

The Relational Model

Lecture 3



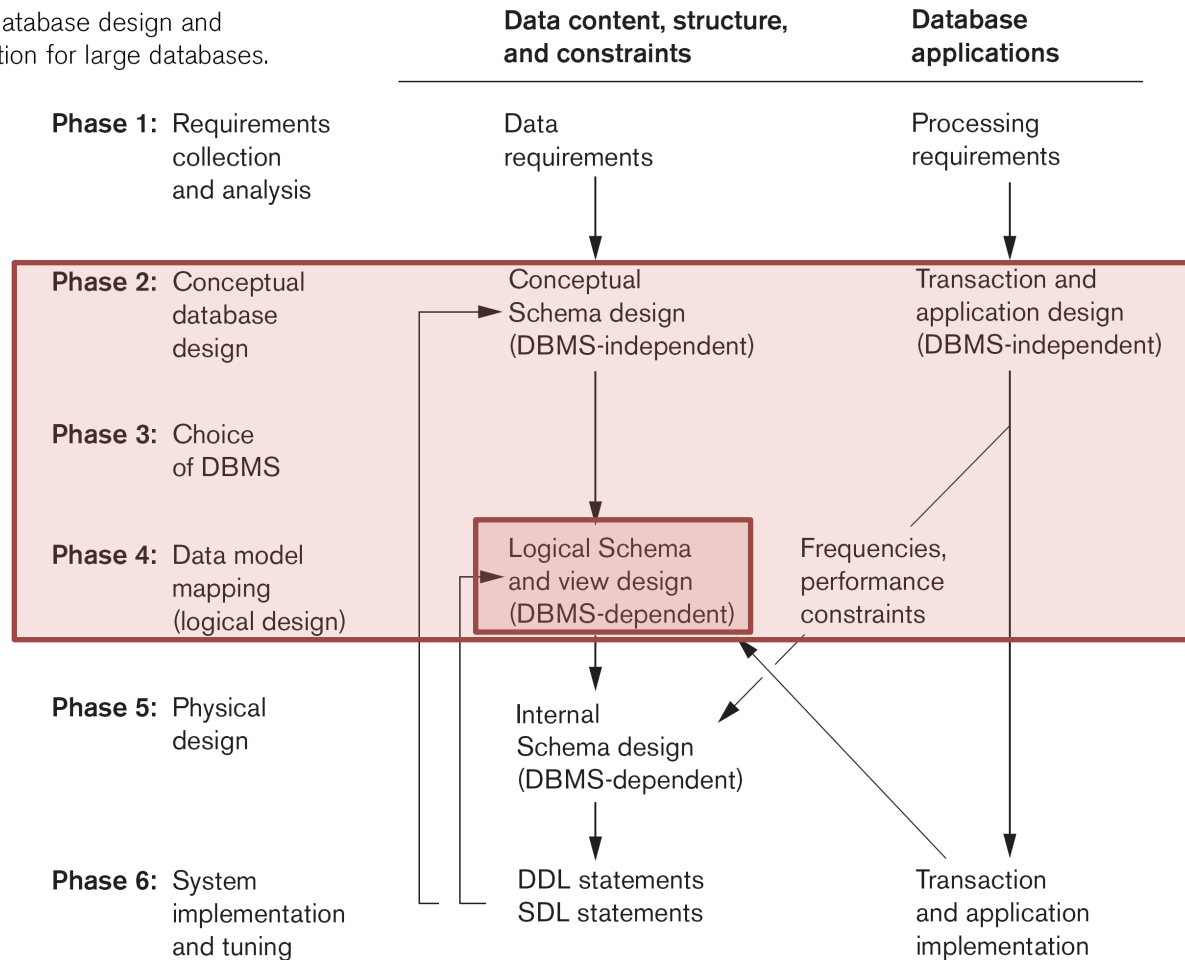
Outline

1. Context
2. Model Concepts
3. Relational Constraints
4. Update Operations
5. Transactions



Database Design and Implementation Process

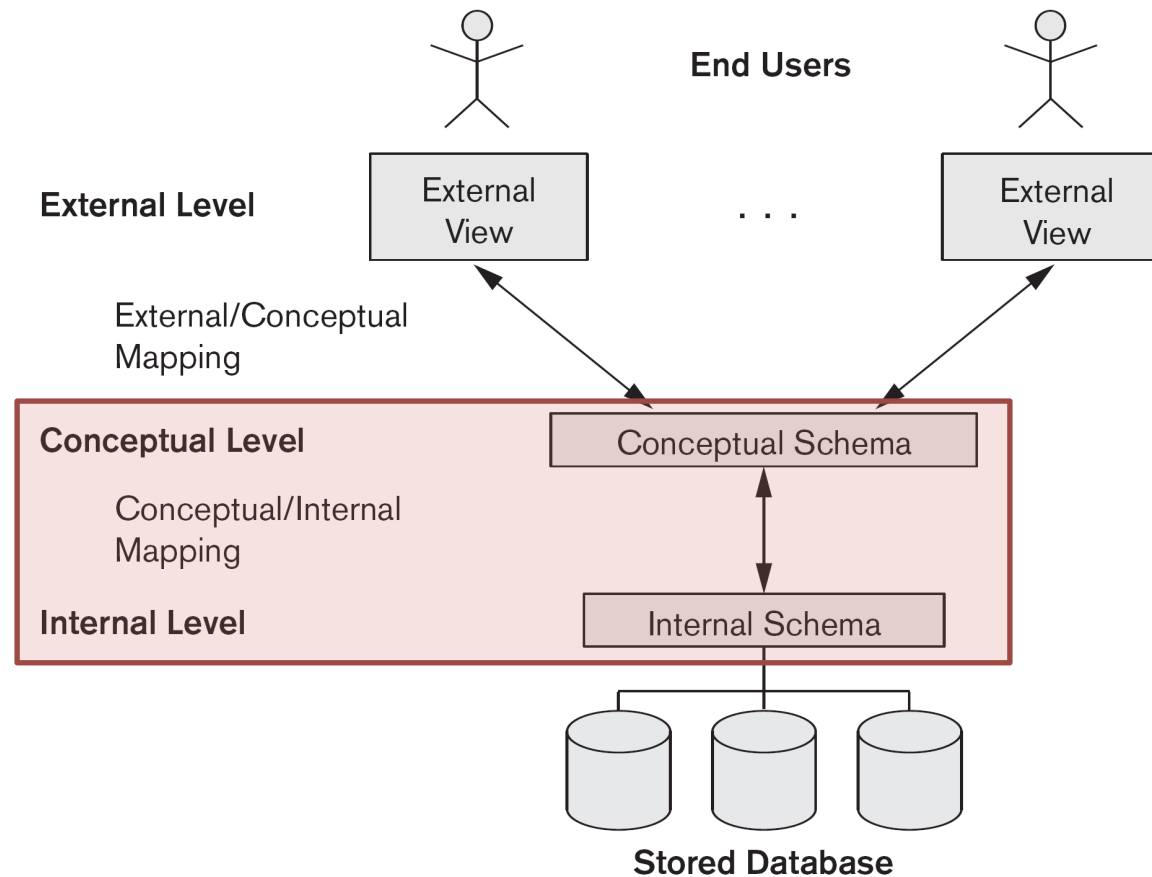
Figure 10.1
Phases of database design and implementation for large databases.



Data Models

Figure 2.2

The three-schema architecture.



The Relational Model

Codd, Edgar F. "A relational model of data for large shared data banks." *Communications of the ACM* 13.6 (1970): 377-387.

