


Setchain: Improving Blockchain Scalability with Byzantine Distributed Sets and Barriers

Margarita Capretto ✉ 

IMDEA Software Institute, Pozuelo de Alarcón, Madrid, Spain

Martín Ceresa ✉ 

IMDEA Software Institute, Pozuelo de Alarcón, Madrid, Spain

Antonio Fernández Anta ✉ 

IMDEA Networks Institute, Leganés, Madrid, Spain

Antonio Russo ✉ 

IMDEA Networks Institute, Leganés, Madrid, Spain

César Sánchez ✉ 

IMDEA Software Institute, Pozuelo de Alarcón, Madrid, Spain

Abstract

Blockchain technologies are facing a scalability challenge, which must be overcome to guarantee a wider adoption of the technology. This scalability issue is due to the use of consensus algorithms to guarantee the total order of the chain of blocks (and of the transactions within each block). However, total order is often overkill, since important advanced applications of smart-contracts do not require a total order among *all* operations. A much higher scalability can potentially be achieved if a more relaxed order (instead of a total order) can be exploited.

In this paper, we propose a distributed concurrent data type, called *Setchain*, which improves scalability significantly. A Setchain implements a *grow-only set object* whose elements are not 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. Therefore, two operations in the same epoch are not ordered, while two operations in different epochs are ordered by their respective epoch number. We present different Byzantine-tolerant implementations of Setchain, prove their correctness and report on an empirical evaluation of a prototype implementation.

Our results show that Setchain is orders of magnitude faster than consensus-based ledgers, since it implements grow-only sets with epoch synchronization instead of total order. Moreover, since the Setchain barriers can be synchronized with the underlying blockchain, Setchain objects can be used as a *sidechain* to implement many smart contract solutions with much faster operations than on basic blockchains.

2012 ACM Subject Classification Security and privacy → Distributed systems security

Keywords and phrases Distributed systems, blockchain, byzantine distributed objects, consensus, Setchain.

1 Introduction

1.1 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

43 functionality, enabling many applications in decentralized finances (DeFi)¹, decentralized
44 governance, Web3, etc.

45 The main element of all distributed ledgers is the “blockchain”, which is a distributed
46 object that contains, packed in blocks, the ordered list of transactions performed on behalf of
47 the users [14, 13]. This object is maintained by multiple servers without a central authority
48 by using consensus algorithms that are resilient to Byzantine attacks.

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

57 Consequently, there is a growing interest in techniques to improve the scalability of
58 blockchains [20, 35]. Approaches include (1) the search for faster consensus algorithms [32], (2)
59 the use of parallel techniques, like sharding [10], (3) building application-specific blockchains
60 with Inter-Blockchain Communication capabilities [34], [19], or (4) extracting functionality
61 out of the blockchain, while trying to preserve the guarantees of the blockchain: the “layer
62 2” (L2) approach [17]. L2 approaches include the computation off-chain of Zero-Knowledge
63 proofs [2], which only need to be checked on-chain (hopefully more efficiently) [1], the adoption
64 of limited (but useful) functionality like *channels* (e.g., Lightning [22]), or the deployment
65 of optimistic rollups (e.g., Arbitrum [18]) based on avoiding running the contracts in the
66 servers (except when needed to annotate claims and resolve disputes).

67 In this paper, we propose an alternative approach to increase blockchain scalability that
68 exploits the following observation. It has been traditionally assumed that cryptocurrencies
69 require total order to guarantee the absence of double-spending. However, many useful
70 applications and functionalities (including cryptocurrencies [16]) can tolerate more relaxed
71 guarantees, where operations are only *partially ordered*. We propose here a Byzantine-fault
72 tolerant implementation of a distributed grow-only set [27, 5], equipped with an additional
73 operation for introducing points of barrier synchronization (where all servers agree on the
74 contents of the set). Between barriers, elements of the distributed set can temporarily be
75 known by some but not all servers. We call this distributed data structure a Setchain. A
76 blockchain \mathcal{B} implementing Setchain (as well as blocks) can align the consolidation of the
77 blocks of \mathcal{B} with barrier synchronizations, obtaining a very efficient set object as side data
78 type, with the same Byzantine-tolerance guarantees that \mathcal{B} itself offers.

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

81 1.1.0.1 A Completely off-chain implementation

82 The major drawback is that from the point of view of the underlying blockchain the resulting
83 implementation does not have the trustability and accountability guarantees that blockchains
84 offer. One example of this approach are *mempools*. Mempools (short for memory pools)
85 are a P2P data type used by most blockchains to maintain a set of pending transactions.

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/>.

86 Mempools fulfill two objectives: (1) to prevent distributed attacks to the servers that mine
 87 blocks and (2) to serve as a pool of transaction requests from where block producers select
 88 operations. Nowadays, mempools are receiving a lot of attention, since they suffer from
 89 lack of accountability and are a source of attacks [26, 25], including front-running [9, 24, 30].
 90 Our proposed data structure, Setchain, offers a much stronger accountability, because it is
 91 resilient to Byzantine attacks and the contents of the set that Setchain maintains is public
 92 and cannot be forged.

93 1.1.0.2 Completely on-chain solution

94 Consider the following implementation (in a language similar to Solidity), where `add` is used
 95 to add elements, and `epochinc` to increase epochs.

```

96 contract Epoch {
97     uint public epoch = 0;
98     set public the_set = emptyset;
99     mapping(uint => set) public history;
100     function add(elem data) public {
101         the_set.add(data);
102     }
103     function epochinc() public {
104         history[++epoch] = the_set.setminus(history);
105     }
106 }
107 }
108
```

109 Since `epoch`, `the_set`, and `history` are defined `public`, there is an implicit getter function
 110 for each of them². One problem of this implementation is that every time we add an
 111 element, `the_set` gets bigger, which can affect the required cost to execute the contract. A
 112 second more important problem is that adding elements is *slow*—as slow as interacting with
 113 the blockchain—while our main goal is to provide a much faster data structure than the
 114 blockchain.

115 Our approach is faster, and can be deployed independently of the underlying blockchain
 116 or synchronized with the blockchain nodes. Thus, it lies between of these two extremes.

