

Introduction to Databases

Motivation
Historical
Classical Models
E/R Model

Robert Keller
21 January 2004

1

Definition: Database

- ◆ A collection of data items
 - ▶ structured so as to provide efficient access to the items based on partial knowledge of those items, and
 - ▶ stored persistently and independently of the application(s) accessing those items.

2

Typical Characteristics of Databases

- ◆ Large volume of data
- ◆ Persistent storage
- ◆ Non-volatile media (vs. main memory)
- ◆ Logging and recovery
- ◆ Access control
- ◆ Concurrency control
- ◆ Transaction-oriented
- ◆ Languages for querying and updating

3

Definition: Enterprise

- ◆ Databases are not just things for storing data.
- ◆ They are also a **model** of a set of **real-world** activities, the running of which they are intended to facilitate.
- ◆ These activities are collectively called the "**Enterprise**". (The term is used even for non-business applications.)

4

Definition: Database System

- ◆ A computer application or related set of applications that makes essential use of a database.

5

Examples of Database Systems

- ◆ Business data
 - ▶ Payroll, orders, reservations, catalogs
- ◆ Genealogies
- ◆ Legal cases, patents
- ◆ Library catalogs
- ◆ Meeting schedules
- ◆ Museum collections
- ◆ Music collections
- ◆ Scientific experiment results
- ◆ Software code versions
- ◆ Trouble reports and resolutions

6

Definition: Database Management System (DBMS)

- ◆ A software system that provides infrastructure for creating database systems.
- ◆ Infrastructure:
 - ▶ Query language(s)
 - ▶ API
 - ▶ Operating system/file system interfaces
 - ▶ Mechanisms for transactions, recovery, logging, etc.

7

Examples of Database Management Systems

- ◆ Oracle
- ◆ DB2 (IBM)
- ◆ MySQL
- ◆ PostgreSQL
- ◆ ObjectStore
- ◆ Poet

8

Definition: Database Administrator

- ◆ Someone endowed with special powers for a database system:
 - ▶ Create new user.
 - ▶ Set privileges.
 - ▶ Modify the schema.

9

Definition: Schema

- ◆ The structural specification for a database.
- ◆ In some cases, the schema itself can be stored as a meta-database.
- ◆ Usually there is a special language for defining the schema, called a **Data Definition Language (DDL)**, in contrast to the **Data Manipulation Language (DML)**.

10

Related Areas

- ◆ Data Structures
- ◆ Logic
- ◆ Software Engineering
- ◆ Operating Systems (including distributed)
- ◆ Information Retrieval
- ◆ Data Mining, machine learning, inductive logic programming

11

Societal Issues

- ◆ Privacy
- ◆ Copyright, Intellectual Property
- ◆ Misinformation
- ◆ Legal liability
- ◆ Global Economy
- ◆ Standardization

12

Prehistory

- ◆ Punched cards, paper tape
- ◆ Magnetic tape

13

Paleozoic Era

- ◆ Data is in one or more files.
- ◆ Files are loaded entirely, or piecemeal, into main memory.
- ◆ Files are rewritten to commit changes.
- ◆ What problems do you see with this approach?

14

Mesozoic Era

- ◆ Sequential file, say ordered by key field
- ◆ Indexed sequential file, disks
- ◆ "Inverted" file
(keys on other fields)

15

Cenozoic Era

- ◆ Data resides in directly-accessible non-volatile storage hierarchy.
- ◆ Units of data are moved to main memory on demand, then written back to storage unit by unit as the application progresses.

16

A Fundamental Dichotomy

- ◆ **Physical** access (where the data are on disk, how they are connected, how are they transported, etc.)
- ◆ **Logical** structure (what are the essential logical connections between data units)

17

Evolution of Database Models

- ◆ Early models (ca 1960's)
 - ◆ Hierarchical (IBM)
 - ◆ Navigational or "Network" (Codasyl DBTG)
- ◆ Relational (1969 to present)
- ◆ Entity/Relationship model (1976)
- ◆ Deductive databases (1980's)
- ◆ Object-Oriented DB's (1990's to present)
- ◆ World-wide-web
- ◆ Peer-to-peer databases
- ◆ XML databases

18

Main Tensions

- ◆ Provide human ease in structuring and querying, while simultaneously providing efficient access and integration with applications.
- ◆ Provide multiple, distributed, access paths, while simultaneously maintaining security and integrity of data.

