Setchain: Improving Blockchain Scalability with Byzantine Distributed Sets and Barriers

Margarita Capretto*, Martín Ceresa*, Antonio Fernández Anta†, Antonio Russo†, César Sánchez*

*IMDEA Software Institute, Pozuelo de Alarcón, Madrid, Spain

†IMDEA Networks Institute, Leganés, Madrid, Spain

Email:{margarita.capretto,martin.ceresa,antonio.fernandez,antonio.russo,cesar.sanchez}@imdea.org

Abstract—**Blockchain technologies are facing a scalability challenge, which must be overcome to guarantee a wider adoption of the technology. This scalability issue is mostly caused by the use of consensus algorithms to guarantee the total order of the chain of blocks (and of the operations within each block). However, total order is often overkilling, since important advanced applications of smart-contracts do not require a total order of *all* the operations. Hence, if a more relaxed partial order (instead of a total order) is allowed under certain safety conditions, a much higher scalability can be achieved.**

In this paper, we propose a distributed concurrent data type, called *Setchain*, that allows implementing this partial order and increases significantly blockchain scalability. A Setchain implements a *grow-only set object* whose elements are not totally ordered, unlike conventional blockchain operations. When convenient, the Setchain allows forcing a synchronization barrier that assigns permanently an epoch number to a subset of the latest elements added. With the Setchain, operations in the same epoch are not ordered, while operations in different epochs are. We present different Byzantine-tolerant implementations of Setchain, prove their correctness and report on an empirical evaluation of a direct implementation.

Our results show that Setchain is orders of magnitude faster than consensus-based ledgers to implement grow-only sets with epoch synchronization. Since the Setchain barriers can be synchronized with block consolidation, Setchain objects can be used as a *sidechain* to implement many smart contract solutions with much faster operations than on basic blockchains.

Index Terms—**Distributed systems, blockchain, byzantine distributed objects, consensus, Setchain.**

I. INTRODUCTION

A. The Problem

Distributed ledgers (also known as *blockchains*) were first proposed by Nakamoto in 2009 [21] in the implementation of Bitcoin, as a method to eliminate trustable third parties in electronic payment systems. Modern blockchains incorporate smart contracts [28], [33], which are state-full programs stored in the blockchain that describe the functionality of the transactions, including the exchange of cryptocurrency. Smart contracts allow to describe sophisticated functionality, enabling many applications in decentralized finances (DeFi)[1], decentralized governance, Web3, etc.

The main element of all distributed ledgers is the "blockchain," which is a distributed object that contains,

packed in blocks, the ordered list of transactions performed on behalf of the users [14], [13]. This object is maintained by multiple servers without a central authority by using consensus algorithms that are resilient to Byzantine attacks.

However, a current major obstacle for a faster widespread adoption of blockchain technologies, and the deployment of new applications based on them, is their limited scalability, due to the delay introduced by Byzantine consensus algorithms [8], [31]. Ethereum [33], one of the most popular blockchains, is limited to less than 4 blocks per minute, each containing less than two thousand transactions. Bitcoin [21] offers even lower throughput. These figures are orders of magnitude slower than what many decentralized applications require, and can ultimately jeopardize the adoption of the technology in many promising domains. This limit in the throughput of the blockchain also increases the price per operation, due to the high demand to execute operations.

Consequently, there is a growing interest in techniques to improve the scalability of blockchains [20], [35]. Approaches include (1) the search for faster consensus algorithms [32], (2) the use of parallel techniques, like sharding [10], (3) building application-specific blockchains with Inter-Blockchain Communication capabilities [34], [19], or (4) extracting functionality out of the blockchain, while trying to preserve the guarantees of the blockchain (an approach known as "layer 2" [17]). Layer 2 (L2) approaches include the computation off-chain of Zero-Knowledge proofs [2], which only need to be checked on-chain (hopefully more efficiently) [1], the adoption of limited (but useful) functionality like *channels* (see, e.g., Lightning [22]), or the deployment of optimistic rollups (e.g., Arbitrum [18]) based on avoiding running the contracts in the servers (except when needed to annotate claims and resolve disputes).

In this paper, we propose an alternative approach to increase blockchain scalability that exploits the following observation. It has been traditionally assumed that cryptocurrencies require total order to guarantee the absence of double-spending. However, in reality, many useful applications and functionalities (including cryptocurrencies [16]) can tolerate more relaxed guarantees, where operations are only *partially ordered*. Hence, we propose a Byzantine-fault tolerant implementation of a distributed grow-only set, equipped with the additional operation of introducing points of synchronization (where all servers agree on the contents of the set). Between barriers,

---

[1]As of December 2021, the monetary value locked in DeFi was estimated to be around $100B, according to Statista https://www.statista.com/statistics/1237821/defi-market-size-value-crypto-locked-usd/.

elements of the distributed set can temporarily be known by some but not all servers. We call this distributed data structure a Setchain. A blockchain $\mathcal{B}$ implementing Setchain (as well as blocks) can align the consolidation of the blocks of $\mathcal{B}$ with synchronizations, obtaining a very efficient set object as side data type, with the same Byzantine-tolerance guarantees that $\mathcal{B}$ itself offers.

There are two extreme implementations of a transaction set with epochs (like Setchain) in the context of blockchains:

a) Completely off-chain solution: One could implement a transaction set totally off-chain (either centralized or distributed), but the resulting implementation does not have the trustability and accountability guarantees that blockchains offer. One example of such an attempt to implement this kind of off-chain data types is *mempools*. Mempools (short for memory pools) are a P2P data type used by most blockchains to maintain a set of pending transactions. Mempools fulfill two objectives: (1) to prevent distributed attacks to the servers that mine blocks and (2) to serve as a pool of transaction requests from where block producers select operations. Nowadays, mempools are receiving a lot of attention, since they suffer from lack of accountability and are a source of attacks [26], [25], including front-running [9], [24], [30]. Our proposed data structure, Setchain, offers a much stronger accountability, because it is resilient to Byzantine attacks and the contents of the set that Setchain maintains is public and cannot be forged.

b) Completely on-chain solution: One could use a smart-contract implementing the Setchain data type. For example, consider the following implementation (in a language similar to Solidity), where **add** is used to add elements, and **epochinc** to increase epochs.

```
contract Epoch {
  uint public epoch = 0;
  set public the_set = emptyset;
  mapping(uint => set) public history;
  function add(elem data) public {
    the_set.add(data);
  }
  function epochinc() public {
    history[++epoch] = the_set.setminus(history);
  }
}
```

Since `epoch,the_set`, and `history` are defined **public** there is an implicit getter function for each of them[2]. One problem of this implementation is that every time we add an element, `the_set` gets bigger, which can affect the required cost to execute the contract. A second more important problem is that adding elements is *slow*—as slow as interacting with the blockchain—while our main goal is to provide a much faster data structure than the blockchain.

Our approach lies in between these two extremes. For any given blockchain $\mathcal{B}$, we propose an implementation of Setchain that (1) is much more efficient than implementing and executing operations directly in $\mathcal{B}$; (2) offers the same decentralized guarantees against Byzantine attacks than $\mathcal{B}$, and

[2]In a public blockchain this function is not needed, since the set of elements can be directly obtained from the state of the blockchain.

(3) can be synchronized with the evolution of $\mathcal{B}$, so contracts could potentially inspect the contents of the Setchain. In a nutshell, these goals are achieved by using faster operations for the coordination among the servers (namely, reliable broadcast) for non-synchronized element insertions, and use only a consensus like algorithm for epoch changes.

