

Database System Concepts and Architecture

406.426 Design & Analysis of Database Systems

Jonghun Park

jonghun@snu.ac.kr

Dept. of Industrial Engineering

Seoul National University

outline

- data models and their categories
- history of data models
- schemas, instances, and states
- three-schema architecture
- data independence
- DBMS languages and interfaces
- database system utilities and tools
- centralized and client-server architectures
- classification of DBMSs



data models

- DB approach provides some level of **data abstraction** by **hiding details of data storage** that are not needed by most DB users
- data model: a collection of concepts to describe the **structure** of a database, and certain **constraints** that the database should obey
- structure of a DB: data types, relationships, and constraints
- data model operations
 - **basic** operations: for specifying **retrievals** and **updates** on the DB
 - **user-defined** operations: for specifying the **behavior** of a database application (e.g., COMPUTE_GPA)

categories of data models

- **conceptual** (high-level, semantic) data models
 - provide concepts that are close to the way many users perceive data
 - uses concepts such as **entities**, **attributes**, and **relationships**
- **physical** (low-level, internal) data models
 - provide concepts that describe details of **how data is stored** in the computer
 - example: record formats, record orderings, access paths
- **representational** (implementation) data models
 - provide concepts that fall between the above two, balancing user views with some computer storage details
 - example: **relational**, **network**, and **hierarchical** models

schemas versus instances

- database schema (aka **intension**):
 - description of a database
 - includes descriptions of the database **structure** and the **constraints** that should hold on the database
 - not expected to change frequently
- schema diagram
 - a diagrammatic display of (some aspects of) a database schema
- schema construct
 - a **component** of the schema or an **object** within the schema
 - e.g., STUDENT, COURSE.
- database state (aka **extension** of a schema)
 - the actual data stored in a database at a particular moment in time.
 - also called the current set of occurrences or instances

example of a schema diagram

STUDENT

Name	Student_number	Class	Major
------	----------------	-------	-------

COURSE

Course_name	Course_number	Credit_hours	Department
-------------	---------------	--------------	------------

PREREQUISITE

Course_number	Prerequisite_number
---------------	---------------------

SECTION

Section_identifier	Course_number	Semester	Year	Instructor
--------------------	---------------	----------	------	------------

GRADE_REPORT

Student_number	Section_identifier	Grade
----------------	--------------------	-------

Figure 2.1

Schema diagram for the database in Figure 1.2.

