

RELATIONAL DATABASE MANAGEMENT SYSTEM (RDBMS) CONCEPTS AND SQL

Venkatesh Vinayakarao

venkateshv@cmi.ac.in

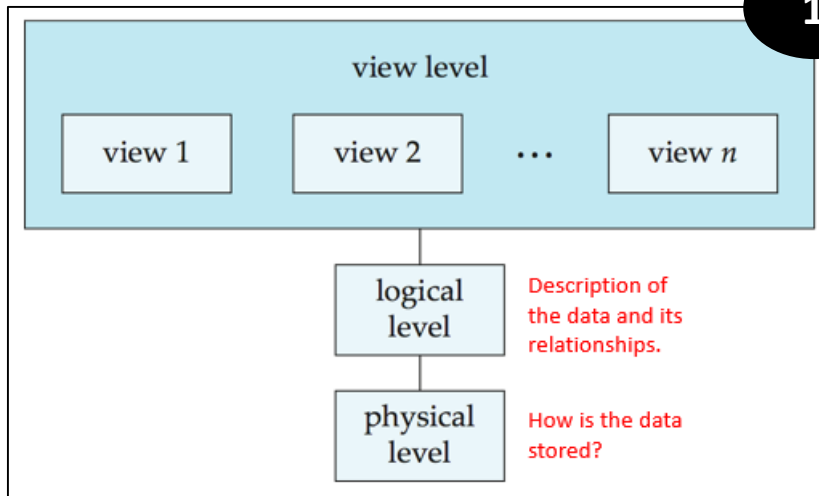
<http://vvtesh.co.in>

Chennai Mathematical Institute

The primary goal of a DBMS is to provide a way to store and retrieve database information that is both *convenient* and *efficient*. - Silberschatz, Korth and Sudarshan.

Review

1



2

Relational Data Model

$$\{\bullet, \blacktriangle, \blacklozenge\}_{\text{set A}} \times \{\triangle, \blacktriangle\}_{\text{set B}} = \{(\bullet, \triangle), (\bullet, \blacktriangle), (\blacktriangle, \triangle), (\blacktriangle, \blacktriangle), (\blacklozenge, \triangle), (\blacklozenge, \blacktriangle)\}$$


set of all ordered pairs, $A \times B$

Relation $R = \{(\bullet, \triangle)\}$

$R(\text{id}, \text{name}) \subseteq \text{id} \times \text{names}$

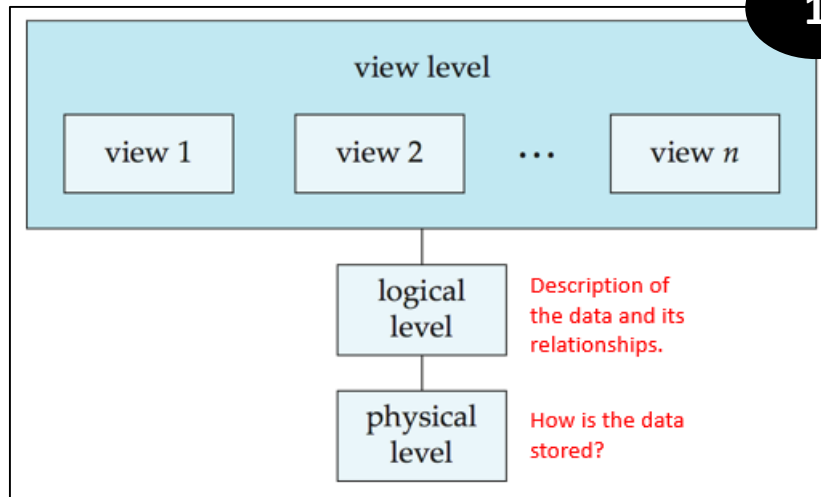
id	name	id	name
1	sd	1	sd
2	vv	1	vv
		2	sd
		2	vv
		3	sd
		3	vv

