

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



Q&A

Part I

Data Management

Literature:

- Silberschatz et al. *Database System Concepts*. McGraw Hill, Sixth Edition, 2010.
- Wiese. *Advanced Data Management*. De Gruyter Oldenbourg, 2015.

Credits: These slides are partially based on slides of other lectures.

- Nikolaus Augsten, University of Salzburg, Austria (Credits also for valuable discussions and guidance).
- Andrew Pavlo, Carnegie Mellon University (CMU), USA.

Introduction

Motivation

The amount of data is growing rapidly in many different domains.

We do not collect data to simply store it, but we want to

The amount of data is growing rapidly in many different domains.

We do not collect data to simply store it, but we want to

- access it fast at any time and from any place,
- search it (for exact and similar patterns),
- aggregate it (partially),
- join it with other data,
- have multiple users working on the data concurrently,
- make sense out of it.

“But I can organize my data without the overhead of a dedicated system!”

Theoretically yes, but this implies:

“But I can organize my data without the overhead of a dedicated system!”

Theoretically yes, but this implies:

- Organizing your data in multiple independent (plain) files.
- Other users may have trouble to understand your organization.
- Taking care of all the requirements yourself.
- You drop flexibility and scalability.

“But I can organize my data without the overhead of a dedicated system!”

Theoretically yes, but this implies:

- Organizing your data in multiple independent (plain) files.
- Other users may have trouble to understand your organization.
- Taking care of all the requirements yourself.
- You drop flexibility and scalability.

What can possibly go wrong?

Example Application

Manage your movies collection in CSV files using Python3.

Movies: (name;year;runtime;genre;)

```
The Godfather;1972;177;Mafia;  
The Avengers;2012;143;Sci-Fi;  
Batman - The Dark Knight;2008;152;Action;  
American Psycho;2000;101;Thriller;  
Home Alone;1990;103;Comedy;
```

Actors: (name;birthyear;knownfor;)

```
Marlon Brando;1924;The Godfather;  
Al Pacino;1940;The Godfather;  
Macaulay Culkin;1980;Home Alone;  
Joe Pesci;1943;Home Alone;  
Robert Downey Jr.;1965;The Avengers;  
Mark Ruffalo;1967;The Avengers;  
Christian Bale;1974;Batman - The Dark Knight;  
Christian Bale;1974;American Psycho;  
Heath Ledger;1979;Batman - The Dark Knight;  
Willem Dafoe;1955;American Psycho;
```

Example Application

Query: Find all movies of “Christian Bale”.

Example Application

Query: Find all movies of “Christian Bale”.

```
if __name__ == "__main__":  
    with open("actor.csv", "r") as fin:  
        for line in fin.readlines():  
            parts = line.split(";")  
  
            name = parts[0]  
            birthyear = int(parts[1])  
            knownfor = parts[2]  
  
            if name == "Christian_Bale":  
                print(knownfor)
```

Redundancy and Inconsistency:

- Several copies of a datum may exist (possibly stored differently).
- **Redundancy:** Higher storage requirements. How about accessing the data?

Redundancy and Inconsistency:

- Several copies of a datum may exist (possibly stored differently).
- **Redundancy:** Higher storage requirements. How about accessing the data?
⇒ Multiple file accesses (which is slow).
- **Inconsistency:** What happens if you have to update the data?

Redundancy and Inconsistency:

- Several copies of a datum may exist (possibly stored differently).
- **Redundancy:** Higher storage requirements. How about accessing the data?
⇒ Multiple file accesses (which is slow).
- **Inconsistency:** What happens if you have to update the data?
⇒ We must not forget a single copy.
- **Goal:**

Redundancy and Inconsistency:

- Several copies of a datum may exist (possibly stored differently).
- **Redundancy:** Higher storage requirements. How about accessing the data?
⇒ Multiple file accesses (which is slow).
- **Inconsistency:** What happens if you have to update the data?
⇒ We must not forget a single copy.
- **Goal:** Minimize redundancy and prevent inconsistency.

Redundancy and Inconsistency:

- Several copies of a datum may exist (possibly stored differently).
- **Redundancy:** Higher storage requirements. How about accessing the data?
⇒ Multiple file accesses (which is slow).
- **Inconsistency:** What happens if you have to update the data?
⇒ We must not forget a single copy.
- **Goal:** Minimize redundancy and prevent inconsistency.

Example: Update “The Avengers” to “Marvel’s The Avengers”.

Example Application

Query: Update “The Avengers” to “Marvel’s The Avengers”.

Movies: (name;year;runtime;genre;)

The Godfather;1972;177;Mafia;
Marvel's The Avengers;2012;143;Sci-Fi;
Batman - The Dark Knight;2008;152;Action;
American Psycho;2000;101;Thriller;
Home Alone;1990;103;Comedy;

Actors: (name;birthyear;knownfor;)

Marlon Brando;1924;The Godfather;
Al Pacino;1940;The Godfather;
Macaulay Culkin;1980;Home Alone;
Joe Pesci;1943;Home Alone;
Robert Downey Jr.;1965;The Avengers;
Mark Ruffalo;1967;The Avengers;
Christian Bale;1974;Batman - The Dark Knight;
Christian Bale;1974;American Psycho;
Heath Ledger;1979;Batman - The Dark Knight;
Willem Dafoe;1955;American Psycho;

Data Access and Analysis:

- We want to analyze our data. How to link related data?

Data Access and Analysis:

- We want to analyze our data. How to link related data?
⇒ Each analysis requires a tailored program.
- **Goal:**

Data Access and Analysis:

- We want to analyze our data. How to link related data?
⇒ Each analysis requires a tailored program.
- **Goal:** Generic analysis and linkage of related data.

Data Access and Analysis:

- We want to analyze our data. How to link related data?
⇒ Each analysis requires a tailored program.
- **Goal:** Generic analysis and linkage of related data.

Example: List all movies of “Christian Bale” and with its runtime.

Example Application

Query: List all movies of “Christian Bale” and with its runtime.

Movies: (name;year;runtime;genre;)

The Godfather;1972;177;Mafia;
Marvel's The Avengers;2012;143;Sci-Fi;
Batman - The Dark Knight;2008;152;Action;
American Psycho;2000;101;Thriller;
Home Alone;1990;103;Comedy;

Actors: (name;birthyear;knownfor;)

Marlon Brando;1924;The Godfather;
Al Pacino;1940;The Godfather;
Macaulay Culkin;1980;Home Alone;
Joe Pesci;1943;Home Alone;
Robert Downey Jr.;1965;The Avengers;
Mark Ruffalo;1967;The Avengers;
Christian Bale;1974;Batman - The Dark Knight;
Christian Bale;1974;American Psycho;
Heath Ledger;1979;Batman - The Dark Knight;
Willem Dafoe;1955;American Psycho;

Data Integrity Issues:

- Updates may violate the integrity of your data.
- How do you ensure data integrity?

Data Integrity Issues:

- Updates may violate the integrity of your data.
- How do you ensure data integrity?
 - ⇒ Each single application must respect all consistency constraints.
- **Goal:**

Data Integrity Issues:

- Updates may violate the integrity of your data.
- How do you ensure data integrity?
 - ⇒ Each single application must respect all consistency constraints.
- **Goal:** Global definition and monitoring of consistency constraints.

Data Integrity Issues:

- Updates may violate the integrity of your data.
- How do you ensure data integrity?
 - ⇒ Each single application must respect all consistency constraints.
- **Goal:** Global definition and monitoring of consistency constraints.

Example: Insert “Christian Bale” as actor in “The Avengers”.

Concurrency Issues:

- Multiple users should be able to access and update the data simultaneously. How do you ensure consistency over all applications that access the data?

Concurrency Issues:

- Multiple users should be able to access and update the data simultaneously. How do you ensure consistency over all applications that access the data?
- **Anomalies:** Inconsistencies, e.g., lost updates.

Concurrency Issues:

- Multiple users should be able to access and update the data simultaneously. How do you ensure consistency over all applications that access the data?
- **Anomalies:** Inconsistencies, e.g., lost updates.
- **Efficiency:** If one user locks a file, then the other user must wait.
- **Goal:**

Concurrency Issues:

- Multiple users should be able to access and update the data simultaneously. How do you ensure consistency over all applications that access the data?
- **Anomalies:** Inconsistencies, e.g., lost updates.
- **Efficiency:** If one user locks a file, then the other user must wait.
- **Goal:** Out-of-the-box multi-user operation without anomalies.

Concurrency Issues:

- Multiple users should be able to access and update the data simultaneously. How do you ensure consistency over all applications that access the data?
- **Anomalies:** Inconsistencies, e.g., lost updates.
- **Efficiency:** If one user locks a file, then the other user must wait.
- **Goal:** Out-of-the-box multi-user operation without anomalies.

Example: Users A inserts “Chris Evans” as actor while user B inserts “Scarlett Johansson” as actress in “The Avengers” simultaneously.

Example Application

Scenario: Users A inserts “Chris Evans” as actor while user B inserts “Scarlett Johansson” as actress in “The Avengers” simultaneously.

Actors: (name;birthyear;knownfor;)

```
Marlon Brando;1924;The Godfather;  
Al Pacino;1940;The Godfather;  
Macaulay Culkin;1980;Home Alone;  
Joe Pesci;1943;Home Alone;  
Robert Downey Jr.;1965;The Avengers;  
Mark Ruffalo;1967;The Avengers;  
Christian Bale;1974;Batman - The Dark Knight;  
Christian Bale;1974;American Psycho;  
Heath Ledger;1979;Batman - The Dark Knight;  
Willem Dafoe;1955;American Psycho;  
Scarlett Johansson;1984;The Avengers;
```

Atomicity and Recovery:

- Data must neither be lost nor inconsistent when the system crashes.

Atomicity and Recovery:

- Data must neither be lost nor inconsistent when the system crashes.
- **Atomicity:** Data may be inconsistent if an operation is only applied partially
⇒ Execute an operation in an all-or-nothing manner.

Atomicity and Recovery:

- Data must neither be lost nor inconsistent when the system crashes.
- **Atomicity:** Data may be inconsistent if an operation is only applied partially
⇒ Execute an operation in an all-or-nothing manner.
- **Recovery:** Backup of data may not reflect the latest state.
- **Goal:**

Atomicity and Recovery:

- Data must neither be lost nor inconsistent when the system crashes.
- **Atomicity:** Data may be inconsistent if an operation is only applied partially
⇒ Execute an operation in an all-or-nothing manner.
- **Recovery:** Backup of data may not reflect the latest state.
- **Goal:** Prevent data loss and inconsistencies by design.

Atomicity and Recovery:

- Data must neither be lost nor inconsistent when the system crashes.
- **Atomicity:** Data may be inconsistent if an operation is only applied partially
⇒ Execute an operation in an all-or-nothing manner.
- **Recovery:** Backup of data may not reflect the latest state.
- **Goal:** Prevent data loss and inconsistencies by design.

Example: Update all “The Avengers” to “Marvel’s The Avengers”. Your system crashes in between “Mark Ruffalo” and “Scarlett Johansson”.

Example Application

Scenario: Update all “The Avengers” to “Marvel’s The Avengers”. Your system crashes in between “Mark Ruffalo” and “Scarlett Johansson”.

Actors: (name;birthyear;knownfor;)

```
Marlon Brando;1924;The Godfather;  
Al Pacino;1940;The Godfather;  
Macaulay Culkin;1980;Home Alone;  
Joe Pesci;1943;Home Alone;  
Robert Downey Jr.;1965;Marvel's The Avengers;  
Mark Ruffalo;1967;Marvel's The Avengers;  
Christian Bale;1974;Batman - The Dark Knight;  
Christian Bale;1974;American Psycho;  
Heath Ledger;1979;Batman - The Dark Knight;  
Willem Dafoe;1955;American Psycho;  
Scarlett Johansson;1984;The Avengers;
```

Other Issues:

- **Efficiency:** Efficient algorithms are required to analyze large amounts of data.
- **General-purpose:** The problems of application developers will partially overlap.
- **Security issues:** Flexible and fine-grained access rights for multiple users.

A **database management system** (DBMS) is

- (i) a collection of interrelated data, the **database**, and
- (ii) a set of programs to access the data.

In other words, you do not have to care about how to store the data (physically), how to analyze it (efficiently), how to (partially) update data, or how to deal with multiple users. A DBMS organizes all this for you.

DBMSs are at the core of many applications.

When Not to Use a DBMS

- The data are too complex to model it.
- Specific requirements like real-time queries or special operations.
- The overhead of a DBMS is too high or unnecessary.
- No or low return on investment (ROI).

One Size Fits All?

General-Purpose DBMS

A DBMS that tries to fit as many application scenarios as possible with a single system. This implies higher complexity but also a large user base.

Examples:

- PostgreSQL (open source)
- MySQL (open source)
- MonetDB (open source)
- SQLite (open source)
- IBM DB2 (closed source)
- Oracle Database (closed source)
- Microsoft SQL Server (closed source)
- ...

“But this sounds good, no?”

“But this sounds good, no?”

Problems:

- Unnecessary overhead (e.g., recovery or strong consistency)
- Limited performance
- Application-specific operations are not supported natively
- Limited flexibility

“One Size Fits All”: An Idea Whose Time Has Come and Gone

Michael Stonebraker¹ and Ugur Cetintemel (2005)

“A one size fits all database doesn’t fit anyone”

Werner Vogels² (2018)

¹Database Systems Researcher at the MIT. Won the Turing Award in 2014.

²Computer Scientist and CTO at Amazon.

A DBMS that is tailored to fit a specific purpose best, i.e., provide all the functionality that is required while also providing the best performance and flexibility (with respect to the specific application domain).

Synonyms: Specialized DBMS, purpose-built DBMS.

Temporal Data: A temporal DBMS is optimized to manage and analyze data that references time (i.e., they are timestamped). For example, a time series $x = \langle x_{t_1}, x_{t_2}, \dots, x_{t_n} \rangle$ is often a sequence of n data points that are spaced at strictly increasing times ($t_i < t_{i+1}$ with $i = 1, \dots, n - 1$).

