6: Relational Data Model

where we're going

conceptual modeling
technology independent

logical modeling
technology considerations, eliminate redundancies

some terms to remember

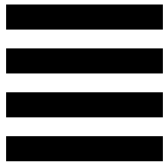

database

collection of relations

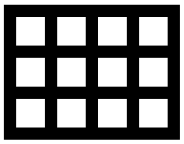

relation

two-dimensional table, consisting of all the tuples

tuple

row in the table, related data values

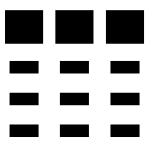

attribute

column in the table

more about relations

relation consists of a heading and a body

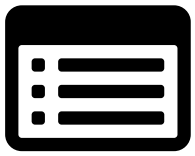

heading

relation schema, schema, intension, relvar

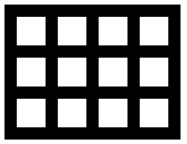

body

extension

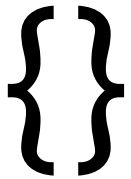

domain

set of possible atomic values for an attribute

more about relations

relation vs. relation schema vs. relational schema

relation

table of data

relation schema

heading of the table

relational schema

all the relation schemas of the database

more about relations

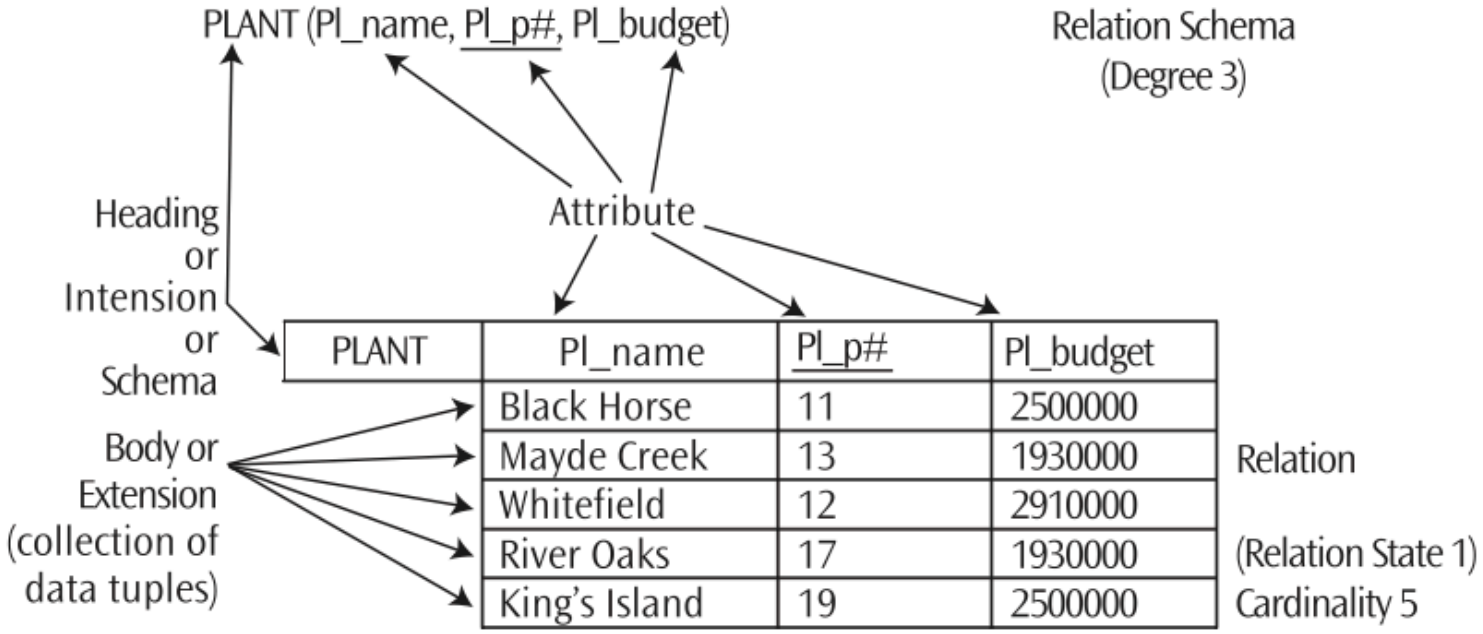

PLANT	Pl_name	Pl_p#	Pl_Budget
	Black Horse	11	2370000
	Whitefield	12	2110000
	River Oaks	17	1930000
	King's Island	19	2110000

