



CFT-Forensics: High-Performance Byzantine Accountability for Crash Fault Tolerant Protocols

Weizhao Tang  

Carnegie Mellon University, Pittsburgh, PA, USA

Peiyao Sheng  

University of Illinois Urbana-Champaign, IL, USA

Ronghao Ni  

Carnegie Mellon University, Pittsburgh, PA, USA

Pronoy Roy 

Carnegie Mellon University, Pittsburgh, PA, USA

Xuechao Wang  

Hong Kong University of Science and Technology, Guangzhou, China

Giulia Fanti  

Carnegie Mellon University, Pittsburgh, PA, USA

Pramod Viswanath  

Princeton University, NJ, USA

Abstract

Crash fault tolerant (CFT) consensus algorithms are commonly used in scenarios where system components are trusted – e.g., enterprise settings and government infrastructure. However, CFT consensus can be broken by even a single corrupt node. A desirable property in the face of such potential Byzantine faults is *accountability*: if a corrupt node breaks the protocol and affects consensus safety, it should be possible to identify the culpable components with cryptographic integrity from the node states. Today, the best-known protocol for providing accountability to CFT protocols is called PeerReview; it essentially records a signed transcript of all messages sent during the CFT protocol. Because PeerReview is agnostic to the underlying CFT protocol, it incurs high communication and storage overhead. We propose CFT-Forensics, an accountability framework for CFT protocols. We show that for a special family of *forensics-compliant* CFT protocols (which includes widely-used CFT protocols like Raft and multi-Paxos), CFT-Forensics gives provable accountability guarantees. Under realistic deployment settings, we show theoretically that CFT-Forensics operates at a fraction of the cost of PeerReview. We subsequently instantiate CFT-Forensics for Raft, and implement Raft-Forensics as an extension to the popular nuRaft library. In extensive experiments, we demonstrate that Raft-Forensics adds low overhead to vanilla Raft. With 256 byte messages, Raft-Forensics achieves a peak throughput 87.8% of vanilla Raft at 46% higher latency (+44 ms). We finally integrate Raft-Forensics into the open-source central bank digital currency OpenCBDC, and show that in wide-area network experiments, Raft-Forensics achieves 97.8% of the throughput of Raft, with 14.5% higher latency (+326 ms).

2012 ACM Subject Classification Security and privacy → Distributed systems security; Networks → Security protocols

Keywords and phrases CFT Protocols, forensics, blockchain

Digital Object Identifier 10.4230/LIPIcs.AFT.2024.3

Related Version *Full Version*: <https://arxiv.org/abs/2305.09123> [57]

Supplementary Material

Software (Source Code): <https://github.com/proy-11/NuRaft-Forensics.git>
archived at `swh:1:dir:7031137a1ba7c969f4b65c8c8b92bbcaaf10d9f3`

Software (Source Code): <https://github.com/WeizhaoT/Raft-Forensics-Simulator>
archived at `swh:1:dir:e86133ad17b9089701023d24e3db7cce362499d9`

Funding *Weizhao Tang, Ronghao Ni, Pronoy Roy and Giulia Fanti*: This work was supported in part by the National Science Foundation under grants CNS-2325477, CIF-1705007, and CCF-2338772, and the Air Force Office of Scientific Research under award number FA9550-21-1-0090. We also thank Chainlink Labs, Ripple Labs, and IC3 industry partners for their generous support, as well as Bosch, the Sloan Foundation, Intel, and the CyLab Secure Blockchain Initiative.

Xuechao Wang: This work is supported in part by a gift from Stellar Development Foundation and



© Weizhao Tang, Peiyao Sheng, Ronghao Ni, Pronoy Roy, Xuechao Wang, Giulia Fanti, and Pramod Viswanath;

licensed under Creative Commons License CC-BY 4.0

6th Conference on Advances in Financial Technologies (AFT 2024).

Editors: Rainer Böhme and Lucianna Kiffer; Article No. 3; pp. 3:1–3:25

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

by the Guangzhou-HKUST(GZ) Joint Funding Program (No. 2024A03J0630).

Pramod Viswanath: This work is supported in part by NSF CNS-2325477, ARO W911NF2310147 and C3.AI.

Acknowledgements We wish to thank Chris Meiklejohn and Heather Miller for their valuable insights and advice on this project. We also thank Sam Stuewe and the MIT Digital Currency Initiative for their feedback and insights regarding integration with OpenCBDC and applications to central bank digital currency.

1 Introduction


In the theory and practice of distributed systems, crash fault tolerance plays a central role [42]. Crash fault tolerant (CFT) protocols allow a system to come to consensus on a log of events even in the presence of nodes that may crash, but otherwise follow the protocol [58, 47, 24, 30]. CFT systems are widely deployed in enterprise systems and support various high-profile services [27, 10, 5, 30, 21]. For example, prevalent systems like etcd [18], CockroachDB [55] and Consul [28] employ CFT protocols like Raft [47]. CFT protocols are also widely-used in security-sensitive critical infrastructure [49, 27], including prospective Central Bank Digital Currencies (CBDCs) [41, 44].

CFT protocols provide theoretical correctness guarantees under the assumption that at least a certain fraction of nodes follow the protocol, and remaining nodes may suffer from crashes. However, these assumptions can be broken in practice. For instance, an agent could be Byzantine, meaning that it can misbehave arbitrarily, e.g., by delaying or tampering with messages. In such cases, consensus can be trivially broken.

One possible solution is to replace the CFT protocol with a *Byzantine fault tolerant* (BFT) protocol, which guarantees consensus under not only crash faults, but also under Byzantine faults [9, 40, 7, 2, 62, 20, 22, 25]. This is a viable solution, though swapping out consensus protocols may be impractical for organizations that have already built infrastructure around a particular CFT system.

In this paper, we explore a complementary approach to managing Byzantine faults: *accountability*. That is, in the case of Byzantine faults in a CFT protocol, can an auditor with access to locally-stored protocol states identify *which* node(s) were responsible for the misbehavior, with cryptographic guarantees? In particular, we want to provide this guarantee by making minimal changes to an existing system and protocol, rather than completely replacing the consensus mechanism.

Accountability for BFT protocols has been studied systematically very recently, both as an intrinsic attribute of existing protocols [50, 45, 46] and as an important feature in the design of new protocols [6, 54, 11, 51]. However, there is comparatively little work on CFT protocols that incorporate accountability for Byzantine faults [26, 24]. An important prior work called PeerReview tackled this problem in the context of general CFT protocols [26]. PeerReview works by producing a signed transcript of every message that is sent in the protocol. Being a general-purpose protocol, it does not always achieve competitive performance with the underlying CFT protocol (details in §6.2). Hence, to our knowledge, existing work on accountability for CFT protocols either: (1) is very general, and thus incurs high performance overhead when applied to specific CFT protocols (i.e., PeerReview [26]), and/or (2) does not include a full implementation-based evaluation to measure the practical effect of accountability [26, 24].



■ **Figure 1** Bandwidth-latency tradeoffs of Raft vs Raft-Forensics over 4 nodes at message size of 256 Bytes.

Our goal in this work is to design a practical accountability framework that incurs low communication and storage overhead by exploiting the structure of the underlying protocol, unlike PeerReview. Crucially, despite exploiting protocol structure, we want the framework to be broadly applicable to common CFT protocols and backwards-compatible with existing systems. To this end, our contributions are threefold:

- **Accountably-Safe Consensus:** We first formally define a subclass of CFT protocols called forensics-compliant protocols, which includes two of the most widely-used CFT protocols in use today: Raft [47] and Paxos [34, 29]¹. Intuitively, the defining feature of this class is that its protocols cycle between two phases: log replication and leader election, and each phase satisfies some formal properties (defined in §4). We then propose CFT-Forensics, a lightweight modification to forensics-compliant CFT protocols that *provably* guarantees to expose at least one node that committed Byzantine faults when consensus is violated. Note that we cannot guarantee to detect more than one Byzantine node, as only one malicious node is needed to break CFT consensus; however, for certain classes of attacks involving multiple Byzantine nodes, we are able to detect multiple misbehaving nodes.
- **Theoretical Efficiency Comparison:** We theoretically analyze the communication and computational overhead of CFT-Forensics compared to the most relevant prior work in this space, PeerReview. We show that CFT-Forensics has (amortized) vanishing storage overhead compared to the baseline protocol in practical scenarios, while PeerReview has overhead that grows linearly with the logs. In addition, during log replication, the communication overhead of CFT-Forensics is 58% lower than PeerReview.
- **Empirical Performance Evaluation on Raft:** We implement Raft-Forensics, an instantiation of CFT-Forensics for the Raft protocol. Our implementation is built on a fork of nuRaft, a popular C++ implementation of Raft. We evaluate its performance compared to Raft, both in benchmark experiments and in a downstream application – specifically, OpenCBDC [41] – an open-source central bank digital currency (CBDC) implementation that uses nuRaft. In benchmark experiments, we observe in Fig. 1 that CFT-Forensics achieves performance close to vanilla Raft (experimental details in §7). For instance, in end-to-end experiments, it achieves a maximum throughput that is 87.8% the maximum throughput of vanilla Raft, at 46% higher confirmation latency (44 ms). In our OpenCBDC experiments over a wide-area network, Raft-Forensics achieves 2.2% lower throughput at 14.4% higher latency (326 ms) than vanilla Raft.

¹ For notational brevity, we use the name “Paxos” to refer to variants of the Paxos algorithm that are sometimes referred to as multi-Paxos to distinguish from the original single-decree Paxos [34, 29].