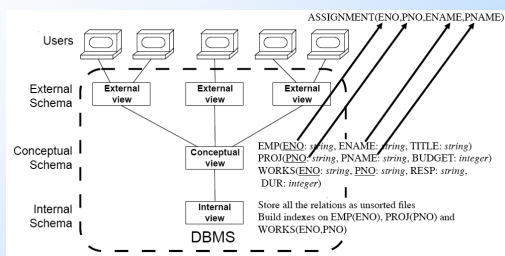
19

Early Attempt at Clarification: ANSI/SPARC Model (1977)

- ◆ ANSI = "American National Standards Institute"
- ◆ SPARC = "Standards Planning And Requirements Committee" (of ANSI)
- ◆ Clear separation of logical "views" from physical structure:
 - ▶ Internal view (physical, **changeable**)
 - ▶ Conceptual view (logical, system-wide, **stable**)
 - ▶ External views (logical, **per-application**)
- ◆ Proposal, rather than an actual standard, but still used as a software pattern, e.g. IHFS ([Integrated Hierarchical File System](#))

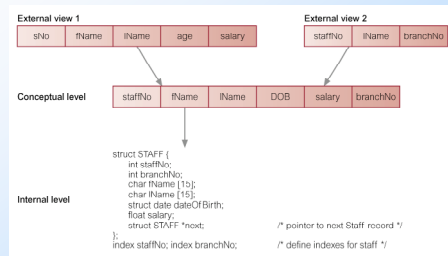
20

ANSI/SPARC Architecture



21

Differences between Three Levels of ANSI-SPARC Architecture



22

Classical Data Models

- ◆ Hierarchical Model
- ◆ Navigational or "Network" Model
- ◆ Relational Model

23

The Hierarchical Data Model

- ◆ Data is organized as a **forest**, with records as nodes.
- ◆ A record can contain **repeating groups**, each consisting of multiple records of another type.
- ◆ For example, a **department** record could have a repeating group **projects**, **projects** could have a repeating group **employee**, etc.

24

Instances of the Hierarchical Model

- ◆ IBM's IMS (Information Management System, 1968, still in use with System 390) was perhaps the original
 - ▶ Allowed virtual links, similar to unix file system, so not strictly a hierarchy
 - ▶ Accessibility integrity: No record could exist without its parent also existing.
 - ▶ [Tutorial](#)
- ◆ MUMPS (used in medical centers)
- ◆ Filemaker Pro (used in HMC administration)


25

The Navigational or "Network" Model

- ◆ A fore-runner of Object-Oriented Databases
- ◆ Data stored in records, which were strongly-typed.
- ◆ Records grouped together in "sets".
- ◆ Each "set" was like a circular list and had an owner record.
- ◆ Through commands, the programmer can "navigate" the database, by
 - ▶ Moving from one member of a set type to the next.
 - ▶ Moving from a member to the owner.
- ◆ Links provided a way to get from one type of set to another.
- ◆ Linked lists were the obvious physical model.

26


The Navigational or "Network" Model



The records themselves are stored on disk. There could be one current member of a given type set in memory.

27

Each member of a set of one type could be the owner of a set of another type, and son on.




28

Intersecting tree-like hierarchies

- ◆ A record can be a member of at most one set of a given type, but it can be a member of other sets of different types.
- ◆ Consequently, the data organization, while somewhat hierarchical, is not strictly a tree.

29

Intersecting tree-like hierarchy




30

Navigation

- ◆ Through a process of navigation, the programmer would fetch current records of one or more types, collecting items for the response to a query.
- ◆ Explicit pointers were avoided by having fetch of the next record be relative to current records.
- ◆ The model was interfaced to specific programming languages (such as COBOL, Fortran).
- ◆ The process is fascinating, but ultimately somewhat cumbersome.

31

The Navigational or "Network" Model



Charles Bachman initiated the DBTG.
He received the ACM Turing Award in 1973.