Relation
(Relation State 2)
Cardinality 4

Note: Both Pl_name and Pl_p# are candidate keys of PLANT, while Pl_p# is the *chosen* primary key.

more about relations

a relation is a two dimensional table with a heading and a body

UNIVERSAL RELATION SCHEMA

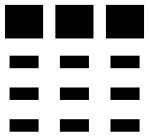

attributes

atomic, have unique names, values come from same domain

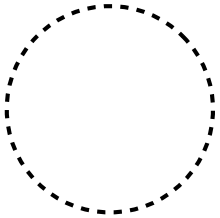

derived attributes

not captured in relation schema

tuples

order of arrangement is immaterial, must have unique identifier

naming convention

for attributes

up to 3
letter prefix

1st letter capitalized

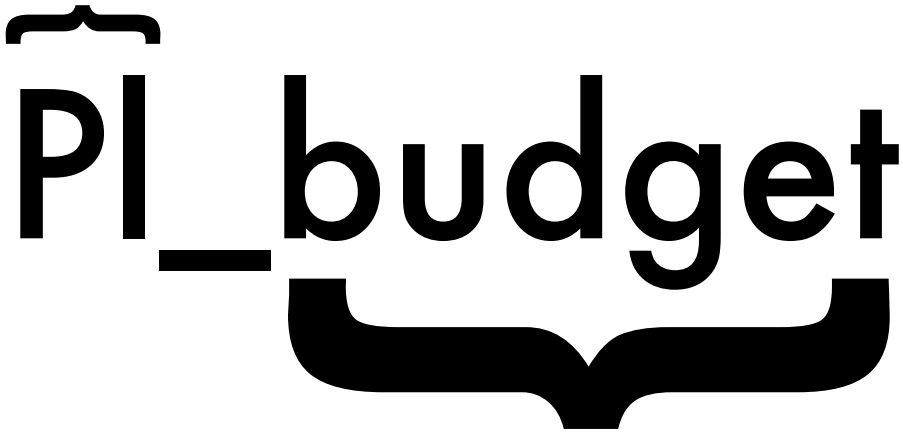

only lowercase letters, a pound sign (#), and underscore characters, and corresponds to name of the attribute in conceptual data model

data integrity constraints

rules that govern the behavior of data at all times in a database

inherent model-based

driven by modeling grammar

schema or declarative-based

domain, key, relationship structural, entity integrity,
referential integrity, functional dependency

semantic integrity constraints

application-based or DBMS-based
procedural constraints

types of data integrity constraints

state constraints

declarative and procedural constraints that every valid state of a database must satisfy

valid relationship statuses: Single, Married, Widowed, Divorced

transition constraints

procedural constraints that defined legal transitions of state

Married to Divorced or Widowed, but not Single

concept of unique identifiers

superkey

set of one or more attributes, which taken collectively, uniquely identifies a tuple of a relation {uniqueness property}

candidate key

a superkey with no proper subset that uniquely identifies a tuple of a relation {uniqueness property + irreducibility}

primary key

a candidate key with no missing values for the constituent attributes {uniqueness property + irreducibility + entity integrity constraint}

