



**CS425 – Summer 2016**  
**Jason Arnold**  
**Chapter 2: Intro to Relational Model**

**Modifies from:**

**Database System Concepts, 6<sup>th</sup> Ed.**

©Silberschatz, Korth and Sudarshan

See [www.db-book.com](http://www.db-book.com) for conditions on re-use



# Textbook: Chapter 2



# Example of a Relation

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>
10101	Srinivasan	Comp. Sci.	65000
12121	Wu	Finance	90000
15151	Mozart	Music	40000
22222	Einstein	Physics	95000
32343	El Said	History	60000
33456	Gold	Physics	87000
45565	Katz	Comp. Sci.	75000
58583	Califieri	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

attributes  
(or columns)

tuples  
(or rows)



# Attribute Types

- n The set of allowed values for each attribute is called the **domain** or **data type** of the attribute
- n Attribute values are (normally) required to be **atomic**; that is, indivisible
  - | E.g., integer values
  - | E.g., not address (street, city, zip code, state, country)
- n The special value ***null*** is a member of every domain
  - | Means *unknown* or *not applicable*
- n The null value causes complications in the definition of many operations
  - | Will be detailed later



# Relations are Unordered

- n A relation is a **set** -> the elements of a set are not ordered per se
- n From a practical perspective:
  - n Order of tuples is irrelevant (tuples may be stored in an arbitrary order)
- n Example: *instructor* relation with unordered tuples

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>
22222	Einstein	Physics	95000
12121	Wu	Finance	90000
32343	El Said	History	60000
45565	Katz	Comp. Sci.	75000
98345	Kim	Elec. Eng.	80000
76766	Crick	Biology	72000
10101	Srinivasan	Comp. Sci.	65000
58583	Califieri	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000



# Database

- n A **database schema**  $S$  consists of multiple relation schema
- n A **database instance**  $I$  for a schema  $S$  is a set of relation instances
  - | One relation for each relation schema in  $S$
- n Information about an enterprise is broken up into parts

*instructor*

*student*

*advisor*

- n Bad design:
  - univ (instructor -ID, name, dept\_name, salary, student\_Id, ..)*results in
  - | repetition of information (e.g., two students have the same instructor)
  - | the need for many null values (e.g., represent an student with no advisor)
- n Normalization theory (Chapter 7) deals with how to design “good” relational schemas avoiding these problems



# Bad Design Example Revisited

- n **Example:** Changing the budget of the 'Physics' department
  - | Updates to many rows!
    - ▶ Easy to break **integrity**
    - ▶ If we forget to update a row, then we have multiple budget values for the physics department!
- n **Example:** Deleting all employees from the 'Physics' department
  - | How to avoid deleting the 'Physics' department?
  - | Dummy employee's to store departments?
    - ▶ This is bad. E.g., counting the number of employees per department becomes more involved

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci.	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000



