RoBFT: Robust Byzantine Fault Tolerance for Client-centric Mobile Web Applications

Anonymous Author(s)

***Abstract*—Part of the web is shifting to a client-centric, decentralized model where web clients become the leading execution environment for application logic and data storage. However, current solutions to build decentralized web applications with multiple distrusting parties often involve a decentralized backend of servers running a BFT protocol between them. In this paper, we present RoBFT, a purely browser-based platform for decentralized BFT consensus in client-centric, community-driven web applications. We propose a novel, optimistic, leaderless consensus protocol, tolerating Byzantine replicas, combined with a robust and efficient state-based synchronization protocol. This protocol makes RoBFT well suited for the decentralized client-centric web and its dynamic nature with many network disruptions or node failures. Using a state-based protocol, no transaction log is stored, keeping the storage footprint small for client-centric devices.**

I. Introduction

Browsers and client-side web technologies offer increasing capabilities to enable fully client-side web applications that can operate independently and in a stand-alone fashion, in contrast to the server-centric model [1], [2]. Mobile apps are also more and more purely web-based clients, where the execution environment is just a browser-based container for a mobile web app. Web 3.0 can be defined as the decentralized web where users are in control of their data, and that replaces centralized intermediaries with decentralized networks and platforms. Community-driven, decentralized networks can open the road to many use cases for the sharing economy [3] or shared loyalty programs for local communities [4]. Such client-centric collaborations can, for example, enable a small network of merchants in a local shopping street, or at a farmer's market to set up a shared loyalty program between the merchants in an ad-hoc fashion. These small-scale, specialized collaborative networks can empower motivated citizens to bring value to their local community, without involving an incumbent big-tech company that can change the rules unilateral at any moment.

However, current state-of-the-art peer-to-peer data synchronization frameworks for the browser such as Legion [5], Automerge [6], and OWebSync [7] focus on full replication and eventual consistency between trusted clients. Each replica can modify all data, and all modifications are automatically replicated to all replicas. These protocols lack Byzantine Fault Tolerance (BFT). Yet, they are easy to set up and *trusted* parties can quickly use these to synchronize and modify a shared data set between them.

Decentralized interactions between *distrusting* parties can be enabled by using a classical BFT consensus protocol such as PBFT [8], BFT-SMaRt [9], Tendermint [10], Algorand [11], Ouroboros [12], or HotStuff [13]. These classical BFT protocols are very fast and have a high throughput, but typically assume server-to-server communication with low-latency network connections, and assume every node is connected to all other nodes. Nakamoto consensus [14], used in several blockchains such as Bitcoin and Ethereum [15], relaxes this requirement and only requires a loosely coupled network. However, blockchains based on Nakamoto consensus are too slow for many use cases. They need minutes, or even an hour, to confirm a transaction with high probability. Moreover, they consume a large amount of energy and need a lot of processing power. At last, Avalanche consensus [16] tries to solve the scalability problem by using the concept of metastability. Only a small subset of replicas need to be sampled to reach consensus. However, you still need a connection to every other replica, as the replicas that you need to sample change continuously.

Ultimately, a decentralized mobile web application should be able to run in a robust and resilient way over a network of online client devices such as smartphones. Such devices have a permanent yet unstable internet connection over a data subscription, and are operational and reactive most of the time. However, the existing BFT consensus protocols are designed for more server-like infrastructure that has lots of processing power, storage space, and a stable, low-latency network connection. The motivated citizens in our envisioned use cases do not have this kind of knowledge, budget, and infrastructure available to set up a private network of servers running a BFT protocol between them. They rather want to use their existing hardware such as a low-end computer, or even a mobile device. They could use a public blockchain network, at the cost of paying a fee for every transaction, which lowers the economic viability of this approach. A private network between the citizens without fees is more suitable. This also has the advantage that not all data is publicly readable by the whole world.

In this paper, we present RoBFT: Robust BFT, a novel peer-to-peer data synchronization framework for decentralized web applications between mistrusting parties. RoBFT combines the efficient operation and lightweight setup of a peer-to-peer data synchronization framework with the resilience and fault tolerance of a BFT consensus protocol. The novel BFT protocol, optimized for unstable network conditions, does not require that all replicas are connected to each other. It also does not rely on a leader, removing the need for a costly leader-election procedure when this leader is malicious or loses its network connection temporary. The latter scenario is

common in our target environment. Each browser replica only maintains the current authenticated state, and does not need to keep track of an operation log or transaction history, keeping the storage footprint small. To further reduce the storage and bandwidth requirements, we use an aggregate signature scheme called BLS [17]. This also reduces the computational requirements when all replicas are honest, as only a single aggregate signature has to be verified. The authenticated state and consensus votes are replicated over multiple hops using a gossip protocol.

To summarize, RoBFT combines the following contributions in a browser-based middleware:

1) Lightweight, leaderless, client-centric Byzantine fault tolerant consensus.
2) Resilient and robust, state-based synchronization of both the data and the votes for the consensus protocol using state-based CRDTs and Merkle-trees.
3) Compact storage of signatures using the BLS signature scheme, with delayed verification and aggregation.

Our evaluation, using our application use case of a shared loyalty program between small-scale merchants, shows that RoBFT is a practical solution for these kinds of community-driven use cases. RoBFT achieves transaction finality in the order of seconds, even in networks with 100 browser clients, or in unstable network conditions.

This paper is structured as follows. Section II presents a motivational use case. Section III presents RoBFT's lightweight BFT consensus protocol and the state-based replication strategy. The detailed web-based middleware architecture of RoBFT is elaborated in Section IV. Our evaluation in Section V focuses on many aspects of performance in both the optimistic scenario as well as more realistic and even Byzantine scenarios. Section VI elaborates on important related work. We conclude in Section VII.

II. Motivation

We describe an initial use case that would benefit from the lightweight, robust consensus offered by RoBFT. The use case involves business transactions happening in real life and needs interactive performance and robustness, rather than high throughput or scalability. We then formulate our vision on decentralized web applications.

Loyalty programs. Integrated loyalty programs can be more effective than traditional loyalty programs that are limited to a single company [18]. Think about airlines that award *miles* which can be redeemed with several partners. Such collaborations usually introduce an extra trusted intermediary and add more layers of management and operational logistics. This trusted party can charge high transaction costs to be part of the integrated network. For small merchants on a farmer's market or in a local shopping street, this operational overhead is too much of a burden. A decentralized peer-to-peer network can enable fast and secure creation, redemption, and exchange of loyalty points across different merchants.

Vision. We envision that communities will be able to use RoBFT as a platform to explore new applications and use cases that were previously not feasible. While our initial proof-of-concept implementation is targeting the browser, the techniques explained in this paper can be easily ported towards native mobile and lightweight desktop applications. RoBFT does not need any complex infrastructure, and it currently provides a simple JavaScript-based API, which allows many developers to start developing decentralized applications. Those decentralized applications can be made open source, which allows many people to verify and vouch for them. Local communities who want to set up a decentralized application between the local participants, can use such an application and do not need to concern themselves with a complex infrastructure setup to run the application. Nor do they need to rely on a third party general purpose public blockchain network.

III. RoBFT protocol

This section explains the state-based consensus protocol used in RoBFT. First, it describes the adversary model and its properties. Then it explains the protocol specification.

A. System model

We assume a partially synchronous network [19]. Messages can be delayed, dropped or delivered out of order. An adversary might corrupt up to f replicas of the $n \geq 3f + 1$ total replicas. They can deviate from the protocol in any arbitrary way. Such replicas are called Byzantine, while the replicas that are strictly following the protocol are called honest. We assume attackers are computationally bounded and it is infeasible to forge the used asymmetric signatures or find collisions for the used cryptographic hash functions.

We address in this paper a replicated key-value store for which replicas coordinate agreement using a Byzantine Fault Tolerant consensus protocol, such that the following classical properties hold [20]:

- *Termination*: Every correct replica eventually decides some value.
- *Validity*: If all replicas are correct and propose the same value v, then no correct replica decides a value different from v; furthermore, if all replicas are correct and some replica decides v, then v was proposed by some replica.
- *Agreement*: No two correct replicas decide differently.
- *Integrity*: No correct replica decides twice.

