

Shoal: Improving DAG-BFT Latency And Robustness

Alexander Spiegelman
Aptos

Balaji Arun
Aptos

Rati Gelashvili
Aptos

Zekun Li
Aptos

Abstract

The Narwhal system is a state-of-the-art Byzantine fault-tolerant scalable architecture that involves constructing a directed acyclic graph (DAG) of messages among a set of validators in a Blockchain network. Bullshark is a zero-overhead consensus protocol on top of the Narwhal’s DAG that can order over 100k transactions per second. Unfortunately, the high throughput of Bullshark comes with a latency price due to the DAG construction, increasing the latency compared to the state-of-the-art leader-based BFT consensus protocols.

We introduce Shoal, a protocol-agnostic framework for enhancing Narwhal-based consensus. By incorporating leader reputation and pipelining support for the first time, Shoal significantly reduces latency. Moreover, the combination of properties of the DAG construction and the leader reputation mechanism enables the elimination of timeouts in all but extremely uncommon scenarios in practice, a property we name “prevalent responsiveness” (it strictly subsumes the established and often desired “optimistic responsiveness” property for BFT protocols).

We integrated Shoal instantiated with Bullshark, the fastest existing Narwhal-based consensus protocol, in an open-source Blockchain project and provide experimental evaluations demonstrating up to 40% latency reduction in the failure-free executions, and up-to 80% reduction in executions with failures against the vanilla Bullshark implementation.

CCS Concepts: • Security and privacy → Distributed systems security.

Keywords: Consensus Protocol, Byzantine Fault Tolerance

ACM Reference Format:

Alexander Spiegelman, Balaji Arun, Rati Gelashvili, and Zekun Li. 2023. Shoal: Improving DAG-BFT Latency And Robustness .

1 Introduction

Byzantine fault tolerant (BFT) systems, including consensus protocols [13, 23, 24, 29] and state machine replication [7, 10, 26, 42, 46], have been a topic of research for over four decades as a means of constructing reliable distributed systems. Recently, the advent of Blockchains has underscored the significance of high performance. While Bitcoin handles approximately 10 transactions per second (TPS), the proof-of-stake committee-based blockchains [38–41, 43, 44] are now engaged in a race to deliver a scalable BFT system with the utmost throughput and minimal latency.

Historically, the prevailing belief has been that reducing communication complexity was the key to unlocking high performance, leading to the pursuit of protocols with linear communication. However, this did not result in drastic enough improvements in the throughput, falling significantly short of the current blockchain network targets. For example, the state-of-the-art Hotstuff [46] protocol in this line of work only achieves a throughput of 3500 TPS [3].

A recent breakthrough, however, stemmed from the realization that data dissemination is the primary bottleneck for leader-based protocols, and it can benefit from parallelization [4, 17, 37, 45]. The Narwhal system [17] separated data dissemination from the core consensus logic and proposed an architecture where all validators simultaneously disseminate data, while the consensus component orders a smaller amount of metadata. A notable advantage of this architecture is that not only it delivers impressive throughput on a single machine, but also naturally supports scaling out each blockchain validator by adding more machines. The Narwhal paper [17] evaluated the system in a geo-replicated environment with 50 validators and reported a throughput of 160,000 TPS with one machine per validator, which further increased to 600,000 TPS with 10 machines per validator.

These numbers are more in line with the ambitions of modern blockchain systems. Consequently, Narwhal has garnered significant traction within the community, resulting in its deployment in Sui [44] and ongoing development in Aptos [39] and Celo [40].

Developing a production-ready reliable distributed system is challenging, and integrating intricate consensus protocols only adds to the difficulty. Narwhal addresses this issue by abstracting away networking from the consensus protocol. It constructs a non-equivocating round-based directed acyclic graph (DAG), a concept initially introduced by Aleph [21]. In this design, each validator contributes one vertex per round, and each vertex links to $n - f$ vertices in the preceding round. Each vertex is disseminated via an efficient reliable broadcast implementation, ensuring that malicious validators cannot distribute different vertices to different validators within the same round. With networking abstraction separated from the details of consensus, the DAG can be constructed without contending with complex mechanisms like view-change or view-synchronization.

During periods of network asynchrony, each validator may observe a slightly different portion of the DAG at any

given time. However, the structure facilitates a simpler ordering mechanism compared to monolithic BFT protocols. In DAG-based consensus protocols, vertices represent proposals, edges represent votes, and the concept of quorum intersection guarantees that validators can consistently order all DAG vertices. This provides efficient consensus because ordering is done via local computation only, without any additional communication cost.

Narwhal-based consensus protocols. As discussed, the idea shared by Narwhal-based consensus protocols is to interpret the DAG structure as the consensus logic [17, 27, 35, 36], but they differ in the networking assumptions and the number of rounds required for vertex ordering. However, all three protocols share a common structure. Prior to the protocol initiation, there is an a-priori mapping from specific rounds to leaders shared among all validators. In the asynchronous protocols (DAG-Rider and Tusk), this mapping to the sequence of leaders is hidden behind threshold cryptography and revealed throughout the protocol. We use the term *anchor* to refer to the vertex associated with the round leader in each relevant round.

The DAG local ordering process by each validator is divided into two phases. First, each validator determines which anchors to order (the rest are skipped). Then, the validators sequentially traverse the ordered anchors, deterministically ordering all DAG vertices contained within the causal histories of the respective anchors. The primary considerations that affect the protocol latency are as follows

1. **Bad leaders.** When a validator is malicious or not fast enough, its vertex may not be included in the DAG. In the case of leaders, the absence of anchors affects the ordering latency of all vertices in previous rounds that are not already ordered. These vertices can only be ordered as a part of a causal history of a future anchor, directly impacting their latency.
2. **Sparse anchors.** In Narwhal-based consensus protocols, not every round includes an anchor. Consequently, vertices located farther from the next anchor must wait for additional rounds before they can be ordered.

Shoal framework. This paper presents Shoal: a framework addressing the aforementioned challenges by incorporating leader reputation and pipelining mechanisms into all Narwhal-based consensus protocols. So far, all available open-source implementations of Narwhal and Bullshark, including Meta ¹, and the production deployment on Sui ² lack these features, while our evaluations demonstrate they can provide significant performance improvements.

Leader reputation is an often overlooked concept in theoretical research, yet it holds crucial importance for practical performance. In practice, Byzantine failures are rare

due to robust protection and economic incentives for validators to adhere to the protocol. (Moreover, Narwhal-based DAG constructions, which provide non-equivocation, significantly reduce the range of potential Byzantine behavior). Thus, the most common failure scenarios in Blockchain (esp. in Narwhal-based) systems involve validators who struggle to keep up, which can occur due to temporary crashes, slower hardware, or geographical distance. If unresponsive validators repeatedly become leaders, progress is inevitably impeded and degrades system performance. The leader reputation schemes select leaders based on the history of their recent activity, as introduced in Diem [42] and later formalized in [16].

In the context of Narwhal-based consensus, pipelining means having an anchor in every round, which would result in improved latency for non-anchor vertices.

The main challenge. While the ability to order the DAG locally, without extra communication contributes to the scalability of Narwhal-based consensus, it poses a significant challenge to supporting leader reputation and pipelining.

The leader reputation problem is simpler to solve for monolithic BFT consensus protocols. While the validators may disagree on the history that determines the next leader’s identity, the worst that can happen is a temporary loss of liveness until view synchronization, i.e. the quorum of validators can eventually recover by agreeing on a fall-back leader. This exact method was utilized in [16], electing the fall-back leaders by a simple round-robin.

In contrast, when all communication is done upfront for building the DAG, the safety of a consensus protocol relies on a key property of the local computation that all validators will decide to order the same set of anchors. This must hold despite the local views of the DAG possibly differing among the validators across multiple rounds. Hence, selecting the round leaders dynamically based on reputation (as opposed to the a-priori mapping) seems impossible due to a circular dependency: we need to agree on mapping to solve consensus, but we need consensus to agree on a new mapping.

For pipelining, even if all validators agree on the mapping, they also must agree on whether to order or skip each anchor. Our attempts to solve the problem by delving into the inner workings of the protocol and exploring complex quorum intersection ordering rules have not been fruitful. Intuitively, this is because consensus requires a voting round after each anchor proposal and the next anchor should link to the decisions (votes) on the previous one.

Our solution. In Shoal, we lean into the power of performing computations on the DAG, in particular the ability to preserve and re-interpret information from previous rounds. For leader reputation, this allows bootstrapping the seemingly circular dependency on consensus, while for pipelining, it allows combining multiple instances of the protocol in a suitable manner. In fact, Shoal runs multiple instances of the

¹<https://github.com/facebookresearch/narwhal/blob/main/consensus/src/lib.rs>

²<https://github.com/MystenLabs/sui/blob/main/narwhal/consensus/src/bullshark>

protocol one after the other, where the trick is to agree on the switching point based on the following observation:

For any Narwhal-based consensus protocol, since all validators agree on which anchors to order vs skip, they in particular agree on the first ordered anchor.

