

Twitter workload for NoSQL databases

University of Minho



- Social networks applications have taken a big growth.
- MySpace, Facebook, Twitter, Hi5, Orkut, Bebo, LinkedIn, PatientsLikeMe, Yahoo!360.
- Are in top of the sites with more traffic and have millions of users worldwide.



- User's actions and preferences may affect many users in their network.
- Pose new challenges to current database servers.
- Use of centralised RDBMS or even a replicated DBMS is a major bottleneck.
- Social applications are thus exploiting NoSQL databases.



- No benchmarks mimicking the workload of a social network.
- Existent NoSQL benchmarks are naive.
- Standard benchmarks (like TPC-C,...) not suited for large scale storage system.



- Create a benchmark based on a twitter alike application:
 - Measure performance (throughput, latency, ...).
 - Behaviour of databases in face of faults and scalability
 - Workload to simulate as close to real the use of the application.
- Event based API that allows to evaluate real and simulated NoSQL databases.
- Architecture compatible with Cloud Environment.



Workload entities

Users

key	value
userID	name, password, creation date, followers, following and lastTweetID
...	...

...	...
-----	-----

Tweets

key	value
userID-tweetID	tweet
...	...

...	...
-----	-----

FriendsTimeLine

key	value
userID	List<date:tweetID>
...	...

...	...
-----	-----

Tags

key	value
tag	List<tweets>
...	...

...	...
-----	-----



- Social graph
 - small-world network (due to high clustering and small diameter).
 - power-law distribution (few nodes have high degree, while the majority of nodes have small degree).
 - scale-free.
- Initial tweets per user.



- statuses_user_timeline (userID)
- statuses_friends_timeline (userID)
- statuses_mentions (userID)
- search_contains_hashtag
- statuses_update (userID)
- friendships_destroy (userID)
- friendships_create (userID)



● statuses_update

- Find the next tweet ID for the user
- Add the tweet to Tweets entity
- for each user's follower update its timeline
- update user's timeline.

● friendships_create and friendships_destroy

- update the list of followers for the user
- update the list of people following the new or old followed user
- recompute the user's timeline.



- The interleaving of operations take into account previous studies and discussions that took place during Twitter's Chirp conference (the Twitter official developers conference).
- Defines a think-time between operations.
- Next operation is randomly chosen with following probabilities per operation:



Ops Probabilities

Operation	Probability
search_contains_hashtag	15%
statuses_mentions	25%
statuses_friends_timeline	50%
statuses_update	5%
friendships_create	2.5%
friendships_destroy	2.5%