32

The Relational Model

- ◆ Relational is the dominant database model.
- ◆ Introduced by Edgar ("Ted") Codd of IBM in 1969. Codd received the Turing Award in 1981.



Edgar Codd, 1923-2003

33

A Few of Codd's 12 Rules

- ◆ *Rule 1: The Information Rule*
All data should be presented to the user in table form.
- ◆ *Rule 2: Guaranteed Access Rule*
All data should be accessible without ambiguity.
- ◆ *Rule 3: Systematic Treatment of Null Values*
A field should be allowed to remain empty.
- ◆ *Rule 8: Physical Data Independence*
The user should be isolated from the physical method of storing and retrieving information from the database.

34

The Relational Model

- ◆ Completely separates logical from physical.
- ◆ Data are stored in **relations** ("tables").
- ◆ Each table is a *set* of **tuples**.
- ◆ Queries based on predicate calculus or a functional formulation known as "relational algebra".
- ◆ Note that using logic for queries was in use within AI in 1968 (Cordell Green) or earlier.

35

Sample Relational DB

Computer relation

PropNo	Make	Model	AcqDate	Memory	HD
7	SGI	PowerFlop	4/1/1999	256M	10G
13	IBM	MegaSwap	12/31/2000	128M	100G
42	Dell	MediaBuster	6/12/2002	512M	30G

Location relation

PropNo	Bldg	Room	PrincUser
7	Olin	3006	Sam
13	Beckman	2001	Alice
42	Linde	8192	Leslie

36

More Relational Nomenclature

- ◆ The column headings are (names of) **Attributes**.
- ◆ A query that makes use of attributes common to multiple tables is called a **Join** query:
"Find the Bldg and Room of all computers made by Dell."

37

Advantages of Relations

- ◆ Very general: functional data is a special case of relational.
- ◆ Fairly neutral about what kind of data are values of attributes.
- ◆ All types of attributes treated uniformly.

38

Perceived Difficulties with Relations

- ◆ All types of attributes are treated uniformly. So need to be prepared for queries that might **not** be likely or **meaningful**.
- ◆ Some queries can be very **expensive** to evaluate, depending on how the relations are set up initially.
- ◆ Unless the database is **normalized**, which requires some care in design, there can be much replication of information and other bad things ("anomalies"). The database might not be a good model of the enterprise it is intended to serve.

39

The Entity/Relationship Model

- ◆ This model was introduced in 1976 by Peter Chen, a professor at Louisiana State Univ.
- ◆ Its purpose was to make the connection with the enterprise clearer.
- ◆ It also was to be an implementation-neutral model that could be used for relational, navigational, hierarchical, and possibly other models.
- ◆ The UML of software engineering was evidently influenced by the E/R model, but E/R is not strictly a subset of UML.



40

Use of the E/R Model

- ◆ While it is easy to create a relational database, it is also easy to "get it wrong".
- ◆ The E/R model is typically used as the basis for a rational **database design** document.
- ◆ There is a graphical notation for it.

41

Ullman Slides

Many of the following slides are from Prof. Jeffrey Ullman, the first author of our text. The original version of these slides is available on the web with audio voice-over. I have freely made modifications to suit my own tastes in presentation.

42

Purpose of the E/R Model

- ◆ The E/R model allows us to sketch the design of a database informally.
- ◆ Designs are pictures called *entity-relationship diagrams*.
- ◆ Fairly mechanical ways to convert E/R diagrams to real implementations like relational databases exist.

43

Entity Sets

- ◆ *Entity* = "thing" or object.
- ◆ *Entity set* = collection of similar entities.
 - ▶ Similar to a class in object-oriented languages.
- ◆ *Attribute* = property of an entity.
 - ▶ Generally, all entities in a set have the same types of attributes.
 - ▶ Attributes have simple values, e.g. integers, character strings, dates, images.


44

E/R Diagrams

- ◆ In an entity-relationship diagram, each **entity set** is represented by a **rectangle**.
- ◆ Each **attribute** of an entity set is represented by an **oval**, with a line to the rectangle representing its entity set.