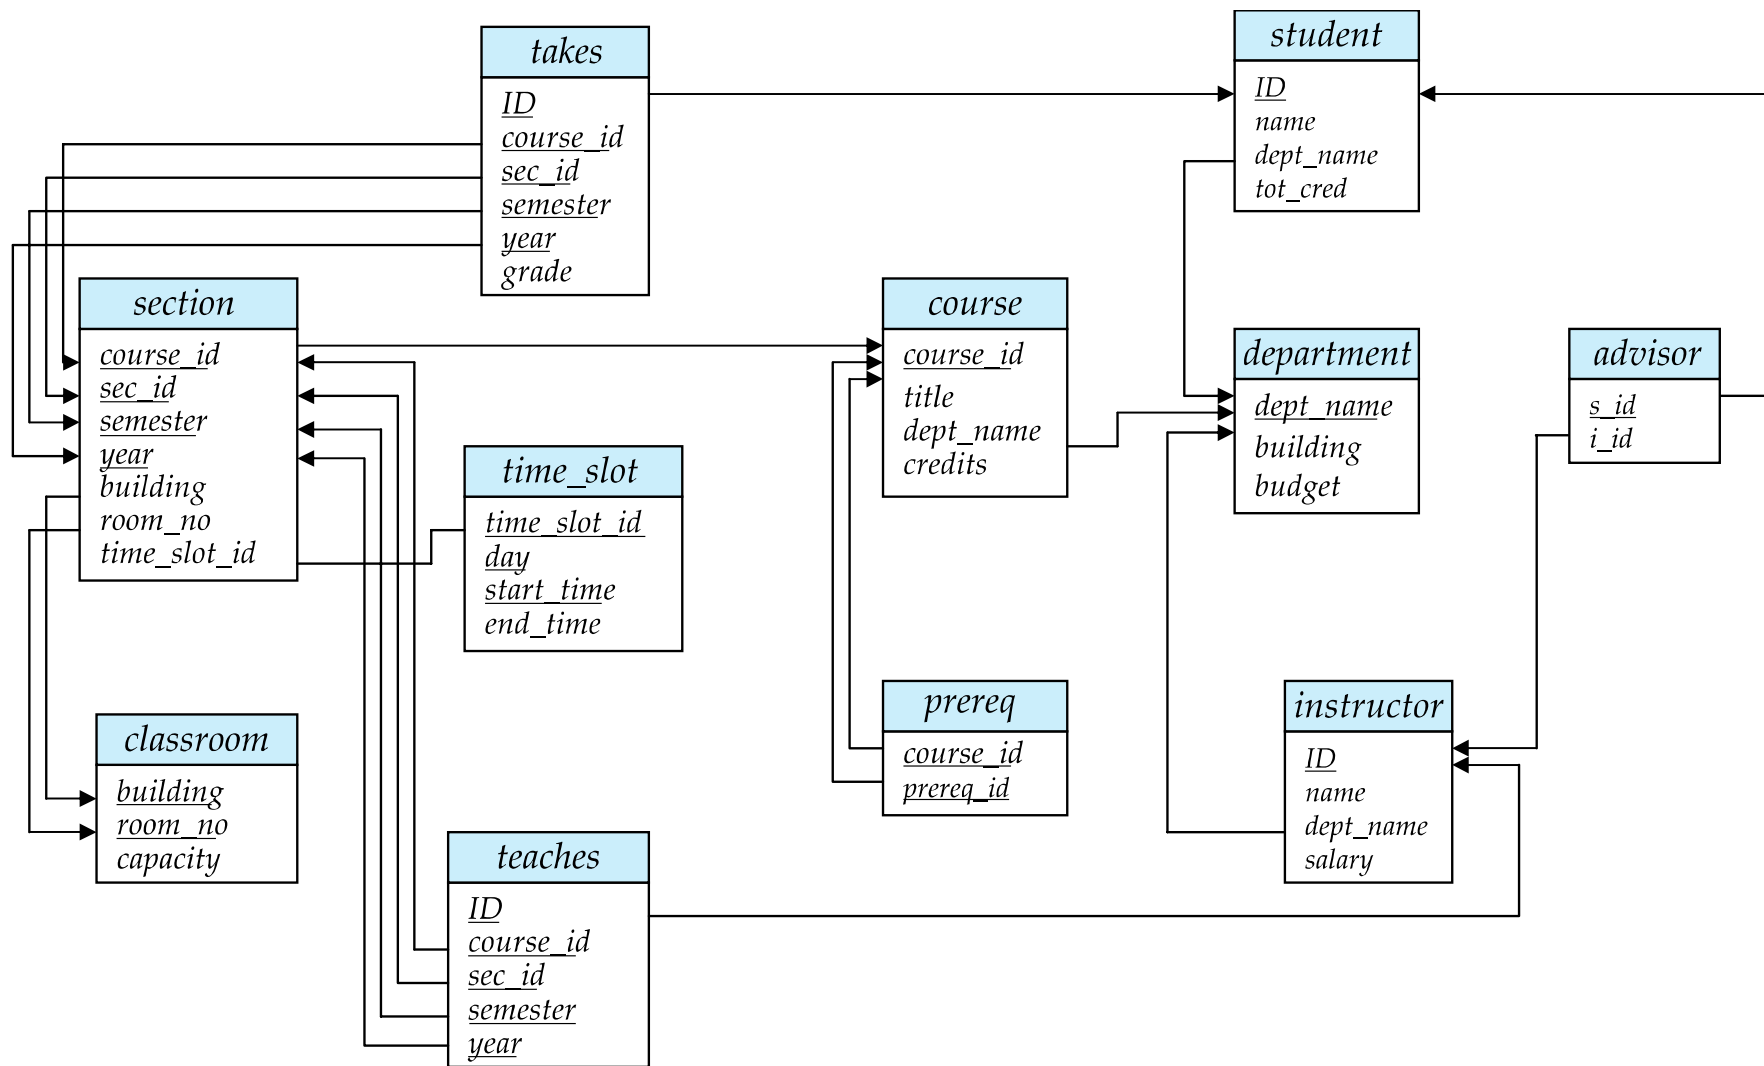
# Keys

- n Let  $K \subseteq R$
- n  $K$  is a **superkey** of  $R$  if values for  $K$  are sufficient to identify a unique tuple of each possible relation  $r(R)$ 
  - | Example:  $\{ID\}$  and  $\{ID, name\}$  are both superkeys of *instructor*.
- n Superkey  $K$  is a **candidate key** if  $K$  is minimal (no subset of  $K$  is also a superkey)
  - Example:  $\{ID\}$  is a candidate key for *Instructor*
- n One of the candidate keys is selected to be the **primary key**.
  - | which one? -> domain specific design choice
- n **Foreign key** constraint: Value in one relation must appear in another
  - | **Referencing** relation
  - | **Referenced** relation





# Schema Diagram for University Database





# Relational Query Languages

- n Procedural vs non-procedural (**declarative**)
- n “Pure” languages:
  - | Relational algebra
  - | Tuple relational calculus
  - | Domain relational calculus
- n Expressive power of a query language
  - | What queries can be expressed in this language?
- n Relational algebra:
  - | Algebra of relations -> set of operators that take relations as input and produce relations as output
  - | -> **composable**: the output of evaluating an expression in relational algebra can be used as input to another relational algebra expression
- n Now: First introduction to operators of the relational algebra