All writes to a key-value pair are atomic, meaning that only a single state transition can happen at any time. Extra application-level conditions can be applied to limit who can write to it, and which values are acceptable given the previous value. RoBFT does not use a leader to coordinate the protocol, removing a common single-point-of-failure compared to many existing BFT protocols. In such leader-based protocols, the failure of a leader leads to a long delay before consensus can be reached. The set of replicas is fixed, and changes to the replica set have to be made outside the protocol, e.g., by halting the protocol, updating the set of replicas on all replicas, and start the protocol again. Consensus is reached for each key-value pair separately, which means that each key has its own instance of the RoBFT protocol.

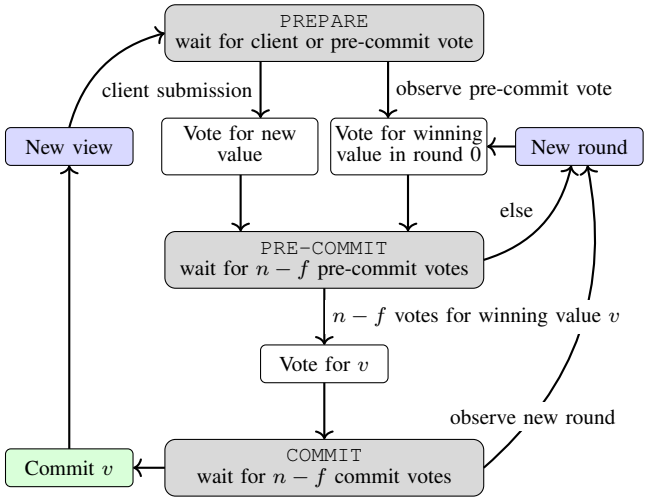

Fig. 1. State transition diagram of the RoBFT consensus protocol.

B. Protocol specification.

The specification of the protocol is shown in Algorithm 1. The state of a replica consists of three parts. The first part is the current value and a quorum certificate. The quorum certificate contains signatures of a supermajority of $n - f$ replicas, and proves the validity of the value. The second part is a map, which maps rounds to a collection of votes for the next value. In each round, there can be multiple proposed values. The third part consists of a new proposed value and a partial quorum certificate for that value. This state is shown at the first 5 lines of Algorithm 1.

Consensus is reached in two steps, first a supermajority needs to be reached in the last round of the $votes$, then a supermajority needs to be reached for the next quorum certificate. The first step will establish a resilient quorum, while the second step will guarantee that sufficiently many replicas know that such a quorum has been achieved. The flow of the protocol is shown in Fig. 1.

1) Proposing new values: To write a new value, a replica has to propose a new value to the other replicas. This process is the PREPARE phase in Algorithm 1. The proposing replica adds the new value and its vote to round 0 of $votes$. As the protocol is leaderless, any replica can be a proposing replica and multiple replicas can propose a new value simultaneously. Replicas are only allowed to vote once in each round for each view, so if the replica already voted for another value in that round, it will have to wait until consensus is reached for the current set of $votes$, and propose the new value in the next view.

2) Consensus: Consensus about which value will be accepted in a view is reached in two phases, called PRE-COMMIT and COMMIT in Algorithm 1. Honest replicas will always vote for the value with the most votes in round 0. If a round has reached a supermajority of votes for a single value, then no new round can be started anymore, and the replicas will start creating a new quorum certificate. If a supermajority of the replicas have voted in a round, but not a single value reaches

Algorithm 1 Basic protocol.

1: $value \leftarrow \perp$ ▷ Current accepted value
2: $qc \leftarrow \perp$ ▷ Quorum certificate for $value$
3: **for** $v \leftarrow 1, 2, 3, \ldots$ **do** ▷ view
4: $votes \leftarrow \emptyset$ ▷ $round \mapsto votesInRound$
5: $qc' \leftarrow \emptyset$ ▷ Next quorum certificate
▷ PREPARE phase
6: **as** a proposing replica:
7: **wait** for value $value'$ from client
8: $votes[0] \leftarrow \{\text{VOTE}(v, 0, value', \text{PRE-COMMIT})\}$
9: **as** a non-proposing replica:
10: **wait** for any value in $votes$
11: **for** $r \leftarrow 0, 1, 2, 3, \ldots$ **do** ▷ round
▷ PRE-COMMIT phase
12: **if** $\neg$HASVOTED$(votes[r])$ **then**
13: $value' \leftarrow$ WINNINGVALUE$(votes[0])$
14: $vote \leftarrow \text{VOTE}(v, r, value', \text{PRE-COMMIT})$
15: $votes[r] \leftarrow votes[r] \cup \{vote\}$
16: **wait** for $(n - f)$ votes in $votes[r]$
17: $value' \leftarrow$ WINNINGVALUE$(votes[r])$
18: $valVotes \leftarrow$ VOTESFORVALUE$(votes[r], value')$
19: **if** LEN$(valVotes) \geq (n - f)$ **then**
20: $vote \leftarrow \text{VOTE}(v, r, value', \text{COMMIT})$
21: $qc' \leftarrow qc' \cup \{vote\}$
22: **else**
23: $value' \leftarrow$ WINNINGVALUE$(votes[0])$
24: $vote \leftarrow \text{VOTE}(v, r + 1, value', \text{PRE-COMMIT})$
25: $votes[r + 1] \leftarrow \{vote\}$
26: **continue**
▷ COMMIT phase
27: **wait** for $(n - f)$ votes in qc':
28: **if** LEN$(votes) - 1 > r$ **then**
29: $qc' \leftarrow \emptyset$
30: **continue**
31: $value \leftarrow$ VALUE(qc')
32: $qc \leftarrow qc'$

a supermajority, a new round is started and all replicas can vote again in this new round. The replicas are only allowed to vote on the current winner in round 0 according to their local state. Because each replica might have a different state on the current set of votes in round 0, there can still be multiple values in the next round without any supermajority for a single value.

Another factor is Byzantine nodes trying to halt the system by voting not according to the rules. However, the set of possible values to vote on gets smaller with every round, and eventually the view of all the honest replicas on the votes in round 0 will become the same, and the winning value can be chosen unanimously. The reason for this is that a replica does not simply send a message with his vote to the others, but instead gossips the entire state. This includes all votes for the previous rounds. This means that when two replicas disagree

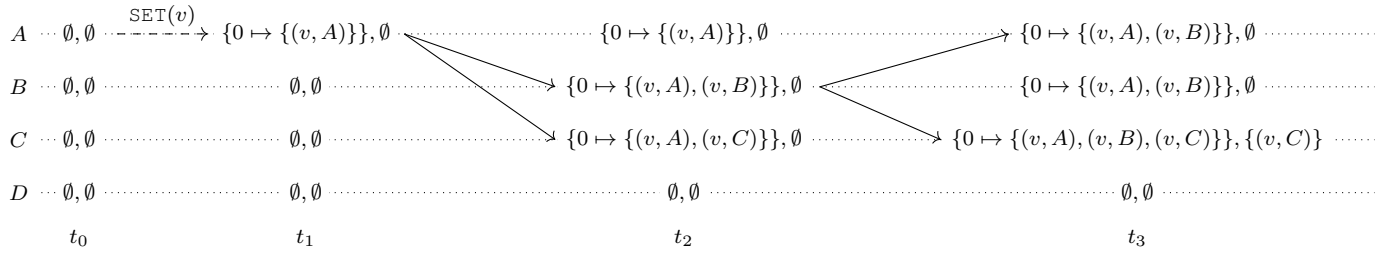

Fig. 2. Example of the state-based synchronization with 4 replicas A, B, C, D. Only the current $votes$ and qc' are shown. Arrows represent a state transfer.

with each other in a certain round, once they communicate with each other, they will learn each other's state. In the next round they will both vote for the same value (as their local state of $votes[0]$ will be the same). Malicious replicas can try to shift the balance to violate liveness, but with each round they have less possibility to do so. Because when they gossip $votes[i]$ they also gossip the previous rounds which should show why they voted on a certain value. If a replica detects that another replica is Byzantine, it will exclude this Byzantine replica permanently, and its votes do not count anymore.