With this observation in mind, each validator can start locally interpreting its view of the DAG by running an instance of its favorite protocol until it determines the first ordered anchor. Since validators agree on this anchor, they can all deterministically start a new protocol instance in the following round. Note that this too, happens locally, from a validator’s perspective, as a part of re-interpreting the DAG. As a result, Shoal ensures the following

1. Leader reputation: validators select new anchors for future rounds based on the information available in the causal history of the ordered anchors.
2. Pipelining: allocate an anchor in the first round of the new instance. That way, if the first anchor in every instance is ordered, we get an anchor in every round, providing the pipelining effect.

Our system and prevalent responsiveness. We implemented Shoal in the open-source codebase of one of the live Blockchain networks and instantiated it with the partially synchronous version of Bullshark³. In this setting, we also discovered a way to eliminate timeouts in all except extremely rare scenarios, a property we refer to as prevalent responsiveness. The design with prevalent responsiveness demonstrates further performance improvements in our evaluations. Added motivation to avoid timeouts in as many situations as possible comes from a purely practical point of view, as (1) when timeouts are common, the duration affects the system performance, but in a way that is non-trivial to configure in an optimal way as it is highly environmentally (network) dependent; and (2) timeout handling is known to add significant complexity to the implementation logic for managing potential state space of validators.

Monolithic leader-based BFT protocols use timeouts to trigger protocol progress every time a leader is faulty or slow, while optimistic responsiveness property, popularized by the HotStuff [46] protocol, effectively eliminates timeout implications in ideal scenarios when the network is synchronous and there are no failures. However, when failures do occur, all validators must still wait until the timeout expires before transitioning to the next leader.

Utilizing the inherent properties of the DAG construction, and leader reputation mechanism, we ensure that Shoal makes progress at network speed under a much larger set of scenarios than optimistically responsive protocols would, which makes Shoal with partially synchronous Bullshark

prevalently responsive. In Shoal, validators do wait for timeouts when a few leaders crash and the corresponding anchors are not ordered. While the FLP [19] impossibility result dictates that there has to be a scenario that requires a timeout, Shoal design aligns this FLP scenario to be extremely improbable in practice (multiple, e.g., 10 consecutive skipped anchors). Conceptually, this is similar to how randomized protocols align FLP scenarios to have 0 probability in solving asynchronous consensus with probability 1 [6].

All available Bullshark implementations use timeouts to ensure honest validators wait for slow anchors even if $2f + 1$ other vertices were already delivered. By eliminating timeouts, Shoal immediately reduces latency when a leader is faulty, as the corresponding anchors would never be delivered and it is best to advance to the next round as fast as possible. If the leader is not crashed and just slower, validators may skip anchors that they could order if they waited a little bit longer. This is however, where the leader reputation mechanism of Shoal shines, filtering out slow validators that constantly delay new rounds and allowing the DAG to proceed at network speed while ordering most anchors.

Our experimental evaluation demonstrates up to 40% reduction in latency against vanilla Bullshark protocol implementation when there are no failures in the system, and up to 80% reduction in latency when there are failures. We provide experiments specifically designed to give insights into the impact of the improvements separately, i.e. pipelining, leader reputation and eliminating the timeouts (prevalent responsiveness).

In summary, the paper focuses on improving latency and robustness in DAG-Based protocols. It provides Shoal, a framework to enhance any Narwhal-based consensus protocol with (1) Leader reputation mechanism that prevents slow, isolated, or crashed validators from becoming leaders, (2) pipelining support that ensures every round on the DAG has an anchor, and (3) eliminating timeouts in many cases further reducing the latency,

The remaining sections of the paper are organized as follows: Section 2 provides background information on DAG-BFT and highlights the main property utilized in this paper. Section 3.1 introduces our pipelining approach, while Section 3.2 presents the leader reputation solution in Shoal. In Section 4, we prove correctness of the proposed framework. Section 5 describes the implementation details and discusses timeouts. Section 6 presents the results of our evaluation. Section 7 discusses related work, and finally, Section 8 concludes the paper.

2 DAG BFT

We start by providing the necessary background on Narwhal-based BFT consensus (Section 2.1) and define a common property (Section 2.2) satisfied by such consensus protocols. We rely on this property while designing Shoal to enhance a

³Shoal of bull sharks.




Figure 1. A possible local view of a round-based DAG. The causal history of the vertex identified by validator 2 in round 2 is highlighted in green.

given baseline protocol with pipelining and leader reputation, thereby reducing latency.

2.1 Background

The concept of DAG-based BFT consensus, initially introduced by HashGraph [5], aims to decouple the network communication layer from the consensus logic. In this approach, each message consists of a collection of transactions and references to previous messages. These messages collectively form an ever-growing DAG, with messaging serving as vertices and references between messages serving as edges.

In Narwhal, the DAG is round-based, similar to Aleph [21]. In this approach, each vertex within the DAG is associated with a round number. In order to progress to round r , a validator must first obtain $n - f$ vertices (from distinct validators) belonging to round $r - 1$. Every validator can broadcast one vertex per round, with each vertex referencing a minimum of $n - f$ vertices from the previous round.

The *causal history* of a vertex v refers to the sub-graph that starts from v . Figure 1 illustrates a validator’s local view of a round-based DAG.

To disseminate messages, Narwhal uses an efficient reliable broadcast implementation that guarantees:

Validity: if an honest validator has a vertex v in its local view of the DAG, then it also has all the causal history of v .

Eventual delivery: if an honest validator has a vertex in round r by validator p in its local view of the DAG, then eventually all honest validators have a vertex in round r by validator p in their local views of the DAG.

Non-equivocation: if two honest validators have a vertex in round r by validator p in their local views of the DAG, then the vertices are identical.

Inductively applying Validity and Non-equivocation, we get:

Completeness: if two honest validators have a vertex v in round r by validator p in their local views of the DAG, then v ’s causal histories are identical in both validators’ local view of the DAG.

In simple words, Narwhal construction guarantees that

1. All validators eventually see the same DAG; and
2. Any two validators that have the same vertex v locally also agree on the whole causal history of v (the contents of vertices and edges between them).

DAG-Rider / Tusk / Bullshark. DAG-Rider, Tusk, and Bullshark are all algorithms to agree on the total order of all vertices in the DAG with no additional communication overhead. Each validator independently looks at its local view of the DAG and orders the vertices without sending a single message. This is done by interpreting the structure of the DAG as a consensus protocol, where a vertex represents a proposal and an edge represents a vote.

DAG-Rider [27] and Tusk [17] are randomized protocols designed to tolerate full asynchrony, which necessitates a larger number of rounds and consequently, a higher latency. Bullshark [36] also provides a deterministic protocol variant with a faster ordering rule, relying on partial synchrony for liveness. While the specific details are not required to understand this paper, next we explain the high-level structure of these protocols and define a property they all share.

2.2 Common framework

Narwhal-based consensus protocols have the following common abstract structure:

1. Pre-determined anchors. Every few rounds (the number depends on the protocol) there is a round with a pre-determined leader. The vertex of the leader is called an *anchor*. In the partially synchronous version of Bullshark, the leaders are a-priori known. In the asynchronous protocols (DAG-Rider, Tusk, asynchronous Bullshark) the leaders are hidden and revealed during the DAG construction.
2. Order the anchors. All validators independently decide which anchors to skip and which to order. The details differ among the protocols, although they all rely on quorum intersection in the DAG structure. The key aspect is that each honest validator locally decides on a list of anchors, and all lists share the same prefix.
3. Order causal histories. Validators process their list of ordered anchors one by one, and for each anchor order all previously unordered vertices in their causal history by some deterministic rule. By Completeness, all validators see the same causal history for any anchor, so all validators agree on the total order.

An illustration of the ordering logic appears in Figure 2.

