

Minimizing Network and Storage Costs for Consensus with Flexible Erasure Coding

Mi Zhang
zhangmi@ict.ac.cn
State Key Lab of Processors,
Research Center for Advanced
Computer Systems,
ICT, CAS

Qihan Kang
kangqihan17@mails.ucas.ac.cn
State Key Lab of Processors,
Research Center for Advanced
Computer Systems,
ICT, CAS, UCAS

Patrick P. C. Lee
pcee@cse.cuhk.edu.hk
The Chinese University of Hong Kong

ABSTRACT

Consensus protocols like Paxos and Raft provide data consistency and fault tolerance for upper-layer distributed services. Log replication in these protocols can be supported by erasure coding, which incurs a lower redundancy ratio than full-copy replication and hence significantly saves network and storage costs for overall performance improvements. However, existing consensus protocols with erasure coding cannot achieve the minimum network and storage costs during log replication. Our observation is that the optimal coding scheme varies with the number of healthy servers in a group, such that the coding scheme with the lowest redundancy ratio in normal cases incurs more network traffic and storage overhead for log replication in the presence of server failures. To this end, we propose FlexRaft, which dynamically adjusts the coding scheme used in Raft based on the server status to always achieve the theoretically minimum redundancy ratio, while maintaining the same liveness as in the original Raft. To address the issue of an inconsistent coding scheme between the leader and its followers, we specify the prerequisite of overwriting a log entry, and also allow the leader and its followers to exactly track the coding scheme used. We further consider how to handle server failures and prove the safety of FlexRaft. We implement a prototype of FlexRaft, atop which we build a distributed key-value store to show its efficacy. Experiments on Alibaba Cloud show that FlexRaft achieves the theoretically minimum network and storage costs in practice, and reduces the commit latency by 44.51% and 19.37% compared with state-of-the-art CRaft and HRaft, respectively.

CCS CONCEPTS

• Computer systems organization → Redundancy; • Information systems → Distributed storage.

KEYWORDS

Raft, Erasure coding, Key-value store

ACM Reference Format:

Mi Zhang, Qihan Kang, and Patrick P. C. Lee. 2023. Minimizing Network

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 the author(s) 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.

ICPP 2023, August 07–10, 2023, Salt Lake City, UT, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0843-5/23/08...\$15.00

<https://doi.org/10.1145/3605573.3605619>

and Storage Costs for Consensus with Flexible Erasure Coding. In *52nd International Conference on Parallel Processing (ICPP 2023)*, August 07–10, 2023, Salt Lake City, UT, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3605573.3605619>

1 INTRODUCTION

Consensus protocols coordinate multiple servers to provide reliable distributed services [16, 17, 21]. As failures are commonplace in distributed environments (e.g., server crashes, network partitioning, and message loss), consensus protocols act as a foundation for building distributed systems with high availability and strong consistency. To protect distributed services against failures, consensus protocols typically replicate commands across multiple servers, so that distributed systems can operate correctly when a minority of servers fail. Paxos [16, 17] and Raft [21] are two widely-adopted consensus protocols in practical distributed systems [6, 9, 23, 25].

To guarantee strong consistency, consensus protocols record data operations as log entries and replicate them to all servers in a group. Each server maintains a log consisting of a sequence of commands in the same order, such that the state machines running in different servers can execute the same commands and output the same results. To tolerate F server failures, consensus protocols need to replicate log entries in $N = (2F + 1)$ servers. This incurs $2F$ times the network traffic from the leader to its followers and N times the storage cost as each server stores a full copy of the log entries. The *redundancy ratio* (i.e., the actual data size divided by the original data size) of full-copy replication is N times. Thus, the high redundancy ratio hinders the distributed systems to achieve low latency and high throughput.

Erasure coding is a well-known low-redundancy approach to achieving fault tolerance in storage systems. Reed-Solomon (RS) codes [22] are the most popular erasure codes deployed in practice. An $RS(k, m)$ code divides a data object into k fixed-size *data chunks*, and performs encoding to generate m *parity chunks*; note that $RS(k, m)$ is equivalent to full-copy replication when $k = 1$. The $k + m$ *coded chunks* (i.e., data and parity chunks) are stored in $k + m$ servers. When a server failure occurs, any k out of the $k + m$ coded chunks suffice to reconstruct the original content. The redundancy ratio of $RS(k, m)$ is $(k + m)/k$, which is only $1/k$ times compared to full-copy replication. Thus, erasure coding significantly reduces the network traffic and storage overhead.

Recent studies apply erasure coding into consensus protocols to save network and storage costs for high performance [14, 20, 27]. RS-Paxos [20] is the first work that adopts erasure coding to replicate the log entries in Paxos. To maintain safety (i.e., never return an

incorrect result), RS-Paxos stores the chunks of a log entry in at least $F + k$ servers, such that the original data can be recovered from the remaining servers under F server failures. RS-Paxos trades *liveness level* (i.e., the number of tolerable failures) for better performance. To maintain liveness level F as in the original Raft, CRaft [27] converts to full-copy replication when there are fewer than $(F + k)$ healthy servers. HRaft [14] optimizes CRaft by replenishing some coded chunks in several healthy servers instead of switching to full-copy replication directly.

However, existing erasure-coded consensus protocols do not specify the optimal *coding scheme* (i.e., the choices of k and m) to minimize both network and storage costs during log replication, thereby failing to achieve the lowest redundancy ratio. The redundancy ratio $(k + m)/k$ depends on the value of k , as the value of $k + m$ is equal to the number of servers in a group (i.e., N) where each server stores a coded chunk of a log entry. When all servers in a group function correctly, we can choose the largest available value of k for log replication to minimize the network and storage costs. However, a larger k also implies that more servers (i.e., $F + k$) are required to store coded chunks for entry commitment. If we use the coding scheme with the largest k for log replication at the beginning, the network and storage costs cannot stay at the minimum when server failures occur. In the presence of server failures, neither CRaft nor HRaft achieves the minimum redundancy ratio, as CRaft degrades to full-copy replication while HRaft needs to store additional coded chunks in the remaining servers. Thus, existing erasure-coded consensus protocols cannot maintain the lowest redundancy ratio for entry commitment all the time.

Our insight is that the optimal coding scheme for entry commitment in consensus protocols has a different value of k depending on the number of healthy servers in a group. For example, as shown in Table 1 in §2.3, the optimal coding scheme (marked with a ★) for a Raft group of five servers is RS(3,2) and RS(2,3) when there is zero and one server failure, respectively. We argue that the coding scheme should be *dynamically* adjusted based on the latest server status (which can be estimated by heartbeat messages between the leader and the followers in Raft [21]). In contrast, both CRaft and HRaft employ a fixed coding scheme regardless of the server status.

In this paper, we propose a flexible erasure coding approach for Raft called FlexRaft, which *provably* minimizes both network and storage costs to commit log entries. FlexRaft varies the coding scheme and maintains the lowest redundancy ratio in the presence of server failures. It sets k as large as possible based on the number of healthy servers in a group. However, the varied coding scheme requires overwriting the coded chunks of a log entry with the new ones, thereby raising new consistency issues between the leader and its followers: 1) for the same log entry, the chunk stored in a follower can be mistakenly overwritten by an old coded chunk, and 2) some followers may not successfully update the old chunks before the log entry is committed. Thus, FlexRaft specifies the prerequisite of overwriting a log entry, and modifies the *AppendEntries RPCs* (i.e., the commands initiated by the leader to replicate log entries (with arguments) or provide heartbeats (without arguments)) to ensure that the chunks stored in the followers are encoded with the same coding scheme as the leader. We consider how to recover a failed server when adopting a varied coding scheme and prove the safety of FlexRaft. We summarize our contributions as follows.