3) Correctness: The integrity and validity properties are trivially satisfied. We can now reformulate the agreement and termination properties more precisely as a safety and liveness property.

Let $\mathfrak{R}$ be a cluster of n replicas with f Byzantine replicas and $n \geq 3f + 1$. RoBFT's correctness is defined by the following two properties:

- *Safety:* If replicas $R_1, R_2 \in \mathfrak{R}$ are able to construct quorum certificates qc_1 for value $value_1$ and qc_2 for value $value_2$ at view v, then $value_1 = value_2$.
- *Liveness:* If an honest replica $R \in \mathfrak{R}$ proposes a new value $value_1$ at view v, eventually a replica will be able to construct a quorum certificate qc for some value at view v.

We proof that RoBFT satisfies these properties in Appendix A.

4) State-based replication protocol: The full state is replicated by using a state-based gossip protocol. A major feature of gossip-based communication is its reliability [21]. Each time a new state is received, the local state is merged with the remote state. This protocol synchronizes data peer-to-peer using state-based Conflict-free Replicated Data Types (CRDTs) [22] combined with a Merkle-tree [23] to efficiently replicate the updated state, similar to Merkle Search Trees [24] or OWebSync [7]. The state of the protocol in Algorithm 1 can be represented as a CRDT: $votes$ and qc' are Grow-only Sets [22], and a state associated with a higher $view$ number overwrites any older state, much similar to a LWWRegister [22]. There are extra constraints imposed on the CRDTs due to the Byzantine nature: not all states are valid, and signatures have to be correct. When a replica receives an invalid state, it will be ignored. The Merkle tree is used to efficiently replicate the state between any two replicas. If the state of two replicas is the same, only the root hash is sent and compared, which limits the network usage. If the states differ, the protocol descends in the tree looking for mismatching hashes to find out which key-value pairs must be synchronized. By using a state-based approach, rather than the operation-based approach of operation-based CRDTs [22], blockchains [14], or traditional BFT protocols, we only need to store the current state together with some metadata. There is no need to store the full log of all operations to later convince replicas that were temporarily offline of the new state. Replicas also do not need to keep track of the state of other replicas, or which messages are already received by which replica. If a new value and quorum certificate with a higher view are received, then the protocol will accept the new state, and the protocol will reset back to line 3 of Algorithm 1 with that newer view. Note that we do not explicitly show the gossiping in Algorithm 1 to keep the algorithm compact. During all phases in the algorithm, the state is continuously replicated to the other replicas. The state-based replication also helps with the consensus protocol. Instead of only sending proposals and decisions to other replicas, the full state of $votes$ and qc' is sent. This approach allows replicas to hold each other accountable when they cast their vote. Their $votes$ should support why they voted for a specific value, otherwise they will be considered Byzantine and excluded from the network.

5) Example: An example of this replication process is shown in Fig. 2. There are four non-Byzantine replicas with an empty set of $votes$ and empty qc' at t_0. The scenario starts at t_1 with replica A proposing a new value v (line 7-8 of Algorithm 1). The state is replicated to the other replicas randomly. In the example, the state is gossiped to replica B and C at t_2, and those replicas merge the received state with their local state. Since B and C did not yet vote in this view and round, they will cast their vote for the current winning value (line 10-15 of Algorithm 1). This process continues at t_3 when replica B sends its state to replica A and C. At t_3, replica C observes that a supermajority of the replicas support value v, and it starts working on a new quorum certificate to determine if at least a supermajority of the replicas also knows about this (line 17-21 of Algorithm 1).

6) Delaying signature verification: For brevity, we did not show the actual verification of signatures in Algorithm 1. However, in the basic protocol, each time a new signature is received, it needs to be verified. This can become quite costly, and therefore RoBFT will use a fast path and delay the verification of any incoming signatures. RoBFT will just accept and replicate them, until a decision needs to be made, such as starting a new round or starting to create a new proposed

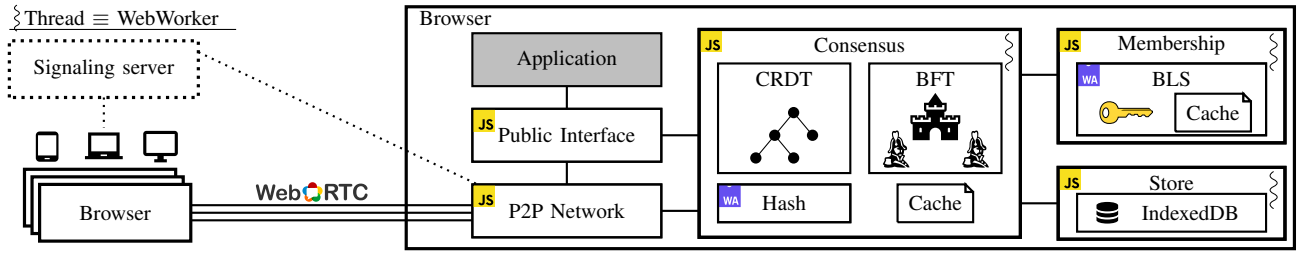

Fig. 3. Browser-based architecture of RoBFT.

quorum certificate. Only then, all signatures will be verified in one batch. If all signatures are valid, the protocol can continue as normal. If there are invalid signatures, then those will be removed and RoBFT will continue to collect more signatures and verify them on arrival. This hybrid approach enables very fast consensus when all replicas are honest, while gracefully degrading to a slower, more costly protocol that can detect which replicas are actively acting Byzantine.

IV. Architecture and Implementation

This section describes the client-centric architecture, deployment, and implementation of RoBFT. This middleware architecture is key to support the BFT consensus and synchronization protocol described in the previous section. RoBFT is fully web-based and written in JavaScript and can execute in any recent browser without any plugins. This section first describes the overall architecture. Then it explains our use of aggregate signatures using BLS to reduce the size of the data.

A. Overall architecture

The RoBFT middleware architecture consists of five main components (Fig. 3): (i) a *public interface* that offers an API for developers, (ii) a *peer-to-peer network* component to communicate directly with other browsers, (iii) a *consensus* component to handle the consensus protocol described in the previous section, (iv) a *membership* component to handle all cryptographic operations, and (v) a *store* component to save all state to persistent storage. The last three components run on a different browser thread by using Web Workers.

(i) Public interface. This component provides an API to application developers to use this middleware. It provides four functions to modify the application state: `GET(key)` returns the current value at the given key, `SET(key, value)` submits a proposal to update the value at the given key, `DELETE(key)` deletes the value at the given key. A tombstone is kept for correct replication, `LISTEN(key, callback)` supports reactive programming by calling the callback with the new value each time a new value for the key is confirmed by the network.

Apart from those functions, the middleware also provides a constructor function to initialize the middleware by passing the following four configuration parameters: the list of all members of the network together with their public key, the private key of the replica, the URL to the signaling server to set up the peer-to-peer connections, and an access-control callback to verify state changes. This access control callback is called before voting for a new proposed value, with both the old and new values as arguments. It should return a `boolean` whether to allow this change or not. This callback enables the implementation of basic access control policies on the values. One example is to embed the public key of the owner into the value and requiring each new value to be signed by the owner. This value can only be changed by the owner, and supports passing ownership by changing the embedded public key.

(ii) Peer-to-peer network. The *P2P Network* component manages the peer-to-peer network and is responsible for the replication of the state-based CRDTs. Many browser-based replicas are connected to each other using WebRTC (Web Real-Time Communications). WebRTC enables a browser to communicate peer-to-peer. However, to set up those peer-to-peer connections, WebRTC needs a signaling server to exchange several control messages. Once the connection is set up, all communication can happen peer-to-peer, without a central server. Another WebRTC peer-connection can also be used as a signaling layer, so once a replica is connected to another one, it can also connect to all of its peers, without the need of a central signaling server. In our adversary model, this server is assumed to be trusted. If this signaling server would be malicious, the safety of the system is not endangered as no actual data is sent to this central server. However, some peers might not be able to join the network and the required supermajority might not be reached, which violates liveness. The use of multiple independent signaling servers can lower the risk of this happening. To defend against an eclipse attack, where few Byzantine neighbors try to surround an honest replica to break liveness, a replica can periodically create new connections to other peers and drop older connections when no updates are being gossiped to them, or when proposals are not being voted on.