The key correctness argument for all the above mention consensus protocols relies on the fact that all validators agree on which anchors to order and which to skip. In particular, they will all agree on the first anchor that no validator skips. More formally, the abstract property of the Narwhal-based consensus protocols that our Shoal framework relies on is the following:

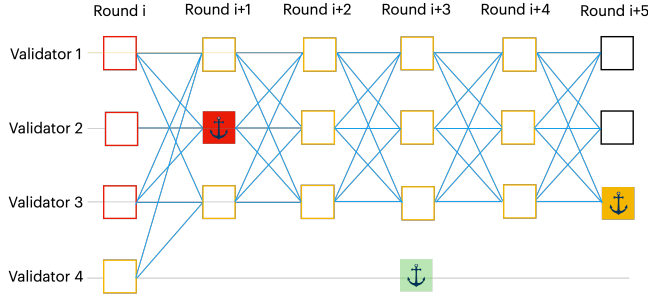


Figure 2. A possible local view of the DAG in the partially synchronous Bullshark protocol. Filled squares represent the pre-defined anchors. In this example, the validator orders the red and yellow anchors, while the green (which is not in the DAG) anchor is skipped. To order the DAG, the validator deterministically orders the red anchor’s causal history (the unfilled red vertices) and immediately after the yellow anchor’s causal history (the unfilled yellow vertices).

Property 1. *Given a Narwhal-based protocol \mathcal{P} , if all honest validators agree on the mapping from rounds to leaders before the beginning of an instance of \mathcal{P} , then they will agree on the first anchor each of them orders during the execution of \mathcal{P} .*

The proof follows immediately from Proposition 2 in DAG-Rider [27] and Corollary C. in Bullshark [35].

3 Shoal

Shoal is protocol agnostic and can be directly applied to all Narwhal-based consensus protocols, i.e., DAG-Rider, Tusk, and Bullshark. It makes no changes to the protocols but rather combines their instances in essentially a “black-box” manner. The entire correctness argument can be derived solely from Property 1.

3.1 Pipelining

A natural progression after the high throughput scalability of BFT consensus achieved by Narwhal is to reduce latency as much as possible. To this end, Bullshark already halved DAG-rider’s latency for ordering anchors from 4 rounds to 2 by adding an optimistic path under the partially synchronous network communication assumption.

Intuitively, it is hard to imagine latency lower than 2 rounds as in the interpretation of the DAG structure as a consensus protocol, one round is needed to “propose” the anchor, while another is needed for “voting”. However, only anchors can be ordered in 2 rounds. The rest of the vertices are ordered as part of the causal history of some anchor and require a minimum latency of 3 or 4 rounds. This is because the vertices in a “voting” round require (minimum) 3 rounds, while vertices that share a round with an anchor have to wait for at least the next anchor to be ordered, thus requiring (minimum) 4 rounds. An illustration of the ordering latency for different vertices appears in Figure 3.

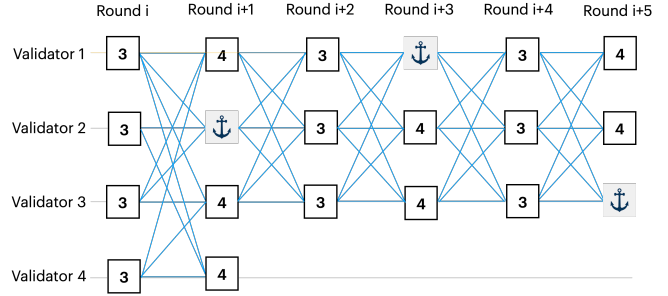


Figure 3. Illustration of the number of rounds required for each vertex in the DAG to be ordered in the best case, according to the Bullshark protocol. The number in each vertex represents its minimum latency. For example, the anchor of round $i + 1$ can be ordered in round $i + 2$, but the other vertices in this round require at least 4 rounds to be ordered.

Ideally, to reduce the latency of ordering vertices we would like to have an anchor in every round. This would allow for non-anchor vertices to be ordered as a part of some anchor’s causal history in each and every round, making latency and throughput of the protocol less spiky. In Bullshark, it would become possible for every non-anchor vertex to be ordered in 3 rounds (see Figure 3), while in DAG-Rider the latency may be reduced from 10 rounds to 7 in expectation.

Solution. Let \mathcal{P} be any Narwhal-based consensus protocol. On a high level, the core technique in Shoal is to execute \mathcal{P} until it, as a consensus protocol, guarantees agreement on some part of the DAG for all validators. Starting from the round following the agreed part of the DAG, all validators can switch over and start executing a new instance of \mathcal{P} (or a different Narwhal-based consensus protocol, if desired) from scratch. While the instances are not executing concurrently, this scheme effectively pipelines the “proposing” and “voting” rounds. As a result in Shoal, in a good case an anchor is ordered in every round.

The pseudocode appears in Algorithm 1. In the beginning of the protocol, all validators interpret the DAG from round 0, and the function F is some pre-defined deterministic mapping from rounds to leaders. Each validator locally runs \mathcal{P} , using F to determine the anchors, until it orders the first anchor, denoted by A in round r . The key is that, by the correctness of \mathcal{P} as stated in Property 1, all validators agree that A is the first ordered anchor (previous anchors are skipped by all validators). Consequently, each validator can re-interpret the DAG from the next round (round $r + 1$) according to a new instance of the protocol \mathcal{P} (or another Narwhal-based protocol) executing from scratch from round $r + 1$.

To order the DAG, much like in the original \mathcal{P} , the validators deterministically order A ’s causal history, and by the Completeness property, arrive at the same total order over the same vertices. Note that without re-interpreting the DAG according to a new instance of \mathcal{P} starting from round $r + 1$,

Algorithm 1 Pipelining in Shoal

```

1:  $current\_round \leftarrow 0$ 
2:  $F : R \rightarrow A$   $\triangleright$  deterministic rounds to anchors mapping
3: while true do
4:   execute  $\mathcal{P}$ , select anchors by  $F$ , starting from
      $current\_round$  until the first ordered (not skipped)
     anchor is determined.
5:   let  $A$  be the first ordered anchor in round  $r$ 
6:   order  $A$ 's causal history according to  $\mathcal{P}$ 
7:    $current\_round \leftarrow r + 1$ 

```

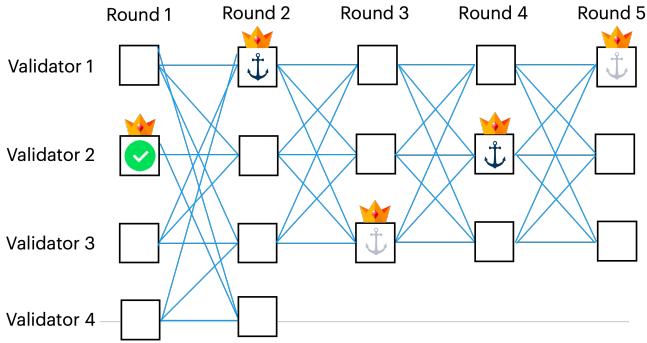


Figure 4. Illustration of Shoal’s pipelining integrated into Bullshark. The vertices that are fixed to be anchors by F are marked by a crown. The protocol starts by interpreting the DAG with anchors in rounds 1, 3, and 5. Bullshark determines that the anchor in round 1, marked by a green checkmark, is the first to be ordered. Then, a new instance of Bullshark starts at round 2 with the anchors marked in rounds 2 and 4.

the next anchor according to the previously executing instance of the protocol would appear in a strictly later round (e.g. $r + 4$ for DagRider and $r + 2$ for Bullshark). The above process can continue for as long as needed. An illustration appears in Figure 4.

Note that in Algorithm 1, function F is fixed and used by each instance of protocol \mathcal{P} . In a true “black-box” implementation, the round numbers could be different from the perspective of the executing protocol instance (i.e. start from 0 for each new instance). However, F is fixed and always assigns the same anchor to any given round r in Shoal regardless of the protocol instance used for this round.

Note that with Shoal, ordering an anchor vertex requires 2 rounds, while all other vertices require 3. In Section A we discuss a potential direction to reduce the latency for non-anchor vertices by treating all vertices as anchors. Intuitively, we can use Property 1 to instantiate a binary agreement to decide whether to commit each vertex individually.

3.2 Leader Reputation

BFT systems are designed to tolerate Byzantine failures in order to provide as strong as possible worst-case reliability guarantees. However, actual Byzantine failures rarely occur in practice. This is because validators are highly secured and

have strong economic incentives to follow the protocol. Slow or crashed leaders are a much more frequent occurrence which can significantly degrade the system performance. In Narwhal-based BFT, if the leader of round r crashes, no validator will have the anchor of round r in its local view of the DAG. Thus, the anchor will be skipped and no vertices in the previous round can be ordered until some later point due to an anchor in a future round.

The way to deal with missing anchors is to somehow ensure that the corresponding leaders are less likely to be elected in the future. A natural approach to this end is to maintain a reputation mechanism, assigning each validator a score based on the history of its recent activity. A validator that has been participating in the protocol and has been responsive would be assigned a high score. Otherwise, the validator is either crashed, slow, or malicious and a low score is assigned. The idea is then to deterministically re-compute the pre-defined mapping from rounds to leaders every time the scores are updated, biasing towards leaders with higher scores. In order for validators to agree on the new mapping, they should agree on the scores, and thus on the history used to derive the scores.

