Mahi-Mahi: Low-Latency Asynchronous BFT DAG-Based Consensus

Philipp Jovanovic[1], Lefteris Kokoris-Kogias[2], Bryan Kumara[3], Alberto Sonnino[1,2], and Pasindu Tennage[4], Igor Zablotchi[2]

[1] University College London (UCL)
[2] Mysten Labs
[3] Turing Institute
[4] EPFL

Abstract. We present Mahi-Mahi, the first asynchronous BFT consensus protocol that achieves sub-second latency in a WAN setting while processing over 100,000 transactions per second. Mahi-Mahi achieves such high performance by leveraging an uncertified structured Directed Acyclic Graph (DAG) to forgo explicit certification. This reduces the number of messages required to commit and the CPU overhead for certificate verification significantly. Mahi-Mahi introduces a novel commit rule that enables committing multiple blocks in each asynchronous DAG round. Mahi-Mahi can be parametrized either with a 5 message commit delay, maximizing the commit probability under a continuously active asynchronous adversary, or with a 4 message commit delay, reducing latency under a more moderate and realistic asynchronous adversary. We demonstrate safety and liveness of Mahi-Mahi in a Byzantine context for all of these parametrizations. Finally, we evaluate Mahi-Mahi in a geo-replicated setting and compare its performance to state-of-the-art asynchronous consensus protocols, showcasing Mahi-Mahi's significantly lower latency.

1 Introduction

Consensus enables a set of replicas to agree on a common value from an initial set of proposals, even in the presence of failures or malicious actors [13]. It is a fundamental primitive for maintaining strong data consistency between replicas in distributed or decentralized systems.

Applications that require Byzantine Fault Tolerant (BFT) consensus [7], such as blockchains [3,7,9,10,29,40], often rely on protocols designed for the partially synchronous network model which aims to approximate mostly benign network conditions and allows the system to perform well under these circumstances. However, protocols designed for partial synchrony lose liveness under asynchronous conditions, which can arise from poor connectivity, an active network adversary, or denial-of-service (DoS) attacks [26]. Asynchronous consensus protocols [13,15,37] address this issue by providing as much liveness as the network connectivity allows. To achieve this, these protocols sacrifice performance during periods of network synchrony, resulting in significantly higher latency compared to their partially synchronous counterparts. While

state-of-the-art partially synchronous protocols can process over $100,000$ transactions per second with sub-second WAN latency [4,5], current asynchronous protocols achieve similar throughput with latencies on the order of seconds [18]. This substantial latency drawback has made asynchronous consensus protocols less attractive for practical deployment.

Dual-mode protocols [25,43] attempt to provide the best of both worlds by operating partially-synchronous consensus by default and reverting to a less performant asynchronous sub-protocol when network conditions become adverse. However, these dual-mode protocols introduce complexity and are prone to errors, as they must maintain two separate protocol stacks and implement mechanisms to detect changing network conditions and switch between the two consensus modes. Additionally, they remain vulnerable to targeted attacks that can cause the protocol to switch constantly between the two modes [43]. Due to these drawbacks, no dual-mode protocol has yet been deployed in a production environment to the best of our knowledge.

We therefore ask if it is possible to design a protocol that can simultaneously: (i) provide liveness under asynchronous network conditions, (ii) achieve performance comparable to state-of-the-art partially-synchronous consensus protocols, and (iii) maintain a simple design that allow for effective security analysis, implementation, and maintenance?

In this paper, we introduce Mahi-Mahi, a novel low-latency and high-throughput asynchronous consensus protocol that simultaneously achieves these goals. Mahi-Mahi accomplishes this through a combination of the following techniques. (1) While state-of-the-art asynchronous protocols, such as Tusk [18], operate over a certified Directed Acyclic Graph (DAG) and attempt to commit one leader block every 9 message delays, Mahi-Mahi utilizes an *uncertified* DAG as its core data structure. This approach eliminates the overhead associated with the reliable broadcast [13] of DAG vertices and allows Mahi-Mahi to commit most blocks with only five message delays, aligning with the theoretical results of Cordial Miners [28]. (2) Mahi-Mahi introduces a novel commit rule that enables the commitment of multiple leader blocks in each DAG round while ensuring safety and liveness in the presence of an asynchronous adversary. (3) Mahi-Mahi also explores more practical network assumptions and can be parameterized to further enhance average-case performance while maintaining liveness against a classic asynchronous adversary.

We implement Mahi-Mahi in Rust and show that it can process an impressive $350,000$ transactions per second in geo-distributed environments with 50 nodes, all while keeping latency below 2 seconds. Additionally, Mahi-Mahi can process $100,000$ transactions per second with latency below 1 second. This achievement sets a new record in the realm of asynchronous consensus protocols and was previously only attainable by partially synchronous protocols [4,5,41]. We further show that Mahi-Mahi maintains the same throughput while improving latency over recent state-of-the-art protocols, Tusk [18] and Cordial Miners [28] – for which we provide the first known implementation – achieving latency reductions of over 70% and 30%, respectively.

Contributions. This paper makes the following contributions:

– We introduce Mahi-Mahi, the first asynchronous consensus protocol capable of committing with sub-second latency while maintaining high throughput. Notably,

Mahi-Mahi is the first DAG-based asynchronous consensus protocol capable of committing multiple leader blocks in each round.
- We provide detailed algorithms and formal security proofs for Mahi-Mahi, demonstrating its safety and liveness under an asynchronous network model.
- We conduct a formal latency analysis of Mahi-Mahi, evaluating its commit probability under various network conditions.
- We present an implementation and evaluation of Mahi-Mahi, comparing it to other state-of-the-art protocols and demonstrating that Mahi-Mahi achieves the lowest commit latency among available asynchronous consensus protocols.

2 System Overview

We present an overview of Mahi-Mahi and the settings in which it operates.

2.1 Threat model, goals, and assumptions

We consider a message-passing system with $n = 3f + 1$ validators processing transactions using the Mahi-Mahi protocol. An adversary can adaptively corrupt up to f validators, referred to as *Byzantine*, who may deviate arbitrarily from the protocol. The remaining validators, called *honest*, follow the protocol. The adversary is computationally bounded, ensuring that standard cryptographic properties such as the security of hash functions, digital signatures, and other primitives hold. Under these assumptions, Mahi-Mahi is *safe* as no two correct validators commit different sequences of transactions. The communication network is asynchronous and messages can be delayed arbitrarily, but messages among honest validators are eventually delivered. Given these conditions Mahi-Mahi is *live*, meaning honest validators eventually commit transactions. We provide proofs in Appendix C.

