



State Machine Replication Scalability Made Simple

Chrysoula Stathakopoulou
IBM Research Europe - Zurich
ETH - Zurich

Matej Pavlovic*
Protocol Labs

Marko Vukolić†
Protocol Labs

Abstract

Consensus, state machine replication (SMR) and total order broadcast (TOB) protocols are notorious for being poorly scalable with the number of participating nodes. Despite the recent race to reduce overall message complexity of leader-driven SMR/TOB protocols, scalability remains poor and the throughput is typically inversely proportional to the number of nodes. We present Insanely Scalable State Machine Replication, a generic construction to turn leader-driven protocols into scalable multi-leader ones. For our scalable SMR construction we use a novel primitive called Sequenced (Total Order) Broadcast (SB) which we wrap around PBFT, HotStuff and Raft leader-driven protocols to make them scale. Our construction is general enough to accommodate most leader-driven ordering protocols (BFT or CFT) and make them scale. Our implementation improves the peak throughput of PBFT, HotStuff, and Raft by 37x, 56x, and 55x, respectively, at a scale of 128 nodes.

CCS Concepts: • Computer systems organization → Dependable and fault-tolerant systems and networks.

Keywords: Byzantine fault tolerance, state machine replication, consensus, scalability

ACM Reference Format:

Chrysoula Stathakopoulou, Matej Pavlovic, and Marko Vukolić. 2022. State Machine Replication Scalability Made Simple. In *Seventeenth European Conference on Computer Systems (EuroSys '22)*, April 5–8, 2022, RENNES, France. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3492321.3519579>

1 Introduction

Considerable research effort has recently been dedicated to scaling state-machine replication (SMR) and total-order

broadcast (TOB) or consensus protocols, fundamental primitives in distributed computing. By scaling, we mean maintaining high throughput and low latency despite a growing number of nodes (replicas) n . Driven by the needs of blockchain systems, particular focus lies on deterministic Byzantine fault-tolerant (BFT) protocols in the eventually or partially synchronous model.

In this model, the classical Dolev/Reischuk (DR) [14] lower bound requires $\Omega(n^2)$ worst case message complexity, which was a focal complexity metric of many subsequent protocols including recent HotStuff [33]. However, we claim message complexity to be a rather poor scalability metric, demonstrated by the fact that HotStuff and other leader-driven protocols scale inversely proportionally to the number of nodes, despite some of them matching the DR lower bound. This is because in leader-driven protocols, the leader has at least $O(n)$ bits to send, even in the common case, yielding n^{-1} throughput scalability.

A recent effort to overcome the single leader bottleneck by allowing multiple parallel leaders (Mir-BFT [31]) in the classical PBFT protocol [10] demonstrates high scalability in practice. Despite certain advantages of PBFT, e.g., being highly parallelizable and designed not to require signatures on protocol messages, among the many (existing and future) TOB solutions, there are none that fit all use cases. HotStuff [33] is the first protocol with linear message complexity both in the common case and in leader replacement, making it suitable for highly asynchronous or faulty networks. On the other hand, other protocols, e.g., Aliph/Chain [19], have optimal throughput when failures are not expected to occur often. Finally, crash fault-tolerant (CFT) protocols such as Raft [28] and Paxos [24] tolerate a larger number of (benign) failures than BFT protocols for the same number of nodes.

Our work takes the Mir-BFT effort one step further, introducing Insanely Scalable SMR, hereinafter referred to as ISS, the first modular framework to make leader-driven TOB protocols scale. ISS is a classic SMR system that establishes a total order of client requests with typical liveness and safety properties, applicable to any replicated service, such as resilient databases or a blockchain ordering layer (e.g., as in Hyperledger Fabric [3]).

Notably, and unlike previous efforts [31][6], ISS achieves scalability without requiring a primary node to periodically decide on the protocol configuration. ISS achieves this by introducing a novel abstraction, Sequenced Broadcast (SB), which requires each instance of an ordering protocol to *terminate* after delivering a finite number of messages. This allows

*Work done while at IBM Research Europe - Zurich

†Work done while at IBM Research Europe - Zurich



This work is licensed under a Creative Commons Attribution International 4.0 License.

EuroSys '22, April 5–8, 2022, RENNES, France

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9162-7/22/04.

<https://doi.org/10.1145/3492321.3519579>

nodes in ISS to decide on the configuration independently and deterministically, without requiring additional communication and without relying on a single primary node. This in turn allows for more flexible and fair leader selection policies. Moreover, it guarantees better resilience against an adaptive adversary that can corrupt the primary node, which changes slowly and can be known in advance with a deterministic round robin rotation [31][6].

ISS implements SMR by multiplexing multiple instances of SB which operate concurrently on a *partition* of the domain of client requests. We carefully select the partition to maintain safety and liveness, as well as to prevent redundant data duplication, which has been shown to be detrimental to performance [31]. This is qualitatively better than related modular efforts [20, 27] which do not provide careful partitioning and load balancing and hence cannot achieve the same scalability and robustness at the same time.

ISS maintains a *contiguous* log of (batches of) client requests at each node. Each position in the log corresponds to a unique sequence number and ISS agrees on the assignment of a unique request batch to each sequence number. Our goal is to introduce as much parallelism as possible in assigning batches to sequence numbers while avoiding *request duplication*, i.e., assigning the same request to more than one sequence number. To this end, ISS subdivides the log into non-overlapping *segments*. Each segment, representing a subset of the log’s sequence numbers, corresponds to an independent instance of SB that has its own leader and executes concurrently with other SB instances. The SB abstraction, moreover, facilitates the reasoning about multiplexing the outputs of multiple instances into a single log, while staying very close to classic definitions of broadcast and thus being easily implementable by existing algorithms.

To prevent the leaders of two different segments from concurrently proposing the same request, and thus wasting resources, while also preventing malicious leaders from censoring (i.e., not proposing) certain requests, we adopt and generalize the rotating bucketization of the request space introduced by Mir-BFT [31]. ISS assigns a different *bucket*—subset of client requests—to each segment. No bucket is assigned to more than one segment at a time and each request maps (through a hash function) to exactly one bucket. ISS periodically changes the bucket assignment, such that each bucket is guaranteed to eventually be assigned to a segment with a correct leader.

To maintain the invariant of one bucket being assigned to one segment, all buckets need to be re-assigned at the same time. ISS therefore uses finite segments that it groups into *epochs*. An epoch is a union of multiple segments that forms a contiguous sub-sequence of the log. After all log positions within an epoch have been assigned request batches, and thus no requests are “in-flight”, ISS advances to the next epoch, meaning that it starts processing a new set of segments forming the next portion of the log.

We implement and deploy ISS on a wide area network (WAN) spanning 16 different locations spread around the world, demonstrating ISSs performance using two different BFT protocols (PBFT[10] and HotStuff[33]) and one CFT protocol (Raft[28]). On 128 nodes ISS improves the performance of the single leader counterpart protocols, PBFT, HotStuff, and Raft, by 37x, 56x, and 55x, respectively.

The rest of this paper is organized as follows. Section 2 presents the theoretical foundation of our work. It models the systems we study (Section 2.1), introduces the SB abstraction (Section 2.2) and describes how we multiplex SB instances with ISS (Sections 2.3 and 2.4). Sections 3 and 4 respectively describe the details of ISS and its implementation. In Section 5 we prove that multiplexing SB instances with ISS implements SMR and in Section 6 we evaluate the performance of ISS. In Section 7 we discuss related work. We conclude in Section 8.

2 Theoretical Foundations

2.1 System Model

We assume a set \mathcal{N} of *node* processes with $|\mathcal{N}| = n$. At most f of the nodes in \mathcal{N} can fail. We further assume a set \mathcal{C} of *client* processes of arbitrary size, any of which can be faulty. Each process is identified by its public key, provided by a public key infrastructure. Unless mentioned otherwise, we assume Byzantine, i.e., arbitrary, faults. Therefore, we require $n \geq 3f + 1$. We further assume that nodes in \mathcal{N} are computationally bounded and cannot subvert standard cryptographic primitives.

Processes communicate through authenticated point-to-point channels. We assume a partially synchronous network [15] such that the communication between any pair of correct processes is asynchronous before an unknown time GST , when the communication becomes synchronous.

Nodes in \mathcal{N} implement a state machine replication (SMR) service to clients in \mathcal{C} . To broadcast request r , a client c triggers an $\langle \text{SMR-CAST} | r \rangle$ event. A client request is a tuple $r = (o, id)$, where o is the request payload, e.g., some operation to be executed by some application, and id a unique request identifier. The request identifier is a tuple $id = (t, c)$ where t is a logical timestamp and c a client identifier, e.g., a client public key. Two client requests $r = (o, id), r' = (o', id')$ are considered equal, we write $r = r'$ and we refer to them as duplicates, if and only if $o = o' \wedge id = id'$.

Nodes assign a unique sequence number sn to r and eventually output an $\langle \text{SMR-DELIVER} | sn, r \rangle$ event such that the following properties hold:

- SMR1 Integrity:** If a correct node delivers (sn, r) , where $r.id.c$ is a correct client’s identity, then client c broadcast r .
- SMR2 Agreement:** If two correct nodes deliver, respectively, (sn, r) and (sn, r') , then $r = r'$.
- SMR3 Totality:** If a correct node delivers request (sn, r) , then every correct node eventually delivers (sn, r) .

SMR4 Liveness: If a correct client broadcasts request r , then some correct node eventually delivers (sn, r) .

2.2 Sequenced Broadcast (SB)

Sequenced Broadcast (SB) is a variant of Byzantine total order broadcast [8] with explicit sequence numbers and an explicit set of allowed messages.

SB is instantiated with a failure detector instance as a parameter. We assume an *eventually strong failure detector* in an environment with Byzantine faults denoted as $\diamond S(bz)$, as defined by Malkhi and Reiter [25].

The failure detector D of the class $\diamond S(bz)$ detects *quiet* nodes. Intuitively, a quiet node is the equivalent to a crashed node in the BFT model, accounting for non-crash faults that are indistinguishable from crashes. For the exact definition we refer the reader to Malkhi and Reiter's work [25].

A failure detector of the $\diamond S(bz)$ class guarantees the following two properties:

Strong Completeness: There is a time after which every quiet node is permanently suspected by every correct node.

Eventual Weak Accuracy: There is a time after which some correct node is never suspected by any correct node.

We can now define Sequenced Broadcast as follows.

