



The picture can't be displayed.

Chapter 1: Introduction

The original presentation is infused with more information and slides by Verena Kantere

Database System Concepts, 6th Ed.

©Silberschatz, Korth and Sudarshan
See www.db-book.com for conditions on re-use

1



Outline

- The Need for Databases
- Data Models
- Relational Databases
- Database Design
- Storage Manager
- Query Processing
- Transaction Manager

2



Database Management System (DBMS)

- DBMS contains information about a particular enterprise
 - Collection of interrelated data
 - Set of programs to access the data
 - An environment that is both *convenient* and *efficient* to use
- Database Applications:
 - Banking: transactions
 - Airlines: reservations, schedules
 - Universities: registration, grades
 - Sales: customers, products, purchases
 - Online retailers: order tracking, customized recommendations
 - Manufacturing: production, inventory, orders, supply chain
 - Human resources: employee records, salaries, tax deductions
- Databases can be very large
- Databases touch all aspects of our lives

3



University Database Example

- Application program examples
 - Add new students, instructors, and courses
 - Register students for courses, and generate class rosters
 - Assign grades to students, compute grade point averages (GPA) and generate transcripts
- In the early days, database applications were built directly on top of file systems

4



Drawbacks of using file systems to store data

- Data redundancy and inconsistency
 - Multiple file formats, duplication of information in different files
- Difficulty in accessing data
 - Need to write a new program to carry out each new task
- Data isolation
 - Multiple files and formats
- Integrity problems
 - Integrity constraints (e.g., account balance > 0) become “buried” in program code rather than being stated explicitly
 - Hard to add new constraints or change existing ones

5



Drawbacks of using file systems to store data (Cont.)

- Atomicity of updates
 - Failures may leave database in an inconsistent state with partial updates carried out
 - Example: Transfer of funds from one account to another should either complete or not happen at all
- Concurrent access by multiple users
 - Concurrent access needed for performance
 - Uncontrolled concurrent accesses can lead to inconsistencies
 - ▶ Example: Two people reading a balance (say 100) and updating it by withdrawing money (say 50 each) at the same time
- Security problems
 - Hard to provide user access to some, but not all, data

Database systems offer solutions to all the above problems

6



Levels of Abstraction

- **Physical level:** describes how a record (e.g., instructor) is stored.
- **Logical level:** describes data stored in database, and the relationships among the data.

type *instructor* = **record**

```
ID : string;  
name : string;  
dept_name : string;  
salary : integer;
```

end;

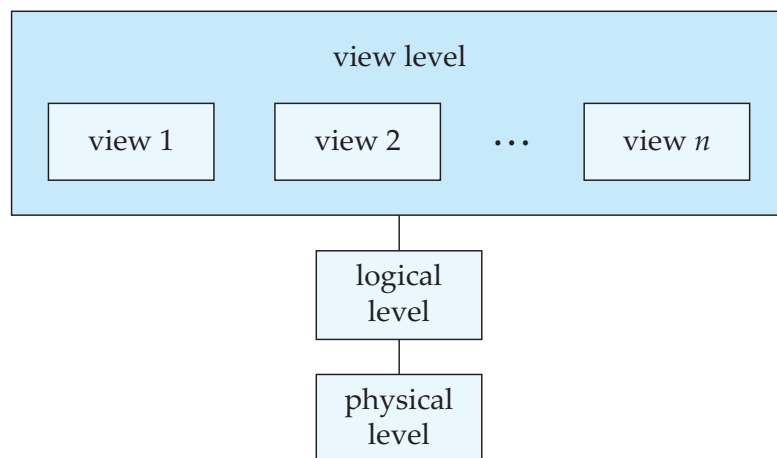
- **View level:** application programs hide details of data types. Views can also hide information (such as an employee's salary) for security purposes.

7



View of Data

An architecture for a database system



8



Instances and Schemas

- Similar to types and variables in programming languages
- **Logical Schema** – the overall logical structure of the database
 - Example: The database consists of information about a set of customers and accounts in a bank and the relationship between them
 - Analogous to type information of a variable in a program
- **Physical schema**– the overall physical structure of the database
- **Instance** – the actual content of the database at a particular point in time
 - Analogous to the value of a variable
- **Physical Data Independence** – the ability to modify the physical schema without changing the logical schema
 - Applications depend on the logical schema
 - In general, the interfaces between the various levels and components should be well defined so that changes in some parts do not seriously influence others.

9



Data Models

- A collection of tools for describing
 - Data
 - Data relationships
 - Data semantics
 - Data constraints
- Relational model
- Entity-Relationship data model (mainly for database design)
- Object-based data models (Object-oriented and Object-relational)
- Semistructured data model (XML)
- Other older models:
 - Network model
 - Hierarchical model

10



Relational Model

- All the data is stored in various tables.
- Example of tabular data in the relational model

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>
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

(a) The *instructor* table



A Sample Relational Database

<i>ID</i>	<i>name</i>	<i>dept_name</i>	<i>salary</i>
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

(a) The *instructor* table

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

(b) The *department* table



Data Definition Language (DDL)

- Specification notation for defining the database schema

Example: **create table** *instructor* (
 ID **char**(5),
 name **varchar**(20),
 dept_name **varchar**(20),
 salary **numeric**(8,2))

- DDL compiler generates a set of table templates stored in a **data dictionary**
- Data dictionary contains metadata (i.e., data about data)
 - Database schema
 - Integrity constraints
 - ▶ Primary key (ID uniquely identifies instructors)
 - Authorization
 - ▶ Who can access what



Data Manipulation Language (DML)

- Language for accessing and manipulating the data organized by the appropriate data model
 - DML also known as query language
- Two classes of languages
 - **Pure** – used for proving properties about computational power and for optimization
 - ▶ Relational Algebra
 - ▶ Tuple relational calculus
 - ▶ Domain relational calculus
 - **Commercial** – used in commercial systems
 - ▶ SQL is the most widely used commercial language



SQL

- The most widely used commercial language
- SQL is NOT a Turing machine equivalent language
- To be able to compute complex functions SQL is usually embedded in some higher-level language
- Application programs generally access databases through one of
 - Language extensions to allow embedded SQL
 - Application program interface (e.g., ODBC/JDBC) which allow SQL queries to be sent to a database

15



Database Design

The process of designing the general structure of the database:

- Logical Design – Deciding on the database schema. Database design requires that we find a “good” collection of relation schemas.
 - Business decision – What attributes should we record in the database?
 - Computer Science decision – What relation schemas should we have and how should the attributes be distributed among the various relation schemas?
- Physical Design – Deciding on the physical layout of the database

16



Database Design (Cont.)

- Is there any problem with this relation?

