Session Guarantees with Raft and Hybrid Logical Clocks

Mohammad Roohitavaf*
Michigan State University
East Lansing, MI, USA
roohitav@cse.msu.edu

Jung-Sang Ahn
eBay Inc.
San Jose, CA, USA
junahn@ebay.com

Woon-Hak Kang
eBay Inc.
San Jose, CA, USA
wokang@ebay.com

Kun Ren
eBay Inc.
San Jose, CA, USA
kuren@ebay.com

Gene Zhang
eBay Inc.
San Jose, CA, USA
genzhang@ebay.com

Sami Ben-Romdhane
eBay Inc.
San Jose, CA, USA

Sandeep S. Kulkarni
Michigan State University
East Lansing, MI, USA
sandeep@cse.msu.edu

ABSTRACT

Eventual consistency is a popular consistency model for geo-replicated data stores. Although eventual consistency provides high performance and availability, it can cause anomalies that make programming complex for application developers. Session guarantees can remove some of these anomalies while causing much lower overhead compared with stronger consistency models. In this paper, we provide a protocol for providing session guarantees for NuKV, a key-value store developed for services with very high availability and performance requirements at eBay. NuKV relies on the Raft protocol for replication inside datacenters, and uses eventual consistency for replication among datacenters. We provide modified versions of conventional session guarantees to avoid the problem of slowdown cascades in systems with large numbers of partitions. We also use Hybrid Logical Clocks to eliminate the need for delaying write operations to satisfy session guarantees. Our experiments show that our protocol provides session guarantees with a negligible overhead when compared with eventual consistency.

CCS CONCEPTS

• **Computing methodologies** → **Distributed algorithms**; • **Information systems** → **Cloud based storage**; **Distributed storage**;

KEYWORDS

Distributed key-value stores, Session Guarantees, Hybrid Logical Clocks, Consistency, Replication, Raft.

* Also with eBay Inc..

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICDCN '19, January 4–7, 2019, Navi Mumbai, India

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6094-4/19/01...\$15.00

<https://doi.org/10.1145/3288599.3288619>

ACM Reference Format:

Mohammad Roohitavaf, Jung-Sang Ahn, Woon-Hak Kang, Kun Ren, Gene Zhang, Sami Ben-Romdhane, and Sandeep S. Kulkarni. 2019. Session Guarantees with Raft and Hybrid Logical Clocks. In *International Conference on Distributed Computing and Networking (ICDCN '19)*, January 4–7, 2019, Navi Mumbai, India. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3288599.3288619>

1 INTRODUCTION

Due to its performance and availability benefits, *eventual consistency* has become a popular consistency model for geo-replicated data stores. The eventual consistency is a weak consistency model that guarantees, in the absence of new writes, all connected replicas will eventually converge to the same state. This weak consistency allows replicas to communicate asynchronously, and allow clients to perform their operations without waiting for the long network latencies among replicas. Despite these benefits, eventual consistency makes the application development complex for the programmers. For instance, a client may not be able to read the value that it has just written, because its write and read operations have been directed to different replicas with different states.

To simplify the design of applications, stronger levels of consistency – causal, sequential etc. – have been proposed. Specifically, causal systems track the causal dependencies of the versions and do not make a version visible in a replica, if some of the causal dependencies of the version are not visible in the replica. We can track causal dependencies explicitly as a list of versions such as in [17], or track them implicitly using a scalar [13], a vector [24], or a matrix [12] timestamps. Although causal consistency provides useful guarantees, it can cause significant performance overhead, especially for partitioned systems. This overhead is not acceptable for certain systems [4].

Session guarantees are consistency models that are weaker than causal consistency but stronger than eventual consistency. Session guarantees include properties such as monotonic-read, monotonic-write, read-your-own-write, and write-follows-reads. To illustrate some of these guarantees, consider a scenario where we have one key k with initial value 0. Client $c1$ writes 1 to k . Without monotonic-read, it is possible that a client $c2$ first reads the value 1 and then the

value 0. Next, suppose that client c_2 writes the value 2 (at another replica). Without monotonic-write, it is possible that the final value of k is 1.

In this paper, we evaluate the cost of providing session guarantees for our partitioned and replicated system that relies on the Raft algorithm [21] for replication inside datacenters and eventual consistency across datacenters. Although Raft guarantees that all replicas inside one datacenter apply writes with the exact same order, it does not guarantee that all replicas provide the same state at all time. For example, it is possible that a client writes on the leader, but does not find its update on one of the replicas that has not committed the update yet. Using session guarantees, we can prevent such anomalies. We use Hybrid Logical Clocks (HLCs) to provide wait-free write operations while satisfying monotonic-writes and write-follows-reads guarantees.

Using partition (aka. sharing) the data, we can increase the scalability of the system, as the size of the data that we can store is not limited by the capacity of a single machine. However, one problem of achieving consistency for partitioned systems at scale is the problem of *slowdown cascades*. Specifically, in partitioned systems delaying the visibility of updates to make sure other partitions are updated enough can create a cascade of slowdowns caused by an only one failed/slow machine. We define modified versions of conventional session guarantees such as those considered in [28]. Our modified versions allow us to separately keep track of session guarantees for each partition. This removes the need for cross-partition communication which in turn eliminates the problem of slowdown cascade. Moreover, using our protocol, session guarantees can be requested by the client on a per-request basis. This allows clients to achieve their desired levels of consistency only when they need them, and enjoy the performance benefit of weaker consistency models when stronger consistency is not necessary for a specific operation.