3



Review

1



2

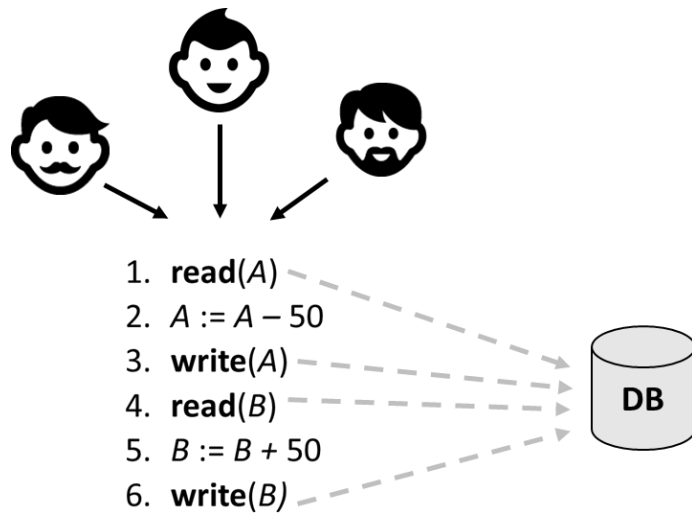
DB Design with ER Model



Representation as a schema

section (course_id, sec_id, semester, year)

Transaction Management

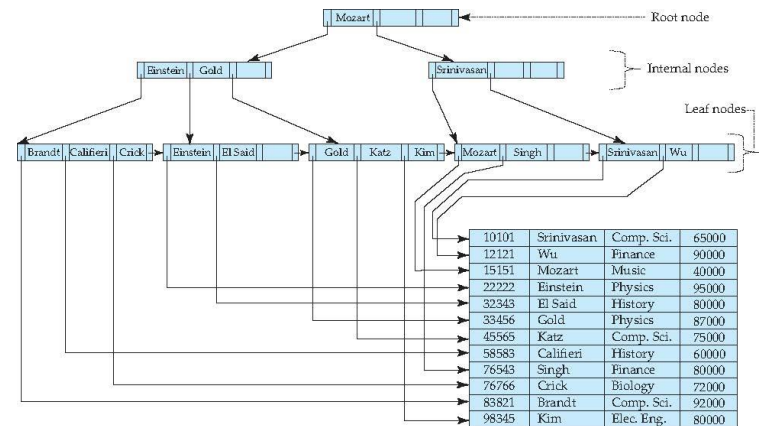
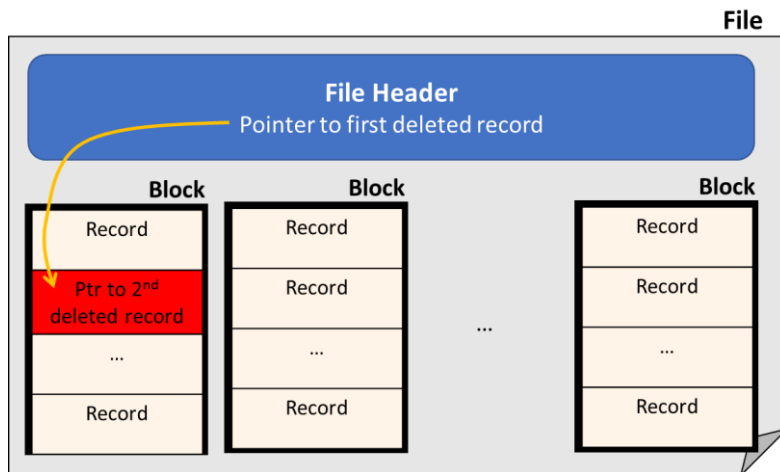
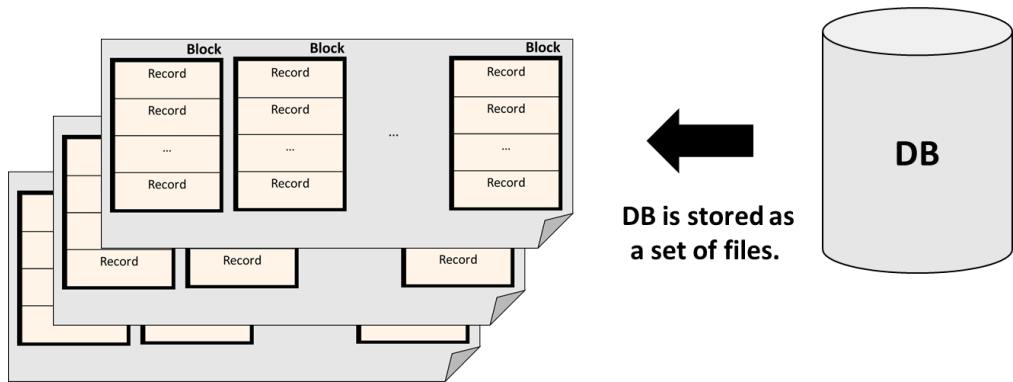