B. Motivation

Setchain potential applications include:

1) Mempool: As mentioned above, user requests to execute operations/trasactions in a blockchain are stored in a mempool before they are chosen by miners. Once mined, the transactions executed are public, but the additional information, including the time of insertions in the mempool, is lost. Recording and studying the evolution of mempools would require an additional object serving as a mempool *log system*. This storage must be fast enough to record every attempt of interaction with the mempool without affecting the underlying blockchain's performance, and thus, it should be much faster than the blockchain itself.

2) Front-running: Mempools encode information about what it is about to happen in blockchains, so anyone observing them can predict the next operations to be mined, and take actions to their benefit. *Front-running* is the action of injecting transactions to be executed before the observed transaction request. This has been reported to be a relevant problem in decentralized exchanges [9], [30]. More specifically, an operation request added to the mempool includes the smart contract to be invoked, the maximum amount of *gas* that the user is willing to pay for the execution (*gas cap*), and a *fee* to pay the miner. Miners create blocks by choosing an attractive subset of transactions from the mempool, attempting to maximize their profit, based on the gas cap and the fee. Hence, the higher the fee and the lower the gas cap, the more likely a request is to be chosen by miners. Users that observe the mempool can inject new operations with higher priority (front-running) by offering a higher fee. Detecting front-running attacks requires a bookkeeping mechanism. As of today, once blocks are mined and consolidated, there is no evidence of the accesses to the mempool. The Setchain data type can serve as a basic mechanism to build a mempool that is efficient and serves as a log of requests.

3) Scalability by L2 Optimistic Rollups: Optimistic rollups, like Arbitrum [18], are based on the idea that a set of computing entities can safely compute outside the blockchain and post updates on the evolution of a smart contract. This is an optimistic strategy where each user can propose what the next state of the contract would be. After some time, the contract on-chain assumes that a given proposed step is correct and executes the effects claimed. A conflict resolution algorithm, also part of the contract on-chain, is used to resolve disputes. This protocol does not require the strict total order that a blockchain guarantees, but only a record of the actions proposed. Conflict resolutions are reduced to checking that the execution of a smart contract would produce the claimed effect, which can be part of the validation of the claim, and

could be performed by the maintainers of the Setchain data type.

4) Sidechain Data: Finally, Setchain can also be used as a general side-chain service used to store and modify data synchronized with the blocks. Applications that require only to update information in the storage space of a smart contract, like digital registries, can benefit from faster (and therefore cheaper) methods to manipulate the storage without invoking expensive blockchain operations.

C. Contributions.

In summary, the contributions of the paper are the following:
- the design and implementation of a side-chain data structure called *distributed Setchain*.
- several implementation of Setchain, providing different levels of abstraction and algorithmic implementation improvements.
- an empirical evaluation of a prototype implementation which suggests that Setchain is several orders of magnitude faster than consensus.

The rest of the paper is organized as follows. Section II contains the preliminaries. Section III describes the intended properties of Setchain. Section IV describes three different implementations of Setchain, and Section V proves the correctness of the fastest algorithm. Section VI discusses an empirical evaluation of our implementations of the different algorithms. Section VII shows how to make the use of Setchain more robust against Byzantine servers. Finally, Section VIII concludes the paper.

II. PRELIMINARIES

In this section, we present the model of computation as well as the building blocks used in our Setchain algorithms.

A. Model of Computation

We consider a distributed system consisting of processes—clients and servers—with an underlying communication graph in which each process can communicate with every other process. The computation proceeds *asynchronously*, and the communication is performed using *message passing*. Each process computes independently and at its own speed, and the internals of each process remain unknown to other processes. Message transfer delays are arbitrary but finite and also remain always unknown to processes. The intention is that servers will communicate among themselves to implement a distributed data type with certain guarantees, and clients can communicate with servers to exercise the data type.

Processes can fail arbitrarily, but the number of failing servers is bounded by f, and the total number of servers, n, is at least $3f + 1$. For clients we assume that any of them can be Byzantine and we explicitly state when we assume a client to be correct. We assume *reliable channels* between non-Byzantine (correct) processes. In other words, no message is lost, duplicated or modified.

Each process (client or server) has a pair of public and private keys. The public keys have been distributed reliably to all the processes that may interact with each other. Therefore, we discard the possibility of spurious or fake processes. We assume that messages are authenticated, so that messages corrupted or fabricated by Byzantine processes are detected and discarded by correct processes [7]. As result, communication between correct processes is reliable but asynchronous.

Finally, we also assume that there is a mechanism for clients to create "valid objects" and for servers to locally check whether an object is valid. In the context of blockchains, using public-key crytography clients can sign well-formed objects and servers can locally and efficiently check the signatures.

B. Building Blocks

We will use four building blocks in our implementations of Setchain: Byzantine Reliable Broadcast [3], [23], Byzantine Atomic Broadcast [11], Byzantine Distributed Grow-Only Sets [5] and Set Binary Consensus, as described in RedBelly [6]. We briefly describe each separately.

1) Byzantine Reliable Broadcast (BRB): The BRB service allows to broadcast messages to a set of processes guaranteeing that messages sent by correct processes are eventually received by *all* correct processes and that all correct processes eventually receive *the same* set of messages. The service provides a primitive BRB.Broadcast(m) for sending messages and an event BRB.Deliver(m) for receiving messages. Some important properties of BRB are:
- **BRB-Validity:** If a correct process p_i executes BRB.Deliver(m) and m belongs to a correct process p_j, then p_j executed BRB.Broadcast(m) in the past.
- **BRB-Termination:** If a correct process p executes BRB.Deliver(m), then all correct processes (including p) eventually execute BRB.Deliver(m).

Note that BRB does not guarantee the delivery of messages in the same order to two different correct participants.

2) Byzantine Atomic Broadcast (BAB): The BAB service extends BRB with an additional guarantee: a total order of delivery of the messages. BAB provides the same operation and event as BRB, which we will rename as BAB.Broadcast(m) and BAB.Deliver(m). In addition to the guarantees provided by BRB services, BAB services also provide:
- **Uniform Total Order:** If processes p and q both BAB.Deliver(m) and BAB.Deliver(m'), then p delivers m before m', if and only if q delivers m before m'.

Solving atomic broadcast has been proven to be as hard as consensus [11], and thus, is subject to the same limitations [15].

3) Byzantine Distributed Grow-only Sets (DSO): Sets are one of the most basic and fundamental data structures in computer science, which typically includes operations for adding and removing elements. Adding and removing operations do not commute, and thus, distributed implementations require additional mechanisms to keep replicas synchronized to prevent conflicting local states. One solution is to allow only additions, forbidding removals. The resulting data-type is called a grow-only set. A grow-only set is a conflict-free

replicated data structure [27] that behaves as a set in which elements can only be added but not removed.

Let A be an alphabet of values. A grow-only set *GS* is a concurrent object maintaining an internal set $GS.S \subseteq A$ offering two operations for any process p:

- $GS.\texttt{add}(r)$: adds an element $r \in A$ to the set $GS.S$.
- $GS.\texttt{get}()$: retrieves the internal set of elements $GS.S$.

Initially, the set $GS.S$ is empty. A Byzantine distributed grow-only set object (DSO) is a concurrent grow-only set implemented in a distributed manner [5] and tolerant to Byzantine attacks. Some important properties of these DSOs are:

- **Byzantine Completeness**: All $\texttt{get}()$ and $\texttt{add}()$ operations invoked by correct clients eventually complete.
- **DSO-AddGet**: All $\texttt{add}(r)$ operations will eventually result in r being in the set returned by *all* $\texttt{get}()$.
- **DSO-GetAdd**: Each element r returned by $\texttt{get}()$ was added using $\texttt{add}(r)$ in the past.