- We analyze the optimal coding scheme for consensus protocols under a different number of server failures. We show that given N' healthy servers in a Raft group ($N' \leq N$), the coding scheme with $k = N' - F$ achieves the lowest redundancy ratio while keeping the liveness level F .
- We propose FlexRaft to dynamically vary the coding scheme used for log replication in the Raft protocol based on the server status. To guarantee that the leader and its followers use the same coding scheme for a log entry, FlexRaft adds a restriction rule to avoid mistakenly updating a coded chunk in a follower, and also updates the *AppendEntries RPCs* so that the leader can exactly track the coding scheme of the stored chunks.
- We handle server failure recovery when using a varied coding scheme, and prove that FlexRaft guarantees Raft safety while maintaining the liveness level F .
- We implement FlexRaft in C++ and build a distributed key-value store using RocksDB [10] atop FlexRaft. Experiments on Alibaba Cloud [2] show that FlexRaft reduces the commit latency by 44.51% and 19.37%, respectively, compared with CRaft and HRaft under two server failures for a group of seven servers.

The source code of our FlexRaft prototype is now available at: <https://github.com/ACS-Storage-Group/FlexRaft-Code>.

2 BACKGROUND AND MOTIVATION

2.1 Basics of Raft Consensus

We first provide the background details of Raft [21]. We consider a Raft group of $N = (2F + 1)$ servers that can tolerate any F server failures ($F \geq 1$). Each server is in one of the following three states: *leader*, *follower*, and *candidate*. In a normal situation, there is one leader in a Raft group and the remaining servers are followers. The leader handles all client requests and replicates *log entries* (i.e., the operations being executed); the followers respond to the requests from the leader and candidates, and redirect the client requests to the leader; the candidate state is used to elect a new leader. Raft divides time into *terms*, numbered with consecutive integers, and each term starts with a *leader election*.

Raft adopts a strong form of leadership to simplify the management of log replication. The leader accepts log entries from clients and replicates them to other servers in the same Raft group by sending *AppendEntries RPCs*. Each log entry tracked by an *index* (a monotonically increasing number) stores an update to the state machine along with the term number when the leader receives the update. When a log entry is replicated to a majority of servers, the leader commits and applies the log entry and its previous log entries, and informs the followers to apply the log entries. Raft guarantees the *leader completeness* property, meaning that the leader at any term has all committed entries in the previous terms.

2.2 Erasure-Coded Consensus Protocols

RS-Paxos [20] is the first protocol to combine erasure coding and consensus to reduce both network and storage costs. It divides the original data of a log entry into k chunks with equal sizes and encodes the k chunks into m parity chunks ($k > 1, m \geq 1$) using RS(k, m). It then sends one chunk to each acceptor for log replication. Any k chunks of data and parity chunks can reconstruct the original log entry. As the chunk size is a fraction (i.e., $1/k < 1$)

of the total data size, RS-Paxos saves network bandwidth cost and disk IOs. RS-Paxos is actually a superset of the vanilla Paxos [16, 17]. It requires that the intersection between the write quorum Q_w and read quorum Q_r should be not less than k to guarantee safety; that is, RS-Paxos should satisfy $Q_r + Q_w - N \geq k$ when using $RS(k, m)$ for storing log entries, such that at least k chunks can be read from Q_r servers to recover the original data. With a larger k , it achieves more network and storage savings, but puts a higher requirement on the write and read quorum. Thus, compared to the vanilla Paxos using full-copy replication, RS-Paxos tolerates fewer failed servers ($< F$) in a group of $N = (2F + 1)$ servers.

CRaft [27] extends Raft with erasure coding like RS-Paxos, while maintaining the same *liveness level* (i.e., the number of failed servers tolerable by a protocol) as Raft. To address the liveness problem of RS-Paxos (i.e., tolerating fewer than F failures), CRaft uses erasure coding and full-copy replication jointly. When at least $F + k$ servers are running normally in a Raft group, CRaft uses $RS(k, m)$ for log replication to reduce network and storage costs; otherwise, it switches to full-copy replication and keeps the liveness level F (i.e., CRaft switches to full-copy replication when the number of failed servers is larger than $N - F - k$). Although CRaft maintains the same liveness as the original Raft, switching to full-copy replication for log replication causes sharp performance degradation when the number of healthy servers reduces to less than $F + k$. When a leader is newly elected in CRaft, the new leader performs a *LeaderPre* operation to deal with any *unapplied* coded chunk (i.e., not yet applied to the state machine) before the leader fully functions. During the *LeaderPre* operation, the leader attempts to recover any unapplied entries in sequence by collecting other coded chunks from its followers, and deletes the entries that cannot be recovered.

HRaft [14] mitigates the sharp performance degradation in CRaft by replenishing some coded chunks in several healthy servers when failures occur. It adjusts the placement of the coded chunks and replicates some coded chunks to healthy servers instead of switching to full-copy replication if the number of failed servers is greater than $N - F - k$. When the leader receives p ($F \leq p \leq F + k - 1$) acknowledgments during log replication (i.e., there are $(N - 1 - p)$ failed servers), it chooses $(F + k - 1 - p)$ coded chunks and replicates them to F healthy servers before committing a log entry. In other words, some servers store multiple coded chunks of a log entry to guarantee safety (i.e., retrieving enough chunks for data reconstruction). Although HRaft avoids switching to full-copy replication, it cannot always maintain the minimum storage and network costs during log replication.

2.3 Motivating Example

We compare the network and storage costs of CRaft and HRaft with different erasure coding schemes under a different number of server failures. We use a group consisting of five servers ($N = 5$), which can tolerate at most two server failures ($F = 2$). The *network cost* is the bandwidth usage from the leader to its followers for a log entry to be committed. As the leader stores a full copy of the log, the *storage cost* is equal to the network cost plus one [27].

Table 1 shows the network and storage costs of CRaft and HRaft with different coding schemes where k ranges from 1 to 3. Here, we denote full-copy replication by $k = 1$. In the normal case that no server fails (i.e., $f = 0$), CRaft and HRaft introduce the same network

Table 1: Performance of CRaft and HRaft using different erasure coding schemes for a group of five servers ($N = 5$).

(a) No server failure ($f = 0$)						
k	Log Replication		Network Cost		Storage Cost	
	CRaft	HRaft	CRaft	HRaft	CRaft	HRaft
1	full-copy	full-copy	4	4	5	5
2	RS(2,3)	RS(2,3)	2	2	3	3
★ 3	RS(3,2)	RS(3,2)	4/3	4/3	7/3	7/3

(b) One server failure ($f = 1$)						
k	Log Replication		Network Cost		Storage Cost	
	CRaft	HRaft	CRaft	HRaft	CRaft	HRaft
1	full-copy	full-copy	3	3	4	4
★ 2	RS(2,3)	RS(2,3)	3/2	3/2	5/2	5/2
3	full-copy	RS(3,2)+	3	5/3	4	8/3

(c) Two server failures ($f = 2$)						
k	Log Replication		Network Cost		Storage Cost	
	CRaft	HRaft	CRaft	HRaft	CRaft	HRaft
★ 1	full-copy	full-copy	2	2	3	3
2	full-copy	RS(2,3)+	2	2	3	3
3	full-copy	RS(3,2)+	2	2	3	3

We label the optimal coding scheme (used by FlexRaft) under a different number of server failures with a ★. $RS(k, m) +$ in HRaft means that it replenishes some coded chunks when using $RS(k, m)$.

and storage costs when using the same coding scheme. Both CRaft and HRaft achieve the minimum network and storage costs with the largest k (i.e., $k = 3$) among all three coding settings. When using $RS(3,2)$ for log replication, the network and storage costs are $4/3$ and $7/3$, respectively. If one server fails (i.e., $f = 1$), CRaft with $k = 3$ switches to full-copy replication as the number of healthy servers in the group is less than $(F + k)$ (i.e., 5). The network and storage costs of CRaft sharply increase to 3 and 4, respectively, the same as the original Raft. Under one server failure, HRaft with $k = 3$ still uses $RS(3,2)$ coding scheme while storing two coded chunks in two servers (denoted by $RS(3,2) +$). The network and storage costs of HRaft gradually increase to $5/3$ and $8/3$, respectively, lower than those of CRaft with $k = 3$ but higher than CRaft/HRaft with $k = 2$. For $k = 2$, both CRaft and HRaft can use $RS(2,3)$ for log replication, which incurs the minimum network and storage costs (i.e., $3/2$ and $5/2$, respectively). When two servers fail (i.e., $f = 2$), CRaft with $k = 2$ and $k = 3$ switches to full-copy replication, while HRaft needs to store more coded chunks to keep the liveness level F . For CRaft and HRaft under two server failures, the network and storage costs are 2 and 3, respectively, equal to that of the original Raft. Thus, the largest k does not guarantee the minimum network and storage costs in all cases; for example, $RS(2,3)$, rather than $RS(3,2)$, incurs the lowest network and storage costs when one server fails. Although HRaft avoids switching to full-copy replication in the presence of server failures, it fails to achieve the minimum network and storage costs using a coding scheme with a fixed k .