Such a mechanism was previously proposed in [16] and implemented in the Diem Blockchain [42] to enhance the performance of Jolteon [22], a leader-based consensus protocol. One important property Jolteon is that Safety is preserved even if validators disagree on the identity of the leader, while liveness is guaranteed as long as they eventually converge. Hence, validators could re-assign the reputation scores every time a new block was committed, even though during asynchronous periods it was possible for different validators to commit the same block in different rounds. Unfortunately, this is not the case for Narwhal-based BFT. If validators disagree on the anchor vertices, they will order the DAG differently and thus violate safety. This makes the leader reputation problem strictly harder in Narwhal-based BFT.

Solution. Shoal constructs a protocol identical to a given Narwhal-based consensus protocol \mathcal{P} , but to support leader reputation anchors are selected according to a function F that takes into account validators’ recent activity, e.g., the number of vertices they have successfully added to the DAG. The function F should be updated as frequently as possible and aim to select validators with a better reputation as leaders more often than their counterparts with a lower reputation.

In Shoal, pipelining and leader reputation can be naturally combined as they both utilize the same core technique of re-interpreting the DAG after agreeing on the first ordered anchor. In fact, the pseudocode for Shoal appears in Algorithm 2 only differs from Algorithm 1 by adding line 8. The idea is that the validators simply need to compute a new mapping, starting from round $r + 1$, based on the causal history of ordered anchor A in round r (which they are guaranteed to agree on by Property 1). Then, the validators start

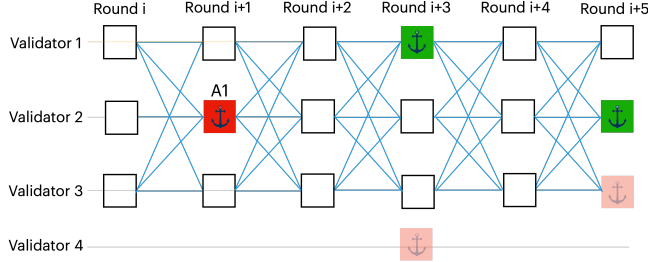


Figure 5. Illustration of Shoal’s leader reputation integrated into Bullshark (no pipelining). First, the DAG is interpreted via the Bullshark protocol and the red anchors. The anchor in round $i+1$, A_1 , is determined to be the first ordered anchor. Then, based on A_1 ’s causal history, new anchors are selected for future rounds (marked in green). Note that validator 4, which had an anchor according to the red selection, no longer has an anchor according to the new mapping (it was not performing well). Then, A_1 ’s causal history is deterministically ordered as in the original Bullshark, and a new instance of Bullshark starts at round $i+2$ based on the green anchors.

executing a new instance of \mathcal{P} from round $r+1$ with the updated anchor selection function F .

Algorithm 2 Shoal

- 1: $current_round \leftarrow 0$
 - 2: $F : R \rightarrow A$ ▶ deterministic rounds to anchors mapping
 - 3: **while true do**
 - 4: execute \mathcal{P} , select anchors by F , starting from $current_round$ until the first ordered (not skipped) anchor is determined.
 - 5: let A be the first ordered anchor in round r
 - 6: order A ’s causal history according to \mathcal{P}
 - 7: $current_round \leftarrow r+1$
 - 8: update F according to A ’s causal history
-

Our solution is protocol agnostic and can be directly applied to all Narwahl-based consensus protocols, i.e., DAG-Rider, Tusk, and Bullshark. An illustration can be found in Figure 5. Shoal makes no changes to the protocols but rather combines their instances, and the entire correctness argument can be derived solely from Property 1.

4 Correctness

To prove the correctness of Shoal (Algorithm 2) we assume that the underlying protocol satisfies Property 1, which we will use inductively.

Lemma 4.1. *Let P be a Narwahl-based DAG-BFT protocol that satisfies Property 1. Let D be a round-based DAG, and assume a known to all function F that maps rounds to anchors. Then all the locally ordered lists of anchors by honest validators executing Shoal with P according to F share the same prefix.*

Proof. Proof is by induction on the ordered anchors.

Base: We need to show that all honest validators agree on the first anchor. Since Shoal starts by running P until the

first anchor is ordered, the base case follows immediately from Property 1.

Step: Assume all honest validators agree on the first k ordered anchors, we need to prove that they agree on anchor $k+1$. First, we show that all honest validators agree on the new function F (Line 8 in Algorithm 2). This holds because the new function F is deterministically computed according to the information in k ’s causal history, and by the Completeness property of the DAG, all honest validators have the same causal history of anchor k in their local view.

Next, let r be the round of anchor k . By the inductive assumption, all honest validators agree on r . Thus, all honest validators start the next instance of P in the same round $r+1$.

Now consider a DAG D' that is identical to D except it does not have the first r rounds. By Property 1, all validators that run P with the new function F on D' agree on the first ordered anchor in D' . Therefore, all validators agree on anchor $k+1$ in D . \square

Theorem 4.2. *Let P be a Narwahl-based DAG-BFT protocol that satisfies Property 1. Shoal with P satisfies total order.*

Proof. By Lemma 4.1, all validators order the same anchors. The theorem follows from the DAG Completeness property as all validators follow the same deterministic rule to order the respective causal histories of the ordered anchors. \square

5 Implementation and Prevalent Responsiveness

We have implemented Narwahl and the partially synchronous version of Bullshark as part of a publicly available open-source blockchain project⁴. This blockchain is live and the process of productionizing our implementation is underway. The code is written in Rust, utilizing Tokio⁵ for asynchronous networking, BLS [8] implemented over BLS12-381 curves for signatures, RocksDB⁶ for persistent data storage, and the Noise⁷ protocol for authenticated messages.

5.1 Vanilla Bullshark

We implemented Bullshark according to [35], but additionally incorporated weak links per [27] in our DAG construction. Observing $n-f$ vertices in a round is sufficient for progressing to the next round. Therefore, without weak links, slow validators may consistently lag behind others in broadcasting their vertices and thus may consistently fail to add their vertices to the DAG. This will incur significant latency for their client transactions. Weak links from a vertex can reference vertices from earlier rounds in addition to the normal (strong) links to $n-f$ vertices from the previous round.

⁴In order to uphold the anonymity requirement of the submission, we do not disclose the name of the blockchain project.

⁵<https://tokio.rs>

⁶<https://rocksdb.org>

⁷https://github.com/noiseprotocol/noise_spec

These weak links are used when establishing the causal history of ordered anchors and thus facilitate the inclusion of transactions contributed by the slow validators into the total order.

We refer to this implementation as Vanilla Bullshark. It is important to note that adding the support for weak links increases the average latency compared to the figures presented in [35], which did not employ the weak links.

5.2 Eliminating Timeouts

The short paper for the stand-alone partially synchronous version of Bullshark [36] assumes the DAG is given and focuses on the ordering of its vertices. On the other hand, full Bullshark is an asynchronous protocol with a fast path under partial synchrony. The full Bullshark paper [35] describes how to build the DAG and in particular, the incorporation of timeouts to support the fast path.

Validators in Bullshark must observe $n - f$ vertices in a round to advance to the next round. Even rounds have anchors, while vertices in odd rounds determine the “voting” pattern. Full Bullshark uses the following timeouts for every validator to support the fast path:

- Even-round: wait until the anchor of the round is delivered (or the timeout expires).
- Odd-round: wait until $2f + 1$ vertices that link to the anchor in the previous round are delivered (or the timeout expires).

The rationale for the above logic is to help order the anchor within 2 rounds. However, part of the contribution of this paper is to eliminate these timeouts in such a way that actually significantly improves latency, according to our evaluation. Having fewer cases where timeouts can occur also inherently simplifies the potential state space and thus, the implementation of the protocol. In Section 6, we refer to even-rounds as *anchor rounds* and to odd-rounds as *vote rounds*.

Vanilla bullshark w/o vote Timeout. In the full Bullshark $2f + 1$ votes are required to order anchors. Without timeouts in odd rounds, a Byzantine adversary can prevent the fast pass from making progress even during synchrony. As long as Byzantine validators deliberately not link to the anchor, and even 1 of their vertices get delivered among the first $2f + 1$ to an honest validator in an odd round, then the honest validator will not be able to order the anchor.

However, we discovered that we can completely eliminate timeouts in odd rounds in the partially synchronous variant of Bullshark. The anchor ordering rule in this case is $f + 1$ votes [36]. As a result, even if f out of the first $2f + 1$ vertices delivered to a validator in a round is from Byzantine validators (and do not link to the anchor), the remaining $f + 1$ vertices will link to the anchor due to the even-round timeout and be sufficient to order it.