The contributions of this paper are as follows:

- We present the design of our key-value store, NuKV. This key-value store is designed for services with very high availability and performance requirements at eBay.
- We evaluate the cost of providing session guarantees in our system that uses Raft protocol for intra-datacenter replication, and eventual consistency for inter-datacenter replication. A client can flexibly access any datacenter it desires. If the client dynamically changes the datacenter, it will receive eventual consistency. However, sequential consistency is provided within a datacenter. Our goal is to provide session guarantees even if clients access data from different datacenters. Hence, we identify the cost of providing such session guarantees.
- We provide an algorithm based on HLCs that eliminates the need for delaying write operations to satisfy monotonic-write and write-follows-read guarantees.
- We provide modified versions of session guarantees that do not cause slowdown cascade for systems with large numbers of partitions.

The rest of paper is organized as follows: In Section 2, we provide a background on log replication via Raft and HLCs. In Section 3, we provide definitions for session guarantees. In Section 4, we provide

the architecture and assumptions of our system. Section 5 provides the protocol. The experimental results are provided in Section 6. Section 7 reviews the related work. Finally, Section 8 concludes the paper.

2 BACKGROUND

In this section, we provide a background on log replication and HLCs. We also explain why we use HLCs [15] instead of conventional logical clocks [16] or physical clocks.

2.1 Log Replication and Raft

While one of the purposes of creating replicas in different datacenters is to reduce the network latency for clients, replicas inside one datacenter are mainly created for fault-tolerance. Specifically, we want updates to be durable and clients to be able to use the data, even if some machine inside a datacenter fails. Since network partitions do not occur inside one datacenter and the network latency is very small, providing high levels of consistency for replicas inside one datacenter is feasible. We use log replication via Raft [21] protocol to make sure all updates are applied with the same order in replicas inside one datacenter. Specifically, replicas maintain a replicated log consisting of the sequence of operations to execute. Each entry on the log is associated with an index which shows its location on the log. Raft [21] guarantees that all replicas find the same operation for any index on their logs.

Raft relies on the leader election for providing log replication. After electing a leader, the leader has the complete responsibility for managing the log. The leader accepts update requests from the clients, adds entries on its log, and forces the log of other servers to follow its log. The leader may crash which cause a new round of leader election. Raft guarantees that the new leader always has the most recent version of the log.

2.2 Hybrid Logical Clocks

Hybrid Logical Clocks (HLCs) [15] allows us to capture *happens-before* relation [16] while assigning timestamps that are very close to the physical clocks. The HLC timestamp of an event e , is a tuple $\langle l.e, c.e \rangle$. $l.e$ is our best approximation of the physical time when e occurs. $c.e$ is a bounded counter that is used to capture causality whenever $l.e$ is not enough to capture causality. Specifically, if we have two events e and f such that e happens-before f and $l.e = l.f$, to capture causality between e and f , we set $c.e$ to a value higher than $c.f$.

Why using HLCs? We want to timestamp versions such that two following requirements are satisfied:

- timestamps are close (within clock drift error) to the physical time when the event of writing occurs, and
- timestamps capture happens-before relation

Logical clocks fail to satisfy the first requirements. Specifically, logical timestamps have no relation to the wall clock time. This causes several problems. First, we cannot use timestamps to provide clients with the version of a data object at a given physical time. More importantly, we cannot resolve conflicting writes based on timestamp. For instance, assume a client writes version v for a data object on replica A . One hour later, another client writes version v' for the same data object on replica B . With logical timestamps,

it is possible that the timestamp assigned to v' is smaller than the version assigned to version v . In this situation, if we use timestamps to select the winner version (i.e. the version with higher timestamp is the winner), v will be selected as the winner which is not desired, as we want the version written one hour later to be the winner.

On the other hand, physical clocks fail to satisfy the second requirement. This is due to the fact that perfect clock synchronization is impossible. Thus, the clock skew between different replicas prevents physical clocks from accurately capturing happens-before relation. For instance, imagine a client writes version v for a data object on replica A . Next, the client writes version v' on replica B . While using logical clocks we can guarantee that the timestamp assigned to v' is higher than that of assigned to v , using physical clocks it is not guaranteed. Specifically, if the clock of replica B is behind A , the timestamp assigned to v' may be smaller than that of v .

HLCs solves both issues explained above, i.e., it allows us to capture the happens-before relation, and at the same time, allows us to provide conflict resolution with respect to the physical time.

3 SESSION GUARANTEES

In this section, we define various session guarantees for key-value stores. We provide modified versions of session guarantees considered in [28]. This modified versions define session guarantees per keys, which allows us to avoid the problem of slowdown cascade [4] by avoiding cross-partition communications. Specifically, to satisfy definitions provided here, different partitions can check session guarantees independently. Thus, failure/slowdown of a partition does not affect the visibility of updates in other partitions.

First, regarding writes and reads of a specific key we consider following definitions:

Definition 3.1. $CommittedWrites(s, k, t)$ is the set of all writes committed for key k at server s at time t .

Definition 3.2. $ClientWrites(c, k, t)$ is the set of all writes done by client c for key k at time t .

Definition 3.3. $ClientsReads(c, k, t)$ is the set of all reads that have written a value for key k read by client c at time t .