4) Set Binary Consensus (SBC): SBC, introduced in Red-Belly [6], is a Byzantine-tolerant distributed problem, similar to consensus. In SBC, each participant proposes a set of elements (in the particular case of RedBelly, a set of transactions). After SBC finishes, all correct servers agree on a set of valid elements which is guaranteed to be a subset of the union of the proposed sets. An intuitive way to understand SBC is that it efficiently runs binary consensus to agree on the sets proposed by each participant, such that if the outcome is positive then the set proposed is included in the final set consensus. Some properties of SBC are:

- **SBC-Termination**: every correct process eventually decides a set of elements.
- **SBC-Agreement**: no two correct process decide different sets of elements.
- **SBC-Validity**: when SBC is used on sets of transactions, the decided set of transactions is a valid non-conflicting subset of the union of the proposed sets.
- **SBC-Nontriviality**: if all processes are correct and propose an identical (valid non-conflicting, if transactions) set of elements, then this set is the decided set.

In [6], the authors show that the RedBelly algorithm, solves SBC in a system with partial synchrony: there is an unknown global stabilization time after which communication is synchronous. Then, they propose to use SBC to replace consensus algorithms in blockchains. They seek to improve scalability, because all transactions to be included in the next block can be decided with one execution of the SBC algorithm. It can be guaranteed that a block contains only valid non-conflicting transactions by applying a deterministic function that totally orders the decided set of transactions, and by using that order remove invalid or conflicting transactions.

Our use of SBC is different from implementing a blockchain. We use it to synchronize the barriers between local views of distributed grow-only sets. To guarantee that all elements are eventually assigned to epochs, we need the following property.

- **SBC-Censorship-Resistance**: there is a time τ after which, if the proposed sets of all correct processes contain

the same element e, then e will be in the decided set. In RedBelly, this property holds because after the global stabilization time, all set consensus rounds decide sets from correct processes.

III. THE SETCHAIN DISTRIBUTED DATA STRUCTURE

Our goal is to devise a distributed Byzantine-fault tolerant data structure, Setchain, that implementing a grow-only set together with synchronization barriers. A key concept of Setchains is the *epoch* number, which is a global counter that the distributed data structure maintains. The "synchronization barrier" is realized as an epoch change: the epoch number is increased and the elements added to the grow-only set since the previous barrier are stamped with the new epoch number.

A. API and Server State of the Setchain

We consider a universe E of elements that client processes can inject into the set. We also assume that servers can locally validate an element $e \in E$. A **Setchain** with elements in E is a distributed data structure where a set of server nodes, $\mathbb{D}$, maintain:

- a set $\texttt{the\_set} \subseteq E$ of elements added;
- a natural number $\texttt{epoch} \in \mathbb{N}$ that denotes the latest epoch;
- a map $\texttt{history} : [1 \dots \texttt{epoch}] \to \mathcal{P}(E)$, that describes the sets of elements that have been stamped with an epoch number ($\mathcal{P}(E)$ denotes the power set of E.)

Each server node $v \in \mathbb{D}$ supports three operations, available to any client process:

- $v.\texttt{add}(e)$: requests that element e is added to the set $\texttt{the\_set}$.
- $v.\texttt{get}()$: returns the values of $\texttt{the\_set}$, $\texttt{history}$, and $\texttt{epoch}$, as seen by v.
- $v.\texttt{epoch\_inc}(h)$ triggers an epoch change (i.e., a synchronization barrier). It must hold that $h = \texttt{epoch} + 1$.

Informally, a client process p invokes a $v.\texttt{get}()$ operation in node v to obtain (S, H, h), which is v's view of set $v.\texttt{the\_set}$ and map $v.\texttt{history}$, with domain $[1 \dots h]$. Process p invokes $v.\texttt{add}(e)$ to insert a new element e in $v.\texttt{the\_set}$, and a $v.\texttt{epoch\_inc}(h+1)$ to request an epoch increment. At server v, the set $v.\texttt{the\_set}$ contains the elements that have been added, including those that have not been assigned an epoch yet, while $v.\texttt{history}$ contains only those elements that have been assigned an epoch. A typical scenario is that an element $e \in E$ is first perceived by v to be in $\texttt{the\_set}$, to eventually be stamped and copied to $\texttt{history}$ in an epoch increment. However, as we will see, some implementations allow other ways to insert elements, in which v gets to know e for the first time during an epoch change. The operation $\texttt{epoch\_inc}()$ initiates the process of collecting elements in $\texttt{the\_set}$ at each node and collaboratively decide which ones are stamped with the current epoch.

Initially, both $\texttt{the\_set}$ and $\texttt{history}$ are empty and $\texttt{epoch} = 0$ in every correct server. Note that client processes can insert elements to $\texttt{the\_set}$ through $\texttt{add}()$, but

only servers decide how to update `history`, which client processes can only influence by invoking `epoch_inc()`.

Different servers may have, at a given point in time, different views of the set `the_set`. The Setchain data structure we propose here only provides eventual consistency guarantees, as defined below.

B. Desired Properties

We specify now properties that a correct implementation of a Setchain must have. We provide here a low-level specification assuming that every client interaction is initiated by a *correct* client and that it interacts with a *correct* server (recall that added elements e can be locally validated by servers). Later, in Section VII, we will describe a protocol that allows correct clients interact with Byzantine servers, which allows to provide a concurrent distributed object specification and implementation, at a price in performance.

We start by requiring from a Setchain that every `add`, `get`, and `epoch_inc` operation issued by a correct client to a correct server eventually terminates.

We say that element e is in epoch h in history H (e.g., returned by a `get` invocation) if $e \in H(h)$. We say that element e is in H if there is an epoch h such that $e \in H(h)$. The first property states that epochs only contain elements coming from the grow-only set.

Property 1 (Consistent Sets): Let $(S, H, h) = v.\text{get}()$ be the result of an invocation to a correct server v. Then, for each $i \leq h, H(i) \subseteq S$.

The second property states that every element added to a correct server is eventually returned in all future gets issued on the same server.

Property 2 (Add-Get-Local): Let $v.\text{add}(e)$ be an operation invoked by a correct client to a correct server v. Then, eventually all invocations $(S, H, h) = v.\text{get}()$ satisfy $e \in S$. The next property states that elements present in a correct server are propagated to all correct servers.

Property 3 (Add-Get): Let v, w be two correct servers, let $e \in E$ and let $(S, H, h) = v.\text{get}()$. If $e \in S$, then eventually all invocations $(S', H', h') = w.\text{get}()$ satisfy that $e \in S'$.

We assume in the rest of the paper that at every point in time, there is a future instant at which `epoch_inc()` is invoked and completed. This is a reasonable assumption in any real practical scenario, since it can be easily guaranteed using timeouts. Then, the following property states that all elements added are eventually assigned an epoch.

Property 4 (Eventual-Get): Let v be a correct server, let $e \in E$ and let $(S, H, h) = v.\text{get}()$. If $e \in S$, then eventually all invocations $(S', H', h') = v.\text{get}()$ satisfy $e \in H'$.

The previous three properties imply the following property.

Property 5 (Get-After-Add): Let $v.\text{add}(e)$ be an operation invoked by a correct client to a correct server v and valid element $e \in E$. Then, eventually all invocations $(S, H, h) = w.\text{get}()$ satisfy that $e \in H$, for all correct servers w.

An element can be at most in one epoch.

Property 6 (Unique Epoch): Let v be a correct server, $(S, H, h) = v.\text{get}()$, and let $i, i' \leq h$ with $i \neq i'$. Then, $H(i) \cap H(i') = \emptyset$.

All correct server processes agree on the content of every epoch (a safety property).

Property 7 (Consistent Gets): Let v, w be correct servers, let $(S, H, h) = v.\text{get}()$ and $(S', H', h') = w.\text{get}()$, and let $i \leq \min(h, h')$. Then $H(i) = H'(i)$.

