Towards Reconfigurable Linearizable Reads

Myles Thiessen
mthiessen@cs.toronto.edu
University of Toronto
Toronto, Canada

Aleksey Panas
aleksey.panas@mail.utoronto.ca
University of Toronto
Toronto, Canada

Guy Khazma
guykhazma@cs.toronto.edu
University of Toronto
Toronto, Canada

Eyal de Lara
delara@cs.toronto.edu
University of Toronto
Toronto, Canada

Abstract

Linearizable datastores are desirable because they provide users with the illusion that the datastore is run on a single machine that performs client operations one at a time. To reduce the performance cost of providing this illusion, many specialized algorithms for linearizable reads have been proposed which significantly improve read performance compared to write performance. The main difference between these specialized algorithms is their performance under different workloads. Unfortunately, since a datastore's workload is often unknown or changes over time and system designers must decide on a single read algorithm to implement ahead of time, a datastore's performance is often suboptimal as it cannot adapt to workload changes.

In this paper, we lay the groundwork for addressing this problem by proposing Chameleon, an algorithm for linearizable reads that provides a principled approach for datastores to switch between existing read algorithms at runtime. The key observation that enables this generalization is that all existing algorithms are specific read-write quorum systems. Chameleon constructs a generic read-write quorum system, by using *tokens* that are *included* to complete write and read operations. This *token quorum system* enables Chameleon to mimic existing read algorithms and switch between them by transferring these tokens between processes.

1 Introduction

A datastore is *linearizable* if it provides users with the illusion that it is run on a single machine that processes client operations one at a time [15]. This property is desirable, as application developers do not need to reason about side effects that can arise from concurrent datastore operations [14, 21].

The de facto standard for providing linearizability is by use of a state machine replication algorithm [27]. These algorithms provide clients with procedures for both writing and reading the application's state. While both write and read operations can be processed with the same procedure, the non-mutating property of read operations allows them to be processed concurrently without coordination between them. This observation has led to the development of many specialized read algorithms which significantly improve read performance compared to write performance [4–8, 13, 17, 24]. These algorithms are of particular practical importance as many real-world workloads are read dominant [3, 9]. For example, Google reported that in their ads workload, read operations outnumber write operations by three orders of magnitude [10]. This prevalence makes read operation performance the primary determinant of end-user latency and overall system throughput.

Existing read algorithms can be classified into one of four categories: leader [6], majority quorum [8, 31], flexible quorum [13, 30], and local [4, 5, 7, 17, 24]. The main difference between algorithms in these categories is their performance in different workloads. For example, in geo-distributed deployments, if most read operations are performed near a distinguished process known as the *leader*, average read latency will be significantly lower using leader reads compared to majority quorum reads. However, majority quorum reads can achieve higher peak throughput comparatively as the leader is not a bottleneck [2, 8]. Moreover, local reads provide the lowest read latency compared to algorithms in all other categories at the cost of increased write latency [24]. Finally, flexible quorum reads increase read latency but reduce write latency in the presence of failures or network partitions.

These differences pose a serious challenge during the design of a linearizable datastore when the expected workload is unknown ahead of time. Even in the unlikely scenario where it is known, designers can only implement a single read algorithm, and as such, performance will degrade if the workload changes. Consequently, to optimize performance when the workload is unknown ahead of time and to adapt to workload changes, a datastore needs to be able to switch between existing read algorithms at run-time.

In this paper, we lay the groundwork for switching between different read algorithms at run-time by presenting *Chameleon*, a generalized and reconfigurable algorithm for linearizable reads. The key observation that enables this generalization is that *all existing algorithms are specific read-write quorum systems* — their correctness is solely based on read and write operations contacting overlapping sets of processes. To guarantee read-write quorum intersection, each process is provided with a set of *tokens* that are *included* to complete both read and write operations. This *token quorum*

