COP 4710: Database Systems Spring 2004

-Day 3 – January 12, 2004 – Introduction to Database Systems

Instructor : Mark Llewellyn markl@cs.ucf.edu CC1 211, 823-2790 http://www.cs.ucf.edu/courses/cop4710/spr2004

School of Electrical Engineering and Computer Science University of Central Florida

COP 4710: Database Systems (Day 3)

Page 1

Database Design (cont.)

- For the system to be acceptable to the end-users, the database design activity is crucial.
- A poorly designed database will generate error that may lead to bad decisions being made, which may have serious repercussions for the organization. On the other hand, a well-designed database produces, in an efficient way, a system that provides the correct information for the decision-making process to succeed.



Roles in the Database Environment

Data and Database Administrators

- The Data Administrator (DA) is responsible for the management of the data resource including database planning, development and maintenance of standards, policies and procedures, and conceptual/logical database design.
- The Database Administrator (DBA) is responsible for the physical realization of the database, including physical database design and implementation, security and integrity control, maintenance of the operational system, and ensuring satisfactory performance of the applications for users. The role of the DBA is more technically oriented than that of the DA.



Roles in the Database Environment (cont.) <u>Database Designers</u>

- In large db design projects, we can distinguish between two types of designers: logical database designers and physical database designers.
 - Logical database designers are concerned with identifying the data (the entities and attributes), the relationships between the data, and the constraints on the data that will be stored in the database.
 - Physical database designers are highly dependent on the target DBMS, and there may be more than one way of implementing a mechanism. The physical db designer must be fully aware of the functionality of the target DBMS.



Roles in the Database Environment (cont.)

Application Developers

• Once the database has been implemented, the application programs that provide the required functionality for the end-users must be implemented. This is the responsibility of the application developers.



Roles in the Database Environment (cont.) End Users

- End users are the "clients" for the database and can be broadly categorized into two groups based upon how they utilize the system.
 - Naïve users are typically unaware of the DBMS. They access the database through specially written application programs which attempt to make the operations as simple as possible. They typically know nothing about the database or the DBMS.
 - Sophisticated users are familiar with the structure of the database and the facilities offered by the DBMS. They will typically use a high-level query language like SQL to perform their required operations and may even write their own application programs.

COP 4710: Database Systems (Day 3)

Page 6

Advantages of DBMSs

control of data redundancy	economy of scale
data consistency	balance of conflicting requirements
more information from same data	improved data accessibility
amount of data available	increased productivity
sharing of data	improved maintenance
improved data integrity	increased concurrency
improved data security	improved backup and recovery
enforcement of standards	improved responsiveness



COP 4710: Database Systems (Day 3)

Disadvantages of DBMSs

complexity	
size	
cost of DBMSs	
additional hardware costs	
cost of conversion	
performance (specific cases)	
higher impact of failure	

COP 4710: Database Systems (Day 3)

Page 8

Three-Levels of Abstraction in a Database System



The External Level

- The external level is the user's view of the database.
- This level describes that part of the database which is relevant to each user.
- The external level consists of a number of different external views of the db. Each user has a view of the "real world" represented in a form that is familiar for that user.
- The external view includes only those entities, attributes, and relationships in the "real world" that the user is interested in. Other entities, attributes, and relationships may exist, but the user will be unaware that they even exist.





The External Level (cont.)

- It is often the case that different external views will have different representations of the same data.
 - Example: one view may represent dates in the form of (month, day, year) while another view may represent dates in the form of (day, month, year).
- Some views may include derived or calculated data. This is data that is not actually stored in the database as such, but created when needed.
 - Example: one view may need to see a person's age. However, this is probably not a stored value in the db since it would require daily updates. Rather, it is probably derived from stored data representing the person's date of birth and the current date.





The Conceptual Level

- The conceptual level is the community view of the database. This level describes *what* data is stored in the database and the relationships among the data.
- This is the level at which the logical structure of the entire database as seen by the DBA is contained. It represents a complete view of the data requirements of the organization that is independent of any storage considerations.
- The conceptual level supports each external view, in that any data available to a user must be contained in, or derivable from, the conceptual level.
- This level contains no storage-dependent details.
 - For example, an entity may be defined as represented by an integer data type at this level, but the number of bytes it occupies is not specified at this level.

COP 4710: Database Systems (Day 3)

Page 12

The Internal Level

- The internal level represents the physical representation of the database on the computer. This level describes *how* the data is stored in the database.
- The internal level describes the physical implementation necessary to achieve optimal runtime performance and storage space utilization.
- It covers the data structures and file organizations used to store the data on the storage devices.
- It interfaces with the OS access methods (file management techniques for storing and retrieving data records) to place the data on the storage devices, build indexes, retrieve the data, and so on.





The Physical Level

- Below the internal level is the physical level that may be managed by the OS under the direction of the DBMS.
- The functions of the DBMS and the OS at the physical level are not clear cut and will vary from system to system.
- Some DBMSs take advantage of many of the OS access methods, while others will use only the most basic ones and create their own file organizations.
- The physical level below the DBMS consists of items only the OS knows, such as exactly how the sequencing is implemented and whether the fields of internal records are stored as contiguous bytes on the disk.