# Selection of tuples

n Relation r

A	B	C	D
$\alpha$	$\alpha$	1	7
$\alpha$	$\beta$	5	7
$\beta$	$\beta$	12	3
$\beta$	$\beta$	23	10

n Select tuples with A=B  
and D > 5

n  $\sigma_{A=B \text{ and } D > 5}(r)$

A	B	C	D
$\alpha$	$\alpha$	1	7
$\beta$	$\beta$	23	10



# Selection of Columns (Attributes)

n Relation  $r$ :

A	B	C
$\alpha$	10	1
$\alpha$	20	1
$\beta$	30	1
$\beta$	40	2

n Select A and C

n Projection

n  $\Pi_{A, C}(r)$

A	C
$\alpha$	1
$\alpha$	1
$\beta$	1
$\beta$	2

=

A	C
$\alpha$	1
$\beta$	1
$\beta$	2



# Joining two relations – Cartesian Product

n Relations  $r, s$ :

A	B
$\alpha$	1
$\beta$	2

$r$

C	D	E
$\alpha$	10	a
$\beta$	10	a
$\beta$	20	b
$\gamma$	10	b

$s$

n  $r \times s$ :

A	B	C	D	E
$\alpha$	1	$\alpha$	10	a
$\alpha$	1	$\beta$	10	a
$\alpha$	1	$\beta$	20	b
$\alpha$	1	$\gamma$	10	b
$\beta$	2	$\alpha$	10	a
$\beta$	2	$\beta$	10	a
$\beta$	2	$\beta$	20	b
$\beta$	2	$\gamma$	10	b



# Union of two relations

n Relations  $r, s$ :

A	B
$\alpha$	1
$\alpha$	2
$\beta$	1

$r$

A	B
$\alpha$	2
$\beta$	3

$s$

n  $r \cup s$ :

A	B
$\alpha$	1
$\alpha$	2
$\beta$	1
$\beta$	3



# Set difference of two relations

n Relations  $r, s$ :

A	B
$\alpha$	1
$\alpha$	2
$\beta$	1

$r$

A	B
$\alpha$	2
$\beta$	3

$s$

n  $r - s$ :

A	B
$\alpha$	1
$\beta$	1



# Set Intersection of two relations

n Relation  $r, s$ :

A	B
$\alpha$	1
$\alpha$	2
$\beta$	1

$r$

A	B
$\alpha$	2
$\beta$	3

$s$

n  $r \cap s$

A	B
$\alpha$	2





# Joining two relations – Natural Join

- n Let  $r$  and  $s$  be relations on schemas  $R$  and  $S$  respectively. Then, the “natural join” of relations  $R$  and  $S$  is a relation on schema  $R \cup S$  obtained as follows:
  - | Consider each pair of tuples  $t_r$  from  $r$  and  $t_s$  from  $s$ .
  - | If  $t_r$  and  $t_s$  have the same value on each of the attributes in  $R \cap S$ , add a tuple  $t$  to the result, where
    - ▶  $t$  has the same value as  $t_r$  on  $r$
    - ▶  $t$  has the same value as  $t_s$  on  $s$



# Natural Join Example

n Relations  $r, s$ :

A	B	C	D
$\alpha$	1	$\alpha$	a
$\beta$	2	$\gamma$	a
$\gamma$	4	$\beta$	b
$\alpha$	1	$\gamma$	a
$\delta$	2	$\beta$	b

$r$

B	D	E
1	a	$\alpha$
3	a	$\beta$
1	a	$\gamma$
2	b	$\delta$
3	b	$\epsilon$

$s$

n Natural Join

n  $r \bowtie s$

A	B	C	D	E
$\alpha$	1	$\alpha$	a	$\alpha$
$\alpha$	1	$\alpha$	a	$\gamma$
$\alpha$	1	$\gamma$	a	$\alpha$
$\alpha$	1	$\gamma$	a	$\gamma$
$\delta$	2	$\beta$	b	$\delta$



## Figure in-2.1

Symbol (Name)	Example of Use
$\sigma$ (Selection)	$\sigma_{\text{salary} \geq 85000}(\text{instructor})$
	Return rows of the input relation that satisfy the predicate.
$\Pi$ (Projection)	$\Pi_{ID, salary}(\text{instructor})$
	Output specified attributes from all rows of the input relation. Remove duplicate tuples from the output.
$\bowtie$ (Natural Join)	$\text{instructor} \bowtie \text{department}$
	Output pairs of rows from the two input relations that have the same value on all attributes that have the same name.
$\times$ (Cartesian Product)	$\text{instructor} \times \text{department}$
	Output all pairs of rows from the two input relations (regardless of whether or not they have the same values on common attributes)
$\cup$ (Union)	$\Pi_{name}(\text{instructor}) \cup \Pi_{name}(\text{student})$
	Output the union of tuples from the two input relations.



# End of Chapter 2

**Modifies from:**

**Database System Concepts, 6<sup>th</sup> Ed.**

©Silberschatz, Korth and Sudarshan

See [www.db-book.com](http://www.db-book.com) for conditions on re-use