Requirements:

- Exact/Approximate matching of (parts of) time series.
- Efficient compression mechanisms.
- Serve specific aspects like valid time or transaction time.
- ...

Real-Time Data: A real-time DBMS manages data that is changed continuously. A DBMS that operates in real time answers the queries within a guaranteed time frame (the response time, i.e., it has a deadline).

Requirements:

- Answer every query in a given time frame.
- Query scheduling (or queuing).
- Consistency may not be that important.
- ...

Process Mining Data: Process mining engines manage business event logs. An example event log is the sequence of activities if you place an order in some online shop. These systems are required to analyze large amounts of data in real time.

Requirements:

- Optimized, domain-specific language.
- Real-time performance for best user experience.
- ...

Multiple specific aspects may need to be combined to serve a novel application scenario. This may also result in a new special-purpose DBMS.

A modern application is not monolithic, i.e., different DBMSs may be used to implement different parts of an application.

Towards a One Size Fits All Database Architecture

Jens Dittrich³ and Alekh Jindal. 2011.

One Size Fits all, Again! The Architecture of the Hybrid OLTP&OLAP Database Management System HyPer

Alfons Kemper⁴ and Thomas Neumann⁵. 2011.

³Database Systems Researcher at Saarland University.

⁴Database Systems Researcher at the TU Munich co-author of the book *Datenbanksysteme*.

⁵Database Systems Researcher at the TU Munich.

Q&A

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



- **Motivation** for data management.
- **Example Application** using plain files.
- **Problems** in data management:
 - Redundancy and inconsistency.
 - Data access and analysis.
 - Data integrity.
 - Concurrency.
 - Atomicity and recovery.
 - Efficiency, general-purpose, security.
- **General- and special-purpose DBMS** and examples.
- **One Size Fits All**

Q&A

Database Fundamentals

Data are facts that are to be stored.

Information is data combined with semantics (meaning).

Knowledge is information combined with an application.

Example Application

What are data, information, and knowledge in our example?

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
Batman - The Dark Knight	2008	152	Action
American Psycho	2000	101	Thriller
Home Alone	1990	103	Comedy

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
Macaulay Culkin	1980	Home Alone
Joe Pesci	1943	Home Alone
Robert Downey Jr.	1965	The Avengers
Mark Ruffalo	1967	The Avengers
Christian Bale	1974	Batman - The Dark Knight
Christian Bale	1974	American Psycho
Heath Ledger	1979	Batman - The Dark Knight
Willem Dafoe	1955	American Psycho

A **database** (DB) is a collection of interrelated data.

Metadata provides us with information about the structure of a database. All the metadata are stored in a **catalog**.

A **database system** (DBS) is also referred to as the combination of a database, the corresponding metadata, and a DBMS (which in this case only provides the set of programs). The terms DBS and DBMS are often used interchangeably.

Example Application

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Catalog:

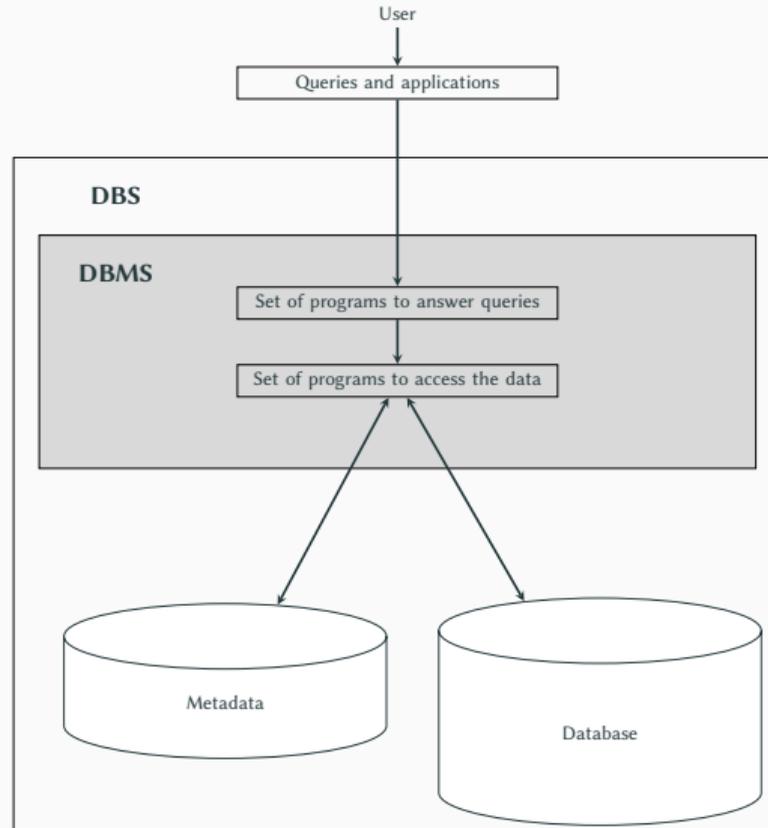
Tables Metadata:

relation	columnCount
Actors	10
Movies	5

Columns Metadata:

columnName	dataType	relation
name	TEXT	Actors
year	INTEGER	Actors
name	TEXT	Movies
...

Database System



A **table** consists of multiple **tuples**, each of which is a sequence of **attributes**. Informally, a tuple (attribute) can be imagined as a row (column) of a table.

A **key** is subset of attributes. A **primary key** is a key of minimum length that uniquely identifies a tuple. A **foreign key** is a reference to a primary key.

Data-Definition Language (DDL): Specify the structure of your data and the consistency constraints that are enforced.

- **Schema:** Describes the structure of your data and how the data are interrelated, e.g., a movie has 4 columns: name, year, runtime, and genre.
- **Consistency Constraints:** Describe integrity constraints that must be satisfied at any given point in time, e.g., the runtime in minutes is an integer > 0 .

Example Application

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Pseudo-DDL (simplified):

```
CREATE TABLE Movies (  
  name TEXT KEY,  
  year INTEGER,  
  runtime INTEGER (> 0),  
  genre TEXT  
)
```

```
CREATE TABLE Actors (  
  name TEXT KEY,  
  birthyear INTEGER,  
  knownfor TEXT REFERENCES(Movies.name)  
)
```

Data-Manipulation Language (DML): Query and manipulate your data.

- **Query Language:** Allows you to query your data without modifying it, e.g., get all movies of “Christian Bale” or get all actors of “The Avengers”.
- **Manipulation Language:** Allows you to modify your data, e.g., insert a new movie, delete an existing movie, update the movies of a particular actor.

A **query** is a statement that requests some information. Informally, your query “asks” and the database system answers by returning the corresponding information.

Caveat: The term query language often refers to both DML parts.

Example Application

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Pseudo-DML (simplified):

```
SELECT knownfor FROM Actors  
WHERE name = 'Christian Bale'
```

```
SELECT name FROM Actors  
WHERE knownfor = 'The Avengers'
```

```
INSERT INTO Movies  
VALUES ('Hulk', 2003, 133, 'Sci-Fi')
```

```
UPDATE Actors  
SET knownfor = 'American Psycho'  
WHERE name = 'Christian Bale'
```

Procedural Languages: Describe what data you want and also how to retrieve those.

Declarative Languages: Describe only what data you are interested in.

Pure Languages: Form the (theoretical) foundation underneath the languages that are used in practice. Examples include relational algebra or tuple calculus.

Our Focus: Declarative languages.

Send someone to the supermarket to get milk.

The SQL Query Language

The **Structured Query Language (SQL)** was developed by IBM and is the de-facto standard language in database systems.

SQL is a declarative query language and includes DDL and DML elements.

The SQL standard (last revision: 2016) comprehensively summarizes all elements.

Inventing new approaches is like *“trying to swim up the Niagara Falls”*.

Michael Stonebraker in Information Age. 2010.

Data Definition:

```
CREATE TABLE Movies (  
  name VARCHAR(50) PRIMARY KEY,  
  year INTEGER NOT NULL,  
  runtime INTEGER CHECK(runtime > 0),  
  genre VARCHAR(50)  
);
```

```
CREATE TABLE Actors (  
  name VARCHAR(50) PRIMARY KEY,  
  birthyear INTEGER,  
  knownfor VARCHAR(50) REFERENCES Movies(name)  
);
```

- **VARCHAR(n), INTEGER:** Domain of a single column (data types).
- **NOT NULL, CHECK:** Constraints on a single column.
- **PRIMARY KEY, REFERENCES:** Constraints on an entire table.

Queries:

```
SELECT knownfor FROM Actors  
WHERE name = 'Christian Bale';
```

```
SELECT name FROM Actors  
WHERE knownfor = 'The Avengers';
```

- **SELECT:** Specifies the column to retrieve.
- **FROM:** Specifies the tables to consider.
- **WHERE:** Specifies the condition(s) the result must satisfy.

Queries:

```
INSERT INTO Movies(name, year, runtime, genre)
VALUES ('Hulk', 2003, 133, 'Sci-Fi');
```

```
UPDATE Actors SET knownfor = 'American Psycho'
WHERE name = 'Christian Bale';
```

- **INSERT INTO ...VALUES:** Adds new tuple to a table based on the given values.
- **UPDATE ...SET:** (Partially) changes the values of a tuple.

Abstraction: Hide the complexity of the system while providing all the functionality (from people without deep computer science background). Everyone should be able to use a database system.

3 Levels of Data Abstraction (bottom-up):⁶

1. **Physical:** How to store the data (e.g., on hard disk).
2. **Logical:** What data are stored and how are they related.
3. **View:** Specific views on the data (e.g., on a specific part of the data).

⁶Cf. also ANSI/SPARC architecture.

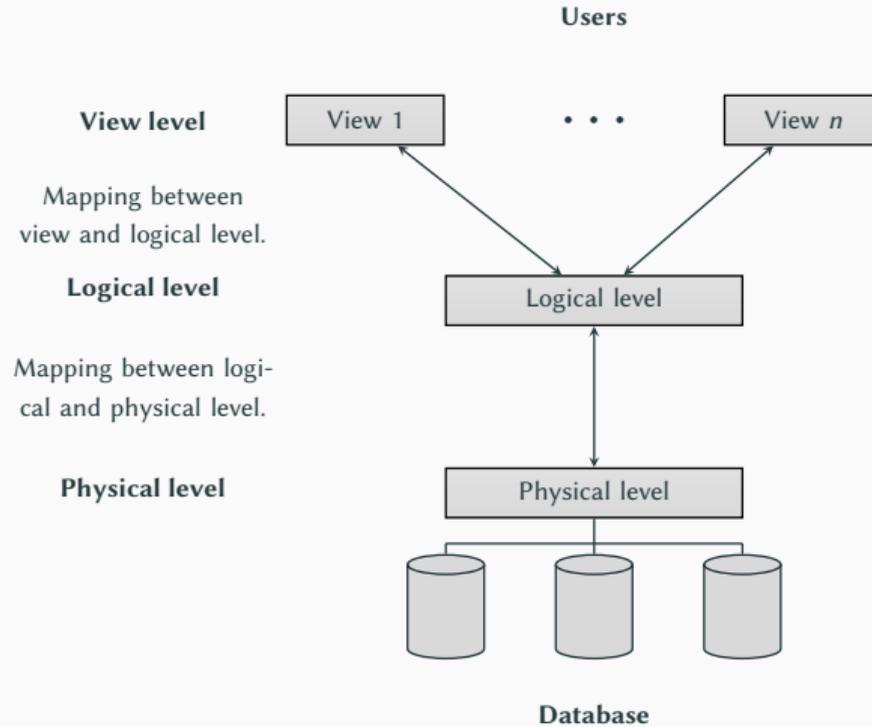
3 Levels of Data Abstraction

The **physical level** defines the data structures that are used to store and access the data physically. Examples are tables or auxiliary structures like indexes.

The **logical level** defines the schemata and constraints of the entire database. Physical structures may be used underneath, but the user does not have to know them.
⇒ **physical data independence.**

The **view level** reduces the complexity by providing only information that is necessary for the respective user. Irrelevant data are not shown.

3 Levels of Data Abstraction



Mappings are used to link the abstraction levels.

Instance vs. Schema

A **schema** describes the overall structure of the data (often using tables) and is typically stable (i.e., rarely modified).

An **instance** is the information that is stored at a particular point in time. The instance may be frequently subject to changes.

Each level has its own schema with the logical schema being the most important one.

A **valid instance** satisfies all structural requirements and consistency constraints.

Instance vs. Schema

Schemata:

Movies:

name	year	runtime	genre
------	------	---------	-------

Actors:

name	birthyear	knownfor
------	-----------	----------

Instances:

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Logical Data Independence: The ability to update the logical schema transparently, i.e., no change on the view level is required.

Physical Data Independence: The ability to update the physical schema transparently, i.e., no change on the logical level is required.

Benefits:

- Only the mapping between the levels need to be adapted.
- No change in the application required (it operates on the views).

Q&A

Assignment 1

Assignment 1

Summary:

- **Where:** Release and submission via Blackboard.
- **When:** March 22 – April 19, 2021 (resp. late: April 26, 2021).
- **What to do:** (a) Set up a relational database locally (PostgreSQL; data will be provided), (b) get familiar with SQL and learn how to execute queries (most queries will be given), (c) write a small Python3 application that executes the queries, and (d) answer questions regarding the assignment.
- **What to submit:** The Python3 code and the answers to the questions.
- **Grading:** 55% Python3 code, 45% questions (incl. the meeting).

Q&A

Happy Eastern!

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



Recap

- **Motivation** for data management.
- **Example Application** using plain files.
- **Problems** in data management:
 - Redundancy and inconsistency.
 - Data access and analysis.
 - Data integrity.
 - Concurrency.
 - Atomicity and recovery.
 - Efficiency, general-purpose, security.
- **General- and special-purpose DBMS** and examples.
- **One Size Fits All**

Recap

- Basics of **database terminology**.
- **Data-Definition** vs. **Data-Manipulation languages**.
- **Procedural** vs. **declarative languages**.
- Basics of **SQL**.
- 3 levels of **data abstraction** (physical, logical, view)
- **Instances** vs. **schemata**.
- **Data independence** (logical and physical).
- **Assignment 1**.