<i>ID</i>	<i>name</i>	<i>salary</i>	<i>dept_name</i>	<i>building</i>	<i>budget</i>
22222	Einstein	95000	Physics	Watson	70000
12121	Wu	90000	Finance	Painter	120000
32343	El Said	60000	History	Painter	50000
45565	Katz	75000	Comp. Sci.	Taylor	100000
98345	Kim	80000	Elec. Eng.	Taylor	85000
76766	Crick	72000	Biology	Watson	90000
10101	Srinivasan	65000	Comp. Sci.	Taylor	100000
58583	Califieri	62000	History	Painter	50000
83821	Brandt	92000	Comp. Sci	Taylor	100000
15151	Mozart	40000	Music	Packard	80000
33456	Gold	87000	Physics	Watson	70000
76543	Singh	80000	Finance	Painter	120000

17



Design Approaches

- Need to come up with a methodology to ensure that each of the relations in the database is “good”
- Two ways of doing so:
 - Entity Relationship Model
 - Models an enterprise as a collection of *entities* and *relationships*
 - Represented diagrammatically by an *entity-relationship diagram*
 - Normalization Theory
 - Formalize what designs are bad, and test for them

18



Object-Relational Data Models

- Relational model: flat, “atomic” values
- Object Relational Data Models
 - Extend the relational data model by including object orientation and constructs to deal with added data types.
 - Allow attributes of tuples to have complex types, including non-atomic values such as nested relations.
 - Preserve relational foundations, in particular the declarative access to data, while extending modeling power.
 - Provide upward compatibility with existing relational languages.



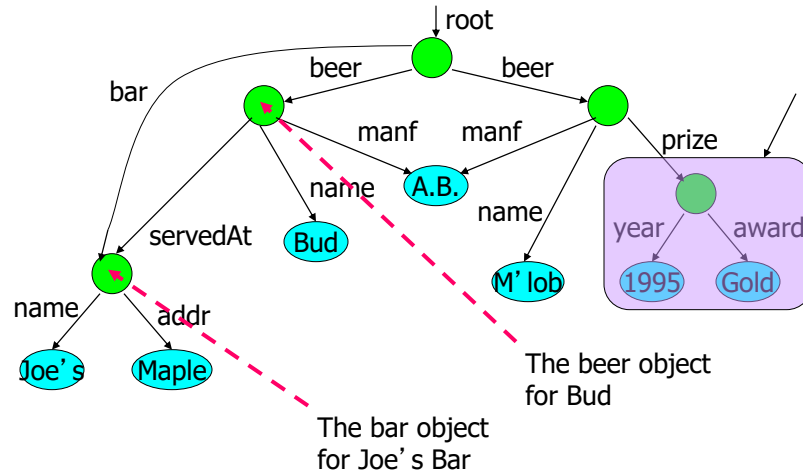
XML: Extensible Markup Language

- Defined by the WWW Consortium (W3C)
- Originally intended as a document markup language not a database language
- The ability to specify new tags, and to create nested tag structures made XML a great way to exchange **data**, not just documents
- XML has become the basis for all new generation data interchange formats.
- A wide variety of tools is available for parsing, browsing and querying XML documents/data

- XML is a language for semistructured data



Graph for semi-structured data



21



Example: Well-Formed XML

```
<? XML VERSION = "1.0" STANDALONE = "yes" ?>
```

```
<BARS>
```

```
<BAR><NAME>Joe's Bar</NAME>
```

```
<BEER><NAME>Bud</NAME>
<PRICE>2.50</PRICE></BEER>
```

```
<BEER><NAME>Miller</NAME>
<PRICE>3.00</PRICE></BEER>
```

```
</BAR>
```

```
<BAR> ...
```

```
</BARS>
```

22



Database Engine

- Storage manager
- Query processing
- Transaction manager

23



Storage Management

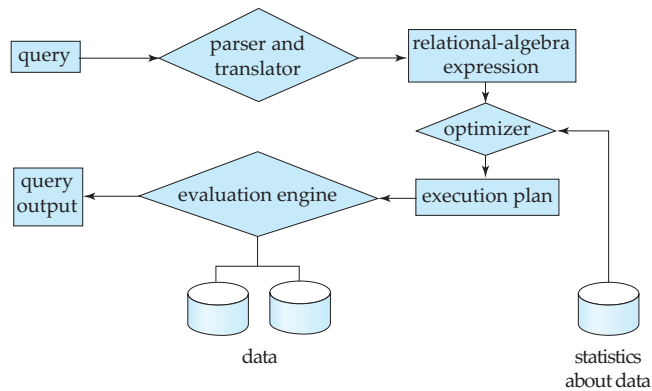
- **Storage manager** is a program module that provides the interface between the low-level data stored in the database and the application programs and queries submitted to the system.
- The storage manager is responsible to the following tasks:
 - Interaction with the OS file manager
 - Efficient storing, retrieving and updating of data
- Issues:
 - Storage access
 - File organization
 - Indexing and hashing

24



Query Processing

1. Parsing and translation
2. Optimization
3. Evaluation



25



Query Processing (Cont.)

- Alternative ways of evaluating a given query
 - Equivalent expressions
 - Different algorithms for each operation
- Cost difference between a good and a bad way of evaluating a query can be enormous
- Need to estimate the cost of operations
 - Depends critically on statistical information about relations which the database must maintain
 - Need to estimate statistics for intermediate results to compute cost of complex expressions

26



Transaction Management

- What if the system fails?
- What if more than one user is concurrently updating the same data?
- A **transaction** is a collection of operations that performs a single logical function in a database application
- **Transaction-management component** ensures that the database remains in a consistent (correct) state despite system failures (e.g., power failures and operating system crashes) and transaction failures.
- **Concurrency-control manager** controls the interaction among the concurrent transactions, to ensure the consistency of the database.

27



Concurrency Control

- Concurrent execution of user programs is essential for good DBMS performance.
 - Because disk accesses are frequent, and relatively slow, it is important to keep the CPU humming by working on several user programs concurrently.
- Interleaving actions of different user programs can lead to inconsistency: e.g., check is cleared while account balance is being computed.
- DBMS ensures such problems don't arise: users can pretend they are using a single-user system.

28



Transaction: An Execution Unit of a DB

- Key concept is transaction, which is an *atomic* sequence of database actions (reads/writes).
- Each transaction, executed completely, must leave the DB in a consistent state if DB is consistent when the transaction begins.
 - Users can specify some simple integrity constraints on the data, and the DBMS will enforce these constraints.
 - Beyond this, the DBMS does not really understand the semantics of the data. (e.g., it does not understand how the interest on a bank account is computed). *Why not?*
 - Thus, ensuring that a transaction (run alone) preserves consistency is ultimately the *user's* responsibility!



Scheduling Concurrent Transactions

- DBMS ensures that execution of $\{T_1, \dots, T_n\}$ is equivalent to some serial execution $T_1' \dots T_n'$.
 - Before reading/writing an object, a transaction requests a lock on the object, and waits till the DBMS gives it the lock. All locks are released at the end of the transaction. (Strict 2PL locking protocol.)
 - **Idea:** If an action of T_i (say, writing X) affects T_j (which perhaps reads X), one of them, say T_i , will obtain the lock on X first and T_j is forced to wait until T_i completes; this effectively orders the transactions.
 - What if T_j already has a lock on Y and T_i later requests a lock on Y ? **What is it called? What will happen?**



Ensuring Atomicity

- DBMS ensures *atomicity* (all-or-nothing property) even if system crashes in the middle of a Xact.
- **Idea:** Keep a *log* (history) of all actions carried out by the DBMS while executing a set of Xacts:
 - **Before** a change is made to the database, the corresponding log entry is forced to a safe location. (*WAL(Write-Ahead-Log) protocol*.)
 - After a crash, the effects of partially executed transactions are *undone* using the log. (Thanks to WAL, if log entry wasn't saved before the crash, corresponding change was not applied to database!)

