

1. Summary

In this project, I design a Load Balancer for Key-value Database (LBKVD for short). LBKVD is written in java, and uses network to communicate with database instance through database defined protocol. The main goal is to scale database service at least linearly as more database instance as added.

2. Why does LBKVD matters

Think about thi senario. Assume you develop a social website for pets, especially for cats and dogs. Because you don't have money, you decided to use a light-weight open source database in backend. At the beginning, you website is not well-known, only pets of your friends use you service, so pressure on database is not that critical. However, you website becomes popular and users grow exponentially.

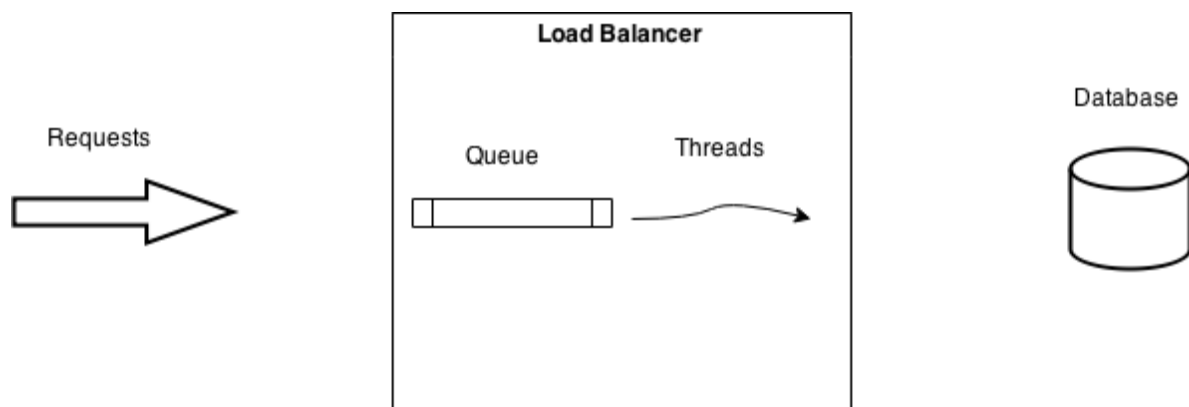
At first, you are very happy because you will become next zark muckerberg. Then you realize a problem: you need to scale your database to maintain your website, before you can get money from investors, and become a billionaire. For now, you are poor, so still need to use open-source database that used in the past. In the case, how can you scale your website?

Then answer is: Load Balancer for Key-Value Database!

3. Experiment

3.1 Description

The basic structure of my LBKVD is fixed, no matter what kind of schedule policy I use. Please see below:



So basicly, LBKVD creates queues to contain coming requets, and also creates worker threads to fetch requests from queues and send to databases. Worker threads are also responsible for responding requests.

Therefore, based on structure of LBKVD, load balancer schedule policy is actually a strategy to decide how to utilize queues and worker threads, to assign jobs to database instances.

3.2 Trace file, software and hardware

The trace file I created contains 300,000, 4KB insert operations with unique keys.

The cluster that I used for testing, has four nodes. Each node has two, quad-core Xeon E5345 processors (2.33GHz, 8M L2 cache, no hyper-threading), 16GB memory and hard disk(15000RPM).

The Key-Value database, is actually a document database, named EmeraldDB.

Following is my policy design.

3.3 Navie assignment

There is only one queue in LBKVD, and one worker thread for each database node. LBKVD assign requests to each worker thread one by one. So every database node will be assigned equal amount of work.

The problem of this assignment, is that it is so simple. There is no multi-threading to hidden latency, and there is no flexible schedule on work assignment. What if a node is idle while another one is busy? So I actually never seriously implement this policy.

3.4 Equal assignment

There is one queue and one worker thread for each database node. Work assignment is dynamic, i.e. LBKVD assign a job based on current workload of each node. LBKVD can monitor length of queue of each node, and pick up the one with shortest queue depth.