2 Related Work

CFT protocols

CFT protocols are designed to handle crash faults, where nodes may fail but do not exhibit malicious behavior. Paxos [36] is a foundational CFT protocol, with many variants [38, 34, 15, 35, 4, 37, 43, 58, 29]. Raft [47] is a CFT protocol that aims to provide a more understandable and easier-to-implement alternative to Paxos [58], where HovercRaft [31] further improves its performance. Viewstamped Replication revisited [39] is a CFT protocol that displays elements of both Paxos and Raft. Both Raft [55, 48, 1] and Paxos [8, 5, 53, 14] are widely-used in practice.

Accountability

Accountability allows protocols to identify and hold misbehaving participants responsible when security goals are compromised [33, 32]. In the context of fault-tolerant protocols, accountability allows a protocol to identify culpable participants when security assumptions are violated and demonstrate their misconduct. Recent works [50, 16] have examined several widely used BFT protocols and assessed their inherent accountability levels without altering the core protocols. In addition, some other works [23, 3, 59] have improved the performance of existing BFT protocols by excluding culpable participants from the consensus with accountability. Since CFT protocols are explicitly designed to handle only crash faults, integrating accountability offers a lightweight enhancement to detect Byzantine actors.

One prior work [24] explored the accountability of the Hyperledger Fabric blockchain, which features a pluggable consensus mechanism. This study conducted a case analysis of incorporating accountability into a Hyperledger Fabric system underpinned by a CFT protocol, Apache Kafka [21] (called Fabric*). However, this work treats the consensus module as a cluster, offering accountability only at the level of the entire consensus group (not individual nodes within the group). In contrast, we aim to identify and attribute Byzantine faults to individual misbehaving consensus replicas participants. Fabric* introduces two primary modifications. First, parties must sign every message they send. Second, it enforces a deterministic block formation algorithm to eliminate ambiguity. However, these changes are neither necessary nor sufficient for ensuring accountability in the CFT protocols we study. In addition, Fabric* does not empirically evaluate their system, whereas we evaluate performance both theoretically and empirically.

PeerReview [26] builds a framework for accountability that applies to general distributed systems. Although it accounts for Byzantine faults in CFT protocols as CFT-Forensics does, it has substantially higher overhead communications and space requirements than CFT-Forensics, which we discuss in §6 in detail. PeerReview requires nodes to audit each other, instead of assuming arbitrarily many central auditors as we do (§4). To address this difference, we disable inter-node auditing in PeerReview, which still incurs substantially higher communication and memory overhead than CFT-Forensics.

3 Setup

We study consensus protocols that solve the crash-fault tolerant state machine replication (CFT-SMR) problem over partially synchronous networks. Precisely, we consider a setting with n servers (also known as nodes) and arbitrarily many clients. For the vanilla CFT-SMR setting, we assume that at most f out of the n nodes can suffer *crash failures*, where they

stop working without resuming at an arbitrary and unpredictable moment. Each node u maintains a state machine SM^u and an append-only log list logs . We let $u[i]$ denote the i -th entry in u 's logs . The goal of CFT-SMR is for the nodes to maintain consistent state machines SM^u with each other (Definition 1). SM^u maintains a local state s initialized to $s_0 = \perp$ and a deterministic function ϕ . logs are sequential inputs to SM^u generated from client requests, which results in state transition

$$\text{SM}^u.s_i = \text{SM}^u.\phi(\text{SM}^u.s_{i-1}, u[i]), \quad \forall i \in \mathbb{Z}_{>0}.$$

The network is partially synchronous, meaning that there exists a global stabilization time (GST) and a constant time length Δ , such that a message sent at time t is guaranteed to arrive at time $\max\{\text{GST}, t\} + \Delta$. GST is unknown to the system designer and is not measurable by any component of the system.

► **Definition 1** (CFT(-SMR) Protocol). *In the setting above, a consensus protocol \mathcal{P} is f -CFT(-SMR) if f nodes can fail by crash, and the following three properties are satisfied.*

1. **Safety:** *If E is the i -th entry of a correct node's log, then no other correct node has $E' \neq E$ at index i .*
2. **Liveness:** *If a correct client submits a request r , then eventually all non-faulty nodes will (1) have a log entry E at index i handling r (2) there exists a log entry at all previous positions $j < i$.*
3. **Validity:** *Each entry in the log of a correct replica can be uniquely mapped to a command proposed by a client request.*

In the remainder of the paper, we study f -CFT protocols with $f = \lfloor (n-1)/2 \rfloor$ and focus on the **boldfaced** safety property. These protocols tolerate f crash failures, but are typically vulnerable under even one *Byzantine failure*, where a node arbitrarily deviates from the stipulated protocol (§5.1).

We formalize our threat assumptions below.

3.1 Threat Model

In addition to the f nodes with crash failures, we further assume the existence of $b \geq 1$ nodes that execute Byzantine faults. We assume $b \leq n-2$ to avoid a trivial problem with at most one honest node. The Byzantine nodes are capable of accessing states of honest nodes and collaboratively determining whether, when, and what to send to every honest node. However, they cannot influence the honest nodes or the communication between them.

Auditor

To identify the adversary, we introduce *auditors* in addition to the clients and (server) nodes in the SMR model. Any actor with access to the full states of any node (details in §5.2.1) can be an auditor. Each auditor works independently. If an auditor requests information, honest nodes always provide their information to the auditor; a Byzantine node can respond arbitrarily. Generally, the information is transmitted through the partially synchronous network, so it is guaranteed to arrive at the auditors eventually. The information may also be collected physically, if the entire system is maintained by a centralized party such as a central bank. Any auditor may determine the safety of the system by checking data legitimacy and consistency among the nodes, as a function of the received state information. Our main goal is to define modifications to the consensus protocol and an auditing algorithm that jointly enable an auditor to *eventually* uncover the identity of the adversarial node if the state machine safety property is violated.

We assume the simplest setting with a single trusted auditor who is unable to directly influence the system. Notably, there are various alternative auditing designs. For example, we can introduce additional independent auditors to trade communication complexity for robustness. We can also allow auditors to create checkpoints for the nodes. Since checkpoints influence the system, the auditors should also be coordinated by a distributed consensus for security. We leave the design of a trustless auditing system to future work.

3.2 The Accountability Problem

If even a single node is Byzantine, CFT protocols are vulnerable to safety violations (examples in Section 5.1). As a result, we want to identify the party responsible for a safety violation using an auditing algorithm. If such an algorithm exists, we say the protocol has *accountability*.

► **Definition 2 (Accountability).** *Let \mathcal{P} denote an f -CFT-SMR consensus protocol. \mathcal{P} has accountability if there exists a polynomial-time auditing algorithm \mathcal{A} s.t.*

1. \mathcal{A} takes the states of \mathcal{P} as input.
2. If safety (Def. 1) is violated, \mathcal{A} outputs a non-empty set of nodes and irrefutable proof that each member of the set violated protocol. Otherwise, \mathcal{A} outputs \perp .

4 Forensics-Compliant Protocols: A Family of CFT Protocols

Modifying an arbitrary CFT-SMR protocol under a general workflow without context can be challenging. To address this, we define a family of CFT-SMR consensus protocols named *forensics-compliant*, which are provably modifiable for accountability under our general framework CFT-Forensics (Def. 2 and Theorem 11). At a high level, a forensics-compliant protocol is leader-based (Property 3). It can be described by a set of *procedures*, which is partitioned into *log replication* and *leader election*² with each satisfying necessary properties. Both Raft [47] and Paxos [58], two dominant CFT protocols in practice [29], are forensics-compliant protocols. Additionally, forensics-compliant protocols include Viewstamped Replication revisited [39] and simple variants such as HovercRaft [31].

4.1 Setup

We start with an f -CFT-SMR protocol. In the protocol, each entry in the log has two possible states: *committed* and *uncommitted*. If an entry is committed, the content in the entry will not be changed in the future and can be applied to the state machine. If a prefix in the log is committed, then all entries in the prefix is considered committed. The largest index of committed entries is called *the last commit index*, denoted as `iCommit`.

Let there be global notion of time $T = [0, \infty)$, which is unknown to any of the nodes. For simplicity, we let $u[i]$ denote the i -th log entry in node u 's logs, and $u.\text{loglen}$ denote the length of u 's logs, which equals the index of u 's last entry. For a given entry E with index i , we say a node v *owns* E if $v[i] = E$. Furthermore, we let $E_{i:j}$ denote a sequence of consecutive entries $\{E_k | k \in [i, j], E_k.\text{index} = k\}$. Throughout the paper, we use **colored monospace** text to denote protocols and methods that appear in pseudo-code.

² We use the terminology of Raft for clarity.

Leader-Based

Let S denote the set of nodes with $|S| = n$. A forensics-compliant protocol must satisfy the *leader-based property* (Property 3).

► **Property 3** (Leader-Based). *At any time $t \in T$, each node $u \in S$ identifies a leader $L_u(t) \in \bar{S} \triangleq S \cup \{\perp\}$. For each u , there exists an interval partition of $T = \bigcup_{i=1}^{\infty} [t_{i-1}, t_i)$ and a sequence of nodes $\{\ell_i \in \bar{S}\}_{i=1}^{\infty}$ where $t_0 = 0$, and for all $i \in \mathbb{Z}_{>0}$,*

$$t_{i-1} < t_i, \quad \ell_i \neq \ell_{i+1}; \quad L_u(t) = \ell_i, \forall t \in (t_{i-1}, t_i).$$

If $L_u(t) = u$ for all $t \in [\underline{t}, \bar{t})$, u is called a leader during the leadership $[\underline{t}, \bar{t})$. Otherwise, if $L_u(t) = \ell \notin \{u, \perp\}$, u is called a follower identifying ℓ . Only a leader can propose a log entry. At time \underline{t} when ℓ starts being a leader, it assigns a unique term to itself which is fixed until it stops being a leader at \bar{t} . Hence, the term can be regarded as an attribute of a leadership during $[\underline{t}, \bar{t})$. For node u to identify ℓ , u must receive a message from ℓ that includes ℓ 's term. u sets its term equal to ℓ 's term as soon as it starts identifying ℓ .