T_1	T_2
read (A) write (A)	
	read (A) write (A)
read (B) write (B)	
	read (B) write (B)

Also written as:


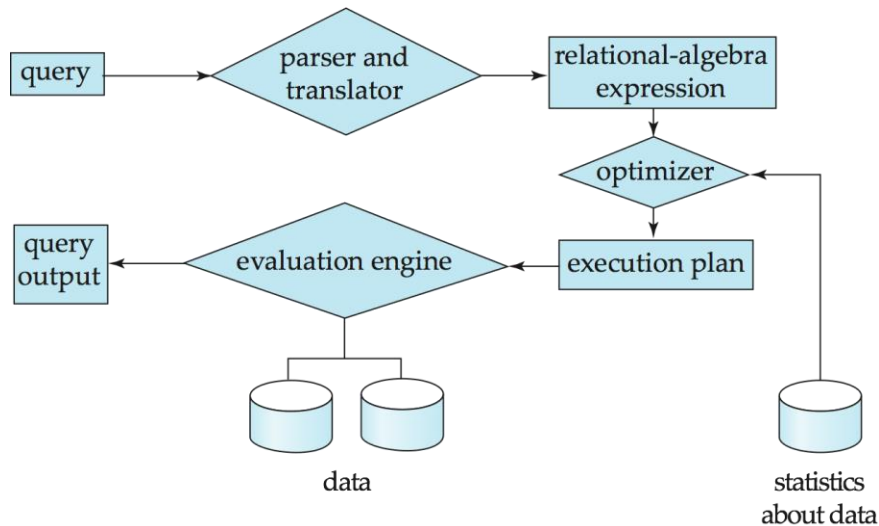
$r_1(A); w_1(A); r_2(A); w_2(A); r_1(B); w_1(B); r_2(B); w_2(B);$