117 For any given blockchain \mathcal{B} , we propose an implementation of Setchain that (1) is much
 118 more efficient than implementing and executing operations directly in \mathcal{B} ; (2) offers the same
 119 decentralized guarantees against Byzantine attacks than \mathcal{B} , and (3) can be synchronized with
 120 the evolution of \mathcal{B} , so contracts could potentially inspect the contents of the Setchain. In a
 121 nutshell, these goals are achieved by using faster operations for the coordination among the
 122 servers (namely, reliable broadcast) for non-synchronized element insertions, and use only a
 123 consensus like algorithm for epoch changes.

124 1.2 Motivation

125 The potential applications that motivate the development of Setchain include:

126 1.2.1 Mempool

127 User transaction requests are nowadays stored in a mempool before they are chosen by
 128 miners, and once mined the information is lost. Recording and studying the evolution of
 129 mempools would require an additional object serving as a mempool *log system*, which must

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.

130 be fast enough to record every attempt of interaction with the mempool without affecting
 131 the underlying blockchain’s performance. Setchain as a sidechain can be used to implement
 132 one such trustable log system.

133 1.2.2 Scalability by L2 Optimistic Rollups

134 Optimistic rollups, like Arbitrum [18], use the fact that computation can be done outside the
 135 blockchain, posting on-chain only claims about its evolution. In this optimistic strategy users
 136 can propose the next state of the “contract.” After some time, the arbitrator smart contract
 137 on-chain assumes that a given proposed step is correct, and executes the annotated effects.
 138 A conflict resolution algorithm, also part of the contract on-chain, is used to resolve disputes.
 139 This protocol does not require a strict total order, but only a record of the actions proposed.
 140 Moreover, conflict resolutions can be reduced to claim validations, which could be performed
 141 by the maintainers of the Setchain.

142 1.2.3 Sidechain Data

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

148 1.3 Contributions.

149 In summary, the contributions of the paper are the following:

- 150 ■ the design and implementation of a side-chain data structure called *distributed Setchain*,
- 151 ■ several implementations of Setchain, providing different levels of abstraction and al-
 152 gorithmic implementation improvements,
- 153 ■ an empirical evaluation of a prototype implementation, which suggests that Setchain is
 154 several orders of magnitude faster than consensus.

155 2 Preliminaries

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

158 2.1 Model of Computation

159 A distributed system consists of processes—clients and servers—with an underlying com-
 160 munication graph in which each process can communicate with every other process. The
 161 communication is performed using message passing. Each process computes independently
 162 and at its own speed, and the internals of each process remain unknown to other processes.
 163 Message transfer delays are arbitrary but finite and also remain always unknown to processes.
 164 The intention is that servers will communicate among themselves to implement a distributed
 165 data type with certain guarantees, and clients can communicate with servers to exercise the
 166 data type.

167 Processes can fail arbitrarily, but the number of failing servers is bounded by f , and
 168 the total number of servers, n , is at least $3f + 1$. We assume *reliable channels* between
 169 non-Byzantine (correct) processes, so no message is lost, duplicated or modified. Each process

170 (client or server) has a pair of public and private keys. The public keys have been distributed
 171 reliably to all the processes that may interact with each other. Therefore, we discard the
 172 possibility of spurious or fake processes. We assume that messages are authenticated, so
 173 that messages corrupted or fabricated by Byzantine processes are detected and discarded
 174 by correct processes [7]. As result, communication between correct processes is reliable
 175 but asynchronous by default. However, for the set consensus service we use as a basic
 176 building block, partial synchrony is required [6, 15], as presented below. Observe that this
 177 requirement is only for the messages and computation of the protocol implementing this
 178 service. Finally, we assume that there is a mechanism for clients to create “valid objects” that
 179 servers can check locally. In the context of blockchains this is implemented using public-key
 180 cryptography.

181 2.2 Building Blocks

182 We will use four building blocks to implement Setchain:

183 2.2.1 Byzantine Reliable Broadcast (BRB)

184 The BRB service [3, 23], allows to broadcast messages to a set of processes guaranteeing
 185 that messages sent by correct processes are eventually received by *all* correct processes
 186 and that all correct processes eventually receive *the same* set of messages. The service
 187 provides a primitive BRB.Broadcast(m) for sending messages and an event BRB.Deliver(m)
 188 for receiving messages. Some important properties of BRB are:

- 189 ■ **BRB-Validity:** If a correct process p_i executes BRB.Deliver(m) and m was sent by a
 190 correct process p_j , then p_j executed BRB.Broadcast(m) in the past.
- 191 ■ **BRB-Termination:** If a correct process p executes BRB.Broadcast(m), then all correct
 192 processes (including p) eventually execute BRB.Deliver(m).

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

195 2.2.2 Byzantine Atomic Broadcast (BAB)

196 The BAB service [11] extends BRB with an additional guarantee: a total order of delivery of
 197 the messages. BAB provides the same operation and event as BRB, which we will rename
 198 as BAB.Broadcast(m) and BAB.Deliver(m). In addition to validity and termination, BAB
 199 services also provide:

- 200 ■ **Total Order:** If two correct processes p and q both BAB.Deliver(m) and BAB.Deliver(m'),
 201 then p delivers m before m' , if and only if q delivers m before m' .

202 BAB has been proven to be as hard as consensus [11], and thus, is subject to the same
 203 limitations [15].

204 2.2.3 Byzantine Distributed Grow-only Sets (DSO) [5]

205 Sets are one of the most basic and fundamental data structures in computer science, which
 206 typically include operations for adding and removing elements. Adding and removing opera-
 207 tions do not commute, and thus, distributed implementations require additional mechanisms
 208 to keep replicas synchronized to prevent conflicting local states. One solution is to allow
 209 only additions. Hence, a grow-only set is a set in which elements can only be added but not
 210 removed (implementable as a conflict-free replicated data structure [27]).

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

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

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

- 218 ■ **Byzantine Completeness**: All $get()$ and $add()$ operations invoked by correct processes
 219 eventually complete.
- 220 ■ **DSO-AddGet**: All $add(r)$ operations will eventually result in r being in the set returned
 221 by *all* $get()$.
- 222 ■ **DSO-GetAdd**: Each element r returned by $get()$ was added using $add(r)$ in the past.