alternate key

any candidate key that is not serving the role of the primary key

unique identifiers visualized

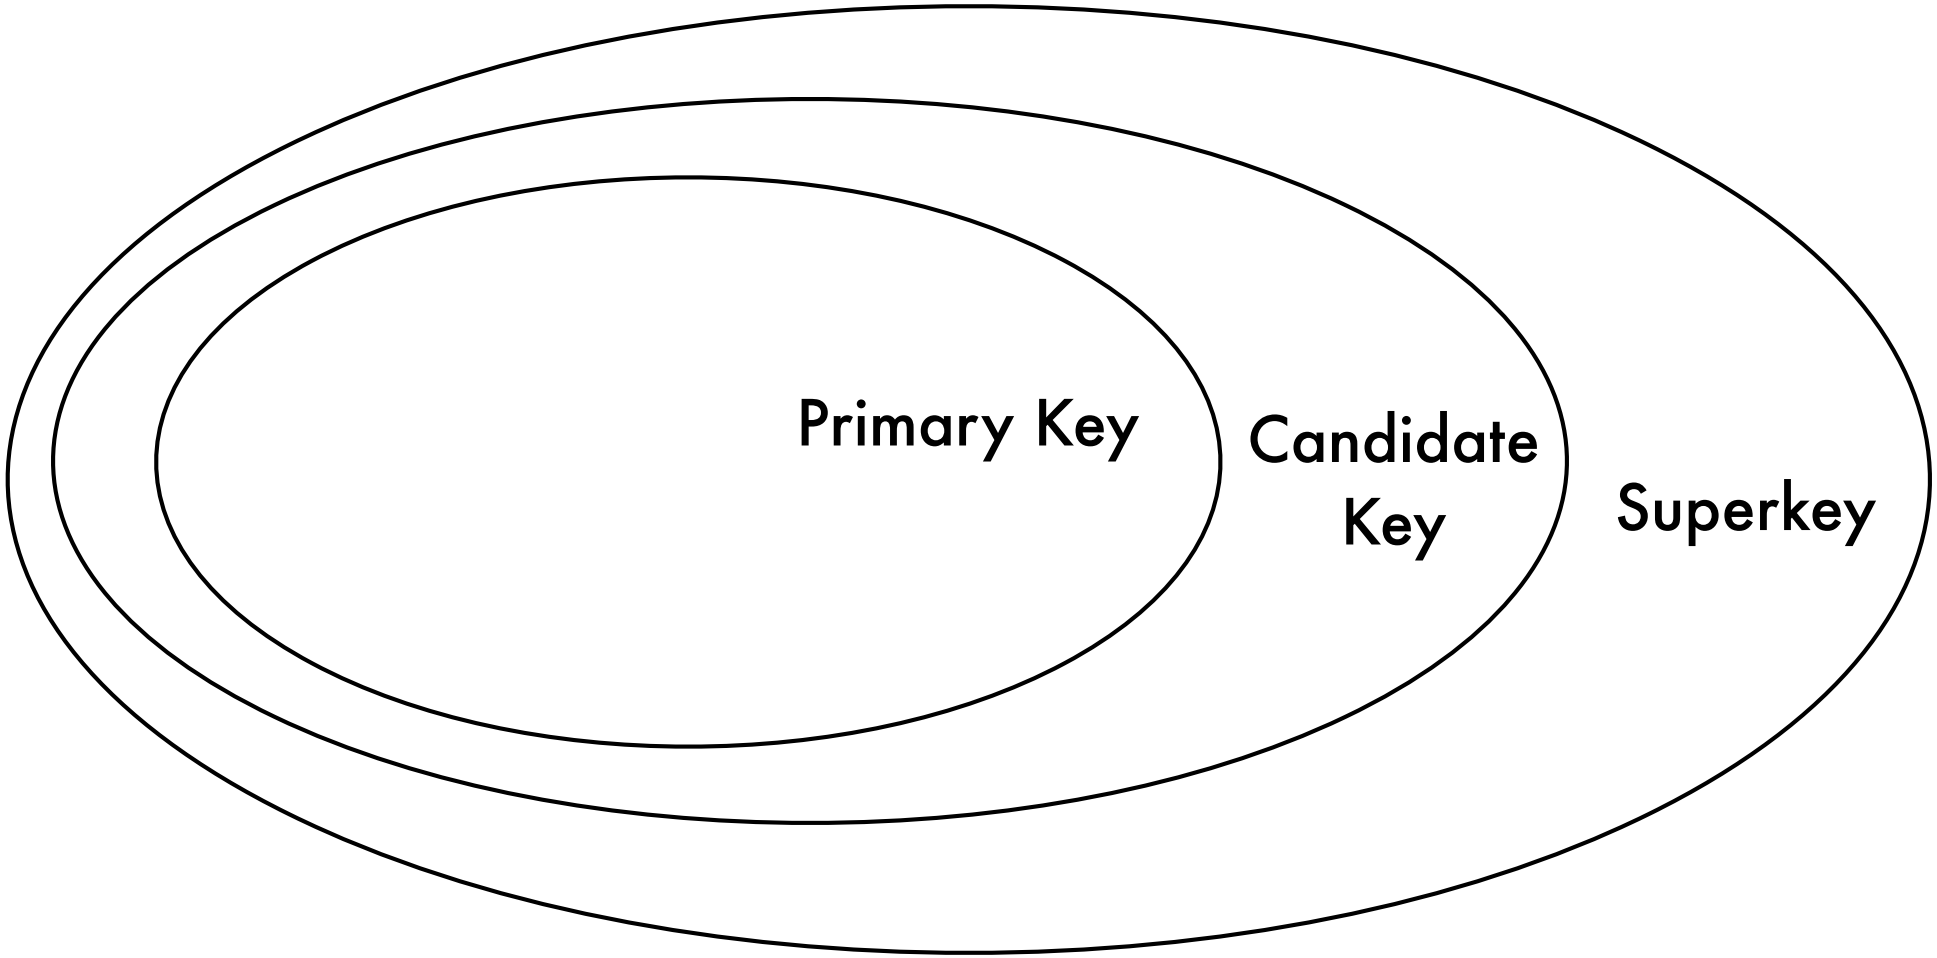

subsets and proper subsets

Attribute(s)	Subset	Proper Subset(s)
Pl_name	Pl_name	<none>
Pl_p#	Pl_p#	<none>
Pl_budget	Pl_budget	<none>
{Pl_name, Pl_p#}	{Pl_name, Pl_p#} Pl_name Pl_p#	Pl_name Pl_p#
{Pl_name, Pl_budget}	{Pl_name, Pl_budget} Pl_name Pl_budget	Pl_p# Pl_budget
{Pl_name, Pl_p#, Pl_budget}	{Pl_name, Pl_p#, Pl_budget} {Pl_name, Pl_budget} {Pl_name, Pl_p#} {Pl_p#, Pl_budget} Pl_name Pl_p# Pl_budget	{Pl_name, Pl_budget} {Pl_name, Pl_p#} {Pl_p#, Pl_budget} Pl_name Pl_p# Pl_budget

Note: A set is a subset of itself, however, a set is not a proper subset of itself.

a sample relation instance

pg. 249

PRESCRIPTION-A

Rx_rx#	Rx_pat#	Rx_medcode	Rx-dosage
A100	7642	PCN	3
A103	4678	TYL	2
A102	4772	CLR	2
A101	6742	ASP	2
A104	4772	ZAN	3
A105	7456	CLR	2
A107	2222	TYL	2
A106	4772	VAL	2
A108	7384	CLR	3
A109	7384	ZAN	2
A110	7642	VAL	2

concept of key/non-key attributes

key attribute

any attribute that is a proper subset of a candidate key

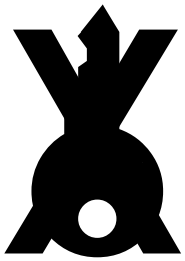

non-key attribute

any attribute that is not a subset of a candidate key

an attribute is either a **key attribute**, a **non-key attribute** or a **candidate key**

