

# Big Data Analytics

## B. Distributed Storage / B.3 NoSQL Databases

Lars Schmidt-Thieme

Information Systems and Machine Learning Lab (ISMLL)  
Institute of Computer Science  
University of Hildesheim, Germany

# Syllabus

- |            |      |  |
|------------|------|--|
| Tue. 4.4.  | (1)  | 0. Introduction  |
|            |      | <b>A. Parallel Computing</b>                           |
| Tue. 11.4. | (2)  | A.1 Threads  |
| Tue. 18.4. | (3)  | A.2 Message Passing Interface (MPI)                    |
|            |      | <b>B. Distributed Storage</b>                          |
| Tue. 2.5.  | (4)  | B.1 Distributed File Systems                           |
| Tue. 9.5.  | (5)  | B.2 Partitioning of Relational Databases               |
| Tue. 16.5. | (6)  | B.3 NoSQL Databases                                    |
| Tue. 23.5. | (7)  | A.3 Graphical Processing Units (GPUs)                  |
|            |      | <b>C. Distributed Computing Environments</b>           |
| Tue. 30.5. | (8)  | C.1 Map-Reduce   |
| Tue. 6.6.  | —    | — Pentecoste Break —                                   |
| Tue. 13.6. | (9)  | C.2 Resilient Distributed Datasets (Spark)             |
|            |      | <b>D. Distributed Machine Learning Algorithms</b>      |
| Tue. 20.6. | (10) | D.1 Distributed Stochastic Gradient Descent            |
| Tue. 27.6. | (11) | D.2 Distributed Matrix Factorization                   |
| Tue. 4.7.  | (12) | D.3 Alternating Direction Method of Multipliers (ADMM) |

# Outline

1. Introduction
2. Key-Value Stores
3. Document Databases
4. Document Databases: Partitioning
5. Graph Databases
6. Column Databases and Object Databases

# Outline

1. Introduction
2. Key-Value Stores
3. Document Databases
4. Document Databases: Partitioning
5. Graph Databases
6. Column Databases and Object Databases

# Relational vs Big-Data Technologies

## ▶ Structure:

- ▶ relational data bases: data is **structured**.
- ▶ big data applications: data often is **raw**.

## ▶ Process:

- ▶ relational data bases: initially created for **transactional processing**.
- ▶ big data applications: **analytical processing**.

## ▶ Entities:

- ▶ relational data bases: big if it has **many entities** (rows) of a kind.
- ▶ big data applications: the number of entities is not necessary high, the amount of information that exist for entities may be large.

## ▶ horizontal scaling:

- ▶ relational data bases: costs for adding new nodes is high.
- ▶ big data applications: scaling should be accomplished inexpensively by **adding new nodes**.

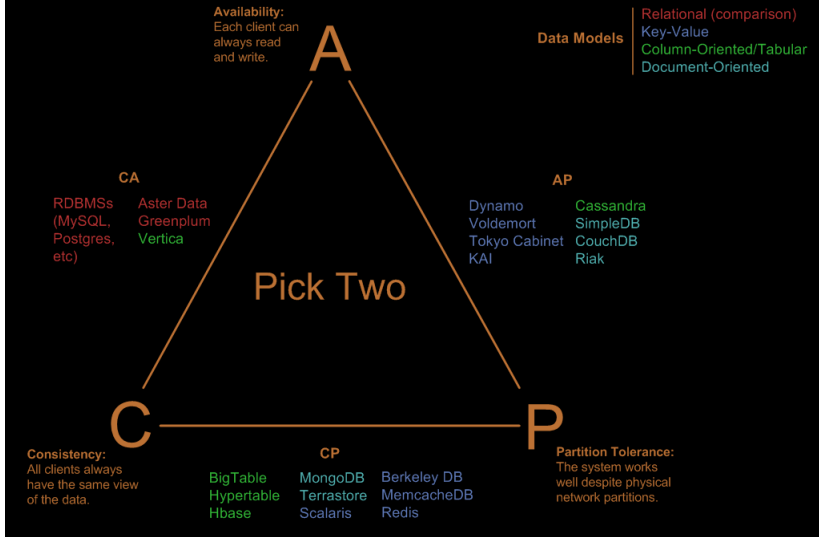
# NoSQL Databases

- ▶ A NoSQL or Not Only SQL database provides a mechanism for storage and retrieval of data that is modeled in means **other than the tabular relations** used in relational databases.
- ▶ Motivations for this approach include
  - ▶ **simplicity of design** and
  - ▶ **horizontal scaling**.
- ▶ The data structure differs from the RDBMS, and therefore
  - ▶ some operations are faster in NoSQL, and
  - ▶ some operations are faster in RDBMS.
- ▶ Most NoSQL stores lack true ACID transactions.