The two previous properties imply that no element can be in two different epochs even if the history sets are obtained from `get` invocations to two different (correct) servers. Then, Property 7 essentially states that the histories returned by two `get` invocations to correct servers are one the prefix of the other. For comparison, it is not necessarily true for the `the_set` part of `get` that two invocations return sets that are contained one in the other. The reason is that the Setchain data structure allows fast insertion of elements that takes time to propagate to all correct nodes, so if two elements e and e' are inserted at two different correct servers, some correct servers may know e but not e', and vice versa.

Finally we require that every element in the history comes from the result of a client adding the element.

Property 8 (Add-before-Get): Let v be a correct server, $(S, H, h) = v.\text{get}()$, and $e \in S$. Then, there was an operation $w.\text{add}(e)$ in the past.

IV. Implementations

In this section, we describe implementations of the Setchain data structure that satisfy the properties in Section III. We first describe a centralized sequential implementation, and then three distributed implementations. The first distributed implementation is built using a Byzantine distributed grow-only set object (DSO) to maintain `the_set`, and Byzantine atomic broadcast (BAB) for epoch increments. The second distributed implementation is also built using DSO, but it uses Byzantine reliable broadcast (BRB) to announce epoch increments and set binary consensus (SBC) for epoch changes. Finally, the third implementation uses local sets, BRB for broadcasting elements and epoch increment announcements, and SBC for epoch changes.

A. Sequential Implementation

Alg. 0 shows a centralized solution, which maintains two local sets, `the_set` and `history`, both initialized as empty sets. The set `the_set` records all added elements and `history` is implemented as a collection of pairs $\langle h, A \rangle$ where h is an epoch number and A is a set of elements. We use `history`(h) to refer to the set A in the pair $\langle h, A \rangle \in$ `history`. In this implementation, it is easy to maintain a local copy for `the_set` because there is a single node maintaining the Setchain. Additionally, the implementation keeps a natural number `epoch` that is incremented each time there is a new epoch. The implementation of the data structure is as follows: Add(e) checks that element e is valid and adds it to `the_set`. Get() returns (`the_set`, `history`, `epoch`). Returning the latest epoch number, `epoch`, allows clients invoke EpochInc(h) with $h = $ `epoch` $+ 1$.

B. Distributed Implementations

We present three distributed algorithms beginning with a data structure that uses off-the-self existing building blocks. We then present more efficient implementations.

1) First approach. DSO and BAB: Alg. 1 uses two external services: a DSO and BAB. We denote messages with the name of the message followed by its content as in "$epinc(h, proposal, i)$". The variable `the_set` is not a local set anymore, but a DSO initialized empty with Init() in line 2. The function Get() invokes the DSO Get() function (line 4) to fetch the set of elements. The function EpochInc(h) triggers the mechanism required to increment an epoch and reach a consensus on which elements should be in it. This process begins by computing a local *proposal* set, of those elements added but that have not been stamped with an epoch yet (line 14). The *proposal* set is then broadcasted using a BAB service alongside the epoch number h and the server node id i (line 15). Then, the server waits to receive exactly $2f + 1$ proposals, and keeps the set of elements E present in at least $f + 1$ proposals, which guarantees that each element $e \in E$ was proposed by at least one correct server. The use of BAB guarantees that every message sent by a correct server eventually reaches every other correct server in the *same order*, so all correct servers use the same set of $2f + 1$ proposals. Therefore, all correct servers arrive to the same conclusion, and the set E is added as epoch h in `history` in line 21.

Alg. 1, while easy to understand and prove correct, is not efficient. To start, in order to complete an epoch increment, it requires at least $3f + 1$ calls to EpochInc(h) to different servers, so at least $2f + 1$ proposals are received (the f Byzantine severs may not propose anything). Another source of inefficiency comes from the use of off-the-shelf building blocks. For instance, every time a DSO Get() is invoked, many messages are exchanged to compute a reliable local view of the set [5]. Similarly, every epoch change requires a DSO Get() in line 14 to create a proposal. Additionally, line 17 requires waiting for $2f + 1$ atomic broadcast deliveries to take place. The most natural implementations of BAB services solve one

---

Algorithm 0 Single server implementation.

1: **Init:** epoch $\leftarrow 0$,   history $\leftarrow \emptyset$
2: **Init:** the_set $\leftarrow \emptyset$
3: **function** GET( )
4:     **return** (the_set, history, epoch)
5: **function** ADD(e)
6:     **assert** $valid(e)$
7:     the_set $\leftarrow$ the_set $\cup \{e\}$
8: **function** EPOCHINC(h)
9:     **assert** $h \equiv$ epoch $+ 1$
10:     $proposal \leftarrow$ the_set $\setminus \bigcup_{k=1}^{\text{epoch}}$ history(k)
11:     history $\leftarrow$ history $\cup \{\langle h, proposal \rangle\}$
12:     epoch $\leftarrow$ epoch $+ 1$

Algorithm 1 Server i implementation using DSO and BAB

1: **Init:** epoch $\leftarrow 0$,   history $\leftarrow \emptyset$
2: **Init:** the_set $\leftarrow$ DSO.Init()
3: **function** GET( )
4:     **return** (the_set.Get(), history, epoch)
5: **function** ADD(e)
6:     **assert** $valid(e)$
7:     the_set.Add(e)
12: **function** EPOCHINC(h)
13:     **assert** $h \equiv$ epoch $+ 1$
14:     $proposal \leftarrow$ the_set.Get() $\setminus \bigcup_{k=1}^{\text{epoch}}$ history(k)
15:     BAB.Broadcast($epinc(h, proposal, i)$)
16: **upon** (BAB.Deliver($epinc(h, proposal, j)$)
17:     from $2f + 1$ different servers j for the same h) **do**
18:     **assert** $h \equiv$ epoch $+ 1$
19:     $E \leftarrow \{e : e \in proposal$
20:         for at least $f + 1$ different $(h, proposal, j)\}$
21:     history $\leftarrow$ history $\cup \{\langle h, E \rangle\}$
22:     epoch $\leftarrow$ epoch $+ 1$
23: **end upon**

Algorithm 2 Server i implementation using DSO, and reliably broadcast (BRB) and set consensus (SBC) (Red Belly primitives).

11: ...   ▷ Get and Add as in Alg. 1
12: **function** EPOCHINC(h)
13:     **assert** $h \equiv$ epoch $+ 1$
14:     BRB.Broadcast($epinc(h)$)
15: **upon** (BRB.Deliver($epinc(h)$) and $h <$ epoch $+ 1$) **do**
16:     **drop**
17: **end upon**
18: **upon** (BRB.Deliver(h) and $h \equiv$ epoch $+ 1$) **do**
19:     **assert** $prop[h] \equiv null$
20:     $prop[h] \leftarrow$ the_set.Get() $\setminus \bigcup_{k=1}^{\text{epoch}}$ history(k)
21:     SBC[h].Propose($prop[h]$)
22: **end upon**
23: **upon** (SBC[h].SetDeliver($propset$)
24:     and $h \equiv$ epoch $+ 1$) **do**
25:     $E \leftarrow \{e : e \in$ at least $f + 1$ different $propset[j]\}$
26:     history $\leftarrow$ history $\cup \{\langle h, E \rangle\}$
27:     epoch $\leftarrow$ epoch $+ 1$
28: **end upon**

---

consensus per message delivered (see Fig. 7 in [4]), which would make this algorithm very slow. We solve these problems in two alternative algorithms.

2) Second approach. Avoiding BAB: Alg. 2 improves the performance of Alg. 1 in several ways. First, it uses BRB to propagate epoch increments, so a client does not need to contact more than one server. Second, the use of BAB and the wait for the arrival of $2f + 1$ messages in line 17 of Alg. 1 is replaced by using a SBC algorithm, which allows solving

Algorithm 3 Server implementation using a local set, Byzantine reliable broadcast (BRB) and Byzanting set consensus (SBC), (Red Belly primitives).