Q&A

(Declarative) Query Processing

Example SQL Query (+ Result):

```
SELECT knownfor FROM Actors
WHERE name = 'Christian Bale';
```



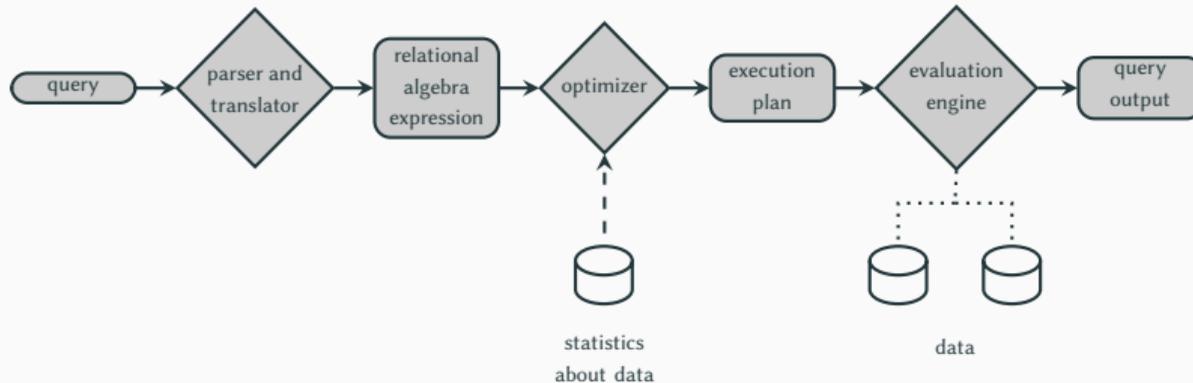
knownfor
Batman - The Dark Knight
American Psycho

Query processing describes the process of extracting data from a database. In other words: What happens in a database when we issue a (declarative) query?

(Declarative) Query Processing

On a high level, **three major components** are used to process a query:

1. **Parser:** Translates the query into an internal representation.
2. **Optimizer:** Choses the most efficient way to evaluate the query.
3. **Evaluation Engine:** Executes the evaluation plan and returns the result.



(Declarative) Query Processing

An evaluation plan typically consists of multiple operation. Optimization is done based on the **estimated costs** of all involved operations.

For a given query, multiple valid **evaluation plans** may exist and must be **compared efficiently** (with respect to their estimated costs).

The estimated costs consider **many different factors** including access to hard disk, time to execute the query on the CPU, or network communication costs.

Data Models

Types of Data Models

3 types of data models that are somewhat related to the 3 levels of data abstraction.

Conceptual data models: High level, i.e., only the schema is reflected but no instances. Related to the view level. Examples include Entity-Relationship (ER) and Unified Modeling Language (UML) models.

Logical data models: Depicts the instances and can be used to implement a database. Related to the logical level. Examples include the relational and the object-based models.

Physical data models: Low level, i.e., as close to the physical storage as possible. Related to the physical level and is typically system-specific.

Types of Data Models – Examples

Logical Model (Relational):

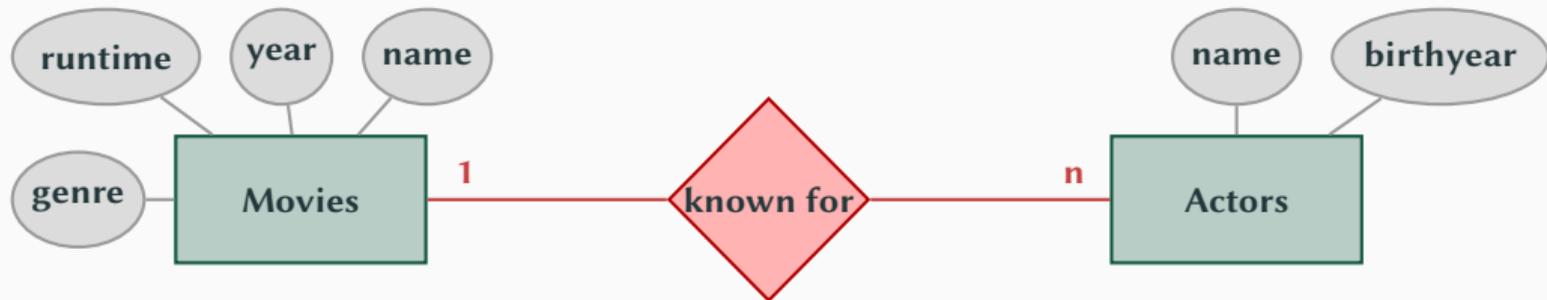
Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Conceptual Model (ER):



Relational Data Model

Intuitive and **widely used** model. An example of **record-based** models.

A **collection of relations (tables)** stores records of data as rows. A **tuple (row)** is an **entity** of the real world, an **attribute (column)** is a **property** of an entity. The **structure** of a record is **fixed**.

A **relation** has a **name** and a **set of (unique) attributes**. An **attribute** has a **name** and a **predefined domain**, i.e., values originate from a specific domain.

Tables are **filled row-wise** and a **row** represents the **state of an entity**.

Relational Data Model

The set of columns is called **relation schema**, and the set of relation schemata over all tables is called the **database schema**.

Intrarelational constraints: Dependencies inside a single table, Σ_i

Interrelational constraints: Dependencies between different tables, Σ

Relation schema: $R_i = (\{A_{i1}, A_{i2}, \dots, A_{im}\}, \Sigma_i)$ ⁷.

Database schema: $D = (\{R_1, R_2, \dots, R_n\}, \Sigma)$.

⁷ A_{ij} ... Name of the j -th attribute (column) of the i -th relation (table).

Relational Data Model – Example

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Relation schemata:

Movies = ($\{\text{name, year, runtime, genre}\}, \{\text{name, year} \rightarrow \text{runtime, genre}\}$)

Actors = ($\{\text{name, birthyear, knownfor}\}, \{\text{name, birthyear} \rightarrow \text{knownfor}\}$)

Database schema:

MovieStore = ($\{\text{Movies, Actors}\}, \{\text{Actors.knownfor} \subseteq \text{Movies.name}\}$)

Normalization: Reduce anomalies (which lead to inconsistencies) by distributing attributes among tables and linking them using **foreign key constraints**.

Referential Integrity: Values of foreign keys exist as values in the referenced table, i.e., the referenced table contains at least one tuple that holds the value of the foreign key.

Relational Data Model – Example

Movies:

id	name	year	runtime	genre
1	The Godfather	1972	177	Mafia
2	The Avengers	2012	143	Sci-Fi
3	Batman - The Dark Knight	2008	152	Action
4	American Psycho	2000	101	Thriller
5	Home Alone	1990	103	Comedy

Actors:

id	name	birthyear	knownfor
101	Marlon Brando	1924	The Godfather
102	Al Pacino	1940	The Godfather
103	Macaulay Culkin	1980	Home Alone
104	Joe Pesci	1943	Home Alone
105	Robert Downey Jr.	1965	The Avengers
106	Mark Ruffalo	1967	The Avengers
107	Christian Bale	1974	Batman - The Dark Knight
108	Christian Bale	1974	American Psycho
109	Heath Ledger	1979	Batman - The Dark Knight
110	Willem Dafoe	1955	American Psycho

Relational Data Model – Example

Movies:

id	name	year	runtime	genre
1	The Godfather	1972	177	Mafia
2	The Avengers	2012	143	Sci-Fi
3	Batman - The Dark Knight	2008	152	Action
4	American Psycho	2000	101	Thriller
5	Home Alone	1990	103	Comedy

Actors:

id	name	birthyear
...
106	Mark Ruffalo	1967
107	Christian Bale	1974
108	Heath Ledger	1979
109	Willem Dafoe	1955

knownFor:

movie-id	actor-id
...	...
2	106
3	108
3	107
4	107
4	109

Drawbacks:

- Relations may not be optimal to represent the data.
- Everything is a relation (semantic overloading).
- Homogeneous structure of data.
- Limited flexibility and data types.

Drawbacks and new challenges gave rise to **non-relational data models**.

Object-Based Data Model

Many programming languages are object-oriented, i.e., based on the concept of objects.

Objects contain data and provide functionality, and interact with each other (e.g., like human beings in the real world).

Example Objects:



Three Options:

1. **Object-relational data model:** Extends relational data model with object-oriented features.
2. **Object-relational mapping (ORM):** Maps objects to tuples in the relational model and handles the translation.
3. **Object-oriented data model:** Implements an object-based data model natively.

Key-Value Data Model

Prototype of a **schemaless** model where each tuple is a pair of two strings (k, v) . k is a **unique key** that is associated with an arbitrary **value** v .

Data are accessed using the key and values are not interpreted by any means.

Example operations:

- `put(k, v)`: Inserts a new pair key-value pairs (k, v)
- `get(k)`: Retrieves the value associated with k .
- `delete(k)`: Deletes the tuple(s) associated with k .

Caveat: In programming often referred to as hash table or dictionary (Python).

Key-Value Data Model

Only keys can be searched, **values cannot be searched**. Combinations of values must be done by the application that accesses the data.

Easy to distribute \Rightarrow Good for **data-intensive applications**.

Example Key-Value Store:

UserID		Shopping Cart
1002	→	Shoe, Jordans, red, 37 # Computer game “Diablo”, Blizzard # Headset, Razer, Kraken Kitty
1003	→	Wilson American Football, NFL, Replica # Hail Mary, Gloves Receiver, 2.0, Black & White
1004	→	Book “Database System Concepts”, 7th Edition, Silberschatz
...		...

Document-Based Data Model

A document stores data in a semi-structured and nested text format (e.g., XML or JSON).

Each document has a unique identifier, but the value is a document structured in a specific format that is interpretable (as opposed to a key-value store).

JavaScript Object Notation (JSON): Human-readable text format for data structures. A JSON document consists of possibly nested key-value pairs.

Document-Based Data Model

A JSON document consists of a JSON object, which is enclosed by curly braces, { . . . }. Inside, keys and the corresponding values are separated by a colon, and a value can be a JSON object itself.

Example Document Store:

UserID	Shopping Cart
1002	{ "1": { "type": "Shoe", "name": "Jordans", "color": "red", "size": 37 }, "2": { "type": "Computer game", "name": "Diablo", "publisher": "Blizzard" }, "3": { "gear": "Headset", "producer": "Razer", "model": "Kraken Kitty" } }
...	...

Graph-Based Data Model

Graphs: Informally, graphs are structures that represent data (as nodes) and their interrelation (as edges in between) by design. Both nodes and edges may carry information. Efficient graph operations are supported natively.

Data and their relationship are distinct naturally (cf. semantic overloading of the relational data model).

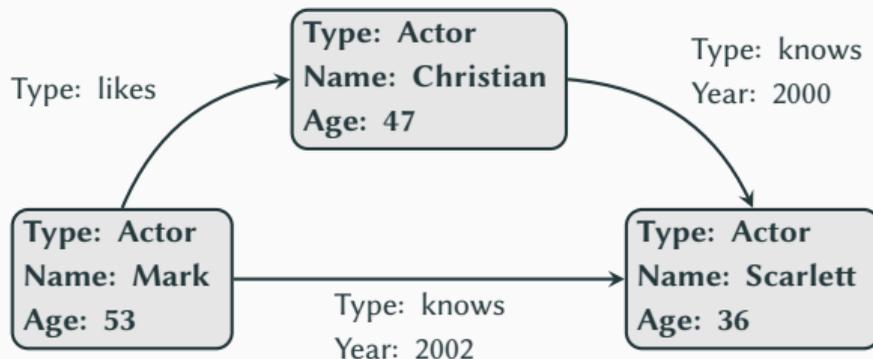
Example Graph:



Graph-Based Data Model

Property Graph Model: Multiple types for nodes/edges (multi-relational), each of which may contain multiple properties (attributes) as name-value pairs (similar to key-value pairs).

Example Property Graph:



Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



- **(Declarative) query processing:** Parser, optimizer, and evaluation engine.
- **Types of data models:** Conceptual, logical, and physical.
- **Relational data model** and its disadvantages.
- **Other data models:**
 - Object-based data models.
 - Key-value data model.
 - Document-based data model.
 - Graph-based data model.

Extensible Record Data Model

Also referred to as **wide column data model** and a generalization of Google's BigTable⁸ system.

Table cells are represented as (2-dimensional) key-value pairs, with row (unique) and column (repeatedly) being the keys.

Example: Access table cell (4:year).

id	name	year	runtime	genre
1	The Godfather	1972	177	Mafia
2	The Avengers	2012	143	NULL
3	Batman - The Dark Knight	2008	152	Action
4	American Psycho	2000	101	Thriller
5	Home Alone	1990	103	Comedy

⁸Chang et al. *Bigtable: A Distributed Storage System for Structured Data*. OSDI, 2006.

Extensible Record Data Model

Also referred to as **wide column data model** and a generalization of Google's BigTable⁸ system.

Table cells are represented as (2-dimensional) key-value pairs, with row (unique) and column (repeatedly) being the keys.

Example: Access table cell (4:year).

id	name	year	runtime	genre
1	The Godfather	1972	177	Mafia
2	The Avengers	2012	143	NULL
3	Batman - The Dark Knight	2008	152	Action
4	American Psycho	2000	101	Thriller
5	Home Alone	1990	103	Comedy

⁸Chang et al. *Bigtable: A Distributed Storage System for Structured Data*. OSDI, 2006.

Extensible Record Data Model

Column Families: Group columns that are accessed simultaneously.

1	<table><thead><tr><th>name</th></tr></thead><tbody><tr><td>The Godfather</td></tr></tbody></table>	name	The Godfather	<table><thead><tr><th>year</th></tr></thead><tbody><tr><td>1972</td></tr></tbody></table>	year	1972	<table><thead><tr><th>runtime</th></tr></thead><tbody><tr><td>177</td></tr></tbody></table>	runtime	177	<table><thead><tr><th>genre</th></tr></thead><tbody><tr><td>Mafia</td></tr></tbody></table>	genre	Mafia
name												
The Godfather												
year												
1972												
runtime												
177												
genre												
Mafia												
2	<table><thead><tr><th>name</th></tr></thead><tbody><tr><td>The Avengers</td></tr></tbody></table>	name	The Avengers	<table><thead><tr><th>year</th></tr></thead><tbody><tr><td>2012</td></tr></tbody></table>	year	2012	<table><thead><tr><th>runtime</th></tr></thead><tbody><tr><td>143</td></tr></tbody></table>	runtime	143			
name												
The Avengers												
year												
2012												
runtime												
143												
...												