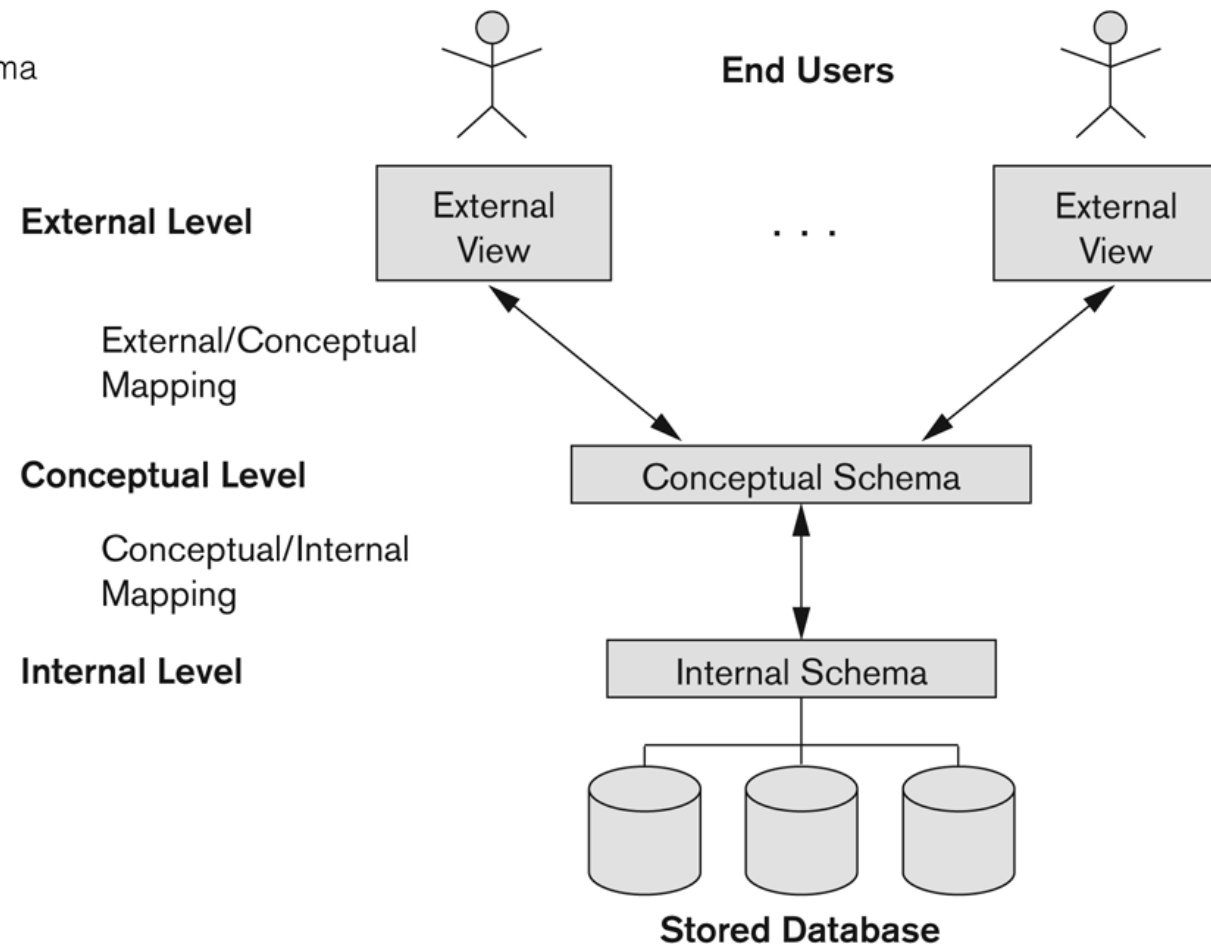
database schema vs. database state

- **DB state**: refers to the **content** of a database at a moment in time
- **initial** database state: refers to the database when it is loaded
- **valid** state: a state that satisfies the **structure** and **constraints** specified in the schema
- **meta-data**: the descriptions of the **schema constructs** and **constraints**
- schema evolution
 - most modern DBMS include some operations for schema evolution that can be applied while the database is operational

three-schema architecture

- goal: to separate the user applications and the physical database
- benefits of “layering”

Figure 2.2
The three-schema architecture.



three-schema architecture

- internal level: has an **internal schema** which describes the **physical storage structure** of the DB
- conceptual level
 - has a **conceptual schema**, which describes the **structure of the whole database** for a community of users
 - conceptual schema **hides the details of physical storage** structures and concentrates on describing **entities, data types, relationships, user operations, and constraints**
- external level (view level)
 - includes a number of external schemas or **user views**