31



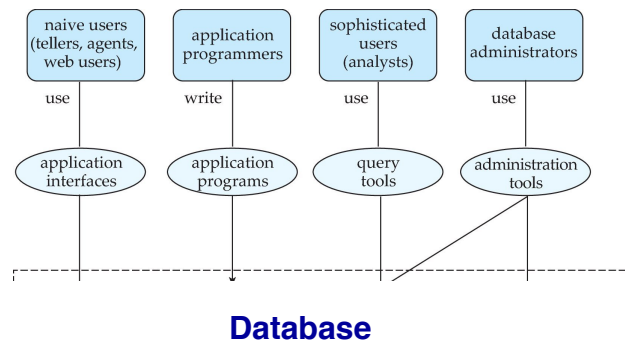
The Log

- The following actions are recorded in the log:
 - *Ti writes an object:* the old value and the new value.
 - ▶ Log record must go to disk *before* the changed page!
 - *Ti commits/aborts:* a log record indicating this action.
- Log records chained together by Xact id, so it's easy to undo a specific Xact (e.g., to resolve a deadlock).
- Log is often *duplexed* and *archived* on "stable" storage.
- All log related activities (and in fact, all CC related activities such as lock/unlock, dealing with deadlocks etc.) are handled transparently by the DBMS.

32



Database Users and Administrators



33



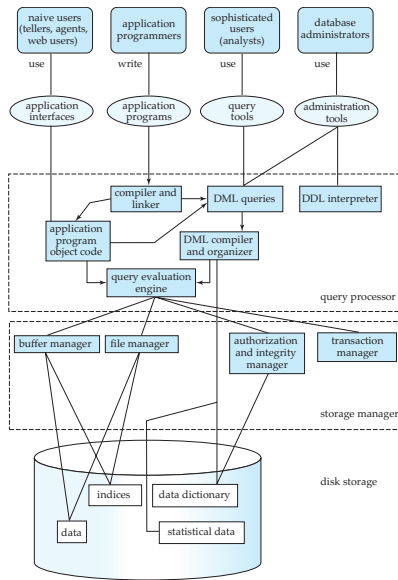
Databases make these folks happy ...

- End users and DBMS vendors
- DB application programmers
 - e.g. webmasters
- Database administrator (DBA)
 - Designs logical /physical schemas
 - Handles security and authorization
 - Data availability, crash recovery
 - Database tuning as needs evolve
 - *Must understand how a DBMS works!*

34



Database System Internals



Database System Concepts - 6th Edition

1.35

©Silberschatz, Korth and Sudarshan

35



Database Architecture

The architecture of a database systems is greatly influenced by the underlying computer system on which the database is running:

- Centralized
- Client-server
- Parallel (multi-processor)
- Distributed

Database System Concepts - 6th Edition

1.36

©Silberschatz, Korth and Sudarshan

36



History of Database Systems

- 1950s and early 1960s:
 - Data processing using magnetic tapes for storage
 - ▶ Tapes provided only sequential access
 - Punched cards for input
- Late 1960s and 1970s:
 - Hard disks allowed direct access to data
 - Network and hierarchical data models in widespread use
 - Ted Codd defines the relational data model
 - ▶ Would win the ACM Turing Award for this work
 - ▶ IBM Research begins System R prototype
 - ▶ UC Berkeley begins Ingres prototype
 - High-performance (for the era) transaction processing

37



History (cont.)

- 1980s:
 - Research relational prototypes evolve into commercial systems
 - ▶ SQL becomes industrial standard
 - Parallel and distributed database systems
 - Object-oriented database systems
- 1990s:
 - Large decision support and data-mining applications
 - Large multi-terabyte data warehouses
 - Emergence of Web commerce
- Early 2000s:
 - XML and XQuery standards
 - Automated database administration
- Later 2000s:
 - Giant data storage systems
 - ▶ Google BigTable, Yahoo PNuts, Amazon, ..

38




History (cont.)

- 1990-2000s (Meta-relational era)
 - It is the era of **COMPLEX ENTITIES** (engineering objects, multimedia, software objects)
 - Object-Relational Database Systems
 - Active Databases, Intelligent Systems, Multimedia databases
 - Client-Server
 - Data Warehouses
 - Data Mining
 - Distributed Databases, parallelization
 - ▶ DBMS in PCs
 - ▶ DBMS on the Internet (Web-based), Java, XML, ..., Cloud

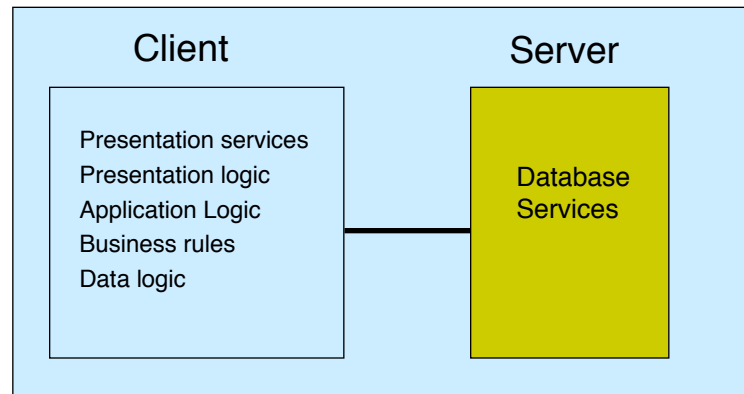


Relational and Object-oriented DBs





First generation of Client-Server



41



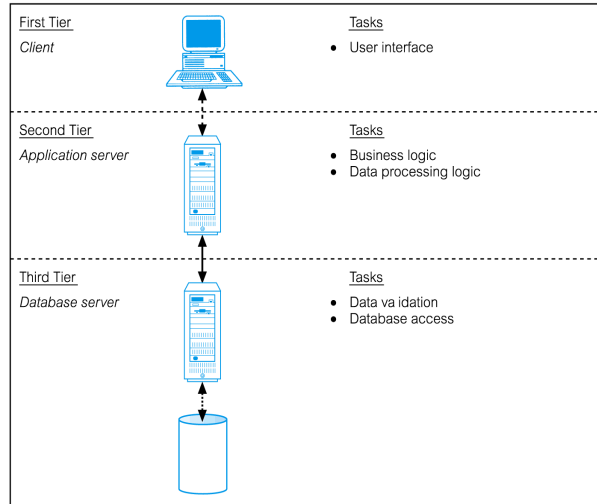
Architectural elements

- **Presentation services**
 - Take inputs and show results
- **Presentation logic**
 - Controls the relation user-application
- **Application logic**
 - Calculations, decisions, actions for the application
- **Business rules**
 - From the whole business
- **Data logic**
 - Creation of queries (e.g. in SQL)
- **Database services**
 - Service reference to data