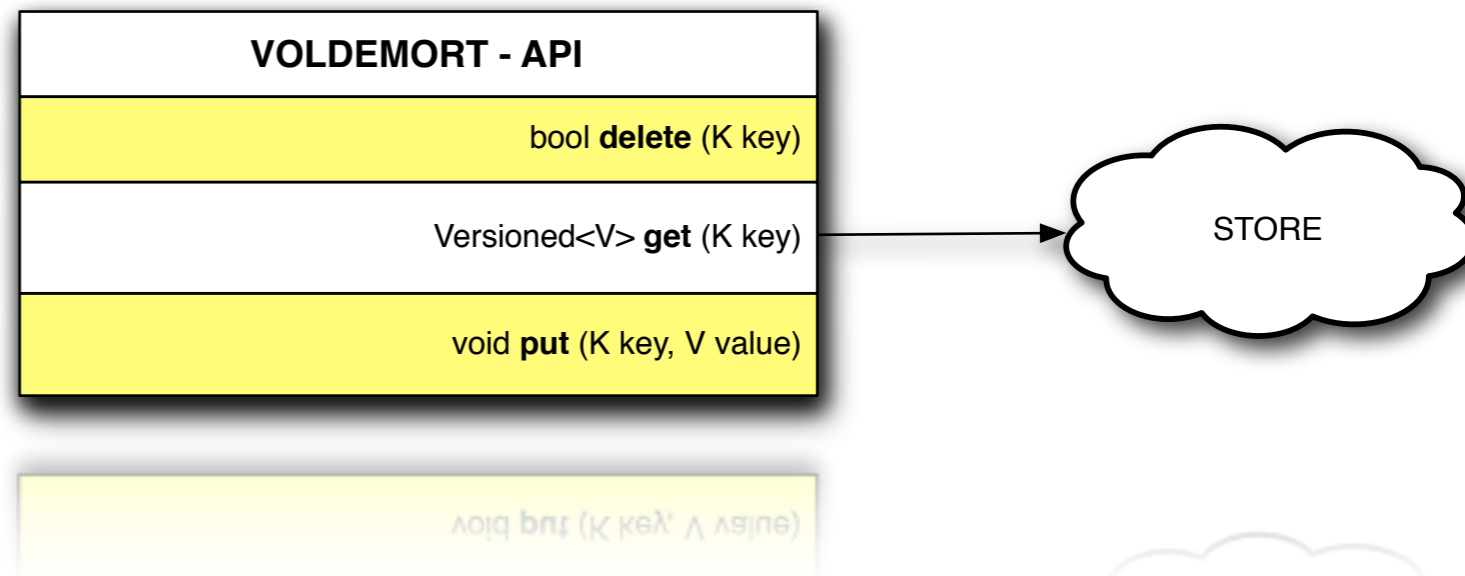


- Workload already implemented for:
 - Voldemort
 - Cassandra
 - Clouder
 - MySQL




Why Voldemort?

- Voldemort represents a family of row-based stores with a simpler data model and API with only puts and gets (e.g. Amazon's Dynamo).
- Simple mapping of key to value.
- Values are treated as opaque array of bytes.



Why Cassandra?