Let M be a set of messages and $S \subseteq \mathbb{N}$ a set of sequence numbers. Only one sender node $\sigma \in \mathcal{N}$ can broadcast messages (we hereon write sb-cast) by invoking $\langle \text{SB-CAST} | sn, m \rangle$ with $(sn, m) \in S \times M$. $\langle \text{SB-DELIVER} | sn, m \rangle$ is triggered at a correct node p when p delivers (we hereon write sb-delivers) message m with sequence number sn .

If a correct node suspects that σ is quiet, all correct nodes are allowed to sb-deliver a special *nil* value $m = \perp \notin M$. If, however, σ is trusted by all correct nodes, all correct nodes are guaranteed to sb-deliver non-nil messages $m \neq \perp$.

SB is explicitly initialized with an $\langle \text{SB-INIT} \rangle$ event. We assume a failure detector list at each correct node which is initially empty. It is only after the invocation of $\langle \text{SB-INIT} \rangle$ that suspecting σ can lead to the \perp value being delivered.

An instance of $SB(\sigma, S, M, D)$ has the following properties:

SB1 Integrity: If a correct node sb-delivers (sn, m) with $m \neq \perp$ and σ is correct then σ sb-cast (sn, m) .

SB2 Agreement: If two correct nodes sb-deliver, respectively, (sn, m) and (sn, m') , then $m = m'$.

SB3 Termination: If p is correct, then p eventually sb-delivers a message for every sequence number in S , i.e., $\forall sn \in S : \exists m \in M \cup \{\perp\}$ such that p sb-delivers (sn, m) .

SB4 Eventual Progress: If some correct node sb-delivers (sn, \perp) for some $sn \in S$, then some correct node suspected σ after having invoked $\langle \text{SB-INIT} \rangle$.

The key difference from TOB is that SB is invoked for an explicit set of sequence numbers S and messages M . Moreover, SB is invoked with a $\diamond S(bz)$ failure detector and correct nodes deliver messages from set M and the special \perp value such that SB terminates for all sequence numbers. The latter

is guaranteed by the \perp value and $\diamond S(bz)$ completeness; if σ is quiet it will eventually be suspected by all correct nodes.

Our technical report[32] shows that SB is implementable with consensus, Byzantine reliable broadcast (BRB), and $\diamond S(bz)$. On a high level, the dedicated sender σ reliably broadcasts a message for each sequence number in S and all correct nodes run consensus for each sequence number in S to decide if they can deliver a message broadcast by σ or if σ is quiet. In the first case they sb-deliver the message and in the second case they sb-deliver \perp . Since both BRB and $\diamond S(bz)$ are implementable with consensus, SB can also be implemented with consensus.

2.3 Multiplexing Instances of SB with ISS

ISS multiplexes instances of SB to implement SMR. Each node maintains a log of ordered messages which correspond to batches of client requests. Each position in the log corresponds to a *sequence number* signifying the offset from the start of the log. The log is partitioned in subsets of sequence numbers called *segments*. Each segment corresponds to one instance of SB. Nodes obtain requests from clients and, after mapping them to a log position using an instance of SB, deliver them together with the assigned sequence number.

ISS proceeds in *epochs* identified by monotonically increasing integer *epoch numbers*. Each epoch e is associated with a set of segments. The union of those segments forms a set $Sn(e)$ of consecutive sequence numbers. Epoch 0 (the first epoch) starts with sequence number 0. The mapping of sequence numbers to epochs is a function known to all nodes with the only requirements being that it is monotonically increasing and that there are no gaps between epochs. More formally, $\max(Sn(e)) + 1 = \min(Sn(e + 1))$. Epoch length can be arbitrary, as long as it is finite. For simplicity, we use a fixed, constant epoch length.

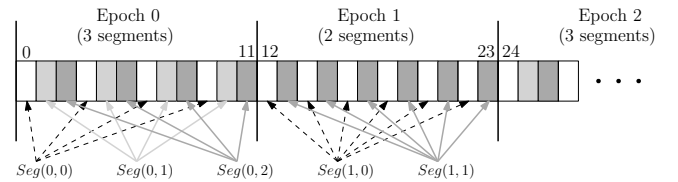


Figure 1. Log partitioned in epochs and segments. In this particular example, each epoch is 12 sequence numbers long. The first epoch has 3 segments while the second epoch only has 2, i.e., $|Leaders(0)| = 3$, $|Leaders(1)| = 2$, $\max(Sn(1)) = 23$ and $\max(Seg(0, 1)) = 10$.

Epochs are processed sequentially, i.e., ISS first agrees on the assignment of request batches to all sequence numbers in $Sn(e)$ before starting to agree on the assignment of request batches to sequence numbers in $Sn(e + 1)$.

Within an epoch, however, ISS processes segments in parallel. Multiple leaders, selected according to a *leader selection*

policy, concurrently propose batches of requests for different sequence numbers in $Sn(e)$. To this end, ISS assigns a different *leader* node to each segment in epoch e . We refer to the set of all nodes acting as leaders in an epoch as the leaderset of the epoch. The numbers of leaders and segments in each epoch always match. We denote by $Seg(e, i)$ the subset of $Sn(e)$ for which node i is the leader. This means that node i is responsible for proposing request batches to sequence numbers in $Seg(e, i)$. No node other than i can propose batches for sequence numbers in $Seg(e, i)$. Let $Leaders(e)$ be nodes that are leaders in epoch e . We associate sequence numbers with segments in a round-robin way, namely, for $0 \leq i < |Leaders(e)|$,

$$Seg(e, i) \subseteq Sn(e) = \{sn \in Sn(e) \mid i \equiv sn \bmod |Leaders(e)|\}$$

An example with $|Leaders(0)| = 3$ and $|Leaders(1)| = 2$ is illustrated in Figure 1.

In principle, any assignment of sequence numbers to segments is possible and leads to a correct algorithm. We choose the round-robin assignment because it uniformly distributes sequence numbers among instances. Therefore, in a fault-free execution, it is the least likely to create “gaps” in the log, which minimizes the end-to-end request latency.

In order not to waste resources on duplicate requests, we require that a request cannot be part of two batches assigned to two different sequence numbers. We enforce this at three levels: (1) within a segment, (2) across segments in the same epoch, and (3) across epochs.

Within a segment, we rely on the fact that a correct leader will propose (and a correct node, as follower, will accept) only batches with disjoint sets of requests for each sequence number within a segment. Across segments, we partition the set of all possible requests into *buckets* using a hash function and enforce that only requests from different buckets can be used for different segments within an epoch. We denote by \mathcal{B} the set of all possible buckets. We assign a subset of \mathcal{B} to each segment, such that each bucket is assigned to exactly one segment in each epoch. We denote by $Buckets(e, i) \subseteq \mathcal{B}$ the set of buckets assigned to leader i in epoch e . For simplicity, we say that we assign a bucket to a leader i when assigning a bucket to a segment for which i is the leader. Across epochs, ISS prevents duplication by only allowing a node to propose a request batch in a new epoch once it has added all batches from the previous epoch to the log. If a request has been delivered in a batch in the previous epoch, a correct leader will not propose it again (see also Section 3.3). Also, a correct node, as follower, will not accept a proposal which includes a previously delivered request.

In summary, a segment of epoch e with leader i is defined by the tuple $(e, i, Seg(e, i), Buckets(e, i))$. For a set of buckets $B \subseteq \mathcal{B}$, we denote with $batches(B)$ the set of all possible batches consisting of valid (we define request validity precisely later in Section 3.7) requests that map to some bucket in B . For each segment $(e, i, Seg(e, i), Buckets(e, i))$, we use

an instance $SB(i, batches(Buckets(e, i)), Seg(e, i), D)$ of Sequenced Broadcast where D is a failure detector. We say that leader i *proposes* a batch $b \in batches(Buckets(e, i))$ for sequence number $sn \in Seg(e, i)$ if i sb-casts b with a sequence number sn at the corresponding instance of SB. A batch b *commits* with sn (and is added to the log at the corresponding position) at node n when, for the corresponding instance of SB, n sb-delivers b with sn .

During epoch e , all nodes that are leaders in e simultaneously propose batches for sequence numbers in their corresponding segments. ISS multiplexes all segments into the single common log as shown in Figure 1. Each node thus executes $|Leaders(e)|$ SB instances simultaneously, while being a leader for at most one of them.

Epoch e ends and epoch $(e + 1)$ starts when all sequence numbers in $Sn(e)$ have been committed. Nodes keep the old instances active until all corresponding sequence numbers become part of a stable checkpoint. This is necessary for ensuring totality (even for slow nodes which might not have actively taken part in the agreement).

2.4 Assigning Buckets to Segments

ISS partitions the request hash space into buckets which it assigns to leaders/segments and changes this bucket assignment at epoch transitions. At any point in time, a leader can assign sequence numbers only to requests from its assigned buckets. This approach was first used in Mir-BFT [31] to counter request duplication and censoring attacks.

During an epoch, the assignment of buckets to leaders is fixed. To ensure liveness, each bucket must repeatedly be assigned to a correct leader. To this end, ISS re-assigns the buckets on every epoch transition as follows.

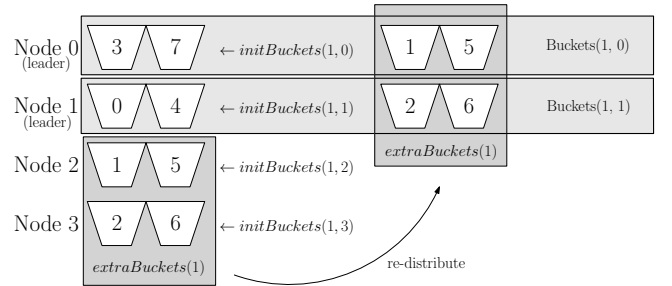


Figure 2. Example assignment of 8 buckets to 2 leaders in a system with 4 nodes in epoch 1.