223 2.2.4 Set Byzantine Consensus (SBC)

224 SBC, introduced in RedBelly [6], is a Byzantine-tolerant distributed problem, similar to
 225 consensus. In SBC, each participant proposes a set of elements (in the particular case of
 226 RedBelly, a set of transactions). After SBC finishes, all correct servers agree on a set of valid
 227 elements which is guaranteed to be a subset of the union of the proposed sets. Intuitively,
 228 SBC efficiently runs binary consensus to agree on the sets proposed by each participant, such
 229 that if the outcome is positive then the set proposed is included in the final set consensus.
 230 Some properties of SBC are:

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

237 The RedBelly algorithm [6] solves SBC in a system with partial synchrony: there is an
 238 unknown global stabilization time after which communication is synchronous. (Other SBC
 239 algorithms may have different partial synchrony assumptions.) Then, [6] proposes to use SBC
 240 to replace consensus algorithms in blockchains, seeking to improve scalability, because all
 241 transactions to be included in the next block can be decided with one execution of the SBC
 242 algorithm. Every server computes the same block by applying a deterministic function that
 243 totally orders the decided set of transactions, removing invalid or conflicting transactions.

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

- 247 ■ **SBC-Censorship-Resistance**: there is a time τ after which, if the proposed sets of all
 248 correct processes contain the same element e , then e will be in the decided set.

249 In RedBelly, this property holds because after the global stabilization time, all set consensus
 250 rounds decide sets from correct processes [6, Theorem 3].

251 3 The Setchain Distributed Data Structure

252 A key concept of Setchain is the *epoch* number, which is a global counter that the distributed
 253 data structure maintains. The synchronization barrier is realized as an epoch change: the
 254 epoch number is increased and the elements in the grow-only set that have not been assigned
 255 a previous epoch are stamped with the new epoch number.

3.1 API and Server State of the Setchain

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

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

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

- $v.add(e)$: requests to add e to **the_set**.
- $v.get()$: returns the values of **the_set**, **history**, and **epoch**, as seen by v .
- $v.epoch_inc(h)$ triggers an epoch change (i.e., a synchronization barrier). It must hold that $h = epoch + 1$.

Informally, a client process p invokes a $v.get()$ operation in node v to obtain (S, H, h) , which is v 's view of set $v.the_set$ and map $v.history$, with domain $[1..h]$. Process p invokes $v.add(e)$ to insert a new element e in $v.the_set$, and $v.epoch_inc(h+1)$ to request an epoch increment. At server v , the set $v.the_set$ contains the knowledge of v about elements that have been added, including those that have not been assigned an epoch yet, while $v.history$ contains only those elements that have been assigned an epoch. A typical scenario is that an element $e \in U$ is first perceived by v to be in **the_set**, to eventually be stamped and copied to **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 $epoch_inc()$ initiates the process of collecting elements in **the_set** at each node and collaboratively decide which ones are stamped with the current epoch.

Initially, both **the_set** and **history** are empty and **epoch** = 0 in every correct server. Note that client processes can insert elements to **the_set** through $add()$, but only servers decide how to update **history**, which client processes can only influence by invoking $epoch_inc()$.

At a given point in time, the view of **the_set** may differ from server to server. The Setchain data structure we propose only provides eventual consistency guarantees, as defined next.

3.2 Desired Properties

We specify now properties of correct implementations of Setchain. We provide first a low-level specification that assumes that clients interact with a *correct* server. Even though clients cannot be sure of whether the server they contact is correct we will see how they can later check and confirm that the operations were successful. These low-level primitives are also used in Section 7 to build a protocol that allows correct clients to perform operations even when they interact with Byzantine servers, at the price of performance.

We start by requiring from a Setchain that every **add**, **get**, and **epoch_inc** operation issued on a correct server eventually terminates. We say that element e is in epoch i in history H (e.g., returned by a **get** invocation) if $e \in H(i)$. We say that element e is in H if there is an epoch i such that $e \in H(i)$. The first property states that epochs only contain elements coming from the grow-only set.

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

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

303 ► **Property 2 (Add-Get-Local).** *Let $v.add(e)$ be an operation invoked to a correct server v .
 304 Then, eventually all invocations $(S, H, h) = v.get()$ satisfy $e \in S$.*

305 The next property states that elements present in a correct server are propagated to all
 306 correct servers.

307 ► **Property 3 (Add-Get).** *Let v, w be two correct servers, let $e \in U$ and let $(S, H, h) = v.get()$.
 308 If $e \in S$, then eventually all invocations $(S', H', h') = w.get()$ satisfy that $e \in S'$.*

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

313 ► **Property 4 (Eventual-Get).** *Let v be a correct server, let $e \in U$ and let $(S, H, h) = v.get()$.
 314 If $e \in S$, then eventually all invocations $(S', H', h') = v.get()$ satisfy that $e \in H'$.*

315 The previous three properties imply the following property.

316 ► **Property 5 (Get-After-Add).** *Let $v.add(e)$ be an operation invoked on a correct server v
 317 with $e \in U$. Then, eventually all invocations $(S, H, h) = w.get()$ satisfy that $e \in H$, for all
 318 correct servers w .*

319 An element can be in at most one epoch, and no element can be in two different epochs even
 320 if the history sets are obtained from `get` invocations to two different (correct) servers.

321 ► **Property 6 (Unique Epoch).** *Let v be a correct server, $(S, H, h) = v.get()$, and let $i, i' \leq h$
 322 with $i \neq i'$. Then, $H(i) \cap H(i') = \emptyset$.*

323 All correct server processes agree on the epoch contents.

324 ► **Property 7 (Consistent Gets).** *Let v, w be correct servers, let $(S, H, h) = v.get()$ and
 325 $(S', H', h') = w.get()$, and let $i \leq \min(h, h')$. Then $H(i) = H'(i)$.*

326 Property 7 states that the histories returned by two `get` invocations to correct servers are one
 327 the prefix of the other. However, since two elements e and e' can be inserted at two different
 328 correct servers—which can take time to propagate—the `the_set` part of `get` obtained from
 329 two correct servers may not be contained in one another.