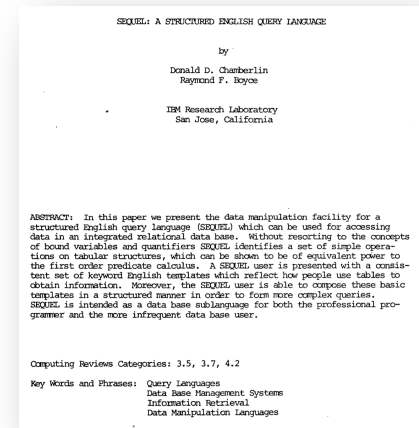
"Future users of large data banks must be protected from having to know how the data is organized in the machine (the internal representation)... Activities of users at terminals and most application programs should remain unaffected when the internal representation of data is changed and even when some aspects of the external representation are changed... In Section 1, inadequacies of [existing] models are discussed. A model based on n-ary relations, a normal form for data base relations, and the concept of a universal data sublanguage are introduced. In Section 2, certain operations on relations (other than logical inference) are discussed and applied to the problems of redundancy and consistency in the user's model."



FYI: SQL

Chamberlin, Donald D., and Raymond F. Boyce. "SEQUEL: A structured English query language." *Proceedings of the 1974 ACM SIGFIDET (now SIGMOD) workshop on Data description, access and control*. ACM, 1974.

“In this paper we present the data manipulation facility for a structured English query language (SEQUEL) which can be used for accessing data in an integrated relational data base. Without resorting to the concepts of bound variables and quantifiers SEQUEL identifies a set of simple operations on tabular structures, which can be shown to be of equivalent power to the first order predicate calculus. A SEQUEL user is presented with a consistent set of keyword English templates which reflect how people use tables to obtain information. Moreover, the SEQUEL user is able to compose these basic templates in a structured manner in order to form more complex queries. SEQUEL is intended as a data base sublanguage for both the professional programmer and the more infrequent data base user.”



Motivation

- A **declarative** method for specifying data and queries
- A **formal** mathematical basis for database systems
- A foundation for efficient and usable database systems



Model Concepts

- A database is a set of named **relations** (tables) and a set of **integrity constraints**
 - Database is in a **valid state** if it satisfies all integrity constraints (else **invalid state**)
- The schema of an n -ary relation is an ordered list of n **attributes** (columns)
 - Mathematically equivalent as a set
- Each attribute has a **domain** (type) of atomic values
 - Related to the 1NF assumption
- The state of the relation is a set of **n -tuples** (rows), each an ordered list of values in the corresponding domain, or **NULL**
 - Mathematically a subset of the Cartesian product of the attribute domains; related to the closed-world assumption
 - Actual implementations loosen the definition to a *bag* of tuples (and hence allow duplicate rows, more later)