1: **Init:** epoch $\leftarrow$ 0,   history $\leftarrow \emptyset$
2: **Init:** the_set $\leftarrow \emptyset$
3: **function** GET( )
4:     **return** (the_set, history, epoch)
5: **function** ADD(e)
6:     **assert** $valid(e)$ and $e \notin$ the_set
7:     BRB.Broadcast(add(e))
8: **upon** (BRB.Deliver(add(e))) **do**
9:     **assert** $valid(e)$
10:     the_set $\leftarrow$ the_set $\cup \{e\}$
11: **end upon**
12: **function** EPOCHINC(h)
13:     **assert** $h \equiv$ epoch $+ 1$
14:     BRB.Broadcast(epinc(h))
15: **upon** (BRB.Deliver(epinc(h)) and $h <$ epoch $+ 1$) **do**
16:     **drop**
17: **end upon**
18: **upon** (BRB.Deliver(epinc(h)) and $h \equiv$ epoch $+ 1$) **do**
19:     **assert** $prop[h] \equiv \emptyset$
20:     $prop[h] \leftarrow$ the_set $\setminus \bigcup_{k=1}^{\text{epoch}}$ history(k)
21:     SBC[h].Propose($prop[h]$)
22: **end upon**
23: **upon** (SBC[h].SetDeliver($propset$) and $h \equiv$ epoch $+ 1$) **do**
24:     $E \leftarrow \{e : e \in propset[j], valid(e) \wedge e \notin$ history$\}$
25:     history $\leftarrow$ history $\cup \{\langle h, E \rangle\}$
26:     the_set $\leftarrow$ the_set $\cup E$
27:     epoch $\leftarrow$ epoch $+ 1$
28: **end upon**

several consensus instances simultaneously.

On a more abstract level, what we want to happen when an EpochInc(h) is triggered is that new elements in the local set the_set of each correct node are stamped with the new epoch number and added to the set history. The DSO guarantees eventual consistency, but at a given moment in time, two invocations to Get may vary from server to server. However, we need to guarantee that for every epoch the set history is the same in every correct server. Alg 1 enforced this using BAB and counting enough received messages to guarantee that every stamped element e was sent by a correct node. Alg. 2 uses SBC to solve several independent consensus instances simultaneously, one on each participant's proposal. Line 14 broadcasts a message inviting servers to begin an epoch change. When a correct server receives an epoch change invitation, it builds a proposed set and sends this proposal to the SBC. Note that there is one instance of SBC per epoch change, identified by h. Finally, with SBC each correct server receives the same set of proposals (where each proposal is a set of elements). Then, every node applies the same function to a

set of proposals reaching the same conclusion on how to update history(h). The function applied is deterministic: preserving only elements e that are present in at least $f + 1$ proposed sets, so these elements are guaranteed to have been proposed by a correct server. Observe that Alg. 2 still requires several invocations of the DSO Get operation to build the local proposal, one at each server.

3) Final approach. BRB and SBC without DSOs: The last approach, shown in Alg. 3, avoids the cascade of messages that DSO Get calls require in Alg. 1 by dissecting the internals of the DSO, and incorporating the corresponding steps in the Setchain algorithm directly. This idea exploits the fact that *a correct* Setchain *server* is a *correct client* of the DSO, and there is no need for the DSO to be defensive (this illustrates that using Byzantine resilient building blocks does not compose efficiently, but exploring this general idea is out of the scope of this paper).

Alg. 3 implements the_set using a local set (line 2) to accumulate a local view of the elements added. Elements newly received (in Add(e)) will be communicated to other server nodes using BRB. At any given point in time two correct servers may have a different local set (due to pending BRB deliveries) but each element added in one will eventually be known to the other. The local variable history is only updated in line 25 with the result of a SBC round. Therefore, all correct servers will agree on the same sets, with all elements being valid, not previously stamped and proposed by some server. Valid elements, by assumption, have been provided by correct clients. Additionally, Alg. 3 updates the_set to account for elements that are new to the node, guaranteeing that all elements in history are also in the_set. Note that this opens the opportunity to add elements to the Setchain by proposing them during an epoch change without broadcasting them before. This is exploited in Section VI to speed up the algorithm even more. As a final note, Alg. 3 allows a Byzantine server to propose elements, which will be accepted as long as they are valid. This is equivalent to a correct client proposing an element through an Add() operation, which is then successfully propagated during the set consensus phase.

V. PROOF OF CORRECTNESS

We prove now the correctness of Alg. 3. We first show that all stamped elements are in the_set.

Lemma 1: For every correct server v, at the end of each function/upon, $\bigcup_h v.\text{history}(h) \subseteq v.\text{the\_set}$.

Proof: Let v be a server. The only way to add elements to $v.\text{history}$ is at line 25, which is followed by line 26 which adds the same elements to $v.\text{the\_set}$. The only other instruction that modifies $v.\text{the\_set}$ is line 10 which only makes the set grow. ∎

Lemma 1 directly implies that Alg. 3 satisfies Prop. 1 (*Consistent Sets*) from Section III.

Lemma 2: Let v be a correct server and e an element in $v.\text{the\_set}$. Then e will eventually be in $w.\text{the\_set}$ for every correct server w.

Proof: Initially, v.the_set is empty. There are two ways to add an element e to v.the_set: (1) At line 10, so e is valid and was received via a BRB.Deliver(add(e)). By Properties **BRB-Validity** and **BRB-Termination** of BRB (see Section II), every correct server w will eventually execute BRB.Deliver(add(e)), and then (since e is valid), w will add it to w.the_set in line 10. (2) At line 26, so element e is valid and was received as an element in one of the sets in *propset* from SBC[h].SetDeliver(*propset*) with $h = v$.epoch $+ 1$. By properties **SBC-Termination SBC-Agreement** and **SBC-Validity** of SBC (see Section II), all correct servers agree on the same set of proposals. Therefore, if v adds e then w either adds it or has it already in its w.history which implies by Lemma 1 that $e \in w$.the_set. In either case, e will eventually be in w.the_set. ∎

Lemma 2, and the code of function Add() and line 4 of function Get() in Alg. 3 imply Prop. 2 (*Add-Get-Local*) and Prop. 3 (*Add-Get*) in Section III. The following lemmas reason about how elements are stamped.

Lemma 3: Let v be a correct server and $e \in v$.history(h) for some h. Then, for any $h' \neq h$, $e \notin v$.history(h').

Proof: It follows directly from the check that e is not injected at v.history(h) if $e \in v$.history in line 25. ∎

Lemma 4: Let v and w be correct servers. At a point in time, let h be such that v.epoch $\geq h$ and w.epoch $\geq h$. Then v.history(h) = w.history(h).

Proof: The proof proceeds by induction on epoch. The base case is epoch $= 0$, which holds trivially since v.history(0) = w.history(0) = $\emptyset$. Variable epoch is only incremented in one unit in line 27, after history(h) has been changed in line 25 when $h =$ epoch$+1$. In that line, v and w are in the same phase on SBC (for the same h). By **SBC-Agreement**, v and w receive the same *propset*, both v and w validate all elements equally, and (by inductive hypothesis), for each $h' \leq$ epoch it holds that $e \in v$.history(h') if and only if $e \in w$.history(h'). Therefore, in line 25 both v and w update history(h) equally, and after line 27 it holds that v.history(epoch) = w.history(epoch). ∎

Lemma 5: Let v and w be correct servers. If $e \in v$.the_set. Then, eventually e is in w.history.