3 DESIGN OF FLEXRAFT

FlexRaft dynamically adjusts the coding scheme for log replication according to the cluster status, so as to minimize both the network and storage costs while maintaining the liveness level F . We first describe the choice of coding schemes in FlexRaft (§3.1). We then explain the issues raised by varying coding schemes and how FlexRaft

addresses them (§3.2). Finally, we introduce how FlexRaft deals with server recovery (§3.3) and provides the safety guarantee (§3.4).

3.1 Choice of Coding Schemes

We first define the terminologies and notations. We consider a Raft group of N servers, where $N = 2F + 1$, so as to tolerate up to F failures. Let N' be the total number of *healthy servers* in the group (i.e., the servers are alive and can communicate with each other and the clients [27]), and f be the number of server failures (i.e., $N' = N - f$). When a Raft group starts, all servers are healthy and the number of healthy servers is equal to the total number of servers in the group (i.e., $N' = N$ and $f = 0$). The number of healthy servers decreases when some servers crash or lose network connection, or increases when the failed servers rejoin the group or some new servers are added to replace the failed servers. Currently, we consider a Raft group of a fixed size where N does not change. The Raft leader can estimate the number of healthy servers and update the value of N' based on the exchanges of the latest heartbeat messages with its followers.

Coding scheme chosen by FlexRaft. To minimize both the network and storage costs, FlexRaft dynamically chooses the coding scheme based on the number of healthy servers remaining in a group. Given N' healthy servers, FlexRaft uses $RS(k, m)$ code where $k = N' - F$ (i.e., $N - f - F$) and $m = N - k$ for log replication. Here, we make $k + m = N$, so that each server in the group can store a coded chunk (either original data chunk or parity chunk) of a log entry. Note that $k = 1$ means that FlexRaft uses full-copy replication. FlexRaft chooses the available largest value of k (i.e., $N' - F$) given the number of healthy servers N' , so as to replicate a log entry with the lowest redundancy ratio. We prove that FlexRaft always minimizes the network and storage costs for log replication below.

THEOREM 1. *When there are N' ($F + 1 \leq N' \leq N$) healthy servers in a Raft group, FlexRaft always minimizes the network and storage costs for log replication.*

PROOF. To maintain safety, at least $F + k$ servers should store the chunks of a log entry before it can be committed, such that there are at least k chunks in any $F + 1$ servers [20, 27]. When N' healthy servers exist in a Raft group, at most N' servers store the chunks of a log entry to commit (i.e., $F + k \leq N'$). Thus, Raft can use $RS(k, m)$, where $k \leq N' - F$, for log replication. The redundancy ratio of $RS(k, m)$ is N'/k . As the network and storage costs decrease with the reduction of redundancy ratio, a larger value of k incurs lower storage and network overhead given a fixed N' . Thus, FlexRaft always minimizes the network and storage costs for log replication by using the largest available k (i.e., $N' - F$). \square

Table 2 shows the coding schemes, log replication methods, network cost, and storage cost of FlexRaft under a different number of server failures where f increases from 0 to F . Under f server failures, FlexRaft uses a coding scheme with $k = N - f - F$ for log replication where the network cost is $(N - f - 1)/k$ and the storage cost is $(N - f - 1)/k + 1$. Note that the network and storage costs increase with the number of server failures. When F servers fail, the only method for log replication is full-copy replication.

Comparison to CRaft and HRaft. FlexRaft always achieves the minimum redundancy ratio, introducing lower network and storage costs than CRaft and HRaft. When the number of healthy servers N'

Table 2: Coding scheme used in FlexRaft to minimize network and storage costs under a different number of server failures.

f	k	Log Replication	Network Cost	Storage Cost
0	$N - F$	$RS(N - F, F)$	$\frac{N-1}{N-F}$	$\frac{N-1}{N-F} + 1$
1	$N - F - 1$	$RS(N - F - 1, F + 1)$	$\frac{N-2}{N-F-1}$	$\frac{N-2}{N-F-1} + 1$
2	$N - F - 2$	$RS(N - F - 2, F + 2)$	$\frac{N-3}{N-F-2}$	$\frac{N-3}{N-F-2} + 1$
...
F	1	full-copy	$N - F - 1$	$N - F$

is equal to or greater than $F + k$, all protocols can employ $RS(k, m)$ for log replication, where the network cost is $(N' - 1)/k$. Suppose that we have Δf additional failed servers, such that the number of healthy servers N' now less than $F + k$ (i.e., $N' = F + k - \Delta f$). The three protocols take different approaches to maintain the liveness level F : i) CRaft converts to full-copy replication, where the network cost increases to $2F - f$ [27]; ii) HRaft keeps using $RS(k, m)$ and stores additional copies of some coded chunks, where the network cost is $(2F - f + F(f - N + F + k))/k$ [14]; iii) FlexRaft reduces the value of k based on the failure status and uses $RS(k - \Delta f, m)$ (for $k - \Delta f \geq 1$), where the network cost is $(F + k - \Delta f - 1)/(k - \Delta f)$. Compared to the network cost when there are $(F + k)$ healthy servers (i.e., $(F + k - 1)/k$), the increased cost of HRaft is $\Delta f(F - 1)/k$, while that of FlexRaft is $\Delta f(F - 1)/(k(k - \Delta f))$. Since we have $k - \Delta f \geq 1$, the increment of FlexRaft is less than that of HRaft. Also, as HRaft incurs less network traffic than CRaft [14], the network cost of FlexRaft is less than that of CRaft as well. Moreover, as the storage cost equals the network cost (from the leader to its followers) plus the storage cost of a full copy in the leader, the storage cost of FlexRaft is also less than those of CRaft and HRaft.

3.2 Varied Coding Scheme

The main idea of FlexRaft is to choose the optimal coding scheme for log replication based on the cluster status. When failures occur during writes, FlexRaft adjusts the value of k if there are not enough servers storing the log entries.

Varying coding scheme during log replication. When FlexRaft decides to vary the coding scheme for a log entry, the leader needs to re-encode the entry with the new coding scheme and send the new chunks to its followers. At the beginning of replicating a log entry, FlexRaft first performs encoding on the data using an initial coding scheme $RS(k, m)$, where k is determined by the latest N' (§3.1). Then the leader distributes the coded chunks to the followers and waits for the responses from the followers. Here, we map the chunk ID to the server ID consistently to determine which server should store which chunk for a log entry. When a follower receives a chunk from the leader, the follower appends the chunk to its log and returns a successful response to the leader. The leader collects the responses from the followers and decides how to perform log replication. If the leader receives at least $F + k - 1$ successful responses, it continues to commit the log entry; otherwise, the leader varies the coding scheme to maintain safety since there are fewer than $F + k - 1$ healthy followers for storing the log entry. FlexRaft then updates N' based on the responses from the followers and switches from the initial coding scheme to the new coding scheme $RS(k', m')$ where $k' = N' - F$ and $m' = N - k'$. The leader re-encodes the log

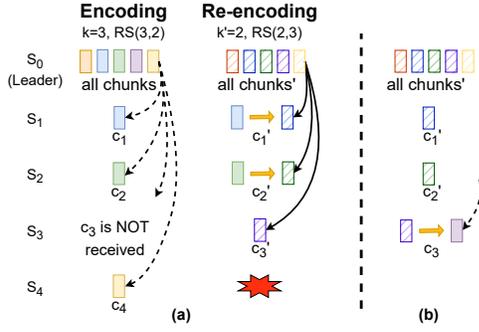


Figure 1: The coding scheme varies during log replication.

entry using the new coding scheme, and sends the new chunks to its followers. Each follower overwrites the old chunks with the new chunks. If the leader receives at least $F + k' - 1$ successful responses, it can commit the log entry; otherwise, the leader continues to adjust the coding scheme by decreasing the value of k and performs the above process until it successfully replicates the log entry. The smallest value of k is one (i.e., full-copy replication).