Baseline Bullshark. FLP impossibility result [19] dictates that any deterministic protocol providing liveness under partial synchrony must use timeouts. In Bullshark, without timeouts in the even rounds, an honest leader that is even slightly slower than the fastest $2f + 1$ validators will struggle to get its anchor linked by other vertices. As a result, the anchor is unlikely to be ordered. The timeout, therefore, ensures that all honest validators link to anchors during periods of synchrony (as long as the leader has not crashed and actually broadcasts the anchor vertex).

Even though timeouts are unavoidable in the worst case, we observe that the DAG construction combined with the leader reputation mechanism allows avoiding them in vast majority of cases in practice. This is in contrast to leader-based monolithic consensus protocols, where timeouts are the only tool to bypass the rounds with bad leaders. Without timeouts, a monolithic protocol could stall forever as there is no other mechanism to stop waiting for a crashed leader. It is also hard to set the timeouts appropriately: conservative timeouts lead to excessive waiting for crashed leaders, while aggressive timeouts lead to bypassing slower validators (and hence unnecessarily failed rounds).

In contrast, the DAG construction provides a “clock” that estimates the network speed. Even without timeouts, the rounds keeps advancing as long as $2f + 1$ honest validators continue to add their vertices to the DAG. As a result, the DAG can evolve despite some leaders being faulty. Eventually, when a non-faulty leader is fast enough to broadcast the anchor, the ordering will also make progress.

Recall that to be ordered, in partially synchronous Bullshark, an anchor needs $f + 1$ votes (links) out of the $3f + 1$ vertices. Therefore, as our evaluation demonstrates, in the failure-free case, most of the anchors are ordered in the next round. The benefit are even more pronounced when there are failures. This is because a crashed validator causes a timeout to expire, stalling the protocol for the entire duration. Without a timer, however, the DAG will advance rounds at network speed and the Bullshark protocol is able to immediately move to the next anchor.

Timeouts as a fallback. By FLP [19] impossibility result, there exists an adversarial schedule of events that can prevent all anchors from getting enough votes to be ordered. This scenario is extremely unlikely to occur in practice, but to be on the safe side, the protocol can deal with it by falling back to using timeouts after a certain amount of consecutive skipped anchors.

5.3 Shoal of Bullsharks

A realistic case in which timeouts can help the performance of a Narwhal-based consensus protocol is when the leader is slower than other validators. Then, as discussed earlier, waiting for an anchor to be delivered even after $2f + 1$ other vertices can allow the anchor to be committed in the next

round. While we eliminated timeouts from partially synchronous Bullshark, note that, due to the leader reputation mechanism, Shoal instantiated with Bullshark does better than repeatedly waiting for the slow leaders. Instead, the leader reputation mechanism excludes (or at least significantly reduces the chances of) slow validators from being selected as leaders. This way, the system takes advantage of the fast validators to operate at network speed.

Prevalent Responsiveness. Shoal provides network speed responsiveness under all realistic failure and network scenarios, a property we name *Prevalent Responsiveness*. Specifically, compared to optimistic responsiveness, Shoal continues to operate at network speed even during asynchronous periods or if leaders fail for a configurable number of consecutive rounds.

We implemented leader reputation and pipelining on top of the Baseline Bullshark and compared it to the baseline (no timeouts) implementation.

Leader reputation logic. As explained in Section 3.2, Shoal ensures all validators agree on the information used to evaluate the recent activity and to bias the leader selection process accordingly towards healthier validators. Any deterministic rule to determine the mapping from rounds to leaders (i.e. the logic in pseudocode Line 8 in Algorithm 2) based on this shared and agreed upon information would satisfy the correctness requirements. Next, we discuss the specific logic used in our implementation.

At any time each validator is assigned either a high or a low score, and all validators start with a high score. After ordering an anchor v , each validator examines v 's causal history H . Every skipped anchor in H is (re-)assigned a low score, and every ordered anchor in H is (re-)assigned a high score. Then, the new sequence of anchors is pseudo-randomly chosen based on the scores, with a validator with a high score more likely to be a leader in any given round. Note that while the validators use the same pseudo-randomness (so that they agree on the anchors), the computation is performed locally without extra communication.

Assigning higher scores to validators whose anchors get ordered ensures that future anchors correspond to faster validators, thus increasing their probability to be ordered. However, we ensure that the low score is non-zero, and thus underperforming validators also get a chance to be leaders. This crucially gives a temporarily crashed or underperforming validator a chance to recover its reputation.

6 Evaluation

We evaluated the performance of the aforementioned variants of Bullshark and Shoal on a geo-replicated environment in Google Cloud. In order to show the improvements from pipelining and leader reputation independently, we also evaluate Shoal PL, which is a Shoal instantiation with

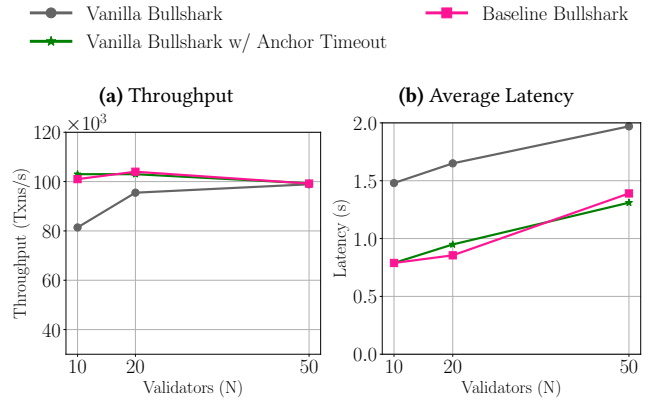


Figure 6. Baseline performance under no failures

only pipelining enabled, and Shoal LR, which is a Shoal instantiation with only Leader Reputation enabled. With our evaluation, we aim to show that (i) Shoal maintains the same throughput guarantees as Bullshark. (ii) Shoal can provide significantly lower latency than Bullshark and its variants. (iii) Shoal is more robust to failures and can improve latency with the help of Leader Reputation.

For completeness, we also compare against Jolteon [22], which is the current consensus protocol of the production system we use. Jolteon combines the linear fast path of Tendermint/Hotstuff with a PBFT style view-change, and as a result, reduces Hotstuff latency by 33%. The implementation extends the original Jolteon protocol with a leader reputation mechanism, which prioritizes well-behaved leaders from previous rounds for future rounds. In addition, to mitigate the leader bottleneck and support high throughput, the implementation uses the Narwhal technique to decouple data dissemination via a pre-step component (called Quorum Store [15]).

We evaluate prevalent responsiveness by presenting experiments that compare variants of Bullshark w.o. timeout in different rounds versus Shoal as discussed in Section 5.

Experimental Setup. Our experimental setup consists of t2d-standard-32 type virtual machines spread equally across three different Google Cloud regions: us-west1, europe-west4, asia-east1. Each virtual machine has 32 vCPUs, 128GB of memory, and can provide up to 10Gbps of network bandwidth. The round-trip latencies are: 118ms between us-west1 and asia-east1, 251ms between europe-west4 and asia-east1, and 133ms between us-west1 and europe-west4. The experiments involve three different values of N (the number of validators): 10, 20, and 50, tolerating up to 3, 6, and 16 failures, respectively.

We only measure the consensus performance to avoid introducing noise from other parts of the production system, such as execution and storage. The transactions are approximately 270B in size. We set a maximum batch size of 5000 transactions. In our experiments, we measure *latency* as the




Figure 7. Baseline performance under failures (N=50)

time elapsed from when a vertex is created from a batch of client transactions to when it is ordered by a validator. The timeouts for moving to the next round, when applicable, are set to 1s, which is less than the 1.5s timeout used by the production Blockchain system we use.

6.1 Baseline Performance

First, we evaluate the performance of the Bullshark variants, namely Vanilla Bullshark, Vanilla Bullshark w/ Anchor Timeouts, and Baseline Bullshark, to align on a baseline performance to evaluate Shoal in the rest of the experiments. The results are in Figures 6 and 7.

Figure 6 shows the throughput and average latencies of the three Bullshark variants as the system size increases. The presence of timeouts in Vanilla Bullshark forces it to build the DAG slowly, which combined with the fact that fewer validators contribute vertices to the DAG when $N = 10$, results in lower throughput than other variants, which have fewer or no timeouts. The latencies for Vanilla Bullshark is up to 88% higher due to the timeouts. Interestingly, the latencies are similar for baseline Bullshark and Vanilla Bullshark w/o Vote timeout in the normal case because there is a trade-off between building a DAG at network-speed while skipping an anchor and waiting slightly longer for the anchor to be part of the votes.

We also evaluated the vanilla variants and the baseline for $N = 50$ and with varying the number of failures, in Figure 7. We observe that Baseline Bullshark provides lower latency than other variants by virtue of being able to build the DAG at network speed skipping failed anchors and ordering using the alive ones. Therefore, in the rest of the section, we use Baseline Bullshark as the baseline to evaluate Shoal.