 - each external schema describes **the part of the database** that a particular user group is interested in and **hides** the rest of the database from that user group
- **mappings** among schema levels are needed to transform requests and data -> overhead of DBMS approach!

data independence

- **data independence**: the capacity to **change the schema at one level** of a database system without having to change the schema at the next higher level
- logical data independence:
 - the capacity to **change the conceptual schema** without having to change the external schemas and their application programs
 - e.g., expand the DB, change constraints, reduce the DB
- physical data independence
 - the capacity to **change the internal schema** without having to change the conceptual schema
 - sometimes done to improve the performance

data independence

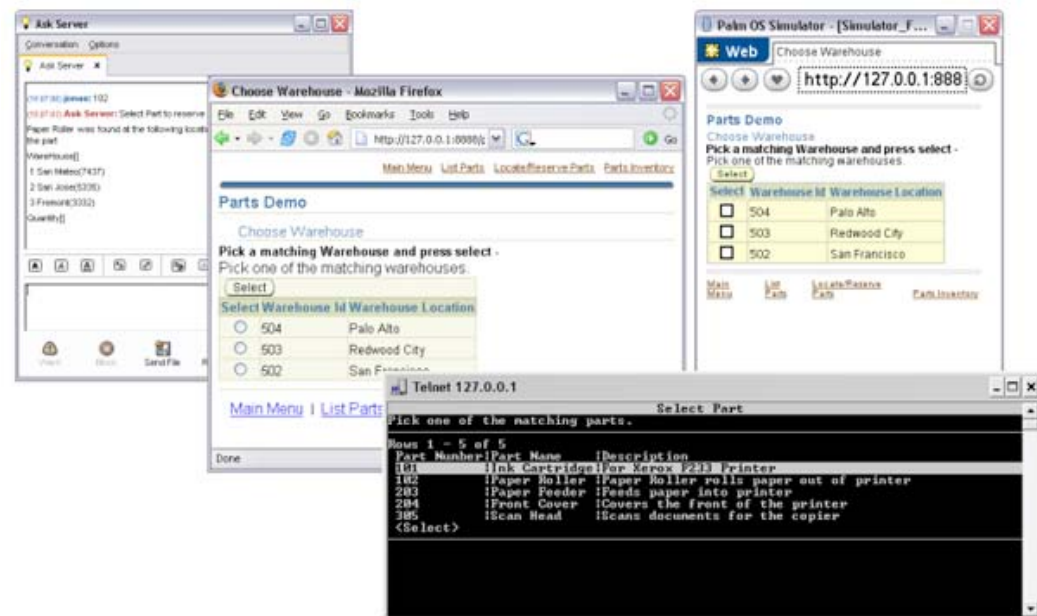
- when a schema at a lower level is changed, **only the mappings** between this schema and higher-level schemas **need to be changed** in a DBMS that fully supports data independence -> **the higher-level schemas themselves are unchanged!**
- hence, the application programs need not be changed since they refer to the external schemas
- however, two levels of mappings create an **overhead** during compilation or execution of a query or program
 - few DBMSs have implemented the full three-schema architecture