We say there exists a global leader $\ell \in S$ of term τ , if there exists a majority subset $M \subseteq S$, such that $\ell \in M$ and for all $u \in M$, $u.\text{term} = \tau$. Since M is the majority, ℓ must be unique at every time, so global leaderships do not overlap in time. We require that the term of a later global leadership must be strictly greater than that of an earlier one.

The full protocol consists of *procedures* that are partitioned into the following subprotocols.

- *Log Replication* is the subprotocol that collects all procedures only executed when the host node u identifies a new leader, i.e., $L_u(t) \neq \perp$.
- *Leader Election* is the subprotocol collecting all the remaining procedures.


Log Replication

On the top level, log replication (Alg. 1) has a main procedure `HandleClientRequest` that is triggered when a leader receives a client request. If a node is not running `HandleClientRequest` or involved in an RPC call within the procedure, it cannot create a new log entry or edit its logs and `iCommit`. It has three steps – log entry creation, replication and commitment.

Creation. When leader ℓ receives a client request, ℓ creates a corresponding log entry E and appends it to the log list. E has 3 attributes – (1) `term`, ℓ 's term; (2) `index`, its index on the log list; and (3) `payload`, which handles the request. We define the *freshness* of a log entry, a log list and a node in Def. 4, and provide an example in Fig. 2.

► **Definition 4** (Freshness). *A log entry E 's freshness is denoted by the tuple $(E.\text{term}, E.\text{index})$. E is as fresh as entry F if their freshness tuples are identical. E is fresher than F if $E.\text{term} > F.\text{term}$ or $E.\text{term} = F.\text{term} \wedge E.\text{index} \geq F.\text{index}$. E is strictly fresher than F if E is fresher than F , and E is not as fresh as F . In contrast, E is staler than F if E is not strictly fresher than F . The freshness of a node or its log list is equivalent to that of the log list's last entry.*

Replication. The procedure of replication can be described by ℓ calling a single RPC `AppendEntries` for each remaining node. Its eventual outcome is a `AppendEntriesResp` message from each callee, which includes a predicate `accept` that indicates whether the replication is successful. In addition, the RPC must satisfy the replication property:



■ **Figure 2** Examples of node freshness. Each box represents an entry containing the entry's term. The entries are **not** necessarily committed. Each world is a possible result of protocol execution by only honest nodes. In all worlds, u 's log list is (unstrictly) fresher than v 's. In World A, u is as fresh as v . In only Worlds B and C, u is *strictly* fresher than v .

► **Property 5** (Replication). *If a follower u replicates $E_{i:j}$ from leader ℓ , u 's term must equal E_j 's term, and it must own $E_{1:i-1}$. Formally, $u.\text{term} = \ell.\text{term} = E_j.\text{term}$ and for all index $k \leq j$, $u[k] = \ell[k]$.*

Commitment. Once ℓ receives `AppendEntriesResp` messages from $(n - f - 1)$ followers with `accept=True`, ℓ commits E_j . Then, ℓ sends a message `InformCommitMsg` including (hash of) E_j to each remaining node u , who also commits E_j if it owns E_j .

■ **Algorithm 1** Log replication of the forensics-compliant family. In persistent storage, a node maintains `term`, `logs`, `iCommit`, `CC` and `LC`, where `iCommit` is the last commit index. The **red lines and variables** are added in CFT-forensics (§5). See §5.2.1 for relevant definitions and our full paper [57] for the complete algorithm.

```

1 Protocol LogReplication(host node  $w$ ):
2   As leader:
3   Procedure HandleClientRequest( $r$ ):
4      $E \leftarrow \text{LogEntry}(\text{term}=\text{term}, \text{index}=\text{loglen} + 1, \text{payload}=\text{Payload}(r))$ 
5      $E.\text{pointer} \leftarrow \text{Hash}(\text{loglen} + 1 || E.\text{payload} || w[\text{loglen}].\text{pointer})$ 
6      $E.\text{stamp} \leftarrow \sigma_w(E.\text{pointer})$ 
7      $E.\text{LC} \leftarrow \text{LC}$ 
8     Append  $E$  to  $w$ 's logs
9      $\text{replicators}, \text{sigs} \leftarrow \{w\}, \{\sigma_w(E.\text{pointer})\}$  // Replication
10    for async  $u \in S - \{w\}$ 
11       $\text{msg} : \text{AppendEntriesResp} \leftarrow \text{async } \text{AppendEntries}([E], \dots)$ 
12      if  $\text{msg}.\text{accept}$ 
13        if not verifySig( $\text{msg}.\text{signature}, E.\text{pointer}, u$ )
14          fail exit
15           $\text{sigs} \leftarrow \text{sigs} \cup \{\text{msg}.\text{signature}\}$ 
16           $\text{replicators} \leftarrow \text{replicators} \cup \{u\}$ 
17    Wait until  $|\text{replicators}| \geq n - f - 1$ 
18    if  $E$  not yet committed // Commitment
19      Commit( $E$ )
20       $\text{CC} \leftarrow i || E.\text{pointer} || \text{replicators} || \text{sigs}$ 
21      for  $u \in S - \{w\}$ 
22        Send InformCommitMsg( $E, \text{CC}$ ) to  $u$ 

```

Leader Election

By definition, leader election is the set of procedures that do not belong to Log Replication, where `CandidateMain` is the main procedure. Without running `CandidateMain` or being involved in an RPC within it, a node cannot identify any leader. Only a candidate within `CandidateMain` can edit its logs and `iCommit`. `CandidateMain` consists of three steps: term switching, candidate qualification, and leadership claim.

Term Switching. At the beginning, the caller ℓ , also called a *candidate*, updates the term to a greater term, which is exactly the term of ℓ 's leadership as the outcome of `CandidateMain`.

Candidate Qualification. This phase is represented by a procedure `Qualification`, which can be completed or *interrupted*. If it is interrupted, `CandidateMain` is also interrupted. Otherwise, it satisfies the election property (Property 6) by necessary communications and state modifications.

► **Property 6 (Election).** *If ℓ completes `Qualification` at term τ , there must exist a set of nodes V where $|V| \geq n - f$ and $\ell \in V$, such that*

1. (*Validity*) After `Qualification`, $\ell[j].\text{term} \geq \ell[i].\text{term}$ for all $j > i$.
2. (*Selection*) For every log entry $E = \ell[i]$ after `Qualification`, there exists $u \in V$ such that $u[i].\text{payload} = E.\text{payload}$.
3. (*Freshness*) Let u^i denote an arbitrary node satisfying Selection at index i . Before `Qualification`, let node v be freshest among u . After `Qualification`, ℓ 's log list is no shorter than v 's and for all $i \leq \text{length of } v\text{'s log list}$, $u^i[i].\text{term} \geq v[i].\text{term}$.

Leadership Claim. After `Qualification`, ℓ identifies itself as the leader. Then, it sends a `LeadershipClaim` message including its term τ to each other node. A recipient u identifies ℓ as the leader and sets its own term to τ if τ is greater than u 's own term; otherwise, u ignores the message.

4.2 Summary

► **Definition 7 (Forensics-Compliant Protocols).** *A forensics-compliant protocol is a leader-based (Property 3) f -CFT-SMR protocol (Def. 1). The protocol can be partitioned into two subprotocols – log replication (Alg. 1) and leader election (Alg. 2), such that*

- *Log replication is a set of procedures that can only be executed when a node identifies a leader. If a node identifies itself, it handles client requests with `HandleClientRequest`, where the `AppendEntries` RPC must have the replication property (Property 5).*
- *Leader election is the set of all the remaining procedures, including `CandidateMain`, which is a unique procedure that allows a node to start identifying a leader. In `CandidateMain`, the `Qualification` procedure must satisfy the election property (Property 6).*

In addition, the log list and `iCommit` must not be modified by any procedure that is not mentioned above or explicitly written in the pseudocode.

► **Proposition 8 (Instances of Forensics-Compliant Protocols).** *Both Raft [47] and Paxos [34, 29] are forensics-compliant.*

■ **Algorithm 2** Leader election of the accountable family. The red lines and phrases are specific to our (unoptimized) CFT-forensics (§5). See §5.2.1 for relevant definitions. See our full paper [57] for the complete algorithm and implementation of `validate`.

```

1 Protocol LeaderElection(host node  $w$ ):
2   Procedure CandidateMain(...):
3     term  $\leftarrow$  a new, higher term than term // Term switching
4     Qualification(term, ...) // Candidate qualification; abort on failure
5      $r \leftarrow w \parallel \text{term} \parallel w[\text{loglen}].\text{term} \parallel \text{loglen} \parallel w[\text{loglen}].\text{pointer}$ 
6      $\text{votes} \leftarrow \{w : \sigma_w(\text{Hash}(r))\}$ 
7     for async  $u \in S - \{w\}$ 
8        $\text{msg} \leftarrow \text{Call RPC RequestVote}(u, \text{term})$ 
9       if verifySig(msg.signature, Hash( $r$ ),  $u$ )
10        |  $\text{votes}[u] \leftarrow \text{msg.signature}$ 
11    Wait for  $\text{votes.size} \geq n - f$ 
12     $\text{LC} \leftarrow r \parallel \text{votes.keys} \parallel \text{votes.values}$ 
13    for  $E \in \text{logs}$ 
14      | if  $E.\text{term} = \text{term}$ 
15        | |  $E.\text{LC} \leftarrow \text{LC}$ 
16     $L_w \leftarrow w$  // Leadership Claim
17    for  $u \in S - \{w\}$ 
18      | Send LeadershipClaim(term, LC) to  $u$ 
19  Response HandleClaimLeadershipMsg( $\ell$ , msg):
20    if msg.term  $\leq$  term
21      | fail exit
22     $L_w, \text{term} \leftarrow \perp, \text{msg.term}$ 
23    if not validate(msg.LC)
24      | fail exit
25     $L_w \leftarrow \ell$ 
26    ... // Protocol-specific state updates

```

Proof (Raft)

Raft is originally designed in a very similar philosophy to the forensics-compliant family. It is a leader-based (Property 3) SMR solution by design. The Raft consensus algorithm has two components: log replication and leader election. We define the core procedures in our full paper [57].