what is Rx_rx#?

a candidate key

concept of key/non-key attribute

PRESCRIPTION-B

Rx_rx#	Rx_pat#	Rx_medcode	Rx_dosage
B100	7642	PCN	3
B103	4678	TYL	2
B102	4772	CLR	2
B101	6742	ASP	2
B102	4772	ZAN	**2**
B105	7456	CLR	2
B107	2222	TYL	2
B106	4772	VAL	2
B108	7384	CLR	3
B109	7384	ZAN	2
B100	7642	VAL	2

concept of key/non-key attribute

PRESCRIPTION (Rx_rx#, Rx_pat#, Rx_medcode, Rx_dosage)

	PRESCRIPTION-A		PRESCRIPTION-B	
	Superkey	Candidate Key	Superkey	Candidate Key
Rx_rx#	Yes	Yes	No	No
Rx_pat#	No	No	No	No
Rx_medcode	No	No	No	No
Rx_dosage	No	No	No	No
Rx_rx#, Rx_pat#	Yes	No	No	No
Rx_rx#, Rx_medcode	Yes	No	Yes	Yes
Rx_rx#, Rx_dosage	Yes	No	No	No
Rx_pat#,Rx_medcode	Yes	Yes	Yes	Yes
Rx_pat#, Rx_dosage	No	No	No	No
Rx_medcode, Rx_dosage	No	No	No	No
Rx_rx#, Rx_pat#, Rx_medcode	Yes	No	Yes	No
Rx_rx#, Rx_pat#, Rx_dosage	Yes	No	No	No
Rx_rx#, Rx_medcode, Rx_dosage	Yes	No	Yes	No
Rx_pat#, Rx_medcode, Rx_dosage	Yes	No	Yes	No
Rx_rx#, Rx_pat#, Rx_medcode, Rx_dosage	Yes	No	Yes	No

referential integrity constraint

specified between two relation schemas

foreign key constraint

specific referential integrity constraint

example of referencing a relation schema:
attribute set A2 in relation schema R2
shares the same domain
with a candidate key (A1) of another relation schema R1