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

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

334 Properties 1, 6, 7 and 8 are safety properties. Properties 2, 3, 4 and 5 are liveness
 335 properties.

336 4 Implementations

337 In this section, we describe implementations of Setchain that satisfy the properties in Sec-
 338 tion 3. We first describe a centralized sequential implementation, and then three distributed

■ **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  $h \equiv \text{epoch} + 1$ 
10:  proposal  $\leftarrow$  the_set  $\setminus \bigcup_{k=1}^{\text{epoch}}$  history( $k$ )
11:  history  $\leftarrow$  history  $\cup$  { $\langle h, \text{proposal} \rangle$ }
12:  epoch  $\leftarrow$  epoch + 1

```

339 implementations. The first distributed implementation is built using a Byzantine distrib-
340 uted grow-only set object (DSO) to maintain `the_set`, and Byzantine atomic broadcast
341 (BAB) for epoch increments. The second distributed implementation is also built using
342 DSO, but it uses Byzantine reliable broadcast (BRB) to announce epoch increments and set
343 Byzantine consensus (SBC) for epoch changes. Finally, the third one uses local sets, BRB
344 for broadcasting elements and epoch increment announcements, and SBC for epoch changes.

345 4.1 Sequential Implementation

346 Alg. 0 shows a centralized solution, which maintains two local sets, `the_set`—to record
347 added elements—, and `history`, which is implemented as a collection of pairs $\langle h, A \rangle$ where
348 h is an epoch number and A is a set of elements. We use `history`(h) to refer to the set A
349 in the pair $\langle h, A \rangle \in \text{history}$. A natural number `epoch` is incremented each time there is a
350 new epoch. The operations are: `Add`(e), which checks that element e is valid and adds it to
351 `the_set`, and `Get`(), which returns (`the_set`, `history`, `epoch`).

352 4.2 Distributed Implementations

353 4.2.1 First approach. DSO and BAB

354 Alg. 1 uses two external services: DSO and BAB. We denote messages with the name of
355 the message followed by its content as in “`epinc`($h, \text{proposal}, i$)”. The variable `the_set` is
356 not a local set anymore, but a DSO initialized empty with `Init`() in line 2. The function
357 `Get`() invokes the DSO `Get`() function (line 4) to fetch the set of elements. The function
358 `EpochInc`(h) triggers the mechanism required to increment an epoch and reach a consensus
359 on the elements. This process begins by computing a local `proposal` set, of those elements
360 added but not stamped (line 14). The `proposal` set is then broadcasted using a BAB service
361 alongside the epoch number h and the server node id i (line 15). Then, the server waits to
362 receive exactly $2f + 1$ proposals, and keeps the set of elements E present in at least $f + 1$
363 proposals, which guarantees that each element $e \in E$ was proposed by at least one correct
364 server. The use of BAB guarantees that every message sent by a correct server eventually
365 reaches every other correct server in the *same order*, so all correct servers use the same set
366 of $2f + 1$ proposals. Therefore, all correct servers arrive to the same conclusion, and the set
367 E is added as epoch h in `history` in line 20.

■ **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:      $E \leftarrow \{e : e \in \text{proposal for at least } f + 1 \text{ different } j\}$ 
20:     history  $\leftarrow$  history  $\cup \{h, E\}$ 
21:     epoch  $\leftarrow$  epoch + 1
22: end upon

```

368 Alg. 1, while easy to understand and prove correct, is not efficient. To start, in order to
369 complete an epoch increment, it requires at least $3f + 1$ calls to EpochInc(h) to different
370 servers, so at least $2f + 1$ proposals are received (the f Byzantine servers may not propose
371 anything). Another source of inefficiency comes from the use of off-the-shelf building blocks.
372 For instance, every time a DSO Get() is invoked, many messages are exchanged to compute
373 a reliable local view of the set [5]. Similarly, every epoch change requires a DSO Get()
374 in line 14 to create a proposal. Additionally, line 17 requires waiting for $2f + 1$ atomic
375 broadcast deliveries to take place. The most natural implementations of BAB services solve
376 one consensus per message delivered (see Fig. 7 in [4]), which would make this algorithm
377 very slow. We solve these problems in two alternative algorithms.

378 4.2.2 Second approach. Avoiding BAB

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

383 Ideally, when an EpochInc(h) is triggered unstamped elements in the local **the_set** of
384 each correct server should be stamped with the new epoch number and added to the set
385 **history**. However, we need to guarantee that for every epoch the set **history** is the same in
386 every correct server. Alg 1 enforced this using BAB and counting sufficient received messages.
387 Alg. 2 uses SBC to solve several independent consensus instances simultaneously, one on
388 each participant's proposal. Line 14 broadcasts an invitation to an epoch change, which
389 causes correct servers to build a proposed set and propose it the SBC. There is one instance
390 of SBC per epoch change, identified by h . With SBC each correct server receives the same
391 set of proposals (where each proposal is a set of elements). Then, every node applies the
392 same function to the same set of proposals reaching the same conclusion on how to update
393 **history**(h). The function preserves elements that are present in at least $f + 1$ proposed sets,

■ **Algorithm 2** Server i implementation using DSO, and reliably broadcast (BRB) and set Byzantine consensus (SBC).

```

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

```

394 which are guaranteed to have been proposed by some correct server. Observe that Alg. 2
395 still triggers one invocation of the DSO Get at each server to build the local proposal.

396 4.2.3 Final approach. BRB and SBC without DSOs

397 Alg. 3, avoids the cascade of messages that DSO Get calls require by dissecting the internals
398 of the DSO, and incorporating the internal steps in the Setchain algorithm directly. This
399 idea exploits the fact that a *correct Setchain server* is a *correct client* of the DSO, and there
400 is no need for the DSO to be defensive (this illustrates that using Byzantine resilient building
401 blocks does not compose efficiently, but exploring this general idea is out of the scope of this
402 paper).

403 Alg. 3 implements `the_set` using a local set (line 2). Elements received in `Add(e)` are
404 propagated using BRB. At any given point in time two correct servers may have a different
405 local sets (due to pending BRB deliveries) but each element added in one server will eventually
406 be known to all others. The local variable `history` is only updated in line 25 as a result of a
407 SBC round. Therefore, all correct servers will agree on the same sets formed by unstamped
408 elements proposed by some server. Additionally, Alg. 3 updates `the_set` to account for
409 elements that are new to the server (line 26), guaranteeing that all elements in `history` are
410 also in `the_set`. Note that this opens the opportunity to add elements directly by proposing
411 them during an epoch change without broadcasting them before. This optimization is
412 exploited in Section 6 to speed up the algorithm further. As a final note, Alg. 3 allows
413 a Byzantine server to bypass `Add` to propose elements, which will be accepted as long as
414 the elements are valid. This is equivalent to a client proposing an element using an `Add`
415 operation, which is then successfully propagated in an epoch change.