# NoSQL Databases / Types and Implementations

- ▶ **Key-value:** Dynamo, FoundationDB, MemcacheDB, Redis, Riak
- ▶ **Document:** Clusterpoint, Couchbase, MarkLogic, MongoDB
- ▶ **Graph:** Allegro, Neo4J, OrientDB, Virtuoso
- ▶ **Column:** Accumulo, Cassandra, HBase
- ▶ **Object:** Db4o, ZODB

# Visual Guide to NoSQL Systems



<http://blog.scottlogic.com/2014/08/04/mongodb-vs-couchdb.html>



# Outline

1. Introduction
2. Key-Value Stores
3. Document Databases
4. Document Databases: Partitioning
5. Graph Databases
6. Column Databases and Object Databases

# Key-Value stores

- ▶ Key-Value stores use the **associative array** (dictionary) as their fundamental data model.
  - ▶ Data is represented as a collection of key-value pairs.
- ▶ One of the simplest non-trivial data models.

Example:

```
{  
  "Great Expectations": "John",  
  "Pride and Prejudice": "Alice",  
  "Wuthering Heights": "Alice"  
}
```

# Outline

1. Introduction
2. Key-Value Stores
- 3. Document Databases**
4. Document Databases: Partitioning
5. Graph Databases
6. Column Databases and Object Databases

# Document Databases

- ▶ Data abstraction:
  - ▶ relational databases: "relations" (= "tables").
  - ▶ document databases: "document".
- ▶ Documents are (possibly nested) **dictionaries**.
  - ▶ like Python dicts.
  - ▶ Documents are not required to have all the same fields (aka sections, slots, parts).
  - ▶ Documents are **schemaless**.
- ▶ Documents are identified by an ID, e.g.,
  - ▶ marked by a special type **Objectid**,
  - ▶ stored by a special key **id**.
- ▶ References are modeled by **foreign keys**.
  - ▶ or avoided by using **embedded documents**.

# Example 1 / Schema-less

```
1 {
2   _id: ObjectId(7df78ad8902c),
3   FirstName: "Bob",
4   Age: 35,
5   Address: "5_Oak_St.",
6   Hobby: "sailing"
7 }
```

```
1 {
2   _id: ObjectId(5df78ad8902c),
3   FirstName: "Jonathan",
4   Age: 37,
5   Address: "15_Wanamassa_Point_Road",
6   Languages: [ 'English', 'German' ]
7 }
```

## Example 2 / Foreign Keys

```
1 {
2   _id: ObjectId(5df78ad8902c),
3   FirstName: "Jonathan",
4   Age: 37,
5   Address: "15_Wanamassa_Point_Road",
6   Children: [ ObjectId(5df78ad89020), ObjectId(5df78ad89021),
7             ObjectId(5df78ad89022), ObjectId(5df78ad89023) ]
8 },{
9   _id: ObjectId(5df78ad89020),
10  FirstName: "Michael",
11  Age: 10,
12 },{
13  _id: ObjectId(5df78ad89021),
14  FirstName: "Jennifer",
15  Age: 8,
16 },{
17  _id: ObjectId(5df78ad89022),
18  FirstName: "Samantha",
19  Age: 5,
20 },{
21  _id: ObjectId(5df78ad89023),
22  FirstName: "Elena",
23  Age: 2,
24 }
```

# Example 3 / Embedded Documents

```
1 {
2   _id: ObjectId(5df78ad8902c),
3   FirstName: "Jonathan",
4   Age: 37,
5   Address: "15_Wanamassa_Point_Road",
6   Children: {
7     { FirstName: "Michael", Age: 10 },
8     { FirstName: "Jennifer", Age: 8 },
9     { FirstName: "Samantha", Age: 5 },
10    { FirstName: "Elena", Age: 2}
11  }
12 }
```

# Organization

- ▶ Documents are addressed in the database via a **unique key**.
- ▶ Documents can be retrieved by their
  - ▶ key
  - ▶ content
- ▶ Documents are organized through
  - ▶ Collections
  - ▶ Tags
  - ▶ Non-visible Metadata
  - ▶ Directory hierarchies
  - ▶ Buckets



# RDBMS vs. Document Database Terminology

| relational database | document database  |
|---------------------|--|
| Database            | Database   |
| Table               | Collection   |
| Tuple/Row           | Document   |
| Column              | Field  |
| Primary key         | Primary key<br>(e.g., default key <code>_id</code> in mongodb) |
| Foreign key         | Foreign key  |