DBMS languages

- data definition language (DDL)
 - used by the DBA and database designers to specify the **conceptual, internal, and external schemas** of a database
- storage definition language (SDL)
 - specifies the internal schema
- view definition language (VDL)
 - specifies user views and their mappings to the conceptual schema
- data manipulation language (DML)
 - used to specify database manipulations
 - high-level DML: to specify complex DB operations in a concise manner (used by casual end users)
 - low-level DML: must be embedded in a general-purpose PL (used by programmers)
- **SQL**: represents a combination of DDL, VDL, and DML, as well as statements for constraint specification, schema evolution, and other features

DBMS interfaces

- menu-based interfaces for web clients or browsing
- forms-based interfaces
- graphical user interfaces
- natural language interfaces
- speech IO
- interfaces for parametric users (special interfaces): e.g., function keys
- interfaces for the DBA



DBMS component modules

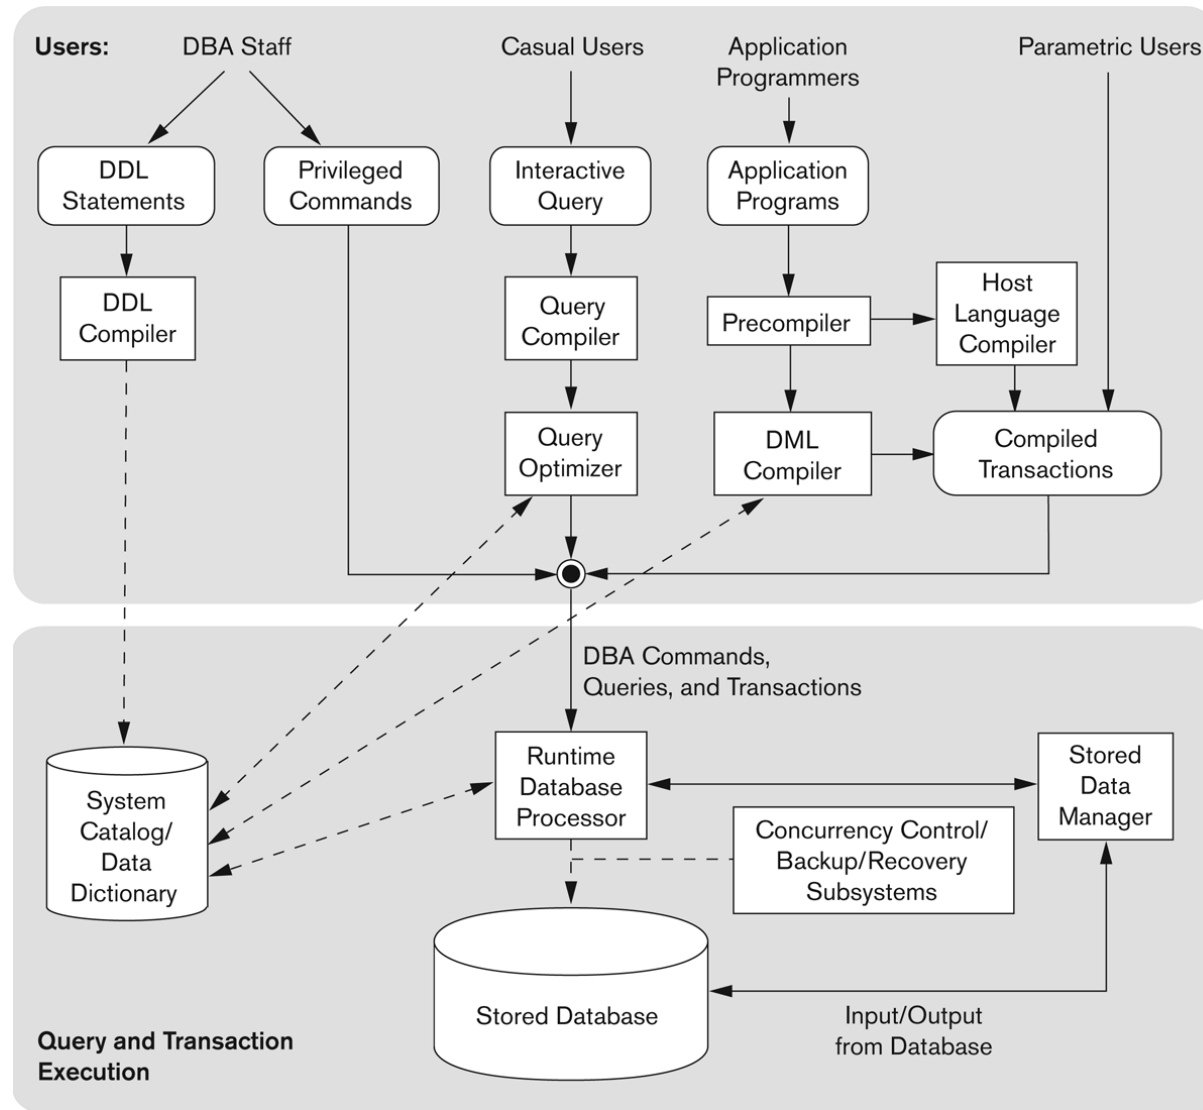


Figure 2.3

Component modules of a DBMS and their interactions.