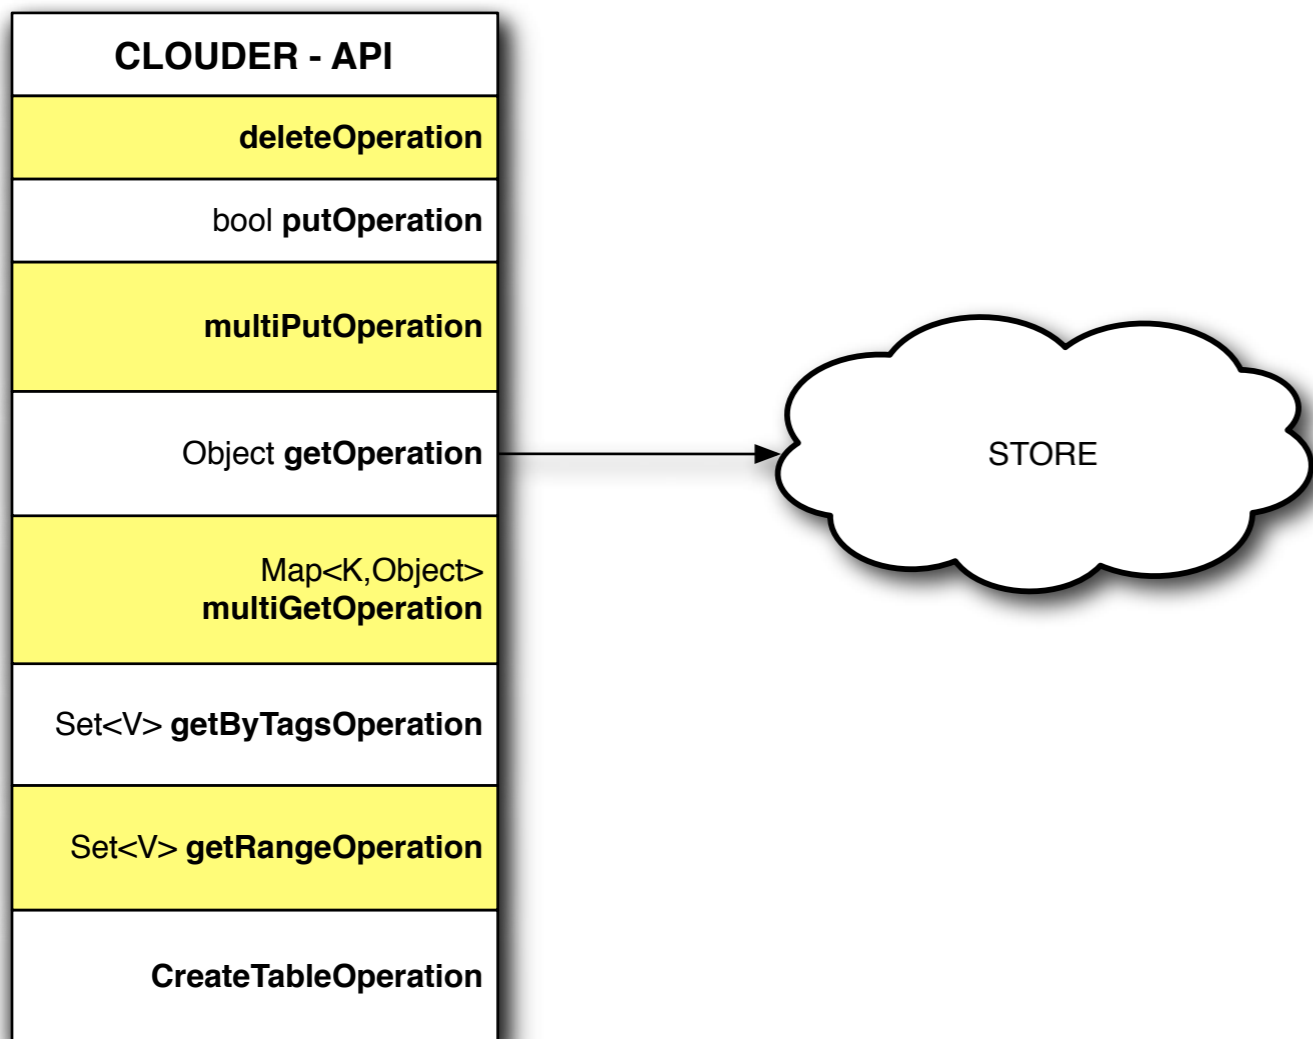
CASSANDRA - API
<code>remove(keyspace, key, column_path, timestamp, consistency_level)</code>
<code>insert(keyspace, key, column_path, value, timestamp, consistency_level)</code>
<code>ColumnOrSuperColumn get(keyspace, key, column_path, consistency_level)</code>
<code>list<ColumnOrSuperColumn> get_slice(keyspace, key, column_parent, predicate, consistency_level)</code>
<code>map<string,ColumnOrSuperColumn> multiget(keyspace, keys, column_path, consistency_level)</code>
<code>map<string,list<ColumnOrSuperColumn>> multiget_slice(keyspace, keys, column_parent, predicate, consistency_level)</code>
<code>list<KeySlice> get_range_slice(keyspace, column_parent, predicate, start_key, finish_key, row_count=100, consistency_level)</code>
<code>i32 get_count(keyspace, key, column_parent, consistency_level)</code>



- Cassandra offers a different data model: it is column oriented (ColumnFamilies, SuperColumns, etc much like Google's BigTable) .
- Offers a higher level API with range operations.



Why Clouder?



- Clouder offers a API with puts, gets as well as search and multi-tuple operations.
- Extends the data model of previous tuple stores with tags, that allows to establish arbitrary relations among tuples.
- Takes advantage of tuple correlation in terms of operations and how partitioning is made.



Why MySQL?

- Provides a baseline to compare with NoSQL databases.
- At the same time, will assess the suitability of RDBMS to today's Social Applications.



- A realistic workload for simulating today's high demanding social applications.
- Easy to adapt to the available databases.
- Compare the different databases, in terms of performance, scalability and fault tolerance.
- Add another implementation: VoltDB



Thank you!