AppendEntries: After a follower receives a list of consecutive log entries (or a single entry), it replicates them if it has the predecessor of the head of the list. Otherwise, it triggers `AppendEntries` recursively to synchronize all uncommitted entries, which guarantees the no-gap property (Property 5).

Qualification: A candidate ℓ in Raft asks voters for votes, and a voter only votes if ℓ 's log list is fresher than its own. This ensures ℓ is fresher than $n - f$ nodes without changing its logs, so `Qualification` RPC satisfies the election property (Property 6).

To summarize, all the RPCs have the required properties, so Raft is forensics-compliant.

Proof (Paxos)

(Multi-)Paxos is an optimized protocol based on a simple array of basic Paxos. Its description varies from paper to paper, so we adopt the version in [29] which enables a clear comparison to Raft. Both the original Paxos [34] and [29]'s variation are leader-based (Property 3). In Paxos, the log replication procedures are identical to those in Raft. Thus, we focus on the leader election subprotocol (details in our full paper [57]).

Unlike Raft, a Paxos voter u always votes for a candidate ℓ with a higher term in [Qualification](#). The vote comes with all u 's entries at ℓ 's uncommitted indices of ℓ . With $n - f - 1$ such votes, at each uncommitted index, ℓ selects the freshest entry it has ever seen. Hence, [Qualification](#) in Paxos also satisfies the election property (Property 6).

To summarize, Paxos (as described in [29]) is also forensics-compliant.

5 CFT-Forensics

Although CFT protocols guarantee safety against crash faults, they are not safety-resilient against even a single Byzantine fault. We first illustrate typical safety attacks. Next, we present CFT-Forensics to endow forensics-compliant protocols with accountability.

5.1 Example Attacks

Recall that in §3, we assumed that $b \in [1, n - 2]$ nodes may behave adversarially. In two examples, we assume $n = 2f + 1$ is odd for simplicity. We show the capabilities of a single attacker Mallory, and the remaining $2f$ nodes are evenly partitioned into X and Y .

► **Example 9** (Proposer's Attack, or Split-Brains). Let Mallory be a corrupt leader. At the same index, Mallory replicates log entries E and $E' \neq E$ to X and Y , respectively. At each side, she commits the corresponding entry with a quorum of $f + 1$ nodes. As a result, the honest nodes in X and Y have different committed log entries at the same index.

► **Example 10** (Voter's Attack). Let Mallory be a corrupt voter who has committed entry E with Y . Nodes in X , however, do not own E . In an election, suppose Carol $\in X$ who earns all f votes from X . When Carol requests vote from Mallory, Mallory votes by simulating itself as a Carol's clone. After being elected, Carol commits $E' \neq E$ at the same index, which conflicts with any honest node in Y .

Surprisingly, these two examples almost exhaustively enumerate the types of safety attacks against forensics-compliant protocols (Theorem 11). This is why forensics-compliant protocols can achieve accountability with much simpler modifications than PeerReview.

5.2 CFT-Forensics Design Overview

We present CFT-Forensics, a framework that enables accountability for forensics-compliant protocols. Here, we present a basic variant of CFT-Forensics, which adds large overhead compared to vanilla CFT protocols; we provide and analyze an optimized variant in §6.1. We use the convention that for a forensics-compliant protocol \mathcal{P} , \mathcal{P} -Forensics denotes the protocol \mathcal{P} augmented with CFT-Forensics (e.g., we implement Raft-Forensics in Section 7).

At a high level, CFT-Forensics adds two central data structures to a forensics-compliant protocol: commitment certificates (CCs) and leader certificates (LCs). A CC irrefutably proves that a quorum of nodes have replicated an entry, and an LC proves *which* quorum of nodes agreed to elect a leader. CFT-Forensics requires each log entry to be signed by its proposer, which provides accountability for a split-brains attack (Example 9). It also requires that each voter signs its vote, for which the voter is forced to take responsibility since the vote exists in a CC or an LC, providing accountability for the voter's attack (Example 10).

Additional Assumptions: Public Key Infrastructure

We assume access to a Public Key Infrastructure (PKI). Each node u has a pair of private and public keys, where the public key is well known, that is, known by all parties in the system, including other nodes and auditors. Node u can use its private key to create an unforgeable signature on (the hash of) an arbitrary message m , denoted by $\sigma_u(\text{Hash}(m))$, and the signature can be verified with u 's public key. A collision-resistant cryptographic hash function Hash is known to all parties. Both signing a message and verifying a signature can be executed in time that is polynomial in message size.

5.2.1 Added States

We first explain the new state that is maintained in CFT-Forensics. CFT-Forensics introduces four new categories of states: hash pointer, proposer stamp, leader certificate (LC) and commitment certificate (CC).

■ **Table 1** Attributes of a CC, an LC and a vote request.

Commitment certificate CC	index		pointer	voters	signatures
	i		h_i	V	$\{\sigma_u(h_i)\}_{u \in V}$
Leader certificate LC			req	voters	signatures
			r	V	$\{\sigma_u(\text{Hash}(r))\}_{u \in V}$
Vote Request	id	term	eterm	end	pointer
	ℓ	$\ell.\text{term}$	$\tau(< \ell.\text{term})$	i	h

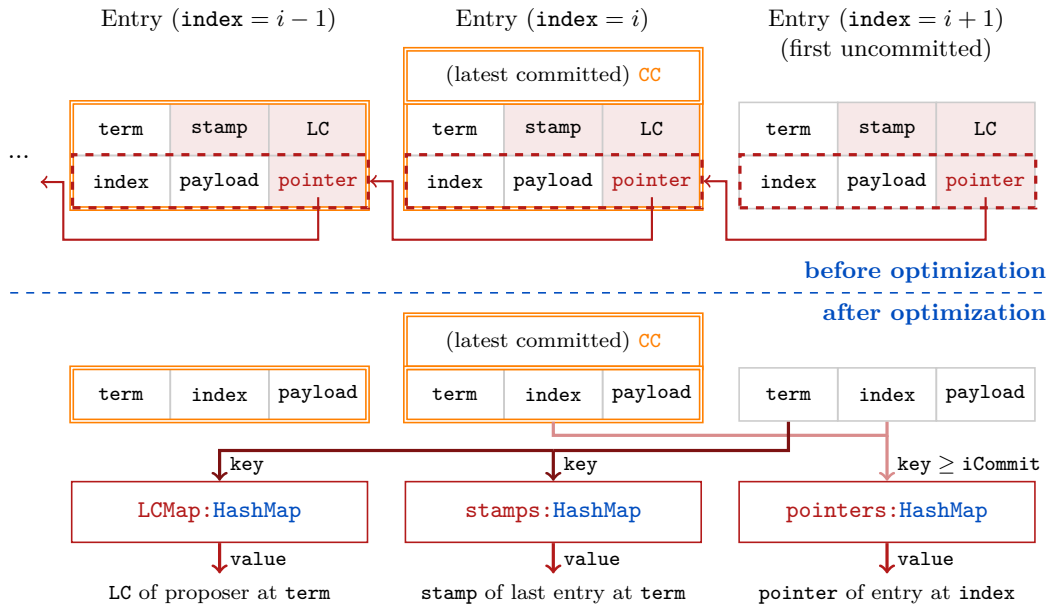
1. **Hash Pointer.** The hash pointer of log entry E_i is denoted by $E_i.\text{pointer}$, where $E_0.\text{pointer} = \perp$. It is a lightweight proof that the host node owns the entire log list from E_1 to E_i . The other hash pointers can be derived by

$$E_i.\text{pointer} = \text{Hash}(i \| E_i.\text{payload} \| E_{i-1}.\text{pointer}), \forall i \in \mathbb{Z}_{>0}. \quad (1)$$

2. **Proposer Stamp.** The (proposer) stamp of log entry E is a digital signature by its proposer ℓ on the hash pointer of E . We denote it by $E.\text{stamp} = \sigma_\ell(E.\text{pointer})$. Should a pair of stamps of E and $E' \neq E$ exist where E is neither an ancestor or descendent of E' and $E.\text{term} = E'.\text{term}$, ℓ must have launched a split-brains attack.
3. **Leader Certificate (LC) of Proposer.** The LC of log entry E , denoted by $E.\text{LC}$, is the LC created by E 's proposer ℓ at term $E.\text{term}$. It collects a quorum of signatures from a set of nodes V on ℓ 's vote request r , where a request includes ID ℓ , term $\ell.\text{term}$, plus the tuple (term, index, hash pointer) of ℓ 's last entry (τ, i, h) . Formally, $r = \ell \| \ell.\text{term} \| \tau \| i \| h$ and $\text{LC} = r \| V \| \{\sigma_u(\text{Hash}(r))\}_{u \in V}$, as shown in Table 1.

In summary, in our basic (un-optimized) CFT-Forensics, a log entry has six attributes (Fig. 3) – **term**, **index**, **payload**, **pointer**, **stamp** and **LC**. In addition, CFT-Forensics requires each node to maintain two independent states – 4) the current leader's LC and 5) the latest CC.

4. **Leader Certificate of Current Leader.** Each node additionally maintains the LC of the current leader it identifies. This LC is not covered above because the current leader may have not proposed any log entry yet.
5. **Commitment Certificate (CC)** Each node only maintains one freshest CC. Like LCs, a CC is a collection of a quorum of signatures on the same log entry. Formally, for a log entry at index i that is replicated to a set of nodes V where $|V| \geq n - f$, we construct a CC following the structure in Table 1. We denote $\text{CC} = i \| h_i \| V \| \{\sigma_u(h_i)\}_{u \in V}$.