Now, using above definitions, we define various per-key session guarantees as follows:

Definition 3.4 (Per-key Monotonic-Read Consistency). Let O be an operation issued by client c on server s at time t that reads the value of key k . Operation O satisfies monotonic-read consistency if any $w \in ClientReads(c, k, t)$ is included in $CommittedWrites(s, k, t')$ where t' is the time when server s actually reads the value for key k .

Monotonic-read consistency is important for the clients, as the lack of monotonic-read causes user confusion by going backward in time. For example, consider a webmail application. Monotonic-read consistency guarantees that the user is always able to see all emails that has seen before. New emails may be added to the mailbox in the next time that the user checks their mailbox, but no email can disappear once the user saw it, unless the user wants to delete it. This is not guaranteed with only eventual consistency.

Monotonic-read consistency guarantees that the client never misses what it has seen so far. On the other hand, read-your-write consistency guarantees that the client never misses what it has written so far. Specifically,

Definition 3.5 (Per-key Read-your-write Consistency). Let O be an operation issued by client c on server s at time t that reads the value of key k . Operation O satisfies read-your-write consistency if any $w \in ClientWrites(c, k, t)$ is included in $CommittedWrites(s, k, t')$ where t' is the time when server s actually reads the value for key k .

For instance, users must be able to see the posts that they posted on their social network page. This seems to be obvious for a centralized system, but unfortunately, with eventual consistency, it may be violated if the client is routed to another replica when it wants to read the data.

Definition 3.6 (Per-key Monotonic-write Consistency). Let O be an operation issued by client c on server s at time t that writes a value for key k . Operation O satisfies monotonic-write consistency if for any server s at any time t' if $CommittedWrites(s, k, t')$ includes O , no client accessing s reads a value written by write $w \neq O$ included in $ClientWrites(c, k, t)$.

Suppose a user updates his/her password two times. Per-key monotonic-write consistency together with eventual consistency guarantees that eventually, the user is able to login to the system with the latest passwords in all replicas. Note that just eventual consistency does not guarantee the second version to be the winner version.

Definition 3.7 (Per-key Write-follows-reads Consistency). Let O be an operation issued by client c on server s at time t that writes a value for key k . Operation O satisfies write-follows-read consistency if for any server s at any time t' if $CommittedWrites(s, k, t')$ includes O , no client accessing s reads a value written by write $w \neq O$ included in $ClientReads(c, k, t)$.

As an example for the necessity of per-key write-follows-reads consistency, consider a shared document in a cloud-based document processing service. Suppose one of the contributors reads the current version of a document and then appends a line to the document. To append the line, the application reads the current value of the data object associated with the content of the document, appends the line, and writes the new content of the document back to the system. The per-key write-follows-reads consistency guarantees that in all replicas, the version of the document with the appended line is the winner version.

With Definitions 3.6 and 3.7, we can trivially satisfy monotonic-write and write-follows-reads consistency by never returning any committed version. To avoid such a trivial implementation, we assume following requirements are implicitly required for all definitions:

- R1 Any write for a key committed on a server will be eventually committed on all servers hosting that key.
- R2 Clients reads only committed versions.

4 SYSTEM ARCHITECTURE

The data of a typical NuKV deployment is fully replicated on D datacenters. Inside each datacenter, the data is hosted by P partitions. For each partition, we create several replicas inside a datacenter. Thus, we have two levels of replication; at one level, we create replicas of the entire data in several datacenters, at another level, we create replicas of partitions inside a datacenter. The replicas of a partition inside a datacenter form a Raft group [21]. Each group has a leader and other replicas are followers. For each partition, L_d denotes the leader at the datacenter d . F_d^i denotes the i th follower of L_d (see Figure 1).

NuKV follows a multi-leader replication, i.e. clients can write to the leader of any datacenter. Followers learn the new updates via the Raft algorithm [21]. An update on a datacenter will be asynchronously replicated to the leaders of other datacenters by a special follower of the Raft group denoted by XC. XC servers are stateless. Thus, in case of failure, they can easily recover and resume sending replicate messages to other datacenters. Data nodes, on the other hand, detect repeated writes received from XC servers, and avoid applying them. We assume FIFO channels between datacenters, i.e., the XC sends updates to the leaders of other datacenters with the same order it reads them from its Raft log. Clients usually perform their operations on their local datacenter. However, it is possible that they access other datacenters too. For read operations, clients may use any replica inside a datacenter, but writes are always on the leaders.

Figure 1: An example of a NuKV deployment with three datacenters each with R . Note that for each partition, we have the architecture shown here.

5 PROTOCOL

In this section, we provide a protocol for providing different levels of consistency defined in Section 3. The read part of the protocol is basically an adoption of the protocol provided in [28] for our architecture with Raft replication and partitioning explained in Section 4. Two basic operations for a key-value store are $PUT(k, v)$ which writes value v for key k and $GET(k)$ which reads the value of key k . Our algorithm may block a GET operation, when the server cannot meet the session guarantee request by the client. For the PUT operations, instead of blocking, we use timestamping with HLC [15] to provide session guarantees.