# Inserting Documents

- ▶ **insert** allows to insert a new document.
- ▶ **insertMany** allows to insert many documents.

```
1 db.inventory.insertMany([
2   { item: "journal", qty: 25, size: { h: 14, w: 21, uom: "cm" }, status: "A" },
3   { item: "notebook", qty: 50, size: { h: 8.5, w: 11, uom: "in" }, status: "A" },
4   { item: "paper", qty: 100, size: { h: 8.5, w: 11, uom: "in" }, status: "D" },
5   { item: "planner", qty: 75, size: { h: 22.85, w: 30, uom: "cm" }, status: "D" },
6   { item: "postcard", qty: 45, size: { h: 10, w: 15.25, uom: "cm" }, status: "A" }
7 ]);
```

Note: Most examples taken from [MongoDB(2017)].

# Query Documents (1/3)

## ► all documents of a collection:

```
1 db.inventory.find()
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9a"), "item" : "journal", "qty" : 25, "size" : { "h" : 14, "w" : 20 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9b"), "item" : "notebook", "qty" : 50, "size" : { "h" : 8.5, "w" : 15 } }
3 { "_id" : ObjectId("59116d5340229e45bb5eea9c"), "item" : "paper", "qty" : 100, "size" : { "h" : 8.5, "w" : 15 } }
4 { "_id" : ObjectId("59116d5340229e45bb5eea9d"), "item" : "planner", "qty" : 75, "size" : { "h" : 22.85, "w" : 15 } }
5 { "_id" : ObjectId("59116d5340229e45bb5eea9e"), "item" : "postcard", "qty" : 45, "size" : { "h" : 10, "w" : 15 } }
```

## ► all documents with a property:

```
1 db.inventory.find({ status: "D" })
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9c"), "item" : "paper", "qty" : 100, "size" : { "h" : 8.5, "w" : 15 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9d"), "item" : "planner", "qty" : 75, "size" : { "h" : 22.85, "w" : 15 } }
```

## ► all documents with several properties:

```
1 db.inventory.find({ status: "D", qty: 100 })
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9c"), "item" : "paper", "qty" : 100, "size" : { "h" : 8.5, "w" : 15 } }
```

## Query Documents (2/3)

- ▶ instead of querying for exact value matches, one can use query operators:
  - ▶ **\$lt**, **\$gt**, **\$lte**, **\$gte**: numerical comparison

```
1 db.inventory.find({ qty: { $gte: 75 } } )
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9c"), "item" : "paper", "qty" : 100, "size" : { "h" : 8.5, "w" : 22.85 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9d"), "item" : "planner", "qty" : 75, "size" : { "h" : 22.85, "w" : 10 } }
```

```
1 db.inventory.find({ qty: { $lte: 75, $gt: 25 } } )
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9b"), "item" : "notebook", "qty" : 50, "size" : { "h" : 8.5, "w" : 22.85 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9d"), "item" : "planner", "qty" : 75, "size" : { "h" : 22.85, "w" : 10 } }
3 { "_id" : ObjectId("59116d5340229e45bb5eea9e"), "item" : "postcard", "qty" : 45, "size" : { "h" : 10, "w" : 15 } }
```

- ▶ **//**: regular expressions.

```
1 db.inventory.find({ item: /er$/ } )
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9c"), "item" : "paper", "qty" : 100, "size" : { "h" : 8.5, "w" : 22.85 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9d"), "item" : "planner", "qty" : 75, "size" : { "h" : 22.85, "w" : 10 } }
```

# Query Documents (3/3)

## ► \$or or queries:

```
1 db.inventory.find({ $or: [ { qty: { $lte: 25 } }, { qty: { $gte: 100 } } ] } )
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9a"), "item" : "journal", "qty" : 25, "size" : { "h" : 14, "w" : 10 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9c"), "item" : "paper", "qty" : 100, "size" : { "h" : 8.5, "w" : 10 } }
```

# Query Documents (4/4)

- fields of nested documents are queried by dot syntax:

```
1 db.inventory.find({ "size.h": { $gte: 14 } } )
```

```
1 { "_id" : ObjectId("59116d5340229e45bb5eea9a"), "item" : "journal", "qty" : 25, "size" : { "h" : 14, "w" : 22.85 } }
2 { "_id" : ObjectId("59116d5340229e45bb5eea9d"), "item" : "planner", "qty" : 75, "size" : { "h" : 22.85, "w" : 14 } }
```

# Advantages and Disadvantages

- ▶ avoiding foreign keys by **embedding documents**:
  - + expensive join operations are not required.
  - possibly redundant information.
- ▶ **schemaless**:
  - + one can add a new field at any time.
  - one never can be sure a field actually has a value.