(iii) Consensus. The *Consensus* component handles the consensus protocol described in Section III. It maintains a Merkle-tree of all key-value pairs and uses the state-based CRDT framework OWebSync [7] to replicate the local state to other replicas using the *P2P Network* component. The Merkle-tree is constructed using the Blake3 cryptographic hash function. For performance reasons, the hash function is implemented in Rust and compiled to WebAssembly.

(iv) Membership. The *Membership* component contains all cryptographic material and is responsible for all cryptographic operations such as signing and verification of signatures.

$\mathbb{G}_0$ and $\mathbb{G}_1$ are two multiplicitive cyclic groups of prime order q. $\mathsf{H_0} : \{0,1\}^* \rightarrow \mathbb{G}_0$ and $\mathsf{H_1} : \{0,1\}^* \rightarrow \mathbb{Z}_q$ are hash functions viewed as random oracles.

1) *Parameters Generation:* $\mathsf{PGen}(\kappa)$ sets up a bilinear group $(q, \mathbb{G}_0, \mathbb{G}_1, \mathbb{G}_t, e, g_0, g_1)$ as described by [25]. e is an efficient non-degenerating bilinear map $e : \mathbb{G}_0 \times \mathbb{G}_1 \rightarrow \mathbb{G}_t$. g_0 and g_1 are generators of the groups $\mathbb{G}_0$ and $\mathbb{G}_1$. It outputs $params \leftarrow (q, \mathbb{G}_0, \mathbb{G}_1, \mathbb{G}_t, e, g_0, g_1)$.
2) *Key Generation:* $\mathsf{KGen}(params)$ is a probabilistic algorithm that take as input the security $params$, generates $sk \xleftarrow{\$} \mathbb{Z}_q$, computes and sets $pk \leftarrow g_1^{sk}$, and outputs (sk, pk).
3) *Signing:* $\mathsf{Sign}(sk, m)$ is a deterministic algorithm that takes as input a secret key sk and a message m. It computes $t \leftarrow \mathsf{H_1}(pk)$, and outputs $\sigma \leftarrow \mathsf{H_0}(m)^{sk \cdot t} \in \mathbb{G}_0$.
4) *Key Aggregation:* $\mathsf{KAgg}(\{(pk_i, r_i)\}_{i=1}^n)$ is a deterministic algorithm that takes as input a set of public key pk and the multiplicity r pairs. It computes $t_i \leftarrow \mathsf{H_1}(pk_i)$, and outputs $apk \leftarrow \prod_{i=1}^n pk_i^{t_i \cdot r_i}$.
5) *(Multi-)Signature Aggregation:* $\mathsf{Agg}(\sigma_1, ..., \sigma_n)$ is a deterministic algorithm that takes as input n signatures. It outputs $\sigma \leftarrow \prod_{i=1}^n \sigma_i$.
6) *Verification:* $\mathsf{Ver}(apk, m, \sigma)$ is a deterministic algorithm that takes as input aggregated public keys $apk \in \mathbb{G}_1$, and the related message m and signature $\sigma \in \mathbb{G}_0$. It outputs $e(g_1, \sigma) \stackrel{?}{=} e(apk, \mathsf{H_0}(m))$.

Fig. 4. Formal specification of the BLS signature scheme.

We use an aggregate signature scheme called BLS [17]. Section IV-B provides more details about the BLS implementation. It is implemented in C and compiled to WebAssembly.

(v) Store. At last, the *Store* component saves all state to the IndexedDB database. IndexedDB is a key-value datastore built inside the browser. Each value and the Merkle-tree are serialized to bytes and stored there under the respective key. This enables users to close the browser and continue afterwards without losing the current state.

B. Aggregate signatures using BLS

The consensus protocol in Section III is resource-intensive with respect to aggregation and verification of digital signatures. Signatures must be continuously collected and verified. This means, in every intermediate state of a transaction, each party needs to keep track of all incoming signatures and verify them to prevent malicious scenarios. Persistence, management, and transmission of these signatures are costly, especially in a browser-based setting. Therefore, our protocol requires short and compact signatures to reduce storage and network footprint. Boneh–Lynn–Shacham (BLS) [17] presented a signature scheme based on bilinear pairing on elliptic curves. The size of a signature produced by BLS is compact since a signature is an element of an elliptic curve group. The aggregation algorithm [26] outputs a single aggregate signature as short and compact as the individual signatures, unlike other approaches that rely on ECDSA, DSA or Schnorr. Other state-of-the-art BFT systems such as SBFT [27] and HotStuff [13] also use aggregate or threshold signatures. However, they use it in a different way. They let the leader compute the aggregate signature. RoBFT uses a different approach, once a proposed quorum certificate has reached a supermajority of the votes, any replica can aggregate these into one single aggregated BLS signature. RoBFT makes a trade-off between performance, bandwidth and storage space. Verifying a single signature is expensive, however, aggregation is cheap in performance. For this reason, RoBFT will delay the verification of the signatures until the latest possible moment (as explained in Section III-B6). Only then the individual signatures are aggregated and verified. If the verification fails, a binary search can be conducted to find the invalid signatures and remove them. This leads to a higher bandwidth usage, compared to always aggregating two shares immediately. But allows for cheaper recovery when a Byzantine replica is sending invalid signatures. Once a signature is aggregated and verified, the individual shares are discarded, saving both bandwidth and storage space.

The standard scheme is vulnerable to rogue public key attacks. The state-of-the-art approach [25] to mitigate such attacks is to compute $(t_1, ..., t_n) \leftarrow \mathsf{H_1}(pk_1, ..., pk_n)$ for each Agg invocation and compute $\sigma \leftarrow \prod_{i=1}^n \sigma_i^{t_i}$, where pk_i is the public key of replica i, $\mathsf{H_1}$ is a hash function, and σ_i is a signature produced by replica i. Although the t_i values can be cached, the computation of σ would be costly. Moreover, Agg does not take as input the same set of public keys at different states of a transaction in our consensus protocol. Therefore, we distribute the computations by moving the calculations of the t_i and $\sigma_i^{t_i}$ values to the signing parties, and as a result, these computations are performed only once. Now, any replica can run Agg by only computing $\sigma_1 ... \sigma_n$. The security properties of BLS remain intact [25], and we obtain more efficient aggregations at scale. We provide the mathematical background and formal specification of the optimized BLS scheme in Fig. 4.

V. Evaluation

We validated the RoBFT middleware with the loyalty points use case presented in Section II. The first subsection presents this validation. Next, we present three different benchmarks with different scales. The first benchmark shows the performance results in the optimistic scenario without network failures or Byzantine failures. The second benchmark evaluates the performance in a more realistic scenario with some network failures. The last benchmark evaluates the performance in the presence of a Byzantine replica.

A. Validation in the loyalty points use case

Integrated loyalty programs can be more effective than traditional loyalty programs that are limited to a single company. Think about airlines that award *miles* which can be redeemed with several partners. Such collaborations usually introduce an extra trusted intermediary and add more layers of management

and operational logistics. This trusted party can charge high transaction costs to be part of the integrated network. For small merchants on a farmer's market or in a local shopping street, this operational overhead is too much of a burden. A decentralized peer-to-peer network can enable fast and secure creation, redemption, and exchange of loyalty points across different merchants.

The deployment of the loyalty points use case consists of three services: a web application running in a browser for each merchant, a web server to serve the static web application files, and a signaling server to set up WebRTC peer-to-peer connections between the browsers. The web server is optional. Every merchant can also store those application files themselves and load them from their local file system. The signaling server is a trusted component. However, if trust is not present, you can set up multiple signaling servers to reduce potential misbehavior. No actual data is sent to the signaling server. It is only used to discover other peers on the network. To have a baseline, we compare RoBFT to two other existing state-of-the-art systems for BFT consensus: BFT-SMaRt [9], [28] and Tendermint [10], [29]. BFT-SMaRt is a more traditional BFT protocol, similar to PBFT [30], where all replicas are connected to each other, and one leader drives the protocol. If that leader fails, a new one will have to be elected before any progress can be made. Tendermint uses gossip for communication between the replicas. There is still a leader, however, that leader changes frequently.