	Baseline	1 Node	2 Node	3 Node	4 Node
Speed up (second)	1x (149s)	0.88x (169s)	1.86x (80s)	2.64x (56s)	3.35x (44s)
Throughput (ops / sec)	2063	1827	3925	5700	7111

Note: Baseline here is trace file tested on one database instance.

I got pretty good result here, except for Equal assignment on 1 Node. The reason is clear, there are queues in equal assignment, but not exist in baseline. Packet requests in queue. fetch and unpacking costs a lot. Things get better when the number of nodes increase.

However, this policy still cannot do the best. For each thread, it read data from disk, put it into memory buffer, and then put to network buffer, and send out, there exists lots of

chance that worker thread is waiting and do nothing, while lots of requests are in the queue. So multi-threading is possible to improve throughput.

3.5 Equal assignment with multi worker threads

There is one queue but multi worker threads for each database node. Work assignment is still the same as Equal assignment. Each database node establish multi connection with LBKVD, and receive multi requests at the same time. There are locks on queues, in order to provide concurrency control under multi-threading situation.

However, how many threads is better? I conducted experiments and get following result:

Time(s)	1 Node	2 Node	3 Node	4 Node
1 thread / Node	1x (169s)	2.11x (80s)	3.01x (56s)	3.84x (44s)
2 threads / Node	1.97x (86s)	3.76x (45s)	4.97x (34s)	6.76x (25s)
3 threads / Node	1.31x (129s)	4.97x (34s)	6.76x (25s)	9.94x (17s)
6 threads / Node	5.12x (33s)	8.45x (20s)	8.45x (20s)	8.04x (21s)

Note: 1 thread/DB is the same as Equal assignment.

Throught(op/sec)	1 Node	2 Node	3 Node	4 Node
1 thread / Node	1x (1827)	2.14x (3925)	3.12x (5700)	3.89x (7111)
2 threads / Node	1.99x (3640)	3.86x (7049)	5.12x (9347)	6.8x (12422)
3 threads / Node	2.74x (5011)	4.96x (9062)	6.81x (12445)	9.34x (17065)
6 threads / Node	4.93x (9018)	7.98x (14588)	8.1x (14717)	7.92x (14476)

Note: 1 thread/DB is the same as Equal assignment.

When create multi-threads for a database instance, things get better compared to one worker thread. Data in the table is using 3 worker threads and one queue for each node. I tried create more threads for each node, however, they cannot beat 3 threads per node. I guess the reason is that, the Load Balancer I ran on the machine that only has two, quad-core Xeon E5345 processors. Xeon E5345 processor doesn't have hype-threading, so the maximum threads it can execute, is 4. CPU cannot afford too many threads, and it seems like too many threads acutally reduce the throughput of Load Balancer. 4 nodes with 3 threads each in Load Balancer has already achieve peek rate.

Still, this assignment has a problem, it releases consistency. Assume here is two operations coming one by one: "INSERT a", "QUERY a". These two operations has the same key, so they are assigned to same node. Because there is a queue and

multi-threading, these two operations may be sent to node at the same time. It is very likely that “QUERY a” arrives before “INSERT a”, then query operation returns nothing.

In order to fix this **weak consistency problem**, finally I come up with a multi threads with multi queues approach.

3.6 Equal assignment with multi worker threads and multi queues

There are multi queues and multi workder threads for each database node. In this policy, because of increase number of queues and worker threads, maybe thirty or more, it is hard to monitor status of each queue and worker threads, so I use **consistenct hashing** here, in order to have a relatively balanced workload assignment under complexity situation. queues for database nodes are mapped to a circle, as well as keys of records.

This approach guarantees that operation serialization. Operations with same keys always sent to the same, only queue. So operations are serialized based on the time they arrive at Load Balancer, while throughput still holds as the same as last policy.