CoreInfo:

1	<table><thead><tr><th>name</th></tr></thead><tbody><tr><td>The Godfather</td></tr></tbody></table>	name	The Godfather	<table><thead><tr><th>year</th></tr></thead><tbody><tr><td>1972</td></tr></tbody></table>	year	1972
name						
The Godfather						
year						
1972						
2	<table><thead><tr><th>name</th></tr></thead><tbody><tr><td>The Avengers</td></tr></tbody></table>	name	The Avengers	<table><thead><tr><th>year</th></tr></thead><tbody><tr><td>2012</td></tr></tbody></table>	year	2012
name						
The Avengers						
year						
2012						
...						

ExtendedInfo:

1	<table><thead><tr><th>runtime</th></tr></thead><tbody><tr><td>177</td></tr></tbody></table>	runtime	177	<table><thead><tr><th>genre</th></tr></thead><tbody><tr><td>Mafia</td></tr></tbody></table>	genre	Mafia
runtime						
177						
genre						
Mafia						
2	<table><thead><tr><th>runtime</th></tr></thead><tbody><tr><td>143</td></tr></tbody></table>	runtime	143			
runtime						
143						
...						

Encourages de-normalization (redundancy) for higher performance and provides **more flexibility**, i.e., every row may be composed of different columns.

The query workload defines how the data is modeled \Rightarrow Know your workload.

Array-Based Data Model

Complex structures (like 2-dimensional satellite or n -dimensional sensor data) are organized along multiple dimensions.

An array cell contains a tuple of a specific length and a tuple element may contain a value or another array (allowing arbitrary nestings).

Specialized array functionality is natively supported (e.g., joins and aggregations).

Q&A

Workloads & Challenges

- New data management **challenges**.
- Basics of **distributed** database systems (also compared to parallel DBS).
- **Consistency guarantees**: ACID vs. BASE.
- The **CAP** Theorem.
- **Workloads**: OLTP vs. OLAP, batch vs. stream processing.

New Data Management Challenges

- Data may be organized in **complex structures** (e.g., graphs).
- Data may be **schemaless** (e.g., not comply to a fixed schema).
- Data may be **sparse** (e.g., values may be non-existent).
- Data or schema may be **constantly changing** (e.g., schema evolution).
- Data may be **distributed over multiple machines**.
 - Access must be **transparent** (i.e., user need not know where data resides).
 - Systems must **scale horizontally** (i.e., new machines come and go).
 - Systems must cope with **large data volumes**.

New Data Management Challenges

NoSQL (Not Only SQL):

- Class of non-relational DBSs.
- Weaker consistency guarantees (e.g., BASE⁹)
- Support for schema independence.
- Highly scalable.

NewSQL:

- Class of relational DBSs.
- Provide NoSQL-like scalability (for OLTP⁹ workloads).
- Retain consistency guarantees of RDBMS (e.g., ACID⁹).

⁹We will cover this term subsequently.

Parallel database management systems (PDBMSs) have multiple processors and hard disks that are connected via a fast interconnection.

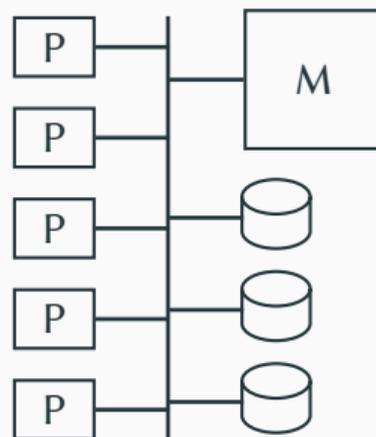
Performance characteristics:

- **Throughput:** Number of tasks (e.g., queries) that can be completed in a given time frame. Example: Queries per seconds.
- **Response time:** Time it takes to complete a single task (e.g., query).

Parallel Database System Architectures

Shared memory DBMSs have many processors and disks that share a common memory (typically via a bus).

- + Efficient communication between processors ($< 1\mu s$).
- Limited scalability (≤ 64 processors; interconnection to memory becomes the bottleneck).

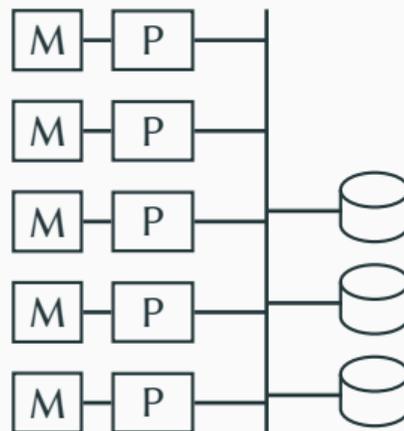


P ... processor, M ... memory, \ominus ... disks.

Parallel Database System Architectures

Shared disk DBMSs have many processors (with isolated memory) that share all disks (typically via a bus).

- + Scale to a larger number of processors.
- Communication between processors is slower (ms; bottleneck now at interconnection to disks).

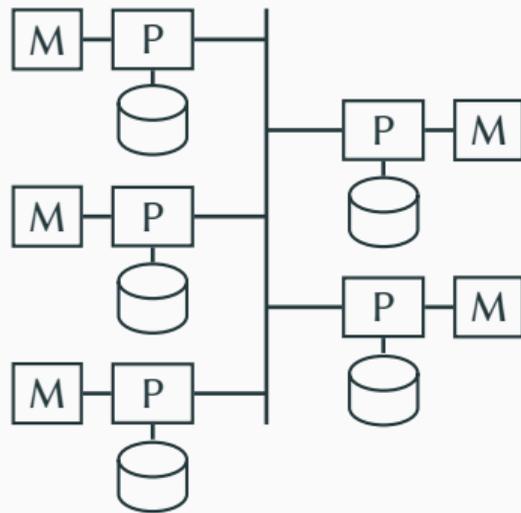


P ... processor, M ... memory,  ... disks.

Parallel Database System Architectures

Shared nothing DBMSs have many processors with isolated memory and disk(s). The combination of a processor with isolated memory and disk(s) is also referred to as **node**.

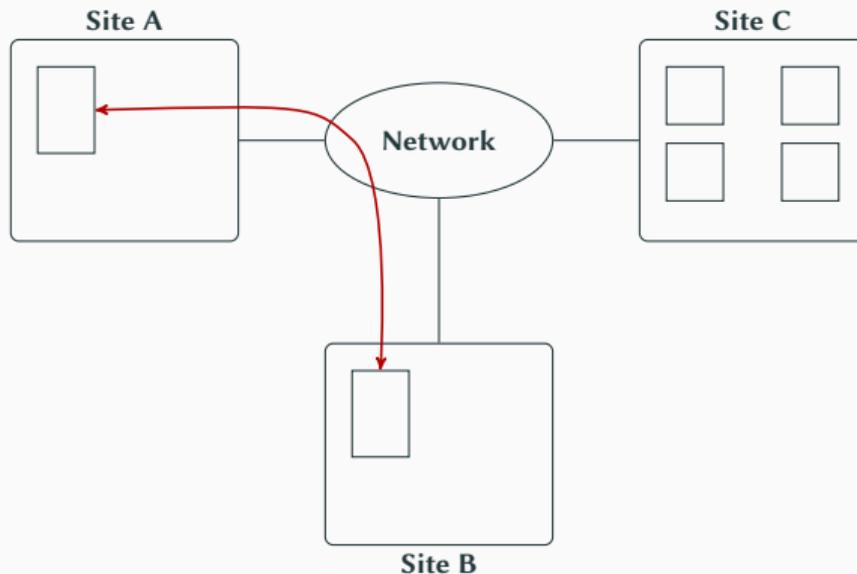
- + Scale to thousands of processors.
- Communication between processors is slow; access to non-local disk data is slow.



P ... processor, M ... memory, ⊕ ... disks.

Distributed Database Systems

Distributed database management systems (DDBMSs) are DBMSs that operate on multiple, geographically separated machines (also called **sites**).



DDBMS vs. Shared-Nothing PDBMS

- Sites within a DDBMS are typically
 - **geographically separated** (i.e., not a single data center)
⇒ lower bandwidth (less throughput), higher latency (higher response time).
 - **separately administered** (i.e., retain some degree of autonomy).
- PDBMS can deal with node failures, whereas **DDBMS** can deal with **failures of entire sites** (e.g., due to natural disasters).
- DDBMS distinguish between **local** and **global transactions**.

Homogeneous vs. Heterogeneous DDBMS

A DDBMS is called **homogeneous** if the nodes share a common global database schema, perform the same tasks (e.g., run the same software), and actively cooperate in processing. Goal: View of a single database.

A DDBMS is called **heterogeneous** if the nodes have different schemata, perform different tasks (e.g., run different software), and may not be aware of other nodes. Goal: Integrate different databases.

A **transaction** refers to a **sequence of operations** that accesses and (possibly) updates various data items. A transaction **transitions the database from one consistent state into another consistent state**.

Example: Transfer EUR 500 from bank account A to bank account B.

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

A DBMS must deal with **two major issues**:

- **System crash**: Software or hardware failures.
- **Concurrency**: Many different users may work at the same time (i.e., multiple transactions are executed).

A DBS implements particular consistency guarantees. Most relational DBMS manage transactions according to the so-called **ACID properties**: **A**tomicity, **C**onsistency, **I**solation, and **D**urability.

The ACID properties are considered **strong consistency guarantees**.

Atomicity: Execute **all operations** of a transaction **or none** of them (*“all or nothing”*).

Example: Transfer EUR 500 from bank account A to bank account B.

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

What happens if the system crashes after step 4?

Consistency: After a transaction, all values in the database are correct (i.e., consistency is preserved).

Example: Transfer EUR 500 from bank account A to bank account B.

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

What is a consistency constraint in the above example?

Consistency: After a transaction, all values in the database are correct (i.e., consistency is preserved).

Example: Transfer EUR 500 from bank account A to bank account B.

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

What is a consistency constraint in the above example? $A + B$ is the same before and after the transaction.

Isolation: A transaction must be unaware of other simultaneously executing transactions (otherwise an inconsistent state may be encountered).

Example: Transfer EUR 500 from bank account A to bank account B.

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

Imagine a second transaction BEGIN; READ(A); READ(B); print(A + B); COMMIT in between steps 4 and 5.

Durability: Successful transaction results persist in the database, even if the system crashes.

Example: Transfer EUR 500 from bank account A to bank account B.

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

If the system crashes and the EUR 500 are lost, nobody is happy.

Documentation -- PostgreSQL 13

Supported Versions: [Current \(13\)](#) / [12](#) / [11](#) / [10](#) / [9.6](#)

Development Versions: [dev](#)

Unsupported versions: [9.5](#) / [9.4](#) / [9.3](#) / [9.2](#) / [9.1](#) / [9.0](#) / [8.4](#) / [8.3](#) / [8.2](#) / [8.1](#) / [8.0](#) / [7.4](#) / [7.3](#) / [7.2](#)

[Prev](#)[Up](#)

3.4. Transactions
Chapter 3. Advanced Features

[Home](#)[Next](#)

3.4. Transactions

Transactions are a fundamental concept of all database systems. The essential point of a transaction is that **it bundles multiple steps into a single, all-or-nothing operation**. The intermediate states between the steps are not visible to other concurrent transactions, and if some failure occurs that prevents the transaction from completing, then none of the steps affect the database at all.

For example, consider a bank database that contains balances for various customer accounts, as well as total deposit balances for branches. Suppose that we want to record a payment of \$100.00 from Alice's account to Bob's account. Simplifying outrageously, the SQL commands for this might look like:

```
UPDATE accounts SET balance = balance - 100.00
  WHERE name = 'Alice';
UPDATE branches SET balance = balance - 100.00
  WHERE name = (SELECT branch_name FROM accounts WHERE name = 'Alice');
UPDATE accounts SET balance = balance + 100.00
  WHERE name = 'Bob';
UPDATE branches SET balance = balance + 100.00
  WHERE name = (SELECT branch_name FROM accounts WHERE name = 'Bob');
```

The details of these commands are not important here; the important point is that there are several separate updates involved to accomplish this rather simple operation. Our bank's officers will want to be assured that either all these updates happen, or none of them happen. It would certainly not do for a system failure to result in Bob receiving \$100.00 that was not debited from Alice. Nor would Alice long remain a happy customer if she was debited without Bob being credited. We need a guarantee that if something goes wrong partway through the operation, none of the steps executed so far will take effect. **Grouping the updates into a transaction gives us this guarantee. A transaction is said to be atomic**; from the point of view of other transactions, it either happens completely or not at all.

We also want a guarantee that once a transaction is completed and acknowledged by the database system, **it has indeed been permanently recorded and won't be lost even if a crash ensues shortly thereafter**. For example, if we are recording a cash withdrawal by Bob, we do not want any chance that the debit to his account will disappear in a crash just after he walks out the bank door. A transactional database guarantees that all the updates made by a transaction are logged in permanent storage (i.e., on disk) before the transaction is reported complete.

Another important property of transactional databases is closely related to the notion of atomic updates: **when multiple transactions are running concurrently, each one should not be able to see the incomplete changes made by others**. For example, if one transaction is busy totalling all the branch balances, it would not do for it to include the debit from Alice's branch but not the credit to Bob's branch, nor vice versa. So transactions must be all-or-nothing not only in terms of their permanent effect on the database, but also in terms of their visibility as they happen. The updates made so far by an open transaction are invisible to other transactions until the transaction completes, whereupon all the updates become visible simultaneously.

In PostgreSQL, a transaction is set up by surrounding the SQL commands of the transaction with `BEGIN` and `COMMIT` commands. So our banking transaction would actually look like:

```
BEGIN;
UPDATE accounts SET balance = balance - 100.00
  WHERE name = 'Alice';
-- etc etc
COMMIT;
```

- Announcement via Blackboard.
- Give anonymous feedback regarding the lecture (up to now).
- Give anonymous feedback regarding assignment 1.