45

Example



- ◆ Entity set Beers has two attributes, name and manf (manufacturer).
- ◆ Each Beer entity has values for these two attributes, e.g. (Bud, Anheuser-Busch)

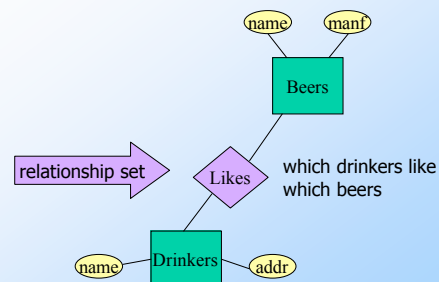
46

Relationships

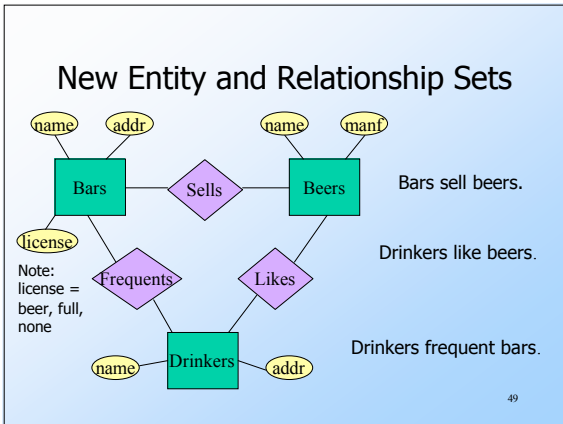
- ◆ A **relationship set** (also called "relationship") connects two or more entity sets.
 - ▶ Similar to an **association** in UML.
- ◆ A relationship set is represented by a **diamond**, with lines to each of the entity sets that are involved.

47

Example



48



Relationship Set

- ◆ The current "value" of an entity set is the set of entities that belong to it.
 - ▶ Example: the set of all bars in our database.
- ◆ The "value" of a relationship set is a set of tuples of currently related entities, one from each of the related entity sets.

50

Relationship vs. Relationship Set

- ◆ The book may use the term "relationship" when "relationship set" is intended.
- ◆ There is no good reason not to have parity in notation, especially since relationships are the more complex of relationships vs. entities.

51

E/R Diagram is for **Types**, not actual entities and relationships

- ◆ Regardless of the number of entities (which beers, which bars, which drinkers) and the number of relationships among those entities, it's still the same diagram.

52

Example

- ◆ For the relationship *Sells*, we might have a relationship set such as:

Bar	Beer
Joe's Bar	Bud
Joe's Bar	Miller
Sue's Bar	Bud
Sue's Bar	Pete's Ale
Sue's Bar	Bud Lite


53

Diagram Instantiation

- ◆ An **instance** or **instantiation** of an ER diagram has actual entities and relationships (instead of entity sets and relationship sets) as nodes.
- ◆ In AI, an instance diagram is called a "Semantic Network".

54

Instance Diagram for Sells



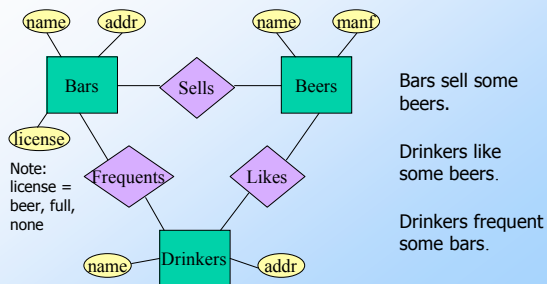
55

Multiway Relationships

- ◆ Sometimes, we need a **relationship** that connects **more than two** entity sets.
- ◆ Suppose that drinkers will only drink certain beers in combination with certain bars.
 - ◆ Our three binary relationships Likes, Sells, and Frequents do not allow us to make this distinction.
 - ◆ But a 3-way relationship would.