Example Relation

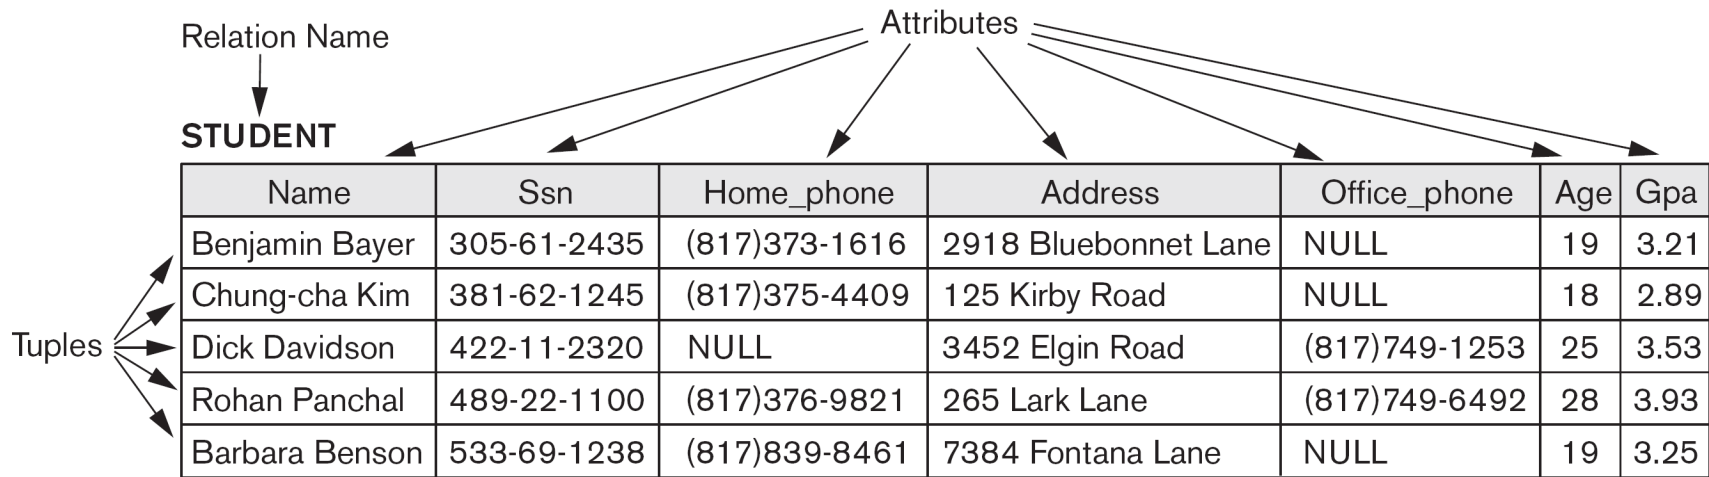


Figure 3.1

The attributes and tuples of a relation STUDENT.

STUDENT(Name, Ssn, Home_phone, Address, Office_phone, Age, Gpa)

dom(Name) = Names

dom(Ssn) = Social_security_numbers

...



Ordering of Tuples

- A relation is formally defined as a set of tuples; thus, there is no inherent order
- The physical representation *will* have an ordering, but the relation definition sets no preference as to this ordering
- As we will discuss later in physical design, indexes may establish an ordering for purposes of query efficiency



Values in Tuples

- Each value must be atomic – no composite or multi-valued attributes
 - Composite: simple component attributes
 - Multi-valued: separate relations
- NULL
 - Several possible meanings (e.g. unknown, not available, does not apply, undefined)
 - Best to avoid, else deal with caution (esp. during queries with filtration/aggregation)



Violation of Atomic (1)

(a)

DEPARTMENT

Dname	<u>Dnumber</u>	Dmgr_ssn	Dlocations

(b)

DEPARTMENT

Dname	<u>Dnumber</u>	Dmgr_ssn	Dlocations
Research	5	333445555	{Bellaire, Sugarland, Houston}
Administration	4	987654321	{Stafford}
Headquarters	1	888665555	{Houston}

(c)

DEPARTMENT

Dname	<u>Dnumber</u>	Dmgr_ssn	<u>Dlocation</u>
Research	5	333445555	Bellaire
Research	5	333445555	Sugarland
Research	5	333445555	Houston
Administration	4	987654321	Stafford
Headquarters	1	888665555	Houston

Figure 15.9

Normalization into 1NF. (a) A relation schema that is not in 1NF. (b) Sample state of relation DEPARTMENT. (c) 1NF version of the same relation with redundancy.



