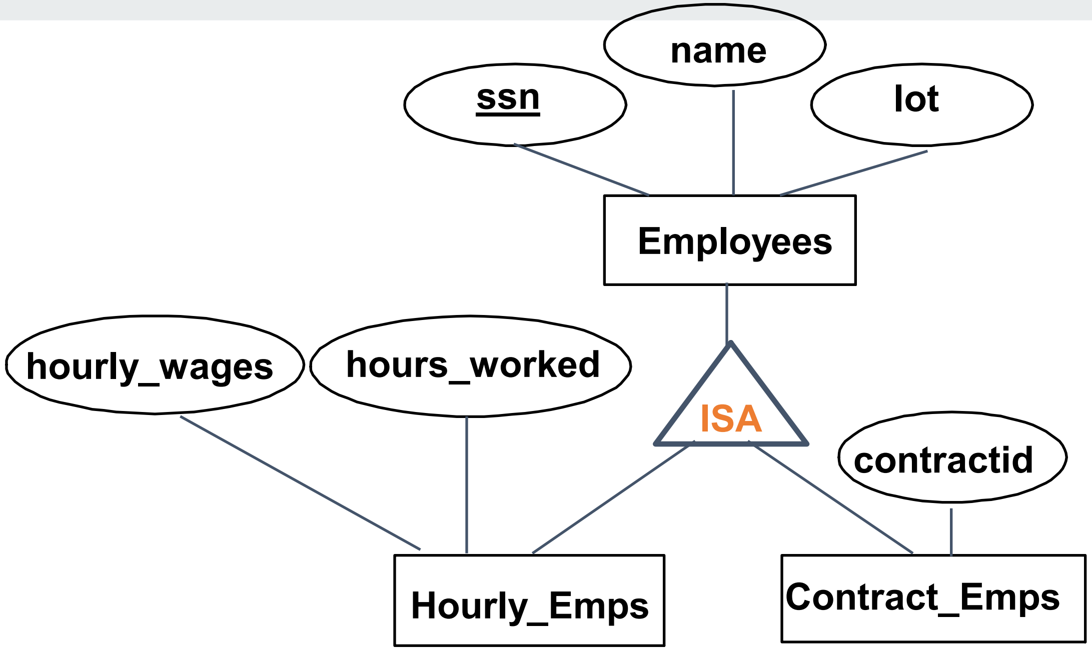

- *Overlap constraints*: Can Uraz be an Hourly Employee as well as a Contract Employee?
- *Covering constraints*: Does every Employee also have to be an Hourly Employee or a Contract Employee?
- Reasons for using ISA:
  - To add descriptive attributes specific to a subclass.
  - To identify entities that participate in a relationship.

Summary

- ER concepts
  - Entity set
    - Attribute
    - Primary key Attribute
    - Multi-valued attribute
    - Derived attribute
    - Weak entity
  - Relationship set
    - Cardinality (1-to-1, 1-to-many, many-to-1, many-to-many)
    - Participation constraint (total participation, partial participation)
    - Weak relation
- There can be several ERDs for a given scenario!

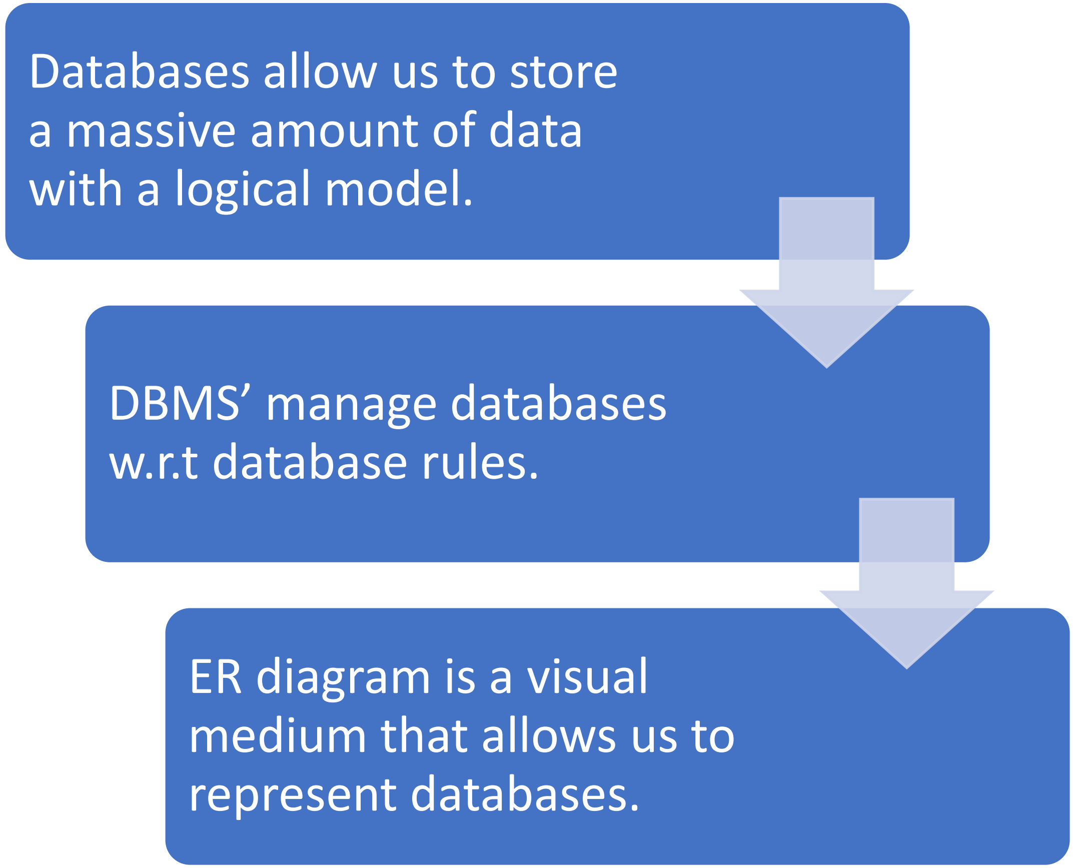

Databases allow us to store a massive amount of data with a logical model.

DBMS’ manage databases w.r.t database rules.

ER diagram is a visual medium that allows us to represent databases.

A nice exam question would be like this.

You are asked to design a database of banks in the UK.

- Customers have names, IN, Account No, Balance
- A customer may have different accounts in different branches
- A branch must have at least one customer
- A customer must have an account in a branch
- A branch has a location and unique location ID
- Draw the ER diagram!