For epoch e , we start by assigning an initial set of buckets to *each node* (leader or not) in a round-robin way. Let $initBuckets(e, i) \subseteq \mathcal{B}$ be the set of buckets initially assigned to each node i , $0 \leq i < n$ in epoch e . We consider the buckets in \mathcal{B} to be numbered, with each bucket having an integer bucket number $b \in \{0, \dots, |\mathcal{B}| - 1\}$. In the following we refer to buckets using b .

$$initBuckets(e, i) = \{b \in \mathcal{B} \mid (b + e) \equiv i \bmod n\}$$

However, not all nodes belong to $Leaders(e)$. Let $extraBuckets(e)$ be the set of buckets initially assigned to non-leaders.

$$extraBuckets(e) = \{b \in \mathcal{B} \mid \exists i : i \notin Leaders(e) \wedge b \in initBuckets(e, i)\}$$

We must re-distribute those extra buckets to the leaders of epoch e . We do this in a round robin way as well. Let $l(e, k)$, $0 \leq k < |Leaders(e)|$ be the k -th leader (in lexicographic order) in epoch e . The $Buckets(e, l(e, k))$ of the k -th leader in e are thus defined as follows.

$$Buckets(e, l(e, k)) = initBuckets(e, l(e, k)) \cup \{b \in extraBuckets(e) \mid (b + e) \equiv k \pmod{|Leaders(e)|}\}$$

An example bucket assignment is illustrated in Figure 2.

With this approach, all buckets are assigned to leaders and every node is eventually assigned every bucket at least through the initial bucket assignment. ISS ensures liveness as long as, in an infinite execution, there is a correct node that (1) eventually stops being suspected forever by every correct node, and (2) is assigned each bucket infinitely often. (1) is satisfied by the properties of the eventually strong failure detector. (2) is satisfied by the bucket re-assignment and the leader selection policy described in Section 3.4.

Figure 3 shows how ISS multiplexes SB instances to build a totally ordered log.

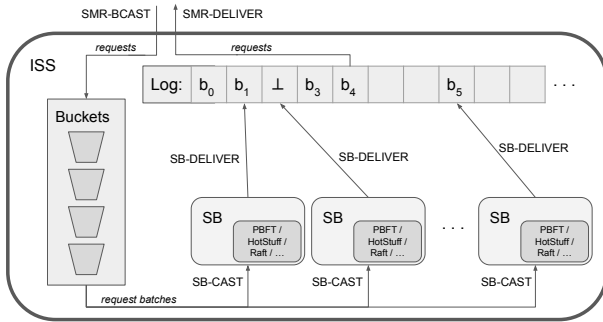


Figure 3. SB invocation in ISS.

3 ISS Algorithm Details

In this section we present the ISS algorithm in detail. The main high-level algorithm that produces a totally ordered log is described in Algorithm 1. For better readability, certain auxiliary functions and the functions related to epoch initialization are presented separately in Algorithms 2 and 3, respectively. The notation $\langle i, e | args \rangle$ corresponds to an event e of instance i with arguments $args$.

As described in the previous sections, ISS proceeds in epochs, each epoch multiplexing multiple segments into a final totally ordered log. We start with epoch number 0 (line

16) and an empty log (line 21). All buckets are initially empty (line 25). Whenever a client submits a new request (line 30), ISS adds the request to the corresponding bucket (line 32).

We assume access across all epochs to a module D that implements an eventually strong failure detector (Section 2.2).

Algorithm 1 Main ISS algorithm for node p

```

1: Implements:
2:   State Machine Replication, instance smr
3:
4: Uses:
5:   Sequenced Broadcast, instance sb( $\sigma, S, M, d$ ) (multiple instances)
6:   with sender  $\sigma$ , sequence numbers  $S$ , messages  $M$ , failure detector instance  $d$ 
7:   Eventually Strong Failure Detector instance  $D$ 
8:
9: Parameters:
10:  numBuckets
11:  batchSize
12:  epochLength
13:  maxBatchSize
14:
15: upon init() do
16:   currentEpoch  $\leftarrow 0$ 
17:   segments  $\leftarrow \emptyset$ 
18:   firstUndelivered  $\leftarrow 0$ 
19:   totalDelivered  $\leftarrow 0$ 
20:   for  $sn \in \mathbb{N}$  do
21:     log[sn]  $\leftarrow none$ 
22:     proposed[sn]  $\leftarrow none$ 
23:   end for
24:   for  $0 \leq b < numBuckets$  do
25:     buckets[b]  $\leftarrow$  empty bucket
26:   end for
27:   initEpoch(0)
28:   runEpoch(0)
29:
30: upon event  $\langle smr, SMR-CAST | req \rangle$  do
31:   if valid(req) then
32:     buckets[hash(req)].add(req)
33:   end if
34:
35: function runEpoch( $e$ ):
36:   for  $sn \in seqNrs(e)$  with  $segOf(sn).leader = p$  do
37:     propose(sn)
38:   end for
39:
40: upon event  $\langle segOf(sn).SB, SB-DELIVER | (sn, batch) \rangle$  do
41:   log[sn]  $\leftarrow batch$ 
42:   if batch  $\neq \perp$  then
43:     for  $req \in batch$  do
44:       buckets[hash(req)].remove(req)
45:     end for
46:   else if proposed[sn]  $\neq none$  then
47:     resurrectRequests(proposed[sn])
48:   end if
49:
50: upon  $\forall sn \in seqNrs(currentEpoch) : log[sn] \neq none$  do
51:   initEpoch(currentEpoch + 1)
52:   runEpoch(currentEpoch + 1)
53:
54: upon log[firstUndelivered]  $\neq none$  do
55:   deliver(sn)
56:   firstUndelivered  $\leftarrow sn + 1$ 
57:

```

3.1 Epoch Initialization

At the start of each epoch e , ISS: (1) calculates $Leaders(e)$, the set of nodes that will act as leaders in e , based on the used leader selection policy (line 98), (2) for each node l in $Leaders(e)$, creates a new segment with leader l (lines 100 and 101), (3) assigns all sequence numbers $Sn(e)$ of epoch

Algorithm 2 Auxiliary functions

```

58: function segOf(sn) :
59:   return seg ∈ segments : sn ∈ seg.seqNrs
60:
61: function seqNrs(e) :
62:   return {i ∈ ℕ |
63:     e · epochLength ≤ i < e · (epochLength + 1)}
64:
65: function propose(sn) :
66:   wait until   segments[sn].batchReady()
67:   or batchTimeout
68:
69:   batch ← cutBatch(segOf(sn).buckets)
70:   trigger (seg.SB, SB-CAST[sn, batch])
71:   proposed[sn] ← batch
72:   if batch ≠ ⊥ then
73:     for req ∈ batch do
74:       buckets[hash(req)].remove(req)
75:     end for
76:   end if
77:
78: function deliver(sn) :
79:   batch ← log[sn]
80:   if batch ≠ ⊥ then
81:     for req ∈ batch do
82:       trigger (smr, SMR-DELIVER|req, totalDelivered)
83:       totalDelivered ← totalDelivered + 1
84:     end for
85:   end if
86:
87: function cutBatch(B) :
88:   return a batch of the maxBatchSize oldest requests in B
89:
90: function resurrectRequests(batch) :
91:   for req ∈ batch do
92:     buckets[hash(req)].add(req)
93:   end for
94:

```

e in a round-robin way to all created segments (line 103), (4) assigns buckets to the created segments as described in Section 2.4 (line 104), and (5) creates an instance of SB for each created segment (line 106).

Algorithm 3 Epoch initialization

```

95: function initEpoch(e) :
96:   currentEpoch ← e
97:
98:   leaders ← leader_selection_policy(e)
99:   for 0 ≤ l < |leaders| do
100:     seg ← new segment
101:     seg.leader ← leaders[l]
102:     seg.seqNrs ← {sn ∈ seqNrs(e) |
103:       sn ≡ l mod |leaders|}
104:     seg.buckets ← Buckets(e, leaders[l])
105:     seg.SB ← sb(leaders[l], batches(seg.buckets), seg.seqNrs, D)
106:     trigger (seg.SB, SB-INIT)
107:     segments ← segments ∪ {seg}
108:   end for
109:
110: function Buckets(e, i) :
111:   for n ∈ N do
112:     initBuckets[n] ← {b < numBuckets |
113:       (e + b) ≡ n mod |N|}
114:   end for
115:   leaders ← epLeaders_POLICY(e)
116:   extraBuckets ← {b : ∃n |
117:     n ∉ leaders ∧ b ∈ initBuckets[n]}
118:   return initBuckets[i] ∪ {b ∈ extraBuckets |
119:     (e + b) ≡ (index of i in leaders) mod |leaders|}
120:

```

3.2 Ordering Request Batches

ISS orders requests in batches, a common technique which allows requests to be handled in parallel, amortizes the processing cost of protocol messages, and, thereby, improves throughput. During an epoch, every node *l* that is the leader of a segment *s* proposes request batches for sequence numbers assigned to *s* (line 37). *l* does so by sb-casting the batches using the instance of SB associated with *s*. Every node then inserts the sb-delivered (sequence number, batch) pairs at the corresponding positions in its copy of the log (line 41). We say the node *commits* the batch with the corresponding sequence number since, once inserted to the log, the assignment of a batch to a sequence number is final.

Proposing Batches. Each node maintains local data structures of *buckets queues*, which store the received and not yet proposed or delivered requests corresponding to the respective bucket. To propose a request batch for sequence number *sn*, *l* first constructs the batch using requests in the bucket queues corresponding to the buckets assigned to *s*. To implement efficient request batching while preserving low latency, *l* waits until at least one of the following conditions is fulfilled: (1) the bucket queues assigned to *s* contain enough requests (more than a predefined *batchSize*) (line 67), or (2) a predefined time elapses since the last proposal (line 68). Under low load, this condition sets an upper bound on the pending¹ latency of requests waiting to be proposed, even if the batch is filling slowly.

l then constructs a *batch* using up to *batchSize* requests (line 69), removes those requests from their bucket queues and proposes the batch by invoking *SB-CAST*(*sn*, *batch*) on the SB instance associated with *s* (line 70).

Every leader also keeps track of the values it proposed for each sequence number (line 71). This, as we explain later, is important in the case of asynchrony.

Assembling the Final Log. Whenever any instance of SB (belonging to any segment) delivers a value associated with a sequence number (line 40) at some node *n*, *n* inserts the delivered value at position *sn* of the log (line 41). If a request batch ≠ ⊥ has been delivered (line 42), *n* removes the contained requests from their corresponding bucket queues (line 44) to avoid proposing them again in a later epoch.

Note that bucket queues are local data structures at each node, and thus each node manages its bucket queues locally. Nodes add all requests they obtain from clients to their local bucket queues, but only propose batches constructed from the queues corresponding to the buckets assigned to their segments. Consequently, a node *n* delivers batches containing requests mapping to other buckets than those *n* uses for proposing. Therefore, to avoid request duplication across epochs, each node must remove all delivered requests from its local bucket queues.

¹Latency from the moment a request is received until it is added in a batch.