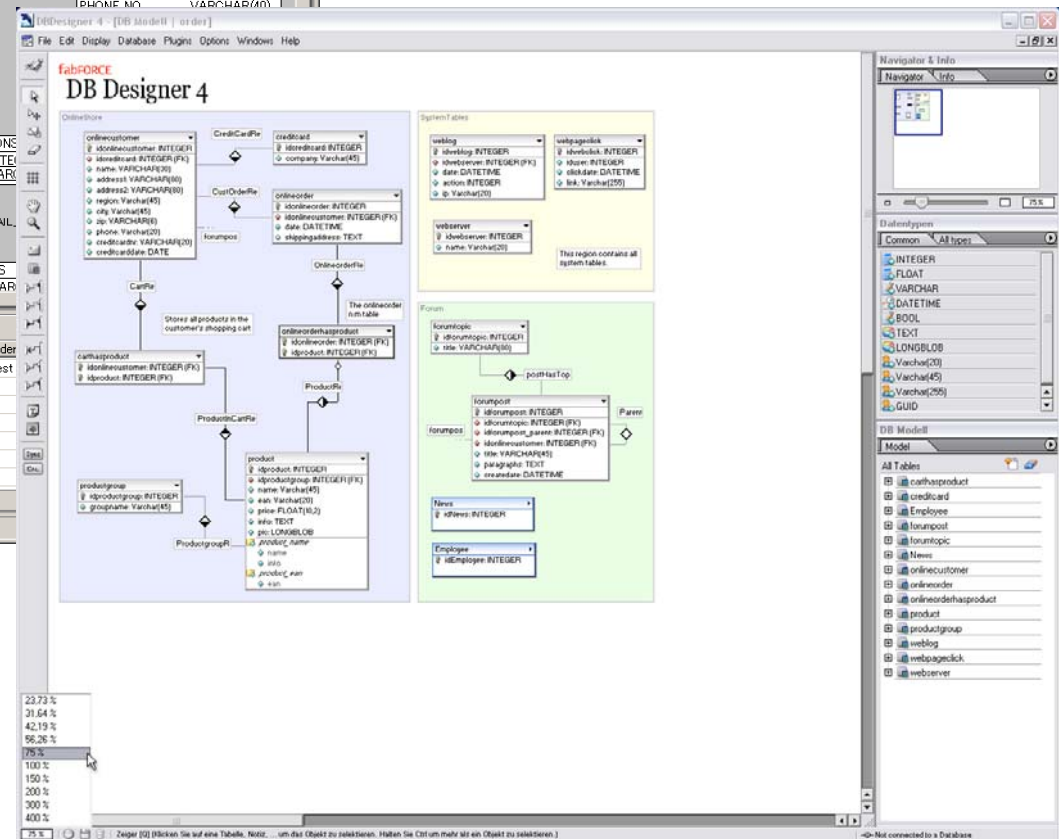
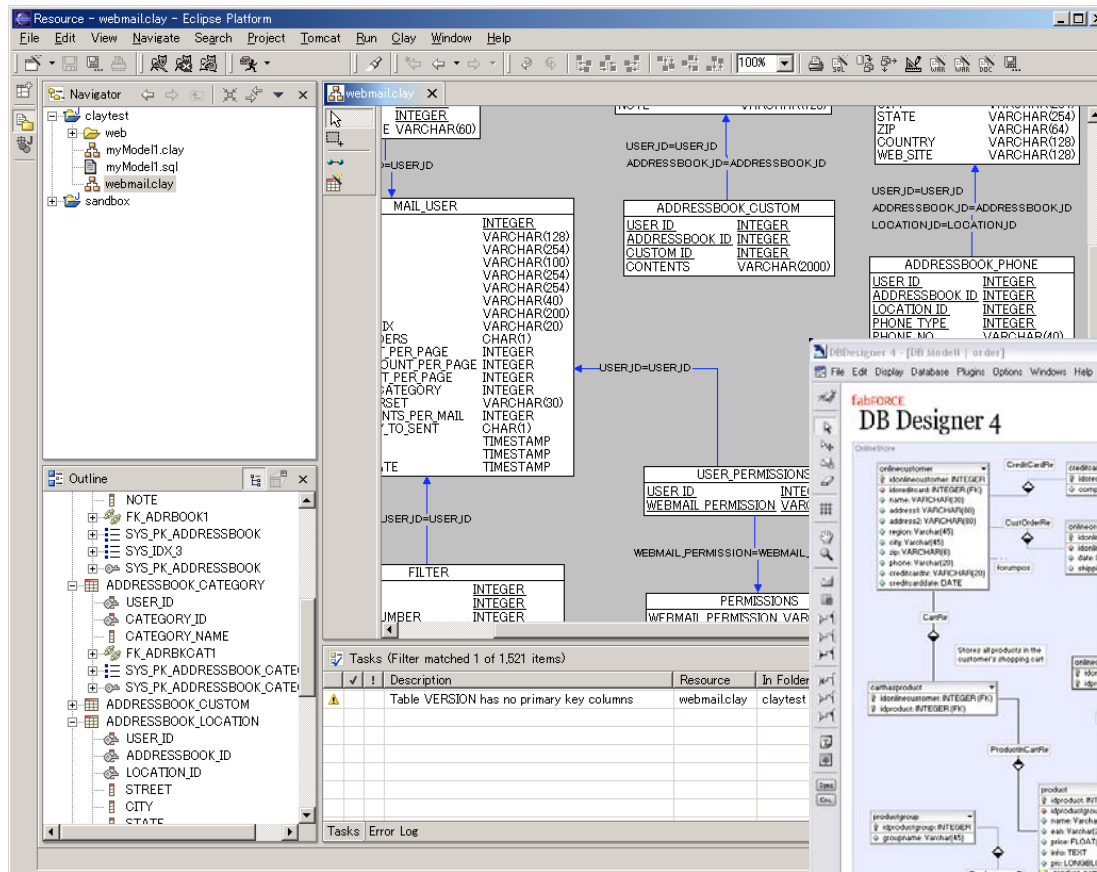
database system utilities