Violation of Atomic (2)

(a)

EMP_PROJ

Ssn	Ename	Projs	
		Pnumber	Hours

(b)

EMP_PROJ

Ssn	Ename	Pnumber	Hours
123456789	Smith, John B.	1	32.5
		2	7.5
666884444	Narayan, Ramesh K.	3	40.0
453453453	English, Joyce A.	1	20.0
		2	20.0
333445555	Wong, Franklin T.	2	10.0
		3	10.0
		10	10.0
		20	10.0
999887777	Zelaya, Alicia J.	30	30.0
		10	10.0
987987987	Jabbar, Ahmad V.	10	35.0
		30	5.0
987654321	Wallace, Jennifer S.	30	20.0
		20	15.0
888665555	Borg, James E.	20	NULL

Figure 15.10
 Normalizing nested relations into 1NF. (a) Schema of the EMP_PROJ relation with a *nested relation* attribute PROJS. (b) Sample extension of the EMP_PROJ relation showing nested relations within each tuple. (c) Decomposition of EMP_PROJ into relations EMP_PROJ1 and EMP_PROJ2 by propagating the primary key.

(c)

EMP_PROJ1

Ssn	Ename
-----	-------

EMP_PROJ2

Ssn	Pnumber	Hours
-----	---------	-------



Types of Relational Constraints

Categories of restrictions on data that can be specified on a relational database:

1. Inherent in the data model (implicit)
2. **Schema-based (explicit)**
3. Application-based (or trigger/assertions)
4. Data dependencies
Relates to “goodness” of database design;
we will revisit in *normalization*



Schema-Based Constraints

- Domain constraints
- Key constraints
- Constraints on NULLs
- Entity integrity
- Referential integrity



Domain Constraints

- Within each tuple, the value of each attribute A must be an atomic value from the domain $\text{dom}(A)$
- More later on standard data types in SQL



Keys

A **key** is a set of attributes that satisfies two properties:

1. Two distinct tuples in any state of the relation cannot have identical values for all the attributes of the key (termed *superkey*)
2. We cannot remove any attributes from the key and still have the uniqueness constraint hold (termed *minimal superkey*)

A relation may have multiple keys (each is a **candidate key**). Relations commonly have a **primary key** (underlined; typically small number of attributes, used to *identify* tuples), and may also have some number of additional **unique keys**.



Constraints on NULLs

- Schema must dictate whether or not a NULL value is allowed for each attribute
- No primary key value can be NULL (**entity integrity constraint**)



Referential Integrity

All tuples in one relation must refer to an *existing* tuple in some relation (or NULL); indicated by directed arc

A foreign key (FK) in R1 *references* R2 if...

1. The attributes in FK have the same domain(s) as the primary key attribute(s) PK of R2
2. A value of FK in a tuple t_1 either is NULL or occurs as a value of PK for some tuple t_2 (t_1 **refers to** t_2)