R2 is the referencing relation, and R1 is the referenced relation
A2 is the foreign key and A1 is the referenced attribute

expressed as an inclusion dependency:
$R2.\{A2\} \subseteq R1.\{A1\}$

foreign key constraint

source schema

Box 1

Bearcat Incorporated is a manufacturing company that has several plants in the northeastern part of the United States. *These plants are responsible for leading different projects that the company might undertake, depending on a plants' function. A certain plant might even be associated with several projects but a project is always under the control of just one plant. Some plants do not undertake any projects at all.* If a plant is closed down, the projects undertaken by that plant cannot be canceled. The project assignments from a closed plant must be temporarily removed in order to allow the project to be transferred to another plant.

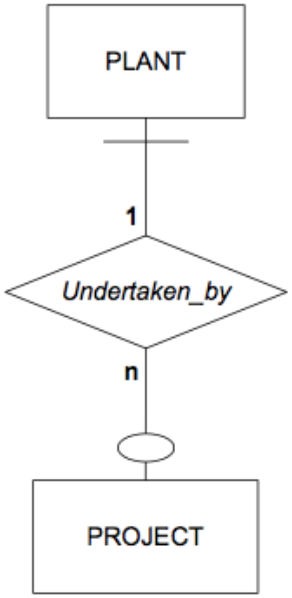

example of a foreign key

version 1

PLANT	Pl_name	Pl_p#	Pl_budget
	Black Horse	11	1230000
	Mayde Creek	13	1930000
	Whitefield	12	2910000
	River Oaks	17	1930000
	King's Island	19	2500000
	Ashton	15	2500000

PROJECT	Prj_name	Prj_n#	Prj_location	*Prj_pl_p#*
	Solar Heating	41	Sealy	11
	Lunar Cooling	17	Yoakum	17
	Synthetic Fuel	29	Salem	17
	Nitro-Cooling	23	Parthi	12
	Robot Sweeping	31	Ponca City	11
	Robot Painting	37	Yoakum	19
	Ozone Control	13	Parthi	19

Version 1

Note: PROJECT.**Prj_pl_p#** is the foreign key referencing PLANT.**Pl_p#**, the primary key of PLANT.

example of a foreign key

version 2

PLANT	Pl_name	Pl_p#	Pl_budget
	Black Horse	11	1230000
	Mayde Creek	13	1930000
	Whitefield	12	2910000
	River Oaks	17	1930000
	King's Island	19	2500000
	Ashton	15	2500000

PROJECT	Prj_name	Prj_n#	Prj_location	*Prj_pl_name*
	Solar Heating	05	Sealy	Black Horse
	Lunar Cooling	17	Yoakum	River Oaks
	Synthetic Fuel	29	Salem	River Oaks
	Nitro-Cooling	23	Parthi	Whitefield
	Robot Sweeping	31	Ponca City	Black Horse
	Robot Painting	37	Yoakum	King's Island
	Ozone Control	13	Parthi	King's Island

Version 2

Note: PROJECT.**Prj_pl_name** is the foreign key referencing PLANT.**Pl_name**, a candidate key of PLANT.