Proof: By Lemma 2 every correct server w will satisfy $e \in w$.the_set at some $t > \tau$. By assumption, there is a new EpochInc() after t (let the epoch number be h). If e is already in history(h') for $h' < h$ we are done, since from Lemma 4 in this case at the end of the SBC phase for h' every correct server node w has e in w.history(h'). If e is not in history at t then, **SBC-Censorship-Resistance** guarantees that the decided set will contain e. Therefore, at line 25 every correct server w will add e to w.history(h). ∎

Lemmas 4 and 5 imply that all elements will be stamped, that is, Prop. 3 (*Eventual-Get*). Prop. 5 follows from Prop. 3. Lemma 3 directly implies Prop. 6 (*Unique Epoch*). Finally, Lemma 4 is equivalent to Prop. 7 (*Consistent Gets*) from Section III.

Finally, we discuss Prop. 8 (*Add-before-Get*). If valid elements can only be created by correct clients, and correct clients only interact with server nodes using the API, then the only way to reveal e to server nodes is through Add(e) and therefore the property trivially holds. If, on the other hand, valid elements can be created by, for example Byzantine server nodes, then server nodes can inject elements in the_set and history of correct servers without using the Add function. They can either execut directly a BRB.Broadcast, like the one in line 7, or even directly injecting them via the SBC. In these cases, Alg. 3 satisfies a weaker version of (*Add-before-Get*) that states that elements returned by Get() are valid and are either added by Add(), by a BRB.Broadcast like that of line 7, or injected in the SBC phase. In fact, we will use this idea to propagate elements directly in the SBC phase to accelerate Alg. 3.

VI. EMPIRICAL EVALUATION

We have implemented the necessary building blocks (1) DSO server code, (2) Reliable Broadcast and (3) Set Binary Consensus servers, and using these building blocks we have implemented Alg. 2 and Alg. 3. Our prototype is written in Golang [12] 1.16 with message passing using ZeroMQ [29] over TCP. Our testing platform uses Docker running on a server with 2 Intel Xeon CPU processors at 3GHz with 36 cores and 256GB RAM, running Ubuntu 18.04 Linux64. Each Setchain server node was wrapped in a Docker container with no limit on CPU or RAM usage. Alg. 2 implements a Setchain and a DSO as two standalone executables that communicate using remote procedure calls on the internal loopback network interface of the Docker container. The RPC server and client are taken from the Golang standard library. For Alg. 3 everything resides in a single executable. For both algorithms, we evaluate two versions, one where each element inserted causes a broadcast and another where servers aggregate locally inserted elements until a maximum message size (of 10^6 elements) or a maximum element timeout (of 5s) is reached.

We evaluate empirically the following hypothesis:
- (H1): The maximum rate of elements that can be inserted is much higher than the maximum epoch rate.
- (H2): Alg. 3 performs better than Alg. 2.
- (H3): The aggregated versions perform better than the basic versions.
- (H4) Silent Byzantine servers do not affect dramatically the performance.
- (H5) The performance does not degrade over time.

To evaluate these hypotheses, we carried out the experiments described below and reported in Fig. 1. In all cases, operations are injected by clients running within the same Docker container. Resident memory was always enough such that in no experiment the operating system needed to recur to disk swapping. All the experiments consider deployments with 4, 7, or 10 server nodes, and each running experiment reported is taken from the average of 10 executions.

We tested first how many epochs per minute our Setchain implementations can handle. In these runs, we did not add any element and we incremented the epoch rate to find out the

(a) Maximum epoch changes

(b) Maximum elements added (no epochs)

(c) Maximum elements added (with epochs)

(d) Silent Byzantine effect in adds

(e) Time to stamp (Alg. 3)
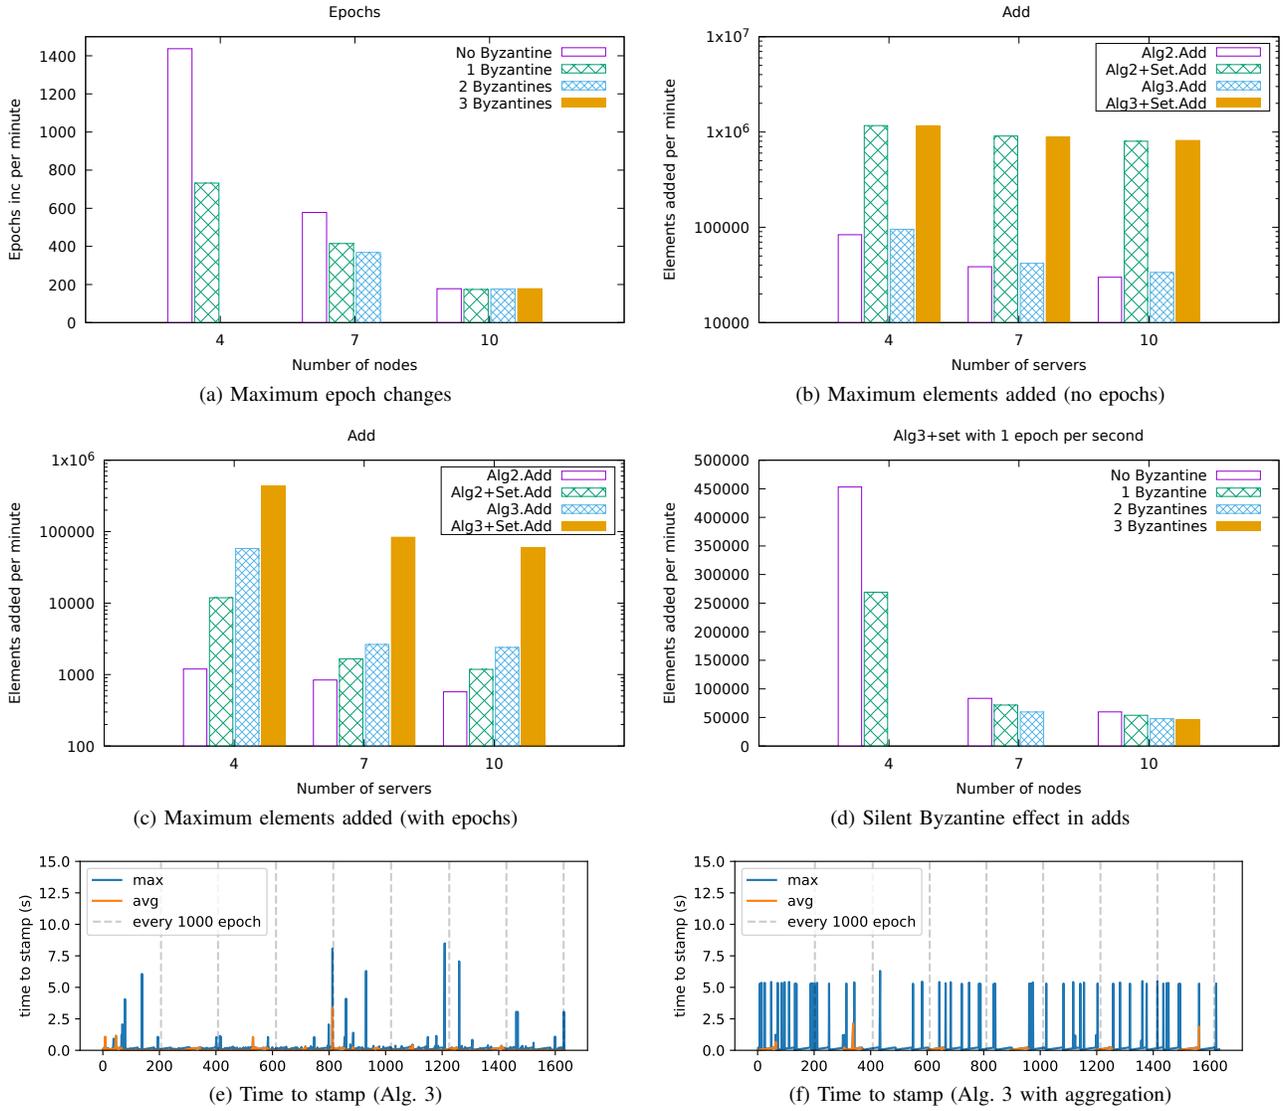
(f) Time to stamp (Alg. 3 with aggregation)

Fig. 1. Experimental results. Alg. 2+set and Alg. 3+set are the versions of the algorithms with aggregation. Byzantine servers are simply silent.

smallest latency between an epoch and the subsequent one. We run it with 4, 7, and 10 nodes, with and without Byzantines servers. This is reported in Fig. 1(a).