(a) Leader reads.   **(b)** Majority quorum reads.   **(c)** Flexible quorum reads.   **(d)** Local reads.
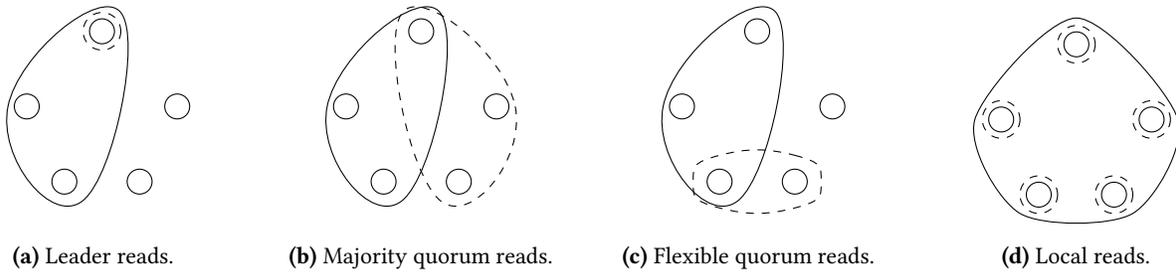
Figure 1. Example read (dashed) and write (solid) quorums of existing algorithms for linearizable reads.

system enables Chameleon to mimic existing read algorithms and switch between them by strategically transferring these tokens between processes.

In the remainder of this paper, we describe Chameleon in detail. Section 2 argues that all existing algorithms are specific read-write quorum systems. Section 3 describes the token quorum system and shows how it can mimic all existing algorithms. Section 4 discusses reconfiguring the token quorum system and tolerating failures. Section 5 discusses related work, and Section 6 presents concluding remarks.

2 Preliminaries

In this section, we discuss the system model, state machine replication, existing algorithms for linearizable reads, and their relationship to read-write quorum systems.

2.1 Model

We assume an asynchronous system consisting of n processes of which at most f may crash where $f < \frac{n}{2}$. In this system, messages passed between processes may be arbitrarily delayed, reordered, or lost entirely. In addition to this model, we require processes to be able to grant *leases* in order to tolerate failures without sacrificing liveliness [12]. Specifically, we require that processes can grant leases *correctly*, that is, the lease granter's perception of the lease's expiration is guaranteed to occur after the holder's perception of the lease's expiration in real-time. This can be done by either assuming processes have access to synchronized clocks with bounded clock skew [12] or by assuming processes are equipped with hardware clocks with bounded drift [22].

2.2 State Machine Replication

The de facto standard for providing linearizable, fault-tolerant replication of client operations across a set of processes is by use of a state machine replication (SMR) algorithm. Linearizability is a consistency model that ensures client operations are executed in some total order that respects their real-time ordering, i.e., if operation o_1 completes before operation o_2 begins, then o_1 must be ordered before o_2 [15].

SMR algorithms provide fault-tolerance by maintaining a copy of the application's state at each process known as their

replica. Replicas are assumed to atomically apply client operations and be deterministic, as in, if two replicas apply the same sequence of operations, they reach the same state [27].

Clients perform operations by submitting them to processes within the system. Once a process receives such a request, it performs either the write or read procedure of the SMR algorithm depending on the submitted operation, and returns the result to the client. Intuitively, the result for a write operation signifies the operation has taken effect and is durable whereas the result of a read operation is the latest version of the application's state.

To provide the real-time ordering guarantee of linearizability, SMR algorithms often rely on a distinguished process known as the *leader* to sequence write operations. The leader does so in two phases: prepare and commit. The prepare phase consists of the leader proposing that a write operation w be assigned to the i-th location in the *log*. The prepare phase completes once a majority of processes accept this proposal ensuring it is *decided*. After this, w can be committed, that is, a process can apply w to their replica after they apply all write operations assigned up to index i.

Read operations are processed by assigning them to the index number of either the latest committed write operation or some concurrent write operation. Once a read operation has been assigned to some index number i, it is completed by executing it against some replica that has applied all write operations with index numbers up to and including i.