Figure 1(a) shows how the coding scheme varies during log replication in a Raft group of five servers. At the beginning, the leader S_0 encodes the log entry using RS(3,2) and sends four chunks to its followers. However, the leader only receives three responses from servers S_1 , S_2 , and S_4 , since S_3 does not receive any chunk (e.g., due to network failures). The leader S_0 then varies the coding scheme to RS(2,3), re-encodes the log entry, and sends the new chunks to its followers. The followers S_1 and S_2 overwrite the old chunk with the new chunk, S_3 stores the new chunk directly, and S_4 does not store the new chunk as it now crashes. After receiving successful responses from servers S_1 , S_2 , and S_3 , the leader S_0 can commit this log entry safely. However, varying the coding scheme raises new challenges to the correctness of log replication.

Challenge 1: the chunk in a follower can be mistakenly overwritten by an old chunk of the same log entry. As the coding scheme varies during log replication in FlexRaft, the servers in a group may store the chunks encoded with different coding schemes for a log entry. Thus, the chunk stored in a follower may be an old one due to chunk overwrites. For example, as shown in Figure 1(b), the followers S_1 , S_2 , and S_3 are expected to store the new chunks c'_1 , c'_2 , and c'_3 , respectively, after the re-encoding with RS(2,3). However, due to network delays, the request containing the old chunk c_3 reaches the follower S_3 after S_3 stores the new chunk c'_3 . S_3 then overwrites c'_3 with c_3 , but the leader has received a successful response from S_3 that it stores c'_3 . In this case, the chunk stored in a follower is inconsistent with that stored in the leader; in other words, a log entry is now encoded by different coding schemes. Thus, FlexRaft should guarantee that the chunks of a log entry stored in the followers are encoded with the same coding scheme as specified by the leader.

Prerequisite of overwriting a log entry. To avoid mistakenly overwriting a chunk in a follower, FlexRaft requires the follower to check the value of k with the chunk before overwriting a log entry: if the value of k with the new chunk is less than the current value of k , the follower overwrites the current chunk with the new one and responds to the leader; otherwise, the follower can ignore

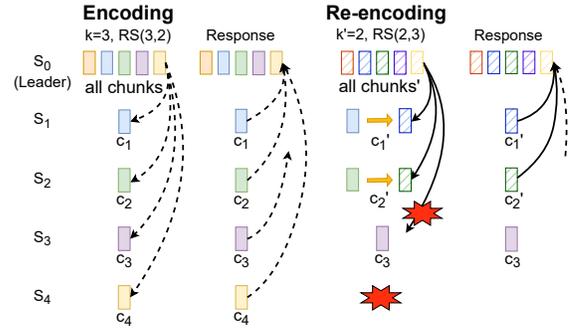


Figure 2: Responses of varied coding schemes during log replication.

this message containing the new chunk. The reason is that the value of k always decreases when the coding scheme varies during log replication, and the new chunk to overwrite is encoded from a coding scheme with a smaller k . Thus, in Figure 1, if the follower S_3 compares the value of k with the chunk c_3 (i.e., $k = 3$) to the one with the chunk c'_3 (i.e., $k = 2$), it will refuse to update the current chunk and the above inconsistent case will not occur.

Challenge 2: some followers may not overwrite the old chunks successfully before the leader commits the log entry. Although each follower checks the value of k before overwriting its chunks, the chunks stored in the follower may still use a different coding scheme from the leader's. One possible case is that the leader requires the followers to overwrite their existing old chunks, yet some followers fail to replace the old chunks with the new chunks. Figure 2 depicts such an example in a group of five servers ($N = 5$). The leader S_0 encodes a log entry with RS(3,2) and sends the coded chunks to its followers. All followers store the chunks successfully and return successful responses to the leader. However, the follower S_3 's response does not reach the leader in time (e.g., due to network failures). The leader S_0 only receives three successful responses from S_1 , S_2 , and S_4 before it decides to vary the coding scheme. Thus, the leader S_0 re-encodes the log entry with RS(2,3) and distributes the new chunks to the followers. As the AppendEntries RPC from the leader to the follower S_3 is lost and the follower S_4 crashes this time, only the followers S_1 and S_2 overwrite the old chunks with the new chunks and return successful responses to the leader again. However, the leader S_0 mistakenly thinks that three followers have stored the new chunks and treats log replication as successful, since it also receives the last successful response from S_3 . That is, the new chunks being re-encoded are stored in only two followers, while S_3 still stores the old chunk. The main reason is that the successful response from a follower to the leader in Raft, which consists of the currentTerm and a true flag if the follower contains the entry matching the index of the previous log entry (i.e., preLogIndex) and the term of prevLogIndex entry (i.e., prevLogTerm), fails to indicate which server stores which coded chunk. Thus, the leader cannot distinguish the exact number of stored chunks belonging to the same coding scheme for a log entry from the responses of the followers.

Updating AppendEntries RPCs. To make a follower's log stay consistent with the leader's, we add *prevK* (i.e., the value of k of the prevLogIndex entry) and *ChunkInfo* (i.e., a tuple of the log index and the value of k) to the AppendEntries RPCs. We include *prevK* in the

request of the AppendEntries RPCs to allow a follower to detect the chunks encoded with a different coding scheme. When receiving an AppendEntries RPC request, the follower checks the value of prevK after checking the term and prevLogTerm. If the follower finds that the value of k of the entry in prevLogIndex does not match prevK (i.e., storing a chunk encoded with an old coding scheme), the follower can return false; the leader then decrements nextIndex and retries to send AppendEntries RPCs. For example, when S_3 rejoins the group as a follower, it can compare prevK and find that the old chunk c_3 should be overwritten. Moreover, we include the ChunkInfo of the log entries in the requests and responses of the AppendEntries RPCs. Thus, the leader sends log entries containing ChunkInfo (included in each log entry) to the followers and each follower responds to the leader with the ChunkInfo of the chunks stored, such that the leader can determine whether the chunks encoded with the latest coding scheme have been stored in enough followers. For example, by checking the ChunkInfo in the received responses, the leader S_0 in Figure 2 can find that one response should not be accounted as a successful response to replicate the new chunks, as the chunk c_3 rather than the new chunk c'_3 is stored in S_3 . Thus, the leader will decrease the value of k (i.e., reducing to full-copy replication) and retry to replicate the log entry. Thus, by adding prevK and ChunkInfo to the AppendEntries RPCs, FlexRaft guarantees that the chunks belonging to a committed log entry are correctly stored in the followers using the same coding scheme.

3.3 Server Recovery

Handling leader failures. When the current leader fails, the process of leader election starts and a new leader is elected. For the committed log entries, the newly elected leader has at least one chunk according to Raft election rules [21], and the leader can recover the complete log entry by collecting other chunks from its followers since we require that any $F + 1$ servers store k chunks. For the unapplied entries, the new leader should perform the LeaderPre operation (§2.2) before it can become a fully-functioning leader [27]. Note that the chunks stored in the new leader may not use the same coding scheme when the log entries are committed, even though the leader’s log is up-to-date. To deal with this issue, the leader recovers the entries among (commitIndex, lastLogIndex), i.e., the entries between the highest index of log entry known to be committed and the index of the last log entry, and replaces the chunks with the complete log entries during the LeaderPre operation. As the leader stores the ChunkInfo for each log entry, the leader can decide whether a log entry is recoverable based on the value of k . For those entries that cannot be recovered, the leader removes them from the log. The leader can process the client requests normally after the LeaderPre operation.