42



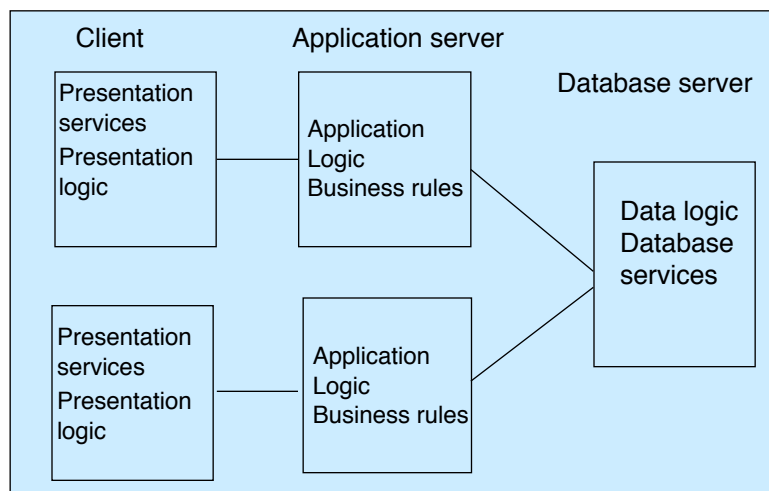
3-Tier architecture



43



Second generation Client-Server



44



Infrastructure and operating system

■ Client

- A processor
- PC (e.g. PowerPC, Power Mac)
- Windows XP, ... NT, ... 2000, Apple Mac, OS/2, Unix


■ Server

- 2 to 100 users
- Up to big machines (SMP (Symmetric Multiprocessors), clusters, MPP (Massively Parallel Processors))
- Windows NT, 2000, Novell Netware, Unix



Today's typical architecture

Clients





Technology for managing data

- **OLAP (Online Analytic Processing)**
 - Analysis in a multi-dimensional space
 - » Sales per (product, customer, time)
 - » Cube
 - drill down, rollup
 - Essbase, Commander, Oracle Express, SAS, Excel, SQL Server,...
- **Data Warehouses**
 - New form for Decision Support Systems
 - Redbrick, Oracle DW, Informix, Sybase, Micro Strategy
- **Data Marts (smaller Data Warehouses)**
- **Data Mining (extraction of information)**

47



OLAP - Applications

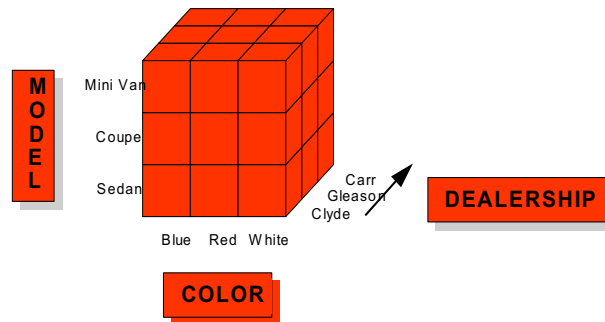
- Classical applications
 - Financial analysis, modeling ,reporting
 - Budgeting
 - Guarantee and control of Quality
 - Financial gain
 - Research analysis
 - Monitoring of products
- OLAP and OLTP cannot work at the same time on the same data!
 - Completely different and opposite requirements

48

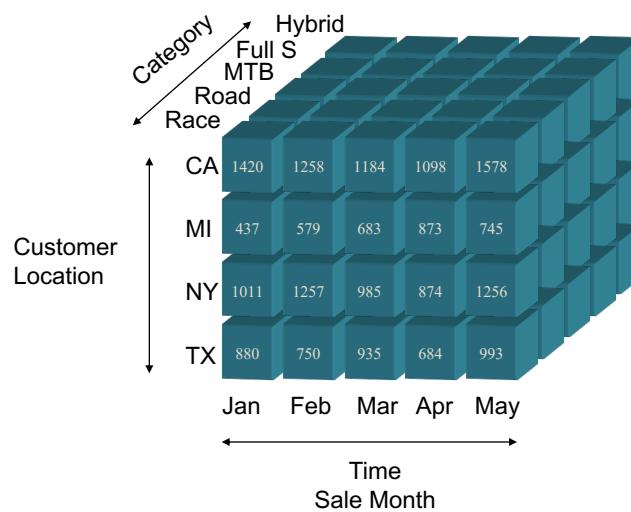


OLAP: The notion of cube

Sales Volumes

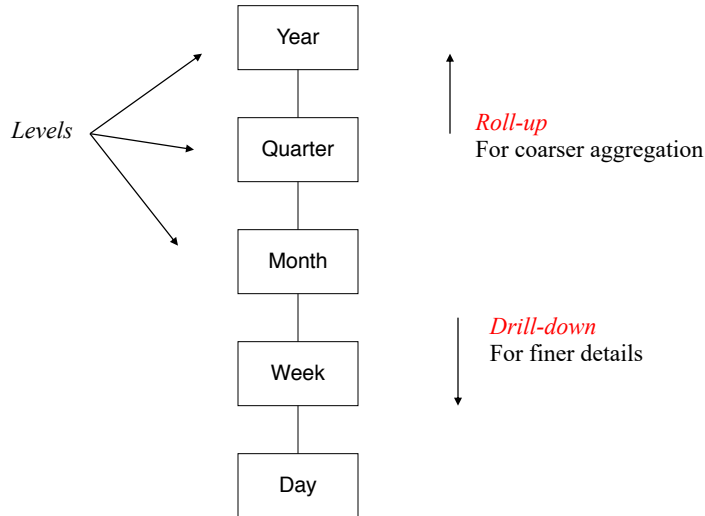


Example of Multidimensional cube





Basic actions / exploration



51



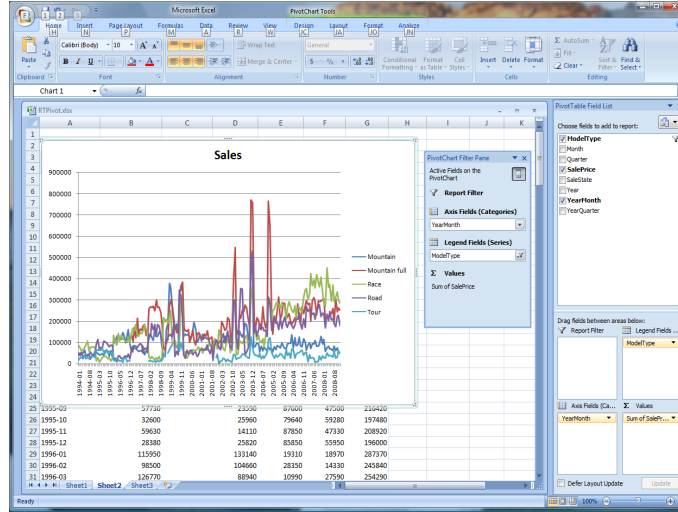
Microsoft Pivot Table

Year	YearMonth	Sum of SalesPrice
1994		\$599,020
1995		\$507,940
1996		\$1,469,870
1997		\$783,820
1998		\$1,366,180
1999		\$1,138,210
2000		\$1,627,070
2001		\$1,221,870
2002		\$1,523,740
2003		\$2,241,600
2004		\$886,390
2005		\$934,540
2006		\$1,205,230
2007		\$1,098,130
2007-01		\$63,590
2007-02		\$88,660
2007-03		\$98,780
2007-04		\$71,210
2007-05		\$92,080
2007-06		\$93,320
2007-07		\$109,880
2007-08		\$83,090
2007-09		\$62,370
2007-10		\$93,430
2007-11		\$128,150
2007-12		\$97,770
2008		\$846,690
Grand Total		\$10,511,290

