



CS425 – Fall 2013 Boris Glavic Chapter 2: Intro to Relational Model

Modifies from:
Database System Concepts, 6th Ed.
©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use



Textbook: Chapter 2

CS425 – Fall 2013 – Boris Glavic

2.2

©Silberschatz, Korth and Sudarshan



Example of a Relation

ID	name	dept_name	salary
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)

CS425 – Fall 2013 – Boris Glavic

2.3

©Silberschatz, Korth and Sudarshan



Attribute Types

- The set of allowed values for each attribute is called the **domain** or **data type** of the attribute
- 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)
- The special value **null** is a member of every domain
 - Means *unknown* or *not applicable*
- The null value causes complications in the definition of many operations
 - Will be detailed later

CS425 – Fall 2013 – Boris Glavic

2.4

©Silberschatz, Korth and Sudarshan



Relation Schema and Instance

- A_1, A_2, \dots, A_n are **attributes names**
 - $R = (A_1, A_2, \dots, A_n)$ is a **relation schema**
- Example:
- instructor* = (*ID*, *name*, *dept_name*, *salary*)
- Formally, given sets D_1, D_2, \dots, D_n of domains a **relation r** (or **relation instance**) is a subset of $D_1 \times D_2 \times \dots \times D_n$
 - Thus, a relation is a **set of n -tuples** (a_1, a_2, \dots, a_n) where each $a_i \in D_i$
 - The current values (**relation instance**) of a relation are often specified in tabular form
 - **Caveat: being a set, the tuples of the relation do not have any order defined as implied by the tabular representation**
 - An element t of r is a **tuple**, represented as a **row** in a table



CS425 – Fall 2013 – Boris Glavic

2.5

©Silberschatz, Korth and Sudarshan



Alternative Definitions

- A relation schema is often defined as a list of attribute-domain pairs
 - That is the data types of each attribute in the relation are considered as part of the relation schema
- Tuples are sometimes defined as functions from attribute names to values (order of attributes does not matter)
- A relation r can be specified as a function
 - $D_1 \times D_2 \times \dots \times D_n \rightarrow \{true, false\}$
 - $t = (a_1, a_2, \dots, a_n)$ is mapped to *true* if t is in r and to *false* otherwise
- These alternative definition are useful in database theory
 - We will stick to the simple definition!



CS425 – Fall 2013 – Boris Glavic

2.6

©Silberschatz, Korth and Sudarshan

Relations are Unordered

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

ID	name	dept_name	salary
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

CS425 – Fall 2013 – Boris Glavic 2.7 ©Silberschatz, Korth and Sudarshan

Database

- A **database schema** *S* consists of multiple relation schema
- A **database instance** *I* for a schema *S* is a set of relation instances
 - One relation for each relation schema in *S*
- Information about an enterprise is broken up into parts
 - instructor*
 - student*
 - advisor*
- 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 a student with no advisor)
- Normalization theory (Chapter 7) deals with how to design "good" relational schemas avoiding these problems

CS425 – Fall 2013 – Boris Glavic 2.8 ©Silberschatz, Korth and Sudarshan

Bad Design Example Revisited

- **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!
- 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

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

CS425 – Fall 2013 – Boris Glavic 2.9 ©Silberschatz, Korth and Sudarshan

Keys

- Let $K \subseteq R$
- *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*.
- 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*
- One of the candidate keys is selected to be the **primary key**.
 - which one? -> domain specific design choice
- **Foreign key** constraint: Value in one relation must appear in another
 - Referencing relation
 - Referenced relation

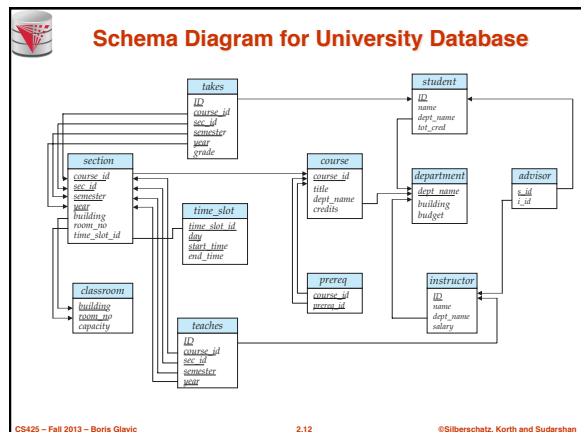
CS425 – Fall 2013 – Boris Glavic 2.10 ©Silberschatz, Korth and Sudarshan

