Complex Systems Theory and Practice

Lecture 7

Piotr Syga

2021-11-26

Relational databases

- based on mathematical model
- reflect the relationships between data sets
- easy normalization for removing the anomalies
- consistency checks(triggers, transactions)
- popularity, multiple dialects following a common standard

Motivation for a change

- SQL databases do not allow flexible data – common schema in tables
- object-oriented concepts added artificially to relational databases
- requirement of atomic types in normalized SQL databases further hinders storing multimedia files
- SQL often duplicates some of the capabilities of programming languages
- SQL was designed with non-distributed environment in mind, having some significant problems in distributed setting

ACID – limitations

Transactions

For keeping the proper state of the database after each transactions we require that all operation are A*tomic* – cannot be performed partialy, that the data is C*onsistent* after the transaction, even if multiple datanodes exist, that all transactions are I*solated* and do not interweave and that all the changes after the transaction ended are D*urable*.

Potential problems

- scalability with respect to the size of data
- scalability with respect to number of users
- synchronization in distributed environments
- vulnerability to errors or access denial

NotOnly SQL

- designed with distributed data in mind
- capability of handling large number of 'records'
- adaptability to cope with lightweight as well as multimedia files
- no consistent schema required
- providing significant horizontal scalability (adding data nodes)

CAP

We replace the ACID requirements with ones adjusted to distributed environment:

- (C*onsistency*) after performing the operation,
- (A*vailability*) of the data without interrupts,
- (P*artition tolerance*) when data nodes cannot exchange messages in specific time periods.

In practice we fulfill two of the requirements, e.g., relational databases(CA), MongoDB/HBase(CP), CouchDB/Cassandra(AP).

NOSQL's downsides

- lack of standardization for all the languages (dialects)
- lack of a universal, declarative query language
- the capabilities of the queries is constrained, often requires additional programming language snippets
- BASE transactions

BASE

In place of a restrictive ACID model, we use BASE model, which allows some leniency:

Basically Available – we have access to the data withou *major* interrupts

Soft state – the data in a node may change even if there is no external data manipulation

Eventual consistency – means that without external data manipulation, given sufficient time, the nodes will synchronize and provide consistent data

Main types of the databases

Depending on the required functionalities and purpose of the database we may organize the data in various ways

- in dictionaries
- wide-column store
- in documents
- in graphs

Key-value store

Storing the data in a dictionary allows significant improvement of data retrieving time*, at the cost of lack of uniform datastructure and data duplication.

- Amazon Dynamo as a cornerstone
- storing data often implemented with hashtables
- fulfills Availability and Partition tolerance
- hashing allows operations on large datasets
- mimics object-oriented programming
- allows storing data of an arbitrary structure**
- strong scalability

Examples: Dynamo, Redis, Riak, Virtuoso

Columnstore

Storing data in columns allows easy aggregation for data analysis, in contrast to RDB the schema is not fixed

- data for the columns are stored in consecutive blocks in a single file
- each column is treated separately
- similar to two-way (2D) dictionary
- in a single column, there is only one datatype
- reducing time overhead on aggregation, BI analysis, data warehousing, however significantly reduced when using multi-column aggregation
- easy and efficient data compression

Examples: Bigtable, Cassandra, HBase, SimpleDB

Documentstore

Storing data in documents allows you to store incomplete or non-atomic data, at the expense of operation time and syntax differences for seemingly similar queries (cf. subdocuments, lists, arrays, sets).

- can be interpreted as a generalization of dictionary (multikey-multivalue store)
- data organized into document collections
- documents do not realize any fixed schema and allow nested (sub)documents, the field scope is (almost) unlimited
- allow connections between document collections
- one may find RDB analogies (a row–document, a column–dictionary, a table–collection,. . .)

Examples: Caché, CouchDB, Elasticsearch, MongoDB

Graphstore

Storing data in graphs allows the use of algorithms known from Graph Theory (checking the degree of a node, checking connectivity, distance and reachability, finding the shortest path between two nodes, . . .) to process the data, at the cost of time-consuming base joining and the requirement to search the entire graph to obtain an (accurate) answer.

- the base consists of ordered pairs of vertices (each pair uniquely identifies an edge)
- both vertices and edges have unique labels or IDs
- vertices represent (often complex) entities
- each vertex knows their neighbors
- one can easily identify the connection between the data
- the cost of traversing an edge (a single hop) is constant
- efficient for inter-connected (dense) data; easy indexing
- one may find RDB analogies (a row –vertex, a table–subgraph, JOIN–edges. . .)
- no universal graph query language (ongoing proposal: ISO/IEC 39075 Information Technology — Database Languages — GQL)

Examples: Amazon Neptune, Janus Graph, Neo4J, TerminusDB

Basics

- data organized in documents⇒collections
- can use indexes to speed up searches (as in RDB)
- flexible data representation
- in a single document all operations are atomic
- easy updates and replacing of documents
- server mirroring
- data sharding and live resharding
- hierarchical data segmentation – Map/Reduce
- supports bindings of many programming languages (np. pyMongo, mongocxx, mongo.jar)
- supports transactions and views
- client-side field level encryption

Current stable version: 5.0.4

BSON

All documents are stored in BSON format.

- the structure of the document is created dynamically and may be easily changed
- each document has its (eventually) unique 12B ID (4B timestamp, randomly generated 5B string, 3B counter)
- 21 datatypes (4 of them deprecated)
- default fields for UUID, MD5 and BSON client-side encryption field
- datatypes may be evoked by the name or their ID
- types might be assigned automatically, based on the value or manually (to verify the type of a field use `instanceof`, `typeof`)

Data bases

Note that if database or a collection does not exist, they are created at their first evocation

- creating or switching the database: `use newDataBase`
- using a different database (or collection from such) without switching: `db.getSiblingDB('secondDB').*` (e.g.,`db.getSiblingDB('secondDB').getCollectionNames()`)
- (in)direct creation of a collection: `db.firstCol.insertOne( { key: val } )` (or `db.createCollection(name, options)`)

The names of both databases and collections are restricted similarly to most programming languages – details. As opposed to SQL by default the names are case sensitive.

Documents

Structure

```
{ field1:  value,
field2:  value,
...
fieldN: value}
```

`_id` field is added automatically with `ObjectId` type.

Documents

Arrays and embedded documents

Fields may use non-atomic types as arrays, e.g., `color: ["white", "red"]`. You may address the field by its position `color.2` or reference all positions (e.g., at updating) with `color.[]`.
Matching: `["white", "red"]` vs `$all: ["white", "red"]`
Non-atomic fields may also be a document, e.g., `phone: {home: "XXX", cell: "YYY"}`. We reference the values hierarchically with dot notation `phone.cell`.
Naturally, one may create arrays of embedded documents and embedded documents that include arrays.

Default size limit of a single document (with embeddings) is 16MB, yet can be modified.