52



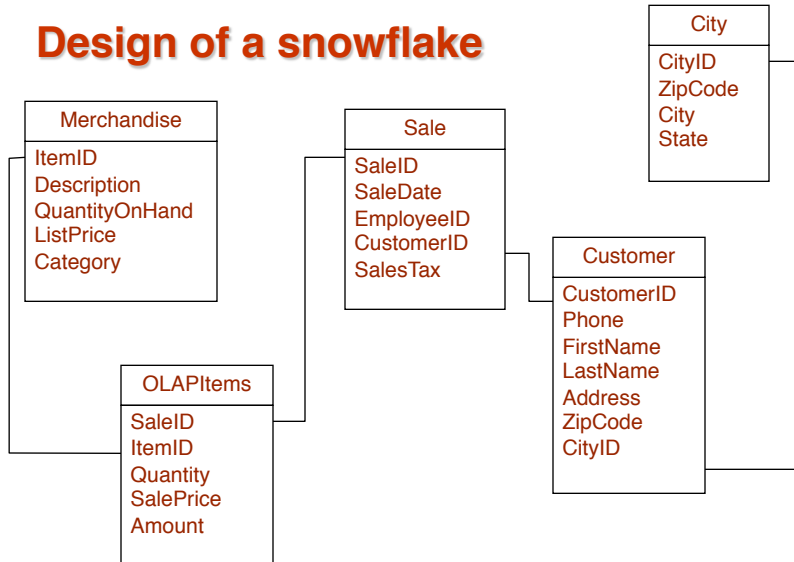
Microsoft Pivot Chart



53



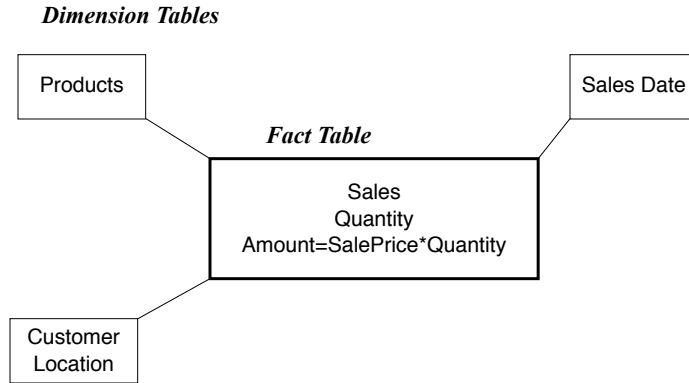
Design of a snowflake



54



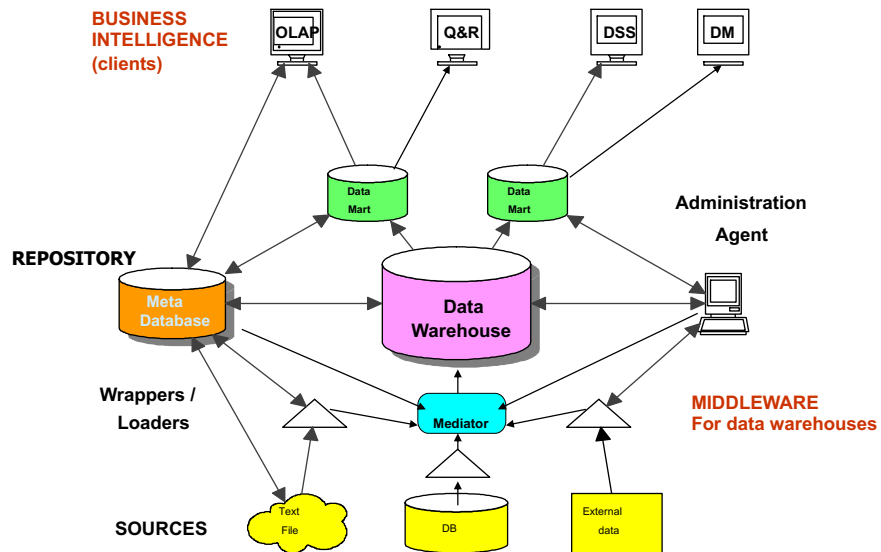
Design of a star



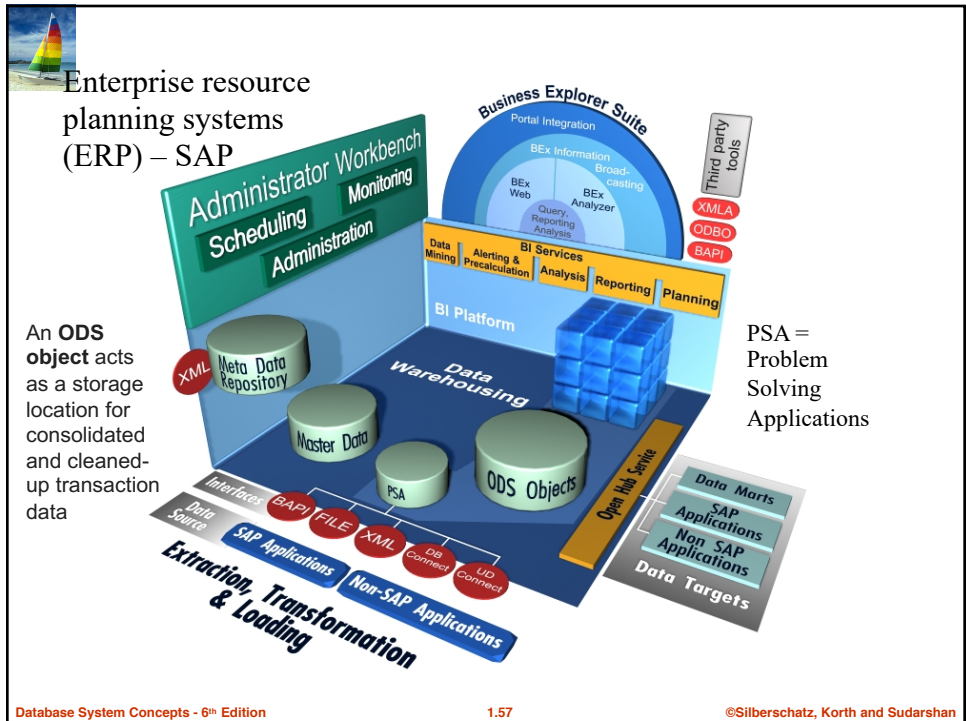
55



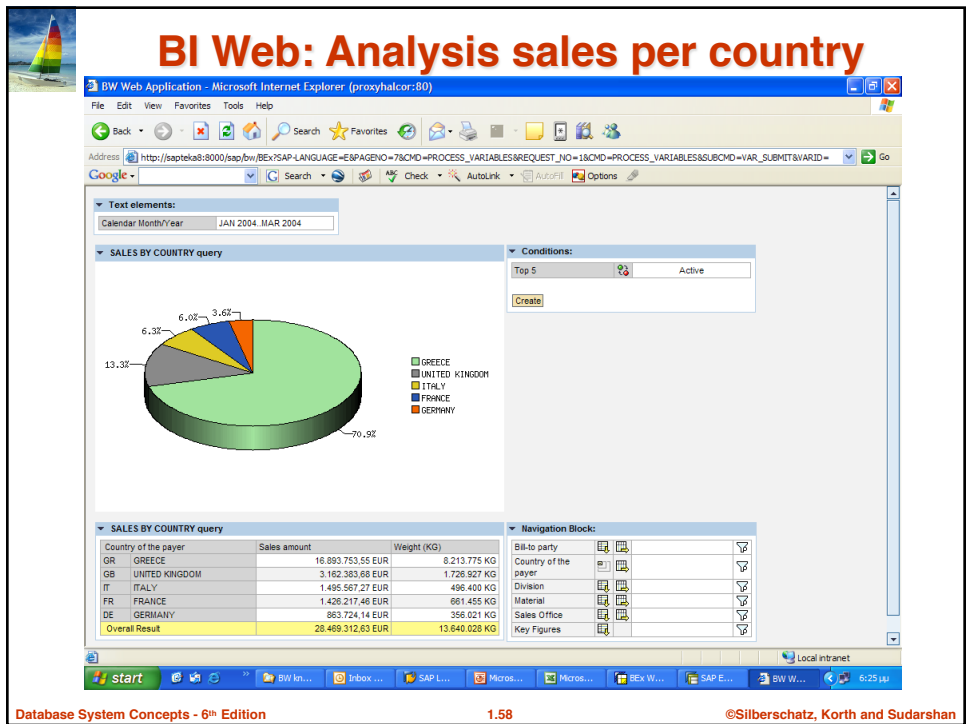
Classical architecture for data warehouse



56



57



58

BI Web: Employee carte

Employee Information

Employee	Entry Date	Organizational Unit	Job	Scale ID	Band / Level	Gender	Age in Years	Length of Service	Gross Salary
Last Name:0003717 First Name:0003717	02.11.2004	BUSINESS SYSTEMS INT	50029221	ACC MOT	2	Female	105,00	1,00	6.300,00 EUR

Personal Information

Address	City	Postal Code	Home Phone	Office Phone	Marital Status
ΔΙΕΥΘΥΝΣΗ ΚΗΡΚΙΔΑΣ 44	ΠΟΛΗ ΜΑΡΟΥΣΙ	11111	TEL0003717	6567	Marr.

Salary Comparison

Employee Gross Salary: 6300 Euro
Average Salary of BUSINESS SYSTEMS INT: 3161,63 Euro

Service Comparison

Employee Length of Service: 4 Years
Average Service of BUSINESS SYSTEMS INT: 4,55 Years

Education

Level	Institute	Country	Title	Certificate	Duration
ABT	ΠΑΝΕΠΙΣΤΗΜΙΟ ΚΡΕΤΤΕ	Greece	Political Science	Final certificate	4 Years

Foreign Languages

Title	Institute	Country	Certificate
ΑΤΤΑΚΑ - Proficiency	British Council	Greece	Final certificate

Database System Concepts - 6th Edition 1.59 ©Silberschatz, Korth and Sudarshan

59

Planning & Budgeting

Enter plan data

Close navigation | Planning profile | Global planning sequences | Set Variables

Planning areas | Technical name

- Actual & Budget data SALES03
 - ΕΙΣΟΤΕΡΙΚΟ PLO1
 - Copy actual to budget PP01
 - ΕΙΣΟΤΕΡΙΚΟ PLO2
 - Copy actual to budget PP01
 - Sales Budget SALES02
 - Εισοτήρι budget - ΕΙΣΟΤΕΡΙΚΟ PLO3
 - Budget - ΕΙΣΟΤΕΡΙΚΟ RPO1
 - Εισοτήρι budget - ΕΙΣΟΤΕΡΙΚΟ PLO4
 - Budget PP02
 - Υπολογισμός αξίας PLO5