B. Test setup.

To test the performance of RoBFT, we implemented the use case and deployed it on the Azure public cloud. We used 21 VMs (Azure F8s v2 with 8 vCPUs and 16 GB of RAM) with one VM acting as a central server running the web server and signaling server. The other VMs are running Chrome browsers inside a Docker container. Each of those VMs holds one to five browser instances for different scales of the benchmarks. To simulate a truly mobile environment, the network is delayed to an average latency of 60 milliseconds using the Linux `tc` tool, which simulates the latency of a 4G network. Every test is executed 10 times to ensure the results are reliable.

We are interested in the time it takes to confirm a transaction, experienced by the browser that submitted the transaction. Each transaction is a group of loyalty points being changed from owner. For example, a merchant gives some loyalty points to a customer or a customer redeems their loyalty points with a merchant. In the evaluation, the browser clients will do one transaction per second. This throughput is more than enough for the local community-scale use cases we envision. We compare the latency and network bandwidth with a different number of browsers. We show a boxplot of the latency results instead of only the average, as all users should experience fast confirmation times, and not only the average user.

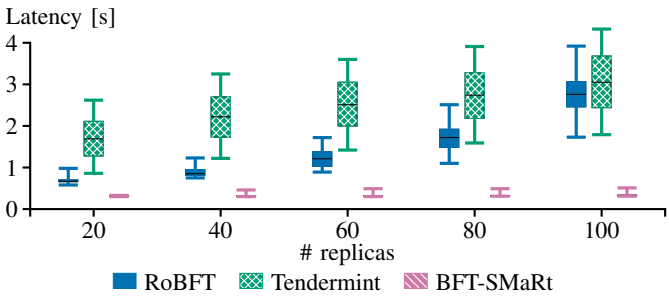

Fig. 5. Latency in the optimistic scenario without failures.

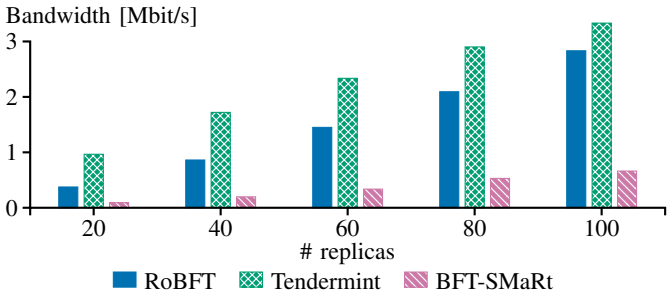

Fig. 6. Network usage in the optimistic scenario without failures.

C. Optimistic scenario

In the optimistic scenario, every replica is honest and no replicas fail, so the fast path can be used. One single aggregate signature is verified only before a decision, avoiding costly signature verifications after every message. As every replica is honest, this aggregate signature is correct and the new value can be accepted by all replicas.

Fig. 5 shows the latency for the different technologies. For the use case of loyalty points, transactions must be confirmed fast, as people are waiting at checkout to receive or redeem loyalty points. RoBFT can confirm transactions within 4 seconds, even with a network of one hundred browsers. BFT-SMaRt can confirm transactions within half a second. This is because all replicas communicate directly with each other. However, having all replicas directly connected to each other is not realistic in a mobile peer-to-peer network. In contrast, RoBFT and Tendermint use gossip and need multiple hops before all replicas are reached. This also causes the increased latency. Furthermore, BFT-SMaRt uses HMAC to authenticate requests, which are an order of magnitude faster than the asymmetric signatures used in RoBFT and Tendermint. We can see a similar pattern in the bandwidth requirements shown in Fig. 6. In the large-scale scenario with 100 browsers, RoBFT uses less than 3 Mbit/s, which is acceptable for a typical mobile network.

D. Realistic scenario

The same benchmark is now repeated with 25% of the replicas failing during the benchmark. A failure is simulated by dropping all network packets to and from that replica. Replicas fail one by one, with a 5-second delay between each failure.

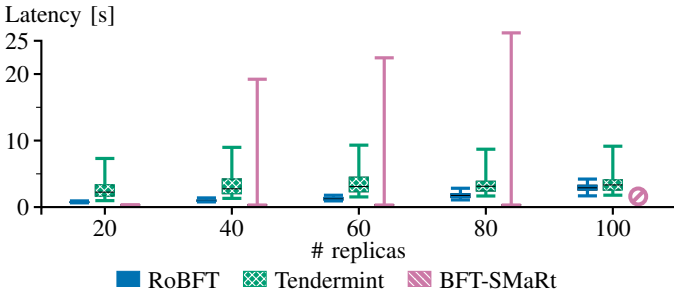

Fig. 7. Latency in the realistic scenario with network failures.

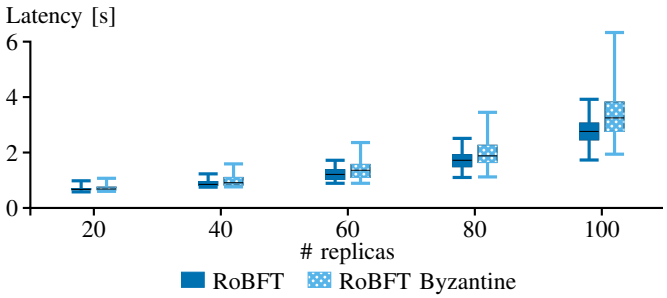

Fig. 8. Comparison of the latency in the normal scenario with one where a Byzantine replica tries to halt the network.

As all systems are Byzantine fault tolerant, they should be able to tolerate up to 33% of the replicas failing or acting Byzantine.

Fig. 7 shows the latency in this scenario. RoBFT is not impacted much by the failing replicas and can still confirm transactions within 5 seconds. The impact on Tendermint is also small, but the latency is doubled to about 10 seconds. BFT-SMaRt however needs to use a costly leader election protocol when the current leader fails. This process takes some time, during which no transaction can be committed. Once a leader is chosen, the same fast performance can be achieved again. This behavior is clearly visible in Fig. 7. The median latency of BFT-SMaRt is not affected by the failures. However, the tail latency increases to 27 seconds for the scenario with 80 replicas. It cannot handle the case with 100 replicas. BFT-SMaRt is unable to handle large network sizes when the latency between the nodes is higher than usual, e.g., in geo-distributed systems or on mobile networks. This has been shown in the literature before [24]. Tendermint does have a leader, but it is rotated round-robin all the time. This makes the failure of a leader less severe, as a new one will quickly be elected anyway.

E. Byzantine scenario

For RoBFT, we performed an extra benchmark with Byzantine replicas. As long as the honest replicas are still using the fast path, the Byzantine replicas will send extra invalid signatures. As the signatures are only verified when a supermajority is reached, the honest replicas only realize this at the end, and they cannot find out which replicas are Byzantine. Once the fast path is disabled, the signatures are verified for every message, so malicious replicas can be detected and excluded from the network. In this case, the Byzantine replicas keep the signature intact to avoid being detected. However, they will try to slow down the consensus by not voting themselves.

The latency in this Byzantine scenario is shown in Fig. 8. RoBFT can handle Byzantine replicas very well for smaller networks, however, for networks of size 100 replicas, the tail latency becomes 7 seconds. Which might already be quite high for the use case of loyalty points. We did not test the effect of Byzantine replicas for BFT-SMaRt or Tendermint. As they do not use a fast path when everyone is honest, the impact is less. However, if the current elected leader happens to be Byzantine, it can delay the consensus until some timers end and a new leader is elected [31].

F. Discussion and conclusions

We have shown that RoBFT can be used for the loyalty points use case with up to 100 different merchants, even when some of them are acting maliciously. RoBFT can achieve similar latencies as other gossip-based BFT protocols, such as Tendermint. Our evaluation also shows the trade-offs that RoBFT makes. In an optimal scenario where there is a good connection available between all replicas and no network disruptions or crashes happen, then a classical leader-based protocol such as BFT-SMaRt will outperform RoBFT. However, as we mention in the introduction, we envision a more ad-hoc network between low-end devices on a residential or even a mobile network, where short-term disruptions are common. Our evaluation shows that RoBFT is very robust against this kind of setting and achieves similar performance as in the optimal scenario: a transaction is always finalized within 5 seconds. A leader-based protocol such as BFT-SMaRt is not well suited. The temporary failure of a leader leads to long commit times, and even total failure for larger network sizes. This leader also needs more resources and a direct connection to every other replica. Keeping 100 WebRTC connections open in a browser, while theoretically possible, drastically reduces performance. However, RoBFT does not impose this, since consensus can be reached gradually over time, as the full state of the proposals and votes propagates through the network. RoBFT can confirm transactions fast, in the order of seconds, without needing a complex back-end setup or wasting a lot of energy. RoBFT has a small storage footprint due to its state-based nature.