5.1 Client-side

Each client c maintains two $D \times P$ matrices: 1) highest read matrix (hrm), and 2) highest write matrix (hwm). $hrm[d, p]$ is the highest Raft log index of the versions written in partition p of datacenter d read by client c . $hwm[d, p]$ is the highest Raft log index of the versions written by client c at partition p of datacenter d . Note that hrm and hwm are only maintained by the client. They are not sent over the network or stored with keys. Thus, their overhead is negligible. The client also maintains two scalars dt_r and dt_w that maintains the highest HLC timestamps of versions read and written by the client, respectively.

Algorithm 1 shows the client-side of our protocol. The client updates its hrm , hwm , dt_r , and dt_w as it reads and writes the data according to the definitions provided above. When the client wants to read a key on partition p , it includes two vectors hrv and hwv with its GET request. The server-side uses these vectors to provide various session guarantees for read operations. Let O be the vector of size D with all entries equal to zero. When the client sends a GET request to partition p , it can send O or the row corresponding to partition p in its hrm , i.e. $hrm[:, p]$ as hrv . Similarly, it can send O or the row corresponding to partition p in its hwm i.e. $hwm[:, p]$ as hwv . By choosing vectors for hrv and hwv , the client can require different session guarantees for its read operations in a per-operation basis as it is shown in Algorithm 1. In this algorithm, monotonic-read-your-write means a session guarantee that requires both monotonic-read and read-your-write guarantees.

When the client wants to write a key on partition p , it includes an integer called dependency time denoted by dt in its PUT request. The server-side uses this value to provide various session guarantees for write operations. The client can control its desired session guarantee for a write operation as it is shown in Algorithm 1 by choosing values for dt . Choosing 0 results in just eventual consistency, dt_w results in monotonic-write, dt_r results in write-follows-reads, and $\max(dt_r, dt_w)$ results in both monotonic-write and write-follows-reads that we refer to by monotonic-write-follows-reads.

5.2 Server-side

Each server maintains a stable vector (sv) with size D . $sv[d]$ in server s is the highest log index of the versions written in datacenter d committed in server s . For each server, PC_d shows the value of the physical clock, and HLC_d shows the value of the hybrid logical clock in L_d .

Algorithm 2 shows the server-side of our protocol. In the GET operation handler, the server blocks if for some i , $sv[i]$ is smaller

Algorithm 1 Client-side

```

1: GET (key  $k$ , ReadConsistencyLevel  $l$ )
2:    $p$  = partition id of  $k$ 
3:   if ( $l$  = EVENTUAL)
4:      $hrv$  =  $hvw$  =  $O$ 
5:   else if ( $l$  = MONOTONIC-READ)
6:      $hvw$  =  $O$ 
7:      $hrv$  =  $hrm[:, p]$ 
8:   else if ( $l$  = READ-YOUR-WRITE)
9:      $hvw$  =  $hwm[:, p]$ 
10:     $hrv$  =  $O$ 
11:   else if ( $l$  = MONOTONIC-READ-YOUR-WRITE)
12:     $hvw$  =  $hwm[:, p]$ 
13:     $hrv$  =  $hrm[:, p]$ 
14:   send  $\langle \text{GETREQ } k, hvw, hrv \rangle$  to server  $i$ 
15:   receive  $\langle \text{GETREPLY } v, dc\_id, log\_idx, t \rangle$ 
16:    $hrm[dc\_id, s] = \max(hrm[dc\_id, s], log\_idx)$ 
17:    $dt_r = \max(dt_r, t)$ 
18: return  $v$ 

19: PUT (key  $k$ , value  $v$ , WriteConsistencyLevel  $l$ )
20:    $s$  = partition id of  $k$ 
21:   if ( $l$  = EVENTUAL)
22:      $dt$  =  $0$ 
23:   else if ( $l$  = WRITE-FOLLOWS-READS)
24:      $dt$  =  $dt_r$ 
25:   else if ( $l$  = MONOTONIC-WRITE)
26:      $dt$  =  $dt_w$ 
27:   else if ( $l$  = MONOTONIC-WRITE-FOLLOWS-READS)
28:      $dt$  =  $\max(dt_r, dt_w)$ 
29:   send  $\langle \text{PUTREQ } k, v, dt \rangle$  to server  $i$ 
30:   receive  $\langle \text{PUTREPLY } dc\_id, log\_idx, t \rangle$ 
31:    $hwm[dc\_id, s] = \max(hwm[dc\_id, s], log\_idx)$ 
32:    $dt_w = \max(dt_w, t)$ 
33: return

```

than $hrv[i]$ or $hvw[i]$ sent by the client. This guarantees that the server has committed all necessary version before reading the value of the requested key. The server, next, returns the version with the highest timestamp among versions written for the requested key.

In the PUT operation handler, the server first updates its HLC using dt sent by the client by calling method *updateHLC* in Line 6. This will call HLC algorithm [15] and guarantees that HLC_d will be higher than any timestamp read/written by the client based on the session guarantee requested by the client. Next, the server creates a new version, timestamps it with the updated HLC_d , and gives it to the Raft algorithm for replication. Once Raft committed, the servers add the version to the version chain of the key and update sv entry for the local datacenter by the Raft log index. After committing the value on the leader, and updating the local entry of sv , the server returns reply to the client and includes the id of the local datacenter together with the sv entry for the local datacenter, and the assigned timestamp. Similar to the leader, other non-XC servers also update their sv upon committing a new version. On the other hand, when an XC server commits version v , if $v.dc_id$ is the same as the datacenter id of the server, i.e. the committed write is a