 - Budget price X Actual Qty PP01
- Planning functions | Technical name
 - Εισοτήρι budget - ΕΙΣΟΤΕΡΙΚΟ PLO3
 - Manual planning 0-MP
 - Manual Planning MP01

Material group | E.E. (Ton) | E.E. (Tunit) | Curmco | TPITEE XOPEE (Ton) | TPITEE XOPEE (Tunit) | Curmco

Material group	E.E. (Ton)	E.E. (Tunit)	Curmco	TPITEE XOPEE (Ton)	TPITEE XOPEE (Tunit)	Curmco
01010	72,000	420,00	EUR	0,000	0,00	USD
01020	20,000	590,00	EUR	62,000	645,00	USD
01030	142,000	540,00	EUR	794,000	710,00	USD
01040	65,000	700,00	EUR	389,000	970,00	USD
01050	122,000	800,00	EUR	457,000	1,120,00	USD
01054	132,000	1,000,00	EUR	231,000	1,300,00	USD
01060	8,000	900,00	EUR	0,000	0,00	USD
01064	21,000	1,550,00	EUR	0,000	0,00	USD
01070	3,000	460,00	EUR	0,000	0,00	USD
01080	51,000	450,00	EUR	0,000	0,00	USD
01090	0,000	0,00	EUR	0,000	0,00	USD
01100	56,000	600,00	EUR	0,000	0,00	USD
01110	3,000	800,00	EUR	25,000	900,00	USD
01120	3,000	800,00	EUR	188,000	900,00	USD
01130	0,000	0,00	EUR	34,000	1,120,00	USD
01200	17,000	750,00	EUR	0,000	0,00	USD
02010	42,000	325,00	EUR	0,000	0,00	USD
02020	173,000	520,00	EUR	225,000	500,00	USD
02030	49,000	615,00	EUR	17,000	670,00	USD
02050	241,000	1,500,00	EUR	38,000	1,885,00	USD
02054	303,000	1,450,00	EUR	0,000	0,00	USD

Database System Concepts - 6th Edition 1.60 ©Silberschatz, Korth and Sudarshan

60



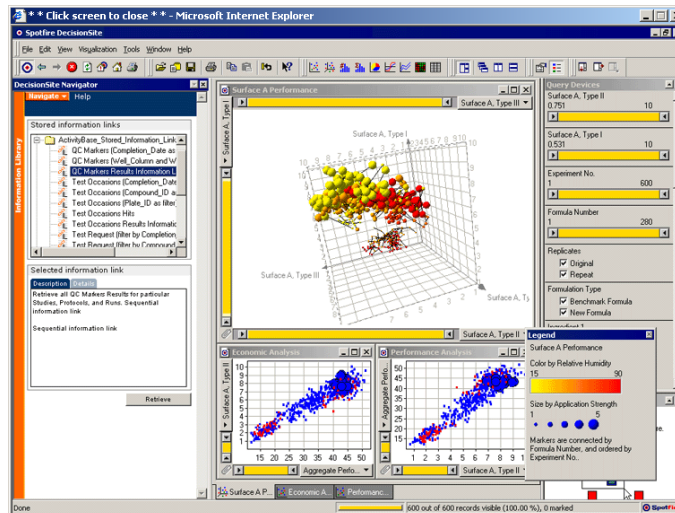
Data Mining

- Automatic analysis of data
- Statistics
 - Correlation
 - Regression (multiple correlation)
 - Clustering
 - Classification
 - Nonlinear relationships
- More automatizations
 - Analysis of the market basket
- Numerical data and non-numerical data
 - Language analysis

61



Data Mining Tools: Spotfire



<http://www.spotfire.com>

62



Market Basket Analysis

What do customers buy at the same time?