416 5 Proof of Correctness

417 We prove now the correctness of Alg. 3. We first show that all stamped elements are in
418 `the_set`, which implies Prop. 1 (*Consistent Sets*).

■ **Algorithm 3** Server implementation using a local set, Byzantine reliable broadcast (BRB) and set Byzantine consensus (SBC).

```

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

```

419 ► **Lemma 1.** For every correct server v , at the end of each function/upon, $\bigcup_h v.\text{history}(h) \subseteq$
420 $v.\text{the_set}$.

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

424 ► **Lemma 2.** Let v be a correct server and e an element in $v.\text{the_set}$. Then e will eventually
425 be in $w.\text{the_set}$ for every correct server w .

426 **Proof.** Initially, $v.\text{the_set}$ is empty. There are two ways to add an element e to $v.\text{the_set}$:
427 (1) At line 10, so e is valid and was received via a BRB.Deliver(add(e)). By Properties
428 **BRB-Validity** and **BRB-Termination** of BRB (see Section 2), every correct server w
429 will eventually execute BRB.Deliver(add(e)), and then (since e is valid), w will add it to
430 $w.\text{the_set}$ in line 10. (2) At line 26, so element e is valid and was received as an element in one
431 of the sets in propset from SBC[h].SetDeliver(propset) with $h = v.\text{epoch} + 1$. By properties
432 **SBC-Termination** **SBC-Agreement** and **SBC-Validity** of SBC (see Section 2), all
433 correct servers agree on the same set of proposals. Therefore, if v adds e then w either adds
434 it or has it already in its $w.\text{history}$ which implies by Lemma 1 that $e \in w.\text{the_set}$. In
435 either case, e will eventually be in $w.\text{the_set}$. ◀

436 Lemma 2, and the code of `Add()` and line 4 of `Get()` imply Prop. 2 (*Add-Get-Local*) and
 437 Prop. 3 (*Add-Get*). The following lemmas reason about how elements are stamped.

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

440 **Proof.** It follows directly from the check that e is not injected at $v.history(h)$ if $e \in$
 441 $v.history$ in line 25. ◀

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

444 **Proof.** The proof proceeds by induction on `epoch`. The base case is `epoch = 0`, which holds
 445 trivially since $v.history(0) = w.history(0) = \emptyset$. Variable `epoch` is only incremented in
 446 one unit in line 27, after `history(h)` has been changed in line 25 when $h = epoch + 1$. In
 447 that line, v and w are in the same phase on SBC (for the same h). By **SBC-Agreement**,
 448 v and w receive the same *propset*, both v and w validate all elements equally, and (by
 449 inductive hypothesis), for each $h' \leq epoch$ it holds that $e \in v.history(h')$ if and only if
 450 $e \in w.history(h')$. Therefore, in line 25 both v and w update `history(h)` equally, and after
 451 line 27 it holds that $v.history(epoch) = w.history(epoch)$. ◀

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

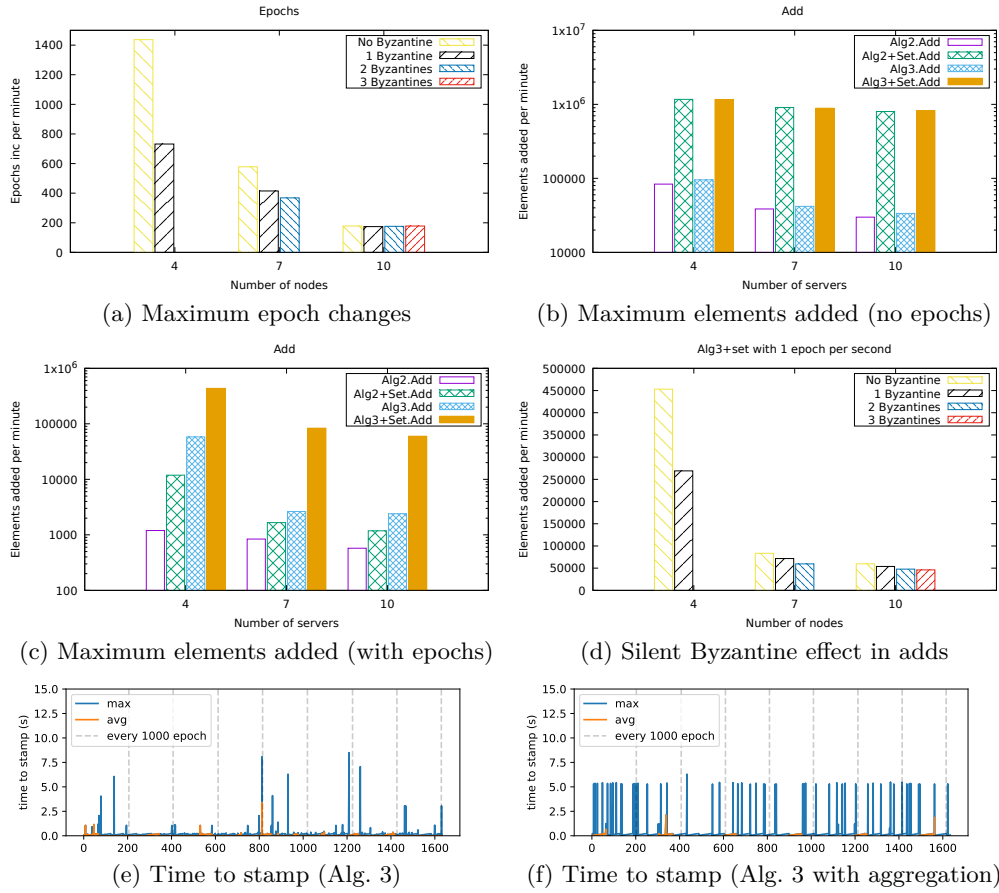
454 **Proof.** By Lemma 2 every correct server w will satisfy $e \in w.the_set$ at some $t > \tau$. By
 455 assumption, there is a new `EpochInc()` after t (let the epoch number be h). If e is already
 456 in `history(h')` for $h' < h$ we are done, since from Lemma 4 in this case at the end of the
 457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

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

463 Finally, we discuss Prop. 8 (*Add-before-Get*). If valid elements can only be created by
 464 clients and added using `Add(e)` the property trivially holds. If valid elements can be created
 465 by, for example Byzantine servers, then they can inject elements in `the_set` and `history`
 466 of correct servers without using `Add()`. They can either execute directly a `BRB.Broadcast`
 467 or directly via the SBC in epoch rounds. In these case, Alg. 3 satisfies a weaker version of
 468 (*Add-before-Get*) that states that elements returned by `Get()` are either added by `Add()`, by
 469 a `BRB.Broadcast` or injected in the SBC phase.