The assessment helps us to improve the course and is very much appreciated!

After-Assignment 1 Meetings

- Available dates/times will be provided via Blackboard (groups).
- Choose a date/time that fits your schedule.
- **One member of your group must enroll** to this date via Blackboard.
- First come first serve.
- Grading will be based on the last submission before the meeting.

Assignment 2

Summary:

- **Where:** Release and submission via Blackboard.
- **When:** April 19 – May 17, 2021 (resp. late: May 24, 2021).
- **What to do:** (a) Set up a document-based database locally (MongoDB; data will be provided), (b) get familiar with JSON and learn how to execute queries (most queries will be given), (c) write a small Python3 application that executes the queries, and (d) answer questions regarding the assignment.
- **What to submit:** The Python3 code and the answers to the questions.
- **Grading:** 55% Python3 code, 45% questions (incl. the meeting).

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



- **New challenges** in data management (NoSQL vs. NewSQL).
- Basics of **parallel database management systems (PDBMS)**.
- Basics of **distributed database management systems (DDBMS)**.
- Basics of **transactions**.
- The **ACID properties** (atomicity, consistency, isolation, durability).

- Give anonymous feedback regarding the lecture (up to now).
- Give anonymous feedback regarding assignment 1.

The feedback helps us to improve the course and is very much appreciated!

Programmers must ensure to properly define the transactions to preserve consistency.

A DBS typically includes a so-called **transaction manager**, which ensures that the transactions comply to the consistency guarantees (e.g., the ACID properties).

A transaction has **committed** if it completes successfully. Otherwise, the transaction is **aborted**. Undoing the changes of an aborted transaction is referred to as **rollback**.

Transaction Management – Atomicity

Before the DBS changes the database, it writes some information (transaction identifier, old/new values) about the changes into a so-called **log file**.

Example: Transfer EUR 500 from bank account A (EUR 800) to account B (EUR 2,000).

Transaction T_1 :

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
- [CRASH]---
7. WRITE(B)
8. COMMIT

Pseudo log file:

T1: OLD(A=800), NEW(A=300)

Database:

UPDATE(A=300)

ROLLBACK T1 on restart

Strict serial execution of concurrent transactions guarantees isolation, but severely limits the performance.

A **concurrency-control scheme** ensures that transactions can execute concurrently (i.e., their operations can be interleaved). Interleaving the operations of a transaction is called **schedule** and may result in a correct database state, or not!

A concurrent schedule is **serializable** if an equivalent serial schedule exists, i.e., the outcome of executing the transactions concurrently is the same as if they would have been executed serially.

Transaction Management – Isolation (Lock-Based)

Example: Transfer EUR 500 from bank account A (EUR 800) to account B (EUR 2,000).

Transaction T_1 :

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)

5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

Transaction T_2 :

1. BEGIN
2. READ(A)
3. $A = A + 1,000$
4. WRITE(A)
5. COMMIT

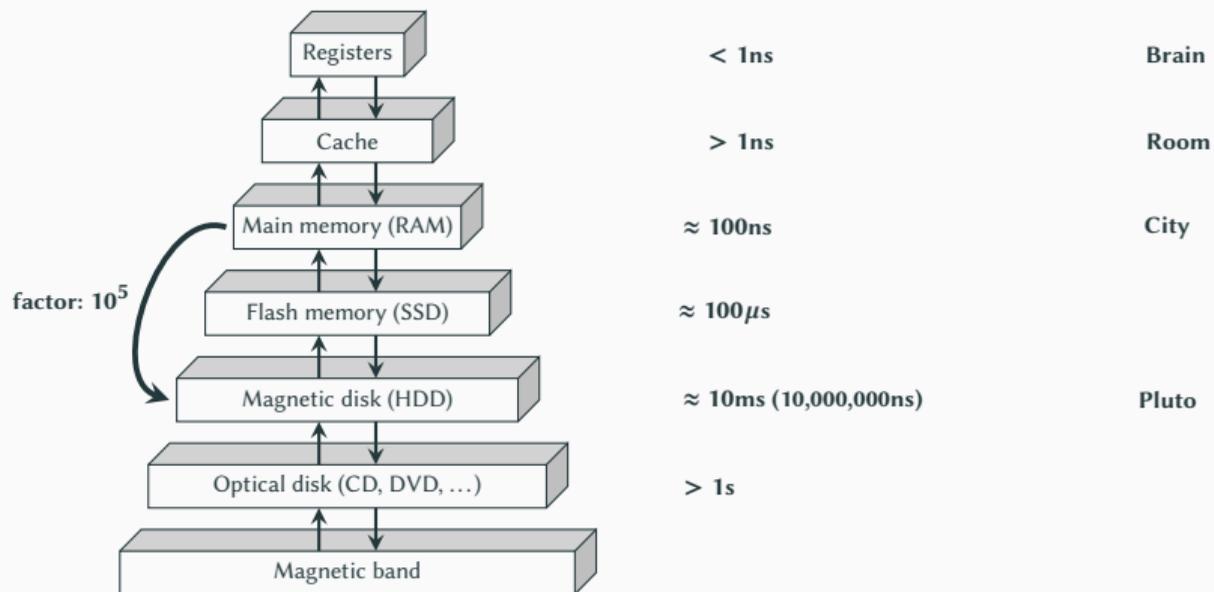
Database:

- | | |
|------------------------------|----------|
| LOCK(A, T1) | ok |
| UPDATE(A=300) | |
| LOCK(A, T2) | conflict |
| LOCK(B, T1) | ok |
| UNLOCK(A, T1), UNLOCK(B, T1) | ok |

Updating the persistent data may be postponed to improve performance, i.e., the data is provisionally updated in main memory (RAM). Data in main memory is volatile, i.e., it is lost on system restart.

Excursion: Nowadays, data can be considered persistent if it is written to hard disk. DBMSs aim to be efficient/fast, thus they try to **avoid** or **postpone expensive/slow operations** like accesses to hard disk. Multiple data structures are maintained in main memory, which are commonly referred to as **buffers**. The content of the buffers is **only written to disk if inevitable**.

Memory Hierarchy



Transaction Management – Durability

Example: Transfer EUR 500 from bank account A (EUR 800) to account B (EUR 2,000).

Transaction T_1 :

1. BEGIN
2. READ(A)
3. $A = A - 500$
4. WRITE(A)
5. READ(B)
6. $B = B + 500$
7. WRITE(B)
8. COMMIT

Pseudo log file:

T1: OLD(A=800), NEW(A=300)

T1: OLD(B=2,000), NEW(B=2,500)

Database:

postponed UPDATE(A=300)

postponed UPDATE(B=2,500)

---[CRASH]---

PSQL Transactions in Action

A DDBMS must consider additional aspects:

- **Distributed data storage** (replication and fragmentation).
- **Distributed transactions** (local and global).

Distributed Data Storage

Assume the movies table of our example database and the relational data model.

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

Distributed Data Storage

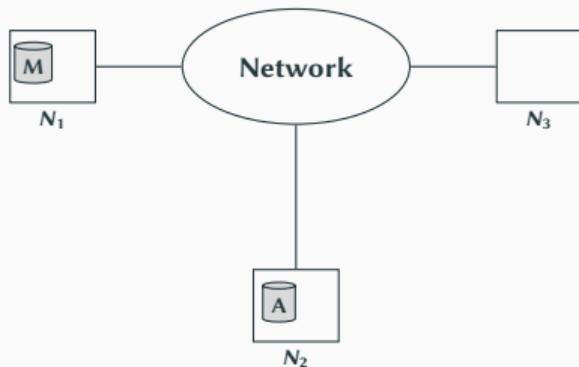
Assume the movies table of our example database and the relational data model.

Movies:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi
...

Actors:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather
...

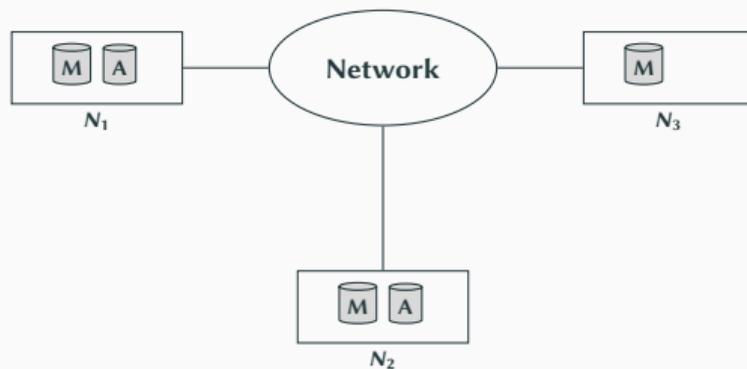


What if N_1 gets disconnected?

Distributed Data Storage – Replication

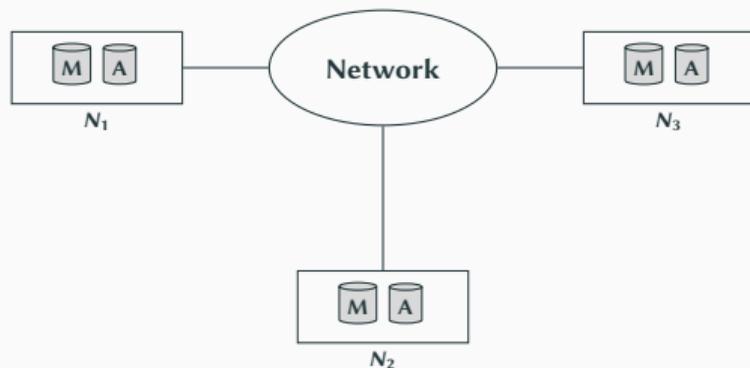
Replication: Data are **replicated** among multiple sites, i.e., a copy (replica) of the same data exists in different sites (intentional redundancy).

Example Replication:



Distributed Data Storage – Replication

Full replication: A copy of a relation is stored at **all sites**.



Fully replicated database: Every site contains a copy of the entire database.

Pros: Higher availability and performance, reduced data transfer.

Cons: Increased update costs and higher complexity of concurrency control.

Replication

A *replica set* in MongoDB is a group of [mongod](#) processes that maintain the same data set. Replica sets provide redundancy and [high availability](#), and are the basis for all production deployments. This section introduces replication in MongoDB as well as the components and architecture of replica sets. The section also provides tutorials for common tasks related to replica sets.

Redundancy and Data Availability

Replication provides redundancy and increases data availability. With multiple copies of data on different database servers, replication provides a level of fault tolerance against the loss of a single database server.

In some cases, replication can provide increased read capacity as clients can send read operations to different servers. Maintaining copies of data in different data centers can increase data locality and availability for distributed applications. You can also maintain additional copies for dedicated purposes, such as disaster recovery, reporting, or backup.

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



- **Transaction management** for ACID guarantees.
- Memory hierarchy.
- **Transactions (ACID) in action** (PostgreSQL).
- **Distributed data storage** (replication).

MongoDB Replication in Action

Fragmentation: Data are **partitioned into fragments** stored in distinct sites, i.e., a specific part of the data is stored in a site.

Horizontal fragmentation: Relation is split row- or tuple-wise, and each row/tuple resides in a separate site.

Vertical fragmentation: Relation is split into subschemata (based on the columns/attributes), and each subschema resides in a separate site.

Example Horizontal Fragmentation

Movies M1:

name	year	runtime	genre
The Godfather	1972	177	Mafia
The Avengers	2012	143	Sci-Fi

Movies M2:

name	year	runtime	genre
Batman - The Dark Knight	2008	152	Action
American Psycho	2000	101	Thriller

Movies M3:

name	year	runtime	genre
Home Alone	1990	103	Comedy

Actors A1:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather

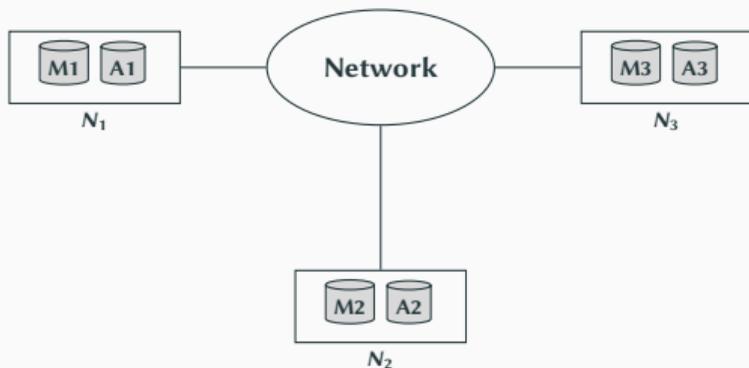
Actors A2:

name	birthyear	knownfor
Macaulay Culkin	1980	Home Alone
Joe Pesci	1943	Home Alone

Actors A3:

name	birthyear	knownfor
Robert Downey Jr.	1965	The Avengers
Mark Ruffalo	1967	The Avengers

Example Horizontal Fragmentation



Actors A1:

name	birthyear	knownfor
Marlon Brando	1924	The Godfather
Al Pacino	1940	The Godfather

Actors A2:

name	birthyear	knownfor
Macaulay Culkin	1980	Home Alone
Joe Pesci	1943	Home Alone

Actors A3:

name	birthyear	knownfor
Robert Downey Jr.	1965	The Avengers
Mark Ruffalo	1967	The Avengers

Horizontal and vertical fragmentation can be **mixed**. In any case, the fragmented data must be **reconstructable**.

Pros: Higher performance (on fragments), better locality (reduced data transfer).

Cons: Increased costs on site failure and if data is retrieved from different sites.

Local transactions access/update data at only one (local) site. A local transaction manager enforces the ACID properties.