■ **Figure 3** Log entry attributes with and without CFT-Forensics; committed blocks are shown with a double gold outline. Our basic (unoptimized) CFT-Forensics (top) adds a hash pointer, a proposer stamp, and a leader certificate LC, all shown in red. We also store a CC only for the latest committed block. Our optimized CFT-Forensics (§6.1) reduces storage costs by storing three hash maps: (1) one containing pointers only for the last committed block and later uncommitted blocks, (2) one storing a single leader certificate LC for every term, and (3) one storing a proposer stamp only for the latest proposed block in the current term.

5.2.2 Modified Procedures

Log Replication

We mark our changes in red in Alg. 1. Upon creation of a log entry E at index i , the leader ℓ correctly attaches the three new states (pointer, stamp and LC). Then it replicates the “enhanced” entry to followers via the `AppendEntries` RPC. Upon receipt, each follower u validates the new states, and eventually puts the entries at their correct indices. As a result of a successful replication, u sends a `AppendEntriesResp` message, which not only includes the predicate `accept`, but also u ’s signature on the last entry E ’s hash pointer.

With $(n - f - 1)$ `AppendEntriesResp` messages, the leader updates its CC by assembling the $n - f$ signatures it has obtained (including its own). To notify followers to commit E , the leader sends a `InformCommit` message which includes CC in addition to E . Upon receipt, a follower commits E if it owns E and the CC passes a follower’s verification.

Leader Election

In the `Qualification` procedure which satisfies the election property (Property 6), if a candidate ℓ ’s logs are changed during `Qualification`, we let ℓ reconstruct every uncommitted entry with the same payload, as if ℓ plans to repropose them. In detail, ℓ a) sets their terms equal to its current term, b) re-derives their hash pointers, c) creates its own stamp for each of them, and d) sets their proposer LCs to its own LC. As a result, the hash pointers will still be correct, and no entry will be overwritten if it has been committed by any node.

Assume that the candidate ℓ passes the [Qualification](#) procedure in a vanilla forensics-compliant protocol. Instead of directly declaring leadership in vanilla, ℓ broadcasts another vote request r based on its current last log entry by calling `RequestVote` RPC. Since ℓ is already qualified, the request deserves at least $n - f$ votes by election property. Each vote from u contains a signature $\sigma_u(\text{Hash}(r))$, proving u 's awareness that ℓ is fresher than itself. After collecting $n - f$ votes, ℓ assembles a leadership certificate (LC) and claims leadership by broadcasting it. Then, each recipient will verify the LC, store it, and identify ℓ as the leader.

In general, we add an additional round of communication to leader election, where the candidate provides information of its last log entry and the voters send signatures. In *passive* leader elections like Paxos, any arbitrary node can be elected under deterministic logic (e.g., under round robin or maximum ID). The new leader must ensure freshness by updating its log entries based on those it receives from the other nodes. As a result, the last log entry is only available after a round of communication, so a second round of signatures is needed. However, it is not needed in *active* elections like Raft, where a node actively seeks leadership candidacy. If each node never modifies its logs during election, then their last entry does not change, and they can collect signatures in just one round of communication.

5.3 Accountability Guarantee

► **Theorem 11.** *If a CFT protocol \mathcal{P} is forensics-compliant, then \mathcal{P} -Forensics achieves accountability (Def. 2).*

Proof Sketch. (Full proof in our full paper [57]) We first establish a map from each term to the LC of that term's leader. If a term is associated with two distinct LCs, we can accuse all voters that contributed signatures to both LCs, as they voted twice at the same term. If this map exists, a term is uniquely used by a leader. Since safety (Def. 1) does not hold, we find the first pair of entries from the logs of two honest nodes that conflict.

If they are of the same term, we discover a *split-brains* attack and we can accuse the leader by its stamps on the conflicting entries or their successors.

If they are of different terms, we discover a voter's attack, which has two possibilities – 1) at least one voter voted for a leader not fresher than itself; and 2) at least one voter replicated and signed an entry at a term less than its term. In this final case, we can accuse all the voters who have signatures in a pair of conflicting CC and LC. ◀

6 Performance Comparison with PeerReview

In this section, we provide a head-to-head comparison of the theoretical overhead costs of CFT-Forensics compared to PeerReview, for the special cases of Raft-Forensics and Paxos-Forensics. We begin by explaining some practical optimizations that reduce the redundancy of CFT-Forensics without affecting accountability, then explain the cost comparison calculations.

6.1 CFT-Forensics State Optimization

The added states in basic CFT-Forensics incur linear overhead in the number of log entries. We next show how to store the new states in independent, more efficient data structures.

Hash Pointer. We let each node u maintain the $u[k].\text{pointer}$ **only** for $k \geq c \triangleq u.i\text{Commit}$ in a hash map `pointers`. This is sufficient for hash pointer reads, which happens only when a node u receives a sequence of entries $E_{i:j}$ to be updated to its logs, plus the preceding pointer $E_{i-1}.\text{pointer}$. We may presume $j > c$ because u rejects updating any

committed entry. Normally, u tells whether $E_{i:j}$ matches its own log list by whether $u[i-1].\text{pointer} = E_{i-1}.\text{pointer}$. If $i \leq c$, u cannot find $u[i-1].\text{pointer}$ in the hash map, but u can alternatively derive $E_c.\text{pointer}$ by (1) and tell whether $u[c].\text{pointer} = E_c.\text{pointer}$. Since `Hash` is collision-resistant, $u[i-1].\text{pointer} = E_{i-1}.\text{pointer}$ is implied by $u[c].\text{pointer} = E_c.\text{pointer}$. Therefore, reduction of committed hash pointers (except the last one) does not affect correctness.

Proposer Stamp. Suppose ℓ has proposed $\{E_k | k \in [i, j]\}$ during a leadership. Since `Hash` is collision resistant, $E_j.\text{pointer}$ effectively represents the entire log list from the head E_1 to E_j . Therefore, the stamp $\sigma_\ell(E_j.\text{pointer})$ proves ℓ has proposed not only E_j , but also E_i, \dots, E_{j-1} . Hence, the stamps on E_i, \dots, E_{j-1} are all redundant, and it suffices to keep only the last stamp E_j within ℓ 's term. Because we only need to maintain one last stamp for each leader, all the stamps can be contained in a hash map `stamps` keyed by term.

Leader Certificate. By design, the LC used for each term is unique. Hence, we may reduce overheads by maintaining the LCs in a hash map `LCmap` keyed by term and valued by LC. Moreover, we may reduce the hash pointer inside the vote request of LC, because the pointer can be derived from the logs.

Summary of Total Spatial Overhead. Let H denote the length of the logs, H' the number of uncommitted entries and Λ the number of global leaderships during which at least one entry is replicated. Our optimized CFT-Forensics substantially reduces total overhead of the three states from $\mathcal{O}(nH)$ to $\mathcal{O}(H' + n\Lambda)$. However, to reduce notations and symbols for better clarity, we continue using the primitive states in the algorithm pseudocode.

6.2 Cost Analysis

Using this optimized implementation, for Raft and Paxos, we compare the overhead space and communication complexities of CFT-Forensics against PeerReview. For log replication, Raft is identical to Paxos, so we merge the comparison in §6.2.1. For leader election, we compare the variants separately in §6.2.2.

PeerReview


PeerReview [26] achieves accountability by logging communication for every message from any node u to another node j , regardless of the underlying consensus protocol. The communication log is an independent data structure introduced by PeerReview. We call such log entries “comm entries”, where each comm entry includes a copy of the message. To make the entire log tamper-evident, a hash pointer is maintained, just as in CFT-Forensics. We assume each comm entry stores a hash pointer, though this storage cost can be reduced by storing one pointer every few blocks, at the expense of time complexity of random access. For every message `msg` sent from u to v , u sends `msg` along with a hash pointer and u 's signature. Then, v replies a hash pointer plus v 's signature to u . Both u and v create a new comm entry including a copy of `msg`. Hence, each message incurs communication overheads of two hash pointers and two signatures.

For auditing, PeerReview allows nodes to supervise each other by forwarding all signatures from a signer to the signer's *witnesses*. For a fair comparison between CFT-Forensics (which has a separate auditor) and PeerReview, we disable witnessing.

6.2.1 Log Replication

■ **Table 2** Complexities of Raft/Paxos, CFT-Forensics and PeerReview in log replication. Π denotes hash size and Σ denotes digital signature size, both in bytes. m denotes number of log messages.