2.3 Algorithms for Linearizable Reads

We now describe the four categories of read algorithms mentioned previously. In leader reads all read operations are forwarded to the network's leader for processing. When the leader receives a read request it assigns it to the highest index that it has sent a commit request for. This guarantees that read operations always observe the latest complete write operations so long as a new leader has not been elected in the meantime. To ensure this doesn't happen, the leader is granted a *leader lease*, guaranteeing at most one process believes they are the leader at any given time. Unlike leader reads, in majority quorum reads, all processes can perform read operations. They do so by assigning read operations

(a) Leader reads. (b) Majority quorum reads. (c) Flexible quorum reads. (d) Local reads.
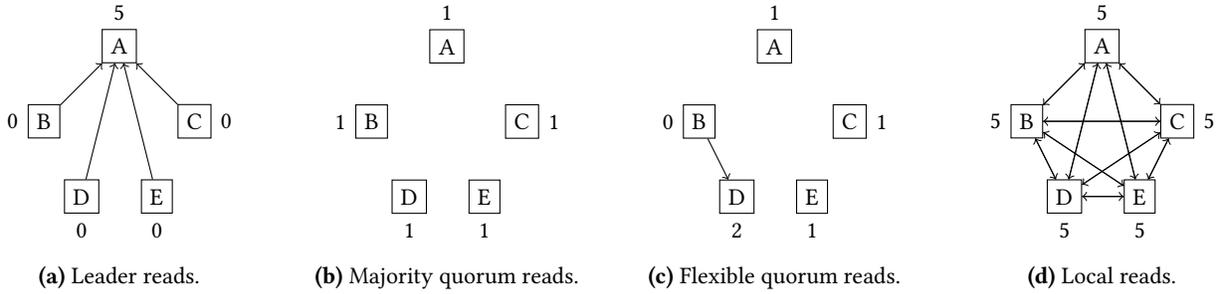
Figure 2. Examples of Chameleon mimicking existing algorithms for linearizable reads by transferring tokens between processes (represented by arrows). The number beside each process represents how many tokens it holds.

to the maximum index a simple majority of processes have received a prepare request for. Flexible quorum reads are similar to majority quorum reads except reads are assigned to the maximum index of an arbitrary read quorum of processes. To enable this flexibility and guarantee the real-time ordering requirement of linearizability, the leader waits to receive acknowledgments from at least one process in every read quorum before committing a write. Similarly to leader reads, to not reduce fault-tolerance, processes that are members of a read quorum are granted *read leases*. Finally, in local reads, all processes assign read operations to an index based on their local perception of the latest assigned write index. Similar to flexible quorum reads, the leader in local reads must receive acknowledgments from all processes that hold a read lease before a write is committed.

2.4 Read-Write Quorum Systems

We now discuss the relationship between existing algorithms for linearizable reads and read-write quorum systems. We do so with the aid of the illustration in Figure 1. A read-write quorum system is defined by a set of read quorums and write quorums such that every read and write quorum intersects. In the context of specialized algorithms for linearizable reads, each read and write quorum is a subset of the n processes in the system. In leader reads, the only read quorum is the leader and any simple majority of processes including the leader is a write quorum. Since at most one process is the leader at any time, every read and write quorum intersects. In majority quorum reads, read and write quorums are any simple majority of processes. Similarly, since any two simple majorities intersect so does every read and write quorum. Flexible quorum reads define the quorum system explicitly. Finally, in local reads, every process is a read quorum and there is a single write quorum containing all processes. Consequently, *all existing algorithms for linearizable reads reduce to specific read-write quorum systems.*

3 Chameleon

We now explain how Chameleon constructs a read-write quorum system using *tokens*, followed by showing how this

enables Chameleon to mimic all existing specialized read algorithms. We then describe how Chameleon performs write and read operations assuming that messages aren't lost, all processes are non-faulty, and the quorum system and leader are fixed followed by sketching why Chameleon is correct. We discuss relaxing these assumptions in Section 4.

3.1 The Token Quorum System

Each token is *owned* by one process and *held* by at most one process at any time. Note, a token's owner *never* changes but its holder can. We say that a token is *transferred* from one process to another. A token is defined as a tuple (l, r) where l is the token's owner and r is some unique integer used to differentiate between tokens owned by the same process. Using these tokens a read quorum is then defined by a set of processes that hold at least *one* token owned by some simple majority of processes. Similarly, a write quorum is defined by a set of processes which is at least a simple majority of processes that collectively hold *every* token owned by some (potentially different) simple majority of processes. Finally, we say that an operation *included* a token if it completes by contacting a quorum that held it.

3.2 Mimicking Existing Algorithms

We now show how strategically assigning which processes hold which tokens enables Chameleon to mimic each of the existing read algorithms discussed in Section 2.