470 6 Empirical Evaluation

471 We have implemented the server code for DSO, BRB and SBC and using these building
 472 blocks we have implemented Alg. 2 and Alg. 3. Our prototype is written in Golang [12]
 473 1.16 with message passing using ZeroMQ [29] over TCP. Our testing platform uses Docker
 474 running on a server with 2 Intel Xeon CPU processors at 3GHz with 36 cores and 256GB
 475 RAM, running Ubuntu 18.04 Linux64. Each Setchain server node was wrapped in a Docker
 476 container with no limit on CPU or RAM usage. Alg. 2 implements a Setchain and a DSO as
 477 two standalone executables that communicate using remote procedure calls on the internal



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

478 loopback network interface of the Docker container. The RPC server and client are taken
 479 from the Golang standard library. For Alg. 3 everything resides in a single executable.
 480 For both algorithms, we evaluate two versions, one where each element inserted causes a
 481 broadcast and another where servers aggregate locally inserted elements until a maximum
 482 message size (of 10^6 elements) or a maximum element timeout (of 5s) is reached. In all cases
 483 elements have 116-126 bytes.

484 We evaluate empirically the following hypothesis:

- 485 ■ (H1): The maximum rate of elements that can be inserted is much higher than the
 486 maximum epoch rate.
- 487 ■ (H2): Alg. 3 performs better than Alg. 2.
- 488 ■ (H3): The aggregated versions perform better than the basic versions.
- 489 ■ (H4) Silent Byzantine servers do not affect dramatically the performance.
- 490 ■ (H5) The performance does not degrade over time.

491 To evaluate these hypotheses, we carried out the experiments described below and reported
 492 in Fig. 1. In all cases, operations are injected by clients running within the same Docker
 493 container. Resident memory was always enough such that in no experiment the operating
 494 system needed to recur to disk swapping. All the experiments consider deployments with 4,
 495 7, or 10 server nodes, and each running experiment reported is taken from the average of 10

496 executions.

497 We tested first how many epochs per minute our Setchain implementations can handle.
498 In these runs, we did not add any element and we incremented the epoch rate to find out
499 the smallest latency between an epoch and the subsequent one. We run it with 4, 7, and 10
500 nodes, with and without Byzantine servers. This is reported in Fig. 1(a).

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

508 In our third experiment, we compare the performance of our implementations combining
509 epoch increments and insertion of elements. We set the epoch rate at 1 epoch change per
510 second and calculated the maximum add ratio. The outcome is reported in Fig. 1(c), which
511 shows that Alg. 3 outperforms Alg. 2. In fact, Alg. 3+set even outperforms Alg. 2+set by a
512 factor of roughly 5 for 4 nodes and by a factor of roughly 2 for 7 and 10 nodes. Alg. 3+set can
513 handle 8x the elements added by Alg. 3 for 4 nodes and 30x for 7 and 10 nodes. The benefits
514 of Alg. 3+set over Alg. 3 increase as the number of nodes increase because Alg. 3+set avoids
515 the broadcasting of elements which generates a number of messages that is quadratic in the
516 number of nodes in the network. This experiment confirms (H2) and (H3). The difference
517 between Alg. 3 and Alg. 2 was not observable in the previous experiment (without epoch
518 changes) because the main difference is in how servers proceed to collect elements to vote
519 during epoch changes.

520 The next experiment explores how silent Byzantine servers affect Alg. 3+set. We
521 implement silent Byzantine servers and run for 4, 7 and 10 nodes with an epoch change ratio
522 of 1 per second, calculating the maximum add rate. This is reported in Fig. 1(d). Silent
523 Byzantine servers degrade the speed for 4 nodes as in this case the implementation considers
524 the silent server very frequently in the validation phase, but it can be observed that this
525 effect is much smaller for larger number of servers, validating (H4).

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

536 Considering that epoch changes is essentially a set consensus, our experiments suggest that
537 inserting elements in a Setchain is three orders of magnitude faster than performing consensus.
538 However, a full validation of this hypothesis would require to fully implement Setchain on
539 performant gossip protocols and compare with comparable consensus implementations.