File Organization, Data Storage and Indexing

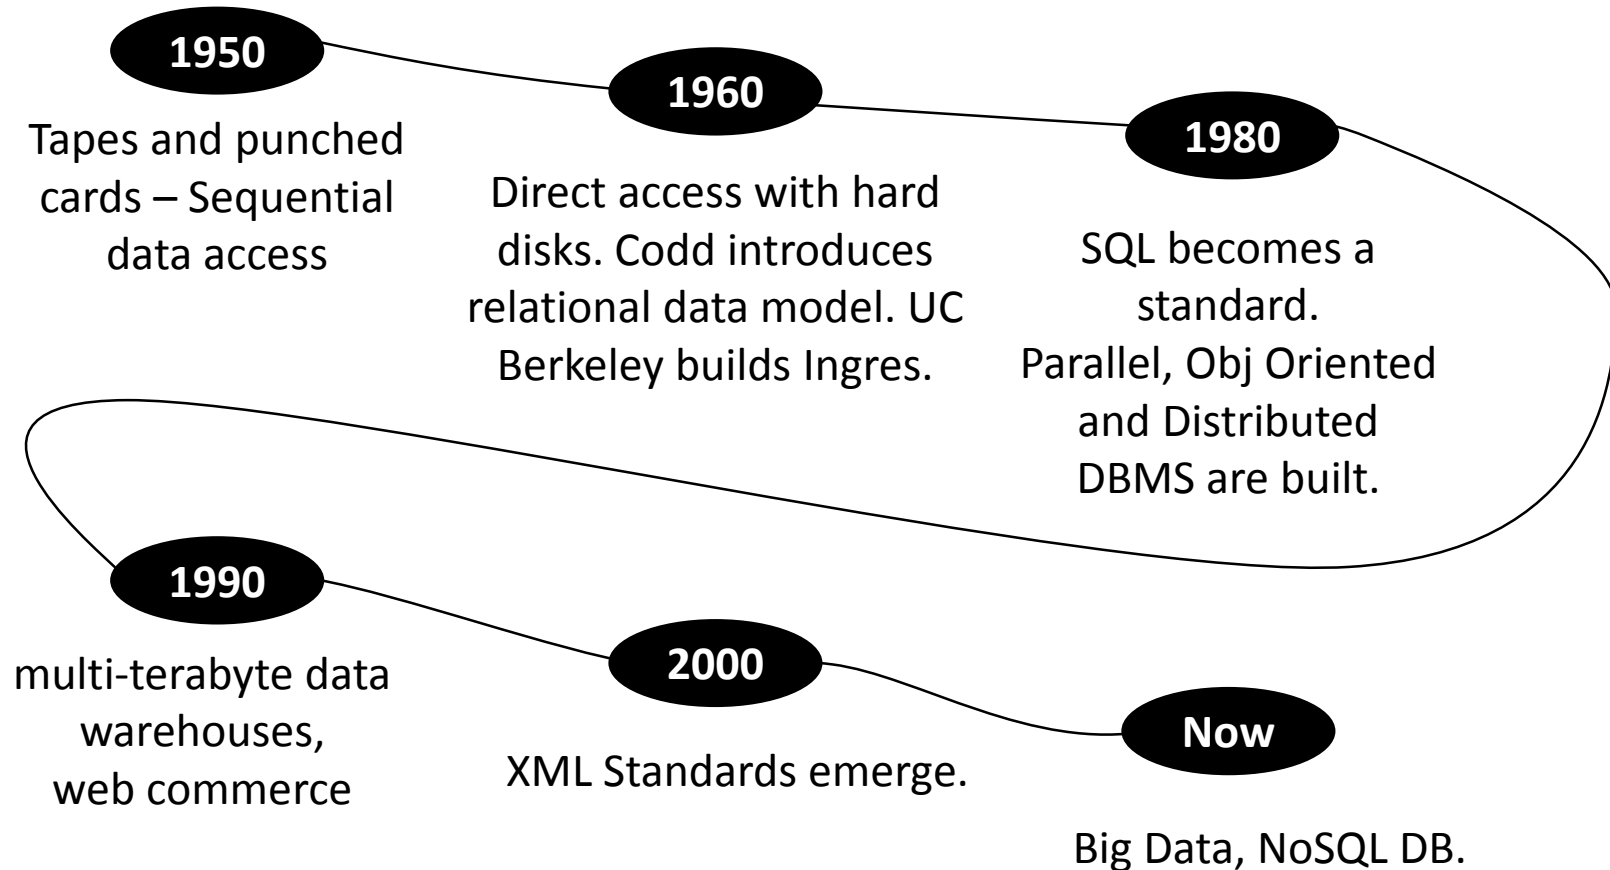


Query Processing

How to effectively execute the query?



History



Tech Quiz (DBMS)

Quiz

- Which of the following is a popular database?
 - MacBook Air
 - Windows
 - Microsoft SQL Server
 - None of the above

A **database-management system**

(DBMS) is a collection of:

1. interrelated data and
2. a set of programs to access those data.

The collection of data is usually referred to as **database**.

Quiz

- Which of these is not a DBMS?
 - MySQL
 - SQL Server
 - Oracle
 - DB2
 - neo4j
 - Ingres
 - Postgres
 - sybase
 - None of the above