Furthermore, we analyze Mahi-Mahi under the *random network model* [18], a variant of the asynchronous network model. While the asynchronous model makes the worst-case assumption that the adversary has perpetually full control over the message schedule (*i.e.,* the order in which messages are received by honest validators), the random network model assumes that the message schedule is random (we give a more concrete definition in Section 2.3). We analyze Mahi-Mahi with parameters optimized for the random network model, representing an average-case evaluation. Our empirical results show that this parameterization generally outperforms a version of Mahi-Mahi configured for maximum resilience against an asynchronous adversary, all while maintaining safety and liveness guarantees.

Mahi-Mahi solves Byzantine Atomic Broadcast (BAB) [16], enabling validators to reach consensus on a sequence of messages necessary for State Machine Replication (SMR). According to the FLP impossibility result [35], BAB cannot be solved deterministically in an asynchronous setting. To address this, we employ a global perfect coin to introduce randomization, similar to previous work [11,14,27,32]. This coin can be constructed using an adaptively secure threshold signature scheme [6,12], with the distributed key setup performed under fully asynchronous conditions [1,2,20,21,30].

More formally, each validator v_k broadcasts messages by invoking $\mathsf{a\_bcast}_k(m, q)$, where m is the message and $q \in \mathbb{N}$ is a sequence number. Every validator v_i has an output $\mathsf{a\_deliver}_i(m, q, v_k)$, where m is the message, q is the sequence number, and v_k is the identity of the validator that initiated the corresponding $\mathsf{a\_bcast}_k(m, q)$. MAHI-MAHI implements a BAB protocol guaranteeing the following properties [27]:

– **Validity:** If an honest participant v_k calls $\mathsf{a\_bcast}_k(m, q)$, then every honest participant v_i eventually outputs $\mathsf{a\_deliver}_i(m, q, v_k)$, with probability 1.
– **Agreement:** If an honest participant v_i outputs $\mathsf{a\_deliver}_i(m, q, v_k)$, then every honest participant v_j eventually outputs $\mathsf{a\_deliver}_j(m, q, v_k)$ with probability 1.
– **Integrity:** For each sequence number $q \in \mathbb{N}$ and participant v_k, an honest participant v_i outputs $\mathsf{a\_deliver}_i(m, q, v_k)$ at most once, regardless of m.
– **Total Order:** If an honest participant v_i outputs $\mathsf{a\_deliver}_i(m, q, v_k)$ and $\mathsf{a\_deliver}_i(m', q', v_k')$ where $q < q'$, all honest participants output $\mathsf{a\_deliver}_j(m, q, v_k)$ before $\mathsf{a\_deliver}_j(m', q', v_k')$.

2.2   Intuition behind the MAHI-MAHI design

MAHI-MAHI builds upon DAG-based consensus protocols that achieve high throughput by processing $O(n)$ blocks per round and fully utilizing network resources. While maintaining these throughput advantages, MAHI-MAHI focuses on reducing latency in asynchronous state machine replication. It introduces novel techniques to decrease the number of message delays required for block commitment and explores more practical network assumptions to further improve average-case performance.

State-of-the-art asynchronous protocols, such as Tusk [18], operate over a certified DAG and try to commit one leader block every three certified rounds, requiring three message delays to certify each round. This results in at least nine message delays. To reduce latency, MAHI-MAHI operates over an uncertified DAG by forgoing the reliable broadcast [13] of DAG vertices, committing most blocks with only five message delays which matches the theoretical results of Cordial Miners [28]. This approach significantly reduces both bandwidth and compute cost, as validators send their blocks to every other validator only once per round, and they avoid the need to verify cryptographic certificates resulting from consistent broadcast.

This, however, creates the first challenge (**Challenge 1**): handling equivocations practically. Unlike certified DAG protocols [18,24,27,43,47], MAHI-MAHI cannot rely on certificates to prevent equivocations, necessitating the design of a novel commit rule to tolerate them. Cordial Miners [28] also face this challenge, but they address it by eventually excluding Byzantine validators that provably equivocate, which can take a long time in asynchrony.

While having five rounds between leaders provides a good probability of committing in asynchronous conditions, it also results in relatively high latency, which is not necessary for ensuring safety. Although it can be shown that we can implement a commit rule that operates in just three message delays (see Appendix C), this approach would not work in our network models as it would (1) introduce significant latency variance and (2) lose liveness in the asynchronous model. Instead, we focus on

addressing (**Challenge 2**): developing a commit rule that effectively reduces average-case latency without sacrificing worst-case liveness. We find that it is possible to reduce the number of rounds to four, achieving a balance between average-case latency in random network conditions and worst-case latency in the classic asynchronous model.

Even with this enhancement, committing only once every four message delays still results in significant latency variance for transactions that are not part of a committed leader block. A primary goal for MAHI-MAHI is to commit multiple blocks in each round, which would help ensure that the system's tail latency aligns more closely with the four-message delay. To achieve this, we need to address (**Challenge 3**): commit every block directly without relying on a sufficient round difference between leader blocks. If MAHI-MAHI were to adopt a traditional recursive commit rule [18,28], which mandates that each leader block always references all previous leader blocks in their causal history, it would at best be able to commit once every four rounds. However, MAHI-MAHI recognizes that this causal reference is only necessary when there is no sufficient evidence to directly commit a block, which is not the typical case (Section 5). This insight indicates that the recursive commit rule used in prior research is overly conservative in its approach to skipping blocks, leading to unnecessary delays, particularly during benign node crashes, which are immediately identifiable. To resolve this issue, we propose a new commit rule capable to promptly determine for each block whether it can be committed or discarded as soon as that decision is evident.

Section 3.2 presents the MAHI-MAHI commit rule that addresses these challenges. As a result, MAHI-MAHI is the first BFT consensus protocol capable of committing multiple blocks per round in the average case, while ensuring both safety and liveness in the asynchronous and random network models.

2.3   Structure of the MAHI-MAHI DAG

We present the structure of the MAHI-MAHI DAG, building an uncertified DAG that offers similar guarantees to a certified DAG, as shown in related work [5,18,28].

The MAHI-MAHI protocol operates in a sequence of logical *rounds*. In each round, every honest validator proposes a unique signed *block*, while Byzantine validators may attempt to equivocate by sending multiple blocks or none at all. During a round, validators receive transactions from users and blocks from other validators, which they refer into their proposed blocks. A block includes hash references to blocks from prior rounds, starting with their most recent block, and adds *fresh transactions* not yet included in preceding blocks. Once a block references at least $2f + 1$ blocks from the previous round, the validator signs it and broadcasts it. Clients send transactions to a validator, who adds them to their blocks. If a transaction does not finalize quickly enough, the client sends it to a different validator.