In our second experiment, we estimated empirically how many elements per minute can be added using our four different implementations of Setchain (Alg. 2 and Alg. 3 with and without aggregation), without any epoch increment. This is reported in Fig. 1(b). In this experiment Alg. 2 and Alg. 3 perform similarly. With aggregation Alg. 2 and Alg. 3 also perform similarly, but one order of magnitude better than without aggregation, confirming (H3). Putting together Fig. 1(a) and (b) one can conclude that sets are three orders of magnitude faster than epoch changes, confirming (H1).

In our third experiment, we compare the performance of our implementations combining epoch increments and insertion of elements. We set the epoch rate at 1 epoch change per second and calculated the maximum add ratio. The outcome

is reported in Fig. 1(c), which shows that Alg. 3 outperforms Alg. 2. In fact, Alg. 3+set even outperforms Alg. 2+set by a factor of roughly 5 for 4 nodes and by a factor of roughly 2 for 7 and 10 nodes. Alg. 3+set can handle 8x the elements added by Alg. 3 for 4 nodes and 30x for 7 and 10 nodes. The benefits of Alg. 3+set over Alg. 3 increase as the number of nodes increase because Alg. 3+set avoids the broadcasting of elements which generates a number of messages that is quadratic in the number of nodes in the network. This experiment confirms (H2) and (H3). The difference between Alg. 3 and Alg. 2 was not observable in the previous experiment (without epoch changes) because the main difference is in how servers proceed to collect elements to vote during epoch changes.

The next experiment explores how silent Byzantine servers affect Alg. 3+set. We implement silent Byzantine servers and run for 4,7 and 10 nodes with a an epoch change ratio of 1 per second, calculating the maximum add rate. This is reported

in Fig. 1(d). Silent byzantine servers degrade the speed for 4 nodes as in this case the implementation considers the silent server very frequently in the validation phase, but it can be observed that this effect is much smaller for larger number of servers, validating (H4).

In the final experiment, we run 4 servers for a long time (30 minutes) with an epoch ratio of 5 epochs per second and add requests to 50% of the maximum rate. We compute the time elapsed between the moment in which the client requests an add and the moment at which the element is stamped. Fig. 1(e) and (f) show the maximum and average times for the elements inserted in the last second. In the case of Alg. 3, the worst case during the 30 minute experiment was around 8 seconds, but the majority of the elements were inserted within 1 sec or less. For Alg. 3+set the maximum times were 5 seconds repeated in many occasions during the long run (5 seconds was the timeout to force a broadcast). This happens when an element fails to be inserted using the set consensus and ends up being broadcasted. In both cases the behavior does not degrade with long runs, confirming (H5).

As final note, in all our experiments all elements are signed and they are added to the Setchain along with their signature in order to simulate a signed element (like a transaction request) in a Blockchain. Therefore, each elements is composed of a constant amount of random bytes, a public key to simulate a sending address and a signature. The Setchain servers verify each element signature before accepting it in the add request. For the implementation, we leveraged the ed25519 crypto library of the Golang standard library.

In summary, considering as baseline the throughput of epochs (which internally performs a set consensus), implementing Setchain is three orders of magnitude more performant than consensus. Also, the performance of the algorithms is resilient to silent Byzantine servers.

VII. DISTRIBUTED PARTIAL ORDER OBJECTS (DPO)

The algorithms presented in Section IV and the proofs in Section V consider the case of a correct client contacting a correct server. Obviuosly, client processes do not know if they are contacting a Byzantine or correct process, so a client protocol is required to encapsulate the details of the distributed system. We describe now such a client protocol inspired by the one for DSO [5]. Implementing a correct client involves the exchange of several more messages than contacting a single server with a request, so we later describe a more efficient alternative.

The general idea of the client protocol is to interact with enough servers to guarantee enough correct nodes are reached to ensure the desired behavior. The API has methods that wait for a result (Get) and methods that do not require a response (EpochInc and Add). Alg. 4 shows the client protocol, which will be executed by correct clients. To contact with at least one correct server, we need to send $f + 1$ messages, as for $\text{Add}(e)$ and $\text{EpochInc}(h)$. Note that each message may trigger different broadcasts.

Algorithm 4 Correct client protocol for DPO (for Alg. 2 and 3).

1: **function** DPO.ADD(e)
2:    **call** Add(e) in $f + 1$ different servers.
3: **function** DPO.GET( )
4:    **call** Get() at least $3f + 1$ different servers.
5:    **wait** $2f + 1$ resp $s.(\texttt{the\_set}, \texttt{history}, \texttt{epoch})$
6:    $S \leftarrow \{e | e \in s.\texttt{the\_set} \text{ in at least } f + 1 \text{ servers } s\}$
7:    $H \leftarrow \emptyset$
8:    $i \leftarrow 1$
9:    $N \leftarrow \{s : s.\texttt{epoch} \geq i\}$
10:   **while** $\exists E : |\{s : s.\texttt{history}(i) = E\}| \geq f + 1$ **do**
11:     $H \leftarrow H \cup \{\langle i, E \rangle\}$
12:     $N \leftarrow N \setminus \{s : s.\texttt{history}(i) \neq E\}$
13:     $N \leftarrow N \setminus \{s : s.\texttt{epoch} = i\}$
14:     $i \leftarrow i + 1$
15:    **return** $(S, H, i - 1)$
16: **function** DPO.EPOCHINC(h)
17:    **call** EpochInc(h) in $f + 1$ different servers.

The wrapper algorithm for function Get can be split in two parts. First, the protocol contacts $3f + 1$ nodes, and waits for at least $2f + 1$ responses, because f Byzantine servers may refuse to respond. The response from server s is $(s.\texttt{the\_set}, s.\texttt{history}, s.\texttt{epoch})$. The protocol then computes S as those elements known to be in $\texttt{the\_set}$ by at least $f + 1$ servers (which includes at least one correct server). To compute the history H, the code goes incrementally epoch by epoch as long as at least $f + 1$ servers within the set N (which is initialized with all the servers that responded with non-empty histories) agree on a set E of elements in epoch i. Note that if at least $f + 1$ servers agree that E is the set of elements in epoch i, then E is indeed the set of stamped elements in epoch i. Then, in the loop we remove from N those servers that either do not know more epochs or that reported something different than E. Once this process ends, the sets S and H, and the latest processed epoch are returned. Note that it is guaranteed that $\texttt{history} \subseteq \texttt{the\_set}$.

Finally, we present an alternative faster optimistic client assuming that servers sign each epoch with a cryptographic signature. Correct servers sign cryptographically a hash of the set of elements in an epoch, and insert this hash in the Setchain as an element. Clients only perform **a single** Add(e) request to one server, hoping it will be a correct server. After some time, the client invokes a Get from **a single** server (which again can be Byzantine) and check whether e is in some epoch signed by (at least) $f + 1$ servers, in which case the epoch is correct and e has been successfully inserted and stamped. Note that this requires only one message per Add and one message per Get. Optimistic clients can repeat the process after a timeout if the contacted server did not respond.