A newly elected leader may need to perform the decoding operation when serving read requests to an object. If the leader only has a coded chunk for a data object to be read, the leader needs to retrieve other coded chunks to decode the original object. The leader then stores a full copy of the recovered object for further reads to avoid decoding every time. That is, the leader handles multiple reads to an object with at most one decoding operation. Note that CRaft and HRaft perform the same read procedure as FlexRaft. These Raft protocols with erasure coding make a trade-off between the network and storage costs during commitment and

the computation overhead during reads.

Recovering a failed follower. When a server resumes normally, FlexRaft restores the committed log entries with the same coding scheme (i.e., the same value of k). For a committed log entry, the recovery cost of FlexRaft is $1/k$ of the original data size. FlexRaft guarantees that the recovered server stores all previous log entries before it is counted as a healthy server to function. After the recovery completes, FlexRaft increases the number of healthy servers (i.e., N') accordingly, and varies the coding scheme with $k = N' - F$ to replicate log entries later.

Discussion. The minimum redundancy ratio is achieved with $RS(F + 1, F)$ code when all servers in the group are healthy. Although FlexRaft minimizes the network and storage costs during log replication, the log entries committed in the presence of failures are encoded with a smaller k than $F + 1$. To further reduce the storage cost, we can convert the log entries stored with a smaller k to the $RS(F + 1, F)$ code when all the servers are healthy ($N' = N$). It needs to re-encode log entries and update the log entries in all servers, which inevitably introduces a large amount of network traffic. Such code conversion actually trades the network bandwidth cost for storage savings. FlexRaft currently focuses on minimizing the network and storage costs to reduce the commit latency. We leave the code conversion as future work.

3.4 Safety Guarantee

We show that FlexRaft guarantees Raft safety by proving the Log Matching Property and the Leader Completeness Property.

THEOREM 2. *Log Matching Property: if two logs contain an entry with the same index and term, then the logs are identical (either a full copy or a coded chunk of the original proposed data) in all entries up through the given index.*

PROOF. The original Raft protocol maintains two properties to constitute the Log Matching Property: 1) if two entries in different logs have the same index and term, then they store the same command, and 2) if two entries in different logs have the same index and term, then the logs are identical in all preceding entries. The first property stands because a leader creates at most one entry with a given log index in a given term, and log entries never change their position in the log. FlexRaft works as in Raft, although it varies the coding scheme when the server status changes. The coded chunk stored with the same index and term may use an old coding scheme (which will be updated during processing AppendEntries RPCs and the LeaderPre phase), but stands for the same command. The second property in Raft is guaranteed by a simple consistency check performed by the AppendEntries RPCs. FlexRaft also follows the AppendEntries RPC rules in the original Raft protocol. When the coding scheme varies for the entry with the same index and term, FlexRaft adds a restriction rule to avoid mistakenly overwriting a log entry, and includes prevK and ChunkInfo in the AppendEntries RPCs to make the followers use the same coding scheme as the leader. Thus, the Log Matching Property still holds where the data is a coded chunk when using erasure coding for the log entry. □

THEOREM 3. *Leader Completeness Property: if a log entry is committed in a given term, then that entry will be present in the logs of the leaders for all higher-numbered terms.*

PROOF. As FlexRaft extends the LeaderPre operation in CRaft, CRaft guarantees that if a log entry e is committed in a given term, then e will be present in the logs of the leaders for all higher-numbered terms, and e will not be deleted in any higher-numbered term’s LeaderPre [27]. We then prove that the committed data is recoverable using the right coding scheme in FlexRaft. For each committed entry, FlexRaft guarantees that there are at least $F + k$ coded chunks stored in the cluster. Thus, the leader can always recover a committed entry by collecting k chunks from other servers. Therefore, the Leader Completeness Property stands in FlexRaft. \square

4 IMPLEMENTATION

We implement a prototype of FlexRaft from scratch in C++. We leverage the ISA-L library [13] for erasure coding operations and use RCF 3.0 [8] for interprocess communication. To evaluate FlexRaft, we also build a distributed key-value store based on RocksDB v7.3.1 [10]. Each server has a constantly running FlexRaft module, a RocksDB engine as a state machine, and a background working thread to apply committed entries to the state machine. The whole system contains about 6K LoC.

Log entry. We modify the structure of the log entry by adding an *EntryType* flag and the *ChunkInfo*, a tuple of (log index, k). The *EntryType* flag in a log entry’s metadata indicates whether the command in this entry is a complete copy or a coded chunk. If this entry is encoded, the *ChunkInfo* identifies the coded chunk uniquely, where k is the encoding parameter used for this entry.

Log replication and commitment. The leader determines the coding scheme for log replication by counting the number of healthy servers. The leader sends heartbeat messages (i.e., empty *AppendEntries* RPCs) to its followers every 100 ms. To track the server status, the leader records the time point when it receives RPC messages (RPC requests and responses) from other servers. If the leader detects that some servers have not sent any messages for some time (e.g., 200 ms), the leader considers these servers as unhealthy and determines the encoding parameter by $k = N' - F$. The leader then encodes the original data into $k + m$ chunks, sends the coded chunk to each healthy follower using *AppendEntries* RPCs, and waits for responses from the followers. If the leader receives more than $F + k - 1$ success responses within a configured time limit (e.g., 1 second), the leader checks the *ChunkInfo* in the responses and commits this entry if at least $F + k$ servers (including itself) successfully store this entry. Otherwise, the leader decreases the encoding parameter k by one, re-encodes the entry, and repeats the above replication process until the log entry gets committed.

Processing in followers. Upon receiving an *AppendEntries* request, the follower checks the metadata of this entry to decide whether to store it or not. For the entry that has the same term and index as the latest one stored in the follower’s log, the follower overwrites the old entry using the newly-received one only if the new entry is encoded with a smaller k ; otherwise, the follower ignores this request. The follower overwrites the old entries by trimming its log and appending the new entries. If the log entry has a different term or index, the follower acts as the original Raft, e.g., appending the new entry. After storing the log entry, the follower returns a response containing the *ChunkInfo* to the leader.

5 EVALUATION

In this section, we show the I/O performance of the key-value store atop FlexRaft, the breakdown performance, the overhead, and the scalability of FlexRaft. We run all experiments on Alibaba Cloud [2]. The cluster consists of at most 11 servers to run key-value service and one individual server to send client requests. Each server is equipped with a Xeon CPU of 8 vCPU, 32 GiB DRAM, and 40 GiB ESSD cloud drive. In each experiment, each Raft server spawns two threads in two vCPUs to execute the Raft protocol and apply log entries, and starts a thread on-the-fly to handle any incoming RPC request. Each client sends read/write requests with a single thread that runs in a vCPU. The network bandwidth is 1 Gbps. We compare FlexRaft with CRaft and HRaft in the normal case and when some servers fail.

5.1 Write Performance

Normal write latency. We first compare the write latency of FlexRaft with CRaft and HRaft in normal cases when no server fails (**Exp#1-2**). The Raft group has five or seven servers (i.e., $N = 5$ or $N = 7$, respectively, where $N = 2F + 1$). For the coding scheme, we configure all possible values of k for CRaft and use the largest k for HRaft; FlexRaft chooses the optimal coding scheme automatically. For $N = 5$, CRaft can use RS(3,2) or RS(2,3) as its coding scheme; for $N = 7$, CRaft can use RS(4,3), RS(3,4), or RS(2,5). HRaft uses the largest value of k , i.e., RS(3,2) for $N = 5$ and RS(4,3) for $N = 7$. The key size is 16 B (same for the following experiments), and the value size varies from 4 KiB to 2 MiB. For each value size, the client generates 10000 PUT requests.