■ **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() in at least  $3f + 1$  different servers.
5:   wait  $2f + 1$  resp  $s$ .(the_set, history, epoch)
6:    $S \leftarrow \{e \mid e \in s.\mathbf{the\_set} \text{ in at least } f + 1 \text{ servers } s\}$ 
7:    $H \leftarrow \emptyset$ 
8:    $i \leftarrow 1$ 
9:    $N \leftarrow \{s : s.\mathbf{epoch} \geq i\}$ 
10:  while  $\exists E : |\{s \in N : s.\mathbf{history}(i) = E\}| \geq f + 1$  do
11:     $H \leftarrow H \cup \{i, E\}$ 
12:     $N \leftarrow N \setminus \{s : s.\mathbf{history}(i) \neq E\}$ 
13:     $N \leftarrow N \setminus \{s : s.\mathbf{epoch} = i\}$ 
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.

```

540 7 Distributed Partial Order Objects (DPO)

541 The algorithms presented in Section 4 and the proofs in Section 5 consider the case of clients
542 contacting a correct server. Obviously, client processes do not know if they are contacting a
543 Byzantine or correct process, so a client protocol is required to encapsulate the details of the
544 distributed system. We describe now such a client protocol inspired by the one for DSO [5],
545 which involves the exchange of several more messages than contacting a single server with a
546 request. We later describe a more efficient “*try and check*” alternative.

547 The general idea of the client protocol is to interact with enough servers to guarantee
548 that some are correct and ensure the desired behavior. The Setchain API has methods that
549 wait for a result (Get) and methods that do not require a response (EpochInc and Add).
550 Alg. 4 shows the client protocol. To guarantee contacting at least one correct server, we need
551 to send $f + 1$ messages. Note that each message may trigger different broadcasts.

552 The wrapper algorithm for function Get can be split in two parts. First, the protocol
553 contacts $3f + 1$ nodes, and waits for at least $2f + 1$ responses (f Byzantine servers may refuse
554 to respond). The response from server s is $(s.\mathbf{the_set}, s.\mathbf{history}, s.\mathbf{epoch})$. The protocol
555 then computes S as those elements known to be in **the_set** by at least $f + 1$ servers (which
556 includes at least one correct server). To compute H , the code goes incrementally epoch by
557 epoch as long as at least $f + 1$ servers within the set N (which is initialized with all the
558 servers that responded with non-empty histories) agree on a set E of elements in epoch i .
559 If $f + 1$ servers agree that E is the set of elements in epoch i , this is indeed the case. We
560 also remove from N those servers that either do not know more epochs or that incorrectly
561 reported something different than E . Once this process ends, the sets S and H , and the
562 latest processed epoch are returned. It is guaranteed that $\mathbf{history} \subseteq \mathbf{the_set}$.

563 We also present an alternative faster optimistic client. In this approach correct servers
564 sign cryptographically a hash of the set of elements in an epoch, and insert this hash in
565 the Setchain as an element. Clients only perform a **single** Add(e) request to one server,
566 hoping it will be a correct server. After waiting for some time, the client invokes a Get from
567 a **single** server (which again can be Byzantine) and check whether e is in some epoch signed
568 by (at least) $f + 1$ servers, in which case the epoch is correct and e has been successfully

569 inserted and stamped. Note that this requires only one message per Add and one message
570 per Get.

571 **8** Concluding Remarks

572 We presented a novel distributed data-type, called Setchain, that implements a grow-only set
573 with epochs, and tolerates Byzantine server nodes. We provided a low-level specification of
574 desirable properties of Setchains and presented three distributed implementations, where the
575 most efficient one uses Byzantine Reliable Broadcast and RedBelly set Byzantine consensus.
576 Our preliminary empirical evaluation suggests that the performance of inserting elements in
577 Setchain is three orders of magnitude faster than with consensus.

578 Future work includes developing the motivating applications listed in the introduction,
579 for example, mempool logs using Setchains, and L2 faster optimistic rollups. We will also
580 study how to equip blockchains with Setchain (synchronizing blocks and epochs) to allow
581 smart-contracts to access the Setchain. An important problem to solve is how clients of the
582 Setchain pay for the usage (even if a much smaller fee than for the blockchain itself).

583 Setchain may be used to implement a solution to front-running. Mempools encode
584 information about what it is about to happen in blockchains, so anyone observing them can
585 predict the next operations to be mined, and take actions to their benefit. *Front-running* is
586 the action of observed transaction request and maliciously inject transactions to be executed
587 before the observed ones [9, 30] (by paying a higher fee to a miner). Setchain can be used
588 to *detect* front-running since it can serve as a basic mechanism to build a mempool that is
589 efficient and serves as a log of requests. Additionally, Setchains can be used as a building
590 block to solve front-running where users encrypt their requests using a multi-signature
591 decryption scheme, where participant decrypting servers decrypt requests after they are
592 chosen for execution by miners once the order has already been fixed.

593 Our Setchain exploits a specific partial orders that relaxes the total order imposed by
594 blockchains. As future we will explore other partial orders and their uses, for example,
595 federations of Setchain, one Setchain per smart-contract, etc.

596 There are also interesting problems for foundational future work. Alg. 3 shows that
597 Byzantine tolerant building blocks do not compose efficiently, because each building is
598 pessimistic and does not exploit the fact that when building a correct sever, the client of
599 the Byzantine tolerant building block is correct. Also, our analysis shows that Byzantine
600 behavior of server nodes can be modeled by a collection of simple interactions with BRB and
601 SBC, so it is possible to model all Byzantine behavior to simplify reasoning.

602 ——— References ———

- 603 1 Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran
604 Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In
605 *Proc. of S&P'14*, pages 459–474, 2014. doi:10.1109/SP.2014.36.
- 606 2 Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive
607 Zero Knowledge for a von Neumann architecture. In *Proc. of USENIX Sec.'14*, pages 781–796.