VI. Related work

Several client-side frameworks for data synchronization between web applications exist: Legion [5], Automerge [6], and OWebSync [7]. They make use of various kinds of Conflict-free Replicated Data Types (CRDTs) [22] to deal with concurrent conflicting operations, and can synchronize data peer-to-peer. They are easy to set up and only require a browser and a peer-to-peer discovery service. However, they assume trusted operation as the default setting. Some work has been done in a semi-trusted setting [32], [33]. None of them can tolerate Byzantine parties.

WebBFT [34] shares a similar vision of client-centric, decentralized web applications. However, they only interface to a backend BFT-SMaRt cluster, instead of running the BFT protocol directly between browsers.

Open or permissionless blockchains such as Bitcoin [14] and Ethereum allow everyone to participate and use Proof-of-Work (PoW) to reach agreement over the ledger. However, PoW has several flaws [35]. PoW uses a lot of processing power and energy [36] and performs poorly in terms of latency. It assumes a synchronous network to guarantee safety. When this assumption is violated, temporary forks can happen in the blockchain as liveness is chosen over safety. Therefore, PoW blockchains do not offer consensus finality, instead one needs to wait for several consecutive blocks to be probabilistically certain that a transaction cannot be reverted. Simplified Payment Verification (SPV) mode [14] for clients can reduce the resource usage at the cost of decentralization.

ByzCoin [37] uses PoW for a separate identity chain to guard against Sybil attacks but uses a BFT protocol to order transactions. ByzCoin makes use of collective signatures (CoSi) [38] and a balanced tree for the communication flow. CoSi makes use of aggregate signatures by constructing a Schnorr multisignature. However, CoSi needs multiple communication round-trips to generate the multi-signature and assumes a synchronous network.

Tendermint [10], [29], used in Cosmos, uses Proof-of-Stake (PoS), where voting power is based on the amount of cryptocurrency owned by each replica. Because block times are short, in the order of seconds, there is a limited number of validators Tendermint can have because finality needs to be reached for each block. It is also not resistant to cartel forming, which allows those with a lot of cryptocurrencies to work together to control the network.

Other protocols use a randomized approach. Ouroboros [12], HoneyBadger [39], Dumbo [40] and BEAT [41] use distributed coin flipping for consensus. HoneyBadger [39] uses threshold encryption [30] for censorship resilience. Algorand [11] uses Verifiable Random Functions [42] to select a random committee for the next round. Avalanche [16], [43] uses meta-stability to reach consensus by sampling other replicas without any leader. While Avalanche is lightweight and scalable, it needs to be able to sample all other validators directly. The number of connections one can open in a browser without performance loss is limited. RoBFT supports propagation of votes over multiple hops.

Permissioned blockchains such as Hyperledger Fabric [44] have closed membership and often use a BFT consensus protocol to order transactions. For example BFT-SMART in HyperLedger Fabric [9], [28]. The first known BFT protocol is Practical Byzantine Fault Tolerance (PBFT) [8]. Other protocols bring improvements to the original PBFT protocol. Zyzzyva [45] uses speculative execution which improves latency and throughput if there are no Byzantine replicas. However, its performance drops significantly if this premise does not hold. 700BFT [46] provides an abstraction for these BFT algorithms. These protocols are targeting a small number of replicas in a local network. They generally work in two phases: the first guarantees proposal uniqueness, and the second guarantees that a new leader can convince replicas to vote for a safe proposal. HotStuff [13] proposed a three-phase protocol to reduce complexity and simplify leader replacement. This makes HotStuff more scalable. All these algorithms use a leader to drive the protocol. When the leader is malicious, the performance can degrade quickly [31]. GeoBFT [47] is a topology-aware, decentralized consensus protocol, designed for geo-distributed scalability. RoBFT does not use a leader and replicas communicate only to a subset of the other replicas using a gossip-like protocol. Another approach is to use a trusted hardware component [48]–[52]. These are faster and less computationally intensive but require specialized hardware to be present. Moreover, trusted execution environments have been broken in the past [53], [54].

AWARE [55] is a variant of BFT-SMaRt that dynamically changes the voting power of a replica depending on its latency over time, decreasing the consensus latency. RoBFT gives every replica equal voting power. In future work, RoBFT could be extended to associate a weight to each vote. While we believe this would be especially beneficial for our target environment with mobile and unreliable clients, special care will have to be given to ensure safety will stay intact.

There are several proposals to improve the performance and response time of BFT. StreamChain [56] reaches consensus over a stream of transactions instead of blocks. Fabric-CRDT [57] uses CRDTs to support concurrent transactions to occur in the same block, using the built-in conflict resolution of CRDTs to resolve the conflict automatically. Other approaches also borrow from CRDTs: PnyxDB [24] supports commuting transactions to be applied out-of-order. A novel design for gossip in Fabric [58] improves the block propagation latency and bandwidth. Other approaches dynamically adapt the number of faults the system can withstand in reaction to threat level changes [59]. While these improvements make BFT faster, none of them try to reduce the infrastructure requirements to be able to easily set up an untrusted peer-to-peer network.

The Lightning Network or state channels for Bitcoin [60] or Ethereum [61], [62] are *off-chain* protocols that run on top of a blockchain. A new state channel between known participants is created by interacting with the blockchain. After its creation, participants can use this channel to execute state transitions by collectively signing the new state. These transactions do not involve the blockchain and have fast confirmation times and no transaction costs. However, state channels assume all participants to be always online and honest. If this is violated, the underlying blockchain needs to be used to resolve the conflict, or a trusted third party can be used [63]. RoBFT uses a similar state-transitioning protocol where only the latest collectively agreed state needs to be stored. However, RoBFT can tolerate both failing and malicious replicas, without resorting to a blockchain or a trusted third party.

VII. Conclusion

In this paper, we presented RoBFT. A browser-based middleware for decentralized, community-driven web applications. RoBFT uses a client-centric, leaderless BFT consensus protocol, combined with a robust and efficient state-based synchronization protocol. RoBFT uses an optimized BLS scheme for efficient computation and storage of signatures. It supports a client-centric, browser-based, state-based, permissioned datastore with a low infrastructure and storage footprint for small-scale, citizen-driven networks. RoBFT offers consistent and robust confirmation times to achieve finality of transactions in the order of seconds, even in failure settings and Byzantine environments. In contrast to traditional blockchains, RoBFT does not store a transaction log or blockchain, keeping the overall storage footprint small.

References

[1] P. Garcia Lopez, A. Montresor, D. Epema, A. Datta, T. Higashino, A. Iamnitchi, M. Barcellos, P. Felber, and E. Riviere, "Edge-centric computing: Vision and challenges," *SIGCOMM Comput. Commun. Rev.*, 2015.

[2] K. Jannes, B. Lagaisse, and W. Joosen, "The web browser as distributed application server: Towards decentralized web applications in the edge," in *EdgeSys*, 2019.

[3] A. Madhusudan, I. Symeonidis, M. A. Mustafa, R. Zhang, and B. Preneel, "SC2Share: Smart contract for secure car sharing," in *ICISSP*, 2019.

[4] K. Jannes, B. Lagaisse, and W. Joosen, "You don't need a ledger: Lightweight decentralized consensus between mobile web clients," in *SERIAL*, 2019.

[5] A. van der Linde, P. Fouto, J. a. Leitão, N. Preguiça, S. Castiñeira, and A. Bieniusa, "Legion: Enriching internet services with peer-to-peer interactions," in *WWW*, 2017.

[6] M. Kleppman and A. R. Beresford, "Automerge: Real-time data sync between edge devices," 2018.