63



Data Mining: Market Basket Analysis

- **Goal: Measure the relation between two objects**
 - What products do customer buy together?
 - Which web pages do users visit together?
- **Classical examples**
 - In convenient stores that are open on Sunday, it was found that customers buy together beers and baby pampers.
 - Amazon.com: shows correlated purchases
- **Strategic use of such information**
 - Decide to place products together in order to increase cross selling
 - Alternatively, put them at the start an the end of the corridor in order for customers to buy even more products

64



DBMS : Old protagonists

- Systems that were popular in the 80s
 - IMS (IBM) -- Hierarchical Model (γλώσσα DL/1)
 - I-D-S (Honeywell) -- Network DBTG (Integrated Data Store)
 - IDMS (Cullinane) - Network (Integrated Data Mgmt System)
 - TOTAL (Cincom) - Network
 - IMAGE (Hewlett-Packard) - Network
 - SYSTEM 2000 (Intel-MRI) - Inverted (ad-hoc model)
Other Inverted: ADABAS (Software AG), Model 204 (CCA)
 - ...



DBMS : Old protagonists

- **SYBASE**
- **INGRES** now it is called Computer Associates-Ask Group and has become an open source system
- Others:
 - **Rdb, Gupta Quadbase, Ralma, Watcom, XDB, ...**



DBMS : Old protagonists

Today there are 3 very popular systems:

- **IBM DB2**, (also in Unix, Linux, Windows)
- **ORACLE 18 (18C)**, almost everywhere – first in the market especially for Unix and large installations
- **Microsoft SQL Server 2017**, in Microsoft platforms
 - **INFORMIX** (Bought by IBM!)



DBMS : Special protagonists

- Massive Parallel Processors (MPP) :
 - **Terradata (the biggest)**, **Tandem (NonStop SQL)**, **Oracle Parallel Server**, **Informix**, **Sybase (Navigator)**, **DB2**, **DEC,...(some already closed)**
- **The biggest in Windows / PC**
 - **MICROSOFT ACCESS**
 - **Powersoft, Gupta...**
 - (a) SQL access πρόσβαση (gateways)
 - (b) Excellent for Client-Server (DBMS)
 - (c) Look very much like the big DBMS



OODBMS Companies

GemStone Systems, Inc.
Hewlett-Packard, Inc. (OpenODB)
IBEX Corporation, SA.
Illustra (Informix, Inc.)
Matisse Software, Inc.
O2 Technology, Inc.
Objectivity, Inc.
Object Design, Inc.
ONTOS, Inc.
POET Software Corporation
UniSQL
Unisys Corporation (OSMOS)
Versant Object Technology

69

69



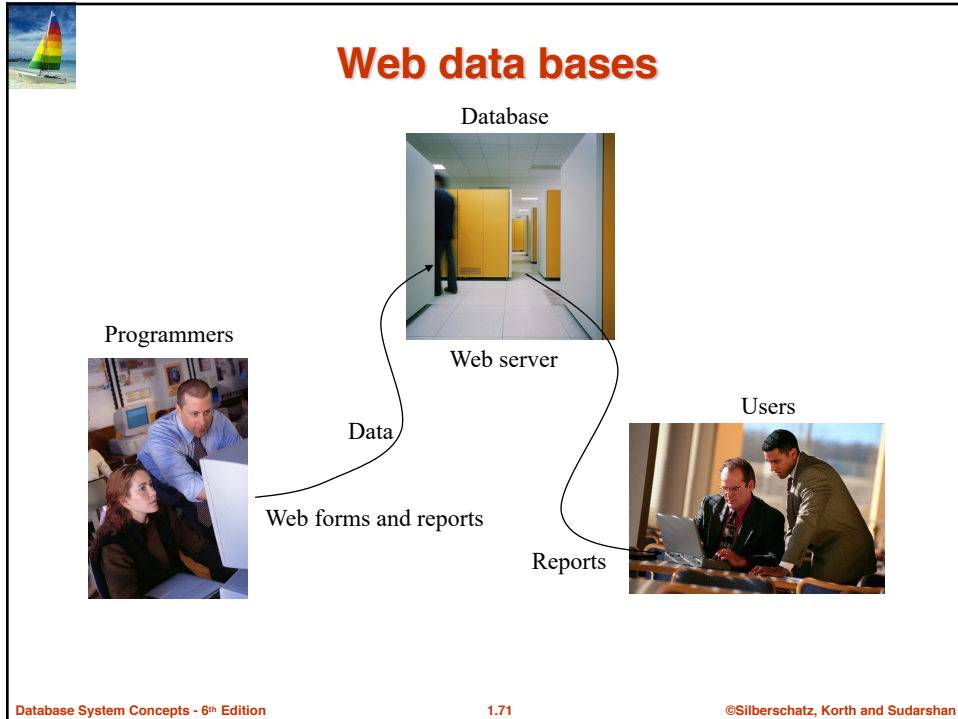
OPEN Source DBMS

- MySQL
- PostgreSQL + EnterpriseDB
- Berkeley DB, Firebird, etc.

There are also memory DBMSs, e.g.:

- SQLite (classical opensource DBMS)

70



71

-
- ## Web data bases
- More than 1B HTML pages, more than 15 TB
 - Extraordinary amount of information
 - Bookstores, restaurants, trip advisors, purchases, news, financial, guides, maps etc etc
 - **Many forms**: text, image, voice, video...
 - **Many formats**: HTML, XML, postscript, pdf, JPEG, MPEG, MP3
 - With special dynamic characteristics
 - More than 1M new pages a day
 - Graph structure among pages
 - 100M queries a day
- Database System Concepts - 6th Edition 1.72 ©Silberschatz, Korth and Sudarshan

72




BIG DATA – Characteristics

- Volume -- In 2000 we had about 800,000 petabytes (PB) stored data in the whole world. We anticipate that this number will reach 35 zettabytes (ZB) until 2020. (Twitter alone produces more than 7 terabytes (TB) per day, Facebook 10 TB, and several enterprises many terabytes an hour.
- Variety -- Structures, semi-structures, unstructured
- Velocity – speed of production (Batch to streams)

73



BIG DATA – Characteristics



Later on:

4th v: Veracity

5th v: Value

Today:

6th v: Variability

7th v: Visualization

Some also add:

8th v: Validity

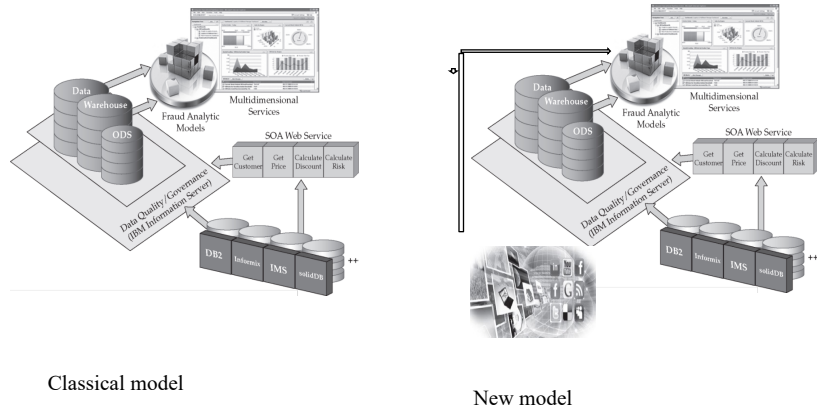
9th v: Vulnerability

10th v: Volatility

74



Usage of 80% of data (non-relational)



75



HADOOP

- Hadoop (<http://hadoop.apache.org/>) is a very important Apache project in Apache Software Foundation written in Java.
- Hadoop is a software environment built on top of a distributed file system that was created specifically for very high intensity processing of data.
- Hadoop was inspired from Google File System (GFS) and the programming environment MapReduce, in which each job is broken into Mapper and Reducer tasks for the processing of data on a cluster of servers for parallelization.

76



Characteristics of NoSQL Systems

- They do not use SQL
- They do not use relations
- They do not support relationships
- They do not support classical transaction properties ACID

- They do not have a formal schema
- They are designed for the support of web applications
- They are designed for the support of very big applications
- They are Open Source



Types of NoSQL Systems (a)

■ DOCUMENT STORES

- **Examples: CouchDB, MongoDB**

```
{  
  "title" : "Foundations of Databases"  
  "rating" : 10  
}
```



Types of NoSQL Systems (b)

■ KEY – VALUE Stores

- Everything is stored in pairs (Key, Value)
- Examples: Memcached, Riak, ...

KEY	VALUE
Name1	bob
Course_44	DBMS
V32_phote	{binary data}
...	



Types of NoSQL Systems (c)

■ GRAPH DATABASES

- Everything is stored as a graph
- Examples: Neo4j, AlegroGraph, DB2 NoSQL, ...





Why NoSQL???

- If you want a FLEXIBLE SCHEMA
- If you have HUGE AMOUNTS OF DATA
- If you are NOT interested in data consistency so much!

81



Cloud Computing with Databases

- Google:** BigTable, for internal storage
AppEngine: <http://code.google.com/appengine/>
Excellent for complex objects
- Generic:** Hadoop (Apache) - Open source for Cloud
- Amazon:**
- | | |
|----------|---|
| S3 | Big files
http://aws.amazon.com/s3/ |
| SimpleDB | Similar to BigTable
http://aws.amazon.com/simplydb/ |
| RDS | Service for relational data
MySQL or Oracle 12g
http://aws.amazon.com/rds/ |
- Microsoft:**
- | | |
|-------|---|
| Azure | SQL Server
http://www.microsoft.com/windowsazure/ |
|-------|---|

82



Advantages of using cloud services for DBMS

- There are no standard costs
 - No cost for infrastructure or software
 - No maintenance
 - Easy management
- Pricing is depends on usage
 - Monthly cost proportional to the DB
 - Monthly cost based on data transfer
- Extensibility
 - Multiple distributed servers
 - Multiple high speed Internet connections
- Reliability
 - Distributed
 - Managed by specialists
 - Security

83



Limitations for cloud DBs

- Cost may become very high
- Then it may be better to buy in-house

84



Cloud Database Pricing

Example: Amazon RDS (MySQL), U.S. East

- 1 Extra large instance
- 20 hours/day
- 20 GB/month at 50 million I/O per month
- 10 GB/month data transfer in
- 500 GB/month data transfer out
- 20 GB/month regional transfer
- => \$616 per month (\$7400/year)

All values are estimates and might not include all fees.

Example: Microsoft SQL Azure Business Edition

- 1 Extra large instance (\$0.96/hour = \$576/month)
- 20 GB/month (\$200/month)
- 10 GB/month data transfer in (\$1/month)
- 500 GB/month data transfer out (\$75/month)
- => \$852 per month (\$10,224/year)

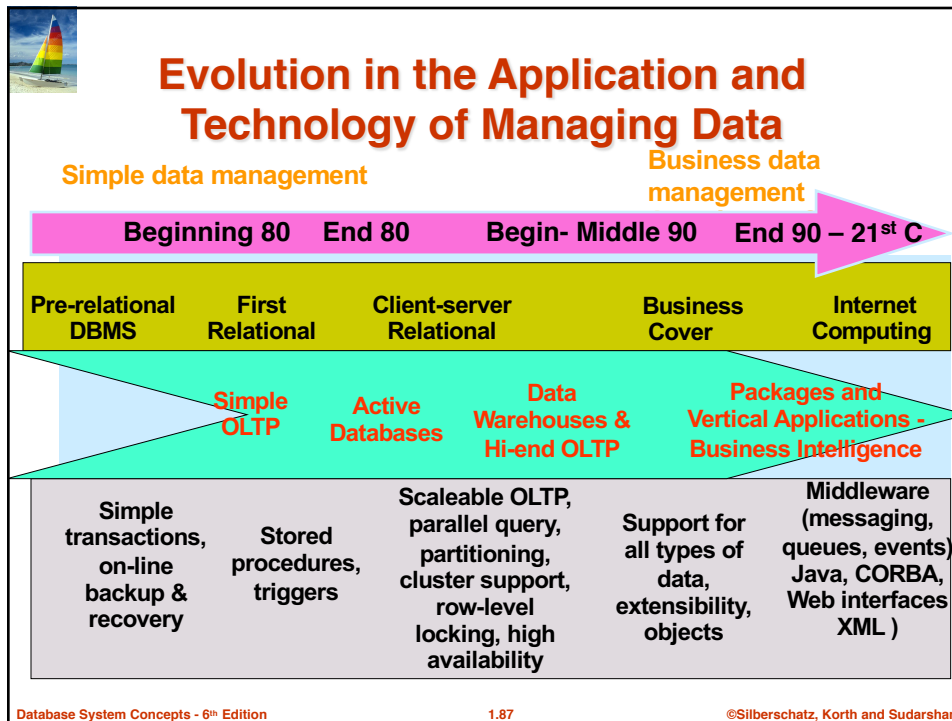
You get a relatively large database with T1-level data transfer for less than 10 percent of the cost of a DBA.



Cloud and Databases



"It was much nicer before people started storing all their data in the Cloud."



87

Final Thoughts

- The DBMS is used for the management of Big data bases.
- Among the advantages is the recovery from system failures, fast application development, concurrency, integrity, security.
- Abstraction levels help in the independence of applications for physical entities.
- DB administrations have very important tasks and are usually well paid.
- DBMS R&D is one of the most attractive computer science domains (perfect combination of theory and practice)

Database System Concepts - 6th Edition 1.88 ©Silberschatz, Korth and Sudarshan

88



End of Chapter 1