VIII. CONCLUDING REMARKS

We presented a novel distributed data-type, called Setchain, that implements a grow-only set with epochs. The data struc-

ture tolerates Byzantine server nodes. We provided a low-level specification of desirable properties of Setchains and presented three distributed implementations, where the most efficient one uses Byzantine Reliable Broadcast and RedBelly set consensus. We also showed an empirical evaluation that suggests that the efficiency of the set (in terms of elements inserted) is three orders of magnitud higher than consensus.

Future work includes developing the motivating applications listed in the introduction, including (1) an implementation of mempools as Setchains for logging, (2) an encrypted version of mempools to prevent front-running, and (3) an improved implementation of L2 optimistic rollups. We will also study how blockchains with a synchronized Setchain as a side-chain can make smart-contracts more efficient. An important problem to solve is, of course, how to make clients of the Setchain pay for the usage (even if a much smaller fee than the blockchain itself).

There is also interesting future work from the point of view of the foundations of distributed systems. As Alg. 3 shows, Byzantine tolerant building blocks do not compose well from the point of view of efficiency. They have to be overly pessimistic and cannot exploit the fact that, when blocks are composed to build a correct node, both the client and the server of the block are same and can hence be considered correct. Studying this as a general principle is an interesting problem for future work. Also, it is easy to see that the Byzantine behavior at the level of Setchain server nodes can be modeled by simple interactions with BRB and SBC so one can build a (non-deterministic) implementation that simulates the behavior of any Byzantine server node. Such a modeling would greatly simplify formal proofs of Byzantine tolerant data structures.

REFERENCES

[1] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *Proc. of the IEEE Symp. on Security and Privacy (SP'14)*, pages 459–474, 2014.

[2] E. Ben-Sasson, A. Chiesa, E. Tromer, and M. Virza. Succinct non-interactive Zero Knowledge for a von Neumann architecture. In *Proc. of the 23rd USENIX Security Symp*, pages 781–796. USENIX, 2014.

[3] G. Bracha. Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75(2):130–143, 1987.

[4] T. D. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, mar 1996.

[5] V. Cholvi, A. Fernández Anta, C. Georgiou, N. Nicolaou, and A. Russo. Byzantine-tolerant distributed grow-only sets: Specification and applications. In *4th International Symposium on Foundations and Applications of Blockchain 2021 (FAB 2021)*, page 2:1–2:19, 2021.

[6] T. Crain, C. Natoli, and V. Gramoli. Red belly: A secure, fair and scalable open blockchain. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 466–483, 2021.

[7] F. Cristian, H. Aghili, R. Strong, and D. Volev. Atomic broadcast: from simple message diffusion to byzantine agreement. In *25th Int'l Symp. on Fault-Tolerant Computing*, pages 431–, 1995.

[8] K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, E. Gün Sirer, D. Song, and R. Wattenhofer. On scaling decentralized blockchains. In *Financial Crypto. and Data Security*, pages 106–125. Springer, 2016.

[9] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. *2020 IEEE Symp. on Security and Privacy (SP)*, pages 910–927, 2020.

[10] H. Dang, T. T. A. Dinh, D. Loghin, E.-C. Chang, Q. Lin, and B. C. Ooi. Towards scaling blockchain systems via sharding. In *Proc. of the 2019 Int'l Conf. on Management of Data*, SIGMOD '19, pages 123—-140. ACM, 2019.

[11] X. Défago, A. Schiper, and P. Urbán. Total order broadcast and multicast algorithms: Taxonomy and survey. *ACM Comput. Surv.*, 36(4):372–421, dec 2004.

[12] A. A. Donovan and B. W. Kernighan. *The Go Programming Language*. Adison-Wesley, 2015.

[13] A. Fernández Anta, C. Georgiou, M. Herlihy, and M. Potop-Butucaru. *Principles of Blockchain Systems*. Morgan & Claypool Publishers, 2021.

[14] A. Fernández Anta, K. Konwar, C. Georgiou, and N. Nicolaou. Formalizing and implementing distributed ledger objects. *ACM Sigact News*, 49(2):58–76, 2018.

[15] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, apr 1985.

[16] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlovic, and D. Seredinschi. The consensus number of a cryptocurrency. In *Proc. of the ACM Symp. on Principles of Distributed Computing, (PODC'19)*, pages 307–316. ACM, 2019.

[17] M. Jourenko, K. Kurazumi, M. Larangeira, and K. Tanaka. Sok: A taxonomy for layer-2 scalability related protocols for cryptocurrencies. *IACR Cryptol. ePrint Arch.*, 2019:352, 2019.

[18] H. Kalodner, S. Goldfeder, X. Chen, S. M. Weinberg, and E. W. Felten. Arbitrum: Scalable, private smart contracts. In *27th USENIX Security Symposium*, pages 1353–1370. USENIX Assoc., 2018.

[19] J. Kwon and E. Buchman. Cosmos whitepaper, 2019.

[20] Z. Mahdi, M. Movahedi, and M. Raykova. Rapidchain: Scaling blockchain via full sharding. In *Proc. of the ACM SIGSAC Conf. on Computer and Comm. Security*, CCS'18, pages 931—-948. ACM, 2018.

[21] S. Nakamoto. Bitcoin: a peer-to-peer electronic cash system, 2009.

[22] J. Poon and T. Dryja. The bitcoin lightning network: Scalable off-chain instant payments, 2016.

[23] M. Raynal. *Fault-Tolerant Message-Passing Distributed Systems: An Algorithmic Approach*. 01 2018.

[24] Robinson, Dan and Konstantopoulos, Georgios. Ethereum is a dark forest, 2020.

[25] M. Saad, L. Njilla, C. Kamhoua, J. Kim, D. Nyang, and A. Mohaisen. Mempool optimization for defending against DDoS attacks in PoW-based blockchain systems. In *2019 IEEE Int'l Conf. on Blockchain and Cryptocurrency (ICBC)*, pages 285–292, 2019.

[26] M. Saad, M. T. Thai, and A. Mohaisen. POSTER: Deterring DDoS attacks on blockchain-based cryptocurrencies through mempool optimization. In *Proc. of the 2018 on Asia Conf. on Computer and Communications Security*, ASIACCS '18, pages 809—811. ACM, 2018.

[27] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. Convergent and Commutative Replicated Data Types. *Bulletin- European Association for Theoretical Computer Science*, (104):67–88, June 2011.

[28] N. Szabo. Smart contracts: Building blocks for digital markets. *Extropy*, 16, 1996.

[29] The ZeroMQ authors. Zeromq, 2021. https://zeromq.org.

[30] C. F. Torres, R. Camino, and R. State. Frontrunner jones and the raiders of the Dark Forest: An empirical study of frontrunning on the Ethereum blockchain. In *30th USENIX Security Symp. (USENIX Security 21)*, pages 1343–1359. USENIX, 2021.

[31] S. Tyagi and M. Kathuria. *Study on Blockchain Scalability Solutions*, page 394–401. ACM, 2021.

[32] K. Wang and H. S. Kim. Fastchain: Scaling blockchain system with informed neighbor selection. In *2019 IEEE International Conference on Blockchain (Blockchain)*, pages 376–383, 2019.

[33] G. Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151:1–32, 2014.

[34] G. Wood. Polkadot: Vision for a heterogeneous multi-chain framework. *White Paper*, 21, 2016.

[35] C. Xu, C. Zhang, J. Xu, and J. Pei. Slimchain: Scaling blockchain transactions through off-chain storage and parallel processing. *Proc. VLDB Endow.*, 14(11):2314–2326, jul 2021.