naming convention

for foreign keys

prefix used for the
attribute names in the
referencing relation schema

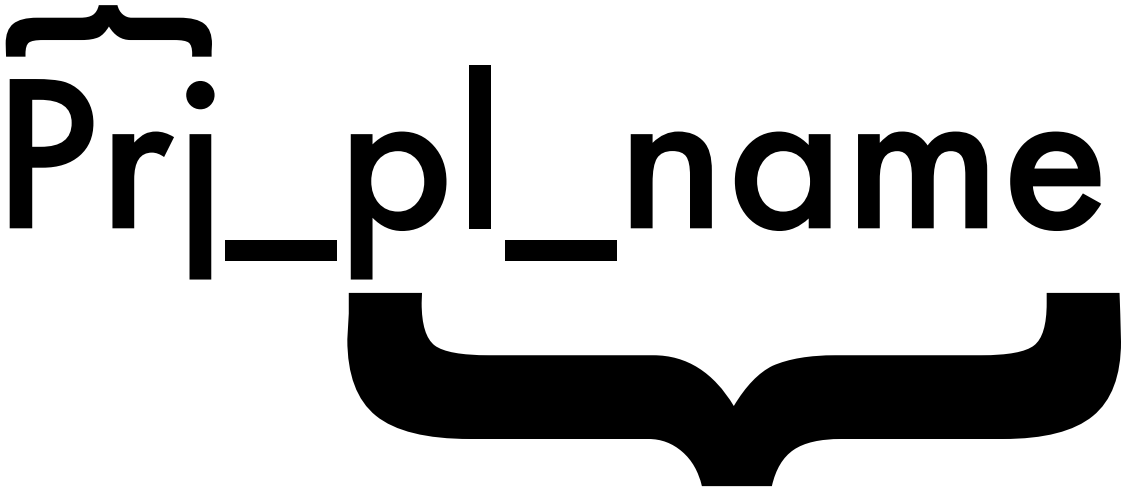

the referenced attribute name

an introduction to relational algebra

σ selection	∩ intersection
π projection	- difference
U union	* natural join

example relational schema

pg. 255

AW_PLANT (Aw_pl_name, Aw_pl_p#, Aw_pl_budget)

TX_PLANT (Tx_pl_name, Tx_pl_p#, Tx_pl_budget)

PROJECT (Prj_name, Prj_p#, Prj_location, *Prj_aw_pl_p#*)

foreign key is italicized

relation instances of the example relational schema

AW_PLANT	**Aw_pl_name**	**Aw_pl_p#**	**Aw_pl_budget**
	Black Horse	11	1230000
	Mayde Creek	13	1930000
	Whitefield	12	2910000
	River Oaks	17	1930000
	King's Island	19	2500000
	Ashton	15	2500000

TX_PLANT	**Tx_pl_name**	**Tx_pl_p#**	**Tx_pl_budget**
	Southern Oaks	16	1230000
	River Oaks	17	1930000
	Kingwood	18	1930000

PROJECT	**Prj_name**	**Prj_n#**	**Prj_location**	***Prj_aw_pl_p#***
	Solar Heating	41	Sealy	11
	Lunar Cooling	17	Yoakum	17
	Synthetic Fuel	29	Salem	17
	Nitro-Cooling	23	Parthi	12
	Robot Sweeping	31	Ponca City	11
	Robot Painting	37	Yoakum	19
	Ozone Control	13	Parthi	19

selection

select a horizontal subset of tuples that satisfy a selection condition from a relation

σ<selection condition>(R)

filters horizontally

Which award-winning plants have a budget that exceeds $2,000,000?

σ<R_aw_pl_budget>2000000> (AW_PLANT)

R_aw_pl_name	R_aw_pl_p#	R_aw_pl_budget
Whitefield	12	2910000
King's Island	19	2500000
Ashton	15	2500000

projection

selects certain attributes from the relation

π<attribute list>(R)

filters vertically

What is the plant number and budget of each award-winning plant?

π<R_aw_pl_p#, R_aw_pl_budget> (AW_PLANT)

R_aw_pl_p#	R_aw_pl_budget
11	1230000
13	1930000
12	2910000
17	1930000
19	2500000
15	2500000