Most Popular Software

357 systems in ranking, November 2019

Rank			DBMS	Database Model	Score		
Nov 2019	Oct 2019	Nov 2018			Nov 2019	Oct 2019	Nov 2018
1.	1.	1.	Oracle	Relational, Multi-model	1336.07	-19.81	+34.96
2.	2.	2.	MySQL	Relational, Multi-model	1266.28	-16.78	+106.39
3.	3.	3.	Microsoft SQL Server	Relational, Multi-model	1081.91	-12.81	+30.36
4.	4.	4.	PostgreSQL	Relational, Multi-model	491.07	+7.16	+50.83
5.	5.	5.	MongoDB	Document, Multi-model	413.18	+1.09	+43.70
6.	6.	6.	IBM Db2	Relational, Multi-model	172.60	+1.83	-7.27
7.	7.	8.	Elasticsearch	Search engine, Multi-model	148.40	-1.77	+4.94
8.	8.	7.	Redis	Key-value, Multi-model	145.24	+2.32	+1.06
9.	9.	9.	Microsoft Access	Relational	130.07	-1.10	-8.36
10.	10.	11.	Cassandra	Wide column	123.23	+0.01	+1.48

Quiz

- _____ is a 1943-born **scientist** specializing in database research. He won the **Turing award** for fundamental contributions to the concepts and practices underlying modern database systems. He is also the **founder of many database companies**, including Ingres Corporation, Illustra, Paradigm4, StreamBase Systems, Tamr, Vertica and VoltDB, and served as chief technical officer of Informix. He is also an **editor** for the book **Readings in Database Systems**. His student is now a **professor at IIT Bombay**.

- **Michael Ralph Stonebraker**

Quiz

- Who takes the credit for inventing the relational data model?

- Edgar Codd

Quiz

- Is SQL relationally complete?

- No!

Quiz

- After reading Edgar F. Codd's seminal papers on the relational model, Stonebraker started work with a colleague, Eugene Wong, to develop an efficient and practical implementation. The result was INGRES, a name that reflected the project's original intention to produce a geographically-oriented system with graphical capabilities.
- What does INGRES stand for?

- “Interactive Graphic and Retrieval System”
- Also the name of a celebrated French painter.

Quiz

- In the early days of Postgres there was no SQL. No not NoSQL, there was not SQL. Postgres had its own query language. It wasn't until 1995 that Postgres received SQL support, and with its addition of SQL support it updated its name to _____.

In his blog, Craig Kerstiens.

- PostgreSQL

Quiz

- Name two top conferences in the field of databases.

- VLDB: *International Conference on Very Large Databases*
- PODS: *Conference on Principles of Database Systems*

Research Trends

- concurrency & recovery, distributed/parallel databases, cloud computing
- data and knowledge integration and exchange, data provenance, views and data warehouses, metadata management
- data-centric (business) process management, workflows, web services
- data management and machine learning
- data mining, information extraction, search
- data models, data structures, algorithms for data management
- data privacy and security, human-related data and ethics
- data streams
- design, semantics, query languages
- domain-specific databases (multi-media, scientific, spatial, temporal, text)
- graph databases and (semantic) Web data
- incompleteness, inconsistency, uncertainty in data management
- knowledge-enriched data management
- model theory, logics, algebras, computational complexity

Research Trends

- Access Methods, Concurrency Control, Recovery, Transactions, Indexing and Search, In-memory Data Management, Hardware Accelerators, Query Processing and Optimization, Storage Management.
- Privacy and Security in Data Management.
- Graph Data Management, Social Networks, Recommendation Systems.
- Data Mining and Analytics, Warehousing.
- Crowdsourcing, Embedded and Mobile Databases, Real-time Databases, Sensors and IoT, Stream Databases.
- Data Models and Query Languages, Schema Management and Design, Database Usability, User Interfaces and Visualization.
- Tuning, Benchmarking, Performance Measurement, Database Administration and Manageability.
- Distributed Database Systems, Cloud Data Management, NoSQL, Scalable Analytics, Distributed Transactions, Consistency, P2P and Networked Data Management, Database-as-a-Service, Content Delivery Networks.
- Provenance and Workflows, Spatial, Temporal, and Multimedia Databases, Scientific and Medical Data Management, Profile-based or Context-Aware Data Management.
- Data Cleaning, Information Filtering and Dissemination, Information Integration, Metadata Management, Data Discovery, Web Data Management, Semantic Web, Heterogeneous and Federated Database Systems.
- Fuzzy, Probabilistic and Approximate Databases, Information Retrieval, Text in Databases.