6.2 Performance of Shoal under fault-free case

We now evaluate the Shoal variants against the baseline under the normal case where there are no failures. The results are in Figure 8. As expected, the throughput of the Shoal variants is similar as the number of validators increases. It can be observed that each variant of Shoal decreases the latency leading to full Shoal protocol. In summary, we observe that

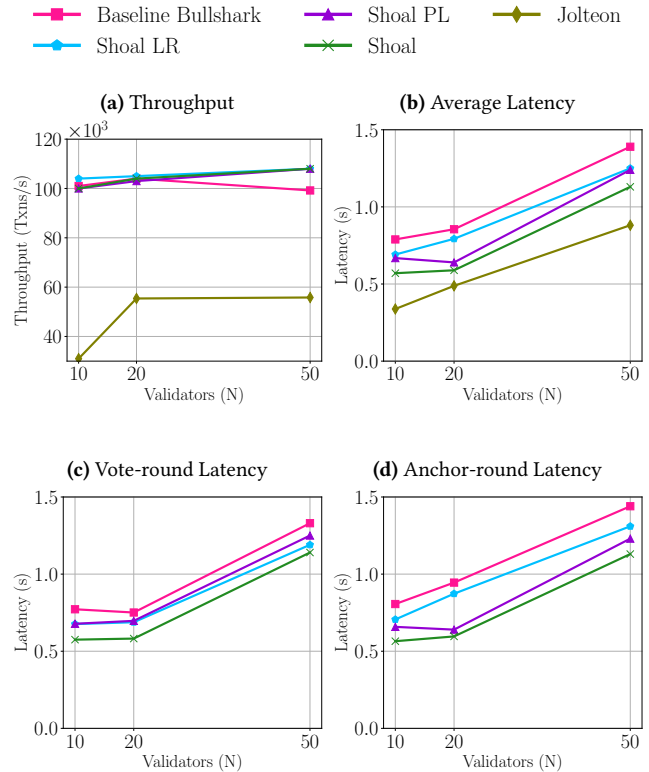


Figure 8. Shoal performance under no failures with 10, 20 and 50 validators.

the Shoal’s average latency decreases by up to 20% compared to Baseline Bullshark.

On the other hand, Jolteon [22], despite its use Narwhal’s data dissemination decoupling, is only able to achieve a peak throughput of less than 60k, about 40% lower than Shoal. This is because under high load leaders become the bottleneck again as they are not able to deal with the required network bandwidth, and as a result, unable to drive progress before timeouts expire. Furthermore, in terms of latency, Jolteon is $\approx 50\%$ better than Vanilla Bullshark, but only $\approx 20\%$ better than Shoal. Note that the latencies presented do not include the pre-step Quorum Store’s latencies, because all the compared protocols include this optimization. However, in the case of Shoal, this latency can be avoided by merging Quorum Store into the DAG construction, as done in Narwhal, which will further close the latency gap from Jolteon.

In Figures 9c and 9d, we distinguish the latencies of transactions in the vote-round vertices from that in anchor-round vertices, in order to show the effect of the pipelining approach. The vote and anchor round latencies for Shoal PL, as well as Shoal, are similar, which helps provide predictable and smooth latency for transactions in real production systems. In contrast, the vote and anchor round latencies for Baseline Bullshark and Shoal LR differ by 5-20% depending on the number of failures.

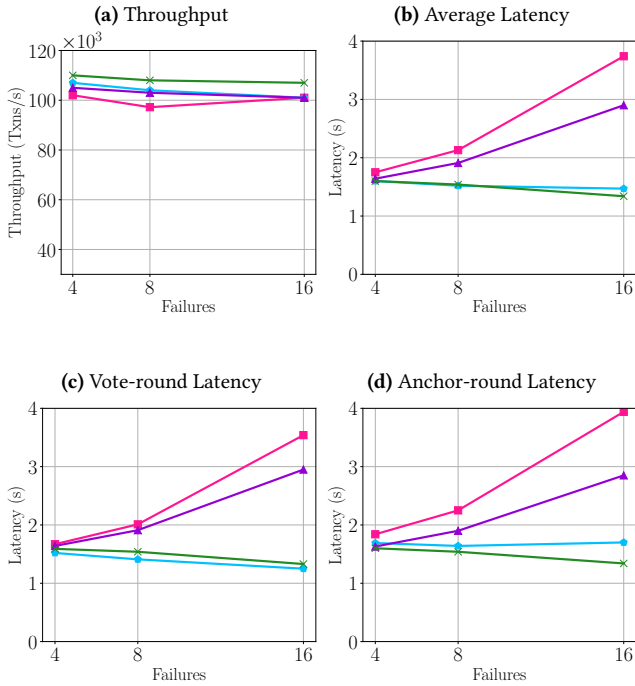


Figure 9. Shoal performance under 4, 8, and 16 failures (N=50)

6.3 Performance of Shoal under faults

Figure 9 shows the behavior of baseline and Shoal variants under faults. For this experiment, $N = 50$ and the failures are increased from 4 to 16 (maximum tolerated). This is the case where the Leader Reputation mechanism helps to improve the latency significantly by reducing the likelihood of failed validators from being anchors. Notice that without Leader Reputation, the latencies of Baseline Bullshark and Shoal PL increases significantly as the number of failures increases. Shoal provides up to 65% lower latencies than Baseline Bullshark under failures.

Figure 10 shows the impact of skipping leaders on the latency by comparing vanilla Bullshark with Shoal on a timeline plot under failures. We have a system of 50 validators, 8 of which have failed. The x-axis represents a part of the experiment time window and the y-axis shows the latency. The presence of timeouts and the need to skip anchors causes vanilla Bullshark’s latency to fluctuate. In our experiment, we observed latency jitter of approximately one second, which makes it impossible to provide predictable latency in production systems. In contrast, Shoal maintains consistent low latency without any jitter.

6.4 Summary

In contrast to Vanilla Bullshark, Shoal provides up to 40% lower latency in the fault-free case and up to 80% lower latency under failures. Furthermore, we show that Shoal




Figure 10. Latency timeline under 8 failures with $N = 50$. The x-axis represents a part of the experiment time window and the y-axis shows the latency.

provides predictable latency and is able to commit at network speed in most cases and without waiting for timeouts.

7 Related work

7.1 BFT systems for Blockchains

Byzantine fault tolerance (BFT) has been an active area of research for over four decades, with a significant body of literature in both theory [11] and systems [1, 7, 13, 29]. With the advent of Blockchain systems in recent years, the focus on performance and scalability has notably increased.

Initial efforts to enhance throughput and scalability attempted to reduce the communication complexity of leader-based eventually synchronous protocols. This resulted in a considerable body of work aiming to achieve communication complexity linear to the number of validators [9, 33, 34, 46]. Despite sound theoretical premise, the practical implications arguably fell short of expectations. An independent evaluation and comparison conducted by [3] revealed that the well-known HotStuff [46] protocol achieved a throughput of only 3,500 TPS on a geo-replicated network.

The practical breakthrough occurred a few years later with the realization that the main bottleneck in BFT systems, particularly those relying on leaders, is data dissemination. Mir-BFT [37] introduced an innovative approach by running multiple PBFT [13] instances in parallel. Independently, Narwhal [17] and later Dispersedledger [45] decoupled data dissemination from the consensus logic. These advancements showcased impressive results, with Narwhal achieving a peak throughput of 160,000 TPS.

There has been systems [18, 22, 25, 32] and theoretical [2, 12, 31] research in asynchronous BFT protocols. However, to the best of our knowledge, no asynchronous protocol is deployed in production in an industrial system. Another appealing property of Narwhal is the support of a partially synchronous [36] as well as asynchronous [17, 27, 35] (as long as randomness is available) protocols, and the ability to easily switch among them.

7.2 Timeouts and responsiveness

The FLP [19] impossibility result states that there is no deterministic consensus protocol that can tolerate a fully asynchronous network. The proof relies on the fact that it is impossible to distinguish between crashed and slow validators during asynchronous periods. The immediate application to partially synchronous networks, therefore, is that all deterministic protocols must rely on timeouts in some way to guarantee liveness against a worst-case adversary. Indeed, to the best of our knowledge, all previous deterministic BFT protocols, including the partially synchronous version of Bullshark [36], relied on timeouts to implement a simple version of a failure detector [14]. This mechanism monitors the leaders and triggers view-changes when timeouts expire, i.e. when faults are suspected.