If the special value \perp has been delivered by SB and, at the same time, n itself had been the leader proposing a batch for sn (line 46), n “resurrects” all requests in the batch it had proposed (line 47) by returning them to their corresponding bucket queues (line 92). This scenario can appear in the case of asynchrony / partitions, where a correct leader is suspected as faulty after having proposed a batch. Such a leader must return the unsuccessfully proposed requests in their bucket queues and, if batches with those requests are not committed by other nodes in the meantime, retry proposing them in a later epoch where it is again leader of a segment with those buckets.

A node considers the ordering of a request finished when it is part of a committed batch with an assigned sequence number sn and the log contains an entry for each sequence number $sn' \leq sn$.

Each request is delivered with a unique sequence number sn_r denoting the total order of the request. sn_r is derived from the sequence number of the batch in which the request is delivered and the position of the request in the batch. Let S_{sn} be the number of requests in a batch delivered with sequence number sn and let r be the k^{th} request in this batch. For each such request r , ISS outputs $\langle \text{SMR-DELIVER} | sn_r, r \rangle$ where:

$$sn_r = k + \sum_{i=0}^{sn-1} S_i \quad (1)$$

3.3 Advancing Epochs

ISS advances from epoch e to epoch $e + 1$ when the log contains an entry for each sequence number in $Sn(e)$ (line 50). This will eventually happen for each epoch at each correct node due to SB *Termination*. Only then does the node start processing messages related to epoch $e+1$ and starts proposing batches for sequence numbers in $Sn(e + 1)$ (lines 51 and 52).

Requiring a node to have committed all batches in epoch e before proposing batches for $e+1$ prevents request duplication across epochs. When a node transitions from e to $e + 1$, no requests are “in flight”—each request has either already been committed in e or has not yet been proposed in $e + 1$.

3.4 Selecting Epoch Leaders

In order to guarantee that each request r submitted by a correct client is ordered (liveness), we must ensure that, eventually, there will be a segment in which r is committed. As implied by the specification of SB, this can only be guaranteed if a correct leader p proposes a batch containing r and the failure detector does not suspect p until r is committed. The choice of epoch leaders is thus crucial.

ISS selects leaders according to a *leader selection policy*, a function known to all nodes that, at the end of each epoch e , determines the set of leaders for epoch $(e + 1)$.

In order to guarantee liveness of the system, the leader selection policy must ensure, for each bucket b , that, in an infinite execution, b will be assigned infinitely many times

to a segment with a correct leader that is not suspected by the failure detector. Weak eventual accuracy (see Section 2.2) guarantees that there exists such a leader. Different leader selection policies pose different trade-offs with respect to performance. For simplicity, we adopt the policy of BFT-Mencius [27] because it trivially ensures liveness by always keeping enough correct nodes in the leaderset.

3.5 Checkpointing and State Transfer

ISS implements a simple checkpointing protocol. Every node i , in each epoch e , when the log contains an entry for each sequence number in $Sn(e)$, broadcasts a signed message $\langle \text{CHECKPOINT}, \max(Sn(e)), \mu(e), \sigma_i \rangle$, where $\mu(e)$ is the Merkle tree root of the digests of all the batches in the log with sequence numbers in $Sn(e)$. Upon acquiring a strong quorum of $2f + 1$ matching CHECKPOINT messages with a valid signature against the sender node’s public key, node i creates a stable checkpoint $\langle \text{STABLE-CHECKPOINT}, \max(Sn(e)), \pi(e) \rangle$, where $\pi(e)$ is the set of $2f + 1$ signatures on the CHECKPOINT messages. At this point, i can garbage collect all segments of epoch e .

When a node i has fallen behind, for example when i starts receiving messages for a future epoch, i performs a state transfer, i.e., it fetches the missing log entries along with their corresponding stable checkpoints which prove the integrity of the data.

ISS checkpointing is orthogonal to any checkpointing and state transfer mechanism pertaining to the SB implementation because SB instances must terminate independently.

3.6 Membership Reconfiguration

A detailed membership reconfiguration protocol is outside the scope of this paper. However, we outline a solution. Thanks to SB 3 (Termination), all correct nodes eventually deliver a value for each sequence number of an epoch. Moreover, thanks to SMR 2 (Agreement), all correct nodes assemble the same log at the end of the epoch. Therefore, the log at the end of the epoch can be used to deterministically make decisions for the next epoch, including decisions about nodes and clients joining/leaving the set of system processes. Such a decision can be based, for example, on a flagged reconfiguration request proposed by a manager process [30] which becomes part of the log.

3.7 Request Handling

A request $r = (o, id)$ with payload p and identifier $id = (t, c)$ sent from a client to a node is wrapped in a signed message. Our implementation represents the client identifier $r.id.c$ with an integer associated with the client’s c public key. The signature is calculated over the request identity $r.id$ and payload $r.o$ to guarantee integrity and authenticity.

Similarly to Mir-BFT [31], clients can submit multiple requests in parallel within a client watermark window, i.e., a contiguous set for the per-client request sequence number

$r.id.t$. ISS advances all clients' watermark widows at the end of each epoch.

Each correct node, upon receiving a valid request, adds the request, based on its identifier, to the corresponding bucket queue. A request is considered *valid* if: (1) it has a valid signature (2) the public key corresponds to a client in the client set C of the system, and (3) is within the client watermarks. Bucket queues are idempotent, i.e., each correct node adds a request to the corresponding bucket queue exactly once. Moreover, the bucket queue implementation maintains a FIFO property to guarantee liveness with the oldest request always being proposed first.

Requests are uniformly distributed to buckets using modulo as a hash function. With $|\mathcal{B}|$ denoting the total number of buckets and $||$ denoting concatenation, each request r maps to a bucket $b = r.id.c || r.id.t \bmod |\mathcal{B}|$.

We exclude the payload of the request from the bucket mapping function to prevent malicious clients from biasing the uniform distribution. In a permissioned system the client cannot assume different identities and may only bias the outcome of the hash function by choosing the request sequence number. However, we limit the available sequence numbers for each client, and, therefore, their ability to bias the request distribution, with the client watermarking mechanism.

Request execution. ISS is oblivious to the payload of the requests for general applicability. Execution is not part of ISS; however, it can be coupled with any application that requires a total order of requests. Moreover, request execution against a state machine is straight-forward. A request, as part of a batch, is considered part of the log (and can be, therefore, executed) once all previous batches are added to the log (see Section 3.2). Therefore, a request can be executed as soon as it is delivered by ISS. This does not require the epoch, in which the request is added to the log, to finish.

4 ISS Implementation

We implement ISS in Go, using gRPC for communication with TLS on nodes with two network interfaces: one for client-to-node and one for node-to-node communication.

Our implementation is highly concurrent: multiple threads are handling incoming client requests, verifying request signatures, sending / receiving messages, and executing various sub-protocols such as checkpointing and fetching missing protocol state. Each SB instance also executes in its own thread. A separate thread orchestrates all of the above.

In the rest of this section, we discuss our implementation of the ordering subprotocols (4.1), the interaction between ISS and its clients (4.2), and crucial technical aspects for achieving robustness and high performance (4.3).

4.1 SB Implementation

In this section we discuss (1) how we implement SB with different leader-driven consensus protocols and (2) adaptations in the leader-driven protocols critical for ISS performance.

All protocol implementations adhere to the following common design principles:

1. We initialize the protocol such that the the first protocol leader is the segment leader (dedicated SB sender) and all other nodes of the system participate as followers.
2. After a leader change, any new leader (including the segment leader if it becomes leader again), only proposes \perp values for any sequence number not initially proposed by the segment leader.²
3. A follower accepts a proposal only if (a) all requests in the batch are valid according to Section 3.7, (b) no request in the batch has previously been proposed in the same epoch or committed in a previous epoch (c) all requests belong to the buckets of the segment, and (d) either the segment leader sb-casts the proposal, or the proposal is \perp .

4.1.1 PBFT. We follow the PBFT protocol as described by Castro and Liskov [10], with a few adaptations.

Our implementation need not deal with timeouts at the granularity of single requests, as PBFT does. To prevent censoring attacks (and thus ensure liveness), a PBFT replica initiates a view change if any request has not been committed for too long. Since ISS prevents censoring attacks by bucket re-assignment, it is sufficient for us to make sure to commit *some batch* before a timeout fires and reset this timer when committing any batch. In the absence of incoming requests, the primary periodically proposes an empty batch to prevent a view change. Moreover, for simplicity, we implement view-change with signatures according to Castro and Liskov [9].

4.1.2 HotStuff. We implement chained HotStuff according to Yin *et al.* [33] with BLS [7] threshold signatures using DEDIS library for Go [1].

In our implementation, each batch corresponds to a HotStuff command, and each segment sequence number corresponds to a HotStuff view. Each segment is implemented as a new HotStuff instance with a new root certificate QC_0 . To ensure that all sequence numbers can be delivered, i.e., to ensure that we can always “flush” the pipeline of chained HotStuff, we extend the segment with 3 dummy sequence numbers corresponding to dummy empty batches which are not added to the ISS log. Figure 4 demonstrates an example of a segment with 3 sequence numbers.

4.1.3 Raft. Briefly, in Raft, nodes set a random timer within a configurable range, which they reset every time they receive a message from the elected leader. If the timer fires, the node advances to a new term (similar to PBFT view) and enters an election phase as a candidate leader. An elected leader periodically sends append-entry requests for new values, possibly also empty, as a heartbeat. The leader collects responses according to which it might resend to the followers any value they declare as missing.

²Enforcing \perp values is necessary for SB to be implementable. Otherwise, both SB *Integrity* and SB *Termination* cannot be satisfied at the same time.

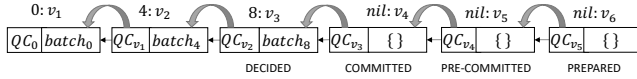


Figure 4. Chained HotStuff execution for a segment with 3 sequence numbers $\{0, 4, 8\}$. When view v_6 is prepared, v_3 is decided and $batch_8$ can be added to the log. QC_{v_i} is the HotStuff quorum certificate for view i (e.g., a threshold signature from $2f + 1$ followers) on the proposal of view v_i .

We implement Raft according to [28] with minor adaptations. We fix the first leader to be the leader of the segment, skipping the election phase. Until the end of the segment, the leader periodically sends append-entry requests containing batches. The leader continues to send empty append-entry requests until the end of the epoch to guarantee that enough nodes have added all the batches of the segment to their log.