# Big Data Document Databases

Document databases are useful for big data for two main reasons:

1. Document databases can be **horizontally partitioned / sharded**.
  - ▶ Documents are distributed over different nodes.
  - ▶ Using a **partition/sharding function**
    - ▶ e.g., a hash function.
  - ▶ works exactly the same ways as for RDBMS.



# Big Data Document Databases

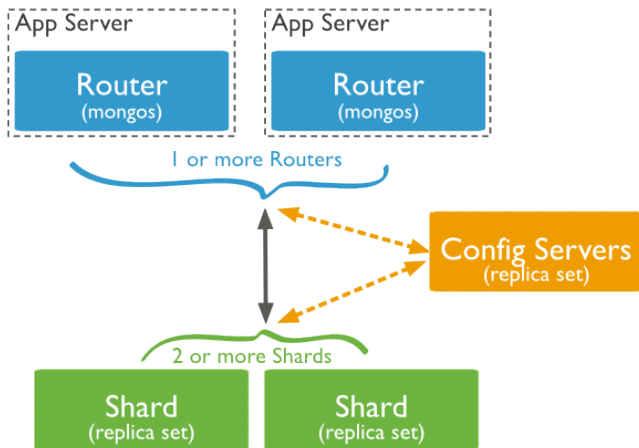
Document databases are useful for big data for two main reasons:

1. Document databases can be **horizontally partitioned / sharded**.
  - ▶ Documents are distributed over different nodes.
  - ▶ Using a **partition/sharding function**
    - ▶ e.g., a hash function.
  - ▶ works exactly the same ways as for RDBMS.
2. Documents are **sparse representations**,
  - ▶ only some fields/keys have values,  
while tuples/rows are **dense**.
  - ▶ all columns have values stored explicitly
    - ▶ also NULL is explicitly stored.

# Outline

1. Introduction
2. Key-Value Stores
3. Document Databases
- 4. Document Databases: Partitioning**
5. Graph Databases
6. Column Databases and Object Databases

# Sharded Document Database / Architecture



[source: <https://docs.mongodb.com/manual/core/sharding-introduction/>]

Mongo DB (from humongous – slang for enormous)

# System Setup / Multiple Hosts

1. install mongodb on all hosts (h0, h1, h2, h3),  
e.g., for OpenSuSE:

```
1 zypper in mongodb
```

2. start a **config server** on one host (e.g., h0):

```
1 mongod --configsvr
```

▶ default port is 27019

3. start a **query router** on one host (e.g., h0):

```
1 mongos --configdb localhost
```

▶ default port is 27017

4. start a couple of **shard servers** on other hosts (e.g., h1, h2, h3):

```
1 mongod --shardsrv
```

▶ default port is 27018

5. add shards:

```
1 mongo --host h0
2 sh.addShard("h1")
3 sh.addShard("h2")
```

# System Setup / Single Hosts (for Testing)

## 1. install mongodb, e.g., for OpenSuSE:

```
1 zypper in mongodb
```

## 2. start a **config server**:

```
1 mongod --configsvr --dbpath db-configsrv/
```

- ▶ config information will be stored in directory db-configsrv/.

## 3. start a **query router**:

```
1 mongos --configdb localhost:27019
```

- ▶ query routers do not require any data path

## 4. start a couple of **shard servers**:

```
1 mongod --port 27021 --dbpath db-shardsrv1/  
2 mongod --port 27022 --dbpath db-shardsrv2/  
3 mongod --port 27023 --dbpath db-shardsrv3/
```

- ▶ use different ports and data paths

## 5. add shards:

```
1 mongo  
2 sh.addShard("localhost:27021")  
3 sh.addShard("localhost:27022")
```

# Database and Collection Setup

## 1. enable sharding per database

```
1 sh.enableSharding("mydb")
```

## 2. sharding a collection requires the sharding key to be indexed:

```
1 db.ijcnn1.createIndex( { _id: 1 } )
```

## 3. shard per collection:

```
1 sh.shardCollection("mydb.ijcnn1", { _id: 1 })
```

## 4. import data

```
1 mongoimport --db mydb --collection ijcnn1 --drop --file ijcnn1.json
```

- ▶ data has to be in json format
- ▶ if a database is imported first and then sharded, it will be distributed across nodes automatically by the balancer.