Figures 3 and 4 show the write latencies of CRaft, HRaft, and FlexRaft under different value sizes when $N = 5$ and $N = 7$, respectively. For $N = 5$, CRaft ($k = 3$) achieves lower latencies than CRaft ($k = 2$), reducing the latencies by 4.05-20.73% since CRaft ($k = 3$) saves the network bandwidth cost. FlexRaft and HRaft achieve the same performance as CRaft ($k = 3$) since both use the largest k to minimize the redundancy ratio. For $N = 7$, CRaft ($k = 4$) achieves the lowest latencies among all three configurations. Compared to CRaft ($k = 3$) and CRaft ($k = 2$), CRaft ($k = 4$) reduces the latencies by up to 18.15% and 35.32%, respectively. The latencies of FlexRaft and HRaft are close to that of CRaft ($k = 4$) under the same value size since they use the same coding scheme with the lowest redundancy ratio. Thus, FlexRaft minimizes the network and storage costs for log replication in normal cases.

Write performance under server failures. We evaluate the write performance under one server failure when $N = 5$ (**Exp#3**). We use the same configuration as Exp#1. Figure 5 plots the write latencies of different value sizes under one server failure. As the number of healthy servers is below $F + k$, CRaft ($k = 3$) converts to full-copy replication and HRaft replicates two additional coded chunks, incurring higher network and storage costs. In this case, FlexRaft varies the coding scheme to RS(2,3), achieving the lowest latencies as CRaft ($k = 2$). Compared to CRaft ($k = 3$) and HRaft, FlexRaft reduces the latencies by at most 32.82% and 8.43%, respectively.

We then measure the write performance under a different number of server failures when $N = 7$ (**Exp#4**). We use the same configuration as Exp#2. Figures 6a and 6b show the write latencies under one server failure and two server failures, respectively. When one

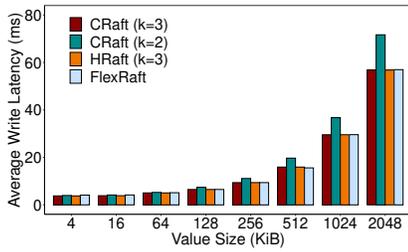


Figure 3: Exp#1 (Write latency in normal cases when $N=5$).

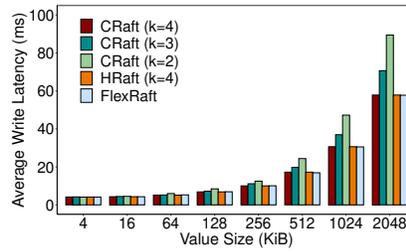


Figure 4: Exp#2 (Write latency in normal cases when $N=7$).

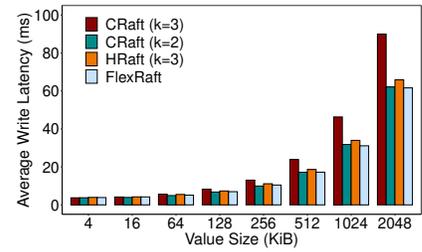
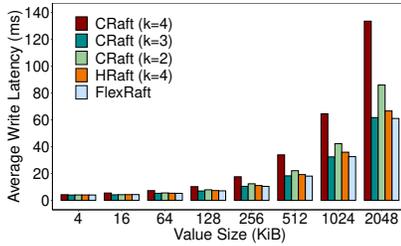
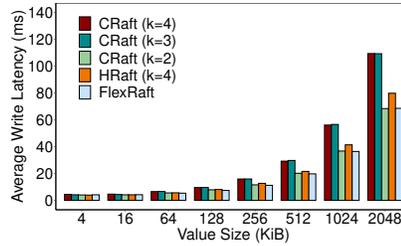


Figure 5: Exp#3 (Write latency under one server failure when $N=5$).



(a) One server failure



(b) Two server failures

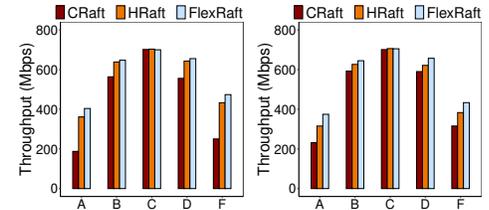
Figure 6: Exp#4 (Write latency under some server failures when $N=7$).

server fails, FlexRaft and CRaft ($k = 3$) achieve the lowest write latencies among all coding schemes, while CRaft ($k = 4$) has the highest latencies as it converts to full-copy replication. FlexRaft reduces the latencies of CRaft ($k = 4$), CRaft ($k = 2$), and HRaft by up to 54.28%, 28.97%, and 9.25%, respectively. When there are two failed servers, FlexRaft uses RS(2,5) for log replication to minimize the redundancy ratio and incurs the minimum network traffic and storage overhead. Both CRaft ($k = 4$) and CRaft($k = 3$) switch to full-copy replication while HRaft stores more coded chunks. Compared to CRaft ($k = 4$) and HRaft, FlexRaft reduces the write latencies by 37.32% and 14.10%, respectively when the value size is 2 MiB.

5.2 Performance Under YCSB Workloads

We compare the performance of FlexRaft to HRaft and CRaft under YCSB [7] workloads when $N = 7$ (Exp#5). We configure the largest value of k ($k = 4$) for CRaft and HRaft. We launch four client threads to generate 10000 requests following a Zipf distribution with a Zipfian constant of 0.99. Here we fix the value size as 2 MiB.

Figure 7 shows the throughput of CRaft, HRaft, and FlexRaft under YCSB workload A (50% reads, 50% updates), B (95% reads, 5% updates), C (100% reads), D (95% read-latest, 5% inserts), and F (50% reads, 50% read-modify-writes), when there is one failed server and two failed servers, respectively. Here, we do not show the results of workload E (95% scans, 5% updates) as our prototype does not support scan operations currently. Compared to CRaft and HRaft, FlexRaft increases the throughput of write-heavy workloads (A and F) by 37.28-115.14% and 9.47-18.76%, respectively, because FlexRaft minimizes the network and storage costs during writes. For read-only workload (C), all three protocols achieve similar throughput as they directly read data from the leader. For read-intensive workloads (B and D), FlexRaft improves the throughput of CRaft and HRaft by 8.90-17.88% and 1.40-5.74%, respectively. Thus, FlexRaft achieves the best performance by reducing the network and storage costs during writes.



(a) One server failure

(b) Two server failures

Figure 7: Exp#5 (Performance under YCSB workloads when $N=7$).

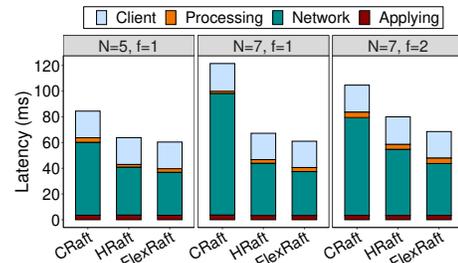


Figure 8: Exp#6 (Breakdown performance of a put operation).

5.3 Microbenchmarks

We study the breakdown performance of a single PUT operation in the presence of server failures (Exp#6). We divide the whole process of a PUT operation into the following parts: 1) the latency of sending a request from the client to the leader, denoted as *client*; 2) the total time spent in handling the request and encoding, called *processing*; 3) the latency of sending AppendEntries RPCs in parallel to replicate an entry, denoted as *network*; 4) the time of applying a log entry to the state machine, denoted as *applying*. Note that the sum of the processing time and network latency is indeed the *commit latency* of a log entry. We fix the value size as 2 MiB and send 10000 PUT requests from the client. We configure the largest value of k for CRaft and HRaft to reduce the redundancy ratio, i.e., $k = 3$ when $N = 5$, and $k = 4$ when $N = 7$.