Keys

- Formally, a set of attributes $K \subseteq R$ is a superkey if for every instance *r* of *R* holds that
 - $\forall t, t' \in r: t.K = t'.K \Rightarrow t = t'$
- A superkey *K* is called a candidate key iff
 - $\forall K' \subset K: K'$ is not a superkey
- A foreign key constraint FK is quartuple (*R*, *K*, *R'*, *K'*) where *R* and *R'* are relation schemata, $K \subseteq R$, K' is the primary key of *R'*, and $|K| = |K'|$
- A foreign key holds over an instance (*r*, *r'*) for {*R*,*R'*} iff
 - $\forall t \in R: \exists t' \in R': t.K = t'.K'$

Theory Warning

CS425 – Fall 2013 – Boris Glavic 2.11 ©Silberschatz, Korth and Sudarshan



Relational Query Languages

- Procedural vs non-procedural (**declarative**)
- "Pure" languages:
 - Relational algebra
 - Tuple relational calculus
 - Domain relational calculus
- Expressive power of a query language
 - What queries can be expressed in this language?
- 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
- Now: First introduction to operators of the relational algebra

CS425 – Fall 2013 – Boris Glavic 2.13 ©Silberschatz, Korth and Sudarshan

Selection of tuples

- Relation r

A	B	C	D
α	α	1	7
α	β	5	7
β	β	12	3
β	β	23	10
- Select tuples with $A=B$ and $D > 5$
 - $\sigma_{A=B \text{ and } D > 5}(r)$

A	B	C	D
α	α	1	7
β	β	23	10

CS425 – Fall 2013 – Boris Glavic 2.14 ©Silberschatz, Korth and Sudarshan

Selection of Columns (Attributes)

- Relation r

A	B	C
α	10	1
α	20	1
β	30	1
β	40	2
- Select A and C
 - Projection
 - $\pi_{A,C}(r)$

A	C
α	1
α	1
β	1
β	2

CS425 – Fall 2013 – Boris Glavic 2.15 ©Silberschatz, Korth and Sudarshan

Joining two relations – Cartesian Product

- Relations r, s :

A	B
α	1
β	2

C	D	E
α	10	a
β	10	a
β	20	b
γ	10	b
- $r \times s$:

A	B	C	D	E
α	1	α	10	a
α	1	β	10	a
α	1	β	20	b
α	1	γ	10	b
β	2	α	10	a
β	2	β	10	a
β	2	β	20	b
β	2	γ	10	b

CS425 – Fall 2013 – Boris Glavic 2.16 ©Silberschatz, Korth and Sudarshan

Union of two relations

- Relations r, s :

A	B
α	1
α	2
β	1

A	B
α	2
β	3
- $r \cup s$:

A	B
α	1
α	2
β	1
β	3

CS425 – Fall 2013 – Boris Glavic 2.17 ©Silberschatz, Korth and Sudarshan

Set difference of two relations

- Relations r, s :

A	B
α	1
α	2
β	1

A	B
α	2
β	3
- $r - s$:

A	B
α	1
β	1

CS425 – Fall 2013 – Boris Glavic 2.18 ©Silberschatz, Korth and Sudarshan

Set Intersection of two relations

Relation r, s :

A	B
α	1
α	2
β	1

A	B
α	2
β	3

r s

$r \cap s$

A	B
α	2

CS425 – Fall 2013 – Boris Glavic 2.19 ©Silberschatz, Korth and Sudarshan

Joining two relations – Natural Join

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

CS425 – Fall 2013 – Boris Glavic 2.20 ©Silberschatz, Korth and Sudarshan

Natural Join Example

Relations r, s :