COMPANY

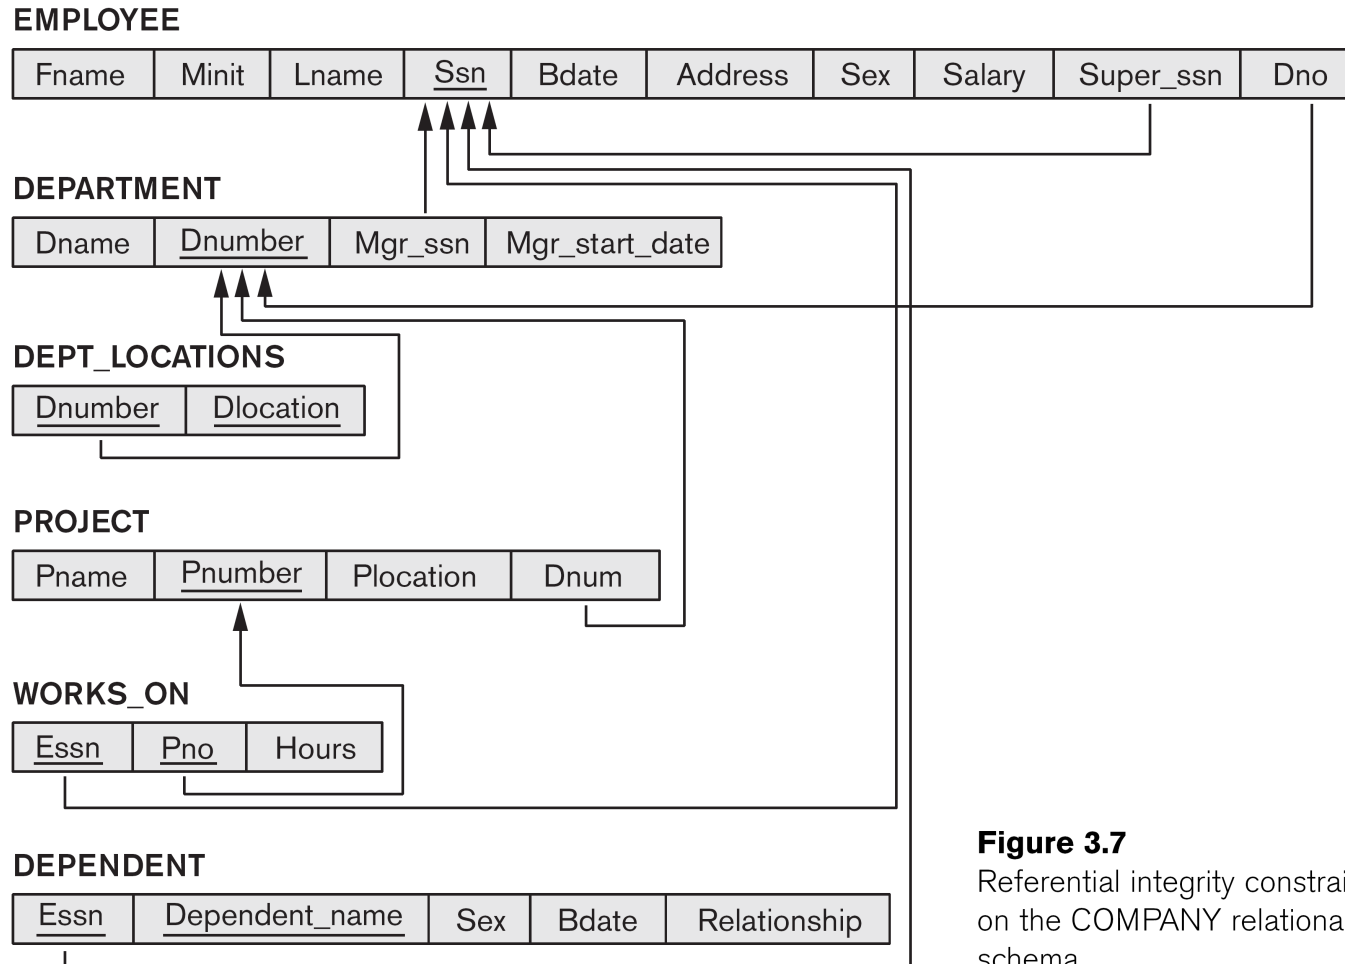


Figure 3.7
Referential integrity constraints displayed on the COMPANY relational database schema.



Update Operations

- Insert
- Delete
- Update

We now examine how these can violate various types of constraints and the resulting actions that can be taken.



Insert

- Domain
 - An attribute value does not appear in the corresponding domain (including NULL)
- Key
 - A key value already exists in another tuple
- Entity
 - Any part of the primary key is NULL
- Referential Integrity
 - Any value of any foreign key refers to a tuple that does not exist in the referenced relation

Typical action: reject insertion



Delete

- Referential Integrity
 - Tuple being deleted is referenced by foreign keys from other tuples

Possible actions:

- Reject deletion
- Cascade (propagate deletion)
- Set default/NULL referencing attribute values (careful with primary key)



Update

- If neither part of primary key nor foreign key, need only check...
 - Domain
- Modifying primary key...
 - Like Delete + Insert
- Modifying foreign key...
 - Like Insert

Actions typically similar to Delete with separate options.



Transaction

- Sequence of database operations, including retrieval and update(s)
- Treated as atomic unit of work
 - Must leave the database in a valid state