Block creation and validation. A block must include at least the following elements: (1) the author A of the block and their signature on the block contents; (2) a round number R; (3) a list of transactions; (4) at least $2f + 1$ distinct hashes of valid blocks from the previous round $R - 1$, along with potentially others from prior rounds; and (5) a share of a global perfect coin. As already mentioned the coin can be reconstructed from any $2f + 1$ shares.
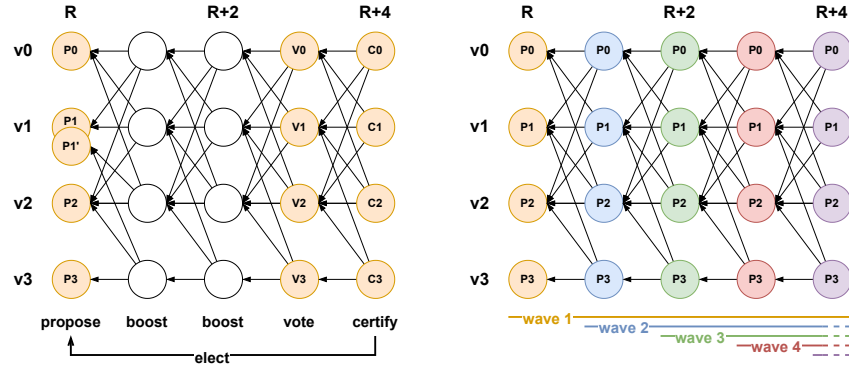
Fig. 1: The structure of the Mahi-Mahi DAG. Left: The structure of a wave, consisting of 5 rounds (Propose, Boost, Boost, Vote, Certify). Right: Waves patterns in the Mahi-Mahi protocol (each round starts a new overlapping wave).

A block is *valid* if: (1) the signature is valid and the author A is part of the validator set; (2) all hashes point to distinct valid blocks from previous rounds, and the sequence of past blocks includes $2f+1$ blocks from the previous round $R-1$; and (3) the share of the global perfect coin is valid[5]. Honest validators only include valid blocks into their DAG and discard invalid ones. Furthermore, honest validators only include hashes of blocks once they have downloaded their entire causal history, ensuring that they have successfully validated the block's causal history.

Rounds and waves. Figure 1 (left) illustrates an example of a Mahi-Mahi DAG with four validators, (v_0, v_1, v_2, v_3) when parametrized to commit in 5 rounds. For the practically efficient 4-round Mahi-Mahi, the second Boost round is omitted.

In its 5-rounds configuration, Mahi-Mahi defines a *wave* of 5 rounds for every block. The first round (Propose) includes the blocks that the wave attempts to commit (P_0, P_1, P_2, P_3) and the equivocating block P_1'. The second and third rounds (Boost) act as a buffer, helping to propagate these blocks to as many validators as possible. In the fourth round (Vote), every block serves as a *vote* for the first block of the Propose that it encounters when performing a depth-first search following the block hash references. In the example shown in this figure, blocks V_0, V_1, and V_2 are votes for P_0, P_1, P_2 (but not for P_1' and P_3), while block V_3 is a vote for P_0, P_1, P_2, and P_3 (but not for P_1'). The procedure IsVote($\cdot$) of Algorithm 3 (Appendix A) formally defines a vote. The fifth round (Certify) reveals which blocks from the Propose round have been implicitly certified. A block from the Propose round is considered *certified* or *has a certificate* if a block from the Certify round contains in its causal history at least $2f + 1$ blocks from the Vote round that are a vote for the block. In this example, blocks C_0, C_1, C_2, and C_3 serve as certificates for P_0, P_1, and P_2. This round also opens the global perfect coin, which the decision rule (Section 3.2) uses to circumvent the FLP result and to commit blocks under asynchrony by electing some of the Propose blocks as *leaders*. Similar

---

[5] Each individual share of the coin can be independently verified if the coin is implemented through a threshold signature.

to related work [18,27] this strategy selects leaders "after the fact" to deter a network adversary from strategically delaying leader blocks so that they are not referenced by blocks of the Vote round.

As illustrated in Figure 1 (right), MAHI-MAHI initiates a new wave every round. The rounds of each wave follow a consistent pattern: Propose round: R, Boost round: $R+1$, Boost round: $R+2$, Vote round: $R+3$, and Certify round: $R+4$. This pattern repeats continuously, with each new round starting a fresh wave. Algorithm 2 of Appendix A formally defines a wave.

Random network model. We analyze MAHI-MAHI in the standard asynchronous network model, as well as in the more practical [5] random network model [18]. In the asynchronous model, the adversary chooses which blocks are received by each honest validator at each round. In contrast, the random network model assumes that at each round $R+1$, an honest validator receives and references valid round-R blocks from a *uniformly random* subset of $2f + 1$ validators. Appendix C provides further details and analyses the commit probability of MAHI-MAHI in both models.

3   The MAHI-MAHI Protocol

We present MAHI-MAHI configured with a wavelength of 5 rounds. A configuration of MAHI-MAHI with a wavelength of 4 rounds operates similarly, but omits one Boost round, and addresses **challenge 2** of Section 2 as we empirically show in Section 5.

3.1   Proposers and anchors

MAHI-MAHI leverages a perfect global coin to define several *leader slots* per round. A leader slot is a tuple (validator, round) and can be either empty or contain the validator's proposal for the respective round. If the validator is Byzantine, the slot may also contain more than one (equivocating) block. In line with related work [5], the slot can assume one of three states: commit, skip, or undecided. All slots are initially set to undecided and the goal of the protocol is to classify them as commit or skip. The number of leader slots instantiated per round and the number of boost rounds can be configured (Section 5 explores different configurations).

3.2   The MAHI-MAHI decision rule

We present the decision rule of MAHI-MAHI leveraging an example protocol run. Appendix A provides detailed algorithms and Appendix B provides a complete step-by-step protocol execution. Figure 2 illustrates an example of a local view of a MAHI-MAHI validator, in a system with four validators, (v_0, v_1, v_2, v_3) and parameterized with two leader slots per round. In this example, we refer to blocks using the notation $B_{(v_i,R)}$, where v_i is the issuing validator and R is the block's round.

All proposer slots are initially in the undecided state. The validator holds the portion of the DAG depicted in Figure 2 and attempts to classify as many blocks in the leader slots as possible as either commit or skip.
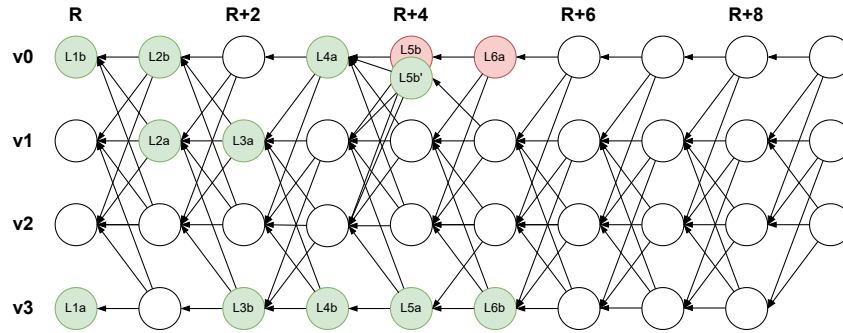
Fig. 2: Example execution with 4 validators, wave length of 5 rounds and 2 leader slots per round.