608 USENIX, August 2014. URL: [https://www.usenix.org/conference/usenixsecurity14/
609 technical-sessions/presentation/ben-sasson](https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/ben-sasson).
- 610 3 Gabriel Bracha. Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75(2):130–143,
611 1987. doi:10.1016/0890-5401(87)90054-X.
- 612 4 Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed
613 systems. *J. ACM*, 43(2):225–267, mar 1996. doi:10.1145/226643.226647.

- 614 5 Vicent Cholvi, Antonio Fernández Anta, Chryssis Georgiou, Nicolas Nicolaou, Michel Raynal,
615 and Antonio Russo. Byzantine-tolerant distributed grow-only sets: Specification and applica-
616 tions. In *Proc. of FAB'21*, page 2:1–2:19, 2021.
- 617 6 Tyler Crain, Christopher Natoli, and Vincent Gramoli. Red belly: A secure, fair and scalable
618 open blockchain. In *Proc. of S&P'21*, pages 466–483, 2021. doi:10.1109/SP40001.2021.00087.
- 619 7 F. Cristian, H. Aghili, R. Strong, and D. Volev. Atomic broadcast: from simple message
620 diffusion to byzantine agreement. In *25th Int'l Symp. on Fault-Tolerant Computing*, pages
621 431–, 1995. doi:10.1109/FTCSH.1995.532668.
- 622 8 Kyle Croman, Christian Decker, Ittay Eyal, Adem Efe Gencer, Ari Juels, Ahmed Kosba, An-
623 drew Miller, Prateek Saxena, Elaine Shi, Emin Gün Sirer, Dawn Song, and Roger Wattenhofer.
624 On scaling decentralized blockchains. In *Financial Crypto. and Data Security*, pages 106–125.
625 Springer, 2016.
- 626 9 Philip Daian, Steven Goldfeder, T. Kell, Yunqi Li, X. Zhao, Iddo Bentov, Lorenz Breidenbach,
627 and A. Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable
628 value, and consensus instability. *Proc. of S&P'20*, pages 910–927, 2020.
- 629 10 Hung Dang, Tien Tuan Anh Dinh, Dumitrel Loghin, Ee-Chien Chang, Qian Lin, and Beng Chin
630 Ooi. Towards scaling blockchain systems via sharding. In *Proc. of SIGMOD'19*, pages 123–140.
631 ACM, 2019. doi:10.1145/3299869.3319889.
- 632 11 Xavier Défago, André Schiper, and Péter Urbán. Total order broadcast and multicast
633 algorithms: Taxonomy and survey. *ACM Comput. Surv.*, 36(4):372–421, dec 2004. doi:
634 10.1145/1041680.1041682.
- 635 12 Alan A.A. Donovan and Brian W. Kernighan. *The Go Programming Language*. Adison-Wesley,
636 2015.
- 637 13 Antonio Fernández Anta, Chryssis Georgiou, Maurice Herlihy, and Maria Potop-Butucaru.
638 *Principles of Blockchain Systems*. Morgan & Claypool Publishers, 2021.
- 639 14 Antonio Fernández Anta, Kishori Konwar, Chryssis Georgiou, and Nicolas Nicolaou. Formaliz-
640 ing and implementing distributed ledger objects. *ACM Sigact News*, 49(2):58–76, 2018.
- 641 15 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed
642 consensus with one faulty process. *JACM*, 32(2):374–382, 1985. doi:10.1145/3149.214121.
- 643 16 Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, and Dragos-Adrian Sered-
644 inschi. The consensus number of a cryptocurrency. In *Proc. of PODC'19*, pages 307–316.
645 ACM, 2019. doi:10.1145/3293611.3331589.
- 646 17 Maxim Jourenko, Kanta Kurazumi, Mario Larangeira, and Keisuke Tanaka. Sok: A taxonomy
647 for layer-2 scalability related protocols for cryptocurrencies. *IACR Cryptol. ePrint Arch.*,
648 2019:352, 2019.
- 649 18 Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S. Matthew Weinberg, and Edward W.
650 Felten. Arbitrum: Scalable, private smart contracts. In *27th USENIX Security Symposium*,
651 pages 1353–1370. USENIX Assoc., 2018. URL: [https://www.usenix.org/conference/
652 usenixsecurity18/presentation/kalodner](https://www.usenix.org/conference/usenixsecurity18/presentation/kalodner).
- 653 19 Jae Kwon and Ethan Buchman. Cosmos whitepaper, 2019.
- 654 20 Zamani Mahdi, Mahnush Movahedi, and Mariana Raykova. Rapidchain: Scaling blockchain
655 via full sharding. In *Proc. of CSS'18*, pages 931–948. ACM, 2018. doi:10.1145/3243734.
656 3243853.
- 657 21 Satoshi Nakamoto. Bitcoin: a peer-to-peer electronic cash system, 2009.
- 658 22 Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant
659 payments, 2016. URL: <https://lightning.network/lightning-network-paper.pdf>.
- 660 23 Michel Raynal. *Fault-Tolerant Message-Passing Distributed Systems: An Algorithmic Approach*.
661 01 2018. doi:10.1007/978-3-319-94141-7.
- 662 24 Robinson, Dan and Konstantopoulos, Georgios. Ethereum is a dark forest, 2020. URL:
663 <https://medium.com/@danrobinson/ethereum-is-a-dark-forest-ecc5f0505dff>.

- 664 25 Muhammad Saad, Laurent Njilla, Charles Kamhoua, Joongheon Kim, DaeHun Nyang, and Aziz
665 Mohaisen. Mempool optimization for defending against DDoS attacks in PoW-based blockchain
666 systems. In *Proc. of ICBC'19*, pages 285–292, 2019. doi:10.1109/BL0C.2019.8751476.
- 667 26 Muhammad Saad, My T. Thai, and Aziz Mohaisen. POSTER: Deterring DDoS attacks on
668 blockchain-based cryptocurrencies through mempool optimization. In *Proc. of ASIACCS'18*,
669 pages 809–811. ACM, 2018. doi:10.1145/3196494.3201584.
- 670 27 Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski. Convergent and Com-
671 mutative Replicated Data Types. *Bulletin- European Association for Theoretical Computer*
672 *Science*, (104):67–88, June 2011. URL: <https://hal.inria.fr/hal-00932833>.
- 673 28 Nick Szabo. Smart contracts: Building blocks for digital markets. *Entropy*, 16, 1996.
- 674 29 The ZeroMQ authors. Zeromq, 2021. <https://zeromq.org>. URL: <https://zeromq.org>.
- 675 30 Christof Ferreira Torres, Ramiro Camino, and Radu State. Frontrunner jones and the raiders
676 of the Dark Forest: An empirical study of frontrunning on the Ethereum blockchain. In *Proc*
677 *of USENIX Sec.'21*, pages 1343–1359, 2021. URL: [https://www.usenix.org/conference/](https://www.usenix.org/conference/usenixsecurity21/presentation/torres)
678 [usenixsecurity21/presentation/torres](https://www.usenix.org/conference/usenixsecurity21/presentation/torres).
- 679 31 Shobha Tyagi and Madhumita Kathuria. *Study on Blockchain Scalability Solutions*, page
680 394–401. ACM, 2021. URL: <https://doi.org/10.1145/3474124.3474184>.
- 681 32 Ke Wang and Hyong S. Kim. Fastchain: Scaling blockchain system with informed neighbor
682 selection. In *Proc. of IEEE Blockchain'19*, pages 376–383, 2019. doi:10.1109/Blockchain.
683 2019.00058.
- 684 33 Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum*
685 *project yellow paper*, 151:1–32, 2014.
- 686 34 Gavin Wood. Polkadot: Vision for a heterogeneous multi-chain framework. *White Paper*, 21,
687 2016.
- 688 35 Cheng Xu, Ce Zhang, Jianliang Xu, and Jian Pei. Slimchain: Scaling blockchain transactions
689 through off-chain storage and parallel processing. *Proc. VLDB Endow.*, 14(11):2314–2326, jul
690 2021. doi:10.14778/3476249.3476283.