

Architectures for massive data management

Key-value stores and Redis

Ioana Manolescu , Silviu Maniu

INRIA Saclay & Ecole Polytechnique

Université Paris-Sud

ioana.manolescu@inria.fr , silviu.maniu@lri.fr

<http://pages.saclay.inria.fr/ioana.manolescu/>

<https://silviu.maniu.info/>

M2 Data and Knowledge, 2018/2019

Université Paris Saclay

Key-value stores


- Relatively recent class of systems, developed as part of the **NoSQL** movement
- Main idea:
 - Trade simplicity for speed and scale
- Extremely simple data model
 - **key**=short byte sequence / integer
 - **value**=byte sequence (may recognize integers)
- No QL. Operations: **PUT(k, v)** and **GET(k,v)**
- **ACID** properties depending on the system; at least atomic PUT and GET
 - Some are in-memory thus no durability at all

Key-value data models

- Simplest model:
 - One key – one value
- Extensions:
 - **Organization:** key-value pairs belong to « collections » or « databases » or « tables »
 - **Multiplicity:** set or list of values
 - **Internal structure:**
 - One key – a list of *attributes*
 - Each attribute has a *name* and a *value / set of values*

Sample key-value data model: DynamoDB

- Provided by Amazon Web Services (AWS)



- Naming may vary (there is no standard). See doc.
- Although it is called « table », *items in the same table may have nothing in common!*
- The interface is very similar to the so-called « Big Tables » (to be seen)

Redis: one of the most popular key-value stores

- **Data model:**
 - Hash (a set of key-value pairs on the same key)
 - List
 - Set
 - Values cannot be lists nor sets (no nesting!)
 - Databases
- **Operations:**
 - Put, get
 - Set operations (union, intersection)
 - List operations: left/right push/pop (→ queue / stack)
 - Arithmetic operations (attempts type conversion to integers)

Redis: lab

- Install Redis, launch it (<https://redis.io/topics/quickstart>)
- Follow the tutorial to learn the available commands (<https://try.redis.io>)
- Write in Redis a toy application of a database of books which can be borrowed in a library.
 - All books have an ISBN, a title and an author.
 - Books may also have other properties, e.g. language, publication year, edition... (up to you)
 - Make books expire after a while (if no one borrows the book)
 - If someone borrows a book (by, e.g., setting a certain field in the data) then make the book refresh its expiry date
 - Make a Redis client subscribe to books by a certain author
 - A tutorial implementing a Twitter clone can be found at <https://redis.io/topics/twitter-clone>

Redis: task

- With the help of a programming language interface to Redis (Python <https://pypi.org/project/redis/>, Java <https://github.com/xetorthio/jedis>, others at <https://redis.io/clients>), create a **publish-subscribe news system**, which:
 - **(server)** indexes a newly published item by the words in its description, publish a channel for each indexed word containing *only* the IDs of the news
 - **(client)** subscribes to news having certain keywords, retrieve the corresponding IDs and show the full text of the news
- **Individual homework:** deadline October 5th, midnight; source code and short report by mail to silviu.maniu@lri.fr