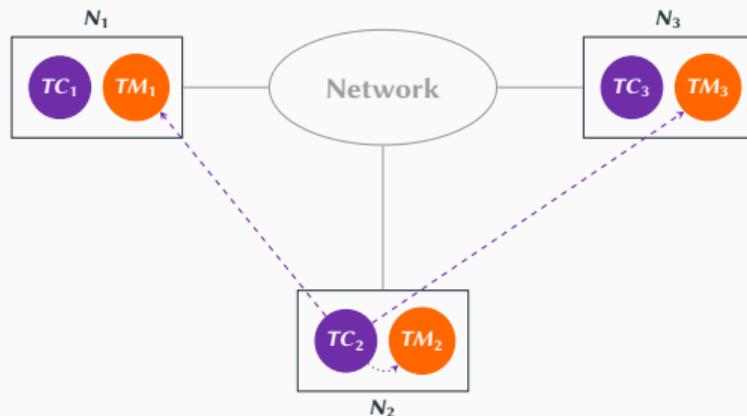
Global transactions access/update data at multiple (local) sites. Local subtransaction are executed at each site. Enforcing the ACID properties is much more complex.

Each site has a **local transaction manager** and a **transaction coordinator**.

Distributed Transactions

The **local transaction manager** ensures that the ACID properties hold for local transactions and maintains the local log files.

The **transaction coordinator** starts transactions that are initiated at a site, distributes the subtransactions to other sites, and ensures that a transaction either executes **at all sites or at none**.



Required because we must ensure **atomicity across all sites**, e.g., we are not allowed to commit a transaction at N_1 but abort the same transaction at N_2 .

List of Protocols:

- Two-Phase Commit (2PC; used in practice).
- Three-Phase Commit (3PC; solves issues of 2PC but impracticable).
- Persistent Messaging (PM).

Distributed Transactions – 2PC in a Nutshell

Transaction T is initiated at site N_i with transaction coordinator TC_i .

Phase 1: TC_i “asks” other participants to **prepare to commit** (and logs it beforehand). Each **transaction manager** determines if **it can commit**, logs it, and “reports” it to TC_i .

Phase 2: If TC_i receives a **single abort** message, then all participants are **informed to abort T** . Otherwise, all participants are **notified to commit T** . Prior to that, TC_i logs the decision locally. The involved transaction managers comply to the decision.

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg



Midterm Exam (German)

- Structured Query Language (SQL).
- Integritätsbedingungen, valide Instanz.
- Datenabstraktion (physische Datenunabhängigkeit).
- Instanz vs. Schema, Primärschlüssel.
- Operationen in einem Key-Value Store.
- Durchsatz in einem parallelen DBMS (PDBMS).

- Warum deklarative Sprachen und nicht direkt Ausführungsplan?
- Prinzip hinter Datenunabhängigkeit. Warum erstrebenswert?
- Wünschenswertes Szenario für Redundanz. Warum?
- Prinzip hinter Transaktionen bzgl. des DB-Zustandes?
- Serieller vs. serialisierbarer Ablaufplan.
- Rolle des Loggings bzgl. ACID. Vorteil der Logs?

- **Replication** exemplified by **MongoDB**.
- **Fragmentation** (horizontal and vertical).
- **Distributed transactions** (local vs. global, transaction manager vs. coordinator).
- **Commit Protocols**, in particular **Two-Phase Commit**.

How to **connect** to a MongoDB replica set **transparently**?

Different notions of consistency to specify desired properties in a DDBMS. Ideally, all updates appear immediately at all sites in the same order (illusion of a single data copy).

Strong consistency refers to this ideal scenario, but this is often expensive (or even impracticable). **Weak consistency** relaxes the consistency constraints to improve the performance or the availability of a DDBMS.

High Availability: A DDBMS with extremely low downtime (about 99.99% available).

In large systems, a failure happens frequently (nodes may be down or the network may partition).

Trade consistency to achieve high availability.

Brewer's CAP Theorem

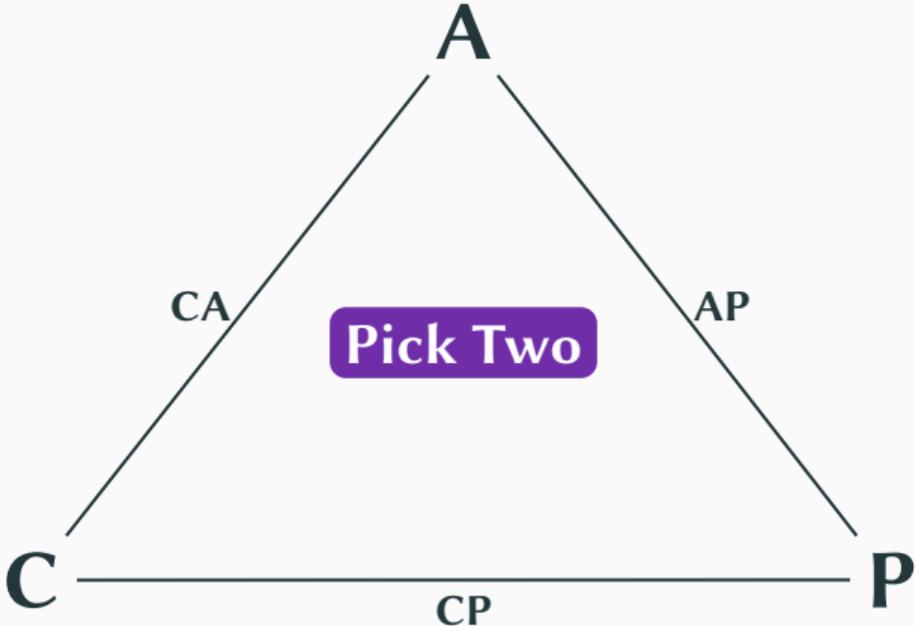
A distributed database system has three properties:

- **Consistency:** All replicated copies are in the same state.
- **Availability:** System runs even in case of failures due to replication.
- **Partition-tolerance:** System runs even if the network is partitioned.

Network partition: Network decomposes into multiple parts/subsystems that cannot reach one another.

CAP Theorem: You can have **at most two** of the three **properties**.

Brewer's CAP Theorem



Brewer's CAP Theorem

AP Systems: Systems stay available in case of a network partition, but inconsistencies may be introduced. These inconsistencies must be resolved once the network partition is resolved.

CP Systems: Systems maintain consistency in case of a network partition, but may be unavailable for some time.

CA Systems: “Ideal” consistency and availability. No availability and consistency guarantees in case of a network partition.

In Practice: Network partitions cannot be avoided \Rightarrow AP or CP.

The BASE Properties

BASE is an **alternative consistency model** that is **not as strict** as the ACID properties and favors **availability over consistency**.

Basically Available: The system appears to **work most of the time**, i.e., reads/writes should be allowed even if the network partitions, but without consistency guarantee.

Soft State: The **state** of the database **may not be precisely defined** all the time, i.e., replicas do not have to be consistent.

Eventually Consistent: Once the network partitioning is resolved, the states of all replicas converge, i.e., all **replicas become consistent eventually**.

Eventual Consistency

No new updates \Rightarrow Writes are propagated to all replicas and all replicas converge towards a common state.

Inconsistent replicas must be identified because two replicas may be updated independently (e.g., version-vector scheme).

Inconsistent updates may need to be **merged** (e.g., in the worst case, human interaction is required – comparable to a merge conflict in git).

With regard to database systems, a **workload** is a set of queries/updates that reflects a **typical usage pattern (load)**.

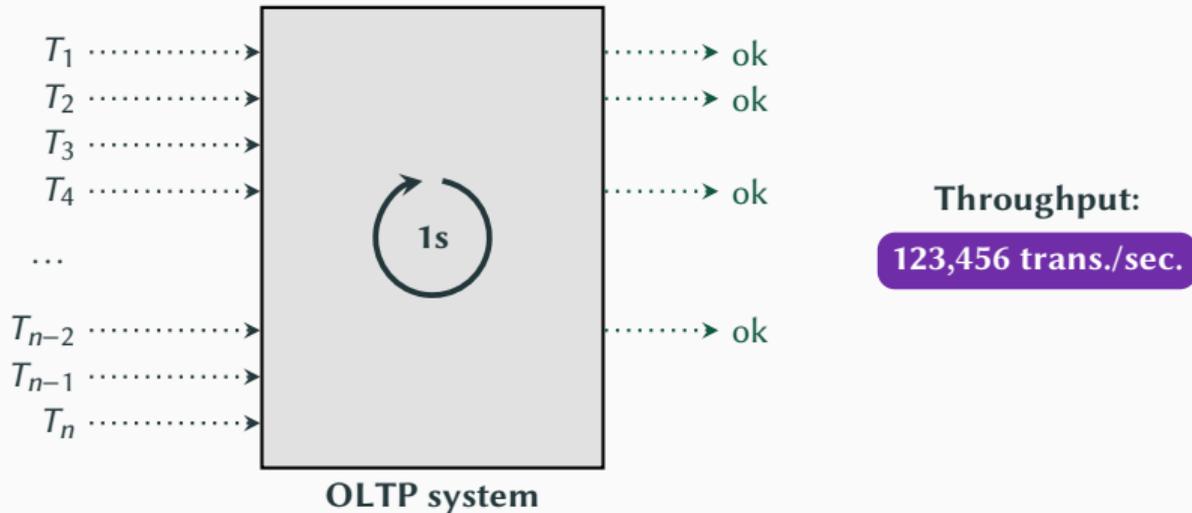
Different database systems perform better/worse on particular workloads.

Transaction Processing Performance Council (TPC): An independent consortium that releases standardized benchmarks for various workloads, the TPC benchmarks ¹⁰.

¹⁰<http://www.tpc.org/>

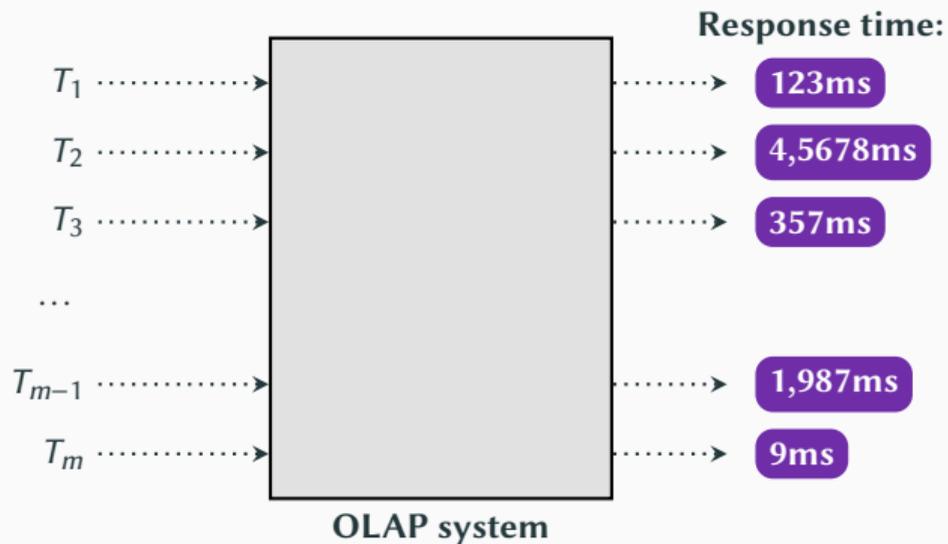
OLTP Workloads

Online Transaction Processing (OLTP): Short-lived read/write transactions with a small footprint (i.e., only a small portion of the data is touched). **Many transactions** must be processed **as fast as possible** (high throughput, low response time). The TPC-C benchmark provides typical OLTP workloads.



OLAP Workloads

Online Analytical Processing (OLAP): Long-running read-only queries that exploratory analyze a large portion of the overall data (to support decision). This often involves complex join operations and the main focus is **low response time**. The TPC-DS benchmark provides typical OLAP workloads. “Data Warehouse”.



Batch Workloads: A **batch** is a large but **bounded static dataset**. Before data can be processed, all data must be completely available (e.g., on hard disk).

Stream Workloads: A **stream** is an **unbounded evolving dataset**. Data items are processed as they stream into the system one after another, i.e., the data does not have to be completely available.

Q&A

Start: 08:15 am CET

Distributed Information Management

Daniel Kocher

Salzburg, Summer term 2021

Department of Computer Sciences
University of Salzburg

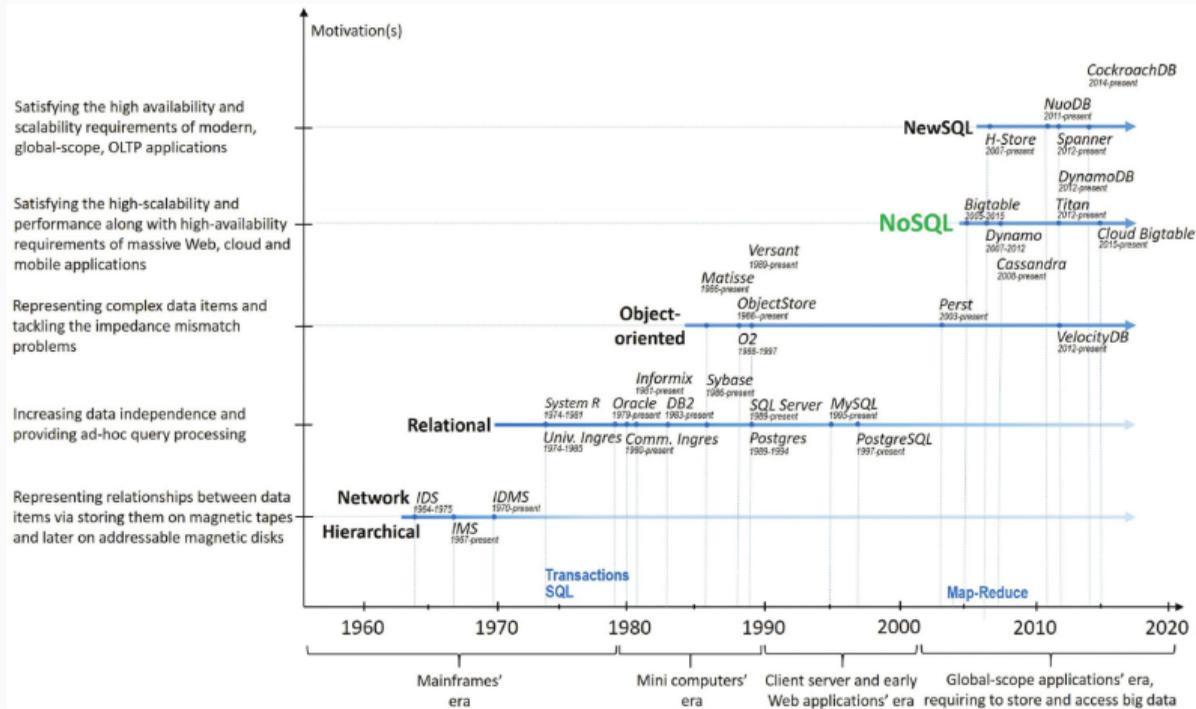


Recap

- Strong vs. weak consistency.
- High availability.
- Brewer's CAP theorem.
- The BASE properties and eventual consistency.
- OLTP vs. OLAP workloads.
- Batch vs. Stream workloads.