projection vs. selection visualized

selection

projection

set theoretic operators

Union

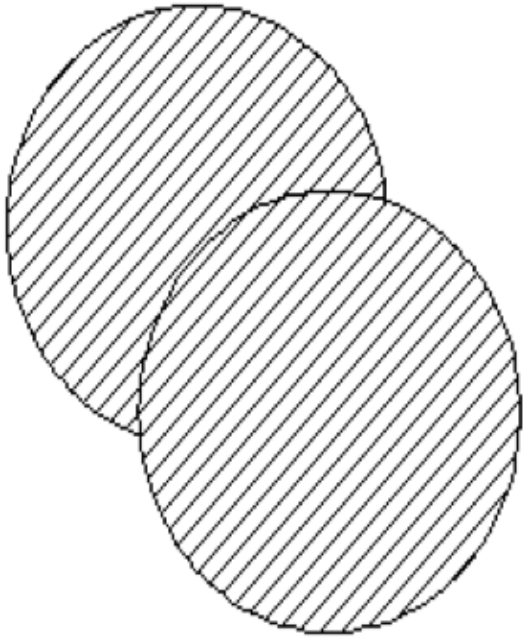

Intersection

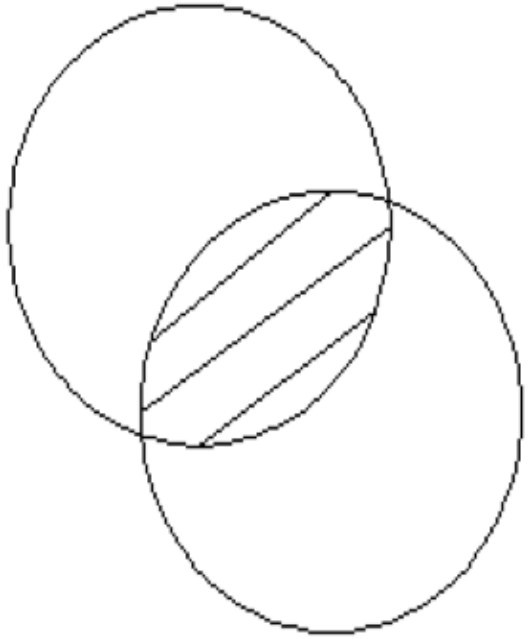

Difference

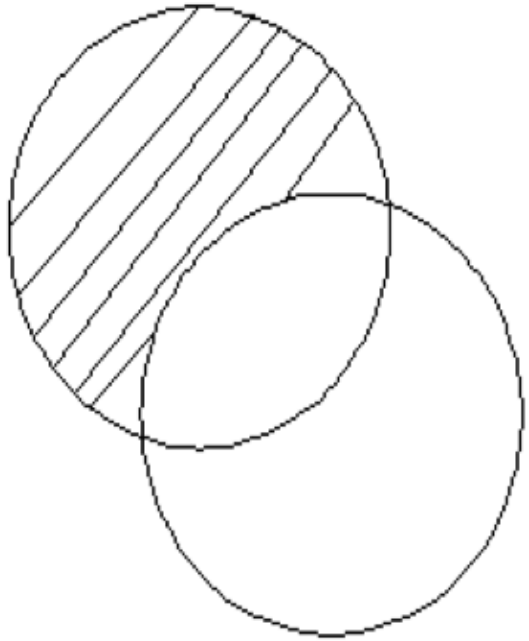

union compatible relations

same degree

(same number of attributes)

same domain

(for each attribute)

union

all the tuples in both relations, duplicates removed

R1 U R2

What plants are located either in Texas or are award-winning plants?

AW_PLANT U TX_PLANT

R_aw_pl_name	R_aw_pl_p#	R_aw_pl_budget
Black Horse	11	1230000
Mayde Creek	13	1930000
Whitefield	12	2910000
River Oaks	17	1930000
King's Island	19	2500000
Ashton	15	2500000
Southern Oaks	16	1230000
Kingwood	18	1930000

intersection

tuples that are found in both relations

R1 ∩ R2

Which award-winning plants are located in Texas?

AW_PLANT ∩ TX_PLANT

R_aw_pl_name	R_aw_pl_p#	R_aw_pl_budget
River Oaks	17	1930000

difference

tuples that are in one relation but not in the other

R1 - R2

Which Texas plants are not award-winning plants?

TX_PLANT - AW_PLANT

R_aw_pl_name	R_aw_pl_p#	R_aw_pl_budget
Southern Oaks	16	1230000
Kingwood	18	1930000

natural join

join two tuples based on a matching attribute

R1 * R2

What are the projects at award-winning plants?

PROJECT * (π<R_aw_pl_p#, R_aw_pl_name>AW_PLANT)

R_prj_name	R_prj_n#	R_prj_location	R_prj_pl_p#	R_aw_pl_name
Solar Heating	41	Sealy	11	Black Horse
Lunar Cooling	17	Yoakum	17	River Oaks
Synthetic Fuel	29	Salem	17	River Oaks
Nitro-Cooling	23	Parthi	12	Whitefield
Robot Sweeping	31	Ponca City	11	Black Horse
Robot Painting	37	Yoakum	19	King's Island
Ozone Control	13	Parthi	19	King's Island

only one copy of the matching attribute (pl_p#) in resulting relation