Step 1: Determine the leader slots. The validator begins by reconstructing the global perfect coin to determine the leader slots for each round. As shown in Figure 1 (Left), the coin shares embedded in round $R + 4$[6] (the Certify round) deterministically establish the leader slots for round R (the Propose round).

In this example, the validator reconstructs the coin from any set of $2f + 1$ blocks from round $R+4$ of a wave, then uses it as a seed to deterministically select two leader slots for round R: L_{1a} and L_{1b}, as illustrated in Figure 2. The coin also imposes an order between these two slots: by convention, L_{1a} is the *first* leader slot and L_{1b} is the *second* leader slot of round R. The validator repeats this process for every subsequent wave, determining leader slots L_{2a} and L_{2b} from the coin shares in round $R + 5$, L_{3a} and L_{3b} from those in round $R + 6$, and so on. The validator then sorts these leader slots in descending order: $[L_{6b}, L_{6a}, L_{5b}, L_{5a}, L_{4b}, L_{4a}, L_{3b}, L_{3a}, L_{2b}, L_{2a}, L_{1b}, L_{1a}]$.

This mechanism of determining multiple, potentially empty, leader slots from a global perfect coin is the first step towards addressing **challenge 3** (Section 2). Even if validators have different views of the DAG, they will still deterministically try to decide the same leader slots, in the same order, for a given round—regardless of whether they have a block for that slot in memory. This enables Mahi-Mahi to achieve low latency by electing more than one leader per round and using these slots to order the blocks they causally refer, as described below.

Step 2: Direct decision rule. The validator attempts to classify each slot, even those without a block as either `commit` or `skip`. To do so, the validator processes each slot individually, starting with the highest (L_{6b}), applying the Mahi-Mahi *direct decision rule*. The validator classifies a block B in a slot as `skip` if it observes $2f +1$ blocks from the subsequent Vote round that do not encounter B when performing a depth-first search following the blocks' references, and as `commit` if it observes $2f + 1$ *certificates* over it. As discussed in Section 2, a certificate over a block B is a block from the Certify round that references at least $2f + 1$ blocks from the Vote round, each of which encounter B when performing a depth-first search starting at the voting block. Otherwise, the validator leaves the slot as `undecided` (for now).

---

[6] Or $R + 3$ when Mahi-Mahi is configured with a wave length of 4 rounds.

In this example, the validator targets L_{6b} first. It observes that $B_{(v_0,R+9)}, B_{(v_1,R+9)}$, and $B_{(v_2,R+9)}$ are certificates for L_{6b}. Therefore, it classifies L_{6b} as commit. Section 5 shows that this scenario is the most common (in the absence of an asynchronous adversary) and results in the lowest latency. The validator then targets L_{6a} and observes that $B_{(v_1,R+8)}, B_{(v_2,R+8)}$, and $B_{(v_3,R+8)}$ do not vote for it. Therefore, it classifies L_{6a} as skip. The presence of $2f+1$ blocks from the Vote round that do not vote for a block ensures that it will never be certified, and will thus never be committed by other validators with a potentially different local view of the DAG. Section 5 shows that this rule allows Mahi-Mahi to promptly skip (benign) crashed leaders to minimize their impact on the protocol's performance.

Malicious validators may attempt to equivocate by creating multiple blocks for the same slot, such as L_{5b} and L'_{5b} in this example. However, the direct decision rule ensures that at most one of these blocks will be classified as commit, while the others will be classified as skip. In this example, the block $B_{(v_0,R+7)}$ is a vote for L_{5b} (and not for L'_{5b}) as it is the first block of the slot encountered when performing a depth-first search starting at $B_{(v_0,R+7)}$ and recursively following all blocks in the sequence of block hashes. Conversely, $B_{(v_1,R+7)}$, $B_{(v_2,R+7)}$, and $B_{(v_3,R+7)}$ are votes for $L_{5b'}$.

This strategy addresses **challenge 1**. Even though Byzantine validators might equivocate by creating multiple blocks per slot, the causal references defined by the DAG allow the validator to interpret blocks from the Certify round as certificates for blocks from the Propose round. Coupled with the rule that honest validators author at most one block per round, this ensures that at most one block per slot receives a certificate, while all possible other equivocating blocks are skipped. In essence, Mahi-Mahi embeds the execution of a Byzantine consistent broadcast [13] into the DAG.

Step 3: Indirect decision rule. In the (rare) case where the direct decision rule cannot classify a slot, the validator uses the Mahi-Mahi *indirect decision rule*. This rule looks at future slots to decide about the current one. First, it finds an *anchor*. This is the first block of the next wave (that is, the earliest slot with a round number $R' > R+4$) that is either still classified as undecided or already classified as commit. If the anchor is undecided, the validator marks the current slot as undecided. If the anchor is commit, the validator checks if it references at least one certificate over the current slot. If it does, the validator marks the current slot as commit. If it does not, the validator marks the current slot as skip. Appendix C shows the direct and indirect decision rules are consistent, namely if one validator direct commits a block no honest validators will indirect skip it (and vice versa).

In this example, the validator fails to classify L_{1a} using the direct decision rule as there is only one certificate for L_{1a} and thus searches for its anchor. Since L_{6a} has been classified as skip, it cannot serve as an anchor; therefore, L_{6b} becomes the anchor for L_{1a}. Given that block $B_{(v_3,R+4)}$, which serves as a certificate for L_{1a}, is referenced in L_{6b}'s causal history, the validator classifies L_{1a} as commit.

This rule is the last step to solving **challenge 3**. It allows the validator to indirectly decide on a block by leveraging the earliest anchors rather than waiting for the next leader slot which may come much later. This enables Mahi-Mahi to eliminate the need for non-leader blocks between leader slots, achieving low latency by electing leader slots in every round.

Step 4: Leader slots sequence. After processing all slots, the validator derives an ordered sequence of the blocks contained in the leader slots. It then iterates over this sequence, committing all slots marked as `commit` and skipping all slots marked as `skip`. This process continues until the validator encounters the first `undecided` slot. As demonstrated in Appendix C, this commit sequence is safe, and eventually, all slots will be classified as either `commit` or `skip`.

In the example shown in Figure 2, the leader sequence output by the validator is $[L_{1a}, L_{1b}, L_{2a}, L_{2b}, L_{3a}, L_{3b}, L_{4a}, L_{4b}, L_{5a}, L'_{5b}, L_{6b}]$. Appendix B provides a detailed walkthrough of the decision rule applied to the example DAG in Figure 2, guiding the reader step-by-step through deriving this commit sequence.

Step 5: Commit sequence. Following the approach introduced by DagRider [27], the validator linearizes the blocks within the sub-DAG defined by each leader block by performing a depth-first search. If a block has already been linearized by a previous leader slot, it is not re-linearized. The validator processes each leader slot sequentially, ensuring that all blocks are included in the final commit sequence in the correct order, according to their causal dependencies. The procedure LinearizeSubDags($\cdot$) of Algorithm 3 (Appendix A) formally describes this process.