4. In the end...

4.1 Best implementation

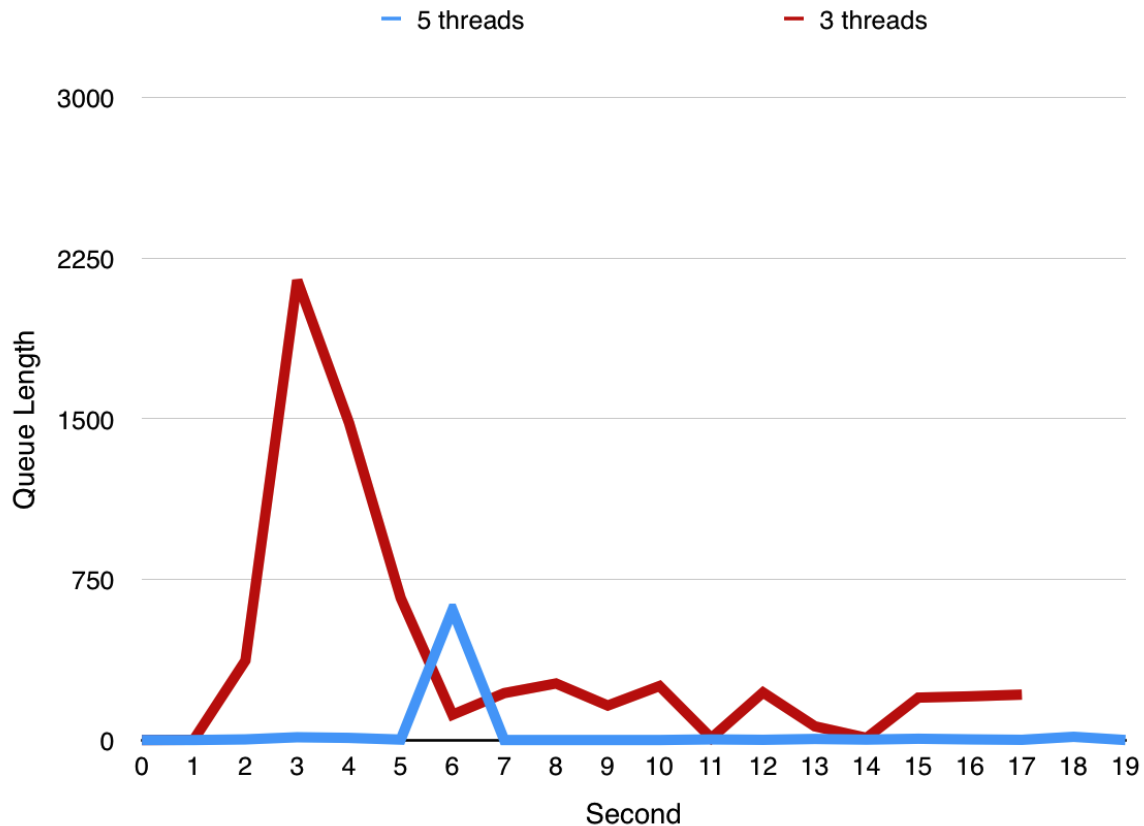
	Trace file on one database instance	4 database instances, each one corresponding to a queue and 3 worker threads
Speed up	1x (149s)	8.52x (17s)
throughput (ops / sec)	2063	17065

Note: Baseline here is trace file tested on one database instance.

Here is best implementation vs baseline, but I believe this should not be the end.....

4.2 Analysis

So, for now, all the effort I made to improve throughput, is based on a single strategy, increasing the number of queues and the number of threads. However, it finally comes to the point, at which throughtput cannot improve anymore. So, what is the bottleneck of this Load Balacner? Here is the queue length I observed when it sets as single queue and multi-threading.



So, from this graph, I can see that more threads efficiently decrease the queue length, that is, reduce the queue time. However, reducing queue time doesn't help improve throughput and reduce execution time. Once I set more than 3 threads for each database instance, job running time becomes slower than 3 threads per instance.

Two reasons possible here. First reason, is that network is the bottleneck. Since more threads reduce the queue length, but cannot accelerate the running time, so it is possible some threads are waiting for sending requests to databases

Second reason, is database it self becomes the bottleneck, because of long time spent on disk I/O.