56

Two-way Relationships



57

Three-Way Relationship



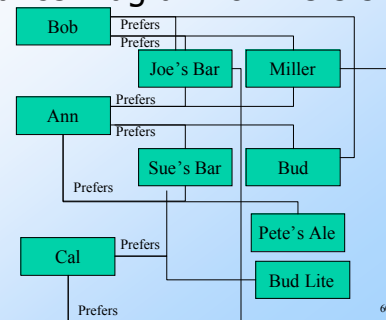
58

A Relationship Set for Prefers

Bar	Drinker	Beer
Joe's Bar	Ann	Miller
Sue's Bar	Ann	Bud
Sue's Bar	Ann	Pete's Ale
Joe's Bar	Bob	Bud
Joe's Bar	Bob	Miller
Joe's Bar	Cal	Miller
Sue's Bar	Cal	Bud Lite

59

Instance Diagram for Prefers



60

Definition: Many-One Relationships

- ◆ Some binary relationships are **many-one** from one or more entity sets to the others.
- ◆ Each entity of the first set of sets is connected to **at most one** entity in the second set.
- ◆ But an entity of the second set can be connected to zero, one, or many entities of the first set.

61

Example: Many-One

- ◆ Assume that a drinker has at most one favorite beer.
- ◆ But a beer can be the favorite of any number of drinkers, including zero.
- ◆ Relationship **Favorite**, from *Drinkers* to *Beers* is thus many-one.

62

Forward Comment

- ◆ Many-one relations capture the idea of a functional dependency between data in the enterprise.
- ◆ These dependencies play an important role in **normalization**, which we discuss later for the relational model.

63

Definition: Many-Many

- ◆ This is the general case.
- ◆ A many-to-many relationship is not many-to-one on any dichotomy of attributes.


64

Definition: One-One Relationships

- ◆ In a one-one relationship between two entity sets, each entity of either entity set is related to **at most one** entity of the other set.
- ◆ Example: Relationship *Best-seller* between entity sets *Manfs* (manufacturer) and *Beers*.
 - ▶ A beer cannot be made by more than one manufacturer, and no manufacturer can have more than one best-seller (assume no ties).

65

In Pictures:



Note that "one-one" is not quite the same as a "one-to-one" function in mathematics.


66

Representing "Multiplicity"

- ◆ Show a many-one relationship by an **arrow toward** the "one" side.
- ◆ Show a one-one relationship by **bi-directional arrows**.
- ◆ In some situations, we can also assert "exactly one," i.e., each entity of one set must be related to exactly one entity of the other set. To do so, we use a **rounded arrow**.

67

Example



68

Example

- ◆ Consider the relationship **Best-seller** between *Manfs* and *Beers*.
- ◆ Some beers are not the best-seller of any manufacturer, so a rounded arrow to *Manfs* would be inappropriate.
- ◆ But a manufacturer has to have a best-seller (we assume they are beer manufacturers).

69

Rounded vs. Ordinary Arrows in the E/R Diagram




70

Attributes on Relationships

- ◆ Sometimes it is useful to attach an **attribute** to a **relationship**.
- ◆ Think of this attribute as a property of tuples in the relationship set [where "property" could be represented as another component of the tuple].

71

Example



Price is a function of both the bar and the beer, not of one alone.


72

Equivalent Diagrams Without Attributes on Relationships

- ◆ Create an entity set representing values of the attribute.
- ◆ Make that entity set participate in the relationship.

73

Example




Definition: Roles

- ◆ Sometimes an entity set appears more than once in a relationship.
- ◆ Label the edges between the relationship and the entity set with names called *roles*.


75

Example: Roles



76

Example: Roles



77

Subclasses

- ◆ Subclass = special case = fewer entities = more properties.
- ◆ Example: Ales are a kind of beer.
 - ▶ Not every beer is an ale, but some are.
 - ▶ Let us suppose that in addition to all the *properties* (attributes and relationships) of beers, ales also have the attribute *color*.


78

Subclasses in E/R Diagrams

- ◆ Assume subclasses form a tree.
 - ▶ i.e., no multiple inheritance.
- ◆ Isa triangles indicate the subclass relationship.
 - ▶ Point to the superclass.

79

Example




80

E/R Vs. Object-Oriented Subclasses

- ◆ In the object-oriented world, objects are in one class only.
 - ▶ Subclasses inherit properties from superclasses.
- ◆ In contrast, E/R entities have components in **all** subclasses to which they belong.
 - ▶ This matters when we convert to relations.