To mimic leader reads, each process owns a single token which is held by the leader. This is illustrated in Figure 2a where the arrows represent the transferring of tokens owned by processes B, C, D, and E to A (the leader), and the number beside each process represents how many tokens it holds. In this scenario, the leader itself is a read quorum as it holds at least one token from a simple majority of processes. Furthermore, a write quorum is any simple majority of processes that includes the leader — the only process that holds every token from a majority of processes.

To mimic majority quorum reads each process owns a single token which is not transferred. Consequently, any simple majority of processes is both a read and write quorum

— the same quorum system produced by majority quorum reads. We illustrate an example of this in Figure 2b.

Chameleon mimics flexible quorum reads through a combination of the previous two approaches, as only one process needs to hold more than one token to achieve a read quorum smaller than a simple majority. We illustrate such an example in Figure 2c where process D holds two tokens. This enables D and processes A, C, or E to perform read operations. Furthermore, in this example, since a write must include every token owned by a simple majority of processes, it is bound to intersect D or A, C, and E. In either case, this guarantees read-write quorum intersection.

Finally, Chameleon mimics local reads by each process owning n tokens and transferring one token to all other processes. Consequently, each process will hold onto one token owned by every process in the system as we illustrate in Figure 2d. This results in each process being a read quorum and the only valid write quorum being all processes. Consequently, the token quorum system can mimic *every* existing specialized algorithm for linearizable reads.

3.3 Performing Client Operations

We now discuss how Chameleon performs operations using the token quorum system.

Writes (Algorithm 1). When a process wishes to perform a write operation, it first forwards it to the leader. Once the leader receives this request, it assigns the write operation to the next index in the log and sends a prepare request to all processes including the write operation and its assigned index number. Upon receipt of this request, a process records this index number if it is the highest they have received a prepare request for. It then sends an acknowledgment to the leader including a description of the current set of tokens they hold. In the system illustrated in Figure 2c, the set of tokens returned per process differs. For example, processes A, C, and E will inform the leader they hold their own token, process B will return saying they do not hold any tokens, and process D will return saying they hold both their own and B's token. Once the leader has sent these prepare requests, it waits to receive acknowledgments from at least a simple majority of processes that collectively hold every token owned by some (potentially different) simple majority of processes. In the system illustrated in Figure 2c, valid sets of acknowledging processes include (A, C, E), (A, D, E), and (C, D, E). Following this, the leader sends commit requests to all processes for this write and returns to the client signifying that the write is completed and durable.

Reads (Algorithm 2). Similarly to write operations, clients submit read operations to any process. Once a process p receives this request, it first computes the closest read quorum to it (R) based on the tokens each process holds. If this quorum happens to be p itself, then p assigns the read to the highest index it has received a prepare request for. Otherwise, p sends a *read* request to all processes in R. For example, if

```
1  procedure write(o):
2      cntr := cntr + 1    /* Local op counter initially 0 */
       /* cntr on lines 3 and 4 is the same as line 2     */
3      send ⟨WRITE, o, cntr⟩ to ℓ
4      wait to receive ⟨WRITE_ACK, cntr⟩ from ℓ

5  upon receiving ⟨WRITE, o, c⟩ from p:
6      i := i + 1           /* Latest log index initially 0 */
7      send ⟨PREPARE, i⟩ to all
       /* Processes that have acknowledged, the tokens they
          have returned, and the current include.        */
8      A := TR := TI := ∅
9       upon receiving ⟨P_ACK, i, T′⟩ from q:
10          A := A ∪ {q} and TR := TR ∪ T′
11          foreach process r in the system do
12              if TR contains k tokens owned by r then
13                  TI := TI ∪ {r}
14      wait for |A| ≥ ⌈n+1/2⌉ and |TI| ≥ ⌈n+1/2⌉
15      send ⟨COMMIT, i, o⟩ to all
16      send ⟨WRITE_ACK, c⟩ to p

17  upon receiving ⟨PREPARE, i⟩ from ℓ:
       /* Maximum prepare index received initially 0     */
18      MaxP := max(MaxP, i)
       /* T is the set of held tokens.                   */
19      send ⟨P_ACK, i, T⟩ to ℓ

20  upon receiving ⟨COMMIT, i, o⟩ from ℓ:
       /* Applies the write o against the local replica once
          all writes up to i − 1 have been applied       */
21      Write(i, o)
```