- help the DBA in **managing** the database system
- loading data stored in files into a database
 - includes data conversion tools
- backing up the database periodically on tape
- reorganizing database file structures
- report generation utilities
- performance monitoring utilities
- other functions, such as sorting, user monitoring, data compression, etc.

other tools

- data repository:
 - used to store other information such as design decisions, application program descriptions, user information, usage standards, etc., as well as **schema descriptions and constraints**
- application development environments and CASE (computer-aided software engineering) tools:
 - provide an environment for **developing database applications** and include facilities that help in many facets of database systems, including DB design, GUI development, querying and updating, and application program development
 - examples: PowerBuilder (Sybase), JBuilder (Borland)
- interface with communication software: APIs

examples



centralized and client-server architectures

- centralized DBMS
 - combines all functions into **single system** including – user application programs and user interface programs, as well as the DBMS functionality
- client / server architecture
 - developed to deal with computing environments in which a large number of PCs, workstations, file servers, printers, database servers, web servers, and other equipment are connected via a network
 - idea: to define **specialized servers** with specific functionalities (e.g., file server, printer server, email server, web server)
 - the client machines provide the user with the **appropriate interfaces** to utilize these servers, as well as with local processing power to run local applications
 - a server is a machine that can provide services to the client machines

a physical centralized architecture

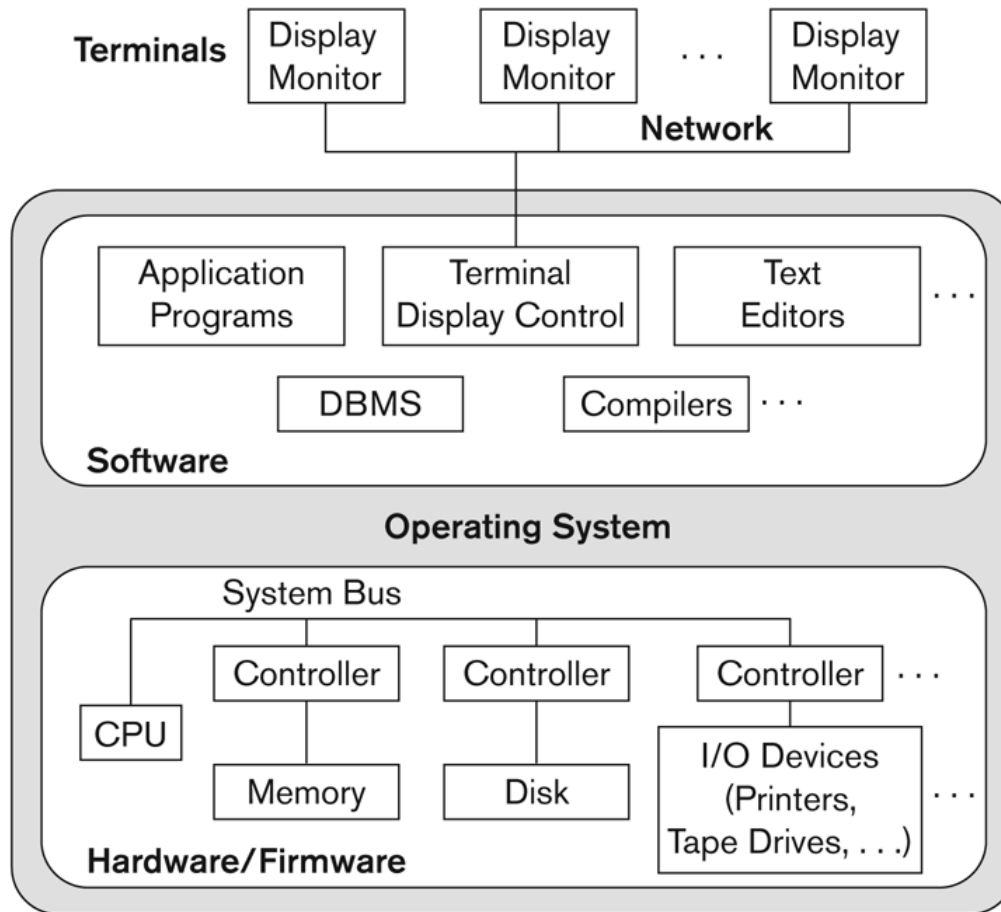


Figure 2.4
A physical centralized architecture.

```

ACADMVS                               Welcome to ASU Computing Services           03/11/99
P0500089                               * This system is only for use authorized by ASU *           13.40.49

                                           STATUS
1.  HELP
2.  Scheduled Service Interruptions
4.  VM/CMS - Information Center (mode ASUAC&D)  ACTIVE
7.  AccessCard                                ACTIVE
8.  Academic TSO (Node ACADMVS)                ACTIVE
11. General Purpose Server

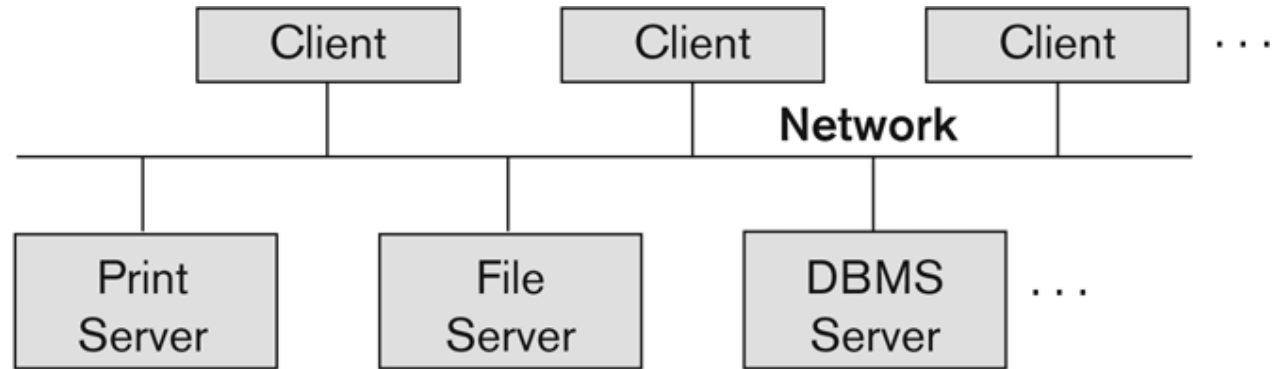
Request ==>                               Enter 1-11 or Press PF1-PF11

=====
                                           NUM           13:24:50 IBM-3278-2
Clear  Erase EOF  New Line  PA1    PA2    PA3
ATTN  Dup        Erase Input Field Mark  Reset  SysReq
    
```




logical two-tier CS architecture

Figure 2.5
Logical two-tier
client/server
architecture.



physical two-tier CS architecture



2-tier CS architectures for DBMSs

- client: **user interface** + **application programs**
- server: **query** and **transaction** functionality (hence the names such as query server, transaction server, and SQL server)
- ODBC
 - provides an **API** which allows client-side programs to call the DBMS, as long as both client and server machines have the necessary software installed
 - cf: JDBC
- advantages: simplicity and seamless compatibility with existing systems

Enter transactions from bank statement

Exit/Cancel Enter

Select Bank Account/Credit Card

Choose the bank account or credit card account which you are reconciling (eg Bank or Credit Card) Bank

Date	Description	Withdrawal (Shown as Debit)	Deposit (Shown as Credit)
------	-------------	-----------------------------	---------------------------

(Note that for a credit card the balance shows the amount owing on the card) Current balance 0.00

Transaction details

Enter the full amount of the transaction including VAT (if applicable) Get date 6 April 2004

Transaction description:

Deposit or Withdrawal?