[7] K. Jannes, B. Lagaisse, and W. Joosen, "OWebSync: Seamless synchronization of distributed web clients," *IEEE Trans. on Parallel and Distributed Systems*, 2021.

[8] M. Castro, B. Liskov *et al.*, "Practical byzantine fault tolerance," in *OSDI*, 1999.

[9] A. Bessani, J. Sousa, and E. E. P. Alchieri, "State machine replication for the masses with BFT-SMART," in *DSN*, 2014.

[10] E. Buchman, J. Kwon, and Z. Milosevic, "The latest gossip on BFT consensus," 2018.

[11] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich, "Algorand: Scaling byzantine agreements for cryptocurrencies," in *SOSP*, 2017.

[12] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *CRYPTO*, 2017.

[13] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, "HotStuff: BFT consensus with linearity and responsiveness," in *PODC*, 2019.

[14] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[15] V. Buterin *et al.*, "A next-generation smart contract and decentralized application platform," ethereum.org, White paper, 2013.

[16] T. Rocket, M. Yin, K. Sekniqi, R. van Renesse, and E. G. Sirer, "Scalable and probabilistic leaderless BFT consensus through metastability," 2019.

[17] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the weil pairing," in *ASIACRYPT*, 2001.

[18] S. Fromhart and L. Therattil, "Making blockchain real for customer loyalty rewards programs," Deloitte, Tech. Rep., 2016.

[19] C. Dwork, N. Lynch, and L. Stockmeyer, "Consensus in the presence of partial synchrony," *J. ACM*, 1988.

[20] C. Cachin, R. Guerraoui, and L. E. T. Rodrigues, *Introduction to Reliable and Secure Distributed Programming*. Springer, 2011.

[21] D. Cason, N. Milosevic, Z. Milosevic, and F. Pedone, "Gossip consensus," in *Middleware*, 2021.

[22] M. Shapiro, N. Perguiça, C. Baquero, and M. Zawirski, "Conflict-free replicated data types," in *SSS*, 2011.

[23] R. Merkle, "A digital signature based on a conventional encryption function," in *CRYPTO*, 1987.

[24] L. Bonniot, C. Neumann, and F. Taïani, "PnyxDB: a lightweight leaderless democratic byzantine fault tolerant replicated datastore," in *SRDS*, 2020.

[25] D. Boneh, M. Drijvers, and G. Neven, "Compact multi-signatures for smaller blockchains," in *ASIACRYPT*, 2018.

[26] D. Boneh, C. Gentry, B. Lynn, and H. Shacham, "Aggregate and verifiably encrypted signatures from bilinear maps," in *EUROCRYPT*, 2003.

[27] G. G. Gueta, I. Abraham, S. Grossman, D. Malkhi, B. Pinkas, M. Reiter, D.-A. Seredinschi, O. Tamir, and A. Tomescu, "SBFT: a scalable and decentralized trust infrastructure," in *DSN*, 2019.

[28] J. Sousa, A. Bessani, and M. Vukolic, "A byzantine fault-tolerant ordering service for the hyperledger fabric blockchain platform," in *DSN*, 2018.

[29] D. Cason, E. Fynn, N. Milosevic, Z. Milosevic, E. Buchman, and F. Pedone, "The design, architecture and performance of the tendermint blockchain network," in *SRDS*, 2021.

[30] V. Shoup, "Practical threshold signatures," in *Eurocrypt*, 2000.

[31] P.-L. Aublin, S. B. Mokhtar, and V. Quéma, "RBFT: Redundant byzantine fault tolerance," in *ICDCS*, 2013.

[32] A. van der Linde, J. a. Leitão, and N. Preguiça, "Practical client-side replication: Weak consistency semantics for insecure settings," *VLDB*, 2020.

[33] M. Barbosa, B. Ferreira, J. a. Marques, B. Portela, and N. Preguiça, "Secure conflict-free replicated data types," in *ICDCN*, 2021.

[34] C. Berger and H. P. Reiser, "WebBFT: Byzantine fault tolerance for resilient interactive web applications," in *DAIS*, 2018.

[35] ——, "Scaling byzantine consensus: A broad analysis," in *SERIAL*, 2018.

[36] K. J. O'Dwyer and D. Malone, "Bitcoin mining and its energy footprint," in *ISSC*, 2014.

[37] E. Kokoris-Kogias, P. Jovanovic, N. Gailly, I. Khoffi, L. Gasser, and B. Ford, "Enhancing bitcoin security and performance with strong consistency via collective signing," in *SEC*, 2016.

[38] E. Syta, I. Tamas, D. Visher, D. I. Wolinsky, P. Jovanovic, L. Gasser, N. Gailly, I. Khoffi, and B. Ford, "Keeping authorities "honest or bust" with decentralized witness cosigning," in *S&P*, 2016.

[39] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, "The honey badger of BFT protocols," in *CCS*, 2016.

[40] B. Guo, Z. Lu, Q. Tang, J. Xu, and Z. Zhang, "Dumbo: Faster asynchronous bft protocols," in *CCS*, 2020.

[41] S. Duan, M. K. Reiter, and H. Zhang, "BEAT: Asynchronous BFT made practical," in *CCS*, 2018.

[42] S. Micali, M. Rabin, and S. Vadhan, "Verifiable random functions," in *FOCS*, 1999.

[43] T. Rocket, "Snowflake to avalanche: A novel metastable consensus protocol family for cryptocurrencies," avalabs.org, White paper, 2018.

[44] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. W. Cocco, and J. Yellick, "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in *EuroSys*, 2018.

[45] R. Kotla, L. Alvisi, M. Dahlin, A. Clement, and E. Wong, "Zyzzyva: Speculative byzantine fault tolerance," in *SOSP*, 2007.

[46] P.-L. Aublin, R. Guerraoui, N. Knežević, V. Quéma, and M. Vukolić, "The next 700 bft protocols," *ACM Trans. Comput. Syst.*, 2015.

[47] S. Gupta, S. Rahnama, J. Hellings, and M. Sadoghi, "ResilientDB: Global scale resilient blockchain fabric," *VLDB*, 2020.

[48] G. S. Veronese, M. Correia, A. N. Bessani, L. C. Lung, and P. Verissimo, "Efficient byzantine fault-tolerance," *IEEE Trans. on Computers*, 2013.

[49] R. Kapitza, J. Behl, C. Cachin, T. Distler, S. Kuhnle, S. V. Mohammadi, W. Schröder-Preikschat, and K. Stengel, "CheapBFT: Resource-efficient byzantine fault tolerance," in *EuroSys*, 2012.

[50] T. Distler, C. Cachin, and R. Kapitza, "Resource-efficient byzantine fault tolerance," *IEEE Transactions on Computers*, 2016.

[51] J. Behl, T. Distler, and R. Kapitza, "Hybrids on steroids: SGX-based high performance BFT," in *EuroSys*, 2017.

[52] J. Liu, W. Li, G. O. Karame, and N. Asokan, "Scalable byzantine consensus via hardware-assisted secret sharing," *IEEE Trans. on Computers*, 2018.

[53] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown: Reading kernel memory from user space," in *USENIX Security*, 2018.
[54] P. Kocher, J. Horn, A. Fogh, , D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," in *S&P*, 2019.
[55] C. Berger, H. P. Reiser, J. Sousa, and A. Bessani, "Resilient wide-area byzantine consensus using adaptive weighted replication," in *SRDS*, 2019.
[56] Z. István, A. Sorniotti, and M. Vukolić, "StreamChain: Do blockchains need blocks?" in *SERIAL*, 2018.
[57] P. Nasirifard, R. Mayer, and H.-A. Jacobsen, "FabricCRDT: A conflict-free replicated datatypes approach to permissioned blockchains," in *Middleware*, 2019.
[58] N. Berendea, H. Mercier, E. Onica, and E. Riviere, "Fair and efficient gossip in Hyperledger Fabric," in *ICDCS*, 2020.
[59] D. S. Silva, R. Graczyk, J. Decouchant, M. Völp, and P. Esteves-Verissimo, "Threat adaptive byzantine fault tolerant state-machine replication," in *SRDS*, 2021.
[60] J. Lind, O. Naor, I. Eyal, F. Kelbert, E. G. Sirer, and P. Pietzuch, "Teechain: A secure payment network with asynchronous blockchain access," in *SOSP*, 2019.
[61] A. Miller, I. Bentov, S. Bakshi, R. Kumaresan, and P. McCorry, "Sprites and state channels: Payment networks that go faster than lightning," in *FC*, 2019.
[62] P. McCorry, C. Buckland, S. Bakshi, K. Wüst, and A. Miller, "You sank my battleship! a case study to evaluate state channels as a scaling solution for cryptocurrencies," in *FC*, 2020.
[63] P. McCorry, S. Bakshi, I. Bentov, S. Meiklejohn, and A. Miller, "Pisa: Arbitration outsourcing for state channels," in *AFT*, 2019.