In this example, L_{1a} and L_{1b} do not define any sub-DAG (the example begins at round R) and are thus directly added to the commit sequence. Next, L_{2a} defines the sub-DAG $\{L_{1b}, B_{(v_1,R)}, B_{(v_2,R)}, L_{2a}\}$, which is linearized as $[B_{(v_1,R)}, B_{(v_2,R)}, L_{2a}]$ since L_{1b} is already part of the commit sequence. The validator continues this process for each leader in the sequence, linearizing the sub-DAGs defined by L_{3b}, then L_{3a} and so forth following the procedure above. The final commit sequence is [L_{1a}, L_{1b}, $B_{(v_1,R)}, B_{(v_2,R)}, L_{2a}, L_{2b}, B_{(v_2,R+1)}, L_{3a}, B_{(v_3,R+1)}, L_{3b}, B_{(v_0,R+2)}, B_{(v_2,R+2)}, L_{4a}$, $L_{4b}, B_{(v_1,R+3)}, B_{(v_2,R+3)}, L_{5a}, L'_{5b}, B_{(v_1,R+4)}, B_{(v_2,R+4)}, L_{6b}$].

4   Implementation

We implemented a networked, multi-core Mahi-Mahi validator in Rust by forking the Mysticeti codebase [31], consisting of about $14,000$ LOC. Our implementation utilizes `tokio` [45] for asynchronous networking and employs raw TCP sockets for communication. We rely on `ed25519-consensus` [46] for asymmetric cryptography and `blake2` [38] for cryptographic hashing. To ensure data persistence and crash recovery, we implemented a Write-Ahead Log (WAL) tailored to the unique requirements of our consensus protocol. Furthermore, we implemented Cordial Miners [28], a state-of-the-art DAG-based asynchronous consensus protocol, using the same system components. This enabled us to to perform a comparative evaluation with Mahi-Mahi, see Section 5. Since the Cordial Miners paper lacks both implementation and evaluation, we believe our implementation and evaluation are additional contributions of our work. We are open-sourcing both our implementations of Mahi-Mahi and Cordial Miners, along with our orchestration tools, to ensure reproducibility of our results[7].

---

[7] `https://github.com/PasinduTennage/mahi-mahi-consensus`   (commit `13a2819800ad3b192ab3c54dfe3ec4d0b34ae361`)

5 Evaluation

We evaluate the throughput and latency of Mahi-Mahi through experiments conducted on Amazon Web Services (AWS), demonstrating its performance improvements over the state-of-the-art. We evaluate Mahi-Mahi with different parametrizations, with a wave length of 4 and 5 and with different numbers of leaders per round.

We compare Mahi-Mahi with Tusk [18], as an example of certified DAG-based consensus protocol, and Cordial Miners [28], as an example of an uncertified DAG-based protocol. We choose these protocols because, to the best of our knowledge, Tusk has shown the highest throughput among all published and implemented asynchronous BFT protocols when evaluated in a geo-distributed environment. Cordial Miners, while lacking an implementation and evaluation, theoretically proves excellent latency bounds and is the protocol most similar to Mahi-Mahi. We also considered a performance comparison with other recent asynchronous consensus protocols, including Pace [48], Fin [22], ParBFT [17], and SQ [44], but ultimately decided against them. The reasons for this decision is that their implementations are either closed-source, only capable of handling a limited number of block proposals (leading to crashes under sustained load), or unable to operate in a WAN environment (resulting in deadlocks after a few seconds).

Our evaluation particularly aims to demonstrate the following claims:

C1 Mahi-Mahi has similar throughput and lower latency than the baseline state-of-the-art protocols when operating in ideal conditions.
C2 Mahi-Mahi scales well by maintaining high throughput and low latency as the number of validators increases.
C3 Mahi-Mahi has a similar throughput to, and lower latency than, Cordial Miners, when operating in the presence of (benign) crash faults.
C4 Mahi-Mahi latency decreases when increasing the number of leader slots per round (up to 3 leaders per round).
C5 Mahi-Mahi parametrized with a wave length of 4 rounds has lower latency in our geo-replicated network than when configured with a wave length of 5 rounds.

Note that evaluating the performance of BFT protocols in the presence of Byzantine faults is an open research question [8], and state-of-the-art evidence relies on formal proofs of safety and liveness (presented in Appendix C). While there is a need to robustly tolerate Byzantine faults, we note that they are rare in observed delegated proof-of-stake blockchains, as compared to crash faults which occur commonly [5].

5.1 Experimental Setup

We deploy Mahi-Mahi on AWS, using `m5d.8xlarge` instances across 5 different AWS regions: Ohio (us-east-2), Oregon (us-west-2), Cape Town (af-south-1), Hong Kong (ap-east-1), and Milan (eu-south-1). Validators are distributed across those regions as equally as possible. Each machine provides 10 Gbps of bandwidth, 32 virtual CPUs (16 physical cores) on a 3.1 GHz Intel Xeon Skylake 8175M, 128 GB memory, and runs Linux Ubuntu server 22.04. We select these machines because they provide decent performance and are in the price range of "commodity servers".
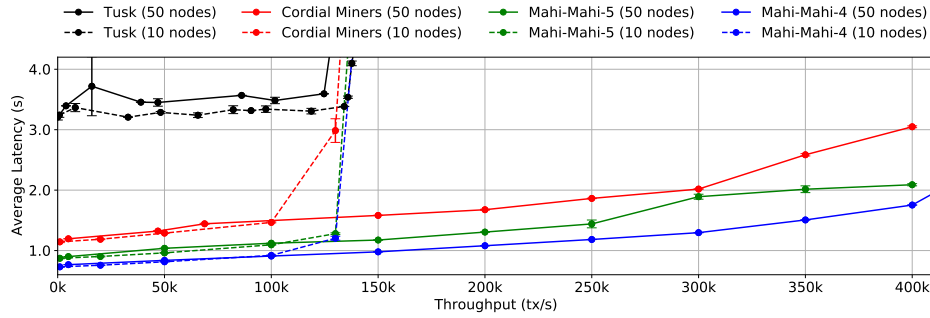
Fig. 3: Comparative throughput-latency performance of Mahi-Mahi, Tusk, and Cordial Miners. WAN measurements with 10 and 50 validators. No validator faults. 512B transaction size.