Schemas, Mappings, and Instances

- The overall description of the database is called the database schema.
- There are three different types of schema in the database and these are defined according to the levels of abstraction of the three-level architecture.
 - At the highest level, there are multiple external schema. Also called subschemas, that correspond to different views of the data.
 - At the conceptual level, there is one conceptual schema, which describes all the entities, attributes, and relationships along with their integrity constraints.
 - At the lowest level of abstraction, there is one internal schema, which is a complete description of the internal model, containing the definition of the stored records, methods of representation, etc..



Schemas, Mappings, and Instances (cont.)

- The DBMS is responsible for mapping between these three types of schema. It must also check the schemas for consistency; in other words, the DBMS must check that each external schema is derivable from the conceptual schema, and it must use the information in the conceptual schema to map between each external schema and the internal schema.
- The conceptual schema is related to the internal schema through a conceptual/internal mapping. This enables the DBMS to find the actual record or combination of records in physical storage that constitute a logical record in the conceptual schema, together with any constraints to be enforced on the operations for that logical record.

Schemas, Mappings, and Instances (cont.)

• Each external schema is related to the conceptual schema by an external/conceptual mapping. This enables the DBMS to map names in the user's view on to the relevant part of the conceptual schema.



Schemas, Mappings, and Instances (cont.)



COP 4710: Database Systems (Day 3)

Page 18

Data Independence

- One of the major objectives of the three-level architecture is provide data independence, which means that the upper levels are unaffected by changes to lower levels.
- There are two types of data independence: logical and physical.
- Logical data independence refers to the immunity of the external schemas to changes in the conceptual schema.
- **Physical data independence** refers to the immunity of the conceptual schema to changes in the external schema.





Data Independence (cont.)



Database Languages

- A data sublanguage consists of two parts: a Data Definition Language (DDL) and a Data Manipulation Language (DML).
- The DDL is used to specify the database schema and the DML is used to both read and update the database.
- These languages are called *data sublanguages* because they do not include constructs for all computing needs such as conditional or iterative statements, which are provided by the high-level programming languages.
- Most DBMSs have a facility for embedding the sublanguage in a highlevel programming language such as COBOL, Pascal, C, C++, Java, or Visual Basic which is then called the *host language*.
- Most data sublanguages also provide a non-embedded or interactive version of the language to be input directly from a terminal.





The Data Definition Language (DDL)

- A Data Definition Language is a language that allows the DBA or user to describe and name the entities, attributes, and relationships required for the application, together with any associated integrity and security constraints.
- The result of the compilation/execution of the DDL statements is a set of tables stored in special files collectively referred to as the system catalog. The system catalog is also commonly referred to as the data dictionary or data directory.



The Data Manipulation Language (DML)

- A Data Manipulation Language is a language that provides a set of operations to support the basic data manipulation operations on the data held in the database.
- DML operations usually include the following:
 - insertion of new data into the database.
 - modification of data stored in the database.
 - retrieval of data contained in the database.
 - deletion of data from the database.
- The part of the DML that involves data retrieval is called a query language.



DMLs (cont.)

- DMLs are distinguished by their underlying retrieval constructs. We can distinguish two basic types of DMLs: procedural and non-procedural.
- Procedural DMLs are languages in which the user informs the system *what* data is required and exactly *how* to retrieve that data.
- Non-procedural DMLs are languages in which the user informs the system *only* of *what* data is required and leaves the how to retrieve the data entirely up to the system.
- It is common for procedural DMLs to be embedded in high-level programming languages.
- Procedural DMLs tend to be more focused on individual records while non-procedural DMLs tend to operate on sets of records.



Fourth Generation Languages

- There is no consensus as to what constitutes a 4GL. In essences it is a shorthand programming language. What requires several hundred lines of code in a 3GL will require only a few lines of code in a 4GL.
- 3GLs are procedural while 4GLs are non-procedural.
- 4GLs include spreadsheets and database languages.
- SQL and QBE are examples of 4GLs.

Data Models

- A data model is an integrated collection of concepts for describing and manipulating data, relationships between data, and constraints on the data in an organization.
- A model is a representation of "real world" objects and events, and their associations. It is an abstraction that concentrates on the essential, inherent aspects of an organization and ignores accidental properties.
- A data model must provide the basic concepts and notations that will allow database designers and end-users unambiguously and accurately to communicate their understanding of the organizational data.



Data Models (cont.)

- A data model can be thought of as comprising three components:
 - 1. A structural part, consisting of a set of rules according to which databases can be constructed.
 - 2. A manipulative part, defining the types of operations that are allowed on the data (this includes operations that are used for updating or retrieving data from the database and for changing the structure of the database).
 - 3. Possibly a set of integrity rules, which ensures that the data is accurate.



Data Models (cont.)

- Looking at the three level architecture, we can identify three different, related data models.
 - 1. An external data model to represent each user's view of the organization.
 - 2. A conceptual data model to represent the logical (or community view) that is DBMS independent.
 - 3. An internal data model to represent the conceptual schema in such a way that it can be understood by the DBMS.



Data Models (cont.)

- There have been many different data models which have been theorized, utilized, developed, and implemented over the years. They fall into three broad categories: object-based, recordbased, and physical.
- There are three principle record-based models: the relational data model, the network data model, and the hierarchical data model. Our focus will be on the relational data model in this course.