This is a deposit into the bank account VAT is applicable

This is a withdrawal from the bank account VAT is not applicable

Select the balancing account

Choose the balancing account for this transaction - for a deposit this may be an Income like "Sales" for example. For a withdrawal it may be an Expenditure, like "Fuel"

Select Account type: All Select Account Name:

Microsoft Access

File Edit View Insert Format Records Tools Window Help

HMMS : Database (Access 2000 file format)

HMMS_HOL : Table

HOL_DATE	HOL_DESCR	UPD	UPDT_DATE	PR	PREV_UPDT
1/1/1999	NEW YEAR'S DAY	KIN	2/22/1999	HM	1/14/1999
1/18/1999	MARTIN LUTHER KING'S BIRTHDAY	KIN	2/22/1999		
4/26/1999	CONFEDERATION DAY				
5/31/1999	NATIONAL MEMORIAL DAY				
7/5/1999	INDEPENDENCE DAY				
9/6/1999	LABOR DAY				
10/11/1999	COLUMBUS DAY				
11/11/1999	VETERAN'S DAY				
11/25/1999	THANKSGIVING				
11/26/1999	ROBERT E. LEE BIRTHDAY				
12/25/1999	CHRISTMAS DAY				
12/27/1999	WASHINGTON BIRTHDAY				
12/31/1999	NEW YEAR'S DAY				
1/17/2000	MARTIN LUTHER KING'S BIRTHDAY				
4/24/2000	CONFEDERATION DAY				
5/29/2000	NATIONAL MEMORIAL DAY				
7/3/2000	INDEPENDENCE DAY	CAI	300 9:22:53 AM		
7/4/2000	INDEPENDENCE DAY	CAI	300 9:23:09 AM		
9/4/2000	LABOR DAY	CAI	300 9:23:25 AM		
10/9/2000	COLUMBUS DAY	CAI	300 9:23:51 AM		
11/10/2000	VETERAN'S DAY	CAI	300 9:24:22 AM		

Record: 14 of 73

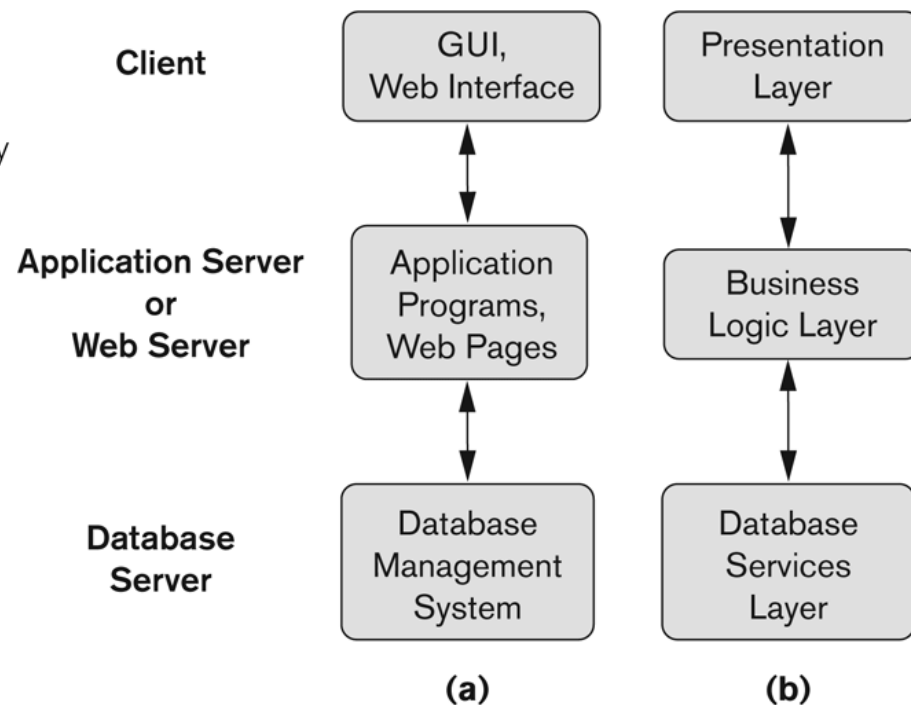
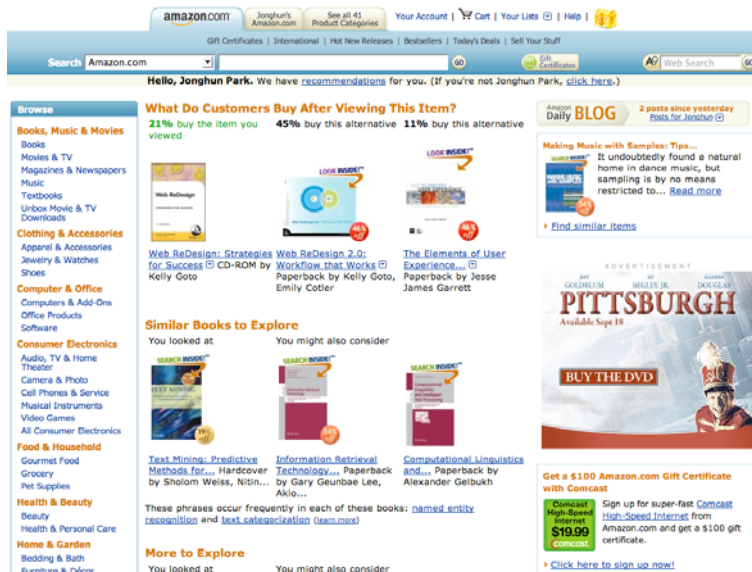
Datsheet View