	Raft/Paxos	CFT-Forensics (ours)	PeerReview
	(Base)	Communication Overhead	
Heartbeat	CONST	0	$2(\Pi + \Sigma)$
AppendEntries	mB	$\Pi + 2\Sigma$	$4(\Pi + \Sigma)$
InformCommit	CONST	$\Pi + (n - f)\Sigma$	0
	(Base)	Storage Overhead	
Heartbeat	0	0	$2(\Pi + \Sigma)$
AppendEntries	mB	0	$2mB + 4(\Pi + \Sigma)$
InformCommit	0	0	0



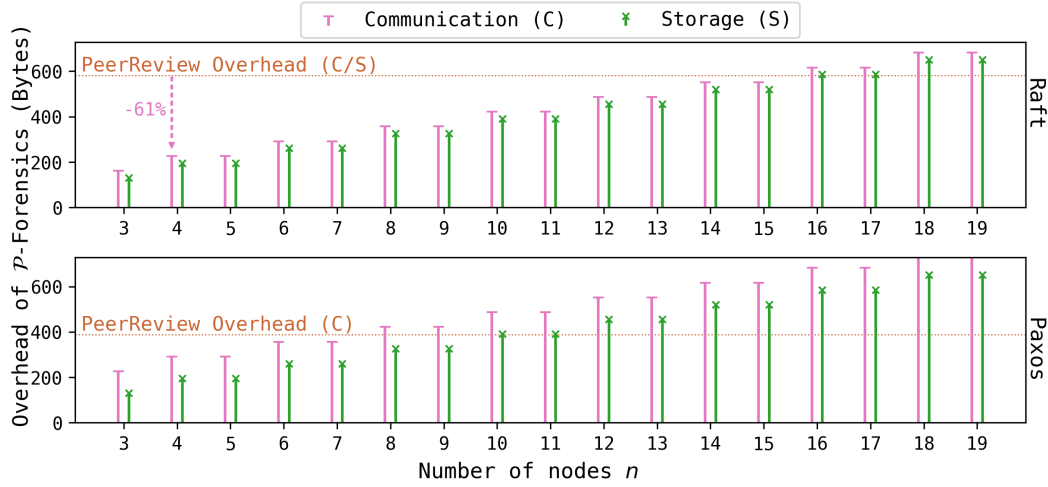
■ **Figure 4** Overhead complexities of CFT-Forensics and PeerReview in log replication when hash size $\Pi = 32$ bytes and digital signatures are $\Sigma = 65$ bytes. The height of the fifth blue bar of “InformCommit (Communication)” is plotted with $(n, f) = (3, 1)$.

Let Π and Σ denote the sizes of a hash and a digital signature, respectively. We choose $\Pi = 32$ bytes and $\Sigma = 65$ bytes for numerical estimation, which are used for Ethereum[61].³ Let B denote the size of a log entry. For messages including a sequence of log entries, we let m denote the number of entries. We assume nodes are up-to-date in term and need only replicate entries of current term. This limits the number of stamps and LCs sent along with the sequence. We also assume that `AppendEntries` complete in a single round, and that `InformCommit` contributes negligible overhead with batched executions.

Table 2 presents the communication and storage complexities of Raft/Paxos, CFT-Forensics and PeerReview in three main log replication RPCs. For our assumed parameter values, we numerically visualize the overheads of the `Heartbeat` and the `AppendEntries` RPCs in Fig. 4. We first observe that CFT-Forensics has **zero** storage overhead in all three RPCs, while PeerReview has a positive overhead for `Heartbeat` and `AppendEntries`. Since message frequency must be lower-bounded by the `Heartbeat` frequency which is typically once every several seconds, CFT-Forensics outperforms PeerReview by saving about 1 KB storage every minute. For communication complexity, we focus on the most frequently-used RPC: (one-round) `AppendEntries`. CFT-Forensics has a $(\Pi + 2\Sigma = 162)$ -byte overhead in communication, which is 58.2% lower than $4(\Pi + \Sigma) = 388$ bytes of PeerReview.

■ **Table 3** Comparison of *overhead* complexities between CFT-Forensics and PeerReview in leader election. $I = 0$ if the candidate’s last committed entry is at the same term as the first entry it receives from the voter; $I = 1$ otherwise. If a voter contributed a signature to the new LC, the LC it receives from the candidate does not need to include its own signature.

		Vanilla (base)	CFT-Forensics (ours) (Overhead)	PeerReview (Overhead)
Raft	Comm.	CONST	$\Pi + (n - f)\Sigma$	$6(\Pi + \Sigma)$
	Storage	0	$(n - f)\Sigma$	$6(\Pi + \Sigma)$
Paxos	Comm.	mB	$\Pi + (n - f + 1)\Sigma$	$4(\Pi + \Sigma)$
	Storage	0	$\tau(n - f)\Sigma$	$2mB + 4(\Pi + \Sigma)$



■ **Figure 5** Overhead complexities of CFT-Forensics and PeerReview in leader election.

6.2.2 Leader Election

Raft-Forensics vs Raft-PeerReview

Now we consider Raft’s leader election. A successful election has three messages between a candidate ℓ and its voter u : 1) ℓ sends vote request to u ; 2) u responds with a vote; and 3) ℓ sends a leadership claim. As shown in Table 3 and Fig. 5, although LC contributes an $\mathcal{O}(n)$ overhead to CFT-Forensics, both complexities are still lower than Raft-PeerReview for $n \leq 15$ (under our assumed parameter values).

Paxos-Forensics vs Paxos-PeerReview

A successful Paxos leader election has two messages between a candidate ℓ and its voter u : 1) ℓ sends its `iCommit` to u ; 2) u responds with all its entries starting with `iCommit + 1`. In Paxos-Forensics, we insert three more messages: 3) ℓ sends a vote request to u ; 4) u responds with a signed vote and 5) ℓ sends an LC to claim leadership. Table 3 lists the overheads for Paxos. We assume that leader elections are rare, so message 2) only includes entries of same term as ℓ [`iCommit`]. By Fig. 5, Paxos has lower communication complexity than Paxos-PeerReview if $n \leq 7$, and on a long enough timescale, its storage complexity is arbitrarily lower than that of Paxos-PeerReview.

³ Ethereum uses Keccak-256 and ECDSA-secp256k1 for hashes and digital signatures, respectively.

7 Empirical Evaluation

We implement Raft-Forensics⁴ in C++ based on nuRaft v1.3 [17] by eBay. With roughly 2,500 lines of code, our implementation fully expands nuRaft with our OpenSSL-based designs in log replication, which correctly reflects the throughput and latency performances between leader elections. We choose the SHA-256 hash function and Elliptic Curve Digital Signature Algorithm (ECDSA) over the secp256r1 curve. For commitment certificates, we used concatenated ECDSA signatures by all the signers.

We evaluate Raft-Forensics in two phases – online phase (§7.1) and offline phase (§7.2). In the online phase, we benchmark the performance of Raft-Forensics over a WAN. In the offline phase, we evaluate the auditing procedure that scans node logs for adversarial behaviors.

7.1 Online Evaluation

Setup on AWS

We evaluate Raft-Forensics over a WAN to demonstrate a geo-redundant deployment for increased resilience [12]. We simulated the WAN environment by deploying Raft-Forensics and other baseline protocols on multiple `c5.large` instances on AWS, where each instance has 2 vCPUs and 4 GB Memory. We ran the experiments on 4 and 16 instances, respectively. Because some typical applications of Raft-Forensics require the nodes to be distributed domestically, we deployed the 16 instances evenly in 8 AWS datacenters in the US, Canada and Europe. For the 4-instance experiments, we deployed the instances in 4 US datacenters.

Baseline Protocols

We compare the performance of Raft-Forensics against Raft [17], using eBay’s NuRaft [17] implementation. We do not directly compare to state-of-the-art BFT protocols in our evaluation because our goal is to propose low-cost solutions that can be easily integrated into existing systems (i.e., the implementation should build upon existing code, and hence be some variant of Raft). Although there exist BFT variants of Raft [56, 60, 13], we were unable to confirm essential theoretical details needed to understand the protocol and guarantees. For completeness, we compare Raft-Forensics against a recent BFT protocol called Dumbo-NG [20] in our full paper [57], though a fair comparison is challenging and not the focus of this work.

Experimental Settings


We benchmark each protocol by two metrics – transaction latency and throughput. Latency is measured by the average time difference between when a transaction is committed by a leader and when it is sent to a node. Throughput is measured by the average number of transactions processed per second during an experiment.

The experiments are configured by two key parameters – transaction size and number of concurrent clients. The transaction sizes range from 256 Bytes to 1 MB. For each transaction size, we sweep the number of concurrent clients sending transactions (in experiments, we let the leader machine spawn transactions). Under each configuration of transaction size and client concurrency, we run all the nodes and client processes simultaneously for 20 seconds.

⁴ <https://github.com/proy-11/NuRaft-Forensics.git>

We measure transaction latency and throughput by the average of five repeated runs to reduce random perturbations. Typically, as the number of clients increases, throughput increases first linearly and then plateaus when the protocol is saturated. In contrast, latency is insensitive to the number of clients before the saturation, but rapidly increases when the bottleneck throughput is reached. We finally evaluate the following quantities:

- **Peak throughput.** We measure the peak throughput of each baseline as the maximum number of transactions processed per second over all numbers of concurrent clients. Fig. 6 presents the performance of all protocols under transaction sizes of 256 Bytes, 4 KiB, and 64 KiB. Compared to Raft, Raft-Forensics has an approximately 10% loss in peak throughput under various transaction sizes, which is caused by the cryptographic operations involved.
- **Latency-Throughput tradeoff.** Under each transaction size, we measure the latency-throughput curve parameterized by number of concurrent clients. Fig. 7 shows the latency-throughput tradeoffs of the two protocols under various transaction sizes. Generally, the tradeoff of Raft-Forensics is only slightly worse than Raft.