Algorithm 1: Write pseudocode for Chameleon assuming no message loss, all processes are non-faulty, a fixed leader ℓ, each process holds a fixed set of tokens T, and each process owns k tokens.

process p is process A in Figure 2c possible R's include (A, C, E), (A, D), (C, D), and (D, E). Once a process receives a read request it returns the highest index it has received a prepare request for and the set of tokens it currently holds. Once p has sent these read requests, it waits to receive acknowledgments from a set of processes that collectively hold at least one token owned by some simple majority of processes. Following this, p assigns the read operation to the maximum index returned from this set of processes. In either case, p then completes the read operation by (1) waiting for its replica to commit all write operations up to and including the assigned index, (2) executing the read operation against its local replica, and (3) returning the result to the client.

Sending read requests to the closest read quorum is similar to the *thrifty* optimization of Paxos [20] which only sends prepare messages to a majority of processes. This reduces the number of messages sent by the leader but increases tail latency when processes don't respond. To avoid this trade-off, R can be replaced with all processes in the system.

```
1   procedure read(o):
2       cntr := cntr + 1 and index := 0
3       R := closest_read_quorum()
4       if R is only the current process then
5       │   index := MaxP
6       else
            /* cntr on lines 7 and 9 is the same as line 2    */
7       │   send ⟨READ, cntr⟩ to all processes in R
8       │   TI := ∅
9       │    upon receiving ⟨R_ACK, cntr, T', MaxP'⟩:
10      │   │   index := max(index, MaxP')
11      │   │   foreach process p in the system do
12      │   │   │   if (p, *) ∈ T' then  TI := TI ∪ {p}
13      │   wait for |TI| ≥ ⌈(n+1)/2⌉
        /* Executes the read request o against the local
           replica once all writes up to and including index
           have been applied and return the result          */
14      return Read(index, o).

15  upon receiving ⟨READ, c⟩ from p:
16  │   send ⟨R_ACK, c, T, MaxP⟩ to p
```
Algorithm 2: Read pseudocode for Chameleon under the same assumptions as Algorithm 1.

3.4 Correctness Sketch

Chameleon is similar to MultiPaxos [19] and Raft [24] in the sense that the leader processes writes in two phases and assigns writes to log indexes in increasing order. Consequently, the argument for Chameleon's correctness is fundamentally similar to these algorithms and as such we focus on its sole difference — how reads are ordered with respect to writes. Specifically, Chameleon must guarantee that if a write operation w assigned to index i completes before some read operation r begins then r is assigned to an index at least i. This follows from our assumption that the token quorum system is fixed and the fact that every read and write quorum has at least one token in their intersection. These together imply that some process p was in the intersection of w's write quorum and r's read quorum. Consequently, since w completed before r began, p processed the prepare request for w before the read request for r, and therefore p would have returned a maximum prepare index of at least i in response to the read request for r. This in turn implies that r is assigned to at least i as required.

4 Discussion

We now discuss mechanisms for reconfiguring the token quorum system and tolerating message loss and failures.

4.1 Reconfiguring the Token Quorum System

Chameleon's novelty lies in its ability to switch between existing algorithms for linearizable reads at runtime. This is achieved by transferring tokens between processes — reconfiguring the token quorum system in the process. To maintain correctness Chameleon must ensure that all read operations observe all previously completed write operations. This can be achieved by ensuring that if a process holds some token t it has a maximum prepare index at least as large as the maximum index number of any completed write operation that included t. We now discuss a synchronous centralized approach to guarantee this. For future work, we plan to explore how to minimize reconfiguration times by designing an asynchronous approach for reconfiguration.

In the centralized approach, the leader appends special entries to the log known as *token configurations* that describe what processes should hold which tokens. Processes then determine the current set of tokens they hold based on the latest configuration. To not impact performance, each process knows the latest configuration locally. To ensure this *local perception* of the token configuration is consistent, processes mark their local perception as invalid once they receive a prepare request for a new token configuration. At this point, processes do not know which tokens they hold and as such, their processing of prepare and read requests is stalled until they know. This occurs after they receive a commit request for the new token configuration, which is only sent once the leader receives acknowledgments from all processes.