Systems Potpourri

Continuous Development



Picture taken from Davoudian et al. *A Survey on NoSQL Stores*. ACM Computing Surveys, 2018.

In-Memory Database Systems: Relies on **main memory (RAM)** for storing the **data** rather than hard disks (HDDs) or SSDs. Very **low response times (microseconds)** but must deal with durability (logs).

Database as a Service (DBaaS): A cloud-based platform (service) that provides computing infrastructure, data storage and **database functionality in the cloud**. The clients (you) do not have to set up their own database system on their own hardware, but just **access and use a database that runs in the cloud**.

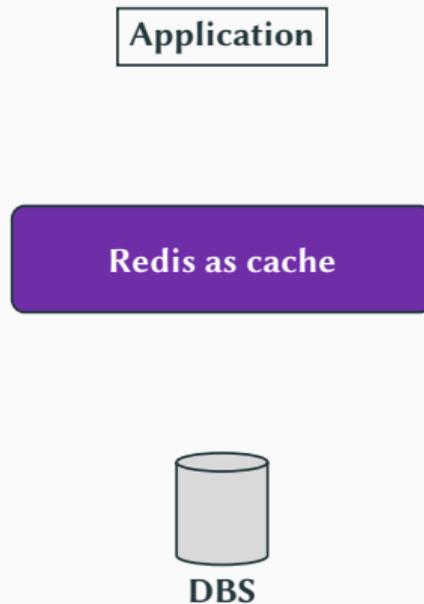
The **RE**mote **DI**ctionary **S**erver ¹¹ is an **in-memory** NoSQL database system based on the **key-value data model** that **supports durability** (i.e., data can be made persistent).

- Redis is a CP system.
- In-memory, i.e., very fast (avg. read/write performance: < 1ms).
- Collection of useful, high-performance data structures (lists, sets, bitmaps, ...).
- Replication and persistence support.
- Supports many programming languages incl. Java, Python, C/C++, and JavaScript.

¹¹<https://redis.io>

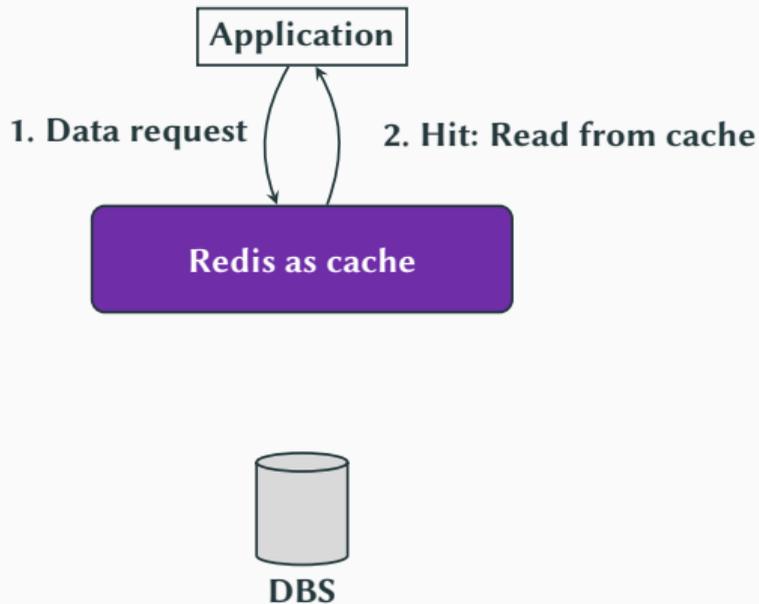
Key-Value Stores – Redis

Use Case: Cache between application and database system.



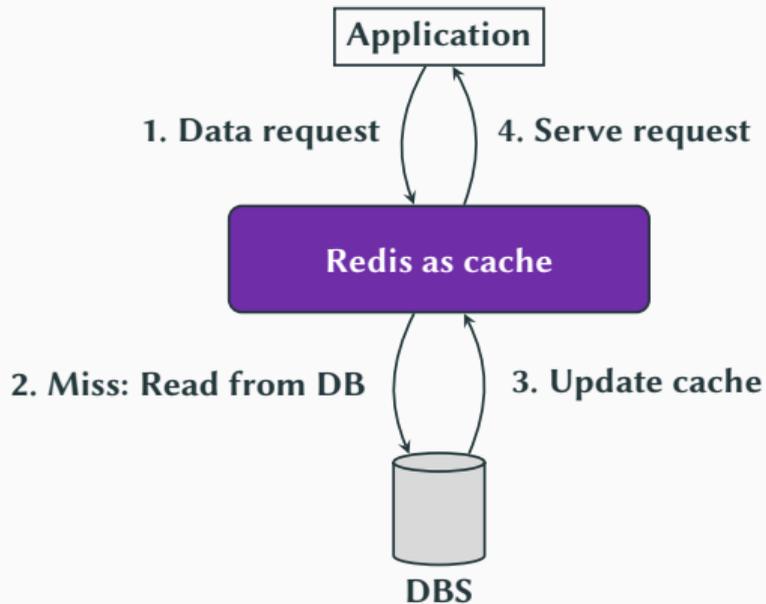
Key-Value Stores – Redis

Use Case: Cache between application and database system.



Key-Value Stores – Redis

Use Case: Cache between application and database system.



Key-Value Stores – Other Systems

Riak KV ¹² is a distributed **persistent** key-value store with support for **advanced data types** (e.g., JSON). It is inspired by Amazon’s DynamoDB ¹³ and provides **high availability** (i.e., is an AP system).

Memcached ¹⁴ (“Mem-Cash-Dee”) is a distributed **in-memory** key-value store with a performance similar to Redis. However, Memcached supports only the **simple key-value data model** and has **no support for durability**. Not a database system!

¹²<https://riak.com/products/riak-kv/>

¹³<https://aws.amazon.com/dynamodb/>

¹⁴<https://memcached.org/>

MongoDB ¹⁵ is a distributed NoSQL database system based on the document-oriented data model with a focus on high availability and horizontal scalability.

- MongoDB is a CP system.
- Uses JSON as document format (specifically, a binary JSON format called BSON).
- Shared-nothing system architecture.
- JSON-based query language (limited support for joins as we know it).
- Replication and sharding (aka fragmentation) support through replica sets.

¹⁵<https://www.mongodb.com/>

Apache CouchDB ¹⁶ is a distributed document store with support for **master-master replication** and a conflict resolution protocol. It is an AP system that uses the plain JSON format.

Couchbase ¹⁷ is a distributed document store with a built-in in-memory cache (memcached) and an SQL-like query language (N1QL). It uses CouchDB as back end.

¹⁶<https://couchdb.apache.org/>

¹⁷<https://www.couchbase.com/>

HandsOn: N1QL¹⁸

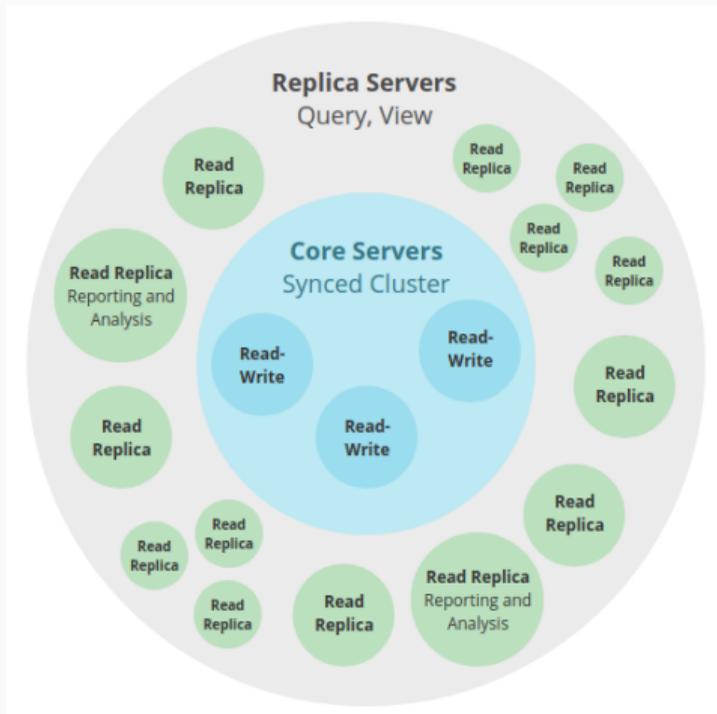
¹⁸<https://query-tutorial.couchbase.com/>

Neo4j ¹⁹ is a NoSQL database system based on the graph data model with ACID guarantees. Neo4j is a **native** graph store.

- Neo4j is a CP system (as of version 3.5).
- Shared-nothing system architecture.
- Set of core servers (serve write+read) and many read replicas.
- Emphasis on read availability (full replication).
- SQL-inspired query language (Cypher) to describe graph patterns (ASCII art).

¹⁹<https://neo4j.com/>

HandsOn: Neo4j



Graph Stores – Other Systems

Tigergraph ²⁰ is a distributed native graph store that supports parallel computation and advanced analytics. It has its own query language GSQL.

ArangoDB ²¹ is a multi-model database system that supports the graph-based model in JSON format.

OrientDB ²² is another multi-model database system that supports the graph-based model. It uses an SQL-like syntax that is extended for graphs.

²⁰<https://www.tigergraph.com/>

²¹<https://www.arangodb.com/>

²²<https://orientdb.org/>

Apache Cassandra ²³ is a distributed NoSQL database system based on the extensible record (or wide column) data model. It is a decentralized and “master-less” DDBMS that supports eventual consistency.

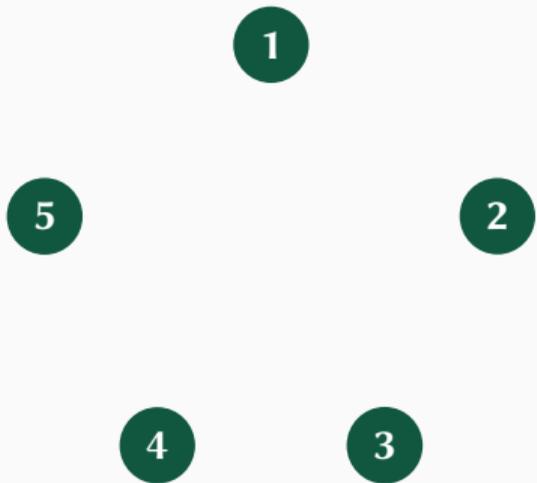
- Apache Cassandra is an AP system (but can be configured as CP system).
- Shared-nothing system architecture.
- Support for nested column families (so-called super column families).
- The Cassandra query language (CQL) resembles SQL.

²³<https://cassandra.apache.org/>

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

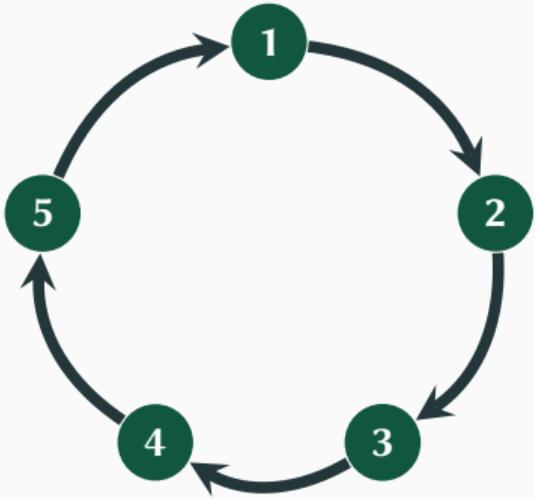
Consistent Hashing (Token Ring)



Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

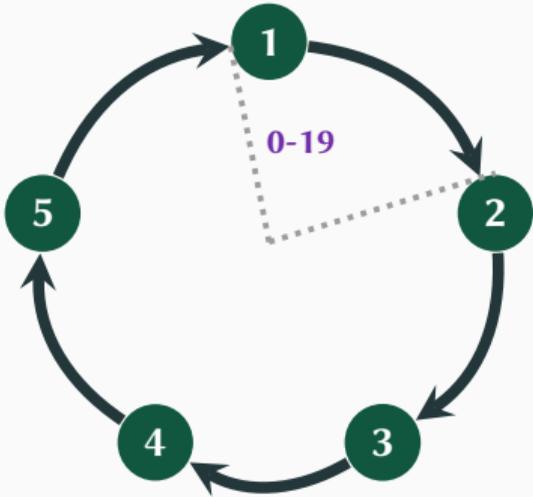
Consistent Hashing (Token Ring)



Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

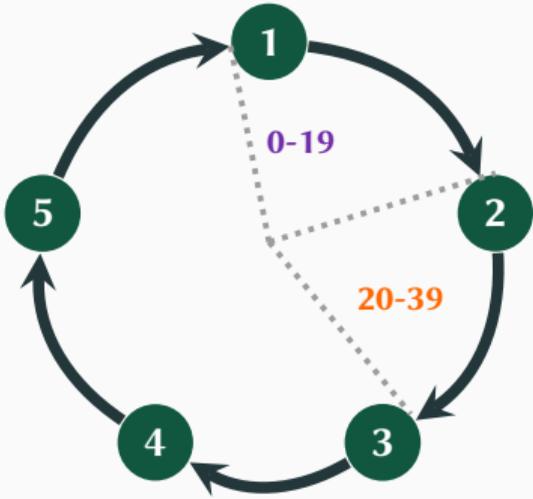
Consistent Hashing (Token Ring)



Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

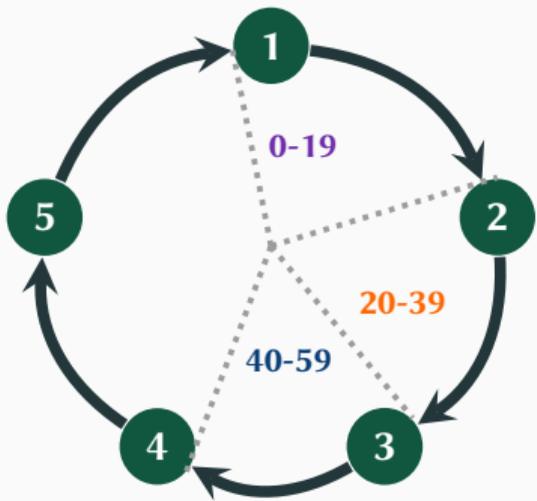
Consistent Hashing (Token Ring)



Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

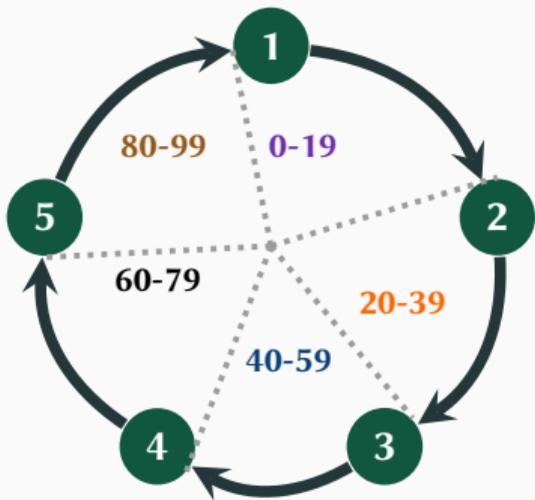
Consistent Hashing (Token Ring)



Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Consistent Hashing (Token Ring)



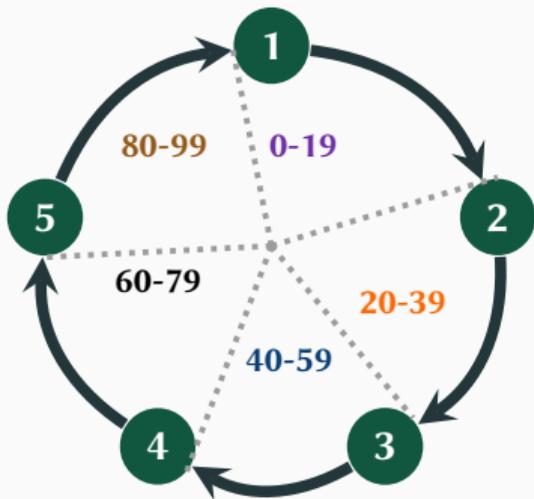
Hash function $h(\cdot)$ determines node.

	Data		$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



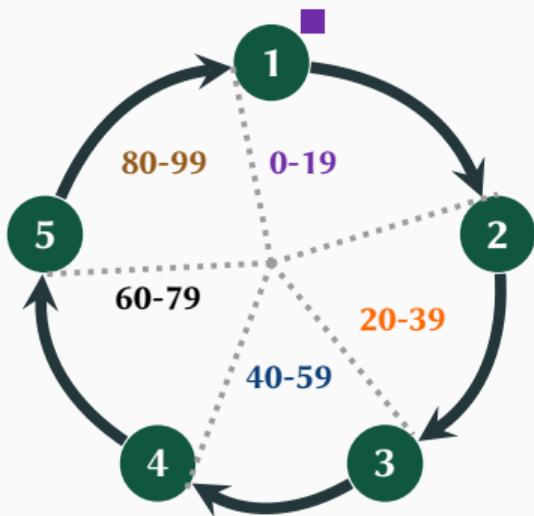
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



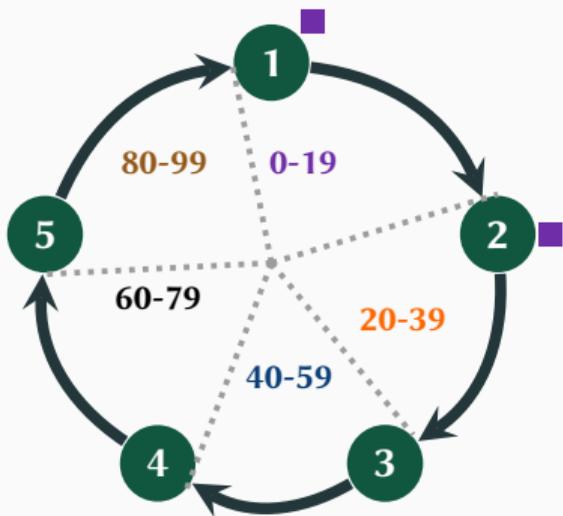
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



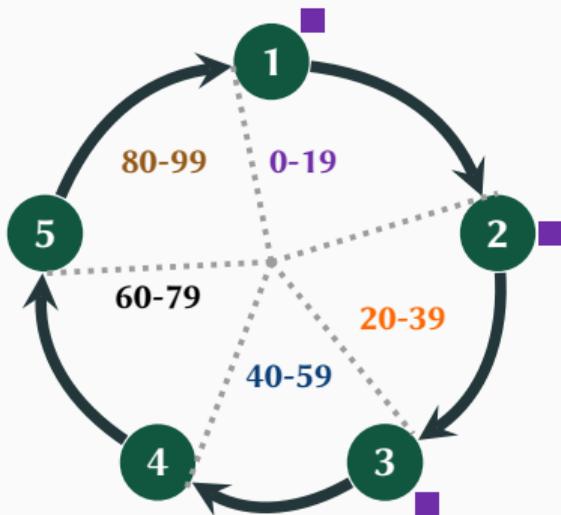
Hash function $h(.)$ determines node.

Data			$h(.)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



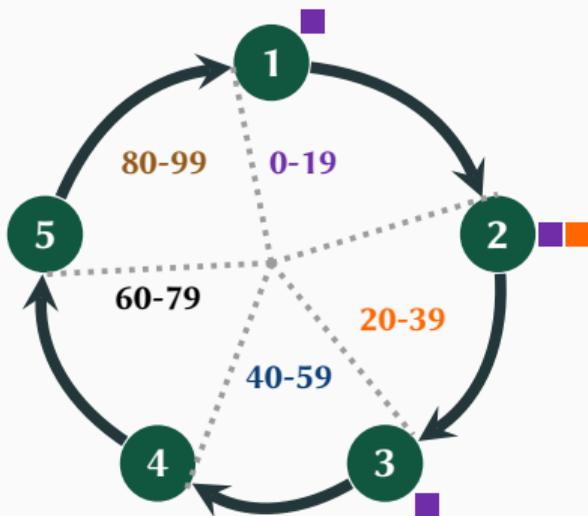
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



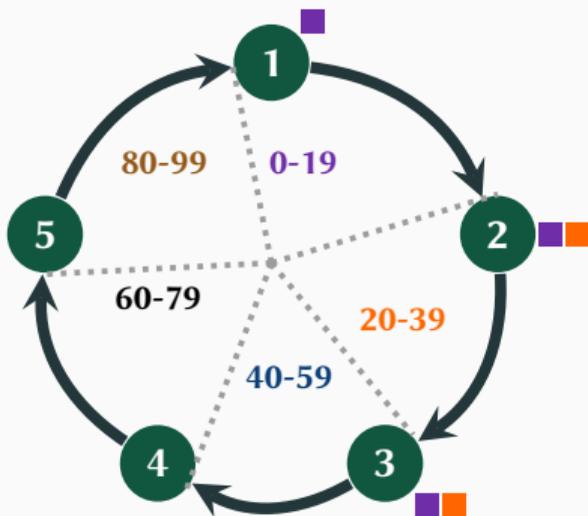
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



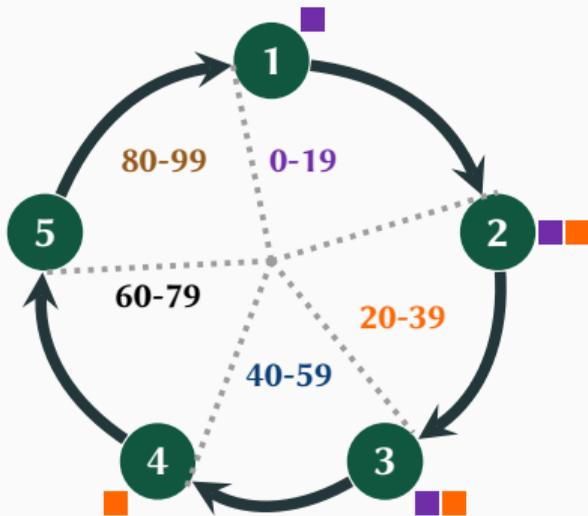
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



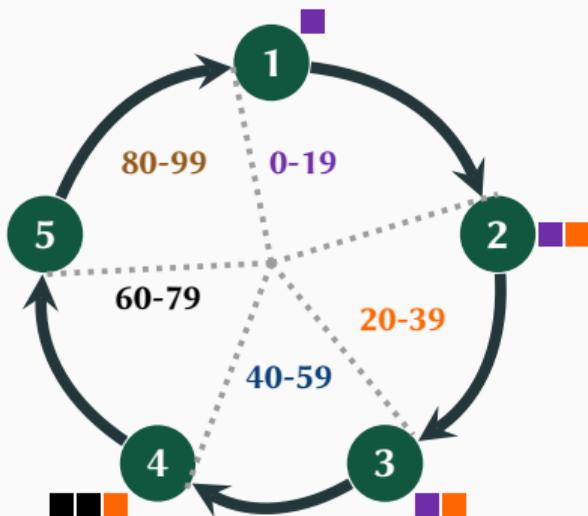
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



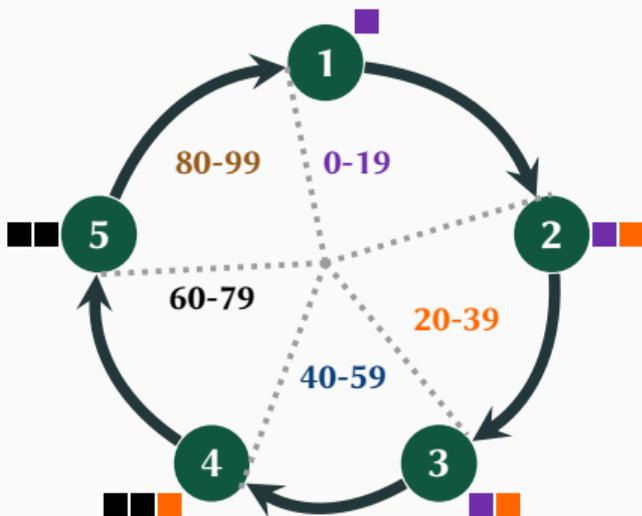
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



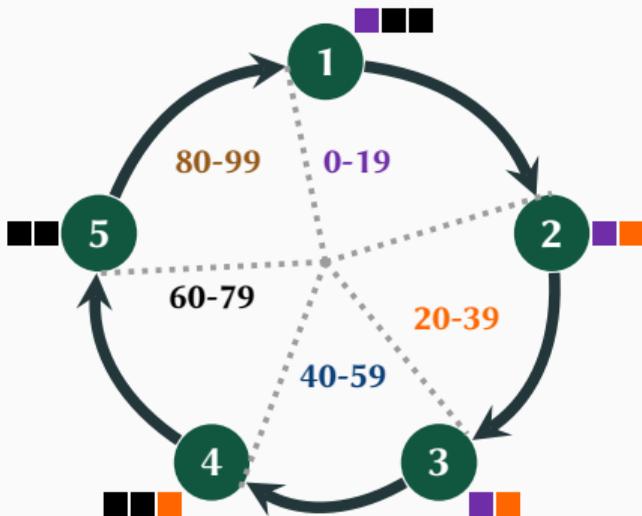
Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Extensible Record Store – Apache Cassandra

Master-less architecture \Rightarrow **No single point of failure.**

Replication Factor (for example, 3)



Hash function $h(\cdot)$ determines node.

Data			$h(\cdot)$	Node
Christian	Bale	1974	18	1
Scarlett	Johansson	1984	21	2
Mark	Ruffalo	1967	67	4
Mark	Wahlberg	1971	67	4

Other Systems

- **RDBMS:** PostgreSQL ²⁴, SQLite ²⁵, Oracle ²⁶, Microsoft SQL Server ²⁷, ...
- **NewSQL Systems:** CockroachDB ²⁸, VoltDB ²⁹, ...
- **In-Memory OTLP & OLAP Systems:** HyPer ³⁰, SAP Hana ³¹, ...
- **Cloud Services:** Amazon DynamoDB ³², Snowflake Cloud Data Warehouse ³³, ...

²⁴<https://www.postgresql.org/>

²⁵<https://www.sqlite.org/>

²⁶<https://www.oracle.com/database/>

²⁷<https://www.microsoft.com/en-us/sql-server/sql-server-2019>

²⁸<https://www.cockroachlabs.com/>

²⁹<https://www.voltdb.com/>

³⁰<https://hyper-db.de/>

³¹<https://www.sap.com/products/hana.html>

³²<https://aws.amazon.com/dynamodb/>

³³<https://www.snowflake.com/>

Use the **database of databases** ³⁴ for a first impression and cross references.

A 37-page **survey on NoSQL database systems** ³⁵

Books on NoSQL database systems ^{36 37}

³⁴<https://dbdb.io/>

³⁵Davoudian et al. *A Survey on NoSQL Stores*. ACM Computing Surveys, 2018.

³⁶Sadalage and Fowler. *NoSQL Distilled – A Brief Guide to the Emerging World of Polyglot Persistence*. Addison-Wesley, 2013.

³⁷Redmond and Wilson. *7 Databases in 7 Weeks - A Guide to Modern Databases and the NoSQL Movement*. Pragmatic Bookshelf, 2012.

Q&A