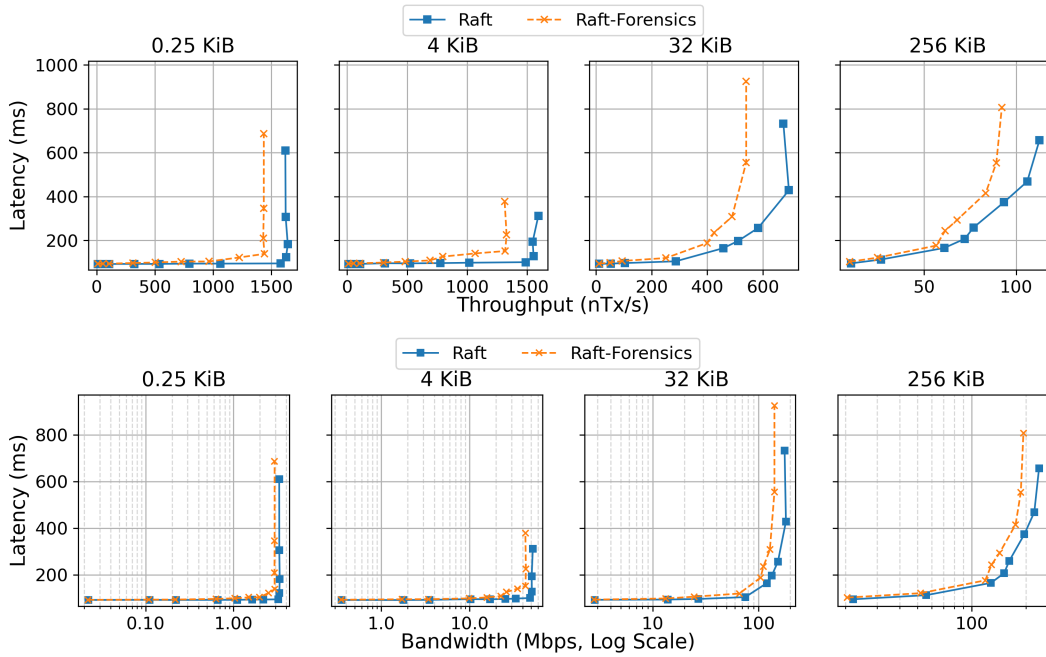
■ **Figure 6** Peak transaction and bandwidth throughputs of consensus algorithms. We plot the error bars with boundaries (mean $\pm 3 \times$ std). ($n = 4$ nodes).

7.2 Offline Evaluation

We next evaluate the offline performance of log auditing. Theorem 11 ensures that we can find at least 1 culprit when State Machine Safety is violated, and we implement an auditing algorithm that finds the culprit. Briefly, the algorithm has two parts: a data legitimacy check and a global consistency check. First, the data legitimacy check verifies the correctness and completeness of the states submitted by each node. For example, the hash pointers must match the logs, and every signature must pass verification. Next, the global consistency check scans for a pair of nodes whose logs are a result of forking. It captures the culprit based on the case discussions in the proof of Theorem 11. See our full paper [57] for the auditing algorithm in detail.

Complexity Analysis

Recall that n denotes the number of nodes. Let H denote the length of the longest chain and Λ the number of elections in total. See our full paper [57] for detailed complexity analysis. The total time complexity of auditing is asymptotically optimal (linear in the size of data $n(H + \Lambda)$), which is required at minimum to ensure data legitimacy), where the complexity



■ **Figure 7** Latency-throughput tradeoff ($n = 4$ nodes). Top row displays throughput in number of transactions per second; bottom row displays throughput in bandwidth.

of global consistency checks does not depend on the chain length H . The linear spatial complexity $\Theta(n(H + \Lambda))$ requires chunked storage of the log chain. For instance, for a chunk size $\Theta(\log H)$, the spatial complexity decreases to $\Theta(n(\log H + \Lambda))$, while the time complexity remains the same. Notably, the time complexity of global consistency check slightly increases to $\mathcal{O}(n(\Lambda + \log^2 H))$, but is still much less than that of legitimacy checks.

Implementation

We implement the auditing algorithm in Python⁵, which can be tested along with a lightweight Raft simulator that achieves better control than the fully-implemented Raft-Forensics in C++ over the leader elections, the adversarial nodes’ behavior and race conditions in general. In particular, it is capable of assigning the adversary to a node and simulating the fork and bad vote attacks in Examples 9 and 10. It ensures that the adversary generates legitimate data to prevent it from being caught before consistency checks. For the best performance in memory usage, it writes the data into chunked files that are available for auditing. In Appendix 7.2, we run benchmarks on the performance of both the data legitimacy and consistency checks of the auditing algorithm. The benchmarks are consistent with our complexity analysis, and demonstrate a significant advantage in chunking data.

Based on the backend software above, we also implement a visualizer based on [52] that demonstrates the attacks and the outputs of the auditing algorithm, including the identity of the culprit and the irrefutable evidence. See our full paper [57] for a screenshot of the visualizer.

⁵ <https://github.com/WeizhaoT/Raft-Forensics-Simulator>

■ **Table 4** Throughput and latency of two different OpenCBDC architectures integrated with Raft and Raft-Forensics (ours), respectively. Each entry is expressed in mean \pm std.

	Throughput (# tx/s)	Latency (ms)
2pc architecture		
Raft	4,800 \pm 14	2,251 \pm 70
Raft-Forensics	4,695 \pm 76	2,577 \pm 252
(% Change)	-105 (-2.2%)	+326 (+14.5%)
atomizer architecture		
Raft	1,284 \pm 56	37,552 \pm 18.75
Raft-Forensics	1,250 \pm 123	40,802 \pm 1,653
(Change)	-34 (-2.6%)	+3,250 (+8.7%)

7.3 Integration with OpenCBDC

Finally, we evaluate the performance of Raft-Forensics integrated into a downstream application: OpenCBDC [41], an open-source implementation of a retail central bank digital currency. OpenCBDC is a good choice because (a) it uses nuRaft, and (b) CBDCs are/will be public infrastructure, so security and performance are paramount. After integrating our Raft-Forensics implementation into OpenCBDC, we deployed our experiments onto `c5n.9xlarge` ec2 instances in AWS over three regions: `us-east-1`, `us-east-2` and `us-west-2`.⁶

We compared Raft-Forensics against Raft in two different OpenCBDC architectures – two-phase-commit (2pc) and atomizer. In both architectures, we replace Raft with Raft-Forensics in every module that is Raft-replicated, i.e., implemented as Raft-variant distributed systems. In the 2pc architecture, we created one generator, one sentinel, three coordinators and three shards, where each coordinator and each shard are Raft-replicated. In the atomizer architecture, we created one watchtower, one watchtower CLI, one sentinel, one archiver, four shards and three atomizers, where only atomizers are Raft-replicated. In both architectures, each Raft-replicated module consists of 3 nodes in 3 different AWS regions.

We used the benchmarking platform [44] of OpenCBDC under default configurations, where load generators produce as much workload as the system can process. The transaction size is 368 bytes. Each experiment lasts 315 seconds and is repeated 3 times. Table 4 shows the throughput and latency of transactions of the entire system. We observe that in practical complex systems like OpenCBDC, Raft-Forensics also performs close to Raft.

8 Discussion and Conclusion

This work is driven by the motivation to improve the Byzantine resistance of CFT protocols by introducing accountability, without sacrificing too much performance. One alternative approach to achieving higher security assurances with CFT protocols involves employing BFT protocols directly. This strategy not only increases tolerance to Byzantine faults but may also inherently include accountability as a bonus feature.

As explained in Section 7, we were unable to directly compare against BFT variants of Raft [56, 60, 13]. Hence, we conducted performance comparisons between Raft-Forensics and leading BFT designs like Dumbo-NG, as detailed in our full paper [57]. Our analysis indicates that, in terms of reducing latency, Raft-Forensics generally surpasses Dumbo-NG,

⁶ Although CFT protocols are often run in the same datacenter, if they are used for critical infrastructure, there will be a need for geographically-distributed deployments for robustness reasons.

though the latter may display competitive or superior throughput for larger transaction volumes. Moreover, Dumbo-NG is optimized for efficiently propagating blocks containing multiple transactions among numerous participants, while Raft variants typically handle single-transaction blocks (as required by SMR) in small-scale distributed systems. As a result, we acknowledge that BFT protocols can indeed be optimized to achieve good performance and replace CFT protocols in applications requiring higher security guarantees, albeit at the cost of increased design complexity and an overhaul of the entire consensus logic. In contrast, accountability may be more suitable for scenarios with moderate security improvement requirements and an emphasis on lightweight changes.

More broadly, accountability need not be viewed as an alternative to Byzantine fault tolerance – it is a complementary, desirable property. For example, all BFT protocols do not inherently offer accountability [50]. We posit that accountability is an important component of distributed system governance – all the more so for geographically-distributed critical infrastructure [19].

References

- 1 Kyle Banker, Douglas Garrett, Peter Bakkum, and Shaun Verch. *MongoDB in Action: Covers MongoDB Version 3.0*. Simon and Schuster, 2016.
- 2 Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. SoK: Consensus in the Age of Blockchains. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, pages 183–198, 2019.
- 3 Michael Ben-Or, Danny Dolev, and Ezra N. Hoch. Simple gradecast based algorithms, 2010. [arXiv:1007.1049](https://arxiv.org/abs/1007.1049).
- 4 Romain Boichat, Partha Dutta, Svend Frølund, and Rachid Guerraoui. Deconstructing Paxos. *SIGACT News*, 34(1):47–67, March 2003. doi:10.1145/637437.637447.
- 5 Mike Burrows. The Chubby Lock Service For Loosely-coupled Distributed Systems. In *Proceedings of the 7th symposium on Operating systems design and implementation*, pages 335–350, 2006.
- 6 Vitalik Buterin and Virgil Griffith. Casper the Friendly Finality Gadget. *arXiv preprint arXiv:1710.09437*, 2017.
- 7 Christian Cachin and Marko Vukolić. Blockchain Consensus Protocols in the Wild. *arXiv preprint arXiv:1707.01873*, 2017.
- 8 Apache Cassandra. Apache Cassandra. *Website*. Available online at <http://planetcassandra.org/what-is-apache-cassandra>, 13, 2014.
- 9 Miguel Castro, Barbara Liskov, et al. Practical Byzantine Fault Tolerance. In *OsDI*, volume 99, pages 173–186, 1999.
- 10 Ben Christensen. Fault Tolerance in A High Volume, Distributed System. Netflix Blog, 2012. URL: <https://netflixtechblog.com/fault-tolerance-in-a-high-volume-distributed-system-91ab4faae74a>.
- 11 Pierre Civit, Seth Gilbert, and Vincent Gramoli. Polygraph: Accountable Byzantine Agreement. *IACR Cryptol. ePrint Arch.*, 2019:587, 2019.
- 12 Team Cloudify. Geo Redundancy Explained, Cloudify. Cloudify Blog, 2021. URL: <https://cloudify.co/blog/geo-redundancy-explained/>.
- 13 Christopher Copeland and Hongxia Zhong. Tangaroa: a byzantine fault tolerant raft. *Stanford University*, 2016.
- 14 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.
- 15 Roberto De Prisco, Butler Lampson, and Nancy Lynch. Revisiting the Paxos Algorithm. *Theoretical Computer Science*, 243(1-2):35–91, 2000.