4.2 Interaction with Clients

Clients submit requests to ISS by sending signed request messages to nodes. When a node delivers a request r as described in Section 3.2, it sends a response message to the client that submitted r . When the client obtains a quorum of responses, it considers the request delivered.

To guarantee Liveness (Section 2.1, SMR4), a client must ensure that at least one correct leader eventually receives the request. A trivial solution is to send the request to all nodes. However, ISS implements an optimistic *leader detection* mechanism to help the clients find the correct leader for each request faster and to better load balance request processing among the nodes. At any point in time, the bucket to which a request belongs is assigned to a single segment. Thus, the client only needs to send its request to the node currently serving as a leader for the corresponding segment.

ISS keeps the clients updated about the assignment of buckets to leaders. At each epoch transition, all nodes send a message with the assignment for the next epoch to all clients. A client accepts such a message once it receives it from a quorum of nodes. The client submits all subsequent requests for this epoch to the appropriate node. Moreover, it resubmits all requests submitted in the past that have not yet been delivered. This guarantees that all correct nodes will eventually receive the request, which ensures liveness.

To make sure that, in most cases, a leader already has a request when it is that leader’s turn to propose it, the client sends its request to two additional nodes that it projects to be assigned the corresponding bucket in the next two epochs.

4.3 Important Technical Aspects

4.3.1 Rate-limiting Proposals in PBFT. PBFT’s ability to send proposals in parallel is instrumental for achieving high throughput. However, as soon as a load spike or a temporary glitch in network connectivity occurs (as it regularly does on the used cloud platform), the leader can end up trying to send

too many batches in parallel. Due to limited aggregate bandwidth, *all* those batches will take longer to transfer, triggering view change timeouts at the replicas.

We address this issue by setting a hard limit on the rate of sending batches “on the wire”, allowing (the most part of) a batch to be transmitted before the transfer of the next batch starts. This measure limits peak throughput but is effective at protecting against unnecessary view changes.

4.3.2 Concurrency Handling. A naive approach to handling requests, where each client connection is served by a thread that, in a loop, receives a request, locks the corresponding bucket queue, adds the request, and unlocks the bucket queue, is detrimental to performance. We attribute this to cache contention on the bucket queue locks.

Still, access to a bucket does have to be synchronized, as adding (on reception) and removing (on commit) requests must happen atomically. At the same time, an efficient lock-free implementation of a non-trivial data structure such as a bucket queue could be a research topic on its own.

We thus dedicate a constant, limited number of threads (as many as there are physical CPU cores) to only adding requests to bucket queues, such that each bucket queue is only accessed by one thread, removing most of the contention. The network-handling threads pass the received requests to the corresponding bucket-adding threads using a lock-free data structure optimized for this purpose (a Go channel).

4.3.3 Deployment, Profiling, and Analysis. ISS comes with tools for automating the deployment of hundreds of experiments across hundreds of geo-distributed nodes on the cloud and for analyzing their outputs. They profile (using `pprof`) the execution at each node, pinpointing lines of code that cause stalling or high CPU load. For example, the above-mentioned cache contention was pointed to by the profiler. They also plot various metrics over time, such as the size of proposed batches, commit rate, or CPU load. Automatic exploration of the multi-dimensional parameter space proved essential for understanding the inner workings of the system.

Using hundreds of cloud machines with hourly billing also incurs significant cost. Automatically commissioning cloud machines only for the time strictly necessary to run our experiments and releasing those resources as soon as possible most likely saved thousands of dollars.

5 Correctness

We now prove that multiplexing SB instances with ISS implements an SMR service as defined in Section 2.1.

SMR1 Integrity: If a correct node delivers (sn, r) , where $r.id.c$ is a correct client’s identity, then client c broadcast r .

Proof. A correct node only delivers request r if it is inserted in the log as part of a committed batch b (line 55). In turn, b is added in the log only upon an event $\langle sb, SB\text{-DELIVER}|sn_b, b \rangle$,

where sb is a $SB(\sigma, M, S, D)$ instance (line 40) and M the set of all possible valid batches in the buckets of the segment. A correct node only invokes $\langle sb, SB\text{-DELIVER}|sn_b, b \rangle$ with a batch b in the set of valid batches M . A validity condition is that every request of the batch has a valid signature (see Section 3.7). Since $r.id.c$ is the only process able to produce a valid signature, c must have broadcast r . \square

SMR2 Agreement: If two correct nodes deliver, respectively, (sn, r) and (sn, r') , then $r = r'$.

Proof. Let r be in batch b committed with sn_b and r' in a batch b' committed with $sn_{b'}$. For r and r' to have the same sequence number, by Equation (1) and by the same log established by SB2 (Agreement), $sn_b = sn_{b'}$. Since b and b' are delivered with the same sequence number, they belong to the same segment S and, thus, also in the same set M for which an instance $SB(\sigma, M, S, D)$ was initialized. Thus, by SB2, $b = b'$, and by Equation (1) $r = r'$. \square

Lemma 5.1. *If a correct node initializes $SB(\sigma, M, S, D)$ then every correct node eventually initializes $SB(\sigma, M, S, D)$.*

Proof. By the ISS algorithm, the SB instances (including their parameters σ, M, S, D) are initialized at the beginning of each epoch e depending only on the state of the node's log at the beginning of e . We prove the lemma by induction on the epoch e . In epoch 0 all nodes have an empty log. Thus, they instantiate SB with the same parameters.

We now assume that in epoch e correct nodes i and j initialize SB with the same parameters. Then, by SB 3 (Termination) and by SMR 2 (Agreement) both i and j advance from epoch e to epoch $e+1$ with the same log. Thus, both i and j initialize SB with the same parameters, i.e., sender σ , message set M and segments S . \square

SMR3 Totality: If a correct node delivers request (sn, r) , then every correct node eventually delivers (sn, r) .

Proof. Let us assume that some correct node i delivers r with sequence number sn . Let j be some other correct node. From Equation (1), request r , delivered by i , uniquely corresponds to some batch b with sequence number sn_b in the log of i which i has committed. Therefore, there exists an instance $SB(\sigma, M, S, D)$ which outputs SB-DELIVER(sn_b, b) at node i . By Lemma 5.1, j also eventually initializes $SB(\sigma, M, S, D)$. Then SB 3 (Termination) guarantees that for each sequence number in S and, therefore, for also sn_b , j delivers a message $m \in M \cup \perp$. Moreover, property SB 2 (Agreement) guarantees that $m = b$. It follows that j delivers r for sn . \square

Lemma 5.2. *In an infinite ISS execution, there exists some correct node which is eventually in the leaderset forever without being suspected.*

Proof. By the Eventual Weak Accuracy of the failure detector, there exists some correct node p that will eventually stop being suspected by any correct node. Eventually, either p will

be in the leaderset forever, or there exists at least $2f + 1$ other nodes that stop being suspected forever – otherwise p would eventually be included in the leaderset. Among those $2f + 1$ nodes there are at least $f + 1$ correct ones. Therefore, in any case, some correct node which stops being suspected remains in the leaderset forever. \square

Lemma 5.3. *Let p be a correct node which is eventually in the leaderset forever without being suspected by any correct node after time t . If p sb-casts a message m after time t , then p eventually sb-delivers m .*

Proof. By SB 3 (Termination), p sb-delivers a message for all sequence numbers for any SB instance. Let us assume, by contradiction, that p sb-casts message m after time t in a SB instance but delivers a set of messages \mathcal{M} for the sequence numbers of the instance, such that $m \notin \mathcal{M}$. By p not being suspected and SB 4 (Eventual Progress), for the SB instances for which p is the dedicated sender, p delivers non \perp values for all sequence numbers. By SB 1 (Integrity), all messages in \mathcal{M} are sb-cast by p . A contradiction to $m \notin \mathcal{M}$, since p , being correct, sb-casts only one message per sequence number. Therefore, p sb-delivers m . \square

SMR4 Liveness: If a correct client broadcasts request r , then some correct node eventually delivers (sn, r) .

Proof. Let us assume by contradiction that r is never delivered by any correct node. This implies that every correct node puts r in their bucket queue (by the correct client re-transmitting r forever, see Section 3.7). Eventually, after some time t , by Lemma 5.2 there will be at least one correct, unsuspected node i in the leaderset forever.

Let $R(r)$ be the set of all requests received by i before receiving r . Let r map to a bucket \hat{b} . If $R(r) = \emptyset$, i.e., r is the oldest request of i , either node i delivers r by time t or r remains in i 's corresponding bucket queue. By bucket re-assignment, i will invoke SB-CAST infinitely many times with batches (messages) containing requests from bucket \hat{b} . Therefore, eventually, i sb-casts a batch b which contains r and by Lemma 5.3 i delivers b . Let sn be the sequence number with which b is sb-delivered. By SB3 (Termination), all correct nodes sb-deliver and add in their log all sequence numbers in the segment of sn . Therefore, by the ISS algorithm (line 55), they deliver all requests in b , including r .

We can prove by induction on the size of $R(r)$ that r is delivered by some correct node. A contradiction to r not being delivered. \square

6 Evaluation

Our implementation is modular, allowing easy switching between different protocols implementing SB. We use 3 well-known protocols for ordering requests: PBFT [10] (BFT), HotStuff [33] (BFT), and Raft [28] (CFT). We evaluate the

impact ISS has on these protocols by comparing its performance to their respective original single-leader versions. In addition, we compare ISS to Mir-BFT [31] which also has multiple leaders. We do not compare, however, to other multi-leader protocols that do not prevent request duplication (e.g., Hashgraph [23], Red Belly [13], RCC [20], OMADA [16], BFT-Mencius [27]). The codebase of these protocols is unavailable or unmaintained. Moreover, Mir-BFT evaluation demonstrates that the performance of this family of protocols deteriorates as the number of nodes increases in the presence of duplicate requests. For the same reason, we also do not compare to trivially running multiple instances of the single leader protocols.

6.1 Experimental Setup

We perform our evaluation in a WAN which spans 16 datacenters across Europe, America, Australia, and Asia on IBM cloud. All processes run on dedicated virtual machines with 32 x 2.0 GHz VCPUs and 32GB RAM running Ubuntu Linux 20.04. All machines are equipped with two network interfaces, public and private, rate limited for repeatability to 1 Gbps; the public one is for request submission and the private one is for node-to-node communication. Clients submit requests with 500 byte payload, the average Bitcoin transaction size [2]. Each node runs on a single virtual machine. Each node setup is uniformly distributed across all datacenters, except for the 4-node setup which spans 4 datacenters, distributed across all 4 continents. We use 16 client machines, also uniformly distributed across all datacenters, each running 16 clients in parallel which communicate independently with the nodes using TLS. We evaluate throughput, i.e., the number of requests the system delivers per second, and end-to-end latency, i.e., the latency from the moment a client submits a request until the client receives $f + 1$ responses.