A	B	C	D
α	1	α	a
β	2	γ	a
γ	4	β	b
α	1	γ	a
δ	2	β	b

B	D	E
1	a	α
3	a	β
1	a	γ
2	b	δ
3	b	ϵ

r s

Natural Join

- $r \bowtie s$


A	B	C	D	E
α	1	α	a	α
α	1	α	a	γ
α	1	γ	a	α
α	1	γ	a	γ
δ	2	β	b	

CS425 – Fall 2013 – Boris Glavic 2.21 ©Silberschatz, Korth and Sudarshan

Figure in-2.1

Symbol (Name)	Example of Use
σ (Selection)	$\sigma_{\text{salary} \geq 85000}(\text{instructor})$ Return rows of the input relation that satisfy the predicate.
Π (Projection)	$\Pi_{ID, \text{salary}}(\text{instructor})$ 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)	$\Pi_{\text{name}}(\text{instructor}) \cup \Pi_{\text{name}}(\text{student})$ Output the union of tuples from the two input relations.

CS425 – Fall 2013 – Boris Glavic 2.22 ©Silberschatz, Korth and Sudarshan



End of Chapter 2

Modifies from:
Database System Concepts, 6th Ed.
©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use

CS425 – Fall 2013 – Boris Glavic 2.24 ©Silberschatz, Korth and Sudarshan

Recap

- Database Schema** (or short schema)
 - Set of **relation schemata**
 - List of **attribute names**
- Database Instance** (or short database)
 - Set of **relations instances**
 - Set of **tuples**
 - List of **attribute values**
- Integrity Constraints**
 - Keys** (Super-, Candidate-, Primary-)
 - For identifying tuples
 - Foreign keys**
 - For referencing tuples in other relations
- Query language**
 - Declarative**
 - Retrieve, combine, and analyze data from a database instance

CS425 – Fall 2013 – Boris Glavic 2.24 ©Silberschatz, Korth and Sudarshan



Outline

- Introduction
- Relational Data Model
- **Formal Relational Languages (relational algebra)**
- SQL
- Database Design
- Transaction Processing, Recovery, and Concurrency Control
- Storage and File Structures
- Indexing and Hashing
- Query Processing and Optimization

CS425 – Fall 2013 – Boris Glavic

2.25

©Silberschatz, Korth and Sudarshan



Figure 2.01

ID	name	dept_name	salary
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	Calliferi	History	62000
76543	Singh	Finance	80000
76766	Crick	Biology	72000
83821	Brandt	Comp. Sci.	92000
98345	Kim	Elec. Eng.	80000

CS425 – Fall 2013 – Boris Glavic

2.26

©Silberschatz, Korth and Sudarshan



Figure 2.02

course_id	title	dept_name	credits
BIO-101	Intro. to Biology	Biology	4
BIO-301	Genetics	Biology	4
BIO-399	Computational Biology	Biology	3
CS-101	Intro. to Computer Science	Comp. Sci.	4
CS-190	Game Design	Comp. Sci.	4
CS-315	Robotics	Comp. Sci.	3
CS-319	Image Processing	Comp. Sci.	3
CS-347	Database System Concepts	Comp. Sci.	3
EE-181	Intro. to Digital Systems	Elec. Eng.	3
FIN-201	Investment Banking	Finance	3
HIS-351	World History	History	3
MU-199	Music Video Production	Music	3
PHY-101	Physical Principles	Physics	4

CS425 – Fall 2013 – Boris Glavic

2.27

©Silberschatz, Korth and Sudarshan



Figure 2.03

course_id	prereq_id
BIO-301	BIO-101
BIO-399	BIO-101
CS-190	CS-101
CS-315	CS-101
CS-319	CS-101
CS-347	CS-101
EE-181	PHY-101

CS425 – Fall 2013 – Boris Glavic

2.28

©Silberschatz, Korth and Sudarshan



Figure 2.04

ID	name	dept_name	salary
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	Calliferi	History	62000
83821	Brandt	Comp. Sci.	92000
15151	Mozart	Music	40000
33456	Gold	Physics	87000
76543	Singh	Finance	80000

CS425 – Fall 2013 – Boris Glavic

2.29

©Silberschatz, Korth and Sudarshan



Figure 2.05

dept_name	building	budget
Biology	Watson	90000
Comp. Sci.	Taylor	100000
Elec. Eng.	Taylor	85000
Finance	Painter	120000
History	Painter	50000
Music	Packard	80000
Physics	Watson	70000

CS425 – Fall 2013 – Boris Glavic

2.30

©Silberschatz, Korth and Sudarshan