- 16 Antonella Del Pozzo and Thibault Rieutord. Fork accountability in tenderbake. In *5th International Symposium on Foundations and Applications of Blockchain 2022 (FAB 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.
- 17 eBay. NuRaft. <https://github.com/eBay/NuRaft/tree/v1.3>, 2017. Accessed on April 19, 2023.
- 18 etcd. Etcd. <https://etcd.io/>, 2023. Accessed on April 19, 2023.
- 19 Mohamed Ezzeldin and Wael E El-Dakhakhni. Robustness of Ontario Power Network under Systemic Risks. *Sustainable and resilient infrastructure*, 6(3-4):252–271, 2021.
- 20 Yingzi Gao, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Dumbo-NG: Fast Asynchronous BFT Consensus with Throughput-oblivious Latency. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 1187–1201, 2022.
- 21 Nishant Garg. *Apache Kafka*. Packt Publishing Birmingham, UK, 2013.
- 22 Rati Gelashvili, Lefteris Kokoris-Kogias, Alberto Sonnino, Alexander Spiegelman, and Zhuolun Xiang. Jolteon and Ditto: Network-adaptive Efficient Consensus with Asynchronous Fallback. In *Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, May 2–6, 2022, Revised Selected Papers*, pages 296–315. Springer, 2022.
- 23 Diana Ghinea, Vipul Goyal, and Chen-Da Liu-Zhang. Round-optimal byzantine agreement. Cryptology ePrint Archive, Paper 2022/255, 2022. URL: <https://eprint.iacr.org/2022/255>.
- 24 Mike Graf, Ralf Küsters, and Daniel Rausch. Accountability in A Permissioned Blockchain: Formal Analysis of Hyperledger Fabric. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 236–255. IEEE, 2020.
- 25 Bingyong Guo, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Speeding Dumbo: Pushing Asynchronous BFT Closer To Practice. *Cryptology ePrint Archive*, 2022.
- 26 Andreas Haeberlen, Petr Kouznetsov, and Peter Druschel. PeerReview: Practical Accountability For Distributed Systems. *ACM SIGOPS operating systems review*, 41(6):175–188, 2007.
- 27 Moin Hasan and Major Singh Goraya. Fault Tolerance in Cloud Computing Environment: A Systematic Survey. *Computers in Industry*, 99:156–172, 2018.
- 28 HashiCorp. Consul. <https://www.consul.io/>, 2023. Accessed on April 19, 2023.
- 29 Heidi Howard and Richard Mortier. Paxos vs Raft: Have We Reached Consensus on Distributed Consensus? In *Proceedings of the 7th Workshop on Principles and Practice of Consistency for Distributed Data*, EuroSys ’20. ACM, April 2020. doi:10.1145/3380787.3393681.
- 30 Patrick Hunt, Mahadev Konar, Flavio Paiva Junqueira, and Benjamin Reed. ZooKeeper: Wait-free Coordination For Internet-scale Systems. In *USENIX annual technical conference*, volume 8, 2010.
- 31 Marios Kogias and Edouard Bugnion. Hovercraft: achieving scalability and fault-tolerance for microsecond-scale datacenter services. In *Proceedings of the Fifteenth European Conference on Computer Systems*, EuroSys ’20, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3342195.3387545.
- 32 Robert Künnemann, Ilkan Esiyok, and Michael Backes. Automated verification of accountability in security protocols. *CoRR*, abs/1805.10891, 2018. arXiv:1805.10891.
- 33 Ralf Küsters, Tomasz Truderung, and Andreas Vogt. Accountability: Definition and Relationship To Verifiability. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 526–535, 2010.
- 34 Leslie Lamport. The Part-time Parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, May 1998. doi:10.1145/279227.279229.
- 35 Leslie Lamport. Paxos Made Simple. *ACM SIGACT News (Distributed Computing Column)* 32, 4 (Whole Number 121, December 2001), pages 51–58, December 2001. URL: <https://www.microsoft.com/en-us/research/publication/paxos-made-simple/>.

- 36 Leslie Lamport. *The part-time parliament*, pages 277–317. Association for Computing Machinery, New York, NY, USA, 2019. doi:10.1145/3335772.3335939.
- 37 Butler Lampson. The ABCD’s of Paxos. In *Proceedings of the Twentieth Annual ACM Symposium on Principles of Distributed Computing*, PODC ’01, page 13, New York, NY, USA, 2001. Association for Computing Machinery. doi:10.1145/383962.383969.
- 38 Butler W Lampson. How To Build A Highly Available System Using Consensus. In *International Workshop on Distributed Algorithms*, pages 1–17. Springer, 1996.
- 39 Barbara Liskov and James Cowling. Viewstamped replication revisited. Technical Report MIT-CSAIL-TR-2012-021, MIT, July 2012.
- 40 Shengyun Liu, Paolo Viotti, Christian Cachin, Vivien Quéma, and Marko Vukolic. XFT: Practical Fault Tolerance Beyond Crashes. In *OSDI*, pages 485–500, 2016.
- 41 James Lovejoy, Madars Virza, Cory Fields, Kevin Karwaski, Anders Brownworth, and Neha Narula. Hamilton: A *High-Performance* Transaction Processor For Central Bank Digital Currencies. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 901–915, 2023.
- 42 Nancy A Lynch. *Distributed Algorithms*. Elsevier, 1996.
- 43 Hein Meling and Leander Jehl. Tutorial Summary: Paxos Explained from Scratch. In *International Conference On Principles Of Distributed Systems*, pages 1–10. Springer, 2013.
- 44 mit dci. Opencbdc-tctl. <https://github.com/mit-dci/opencbdc-tctl>, 2022. Accessed on April 19, 2023.
- 45 Joachim Neu, Ertem Nusret Tas, and David Tse. The Availability-accountability Dilemma and Its Resolution via Accountability Gadgets. In *International Conference on Financial Cryptography and Data Security*, pages 541–559. Springer, 2022.
- 46 Joachim Neu, Ertem Nusret Tas, and David Tse. Accountable Safety Implies Finality. *arXiv preprint arXiv:2308.16902*, 2023.
- 47 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.
- 48 Mohammad Roohitavaf, Jung-Sang Ahn, Woon-Hak Kang, Kun Ren, Gene Zhang, Sami Ben-Romdhane, and Sandeep S Kulkarni. Session Guarantees with Raft and Hybrid Logical Clocks. In *Proceedings of the 20th International Conference on Distributed Computing and Networking*, pages 100–109, 2019.
- 49 Ermin Sakic and Wolfgang Kellerer. Response Time and Availability Study of RAFT Consensus in Distributed SDN Control Plane. *IEEE Transactions on Network and Service Management*, 15(1):304–318, 2017.
- 50 Peiyao Sheng, Gerui Wang, Kartik Nayak, Sreeram Kannan, and Pramod Viswanath. BFT Protocol Forensics. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, pages 1722–1743, 2021.
- 51 Peiyao Sheng, Gerui Wang, Kartik Nayak, Sreeram Kannan, and Pramod Viswanath. Player-replaceability and Forensic Support Are Two Sides of the Same (crypto) Coin. *Cryptology ePrint Archive*, 2022.
- 52 simplespy. DiemForensics. <https://github.com/simplespy/DiemForensics>, 2020. Accessed on April 19, 2023.
- 53 Swaminathan Sivasubramanian. Amazon DynamoDB: A Seamlessly Scalable Non-relational Database Service. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 729–730, 2012.
- 54 Alistair Stewart and Eleftherios Kokoris-Kogia. GRANDPA: A Byzantine Finality Gadget. *arXiv preprint arXiv:2007.01560*, 2020.
- 55 Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis, Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, et al. Cockroachdb: The Resilient Geo-distributed Sql Database. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1493–1509, 2020.

- 56 Dezhi Tan, Jianguo Hu, and Jun Wang. VBBFT-Raft: An Understandable Blockchain Consensus Protocol with High Performance. In *2019 IEEE 7th International Conference on Computer Science and Network Technology (ICCSNT)*, pages 111–115, 2019. doi:10.1109/ICCSNT47585.2019.8962479.
- 57 Weizhao Tang, Peiyao Sheng, Ronghao Ni, Pronoy Roy, Xuechao Wang, Giulia Fanti, and Pramod Viswanath. Cft-forensics: High-performance byzantine accountability for crash fault tolerant protocols, 2024. arXiv:2305.09123.
- 58 Robbert Van Renesse and Deniz Altinbuken. Paxos Made Moderately Complex. *ACM Computing Surveys (CSUR)*, 47(3):1–36, 2015.
- 59 Jun Wan, Atsuki Momose, Ling Ren, Elaine Shi, and Zhuolun Xiang. On the amortized communication complexity of byzantine broadcast. In *Proceedings of the 2023 ACM Symposium on Principles of Distributed Computing*, PODC '23, pages 253–261, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3583668.3594596.
- 60 Zhou Wang, Zhang and Xu. A Byzantine Fault Tolerance Raft Algorithm Combines with BLS Signature. *Journal of Applied Sciences*, 38(1):93, 2020. doi:10.3969/j.issn.0255-8297.2020.01.007.
- 61 Dr. Gavin Wood. Ethereum: A Secure Decentralised Generalised Transaction Ledger (Paris Version). <https://ethereum.github.io/yellowpaper/paper.pdf>, March 2024. (Accessed on 05/22/2024).
- 62 Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. HotStuff: BFT Consensus with Linearity and Responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 347–356, 2019.