6.2 ISS Configuration

After a preliminary evaluation of each protocol, we established a meaningful set of parameters. We do not claim that this set is optimal. Our choice of parameters allows us, however, to demonstrate that ISS makes all three protocols scalable, which is the key contribution of this work. Table 1 summarizes the set of parameters of our evaluation.

	PBFT	HotStuff	Raft
Initial leaderset size	$ N $	$ N $	$ N $
Max batch size	2048	4096	4096
Batch rate	32 b/s	not applicable	32 b/s
Min batch timeout	0 s	1 s	0 s
Max batch timeout	4 s	0	4 s
Min epoch length	256	256	256
Min segment size	2	16	16
Epoch change timeout	10 s	10 s	[10,20) s
Buckets per leader	16	16	16
Client signatures	256-bit ECDSA	256-bit ECDSA	none

Table 1. ISS configuration parameters used in evaluation

For Raft and PBFT we maintain a fixed batch rate. This translates to $O(1/n)$ proposals per leader and $O(n)$ message complexity per bottleneck node³. The choice of a fixed batch rate prevents throughput from dropping due to super-linear message complexity. On the other hand, it introduces higher end-to-end latency as the number of nodes grows, since the batch timeout increases.

The epoch length is kept short: 256 batches per epoch for a batch rate of 32 batches per second yield an epoch duration of approximately 8 seconds in a fault free execution. Shorter epoch length maintains lower latencies in case a fault occurs because bucket re-distribution is executed faster. However, a fixed epoch length yields a shorter segment length as the number of nodes increases. Too short segments for HotStuff and Raft translate to a significant overhead of the dummy/empty batches at the end of the segment. We, therefore, chose a larger minimum segment size for those two protocols.

Batch timeouts should in general be kept small to prevent increasing end-to-end latency. In Raft, however, by design, a leader node re-sends proposals until it has received an acknowledgment from the followers. A very short batch timeout would result in sending proposals too soon and therefore repeating previous proposals. This has a negative impact on throughput because the bandwidth is consumed by unnecessary duplicate proposals. To avoid this phenomenon, we opted for a minimum batch timeout longer than the approximated network round trip. To prevent rate-limiting Raft throughput due to the long batch timeouts, we allowed a large batch size.

HotStuff, on the other hand, is a latency bound protocol as sending a proposal first requires assembling a quorum certificate which depends on the previous proposal. We opted, therefore, for a batch timeout of 0 to allow the leader to send proposals as fast as possible. Similarly to Raft, we allowed a large batch size to prevent rate-limiting the throughput.

6.3 Failure-Free Performance

Figure 5 shows the overall throughput scalability of PBFT, HotStuff, and Raft, with and without ISS, as well as that of Mir-BFT.

We evaluate the scalability of ISS with up to 128 nodes, uniformly distributed across all 16 datacenters. Mir-BFT is evaluated on the same set of datacenters on machines with the same specifications. For an apples-to-apples, comparison, we disabled Mir-BFT optimizations (signature verification sharding and light total order broadcast). Such optimizations could be implemented on top of ISS yielding even better performance. However, this goes beyond the scope of this work. For all protocols we run experiments with increasing the client request submission rate until the throughput is saturated.

In Figure 5 we report the highest measured throughput before saturation. We observe that ISS dramatically improves

³We use the term *bottleneck* node to refer to the node that processes the most messages in each protocol.

the performance of the single leader protocols as the number of nodes grows (37x, 56x and 55x improvement for PBFT, HotStuff, and Raft, respectively, on 128 nodes). This improvement is due to overcoming the single leader bandwidth bottleneck. Moreover, as the number of nodes grows, ISS-PBFT outperforms Mir-BFT. While in theory, in a fault-free execution, we would expect the two protocols to perform the same, we attribute this improvement to the more careful concurrency handling in the ISS implementation.

ISS-PBFT maintains more than 58 kreq/s on 128 nodes. Its performance, though, drops compared to smaller configurations. We attribute this to the increasing number of messages each node processes, which, with a fixed batch rate (Table 1), grows linearly with the number of nodes. We further observe that throughput increases for Raft and HotStuff ISS implementations with the number of nodes, approaching that of ISS-PBFT. While PBFT's watermarking mechanism allows the leader to propose batches in parallel, HotStuff, as explained in Section 6.2, is latency-bound. However, running multiple independent protocol instances with ISS helps improve the overall throughput. Raft, on the other hand, suffers from the redundant re-proposals. While this drawback is mostly hidden in fast LANs with negligible latency, it manifests strongly in a WAN. With more nodes in the leaderset the batch timeout increases and re-proposals are reduced.

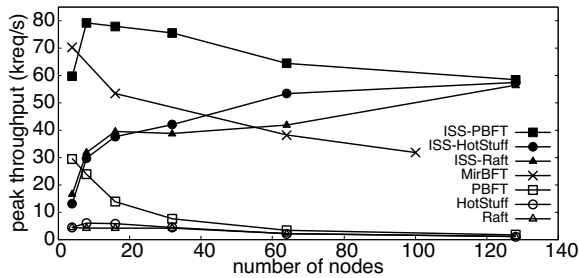


Figure 5. Scalability of single leader protocols, their ISS counterpart, and Mir-BFT.

In Figure 6 we observe that ISS latency grows with the number of nodes. This is due to our choice of a fixed batch rate in order to reduce message complexity and sustain high throughput with an increasing number of nodes.

6.4 ISS Under Faults

In this section we fix PBFT as the protocol multiplexed with ISS and study its performance under crash faults and Byzantine stragglers in a WAN of 32 nodes. The PBFT view change timeout is set to 10 seconds.

6.4.1 Crash Faults. We study two edge cases of faults: (a) one or more leaders crash at the beginning of the first epoch of the execution (*epoch-start*) and (b) one or more leaders crash before sending the proposal for the last sequence number

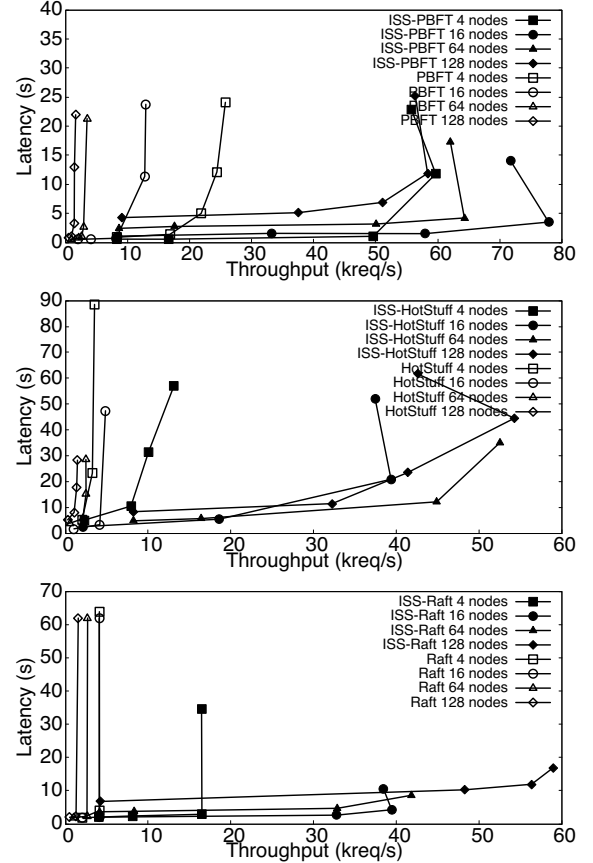


Figure 6. Latency over throughput for increasing load for (a) ISS-PBFT, (b) ISS-HotStuff, and (c) ISS-Raft.

they lead in the first epoch of the execution (*epoch-end*). The *epoch-start* crash fault is a worst-case scenario for the number of proposed sequence numbers in an epoch. The *epoch-end* crash fault is a worst-case scenario for the duration of the epoch; all nodes need to wait for the fault to be detected at the end of the epoch.

Figure 7 shows the impact of crash faults on latency. We see that latency converges towards that of a fault-free execution as we increase the duration of the experiment. This is due to our leader selection policy removing the faulty node from the leaderset once detected. Note that, regardless of their number, epoch-end failures have a stronger impact on latency (as they delay requests in all bucket queues) than epoch-start failures (affecting only the faulty nodes' bucket queues).

Figure 8 shows throughput over time. The short drops to 0 in throughput correspond to the epoch change. We see that an epoch-start fault does not delay the epoch change, as the timer detecting the faulty leader of one segment runs in parallel with other segments. An epoch-end fault, on the other hand, delays the epoch change. However, ISS quickly recovers by ordering more than 170k req/s at the beginning of the second epoch (see the spike in Figure 8(b)).

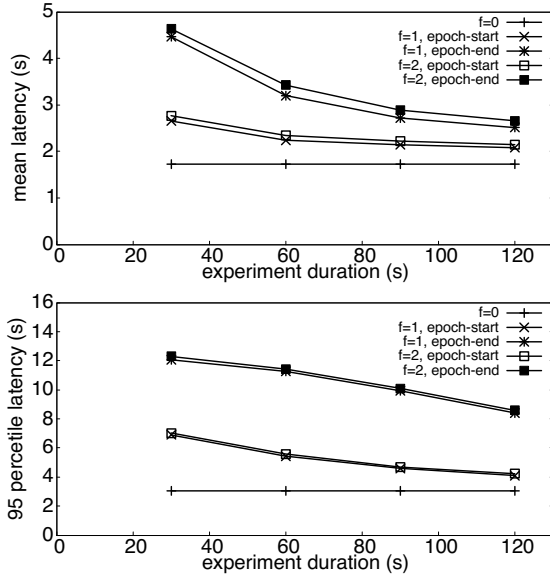


Figure 7. Impact of crash faults on mean (a) and tail (b) end-to-end latency for increasing experiment duration.