Algorithm 2 Server-side

```

1: Upon receive  $\langle \text{GETREQ } k, hvw, hrv \rangle$ 
2:   block while  $\exists i$  s.t. ( $sv[i] < hrv[i] \vee sv[i] < hvw[i]$ )
3:    $v$  = the version for  $k$  with the highest timestamp
4:   send  $\langle \text{GETREPLY } v.value, v.dc\_id, sv[v.dc\_id], v.t \rangle$  to client

5: Upon receive  $\langle \text{PUTREQ } k, value, dt \rangle$  at  $L_d$ 
6:   updateHLC ( $dt$ )
7:    $t = HLC_d$ 
8:   create new version  $v$ 
9:    $v.value \leftarrow value$ 
10:   $v.t \leftarrow t$ 
11:   $v.dc\_id \leftarrow d$ 
12:  append the  $\langle k, v \rangle$  to the Raft log
13:  send  $\langle \text{PUTREPLY } d, sv[d], t \rangle$  to client

14: updateHLC ( $t$ ) at  $L_d$ 
15:   $l' \leftarrow HLC_d.l$ 
16:   $HLC_d.l \leftarrow \max(l', PC_d, t.l)$ 
17:  if ( $HLC_d.l = l' = t.l$ )
18:     $HLC_d.c \leftarrow \max(HLC_d.c, t.c) + 1$ 
19:  else if ( $HLC_d.l = l'$ )  $HLC_d.c \leftarrow HLC_d.c + 1$ 
20:  else if ( $HLC_d.l = l$ )  $HLC_d.c \leftarrow t.c + 1$ 
21:  else  $HLC_d.c \leftarrow 0$ 

22: Upon commit  $\langle k, v \rangle$  with  $log\_idx$  at a non-XC server
23:    $sv[v.dc\_id] = log\_idx$ 
24:   add  $v$  to the version chain of  $k$ 

25: Upon commit  $\langle k, v \rangle$  at  $XC_d$  with  $log\_idx$ 
26:   if  $v.dc\_id = d$ 
27:     send  $\langle \text{REPLICATE } k, v, log\_idx \rangle$  to  $L_{d'}$  for all  $d' \neq d$ 

28: Upon receive  $\langle \text{REPLICATE } k, v, log\_idx \rangle$  at  $L_d$ 
29:   append the  $\langle k, v \rangle$  with  $log\_idx$  to the Raft log

```

local write, it sends a replicate message to the leader of the other groups to propagate a local write to other datacenters. It includes the log index of the operation writing the version. Upon receiving a replicate message, the leaders in other datacenters append the received message to their Raft log. Followers commit versions received from XC servers like normal local versions and update the their sv upon committing in Line 23. If the XC commits a version with a dc_id that is different from the id of the local datacenter, it simply drops the message and does nothing, because the XC of the original datacenter is responsible for propagating the writes to the other datacenters. For the sake of space, the correctness of this protocol is provided in [23].

6 EXPERIMENTAL RESULTS

In this section, we provide the results of benchmarking the protocol defined in Section 5 for the architecture explained in Section 4. In Section 6.1, we provide the experimental setup. In Section 6.2, first, we evaluate the overhead of different levels of session guarantees when clients never change their datacenter. The results

of this case reflect the impact of the delay of the Raft on the performance overhead of session guarantees. Next, we consider the case where clients switch to a remote datacenter for some of their operations. In addition to the Raft delay, the results of this case also show the impact of cross-datacenter propagation delay on the performance overhead of providing session guarantees. In Section 6.3, we investigate the effect of workload characteristics. Section 6.4 provides the overall implication of the results.

6.1 Experimental Setup

NuKV is implemented in C++. For Raft, it uses an enhanced version of [1]. The client and servers are connected by gRPC [3] and we use Google Protocol Buffers [2] for marshaling/unmarshaling the data.

We did our experiments on a deployment including two datacenters, each with three replicas on different machines. Inside a datacenter, we also have an XC server that forwards the log commits to the leader of the other datacenter. One datacenter is located at Phoenix, Arizona, and the other datacenter is located at Salt Lake City, Utah. The average RTT time between these two datacenters is ~15 (ms). We run all replicas and XC servers on machines with the following specification: 4 vCPUs, Intel Core Processor (Haswell) 2.0 GHz, 4 GB memory, 40 GB Storage Capacity running Ubuntu 16.04.3.

We run client threads from both datacenters. We use one machine with following specification for each datacenter to run the client threads: 8 vCPUs, Intel Core Processor (Haswell) 2.0 GHz, 4 GB memory, 40 GB Storage Capacity running Ubuntu 16.04.3.

Each client thread randomly decides to write or read a key. When it decides to write a key, it randomly picks one of the leaders and writes on it. On the other hand, when it decides to read a key, it randomly picks one of the six replicas and reads from it. These random selections are done with the given probabilities. With different probabilities, we can study the effect of different client usage patterns and the workload characteristics on the performance of the system. We consider 16B keys and 64B values.

6.2 Effect of Locality of Traffic

In the first set of experiments, we investigate the effect of locality of traffic on the performance of our key-value store. First, we consider the case where clients only access their local datacenter.

We consider seven cases as shown in Table 1. The cases with *HLC* subscript show the results for our protocol that uses HLC, and cases without *HLC* subscript show results when we do not use HLC. If we do not use HLC then the PUT operation must wait until the current server receives all the relevant updates that must occur before the current PUT operation. However, with HLC, we can simply assign the new PUT operation a higher timestamp to ensure that it will be ordered later than previous writes.