Figure 8 shows the breakdown performance of CRaft, HRaft, and FlexRaft. All three protocols have similar client, processing, and applying latencies, as they go through the same workflow. The network latency, depending on the network bandwidth cost, determines the overall write latency. For $N = 5$, FlexRaft reduces the network latencies of CRaft and HRaft by 41.02% and 10.45% under one server failure, resulting in 39.95% and 8.19% reduction of commit latencies. For $N = 7$, compared to CRaft and HRaft, FlexRaft reduces the commit latencies by 61.20% and 14.26%, respectively

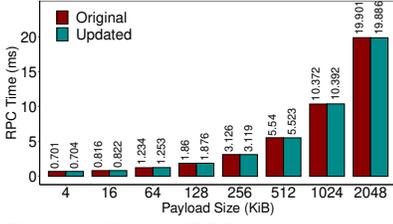


Figure 9: Exp#7 (RPC time of AppendEntries requests).

under one server failure, and 44.51% and 19.37%, respectively under two server failures. The reduction of commit latencies confirms the saving of network traffic; for example, FlexRaft reduces the network bandwidth costs of CRaft and HRaft by 50% and 20% in theory under two server failures when $N = 7$, respectively. Thus, FlexRaft minimizes the network bandwidth and storage costs using a flexible coding scheme for log replication.

5.4 Overhead of FlexRaft

Overhead of updated AppendEntries RPCs. We compare the updated AppendEntries RPCs in FlexRaft to the original RPCs in Raft (Exp#7). We measure the completion time of one AppendEntries RPC (including one request and one response) carrying one log entry, labeled as *RPC time*. We vary the payload size of the log entry from 4 KiB to 2048 KiB. Figure 9 plots the RPC time of the original and the updated AppendEntries. The RPC time of the updated AppendEntries is almost the same as that of the original across different payload sizes. FlexRaft only adds 16 B to the request and 8 B to the response to include the EntryType and ChunkInfo, which are negligible compared to the payload sizes. Thus, the overhead of the updated AppendEntries in FlexRaft is negligible.

Overhead of varying the coding scheme. We study the overhead of varying the coding scheme in FlexRaft (Exp#8). We measure the write latency under a different number of server failures for a group of seven servers (i.e., $N = 7$). Here, we fix the value size as 2 MiB. Figure 10 shows the average commit latencies of Raft, CRaft, HRaft, and FlexRaft over 10000 times. All three protocols with erasure coding achieve the same latency using RS(4,3) in normal cases, reducing the commit latency of Raft by 69.03%. When a server fails during log replication (labeled as $f \rightarrow 1$), CRaft reduces to full-copy replication, while HRaft needs to store one more chunk in three followers; FlexRaft re-encodes the log entry with RS(3,4) and distributes the new chunks to the remaining followers. The additional network costs of CRaft, HRaft, and FlexRaft are 5, 3/4, and 5/3, respectively. Thus, the write latency of CRaft is the largest, which is 11.67% higher than that of Raft and about four times the normal latency with erasure coding. HRaft and FlexRaft increase the normal write latency by 63.45% and 90.28%, respectively, but this only occurs once when the server fails during writes. When the server status remains unchanged under one server failure, FlexRaft achieves the lowest commit latency and reduces the latency of CRaft and HRaft by 61.20% and 14.26%, respectively. Note that CRaft performs the same as Raft since one server fails. When one more server fails during writes (i.e., $f \rightarrow 2$), FlexRaft switches to RS(2,5), resulting in 11.96% higher latency than HRaft. When there are two failed servers, FlexRaft reduces the latencies of CRaft and HRaft by 44.51% and 19.37%, respectively. The commit latency of

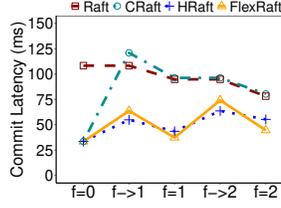


Figure 10: Exp#8 (The overhead of varying the coding scheme).

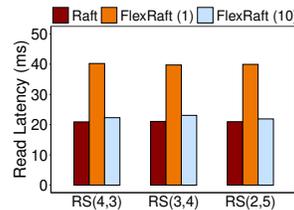


Figure 11: Exp#9 (Decoding overhead during reads).

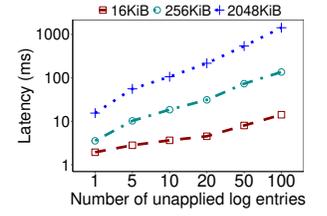


Figure 12: Exp#10 (Duration of the LeaderPre operation).

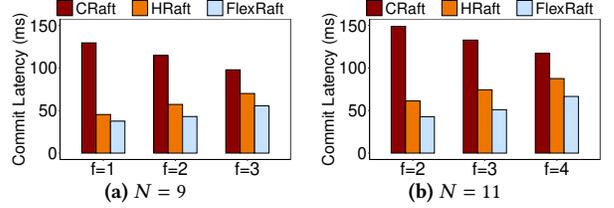


Figure 13: Exp#11 (Commit latency of CRaft, HRaft, and FlexRaft with a larger number of servers in a group).

CRaft under two server failures is lower than that under one failure since CRaft has one fewer follower storing a full copy of the log entries and reduces the network and storage costs. While FlexRaft incurs additional network transfer by varying the coding scheme when a server fails during the log replication, FlexRaft achieves the lowest latency by minimizing the network and storage costs when the cluster status remains stable (i.e., $f = 0, 1, 2$).

Decoding overhead during reads. We compare the read performance of FlexRaft to that of Raft (Exp#9). Figure 11 shows the read latencies of Raft, FlexRaft at the first time to read (FlexRaft (1)), and FlexRaft with 10 read times (FlexRaft (10)) under different coding schemes when $N = 7$. Compared to Raft, FlexRaft (1) increases the read latencies by 88.60-92.35%, while FlexRaft (10) only increases by 4.34-6.62%. Thus, the decoding operation does not harm the read performance too much when the leader does not crash frequently (i.e., the leader already stores a full copy for read requests).

Duration of the LeaderPre operation under the leader failure. We measure the duration of the LeaderPre operation when the leader fails and a new leader is elected; here, we consider $N = 7$ (Exp#10). As the LeaderPre operation only recovers the unapplied log entries, we vary the number of unapplied log entries (from 1 to 100) with different value sizes. Figure 12 shows the duration of the LeaderPre operation, which increases with the number of unapplied log entries. For the small values (e.g., 16 KiB), it takes about 4 ms and 14 ms to recover 10 and 100 log entries, respectively. For the values with a medium size (e.g., 256 KiB), the LeaderPre operation lasts about 18 ms and 136 ms when there are 10 and 100 unapplied log entries, respectively. For the large values (e.g., 2 MiB), it takes about 106 ms to recover 10 entries and 1.4 s to recover 100 entries. Thus, the newly-elected leader can quickly complete the LeaderPre operation, making a limited impact on system availability.

5.5 Scalability of FlexRaft

We evaluate the performance of CRaft, HRaft, and FlexRaft with a larger number of servers in a group (Exp#11). As all three protocols achieve the lowest commit latency using the largest available value of k , we compare their commit latencies under a different number

of server failures. We set the payload size of each log entry as 2 MiB and plot the average commit latencies over 10000 times. Figure 13 shows the commit latencies of CRaft, HRaft, and FlexRaft under a different number of server failures in a group of $N = 9$ and $N = 11$ servers. For a group of 9 servers, FlexRaft reduces the commit latency of CRaft and HRaft by 43.20-71.00% and 17.14-24.87%, respectively. Also, FlexRaft reduces the commit latency of CRaft and HRaft by 43.12-71.45% and 24.01-31.47%, respectively, when more than one server fails in a group of 11 servers.