The optimistic responsiveness property, popularized by HotStuff [46], avoids timeouts in the best-case failure-free scenario. However, when failures do occur, all validators wait until the timeout expires before view-changing to the next leader, introducing a significant slowdown in the protocol execution. Moreover, as discussed in Section 5, setting a proper timeout duration is a non-trivial problem in its own right.

Shoal provides prevalent responsiveness, which is a strictly better property than optimistic responsiveness as it guarantees network speed progress in case of healthy leaders and zero delays in case of failures. Shoal achieves this by relying on the network speed “clock” inherent in the DAG construction itself [20], combined with the leader reputation mechanism. While due to the FLP result, the worst case in which a timeout would be required for maintaining the liveness of the protocol cannot completely be eliminated, Shoal successfully relegates such cases to occur in specific extremely uncommon scenarios from a practical point of view (multiple consecutive unordered anchors).

7.3 DAG-based BFT

DAG-based consensus in the context of BFT was first proposed by HashGraph [5]. The idea is to separate the network communication layer, i.e. efficiently constructing a system that forms a DAG of messages, and the consensus logic that can involve complex pieces such as view-change and view-synchronization. The consensus logic is performed locally, whereby a validator examines its local view of the DAG and orders the vertices without sending any messages. The challenge arises from the asynchronous nature of the network, which may cause different validators to observe slightly different portions of the DAG. To address this, the DAG structure is interpreted as a consensus protocol, wherein a vertex represents a proposal and an edge represents a vote.

Aleph [21] introduced a round-based DAG structure. Such a structure simplifies support for garbage collection and non-equivocation, which in turn simplifies the consensus

logic to order the vertices. Narwhal implements round-based DAG, and three Narwhal-based consensus protocols have been previously proposed. The first is DAG-Rider [27], which introduced a quantum-safe asynchronous protocol with optimal amortized communication complexity and $O(1)$ latency. Tusk [17] improved latency in the best case. An asynchronous version of Bullshark [35, 36] includes a fast path [35], while a stand-alone partially synchronous protocol [36] also exists and is currently deployed in production in Sui [44]. Shoal presents a framework that applies to all Narwhal-based protocols, enhancing their latency through a more efficient ordering rule and a leader reputation mechanism.

An orthogonal theoretical effort [28] trades off the non-equivocation property of the DAG construction (which typically requires reliable broadcast), as well as the separation from the consensus logic, in order to reduce latency.

7.4 Pipelining

To the best of our knowledge, pipelining in the BFT context was first proposed by Tendermint [10], and later utilized in HotStuff [46] and Diem [42]. State machine replication (SMR) systems can be constructed from multiple instances of single-shot consensus [30], e.g. one approach to build Byzantine SMR is by running a PBFT instance [13] for each slot. Tendermint introduced the elegant idea of chaining proposals or piggybacking single-shot instances such that a value for a new slot could be proposed before the value for the previous slot was committed. In this approach, a message in the i^{th} round of the k^{th} instance can be interpreted as a message in round $i - 1$ of instance $k + 1$. While the latency for each instance remains unchanged, clients experience improved latency as their transactions can be proposed earlier.

In DAG-based consensus, the concept of piggybacking proposals is inherent in the design, as each vertex in the DAG links to vertices in previous rounds. However, previous protocols did not allow having an anchor in every round. Shoal framework supports having an anchor in each round in a good case for any Narwhal-based protocol, providing a “pipelining effect”.

7.5 Leader reputation

Leader reputation is often overlooked in theory, yet it plays a crucial role in performance in practice. While Byzantine failures are rare as validators are highly protected, isolated, and economically incentivized to follow the protocol, more common are validators that are unresponsive. This may be because they temporarily crashed, running slow hardware, or are simply located farther away. If a leader/anchor election is done naively, unresponsive validators will unavoidably stall progress and lead to significant performance impact.

A practical approach, implemented in Diem [42] and formalized in [16], is to exclude underperforming validators from leader election. This is achieved by updating the set of candidates after every committed block based on the recent

activity of validators. In a chained protocol, if all validators observe the same committed block, they can deterministically elect future leaders based on the information in the chain. However, in some cases, certain validators may see a commit certificate for a block earlier than others. This can lead to disagreements among validators regarding the list of next leaders, causing a temporary loss of liveness.

For DAG-based protocols, disagreements on the identity of round leaders can lead the validators to order the DAG completely differently. This poses a challenge for implementing leader reputation on the DAG. As evidence, a Narwhal and Bullshark implementation currently deployed in production in Sui blockchain does not support such a feature⁸. Shoal enables leader reputation in Narwhal-based BFT protocols without any additional overhead.

8 Discussion

Shoal can be instantiated with any Narwhal-based consensus protocol, and can even switch between protocols during the DAG retrospective re-interpretation step.

Shoal uniformizes the latency and throughput across the validators and eliminates the use of timeouts except in very rare cases, which contributes to the robustness and performance of the system. Predictable and smooth latency and throughput patterns have major practical benefits for real systems. It facilitates setting up effective monitoring and alerts for anomaly detection. This is crucial for ensuring security and quality of service by enabling timely response and any intervention necessary, be it manual or automated. Predictable consensus throughput also facilitates pipelining the ordering of transactions with other components of the Blockchain, e.g. transaction execution and commit.

Shoal satisfies the property we name prevalent responsiveness, ensuring the worst-case executions that must use timeouts due to the FLP impossibility result are aligned with the improbable (and worst-case) scenarios from the practical standpoint. Moreover, the design without timeouts plays into the strengths of the leader reputation mechanism of Shoal, and as a result, provides further latency improvements.

References

- [1] Michael Abd-El-Malek, Gregory R. Ganger, Garth R. Goodson, Michael K. Reiter, and Jay J. Wylie. 2005. Fault-scalable Byzantine fault-tolerant services. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles 2005, SOSP 2005, Brighton, UK, October 23-26, 2005*, Andrew Herbert and Kenneth P. Birman (Eds.). ACM, 59–74.
- [2] Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. 2019. Asymptotically Optimal Validated Asynchronous Byzantine Agreement. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 337–346.
- [3] Salem Alqahtani and Murat Demirbas. 2021. Bottlenecks in Blockchain Consensus Protocols. *CoRR* abs/2103.04234 (2021).
- [4] Balaji Arun and Binoy Ravindran. 2022. Scalable Byzantine Fault Tolerance via Partial Decentralization. *Proc. VLDB Endow.* 15, 9 (may 2022), 1739–1752. <https://doi.org/10.14778/3538598.3538599>
- [5] Leemon Baird. 2016. The swirls hashgraph consensus algorithm: Fair, fast, byzantine fault tolerance. *Swirls Tech Reports SWIRLDS-TR-2016-01, Tech. Rep* (2016).
- [6] Michael Ben-Or. 1983. Another advantage of free choice (extended abstract) completely asynchronous agreement protocols. In *Proceedings of the second annual ACM symposium on Principles of distributed computing*. 27–30.
- [7] Alysse Bessani, Joao Sousa, and Eduardo EP Alchieri. 2014. State machine replication for the masses with BFT-SMART. In *2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. IEEE, 355–362.
- [8] Dan Boneh, Benjamin Lynn, and Hovav Shacham. 2001. Short signatures from the Weil pairing. In *Advances in Cryptology - CRYPTO 2001 (Lecture Notes in Computer Science, Vol. 2139)*. Springer, 514–532. https://doi.org/10.1007/3-540-44647-8_32
- [9] Manuel Bravo, Gregory Chockler, and Alexey Gotsman. 2022. Making byzantine consensus live. *Distributed Computing* 35, 6 (2022), 503–532.
- [10] Ethan Buchman, Jae Kwon, and Zarko Milosevic. 2018. The latest gossip on BFT consensus. *arXiv preprint arXiv:1807.04938* (2018).
- [11] Christian Cachin, Rachid Guerraoui, and Luis Rodrigues. 2011. *Introduction to reliable and secure distributed programming*. Springer Science & Business Media.
- [12] Christian Cachin, Klaus Kursawe, and Victor Shoup. 2005. Random oracles in Constantinople: Practical asynchronous Byzantine agreement using cryptography. *Journal of Cryptology* 18, 3 (2005), 219–246.
- [13] Miguel Castro and Barbara Liskov. 2002. Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems (TOCS)* 20, 4 (2002), 398–461.
- [14] Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. 1996. The weakest failure detector for solving consensus. *Journal of the ACM (JACM)* 43, 4 (1996), 685–722.
- [15] Brian Cho and Alexander Spiegelman. 2023. Quorum Store: How Consensus Horizontally Scales on the Aptos Blockchain. <https://medium.com/aptoslabs/quorum-store-how-consensus-horizontally-scales-on-the-aptos-blockchain-988866fd5b0> Aptos Labs.
- [16] Shir Cohen, Rati Gelashvili, Lefteris Kokoris Kogias, Zekun Li, Dahlia Malkhi, Alberto Sonnino, and Alexander Spiegelman. 2022. Be aware of your leaders. In *Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, May 2–6, 2022, Revised Selected Papers*. Springer, 279–295.
- [17] George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. 2022. Narwhal and tusk: a dag-based mempool and efficient bft consensus. In *Proceedings of the Seventeenth European Conference on Computer Systems*. 34–50.
- [18] Sisi Duan, Michael K. Reiter, and Haibin Zhang. 2018. BEAT: Asynchronous BFT Made Practical. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM, 2028–2041.
- [19] Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. 1985. Impossibility of Distributed Consensus with One Faulty Process. *J. ACM* 32, 2 (1985), 374–382.
- [20] Bryan Ford. 2019. Threshold logical clocks for asynchronous distributed coordination and consensus. *arXiv preprint arXiv:1907.07010* (2019).
- [21] Adam Gągol, Damian Leśniak, Damian Straszak, and Michał Świątek. 2019. Aleph: Efficient Atomic Broadcast in Asynchronous Networks with Byzantine Nodes. In *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. 214–228.