I have no idea which one is the real reason, because I don't have data. In the next step of this project, I need to collect more data to analyse what is the real bottleneck. One type of data I need to collect is network throughput. Another type of data is database IO.

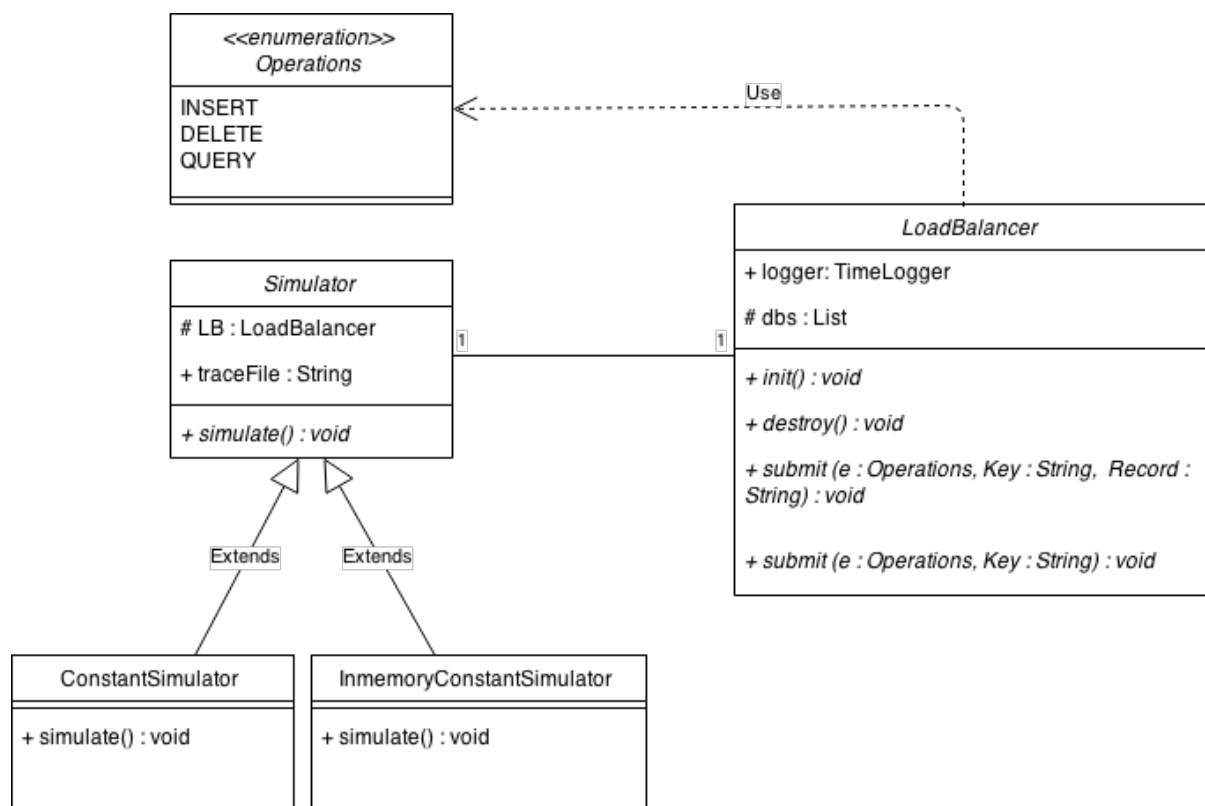
5. Retrospective

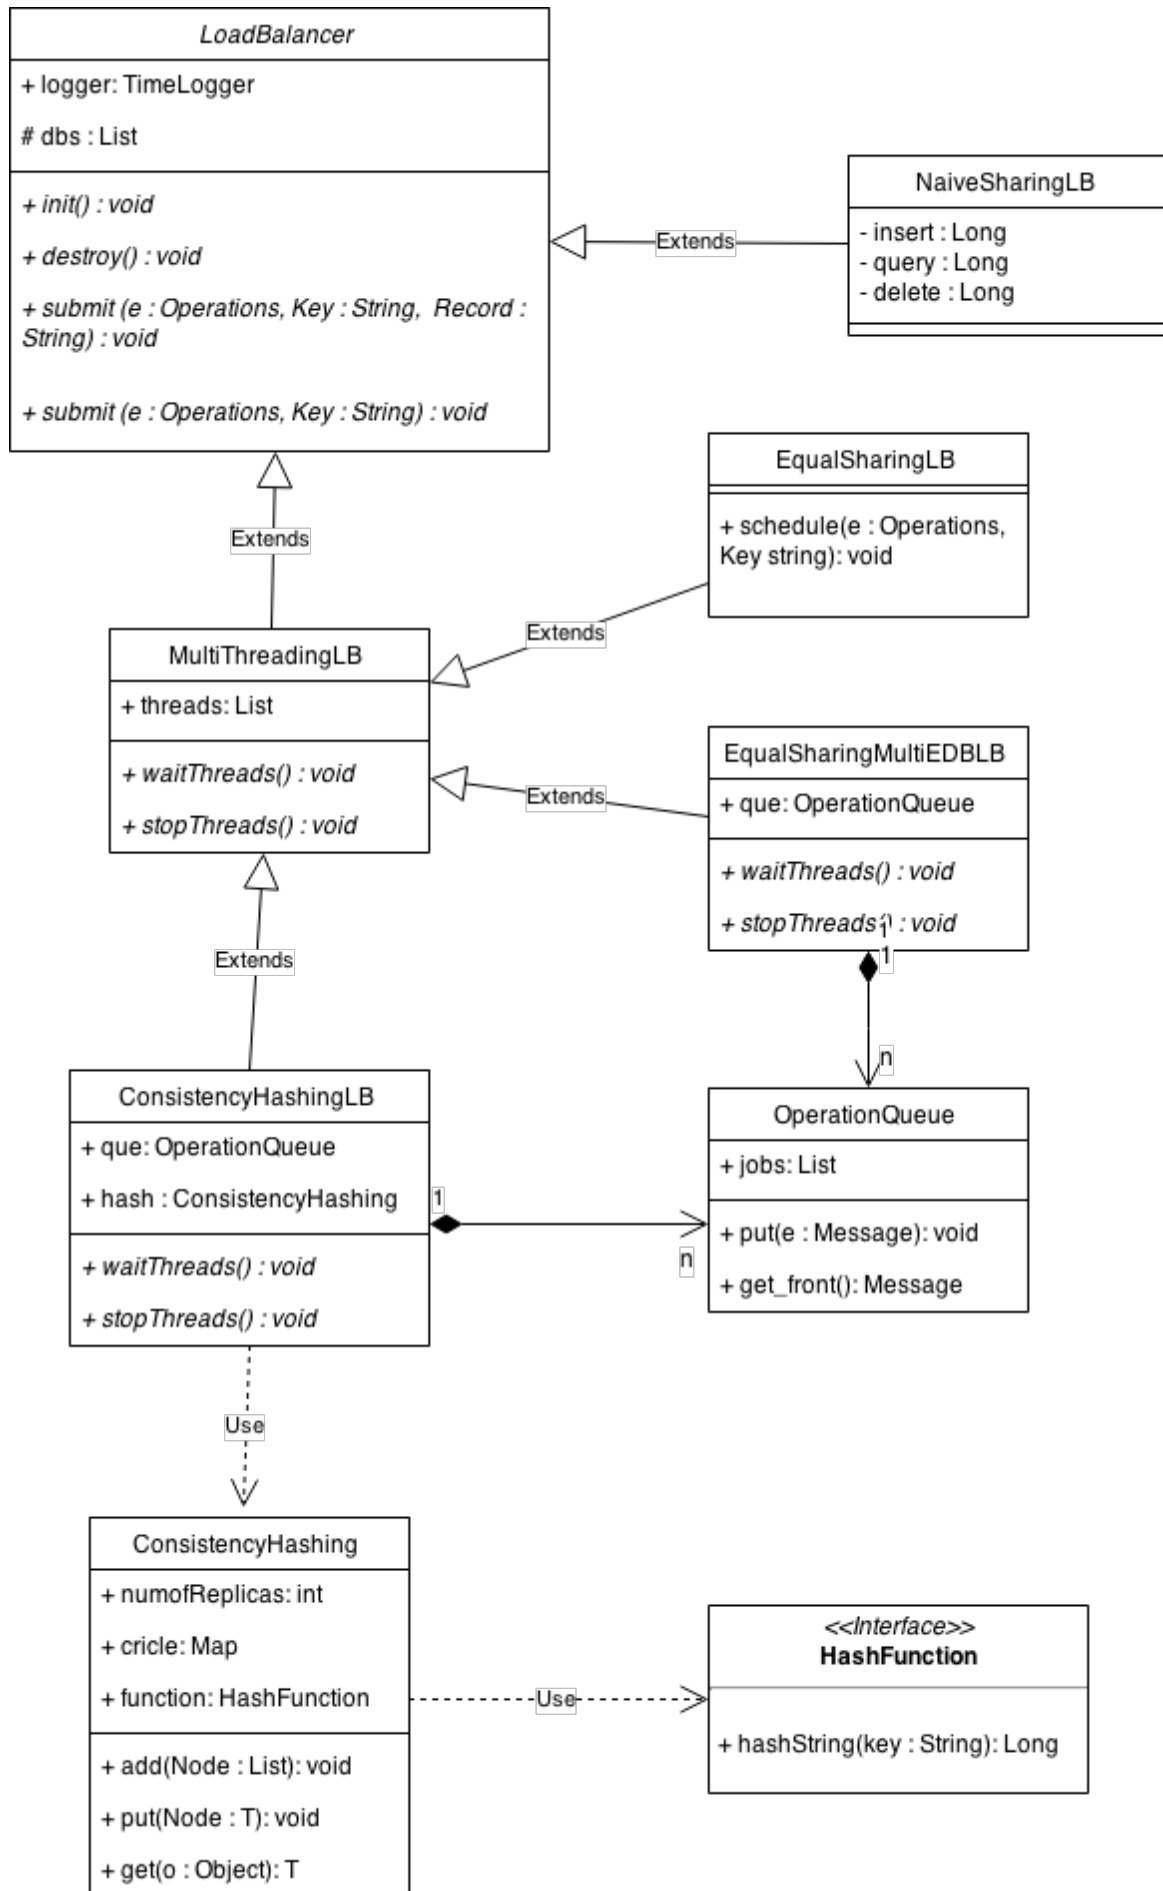
Although I have achieved big breakthrough by incrementally improving design and implementation of LBKVD, I still have ideas that worth to try, which I think can keep improving performance of LBKVD. Here is some ideas.

a. It turns out best number of worker threads is 12 on my testing machine. Increasing this number will not bring any benefit. So this is one of the potential bottlenecks of LBKVD. So, why I need put all the logic of LBKVD on one machine? Why don't I use client-server model for LBKVD? So new design of LBKVD sounds like this: server of LBKVD is in charge of receiving requests. In the meantime, each node will not only install database instance, but also install client of LBKVD. Server can send requests to clients, and each client can launch 12 worker threads and handle requests. Client, instead of server, communicates with database instance.

b. One of the problem of idea a, is that server may become new bottleneck, e.g. network becomes bottleneck. So new idea is peer to peer system. Client can talk to each other, and reduce pressure on server. For example, server may avoid complicated scheduling policy, instead, server sends request to a random client. Client can decide the right node, to which this request should be sent.

6.UML





7. Reference

- [1] Intel® Xeon® Processor E5345 (8M Cache, 2.33 GHz, 1333 MHz FSB), [link](#).
- [2] Document of EmeraldDB, a light-weight NoSQL Database.
- [3] Wikipedia: Consistent hashing, [link](#).