81

Example



82

Definition: Key

- ◆ A key is a **set of attributes** for one entity set such that no two entities in this set agree on all the attributes of the key.
 - ▶ It is allowed for two entities to agree on some, but not all, of the key attributes.
- ◆ We must designate a key for every entity set.


83

Keys in E/R Diagrams

- ◆ **Underline** the key attribute(s).
- ◆ In an Isa hierarchy, only the root entity set has a key, and it must serve as the key for all entities in the hierarchy.


84

Example: *name* is Key for *Beers*



85

Example: a Multi-attribute Key



- Note that *hours* and *room* could also serve as a key, but we must select only one key.

86

Definition: Weak Entity Set

- ◆ Occasionally, entities of an entity set need “help” to identify them uniquely.
- ◆ Entity set *E* is said to be **weak** if in order to identify entities of *E* uniquely, we need to follow one or more many-one relationships from *E* and **include the key of the related entities** from the related entity sets.


87

Example : Weak Entity Set

- ◆ ***name*** is almost a key for football players, but there might be two with the same name.
- ◆ ***number*** is certainly not a key, since players on two teams could have the same number.
- ◆ But ***number***, together with the ***Team*** related to the player by ***Plays-on*** should be unique.

88

In E/R Diagrams



- Double diamond for **supporting** many-one relationship.
- Double **rectangle** for the weak entity set.

89

Weak Entity-Set Rules

- ◆ A weak entity set has one or more many-one relationships to other (supporting) entity sets.
 - ◆ Not every many-one relationship from a weak entity set need be supporting.
- ◆ The **key for a weak entity set** is its own underlined attributes **together with** the keys for the supporting entity sets.
 - ◆ E.g., *player-number* and *team-name* is a key for *Players* in the previous example.

90

Design Guidelines

1. Avoid redundancy.
2. Limit the use of weak entity sets.
3. Don't use an entity set when an attribute will do.


91

Avoiding Redundancy

- ◆ Redundancy occurs when we say the same thing in two different ways.
- ◆ Redundancy wastes space and (more importantly) encourages inconsistency.
 - ▶ The two instances of the same fact may become inconsistent if we change one and forget to change the other, related version.

92


Example: Good



This design gives the address of each manufacturer exactly once.

93


Example: Bad



This design states the manufacturer of a beer twice: as an attribute and as a related entity.

94

Example: Bad



This design repeats the manufacturer's address once for each beer; loses the address if there are temporarily no beers for a manufacturer.


95

Entity Sets Versus Attributes

- ◆ An entity set should satisfy at least one of the following conditions:
 - ▶ It is **more than a name** of something; it has at least one nonkey attribute.
 - or**
 - ▶ It is the "many" in a many-one or many-many relationship.

96


Example: Good



- *Manfs* deserves to be an entity set because of the nonkey attribute *addr*.
- *Beers* deserves to be an entity set because it is the "many" of the many-one relationship *ManfBy*.

97


Example: Good



There is no need to make the manufacturer an entity set, because we record nothing about manufacturers besides their name.

98

Example: Bad



Since the manufacturer is nothing but a name, and is not at the "many" end of any relationship, it should not be an entity set.

99

Exception to the Previous Rule [by RK]

- ◆ We may want to keep something as an entity even if it only has one attribute when:
 - ▶ The something **must stay represented in the database** even when not a current attribute value of any other entity.
 - For example, to offer a menu of choices.
 - ▶ To allow for evolution: We may plan to add more attributes later.

100

Don't Overuse Weak Entity Sets

- ◆ Beginning database designers often overlook that an entity could be a key by itself.
 - ▶ They make all entity sets weak, supported by all other entity sets to which they are linked.
- ◆ In reality, we often create unique ID's for entities and these serve as keys.
 - ▶ Examples include social-security numbers, automobile VIN's etc.

101

When Do We Need Weak Entity Sets?

- ◆ The usual reason is that there is **no global authority** capable of creating unique ID's.
- ◆ Example: it is unlikely that there could be an agreement to assign unique player numbers across all football teams in the world.

102

Additional Design Guidelines

- ◆ We will be able to propound additional guidelines within the context of the more-precise relational model.
- ◆ Many of these are expressible using the concept of functional dependency and normal forms.

103