# Database and Collection Setup

- ▶ production databases also have to be replicated
  - ▶ both, data nodes and config nodes

# Database and Collection Setup

- ▶ production databases also have to be replicated
  - ▶ both, data nodes and config nodes
  
- ▶ to work on sharded big data, one should
  - ▶ **not** query it from the central router nodes
    - ▶ rule of thumb: 1 PB moves through a 1 GB ethernet in ca. 100d.



# Database and Collection Setup

- ▶ production databases also have to be replicated
  - ▶ both, data nodes and config nodes
  
- ▶ to work on sharded big data, one should
  - ▶ **not** query it from the central router nodes
    - ▶ rule of thumb: 1 PB moves through a 1 GB ethernet in ca. 100d.
  - ▶ **not** query it locally on the shard nodes
    - ▶ if data also is replicated, one would need to know which nodes to query
    - ▶ the balancer may move data during query time

# Database and Collection Setup

- ▶ production databases also have to be replicated
  - ▶ both, data nodes and config nodes
  
- ▶ to work on sharded big data, one should
  - ▶ **not** query it from the central router nodes
    - ▶ rule of thumb: 1 PB moves through a 1 GB ethernet in ca. 100d.
  - ▶ **not** query it locally on the shard nodes
    - ▶ if data also is replicated, one would need to know which nodes to query
    - ▶ the balancer may move data during query time
  
- ▶ to work on sharded big data, one should use technologies provided by the database, e.g., map-reduce.
  - ▶ for mongo: mappers and reducers have to be coded in javascript.

# Map-Reduce Example

```
1 {
2   _id: ObjectId("50a8240b927d5d8b5891743c"),
3   cust_id: "abc123",
4   ord_date: new Date("0ct_04,_2012"),
5   status: 'A',
6   price: 25,
7   items: [ { sku: "mmm", qty: 5, price: 2.5 },
8             { sku: "nnn", qty: 5, price: 2.5 } ]
9 }

1 var mapFunction1 = function() {
2     emit(this.cust_id, this.price);
3 };
4
5 var reduceFunction1 = function(keyCustId, valuesPrices) {
6     return Array.sum(valuesPrices);
7 };
8
9 db.orders.mapReduce(
10    mapFunction1,
11    reduceFunction1,
12    { out: "map_reduce_example" }
13 )
```

# Outline

1. Introduction
2. Key-Value Stores
3. Document Databases
4. Document Databases: Partitioning
- 5. Graph Databases**
6. Column Databases and Object Databases

# Graph databases

A graph database is a database that uses graph structures with

- ▶ nodes,
- ▶ edges, and
- ▶ node/edge properties

to represent and store data.

# Graph databases

- ▶ Nodes represent entities such as people, businesses, accounts, or any other item you might want to keep track of.
- ▶ Properties are relevant information that relate to nodes.
- ▶ Edges are the lines that connect nodes to nodes
- ▶ Most of the important information is often stored in the edges.
- ▶ Meaningful patterns emerge when one examines the connections and interconnections of nodes, properties, and edges.

# Graph databases

- ▶ Compared with relational databases, graph databases are often faster for associative data sets
- ▶ They map more directly to the structure of object-oriented applications.
- ▶ As they depend less on a rigid schema, they are more suitable to manage ad hoc and changing data with evolving schemas.
- ▶ Graph databases are a powerful tool for graph-like queries.

# Graph queries

- ▶ Reachability queries
- ▶ shortest path queries
- ▶ Pattern queries



# Outline

1. Introduction
2. Key-Value Stores
3. Document Databases
4. Document Databases: Partitioning
5. Graph Databases
- 6. Column Databases and Object Databases**

# Column Databases

- ▶ Column databases are a tuple (a key-value pair) consisting of three elements:
  - ▶ Unique name: Used to reference the column
  - ▶ Value: The content of the column.
  - ▶ Timestamp: The system timestamp used to determine the valid content.
- ▶ Differences to a relational database

# Example

```
{  
  street: name: "street", value: "1234 x street", timestamp: 123456789,  
  city: name: "city", value: "san francisco", timestamp: 123456789,  
  zip: name: "zip", value: "94107", timestamp: 123456789,  
}
```

# Object Database

- ▶ An object database is a database management system in which information is represented in the form of objects as used in object-oriented programming.
- ▶ Most object databases also offer some kind of query language, allowing objects to be found using a declarative programming approach (OQL)
- ▶ Access to data can be faster because joins are often not needed.
- ▶ Many object databases offer support for versioning.
- ▶ They are specially suitable in applications with complex data.

# References I



MongoDB.

**Mongodb tutorial — query documents, 2017.**

URL <https://docs.mongodb.com/manual/tutorial/query-documents/>.