APPENDIX A
SAFETY AND LIVENESS

This section sketches the proof that the algorithm provides safety and liveness. The protocol described before guarantees both safety and liveness when there are at least $2f+1$ honest replicas available.

A. Safety

Lemma 1 (Safety): Let $\mathfrak{R}$ be a cluster of n replicas with f Byzantine nodes and with $n > 3f$. If replicas $R_1, R_2 \in \mathfrak{R}$ are able to construct quorum certificates qc_1 for value $value_1$ and qc_2 for value $value_2$ at view v, then $value_1 = value_2$.

We will first prove this for the simplified case when both quorum certificates belong to the same round, and we will then prove that once a quorum certificate can be constructed, no more rounds can be started.

Lemma 2: If replicas $R_1, R_2 \in \mathfrak{R}$ are able to construct quorum certificates qc_1 and qc_2 for value $value_1$ and $value_2$ respectively with $qc_{1\ view} = qc_{2\ view}$ and $qc_{1\ round} = qc_{2\ round}$, then $value_1 = value_2$.

Proof: Assume two different replicas R_1 and R_2 have constructed a quorum certificate qc_1 and qc_2 for value $value_1$ and $value_2$ respectively with $qc_{1\ view} = qc_{2\ view}$ and $qc_{1\ round} = qc_{2\ round}$. They are constructed in the same round, so of the n possible votes, at least $n-f$ replicas have voted on $value_1$, and at least $n-f$ replicas have voted on $value_2$. Honest replicas will never vote twice in the same view and round. Therefore, at least $n-2f$ honest replicas have voted on $value_1$ and $n-2f$ *different* honest replicas have voted on $value_2$. In total, we have $(n-2f)+(n-2f)+f \equiv 2n-3f$ replicas that have voted. We defined $n \geq 3f+1$ before, which gives $2n-3f \geq 3f+2 \geq n+1$ replicas. This is a contradiction, there need to be at least $n+1$ replicas to construct two such certificates for different values, however, we only have n replicas. So the two values $value_1$ and $value_2$ have to be equal. □

Lemma 3: If replicas $R_1, R_2 \in \mathfrak{R}$ are able to construct quorum certificates qc_1 and qc_2 for value $value_1$ and $value_2$ respectively with $qc_{1\ view} = qc_{2\ view}$, then $qc_{1\ round} = qc_{2\ round}$.

Proof: Assume two different replicas R_1 and R_2 have constructed a quorum certificate qc_1 and qc_2 for value $value_1$ and $value_2$ respectively with $qc_{1\ view} = qc_{2\ view}$ and $qc_{1\ round} < qc_{2\ round}$. Since qc_1 is accepted, at least $n-f$ replicas vote on the proposed quorum certificate and at least $n-f$ replicas voted on $value_1$ in round $qc_{1\ round}$. The fact that $n-f$ replicas voted on the proposed quorum certificate means that at least $n-2f$ honest replicas observed $n-f$ votes for $value_1$. Of those votes, at least $n-2f$ are coming from honest replicas. The only way to now construct a quorum certificate qc_2 for $value_2$ is to start a new round. To start a new round, a replica needs to not have voted for the proposed quorum certificate qc_1, and observe a different winning value $value_2$. Yet, at least $n-2f$ honest replicas observed that at least $n-2f$ honest replicas think that $value_1$ is the winning value. This leaves only $2f$ replicas who can prefer another value $value_2$. By definition we have $n \geq 3f+1$. This means that in the worst case, $f+1$ honest replicas observe $f+1$ honest replicas thinking $value_1$ is the winning value, together with f Byzantine replicas. Value $value_2$ has only $2f$ supporting replicas, which is not enough to start a proposed quorum certificate. So, at least one replica currently supporting $value_1$ needs to switch votes in a future round. However, once a replica has voted for a proposed quorum certificate, it will not change their opinion unless it is convinced that a new valid round is started. So once $n-2f$ honest replicas are locked on a value, by voting on a proposed quorum certificate, it is impossible that a valid new round can be started. □

B. Liveness

When a new value is proposed, eventually the protocol will end and a valid quorum certificate is created for a new value. This value is not necessarily the first proposed value, and it is not even guaranteed that a specific value ever gets committed as long as other values continue to be proposed. Safety is always chosen over liveness. When there are not enough honest replicas online to reach a supermajority, no consensus can be reached and the protocol will simply block and wait for more votes. However, all those replicas do not need to be online at the same time, since the state is replicated to all available replicas over time, and votes can be verified by all replicas in the end.

Lemma 4 (Liveness): Let $\mathfrak{R}$ be a cluster of n replicas with f Byzantine nodes and with $n > 3f$. If an honest replica $R \in \mathfrak{R}$ proposes a new value at view v, eventually a replica will be able to construct a quorum certificate qc for some value at view v.

Lemma 5: If only a single replica $R \in \mathfrak{R}$ proposes a new value $value_1$, eventually a replica will be able to construct a valid quorum certificate qc.

Proof: As there is only a single proposed value, all honest replicas who observe this will cast their vote for that value. Eventually, one replica will observe $n - f$ votes for $value_1$ and a new proposed quorum certificate qc' will be constructed. Eventually, $n - f$ votes will be cast to this proposed quorum certificate qc' and a valid quorum certificate qc is constructed and $value$ is committed. □

Lemma 6: If x replicas $R_{1..x} \in \mathfrak{R}$ propose values $value_{1..x}$, and no Byzantine replicas vote twice in the same round, eventually a replica will be able to construct a valid quorum certificate qc.

Proof: Either a single value reaches a quorum, in which case the previous lemma holds. Or a split vote occurs and a new round will be started after $n - f$ votes are observed. All replicas will base their vote for this new round on the winning value that they observed from round 0. At least $n - f$ votes are known, and only f votes are still unknown. As long as not all votes are known to all voting replicas, the winning value might change. In each new round, either an unknown vote stays unknown, or it becomes known. In the former case, then the currently known votes will all be the same, and a proposed quorum certificate can be started. In the latter case, one extra vote is known, which might again result in the system ending up in a split vote, and a new round will be started. However, this last case can only happen at most f times. After $f + 1$ rounds, all replicas will have voted in round 0, and every replica will observe the same winning value, and a quorum certificate can be created. □

Lemma 7: If x replicas $R_{1..x} \in \mathfrak{R}$ propose values $value_{1..x}$, eventually a replica will be able to construct a valid quorum certificate qc.

Proof: If no Byzantine replicas vote twice in the same round, or only a single value is proposed, the previous two lemmas hold. If a split vote occurs, a new round will be started after $n-f$ votes are observed. f of those votes might belong to Byzantine replicas who can vote for multiple values. As a new round is only started after $n - f$ votes, a least $n - 2f$ honest votes are observed. No Byzantine replica can send conflicting votes to any of those $n - 2f$ honest replicas, as otherwise those replicas will detect this conflicting vote and exclude the Byzantine replica. If this happens repeatedly, at most f times, all Byzantine replicas are excluded and the previous lemma holds. Moreover, no Byzantine replica can continue to vote on values that are not the winning value. Each replica is only allowed to vote on the winning value or any other value that has at least support from $f + 1$ replicas in the previous round. All honest replicas converge to a single value, even with Byzantine replicas supporting other values. Because the protocol only looks to round 0 to determine the winning value. In the rounds after that, the f Byzantine replicas can support a different value, but if they do, they will be excluded as $f < f + 1$. This means that after at most $2f + 1$ rounds, a proposed quorum certificate can be started, which will be committed. □