Start Microsoft Access 10:45 AM

3-tier CS architectures for web applications

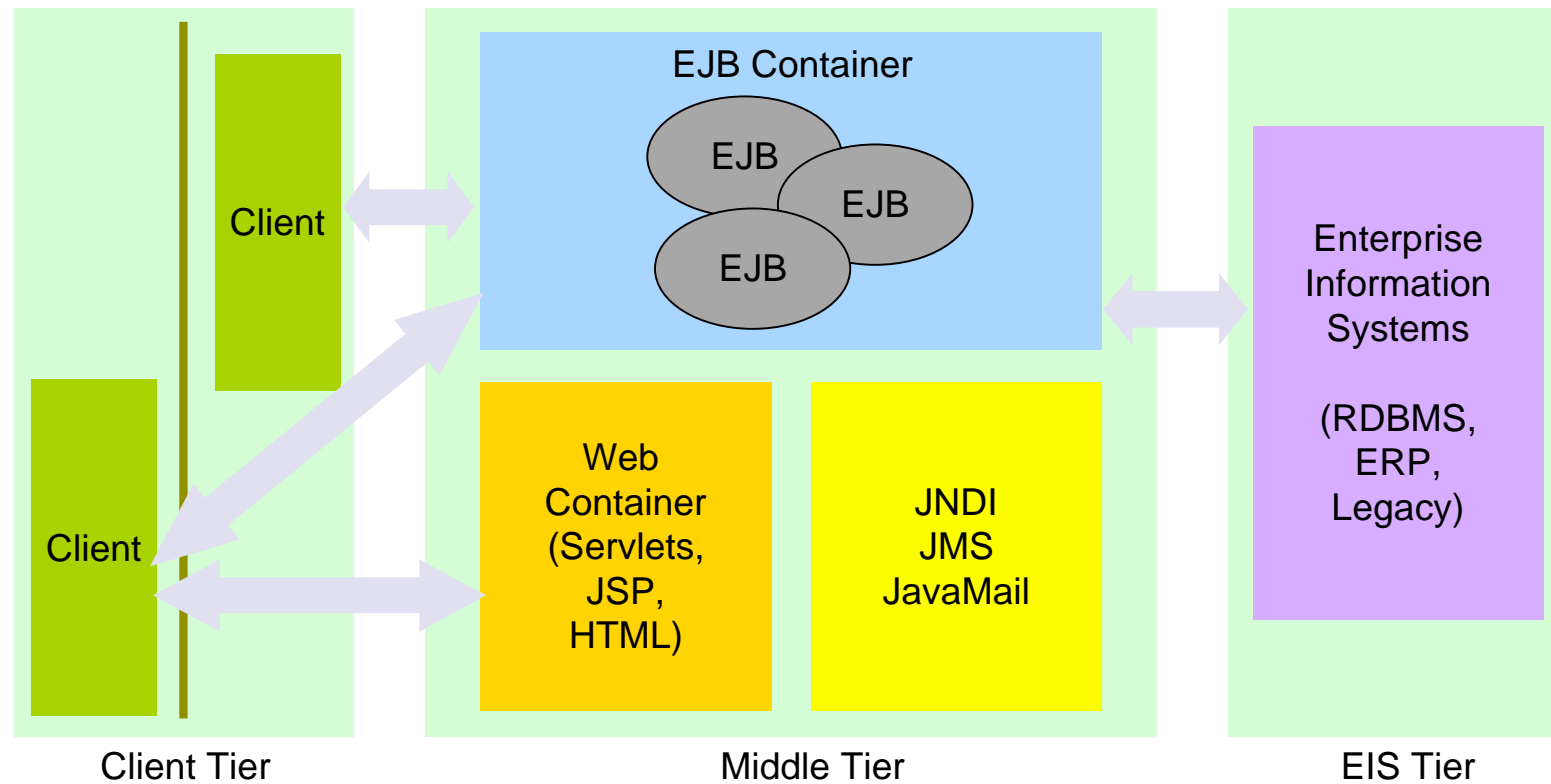
- many **web applications** use 3-tier CS architecture
- adds an intermediate layer between the client and the DB server
- application server: plays an intermediary role by storing **business logic** that are used to access data from the DB server

Figure 2.7
Logical three-tier client/server architecture, with a couple of commonly used nomenclatures.



J2EE platform

- a Java-based reference platform for developing and integrating enterprise applications
- a collection of standards for enterprise services



MySQL (www.mysql.com)

- world's most popular open source database
 - over 100 million copies of its software downloaded or distributed
- MySQL AB
 - the company founded by the creators of the MySQL
 - set up in Sweden by two Swedes and a Finn



classification of DBMSs

- based on the data model used:
 - traditional: relational, network, hierarchical
 - emerging: object-oriented, object-relational, XML
- other classifications:
 - single-user (typically used with micro- computers) vs. multi-user (most DBMSs).
 - centralized (uses a single computer with one database) vs. distributed (uses multiple computers, multiple databases)
 - homogeneous vs. heterogeneous (federated)
 - cost of DBMS
 - general purpose vs. special purpose (e.g., OLTP systems)
 - ...