To guarantee that all read operations observe all previously completed write operations we need to make two additional changes. First, when the leader processes a new token configuration, it (1) waits for all outstanding write operations to complete before proposing the token configuration and (2) stalls the processing of new write operations until the new token configuration has been acknowledged by all processes. Second, when acknowledging read requests, processes return the index of the token configuration it used to compute the set of tokens it holds. The process performing the read operation then keeps track of the highest index returned and counts tokens only from processes that return that index. In the event this is insufficient to cover a read quorum, the process resends read requests until it covers a read quorum. These changes ensure that each token holder has a maximum prepare index of at least as large as the maximum index number of any completed write operation. This is because the new token configuration's index number is higher than all previously completed writes since the leader waited for all outstanding writes to be completed. Furthermore, read operations are bound to observe all previously completed write operations since readers only include tokens from the latest token configuration.

4.2 Tolerating Message Loss & Failures

Like in existing algorithms, message loss can be tolerated by supporting request re-transmission and detecting duplicate operations. Consequently, Chameleon can be adapted to do so using existing techniques such as the ones proposed by

CHT [7]. Similarly, tolerating failures in Chameleon is functionally equivalent to how algorithms for local reads tolerate failures. This is because they both have to guarantee liveliness even when a read quorum is partitioned from the rest of the network. To handle this case, algorithms for local reads grant processes *read leases* which are time-based contracts that while valid allow processes to perform read operations against their local replica. When these leaseholders crash or are partitioned from the rest of the network, the leader will wait for the expiration of their read lease before completing ongoing write operations. This guarantees they are no longer performing reads locally and thus it is safe for the rest of the processes to commit said write. We now discuss how this idea can be adapted to Chameleon to tolerate failures.

To enable fault-tolerant local reading of the token configuration, each process is granted a read lease which is revoked when a process doesn't respond to token configuration changes promptly. Similarly, to tolerate the failure of a token holder, holding a token is attached to a lease. Specifically, if process p holds a token t owned by process q, q will grant p the right to hold t for some lease duration. This enables the leader to revoke all tokens held by some process p by contacting all processes and requesting them to wait for their leases granted to p to expire. Once this occurs, the leader includes these revoked tokens until the write quorum condition is met. Like in algorithms for local reads, this ensures that any read quorum that p was a part of is no longer performing read operations after its leases are revoked.

In addition to this revocation mechanism, once a token has been revoked, the token needs to be associated with a valid prepare index to ensure read operations observe all completed write operations. Like in our reconfiguration mechanism, this can be accomplished by relying on the leader to provide their latest assigned log index to the process which is revoking the token. However, unlike the reconfiguration mechanisms, this requires that at most one process is the leader at any time. This can be accomplished by granting the leader a leader lease. To also handle the failure of the leader, we can adopt the enhanced leader election service proposed by CHT [7] that takes care of both granting the leader a leader lease and electing a leader. For future work, we will formalize the mechanism required to tolerate failures.

5 Related Work

SMR algorithms. The most well known SMR algorithms are MultiPaxos [19] and Raft [26]. Like Chameleon, they are leader-based and process writes in two phases. An alternative to these leader-based algorithms is leaderless algorithms such as EPaxos [23, 28] and Nezha [11]. These algorithms allow all processes to perform write operations without contacting the leader at the expense of needing a *super quorum* of acknowledgments instead of a simple majority. We believe Chameleon can be adapted to work with these algorithms.