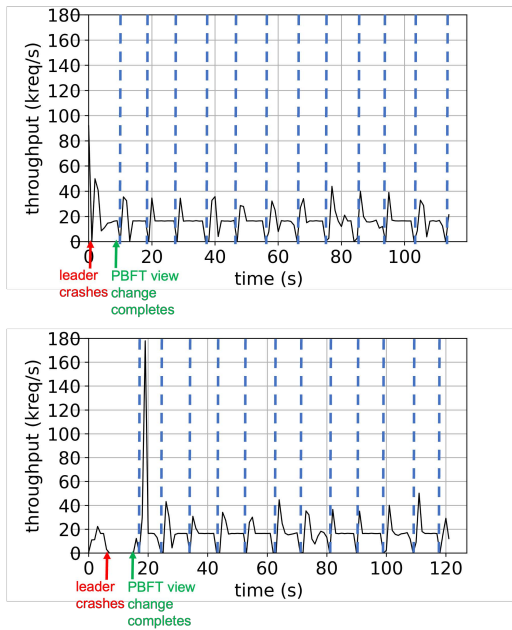


Figure 8. ISS-PBFT throughput average (over 1s intervals) over time with one crash fault (a) at the beginning and (b) at the end of the first epoch. The dashed lines indicate the end of an epoch.

We compare the ISS performance under crash faults to Mir-BFT. In Figure 9 we study a run of Mir-BFT on 32 nodes with a single epoch-start crash fault. Mir-BFT stops processing all proposals during epoch changes, unlike ISS where segments make progress independently. This results in any crash fault

having an impact similar to that of the epoch-end fault for ISS. Moreover, Mir-BFT relies on an epoch primary for liveness. Every time the crashed node becomes epoch primary, it causes an ungraceful epoch change. In Figure 9 this happens around $t = 600$ s. The phenomenon repeats periodically, unlike in ISS, where the faulty leader is permanently removed. Finally, ISS crash fault recovery is more lightweight, since it concerns only the batches of a single segment.

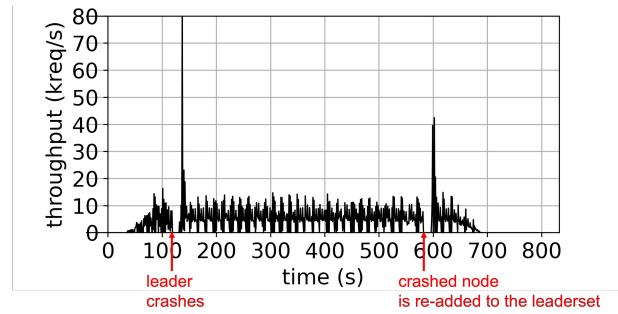


Figure 9. Mir-BFT throughput average (over 1s intervals) over time with one epoch-start fault. Epoch change timeout is at 10s and epoch duration is 256 blocks.

6.4.2 Byzantine Stragglers. A Byzantine straggler delays proposals as much as possible without being suspected as faulty and does not add requests in its proposals to harm latency and throughput. We evaluate latency and throughput with $f = 1$ up to the maximum tolerated number of $f = 10$ stragglers. In our evaluation the straggler sends out an empty proposal every $0.5x$ epoch change timeout.

Figure 10 shows the impact of an increasing number of stragglers. ISS with PBFT reaches from 15% of its maximum throughput with one straggler to 10% of its maximum throughput with 10 stragglers. This, though, translates to maintaining more than 11.4 and 7.9 kreq/s, respectively, on 32 nodes. Mean latency before saturation increases from 14x with one up to 29x with 10 stragglers.

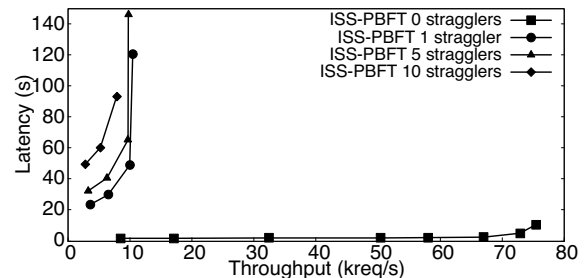


Figure 10. ISS-PBFT latency over throughput for an increasing number of stragglers.

Figure 11 shows how throughput is affected over time with a total submission rate of 16.4kreqs/s. The performance degradation is due to the “holes” in the log temporarily created by the stragglers. Request delivery progresses as fast as the slowest straggler, hence the spikes in the graph. When the straggler’s batch is finally committed, one more batch per leader can be delivered as well (due to the interleaved batch sequence numbers). This is inherent to any SMR protocol [4] until the straggler is removed from the leaderset. Straggler resistance in ISS depends on the underlying SB implementation. A more sophisticated leader selection policy could dynamically detect and remove stragglers from the leaderset. ISS facilitates such dynamic detection by comparing the progress of SB instances, similarly to RBFT [5] but without the need for redundant instances. Alternatively, SB instances could implement Aardvark’s [11] straggler detection mechanism with reducing timeouts. This is promising future work.

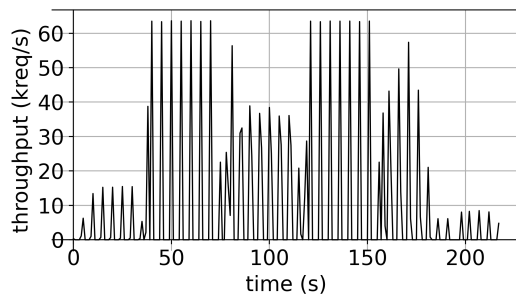


Figure 11. ISS-PBFT throughput average (over 1s intervals) over time with one Byzantine straggler. Each spike (every 5 seconds) corresponds to a group of correct leaders’ batches delivered after the straggler’s batch.

7 Related Work

Consensus under Byzantine faults was first made practical by Castro and Liskov [9][10] who introduced PBFT, a semi-permanent leader-driven protocol. The quadratic message complexity of PBFT across all replicas triggered vigorous research towards protocols with linear message complexity. Ramasamy and Cachin [29] replace reliable broadcast in the common case (fault-free execution) with echo broadcast, achieving common case message complexity $O(n)$ per delivered payload. Echo broadcast is also exploited in [22][18] to achieve linear common case message complexity. Recently, HotStuff [33] introduced a 4th communication round to the 3 message rounds of reliable broadcast, to achieve linear message complexity also for the recovery phase (view-change) of the protocol. Regardless of the improvement of message complexity, all aforementioned protocols have a single leader at a time, permanent or not, limiting throughput scalability.

Mencius[26] introduced multiple parallel leaders, running instances of Paxos [24], to achieve throughput scalability and

low latency in WAN under crash fault assumptions. BFT-Mencius [27] was the first work to introduce parallel leaders under Byzantine faults. BFT-Mencius introduced the Abortable Timely Announced Broadcast communication primitive to guarantee bounded delay after GST. However, BFT-Mencius partitions requests among instances by deterministically assigning clients to replicas, which cannot guarantee load balancing. Moreover, this opens a surface to duplication performance attacks, since malicious clients and replicas can abuse the suggested denial of service mitigation mechanism.

Guerraoui *et al.* [19] also introduced an abstraction which allows BFT instances to abort. The paper uses the abstraction to compose sequentially different BFT protocols, which allows a system to choose the optimal protocol according to network conditions. Our work, on the other hand, composes TOB instances in parallel to achieve throughput scalability.

Mir-BFT [31] is the multi-leader protocol which eliminates request duplication *ante* broadcast, effectively preventing duplication attacks. Later FnF [6] suggested improved leaderset policies for throughput scalability under performance attacks. FnF adopts Mir-BFT’s request space partitioning mechanism for duplication prevention. Similarly, Dandelion[21] leverages the same mechanism to combine Algorand[17] instances. However, Mir-BFT and FnF multiplex PBFT and SBFT [18] instances, respectively, leveraging a single replica in the role of epoch primary. ISS not only eliminates the need for an epoch primary but also provides a modular framework to multiplex any single leader protocol that can implement SB.

Parallel to this work, several works attempt multiplexing BFT instances to achieve high throughput (Redbelly[12], RCC[20], Omada[16]). However, similarly to BFT-Mencius, clients are assigned to primaries, and, after a timeout, a client can change primary to guarantee liveness, again allowing duplication attacks.

8 Conclusion

In this work we introduced ISS, a general construction for efficiently multiplexing instances of leader-based consensus protocols. ISS leverages request space partitioning to prevent duplication, similarly to Mir-BFT, but re-assigns the partitions without the need of a replica to act as a primary, even in the case of faults. To achieve this, we introduced Sequenced Broadcast, a novel abstraction which allows periodically terminating and synchronizing the otherwise independent consensus instances. Our evaluation shows that our careful engineering along with the multi-leader paradigm results in scalable performance for three single-leader protocols (PBFT[10], HotStuff[33], and Raft [28]), outperforming their original designs by an order of magnitude at scale.

References

- [1] Cosmos: A network of distributed ledgers. <https://github.com/dedis/kyber>. Accessed: 30.05.2021.