In the following, *latency* refers to the time elapsed from the moment a client submits a transaction to when it is committed by the validators, and *throughput* refers to the number of transactions committed per second. Each data point is the average latency of 3 runs and the error bars represent one standard deviation (error bars are sometimes too small to be visible on the graph). We instantiate several geo-distributed benchmark clients within each validator submitting transactions in an open loop model, at a fixed rate. We experimentally increase the load of transactions sent to the systems, and record the throughput and latency of commits. As a result, all plots illustrate the steady-state latency of all systems under low load, as well as the maximal throughput they can provide after which latency grows quickly. Transactions in the benchmarks are arbitrary and contain 512 bytes. Unless stated otherwise, we configure Mahi-Mahi with 2 leaders per round. In the following graphs, we refer to Mahi-Mahi with a wave length of 5 as Mahi-Mahi-5 and Mahi-Mahi with a wave length of 4 as Mahi-Mahi-4.

5.2   Benchmark under ideal conditions

We assess the performance of Mahi-Mahi under normal, failure-free conditions in a wide-area network (WAN) environment. Figure 3 presents the performance results of Mahi-Mahi in a geo-replicated setting, comparing both a small committee of 10 validators and a large committee of 50 validators.

For a small committee of 10 nodes, all three systems—Tusk, Cordial Miners, and Mahi-Mahi—reach a peak throughput of approximately 100k-130k transactions per second (tx/s). However, their latencies vary significantly. Tusk and Cordial Miners achieve average latencies of 3.5s and 1.5s, respectively. In contrast, Mahi-Mahi configured with a wave length 5 has a latency of 1.1s, representing a substantial reduction of 68% compared to Tusk and 27% compared to Cordial Miners. Mahi-Mahi with wave length 4 has a latency of 0.9s, representing a substantial reduction of 74% compared to Tusk and 40% compared to Cordial Miners. Tusk's higher latency stems from its certified DAG architecture, requiring at least 9 network messages to commit a block. While Cordial Miners bypasses DAG certification, it can only commit one leader every 5 rounds. In contrast, Mahi-Mahi operating with wave length 5 consistently commits multiple blocks. Mahi-Mahi operating with wave length 4 further reduces latency as it commits blocks after 4 message delays. These results validate our claim **C1**.

For a large committee of 50 nodes, Figure 3 shows that the throughput of Cordial Miners and Мані-Мані exceeds 350,000 transactions per second (tx/s), while Tusk's throughput remains around 125,000 tx/s. This perhaps surprising increase in throughput occurs because our Мані-Мані's validator implementation is optimized for large networks and does not fully utilize all available resources (network, disk, CPU) when deployed with smaller committees. Consequently, adding more validators improves resource multiplexing, boosting Мані-Мані's performance. Additionally, as the committee size grows, the number of blocks per round increases, thus a larger number of blocks are included in the causal history of elected leader blocks, without incurring additional network hops. Unlike Tusk, both Cordial Miners and Мані-Мані experience no significant CPU overhead as the committee size increases, and bandwidth does not become a bottleneck at these throughput levels. However, we do not expect further throughput gains by increasing the committee size beyond 50 nodes (such experiments would be prohibitively expensive). As expected, Cordial Miners and Мані-Мані share nearly identical throughput since both rely on the same DAG implementation, and throughput is determined by the efficiency of the DAG propagation layer.

In terms of latency, Tusk and Cordial Miners achieve average latency of 3.5s and 2.6s, respectively. Мані-Мані parametrized with a wave length of 5 has a latency of 2s (at 350,000 tx/s), which is a 42% reduction compared to Tusk and a 23% reduction compared to Cordial Miners. Мані-Мані with a wave length 4 has a latency of 1.5s, which is a 57% reduction compared to Tusk and a 42% reduction compared to Cordial miners. These results validate our claim **C2**. Comparing the two versions of Мані-Мані in those two experiments also validates our claim **C5**.

5.3   Performance under faults

Figure 4 depicts the performance of all systems when a committee of 10 validators suffers 3 crash-faults (the maximum that can be tolerated for this committee size).

We observe that all three systems achieve a throughput of approximately $35,000$-$40,000$ tx/s. Tusk and Cordial Miners record a latency of around 7s and 1.7s, respectively. Мані-Мані records a latency of 0.95s and 0.85s when running with a wave length 5 and 4, respectively. Despite the presence of faulty validators, the DAG continues to collect and disseminate transactions without significant impact. The reduction in throughput seen in Figure 4, compared to Figure 3, can be attributed to two primary factors: (1) the loss of capacity due to faulty validators, and (2) the higher frequency of missing elected leader blocks, which leads to increased commit delays. As expected, the latency advantage of Мані-Мані over Cordial Miners narrows when operating with a high number of faulty leaders, which cause head-of-line blocking and prevent the protocol from promptly committing future leaders. However, Мані-Мані still maintains a latency advantage of approximately 50% over Cordial Miners, thanks to its direct skip rule (Section 3), which allows Мані-Мані to bypass faulty leaders roughly 2 rounds earlier than Cordial Miners. Thus, our claim **C3** holds.
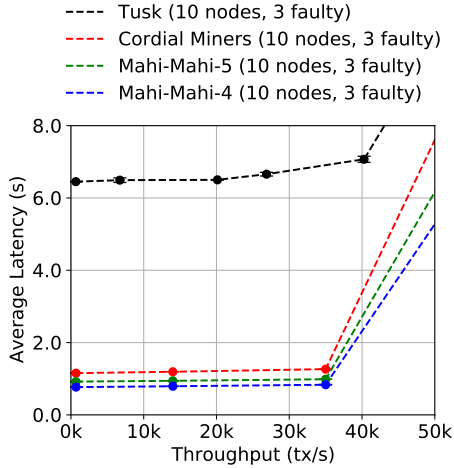
Fig. 4: Comparative throughput-latency of Mahi-Mahi, Tusk, and Cordial Miners. WAN measurements with 10 validators. Three faults. 512B transaction size.
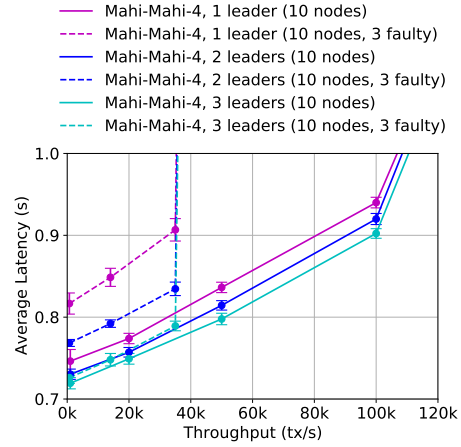
Fig. 5: Impact of the number of leaders per round in Mahi-Mahi. WAN measurements with 10 validators. Zero and three faults. 512B transaction size.