Quorum systems. Flexible [16], Wan [1], and Dynamic Paxos [25] use non-majority quorum systems to ensure that leader election quorums always intersect with prepare quorums. This guarantees that writes are consistent across leaders. In Flexible and Wan Paxos, the quorum system is configurable but must be known a priori. Dynamic Paxos allows for the quorum system to be reconfigured at runtime but is specific to leader election quorums. Specifically, Dynamic Paxos enables the *leader election zone* to move at runtime to tolerate failures or for access locality. Due to the specific use case of this mechanism, it is not capable of mimicking existing algorithms for linearizable reads. Chameleon, on the other hand, can because of its use of the token quorum system. EdgePQR [13] extends Wan Paxos with the ability to perform low-latency reads at edge data centers. To do so, they introduce the idea of *edge quorums* which enables reads to be processed by contacting any majority of processes within an edge datacenter. To guarantee correctness, all write operations are also required to contact a majority of processes in each edge data center. EdgePQR tolerates the failure of edge datacenters without stalling write operations by enabling the addition and removal of edge datacenters at runtime. This is similar to Chameleon's ability to reconfigure the read-write quorum system but only applies to edge quorums and not the entire read-write quorum system. Quoracle [30] quantifies the trade-offs between different read-write quorum systems but does not consider reconfiguration at runtime.

Token virtualization. An alternative interpretation of the token quorum system presented in this paper is that each token is a *virtual* process which is mapped to a single *physical* process through token holding. Quorum intersection in Chameleon's case is then guaranteed by intersecting sufficiently many virtual processes for both reads and writes. This idea of *token virtualization* has parallels to works in the late 80s and early 90s done on fair resource scheduling [18, 29]. In these works, tokens were known as tickets or shares which were given to users to fairly schedule their programs on shared CPU resources. Chameleon differs from these works in its use of tokens as a means to guarantee read-write quorum intersection instead of fair scheduling.

6 Conclusion

In this paper, we presented Chameleon, an algorithm for linearizable reads that lays out a principled approach for datastores to switch between existing read algorithms at runtime. To do so, Chameleon constructs a read-write quorum system using tokens that are included to complete read and write operations. This token quorum system enables Chameleon to switch between existing read algorithms by transferring them between processes. For future work, we plan to develop a complete description of Chameleon along with an evaluation that demonstrates the benefits of switching between existing read algorithms at runtime.

References

[1] Ailidani Ailijiang, Aleksey Charapko, Murat Demirbas, and Tevfik Kosar. Multileader wan paxos: Ruling the archipelago with fast consensus. *arXiv preprint arXiv:1703.08905*, 2017.

[2] Vaibhav Arora, Tanuj Mittal, Divyakant Agrawal, Amr El Abbadi, Xun Xue, et al. Leader or majority: Why have one when you can have both? improving read scalability in raft-like consensus protocols. In *9th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 17)*, 2017.

[3] Berk Atikoglu, Yuehai Xu, Eitan Frachtenberg, Song Jiang, and Mike Paleczny. Workload analysis of a large-scale key-value store. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*, pages 53–64, 2012.

[4] Jason Baker, Chris Bond, James C Corbett, JJ Furman, Andrey Khorlin, James Larson, Jean-Michel Leon, Yawei Li, Alexander Lloyd, and Vadim Yushprakh. Megastore: Providing scalable, highly available storage for interactive services. 2011.

[5] Changyu Bi, Vassos Hadzilacos, and Sam Toueg. Parameterized algorithm for replicated objects with local reads. *arXiv preprint arXiv:2204.01228*, 2022.

[6] Tushar D Chandra, Robert Griesemer, and Joshua Redstone. Paxos made live: an engineering perspective. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, pages 398–407, 2007.

[7] Tushar D Chandra, Vassos Hadzilacos, and Sam Toueg. An algorithm for replicated objects with efficient reads. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 325–334, 2016.

[8] Aleksey Charapko, Ailidani Ailijiang, and Murat Demirbas. Linearizable quorum reads in paxos. In *11th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 19)*, 2019.

[9] Audrey Cheng, Xiao Shi, Aaron Kabcenell, Shilpa Lawande, Hamza Qadeer, Jason Chan, Harrison Tin, Ryan Zhao, Peter Bailis, Mahesh Balakrishnan, et al. Taobench: an end-to-end benchmark for social network workloads. *Proceedings of the VLDB Endowment*, 15(9):1965–1977, 2022.

[10] James C Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, Jeffrey John Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, et al. Spanner: Google's globally distributed database. *ACM Transactions on Computer Systems (TOCS)*, 31(3):1–22, 2013.

[11] Jinkun Geng, Anirudh Sivaraman, Balaji Prabhakar, and Mendel Rosenblum. Nezha: Deployable and high-performance consensus using synchronized clocks. *Proceedings of the VLDB Endowment*, 16(4):629–642, 2022.