Figure 2(a) and Figure 2(b) show the effect of load on the operation latency and throughput, respectively. In all diagrams with throughout, we report the total operations done by the clients accessing one of the datacenters. The results are for 0.5:0.5 put:get ratio, i.e., 50% of operations are PUT, and 50% GET. Note that when clients do not change their datacenters, our system actually provides sequential consistency for keys of each partition thanks to the

Raft protocol [21], i.e. all updates are applied with the same order on all replicas in one datacenter. However, regarding the recency of the updates, it is possible that a client reads a version in one replica, but does not find it on another replica.

As expected, the latency and throughput increase in all cases as we run more client threads. Since propagation delay is very small inside one datacenter, providing session guarantees causes only a negligible overhead such that for some cases we even observed better latency and throughput for stronger guarantees compared with eventual consistency due to the experimental error. However, generally, eventual consistency and M/E , and M/E_{HLC} show better results. In all cases, the additional latency compared with eventual consistency always remains less than 1 (ms).

Next, we consider the case where clients may use the remote datacenter with 10% probability. Figures 3(a) and 3(b) show the effect of load on the operation latency with 0.5:0.5 put:get ratio. Unlike the case with 100% local traffic, the difference between eventual consistency and other cases is clear here. In all cases, eventual consistency has the lowest latency. M/E_{HLC} provides the same latency as that of eventual consistency due to write-free operations by using HLCs. For 40 client threads, M/E , E/M , and M/M requires an additional ~10 (ms) compared with eventual consistency. When we use HLCs, M/E_{HLC} requires no additional latency, and E/M_{HLC} and M/M_{HLC} requires ~7.7 (ms) and ~8.6 (ms) additional latency. This improvement allows us to process 12-160% additional operations (160% occurs for the case M/E where the use of HLC eliminates the delay in write operations thereby allowing clients to issue more operations).

To further analyze the effect of locality of traffic on the performance, we consider the system with various local access probabilities. Figures 4 and 5 show how the average latency and throughput change as we increase the local access probability for 40 client threads in each datacenter. The change in the latency and throughput is not clear, except for case with local access probability 1 where the latency collapses, and throughput increases significantly.

6.3 Effect of Workload

In this section, we want to see how the performance changes for workloads with various natures (e.g. read-heavy, write-heavy). The results are for 40 client threads per datacenter. First, we consider the sticky clients, i.e., 100% local traffic. Figure 6(a) shows the effect of write proportion on the average latency. As write proportion increases the latency increases which is expected, because write operations are expected to be heavier than read operations. Figure 6(b) shows how throughput changes as we change the workload for different levels of consistency. Like Section 6.2, the overhead of session guarantees for sticky clients is negligible. For pure-read workload, there is no meaningful overhead, as there is no update and all replicas provide the same data. Thus, all session guarantees are always satisfied.

Figures 7(a) and 7(b) show the effect of write proportion on the average latency for the case with 10% remote traffic. Again, we see that the difference between eventual consistency and stronger consistency models is more obvious due to inter-datacenter replication latency. Also, like all other results M/E_{HLC} provide the same performance as eventual consistency. For M/M_{HLC} consistency

w/o HLC	w/ HLC	Write Guarantee	Read Guarantee
E		Eventual	Eventual
M/E	M/E_{HLC}	Monotonic-write-follows-reads	Eventual
E/M	E/M_{HLC}	Eventual	Monotonic-read-your-write
M/M	M/M_{HLC}	Monotonic-write-follows-reads	Monotonic-read-your-write

Table 1: Legend for Figures 2- 7

(a) Latency

(b) Latency

Figure 2: The effect of load on the latency and throughput of mixed PUT and GET operations with 100% local traffic

level, the latency drops as we move from 0.75 write proportion to 1. The reason is with write proportion 1, there is no read operation, and for write operations, M/M_{HLC} is the same as M/E_{HLC} . Thus, it does not require any blocking. Figure 7(c) shows how throughput changes as we change the workload for different levels of consistency with 10% remote access.

6.4 Implication of Experimental Results

From these results, we find that if we introduce session guarantees, and the client remains within the same datacenter, the overhead is within experimental error. Furthermore, even if the client changes its datacenter, the cost of increased latency is very small, $\sim 10(ms)$. From these results, we find that the session guarantees considered in this paper can be achieved with a very low cost. Without session guarantees, the application can suffer from issues such as the user writes a new value for a key, but obtains the old value when he reads the key, or the user changes the password twice, but finds that the stored password is not the same as the latest one. Such common but highly undesirable problems with eventual consistency can be eliminated with a very low overhead with session guarantees.

Finally, this low overhead also suggests that it is not worthwhile to attempt session guarantees by having clients cache the data they have read or written, as it complicates the design of clients substantially. This is especially important to support computing-challenged clients.

7 RELATED WORK

Terry et al. provided a protocol for providing session guarantees in [28]. An important improvement over [28] is utilizing HLCs to avoiding blocking write operations. Also, we consider the per-key versions of session guarantees considered in [28]. This allows us to avoid the overhead of inter-partition communication which in turns eliminates the possibility of slowdown cascade [4] for systems with a large number of partitions. The protocol proposed in [28] for providing session guarantees is used in Bayou architecture [9, 22]. Similar approach is provided in [7, 14]. Bermbach et al. [6] have provided a middleware to provide monotonic-read and read-your-write guarantees on top of eventually consistency systems. There are no experimental results provided in existing works such as [6, 7, 9, 14, 22, 28] to evaluate the cost of providing session guarantees. In our work, we investigated the cost of providing session guarantees in our architecture that uses Raft for replication.