6 RELATED WORK

Improving the performance of consensus protocols. Distributed systems employ consensus algorithms to provide high reliability and availability for upper-layer applications. Paxos [16, 17] is one commonly-used consensus protocol in distributed systems, such as Chubby [3] and Spanner [3]. Many variants of Paxos [1, 5, 18, 19, 24, 28] have been proposed in the literature to improve the performance of distributed systems built based on the Paxos algorithm. For example, EPaxos [19, 24] extends vanilla Paxos in a decentralized fashion to achieve optimal commit latency.

Raft [21] is a widely used consensus algorithm designed for easy understanding and implementation, which is equivalent to Multi-Paxos. There have been many optimizations proposed to improve the performance of the original raft protocols in recent years [4, 15, 26]. ParallelRaft [4] realizes a parallelized version of Raft, which removes the original Raft’s strict serialization for high I/O concurrency. HovercRaft [15] extends the Raft protocol by eliminating CPU and I/O bottlenecks to achieve both scalability and fault tolerance. KV-Raft [26] introduces commit return and immediate read into vanilla Raft, to accelerate the write and read performance of distributed key-value storage systems respectively.

Optimizing consensus algorithms with erasure coding. Some studies [14, 20, 25, 27] extend the replication-based Paxos and Raft protocols with erasure coding (based on RS codes [22]) for higher performance with lower overhead. Erasure coding has been widely applied in distributed storage systems [11, 12] to protect data against server failures with a low redundancy ratio. The adoption of erasure coding in consensus algorithms introduces a new approach to improve the overall system performance by reducing the redundancy overhead. RS-Paxos [20] is the first consensus protocol that combines erasure coding into the Paxos protocol, but reduces the liveness level. Pando [25] leverages erasure coding in geo-distributed storage to reduce costs for preserving consistency. CRaft [27] applies erasure coding to Raft, but it degrades to full-copy replication when the number of failed servers exceeds a certain threshold. HRaft [14] addresses the degradation problem of CRaft by maintaining additional coded data in healthy servers instead of switching to full-copy replication when server failure occurs.

7 CONCLUSIONS

This paper proposes FlexRaft which minimizes the network and storage costs by dynamically adjusting the coding scheme used for log replication in Raft. It uses the optimal coding scheme with the available largest k based on the server status. When varying the coding scheme during writes, FlexRaft restricts the overwriting of the coded chunks and updates the AppendEntries RPCs to make

sure that all servers store the right coded chunks.

ACKNOWLEDGMENTS

This work is supported by the Major Research Plan of NSFC (Grant No. 92270202), the Strategic Priority Research Program of CAS (Grant No. XDB44030200), and Research Grants Council of Hong Kong (GRF 14214622 and AoE/P-404/18). The corresponding author is Patrick P. C. Lee.

REFERENCES

- [1] Ailidani Ailijiang, Aleksey Charapko, Murat Demirbas, and Tevfik Kosar. 2019. WPaxos: Wide Area Network Flexible Consensus. *IEEE Transactions on Parallel and Distributed Systems* 31, 1 (2019), 211–223.
- [2] Alibaba. -. Alibaba cloud. <https://us.alibabacloud.com/>
- [3] Mike Burrows. 2006. The Chubby Lock Service for Loosely-Coupled Distributed Systems. In *Proc. of USENIX OSDI*.
- [4] Wei Cao, Zhenjun Liu, Peng Wang, Sen Chen, Caifeng Zhu, Song Zheng, Yuhui Wang, and Guoqing Ma. 2018. PolarFS: An Ultra-low Latency and Failure Resilient Distributed File System for Shared Storage Cloud Database. In *Proc. of VLDB Endowment*.
- [5] Aleksey Charapko, Ailidani Ailijiang, and Murat Demirbas. 2021. PigPaxos: Devouring the Communication Bottlenecks in Distributed Consensus. In *Proc. of ACM SIGMOD*.
- [6] Yu Lin Chen, Shuai Mu, Jinyang Li, Cheng Huang, Jin Li, Aaron Ogus, and Douglas Phillips. 2017. Giza: Erasure Coding Objects Across Global Data Centers. In *Proc. of USENIX ATC*.
- [7] Brian F Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking Cloud Serving Systems with YCSB. In *Proc. of ACM SoCC*.
- [8] Delta V Software. -. RCF. <https://www.deltavsoft.com/>
- [9] etcd. -. etcd. <https://etcd.io/>
- [10] Facebook. -. RocksDB. <https://github.com/facebook/rocksdb/tree/v7.3.1>
- [11] Hadoop. -. HDFS 3.1. <https://hadoop.apache.org/release/3.1.1.html>
- [12] Cheng Huang, Huseyin Simitci, Yikang Xu, Aaron Ogus, Brad Calder, Parikshit Gopalan, Jin Li, Sergey Yekhanin, et al. 2012. Erasure Coding in Windows Azure Storage. In *Proc. of USENIX ATC*.
- [13] Intel. -. ISA-L. <https://github.com/intel/isa-l>
- [14] Yulei Jia, Guangping Xu, Chi Wan Sung, Salwa Mostafa, and Yulei Wu. 2022. HRaft: Adaptive Erasure Coded Data Maintenance for Consensus in Distributed Networks. In *Proc. of IEEE IPDPS*.
- [15] Marios Kogias and Edouard Bugnion. 2020. HovercRaft: Achieving Scalability and Fault-tolerance for Microsecond-scale Datacenter Services. In *Proc. of ACM EuroSys*.
- [16] Leslie Lamport. 1998. The Part-Time Parliament. *ACM Trans. on Comput. Syst.* 16, 2 (1998), 133–169.
- [17] Leslie Lamport et al. 2001. Paxos Made Simple. *ACM Trans. on Sigact News* 32, 4 (2001), 18–25.
- [18] Yanhua Mao, Flavio P. Junqueira, and Keith Marzullo. 2008. Mencius: Building Efficient Replicated State Machines for WANs. In *Proc. of USENIX OSDI*.
- [19] Iulian Moraru, David G Andersen, and Michael Kaminsky. 2013. There Is More Consensus in Egalitarian Parliaments. In *Proc. of ACM SOSP*.
- [20] Shuai Mu, Kang Chen, Yongwei Wu, and Weimin Zheng. 2014. When Paxos Meets Erasure Code: Reduce Network and Storage Cost in State Machine Replication. In *Proc. of ACM HPDC*.
- [21] Diego Ongaro and John Ousterhout. 2014. In Search of an Understandable Consensus Algorithm. In *Proc. of USENIX ATC*.
- [22] Irving S Reed and Gustave Solomon. 1960. Polynomial Codes over Certain Finite Fields. *J. Soc. Indust. Appl. Math.* 8, 2 (1960), 300–304.
- [23] TiKV. -. TiKV: A Distributed Transactional Key-Value Database. <https://tikv.org>
- [24] Sarah Tollman, Seo Jin Park, and John Ousterhout. 2021. EPaxos Revisited. In *Proc. of USENIX NSDI*.
- [25] Muhammed Uluyol, Anthony Huang, Ayush Goel, Mosharaf Chowdhury, and Harsha V Madhyastha. 2020. Near-Optimal Latency Versus Cost Tradeoffs in Geo-Distributed Storage. In *Proc. of USENIX NSDI*.
- [26] Yangyang Wang, Zikai Wang, Yunpeng Chai, and Xin Wang. 2021. Rethink the Linearizability Constraints of Raft for Distributed Key-Value Stores. In *Proc. of IEEE ICDE*.
- [27] Zizhong Wang, Tongliang Li, Haixia Wang, Airan Shao, Yunren Bai, Shangming Cai, Zihan Xu, and Dongsheng Wang. 2020. CRaft: An Erasure-Coding-Supported Version of Raft for Reducing Storage Cost and Network Cost. In *Proc. of USENIX FAST*.
- [28] Hanyu Zhao, Quanlu Zhang, Zhi Yang, Ming Wu, and Yafei Dai. 2018. SDPaxos: Building Efficient Semi-Decentralized Geo-replicated State Machines. In *Proc. of ACM SoCC*.