⁸github.com/MystenLabs/sui/blob/main/narwhal/consensus/src/bullshark.rs

- [22] Rati Gelashvili, Lefteris Kokoris-Kogias, Alberto Sonnino, Alexander Spiegelman, and Zhuolun Xiang. 2022. Jolteon and ditto: Network-adaptive efficient consensus with asynchronous fallback. In *Financial Cryptography and Data Security: 26th International Conference, FC 2022, Grenada, May 2–6, 2022, Revised Selected Papers*. Springer, 296–315.
- [23] Rachid Guerraoui, Nikola Knežević, Vivien Quéma, and Marko Vukolić. 2010. The next 700 BFT protocols. In *Proceedings of the 5th European conference on Computer systems*. 363–376.
- [24] Guy Golan Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. 2019. Sbft: a scalable and decentralized trust infrastructure. In *2019 49th Annual IEEE/IFIP international conference on dependable systems and networks (DSN)*. IEEE, 568–580.
- [25] Bingyong Guo, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. 2020. Dumbo: Faster Asynchronous BFT Protocols. In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9–13, 2020*, Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM, 803–818.
- [26] Rüdiger Kapitza, Johannes Behl, Christian Cachin, Tobias Distler, Simon Kuhnle, Seyed Vahid Mohammadi, Wolfgang Schröder-Preikschat, and Klaus Stengel. 2012. CheapBFT: Resource-efficient Byzantine fault tolerance. In *Proceedings of the 7th ACM european conference on Computer Systems*. 295–308.
- [27] Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. 2021. All You Need is DAG. In *Proceedings of the 40th Symposium on Principles of Distributed Computing (Virtual Event, Italy) (PODC '21)*. Association for Computing Machinery, New York, NY, USA.
- [28] Idit Keidar, Oded Naor, and Ehud Shapiro. 2022. Cordial Miners: A Family of Simple, Efficient and Self-Contained Consensus Protocols for Every Eventuality. *arXiv preprint arXiv:2205.09174* (2022).
- [29] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. 2007. Zyzzyva: speculative byzantine fault tolerance. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*. 45–58.
- [30] Leslie Lamport. 2019. The part-time parliament. In *Concurrency: the Works of Leslie Lamport*. 277–317.
- [31] Yuan Lu, Zhenliang Lu, Qiang Tang, and Guiling Wang. 2020. Dumbo-MVBA: Optimal Multi-Valued Validated Asynchronous Byzantine Agreement, Revisited. In *Proceedings of the 39th Symposium on Principles of Distributed Computing (Virtual Event, Italy) (PODC '20)*. Association for Computing Machinery, New York, NY, USA, 129–138. <https://doi.org/10.1145/3382734.3405707>
- [32] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. 2016. The honey badger of BFT protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. 31–42.
- [33] Oded Naor, Mathieu Baudet, Dahlia Malkhi, and Alexander Spiegelman. 2019. Lumière: Byzantine View Synchronization. *CoRR abs/1909.05204* (2019). [arXiv:1909.05204](http://arxiv.org/abs/1909.05204) <http://arxiv.org/abs/1909.05204>
- [34] Oded Naor and Idit Keidar. 2020. Expected linear round synchronization: The missing link for linear byzantine smr. *arXiv preprint arXiv:2002.07539* (2020).
- [35] Alexander Spiegelman, Neil Giridharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. 2022. Bullshark: Dag bft protocols made practical. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. 2705–2718.
- [36] Alexander Spiegelman, Neil Giridharan, Alberto Sonnino, and Lefteris Kokoris-Kogias. 2022. Bullshark: The Partially Synchronous Version. *arXiv preprint arXiv:2209.05633* (2022).
- [37] Chrysoula Stathakopoulou, Tudor David, and Marko Vukolic. 2019. Mir-BFT: High-Throughput BFT for Blockchains. *CoRR abs/1906.05552* (2019). [arXiv:1906.05552](http://arxiv.org/abs/1906.05552) <http://arxiv.org/abs/1906.05552>
- [38] The Avalanche Team. 2020. Avalanche Networks. <https://www.avax.network/> Accessed: May 2023.
- [39] The Aptos Team. 2022. Aptos Networks. <https://aptosfoundation.org/> Accessed: May 2023.
- [40] The Celo Team. 2020. Celo Networks. <https://celo.org/> Accessed: May 2023.
- [41] The Definity Team. 2018. Definity Networks. <https://definity.network/> Accessed: May 2023.
- [42] The Diem Team. 2021. DiemBFTv4. <https://developers.diem.com/papers/diem-consensus-state-machine-replication-in-the-diem-blockchain/2021-08-17.pdf> Accessed: May 2023.
- [43] The Solana Team. 2020. Solana Networks. <https://solana.com/> Accessed: May 2023.
- [44] The Sui Team. 2023. Sui Networks. <https://sui.io/> Accessed: May 2023.
- [45] Lei Yang, Seo Jin Park, Mohammad Alizadeh, Sreeram Kannan, and David Tse. 2022. {DispersedLedger}:{High-Throughput} Byzantine Consensus on Variable Bandwidth Networks. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. 493–512.
- [46] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan-Gueta, and Ittai Abraham. 2019. HotStuff: BFT Consensus with Linearity and Responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 347–356.

A Multiple Anchors per Round

With pipelining, Shoal introduces an anchor in every round. As a result, in the best case, each anchor requires 2 rounds to commit and while non-anchor vertices require 3 rounds. Next, we present an approach to further optimize the latency for non-anchor vertices, which relies on retrospectively re-interpreting the DAG structure.

We could envision a protocol in which we iterate over more than one vertex in each round in a deterministic order and treat each vertex as an anchor. More specifically, for a vertex v in round r , we consider executing an instance of the underlying Narwhal-based consensus protocol \mathcal{P} (i.e., DAG-Rider, Tusk, and Bullshark) starting from round r with v being the first anchor. This involves re-interpreting the existing DAG structure, and potentially letting it evolve, until a decision of whether v is ordered or skipped is locally made. If v is ordered by \mathcal{P} , then the causal history of v followed by v is added to the ordering determined by the new protocol. Otherwise, v is skipped and the protocol proceeds to considering a new instantiation of \mathcal{P} from the next potential anchor (which may be in the same round).

A pseudocode in which all vertices are considered as anchors appears in Algorithm 3.

In the good case, each vertex that is considered as an anchor can be ordered in 2 rounds. However, the drawback of this approach is that if some validators are slow and a potential anchor takes many rounds to decide whether to skip or order, the progress of the whole protocol will be stalled. This happens because potential anchor vertices must be considered in an agreed-upon and deterministic order.

Algorithm 3 Every vertex as an Anchor

```
1:  $r \leftarrow 0$ 
2: while true do
3:   for each validator  $k$  do
4:     let  $v_{r,k}$  be a vertex by validator  $k$  in round  $r$ 
5:     let  $F_{r,k} : R \rightarrow A$  be a known to all mapping from
6:     rounds to anchors such that  $F_{r,k}(r) = k$ 
7:     execute  $\mathcal{P}$ , select anchors by  $F_{r,k}$ , starting from
8:      $r$  until the first ordered (not skipped) anchor  $A$ 
9:     is determined.
10:    if  $A = v_{r,k}$  then
11:      order  $A$ 's causal history according to  $\mathcal{P}$ 
12:     $r \leftarrow r + 1$ 
```

As a result, a vertex that necessitates more rounds incurs a latency penalty for the subsequent vertices.

The above issue can potentially be mitigated by combining it with a leader reputation mechanism to select the vertices that are considered as potential anchors, making the bad case delays less likely. The other vertices can be ordered based on causal history as previously.