In recent years, causal consistency has gotten significant attention from the community [5, 10–13, 17, 18, 20, 24, 25, 27, 29, 30]. Causal consistency is especially interesting, as it is the strongest consistency model that remains available under network partitions [19]. However, for this availability clients must be sticky [24]. The broad approach for providing causal consistency in these systems is to track the causal dependencies of a version, and check them before making the version visible in another replica. This requires inter-partition communication that results in more overhead. This results in higher visibility latency in causal systems. This delay can also result in slowdown cascade that impacts the application further [4].

Figure 3: The effect of load on the latency and throughput of mixed PUT and GET operations with 10% remote traffic

Figure 4: The effect of local access probability on the latency of mixed PUT and GET operations.

Figure 5: The effect of local access probability on the throughput of mixed PUT and GET operations.

To avoid slowdown cascade, [20] makes updates visible as soon as they arrive a replica, but at the read time informs the client, if it cannot guarantee the causal consistency according to client's session. Although this reduces the update visibility latency, tracking causal dependencies has much more overhead than per-key session guarantees considered in our paper. Basically, to accurately track

causal dependencies of a version, we have to store timestamps of size $O(n)$ where n is the number of all partitions that we can write [8]. For systems with thousands of partitions such as ours, such overhead is not practical. We can reduce the size of metadata at the cost of false positive causal consistency violation detection. Generally, there is a trade-off between the overhead of metadata to

Figure 6: The effect of workload on latency and throughput of mixed PUT and GET operations with 100% local traffic

Figure 7: The effect of workload on latency and throughput of mixed PUT and GET operations with 10% remote traffic

track causal dependencies and accuracy of tracking [26]. To reduce the overhead of causal dependency tracking, session-based causal systems such as [20] assumes more restrictive assumptions. For instance, [20] reduces n by assuming a single-leader replication where clients can only write on a single machine per partition. On the other hand, in our system, a client can write on any leader in any datacenter (e.i., multi-leader). Also, in our protocol, we only need to store a single scalar timestamp with each key. Considering the overhead of causal consistency in terms of storage, network communication, and visibility latency, we decided to use session guarantees instead.

HLCs are used in some existing protocols to overcome the imperfection of clock synchronization. HLCs are used in [24, 25] to provide causal consistency for basic read and write operations for a key-value store. Spirovska et al. [27] have utilized HLCs for transactional causal consistency. HLCs are also used in the context of stronger consistency.

8 CONCLUSION

In this paper, we presented our analysis of the cost associated with providing session guarantees for NuKV, a key-value store that aims to simultaneously provide high availability and consistency for eBay services. In particular, NuKV maintains the data in multiple datacenters and each datacenter contains multiple replicas of the data. Using Raft, NuKV allows clients to obtain sequential consistency by accessing any replica in the given datacenter. At the same, it provides flexibility where a client can access any datacenter. When a client changes its datacenter, NuKV provides eventual consistency.

In this paper, we showed that one could upgrade the guarantees provided to the client when it switches a datacenter at a low cost. Specifically, we showed that the increased latency to provide session guarantees is negligible (less than 10 ms). We demonstrated this for different types of workloads (100% read, read-heavy, write-heavy and 100% write). Our analysis showed that providing monotonic-write-follows-read highly benefits from the use of HLCs [15].

Our design was based on a modified version of session guarantees that enabled us to avoid slowdown cascade where the delay of receiving some data item causes visibility of other data items from being available to the clients in a cascading fashion. This cascading delay is one of the important obstacles in providing an higher levels of consistency such as causal consistency. From this analysis, we argue that providing session guarantees is a significantly less expensive but highly valuable approach for key-value stores.