[12] Cary Gray and David Cheriton. Leases: An efficient fault-tolerant mechanism for distributed file cache consistency. *ACM SIGOPS Operating Systems Review*, 23(5):202–210, 1989.

[13] Joshua Guarnieri and Aleksey Charapko. Linearizable low-latency reads at the edge. In *Proceedings of the 10th Workshop on Principles and Practice of Consistency for Distributed Data*, pages 77–83, 2023.

[14] Haryadi S Gunawi, Mingzhe Hao, Tanakorn Leesatapornwongsa, Tiratat Patana-anake, Thanh Do, Jeffry Adityatama, Kurnia J Eliazar, Agung Laksono, Jeffrey F Lukman, Vincentius Martin, et al. What bugs live in the cloud? a study of 3000+ issues in cloud systems. In *Proceedings of the ACM symposium on cloud computing*, pages 1–14, 2014.

[15] Maurice P Herlihy and Jeannette M Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 12(3):463–492, 1990.

[16] Heidi Howard, Dahlia Malkhi, and Alexander Spiegelman. Flexible paxos: Quorum intersection revisited. *arXiv preprint arXiv:1608.06696*, 2016.

[17] Antonios Katsarakis, Vasilis Gavrielatos, MR Siavash Katebzadeh, Arpit Joshi, Aleksandar Dragojevic, Boris Grot, and Vijay Nagarajan. Hermes: A fast, fault-tolerant and linearizable replication protocol. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 201–217, 2020.

[18] Judy Kay and Piers Lauder. A fair share scheduler. *Communications of the ACM*, 31(1):44–55, 1988.

[19] Leslie Lamport. Paxos made simple. *ACM SIGACT News (Distributed Computing Column) 32, 4 (Whole Number 121, December 2001)*, pages 51–58, 2001.

[20] Leslie Lamport and Mike Massa. Cheap paxos. In *International Conference on Dependable Systems and Networks, 2004*, pages 307–314. IEEE, 2004.

[21] Tanakorn Leesatapornwongsa, Jeffrey F Lukman, Shan Lu, and Haryadi S Gunawi. Taxdc: A taxonomy of non-deterministic concurrency bugs in datacenter distributed systems. In *Proceedings of the twenty-first international conference on architectural support for programming languages and operating systems*, pages 517–530, 2016.

[22] Barbara Liskov. Practical uses of synchronized clocks in distributed systems. In *Proceedings of the tenth annual ACM symposium on Principles of distributed computing*, pages 1–9, 1991.

[23] Iulian Moraru, David G Andersen, and Michael Kaminsky. There is more consensus in egalitarian parliaments. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 358–372, 2013.

[24] Iulian Moraru, David G Andersen, and Michael Kaminsky. Paxos quorum leases: Fast reads without sacrificing writes. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 1–13, 2014.

[25] Faisal Nawab, Divyakant Agrawal, and Amr El Abbadi. Dpaxos: Managing data closer to users for low-latency and mobile applications. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1221–1236, 2018.

[26] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX annual technical conference (USENIX ATC 14)*, pages 305–319, 2014.

[27] Fred B Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys (CSUR)*, 22(4):299–319, 1990.

[28] Sarah Tollman, Seo Jin Park, and John Ousterhout. {EPaxos} revisited. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 613–632, 2021.

[29] Carl A Waldspurger and William E Weihl. Lottery scheduling: Flexible proportional-share resource management. In *Proceedings of the 1st USENIX conference on Operating Systems Design and Implementation*, pages 1–es, 1994.

[30] Michael Whittaker, Aleksey Charapko, Joseph M Hellerstein, Heidi Howard, and Ion Stoica. Read-write quorum systems made practical. In *Proceedings of the 8th Workshop on Principles and Practice of Consistency for Distributed Data*, pages 1–8, 2021.

[31] Jian Yi, Qing Li, Bin Zhang, Yong Jiang, Dan Zhao, Yuan Yang, and Zhenhui Yuan. Gleaning the consensus for linearizable and conflict-free per-replica local reads. In *Proceedings of the 7th Asia-Pacific Workshop on Networking*, pages 143–149, 2023.