- [2] Bitcoin visuals: Transaction sizes. <https://bitcoinvisuals.com/chain-tx-size>, 2019.
- [3] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolic, Sharon Weed Cocco, and Jason Yellick. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto, Portugal, April 23-26, 2018*, pages 30:1–30:15, 2018.
- [4] Karolos Antoniadis, Rachid Guerraoui, Dahlia Malkhi, and Dragos-Adrian Seredinschi. State machine replication is more expensive than consensus. Technical report, 2018.
- [5] Pierre-Louis Aublin, Sonia Ben Mokhtar, and Vivien Quéma. RBFT: redundant Byzantine fault tolerance. In *IEEE 33rd International Conference on Distributed Computing Systems, ICDCS 2013, 8-11 July, 2013, Philadelphia, Pennsylvania, USA*, pages 297–306, 2013.
- [6] Zeta Avarikioti, Lioba Heimbach, Roland Schmid, and Roger Wattenhofer. Fnf-bft: Exploring performance limits of BFT protocols. *CoRR*, abs/2009.02235, 2020.
- [7] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *International conference on the theory and application of cryptography and information security*, pages 514–532. Springer, 2001.
- [8] C. Cachin, R. Guerraoui, and L. Rodrigues. *Introduction to reliable and secure distributed programming*. Springer-Verlag New York Inc, 2010.
- [9] M. Castro and B. Liskov. Practical byzantine fault tolerance. *Operating Systems Review*, 33:173–186, 1998.
- [10] Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Trans. Comput. Syst.*, 20(4):398–461, November 2002.
- [11] Allen Clement, Edmund Wong, Lorenzo Alvisi, Mike Dahlin, and Mirco Marchetti. Making Byzantine fault tolerant systems tolerate Byzantine faults. In *NSDI*, 2009.
- [12] Tyler Crain, Vincent Gramoli, Mikel Larrea, and Michel Raynal. Dbft: Efficient leaderless byzantine consensus and its application to blockchains. In *2018 IEEE 17th International Symposium on Network Computing and Applications (NCA)*, pages 1–8. IEEE, 2018.
- [13] Tyler Crain, Christopher Natoli, and Vincent Gramoli. Evaluating the Red Belly blockchain. *CoRR*, abs/1812.11747, 2018.
- [14] Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for Byzantine agreement. *J. ACM*, 1985.
- [15] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, April 1988.
- [16] Michael Eischer and Tobias Distler. Scalable byzantine fault-tolerant state-machine replication on heterogeneous servers. *Computing*, 101(2):97–118, 2019.
- [17] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling Byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 51–68. ACM, 2017.
- [18] Guy Golan-Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael K. Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. SBFT: A scalable and decentralized trust infrastructure. In *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2019, Portland, OR, USA, June 24-27, 2019*, pages 568–580, 2019.
- [19] Rachid Guerraoui, Nikola Knezevic, Vivien Quema, and Marko Vukolic. The Next 700 BFT Protocols. In *Proceedings of the ACM European conference on Computer systems (EuroSys)*, 2010.
- [20] Suyash Gupta, Jelle Hellings, and Mohammad Sadoghi. RCC: resilient concurrent consensus for high-throughput secure transaction processing. In *37th IEEE International Conference on Data Engineering, ICDE 2021, Chania, Greece, April 19-22, 2021*, pages 1392–1403. IEEE, 2021.
- [21] Kadir Korkmaz, Joachim Bruneau-Queyreyx, Sonia Ben Mokhtar, and Laurent Réveillère. Dandelion: multiplexing byzantine agreements to unlock blockchain performance. *arXiv preprint arXiv:2104.15063*, 2021.
- [22] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. Zyzzyva: speculative Byzantine fault tolerance. In *Proceedings of the Symposium on Operating Systems Principles (SOSP)*. ACM, 2007.
- [23] L. Baird. The Swirls Hashgraph consensus algorithm: Fair, fast, Byzantine fault tolerance. <https://www.swirls.com/downloads/SWIRLDS-TR-2016-01.pdf>, 2016.
- [24] Leslie Lamport. Paxos made simple. *ACM SIGACT News*, 32(4):18–25, 2001.
- [25] Dahlia Malkhi and Michael Reiter. Unreliable intrusion detection in distributed computations. In *Proceedings 10th Computer Security Foundations Workshop*, pages 116–124. IEEE, 1997.
- [26] Yanhua Mao, Flavio P. Junqueira, and Keith Marzullo. Mencius: Building efficient replicated state machines for wans. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation, OSDI'08*, pages 369–384, Berkeley, CA, USA, 2008. USENIX Association.
- [27] Zarko Milosevic, Martin Biely, and André Schiper. Bounded delay in byzantine-tolerant state machine replication. In *IEEE 32nd Symposium on Reliable Distributed Systems, SRDS*, 2013.
- [28] Diego Ongaro and John K. Ousterhout. In search of an understandable consensus algorithm. In *Proc. USENIX Annual Technical Conference*, pages 305–319, 2014.
- [29] HariGovind V. Ramasamy and Christian Cachin. Parsimonious asynchronous byzantine-fault-tolerant atomic broadcast. In James H. Anderson, Giuseppe Prencipe, and Roger Wattenhofer, editors, *Principles of Distributed Systems, 9th International Conference, OPODIS 2005, Pisa, Italy, December 12-14, 2005, Revised Selected Papers*, volume 3974 of *Lecture Notes in Computer Science*, pages 88–102. Springer, 2005.
- [30] Michael K. Reiter. A secure group membership protocol. *IEEE Trans. Software Eng.*, 22(1):31–42, 1996.
- [31] Chrysoula Stathakopoulou, Tudor David, Matej Pavlovic, and Marko Vukolić. Mir-bft: High-throughput bft for blockchains. *arXiv preprint arXiv:1906.05552*, 2019.
- [32] Chrysoula Stathakopoulou, Matej Pavlovic, and Marko Vukolić. State-machine replication scalability made simple (extended version). <https://arxiv.org/abs/2203.05681>, 2022.
- [33] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan-Gueta, and Ittai Abraham. Hotstuff: BFT consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC*, 2019.

A Artifact Appendix

In this document we provide a guide on how to reproduce the results of the main paper.

A.1 Abstract

ISS is a modular framework for implementing, deploying and testing a state-machine replication service. The main task of such a service is maintaining a totally ordered log of client requests. This implementation uses multiple instances of an ordering protocol (an SB implementation) and multiplexes their outputs into the final log. The ordering protocol instances running on each peer are orchestrated by a *manager* module

that decides which instance is responsible for which part of the log, when to execute a checkpoint protocol and which client requests are to be ordered by which ordering instance. The decisions of the manager must be consistent across all peers.

The log is a sequence of *entries*. Each entry has a sequence number (sn) defining its position in the log, and contains a batch of requests. The log is logically partitioned into *segments* - parts of the log attributed to a single instance of an ordering protocol. It is the manager's task to create these segments and to instantiate the ordering protocol for each created segment.

The set of all possible client requests is partitioned (based on their hashes) into subsets called *buckets*. The manager assigns a bucket to each segment it creates. The ordering protocol instance only creates batches of requests using the assigned bucket. It is the manager's task to assign buckets in a way ensuring that no two segments that are being ordered concurrently are assigned the same bucket. This is required to prevent request duplication.

The manager observes the log and creates new segments as the log fills up. When the manager creates a new segment, it triggers the orderer that orders the segment. Ordering a segment means committing new entries with the sequence numbers of that segment. Periodically, the manager triggers the *checkpointer* to create checkpoints of the log. The manager observes the created checkpoints and issues new segments as the checkpoints advance, respecting the watermark window.

A.2 Description & Requirements

A.2.1 How to Access. The code used to produce the results of the experiments is publicly available under Hyperledger Labs ⁴⁵. Importantly, the artifact does not reside in the main branch of the repository. It can be found in the `research-iss` branch. For completeness we mention that main branch started as a production implementation of MirBFT which has started shifting towards ISS.

A.2.2 Hardware Dependencies. We performed our evaluation on a WAN on IBM cloud. All processes ran on a dedicated virtual machine with 32 x 2.0 GHz VCPUs and 32GB RAM running Ubuntu Linux 20.04.

Our WAN spanned in 16 datacenters across Europe, America, Australia, and Asia. In detail, machines were deployed in the following locations: San Jose, Osaka, Amsterdam, Sydney, London, Washington D.C., Chennai, Tokyo, Paris, Dallas, Frankfurt, Milan, Mexico City, Toronto, Montreal, Seoul. For each deployment (16, 32, 64, 128 nodes) the virtual machines were uniformly distributed across all data centers. For the smaller deployments (4, 8 nodes) the machines were located on the first 4 and 8 locations of the list respectively.

⁴<https://github.com/hyperledger-labs/mirbft/tree/research-iss>

⁵Persistent ID:10.5281/zenodo.6344032

Finally, all machines were equipped with two network interfaces, public and private, rate limited for repeatability to 1 Gbps.

A.2.3 Software Dependencies. Go 16+, Python 3

A.2.4 Benchmarks. None.

A.3 Set-up

In the `README`⁶ file of the root directory of the project exist high level information about the architecture of the code. Detailed information on how to deploy the code can be found under the `deployment` directory⁷. In a nutshell, the `deployment` provides scripts that automatically deploy a network of nodes and clients on IBM cloud or locally, runs a set of experiments, analyzes the results and produces a result summary.

For an IBM cloud deployment, the user first needs to set up an account with IBM cloud and register an ssh key with this account. The repository has scripts to help the user set up and initialize the `ibmcloud` cli.

To run a set of experiments the user needs only needs to describe the desired set of experiments in a configuration script. Then the user runs the main script `deploy.sh`. After the result summary is generated, the user can find it in a `.csv` file, or process it with a plotting scrip to visualize the results.

A.3.1 Major Claims.

- We observe that ISS dramatically improves the performance of the single leader protocols as the number of nodes grows: 37x, 56x and 55x improvement for PBFT, HotStuff, and Raft respectively on 128 nodes (Figure 5).
- ISS is robust under crash faults: the impact of the crash fault is limited within the epoch it occurs. Latency remains in the order of seconds even for the affected by the crash fault requests (Figures 7 and 8).

A.3.2 Experiments. ISS configuration includes a wide set of parameters. While the most important are grouped in Table 1, in the artifact we provide the configuration scripts that reproduce the main claims of the paper. Each script generates the configuration for a set of experiments. In the rest of this section we describe the set of experiments each configuration file produced and we link them to the main claims of the main paper.

For each of the experiments listed in Table 2⁸ we performed experiments on 4,8,16,32,64, and 128 nodes with 16 clients for increasing client request rate. By fixing the number of

⁶<https://github.com/hyperledger-labs/mirbft/blob/research-iss/README.md>

⁷<https://github.com/hyperledger-labs/mirbft/blob/research-iss/deployment/README.md>

⁸All scripts referenced in Table 2 can be found under `deployment/scripts/experiment-configuration/`

nodes to 32 we produced the Latency-Throughput plots (Figure 6). We further produced latency throughput plots for each number of nodes - not included in the paper. For each latency - throughput plot the data point at which throughput stops increasing and latency starts increasing significantly represents the maximum throughput. By collecting the maximum throughput data points for all node and protocol configurations we produced Figure 5.

Script	Experiments	Figure
<code>generate-pbft.sh</code>	PBFT	Figure 5, Figure 6a
<code>generate-pbft.sh</code>	ISS-PBFT	Figure 5, Figure 6a
<code>generate-hotstuff.sh</code>	HotStuff	Figure 5, Figure 6b
<code>generate-ISS-hotstuff.sh</code>	ISS-HotStuff	Figure 5, Figure 6b
<code>generate-raft.sh</code>	Raft	Figure 5, Figure 6c
<code>generate-ISS-raft.sh</code>	ISS-RAFT	Figure 5, Figure 6c

Table 2. Failure free performance.

The script `generate-crash-faults.sh` is used to generate experiments for ISS-PBFT with [1, 2] failures at the [beginning, end] of the epoch relevant to Figures 7 and 8. Finally, script `generate-straggler-0.5.sh` is used to generate experiments for ISS-PBFT with a straggler relevant to Figures 10 and 11.