REFERENCES

- [1] Cornerstone. <https://github.com/datatechnology/cornerstone>.
- [2] Google protocol buffers. <https://developers.google.com/protocol-buffers/>.
- [3] grpc. <https://grpc.io/>.
- [4] Philippe Ajoux, Nathan Bronson, Sanjeev Kumar, Wyatt Lloyd, and Kaushik Veeraraghavan. Challenges to adopting stronger consistency at scale. In *HotOS*, 2015.
- [5] Peter Bailis, Ali Ghodsi, Joseph M Hellerstein, and Ion Stoica. Bolt-on causal consistency. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 761–772. ACM, 2013.
- [6] David Bermbach, Jörn Kuhlenskamp, Bugra Derre, Markus Klems, and Stefan Tai. A middleware guaranteeing client-centric consistency on top of eventually consistent datastores. In *Cloud Engineering (IC2E), 2013 IEEE International Conference on*, pages 114–123. IEEE, 2013.
- [7] Jerzy Brzeziński, Cezary Sobaniec, and Dariusz Wawrzyniak. Safety of a server-based version vector protocol implementing session guarantees. In *International Conference on Computational Science*, pages 423–430. Springer, 2005.
- [8] Bernadette Charron-Bost. Concerning the size of logical clocks in distributed systems. *Information Processing Letters*, 39(1):11–16, 1991.
- [9] Alan Demers, Karin Petersen, Mike Spreitzer, Doug Terry, Marvin Theimer, and Brent Welch. The bayou architecture: Support for data sharing among mobile users. In *Mobile Computing Systems and Applications, 1994. WMCSA 1994. First Workshop on*, pages 2–7. IEEE, 1994.
- [10] Diego Didona, Rachid Guerraoui, Jingjing Wang, and Willy Zwaenepoel. Causal consistency and latency optimality: Friend or foe? *arXiv preprint arXiv:1803.04237*, 2018.
- [11] Diego Didona, Kristina Spirovska, and Willy Zwaenepoel. Okapi: Causally consistent geo-replication made faster, cheaper and more available. *arXiv preprint arXiv:1702.04263*, 2017.
- [12] Jiaqing Du, Sameh Elnikety, Amitabha Roy, and Willy Zwaenepoel. Orbe: Scalable causal consistency using dependency matrices and physical clocks. In *Proceedings of the 4th Annual Symposium on Cloud Computing, SOCC '13*, pages 11:1–11:14, New York, NY, USA, 2013.
- [13] Jiaqing Du, Călin Iorgulescu, Amitabha Roy, and Willy Zwaenepoel. Gentlerain: Cheap and scalable causal consistency with physical clocks. In *Proceedings of the ACM Symposium on Cloud Computing, SOCC '14*, pages 4:1–4:13, New York, NY, USA, 2014.
- [14] Anna Kobusinska, Cezary Sobaniec, Marek Libuda, and Dariusz Wawrzyniak. Version vector protocols implementing session guarantees. In *Cluster Computing and the Grid, 2005. CCGrid 2005. IEEE International Symposium on*, volume 2, pages 929–936. IEEE, 2005.
- [15] Sandeep S Kulkarni, Murat Demirbas, Deepak Madappa, Bharadwaj Avva, and Marcelo Leone. Logical physical clocks. In *International Conference on Principles of Distributed Systems*, pages 17–32. Springer, 2014.
- [16] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, July 1978.
- [17] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don't settle for eventual: Scalable causal consistency for wide-area storage with cops. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles, SOSP '11*, pages 401–416, New York, NY, USA, 2011.
- [18] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Stronger semantics for low-latency geo-replicated storage. In *NSDI*, volume 13, pages 313–328, 2013.
- [19] Prince Mahajan, Lorenzo Alvisi, Mike Dahlin, et al. Consistency, availability, and convergence. *University of Texas at Austin Tech Report*, 11, 2011.
- [20] Syed Akbar Mehdi, Cody Little, Natacha Crooks, Lorenzo Alvisi, Nathan Bronson, and Wyatt Lloyd. I can't believe it's not causal! scalable causal consistency with no slowdown cascades. In *NSDI*, pages 453–468, 2017.
- [21] Diego Ongaro and John K Ousterhout. In search of an understandable consensus algorithm. In *USENIX Annual Technical Conference*, pages 305–319, 2014.
- [22] Karin Petersen, Mike Spreitzer, Douglas Terry, and Marvin Theimer. Bayou: replicated database services for world-wide applications. In *Proceedings of the 7th workshop on ACM SIGOPS European workshop: Systems support for worldwide applications*, pages 275–280. ACM, 1996.
- [23] Mohammad Roohitavaf, Jung-Sang Ahn, Woon-Hak Kang, Kun Ren, Gene Zhang, Sami Ben-Romdhane, and Sandeep S Kulkarni. Session guarantees with raft and hybrid logical clocks. *arXiv preprint arXiv:1808.05698*, 2018.
- [24] Mohammad Roohitavaf, Murat Demirbas, and Sandeep Kulkarni. Causalspartan: Causal consistency for distributed data stores using hybrid logical clocks. In *Reliable Distributed Systems (SRDS), 2017 IEEE 36th Symposium on*, pages 184–193. IEEE, 2017.
- [25] Mohammad Roohitavaf and Sandeep Kulkarni. Gentlerain+: Making gentlerain robust on clock anomalies. *arXiv preprint arXiv:1612.05205*, 2016.
- [26] Mohammad Roohitavaf and Sandeep Kulkarni. Toward adaptive causal consistency for replicated data stores. *arXiv preprint arXiv:1803.08609*, 2018.
- [27] Kristina Spirovska, Diego Didona, and Willy Zwaenepoel. Wren: Nonblocking reads in a partitioned transactional causally consistent data store. In *48th International Conference on Dependable Systems and Networks (DSN'18)*, number CONF, 2018.
- [28] Douglas B Terry, Alan J Demers, Karin Petersen, Mike J Spreitzer, Marvin M Theimer, and Brent B Welch. Session guarantees for weakly consistent replicated data. In *Parallel and Distributed Information Systems, 1994., Proceedings of the Third International Conference on*, pages 140–149. IEEE, 1994.
- [29] Zhuolun Xiang and Nitin H Vaidya. Global stabilization for causally consistent partial replication. *arXiv preprint arXiv:1803.05575*, 2018.
- [30] Marek Zawirski, Nuno Preguiça, Sérgio Duarte, Annette Bieniusa, Valter Balesgar, and Marc Shapiro. Write fast, read in the past: Causal consistency for client-side applications. In *Proceedings of the 16th Annual Middleware Conference*, pages 75–87. ACM, 2015.