5.4   Impact of the number of leader slots per round

Finally, we assess the impact of multiple leaders on Mahi-Mahi's performance. We experiment with Mahi-Mahi parametrized with wave lengths of 4 and 5. Figure 5 illustrates how Mahi-Mahi configured with a wave length of 4 rounds performs with 1, 2, and 3 leaders under both normal conditions and scenarios involving 3 crash faults. Figure 7 in Appendix D shows the same experiment with wave length 5. We observe a notable reduction in average latency as the number of leaders increases. Specifically, when the number of leaders rises from 1 to 3, Mahi-Mahi's average latency decreases by approximately 40ms in the ideal scenario, and by approximately 100ms in the crash failure scenario. This improvement arises because having more leaders per round increases the number of blocks committed directly by leaders, rather than through the causal history of previous leader blocks. These findings validate our claim **C4**. Increasing the number of leaders beyond 3 did not further decrease latency. This is due to the higher likelihood of failing to commit via the direct decision rule, which may cause head-of-line blocking and delays the commitment of future leaders.

6   Related Work

We compare Mahi-Mahi with several categories of related work.

Uncertified DAG-based consensus protocols. The system most similar to Mahi-Mahi is Cordial Miners [28]. Like Mahi-Mahi, Cordial Miners operates over an uncertified DAG, where each vertex represents a block that is disseminated with best-effort to all peers [23]. The primary distinction between the two lies in their commit rules. Cordial Miners can commit at most one leader block every five rounds, which leads

to significantly higher latency for transactions not included in that leader block. In contrast, MAHI-MAHI's commit rule allows for a configurable number of blocks to be committed in each round, increasing the number of blocks committed per round and reducing the latency for most transactions. MAHI-MAHI commits more blocks directly through leaders, rather than relying on the causal history of previous leader blocks. Additionally, Cordial Miners does not provide an implementation or evaluation.

Mysticeti [5] is a recent protocol that, like MAHI-MAHI, operates over an uncertified DAG but in a partially synchronous setting. Mysticeti takes advantage of synchronous periods in the network to commit blocks in three rounds, and like MAHI-MAHI, it can commit blocks every round. However, unlike MAHI-MAHI and other asynchronous protocols, Mysticeti completely loses liveness when the network is not synchronous. To maintain liveness in asynchronous conditions, MAHI-MAHI interprets the DAG differently from Mysticeti. Specifically, MAHI-MAHI incorporates a global perfect coin into the protocol and modifies the role of several DAG rounds to ensure that an asynchronous adversary cannot indefinitely manipulate message schedules to prevent block certificates from forming—an issue that can easily arise in Mysticeti [26].

Certified DAG-based consensus protocols. DAG-Rider [27], Tusk [18], and Dumbo-NG [24] are popular asynchronous certified DAG-based consensus protocols that use reliable or consistent broadcast to explicitly certify every DAG vertex [36]. This approach introduces 3 message delays per DAG round but simplifies the commit rule by ensuring that equivocating DAG vertices never occur. However, this method results in significantly higher latency compared to MAHI-MAHI. For instance, DAG-Rider requires at least 12 messages to commit a block, while Tusk and Dumbo-NG require 9 messages. By contrast, MAHI-MAHI can commit in just 4 or 5 message delays when respectively configured with a wave length of 4 and 5. Also, certified DAGs have higher bandwidth and CPU requirements, as validators must disseminate, receive, and verify the cryptographic certificates generated by consistent broadcast. As shown in Section 5, these factors lead to up to 70% higher latency in comparison to MAHI-MAHI.

Sailfish [39], BBCA-Chain [33], Fino [34], Shoal [41], and Shoal++[4] build on the partially synchronous version of Bullshark [42] through various improvements, including the ability to commit more blocks per round and a relaxation of DAG certification requirements. However, these protocols are limited to partially synchronous environments and, unlike MAHI-MAHI, they lose liveness in asynchronous conditions.

Linear-chain protocols. Linear-chain asynchronous protocols such as Das *et al.* [19], Pace [48], FIN [22], and SQ [44] do not leverage an underlying DAG structure. They instead rely on explicit Byzantine consistent broadcast [13] and a common coin to elect a leader, whereas MAHI-MAHI incorporates these components implicitly within the DAG. This leader drives the protocol by constructing a linear chain. Consequently, these protocols do not achieve the same level of throughput and robustness as DAG-based systems [18]. Their contributions instead lie primarily in their theoretical foundations. For example, Das *et al.* introduces a protocol that operates without a trusted setup or the need for public-key cryptography; FIN presents the first constant-time asynchronous consensus (ACS) protocol with $O(n^3)$ messages in both information-theoretic and signature-free settings; and SQ reduces this message complexity to $O(n^2)$.

7   Conclusion

We introduce Mahi-Mahi, an asynchronous consensus protocol achieving a new performance milestone: Mahi-Mahi can process an impressive 350,000 transactions per second in geo-distributed environments with 50 nodes all while keeping latency below 2 seconds, or 100,000 transactions per second with sub-second latency—an achievement that sets a new record in the realm of asynchronous consensus protocols and that was only thought possible for partially-synchronous protocols. The exceptional performance is made possible through a novel commit rule applied over an uncertified DAG that enables commits of multiple leaders every round. This allows Mahi-Mahi to inherit the robustness and throughput inherent in DAG-based protocols, all while establishing a new standard for the latency of asynchronous consensus protocols.

References

1. Abraham, I., Jovanovic, P., Maller, M., Meiklejohn, S., Stern, G.: Bingo: Adaptivity and Asynchrony in Verifiable Secret Sharing and Distributed Key Generation. In: Advances in Cryptology – CRYPTO 2023 (2023)
2. Abraham, I., Jovanovic, P., Maller, M., Meiklejohn, S., Stern, G., Tomescu, A.: Reaching Consensus for Asynchronous Distributed Key Generation. Distributed Computing **36** (2023)
3. Al-Bassam, M., Sonnino, A., Bano, S., Hrycyszyn, D., Danezis, G.: Chainspace: A sharded smart contracts platform. arXiv preprint arXiv:1708.03778 (2017)
4. Arun, B., Li, Z., Suri-Payer, F., Das, S., Spiegelman, A.: Shoal++: High Throughput DAG BFT Can Be Fast! (2024)
5. Babel, K., Chursin, A., Danezis, G., Kokoris-Kogias, L., Sonnino, A.: Mysticeti: Low-Latency DAG Consensus with Fast Commit Path (2024)
6. Bacho, R., Loss, J.: On the Adaptive Security of the Threshold BLS Signature Scheme. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (2022)
7. Bano, S., Sonnino, A., Al-Bassam, M., Azouvi, S., McCorry, P., Meiklejohn, S., Danezis, G.: Consensus in the age of blockchains (2017)
8. Bano, S., Sonnino, A., Chursin, A., Perelman, D., Li, Z., Ching, A., Malkhi, D.: Twins: Bft systems made robust. In: 25th International Conference on Principles of Distributed Systems (OPODIS 2021) (2021)
9. Baudet, M., Danezis, G., Sonnino, A.: Fastpay: High-performance byzantine fault tolerant settlement. In: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies. pp. 163–177 (2020)
10. Blackshear, S., Chursin, A., Danezis, G., Kichidis, A., Kokoris-Kogias, L., Li, X., Logan, M., Menon, A., Nowacki, T., Sonnino, A., et al.: Sui lutris: A blockchain combining broadcast and consensus. arXiv preprint arXiv:2310.18042 (2023)
11. Blum, E., Katz, J., Liu-Zhang, C.D., Loss, J.: Asynchronous byzantine agreement with sub-quadratic communication. In: Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part I 18. pp. 353–380. Springer (2020)

12. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: International conference on the theory and application of cryptology and information security. pp. 514–532. Springer (2001)

13. Cachin, C., Guerraoui, R., Rodrigues, L.: Introduction to reliable and secure distributed programming. Springer Science & Business Media (2011)

14. Cachin, C., Kursawe, K., Shoup, V.: Random oracles in constantipole: practical asynchronous byzantine agreement using cryptography. In: Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing. pp. 123–132 (2000)

15. Chen, J., Sonnino, A., Kokoris-Kogias, L., Sadoghi, M.: Thunderbolt: Causal Concurrent Consensus and Execution (2024)

16. Cristian, F., Aghili, H., Strong, R., Dolev, D.: Atomic Broadcast: From Simple Message Diffusion to Byzantine Agreement. Information and Computation, Volume 118, Issue 1 **118** (1995)

17. Dai, X., Zhang, B., Jin, H., Ren, L.: ParBFT: Faster Asynchronous BFT Consensus with a Parallel Optimistic Path. In: CCS '23: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (2023)

18. Danezis, G., Kokoris-Kogias, L., Sonnino, A., Spiegelman, A.: Narwhal and Tusk: a DAG-based mempool and efficient BFT consensus. In: EuroSys '22: Seventeenth European Conference on Computer Systems, Rennes, France, April 5 - 8, 2022. pp. 34–50. ACM (2022)

19. Das, S., Duan, S., Liu, S., Momose, A., Ren, L., Shoup, V.: Asynchronous Consensus without Trusted Setup or Public-Key Cryptography (2024)

20. Das, S., Xiang, Z., Kokoris-Kogias, L., Ren, L.: Practical Asynchronous High-threshold Distributed Key Generation and Distributed Polynomial Sampling. In: 32nd USENIX Security Symposium (USENIX Security 23). pp. 5359–5376 (2023)

21. Das, S., Yurek, T., Xiang, Z., Miller, A., Kokoris-Kogias, L., Ren, L.: Practical asynchronous distributed key generation. In: 2022 IEEE Symposium on Security and Privacy (SP). pp. 2518–2534 (2022). https://doi.org/10.1109/SP46214.2022.9833584

22. Duan, S., Wang, X., Zhang, H.: FIN: Practical Signature-Free Asynchronous Common Subset in Constant Time. In: CCS '23: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. pp. 815 – 829 (2023)

23. Ford, B.: Threshold Logical Clocks for Asynchronous Distributed Coordination and Consensus. CoRR **abs/1907.07010** (2019)

24. Gao, Y., Lu, Y., Lu, Z., Tang, Q., Xu, J., Zhang, Z.: Dumbo-ng: Fast asynchronous bft consensus with throughput-oblivious latency. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. pp. 1187–1201 (2022)

25. Gelashvili, R., Kokoris-Kogias, L., Sonnino, A., Spiegelman, A., Xiang, Z.: Jolteon and Ditto: Network-Adaptive Efficient Consensus with Asynchronous Fallback. In: Financial Cryptography and Data Security: 26th International Conference, FC 2022 (2022)

26. Giuliari, G., Sonnino, A., Frei, M., Streun, F., Kokoris-Kogias, L., Perrig, A.: An Empirical Study of Consensus Protocols' DoS Resilience. In: Proceedings of the 19th ACM Asia Conference on Computer and Communications Security. pp. 1345–1360 (2024)

27. Keidar, I., Kokoris-Kogias, E., Naor, O., Spiegelman, A.: All You Need is DAG. In: PODC'21: Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing (2021)

28. Keidar, I., Naor, O., Poupko, O., Shapiro, E.: Cordial Miners: Fast and Efficient Consensus for Every Eventuality. In: 37th International Symposium on Distributed Computing (DISC 2023) (2023)

29. Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Syta, E., Ford, B.: OmniLedger: A Secure, Scale-Out, Decentralized Ledger via Sharding. In: IEEE S and P'18: Proceedings of the 39th IEEE Symposium on Security and Privacy. pp. 19–34. IEEE (2018)

30. Kokoris Kogias, E., Malkhi, D., Spiegelman, A.: Asynchronous Distributed Key Generation for Computationally-Secure Randomness, Consensus, and Threshold Signatures. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security. pp. 1751–1767 (2020)
31. Labs, M.: Mysticeti: Low-latency dag consensus with fast commit path. `https://github.com/asonnino/mysticeti` (2024)
32. Loss, J., Moran, T.: Combining asynchronous and synchronous byzantine agreement: The best of both worlds (2018)
33. Malkhi, D., Stathakopoulou, C., Yin, M.: BBCA-CHAIN: One-Message, Low Latency BFT Consensus on a DAG. In: Financial Cryptography and Data Security 2024. pp. 1–18. International Financial Cryptography Association (2024)
34. Malkhi, D., Szalachowski, P.: Maximal extractable value (mev) protection on a dag. In: 4th International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2022). pp. 1–18. Tokenomics (2022)
35. Michael J. Fischer, Nancy A. Lynch, M.S.P.: Impossibility of distributed consensus with one faulty process. Journal of ACM, Volume 32, Issue 2 **32** (1985)
36. Raikwar, M., Polyanskii, N., Müller, S.: SoK: DAG-based Consensus Protocols. In: 2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–18. IEEE (2024)
37. Ren, Z., Cong, K., Pouwelse, J., Erkin, Z.: Implicit Consensus: Blockchain with Unbounded Throughput (2017)
38. RustCrypto: Rustcrypto: Hashes. `https://github.com/RustCrypto/hashes` (2024)
39. Shrestha, N., Shrothrium, R., Kate, A., Nayak, K.: Sailfish: Towards improving latency of dag-based bft. Cryptology ePrint Archive, Paper 2024/472 (2024)
40. Sonnino, A., Bano, S., Al-Bassam, M., Danezis, G.: Replay attacks and defenses against cross-shard consensus in sharded distributed ledgers. In: 2020 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 294–308. IEEE (2020)
41. Spiegelman, A., Arun, B., Gelashvili, R., Li, Z.: Shoal: Improving dag-bft latency and robustness. In: Financial Cryptography and Data Security: 28th International Conference, FC 2024 (2024)
42. Spiegelman, A., Giridharan, N., Sonnino, A., Kokoris-Kogias, L.: Bullshark: the partially synchronous version. arXiv preprint arXiv:2209.05633 (2022)
43. Spiegelman, A., Giridharan, N., Sonnino, A., Kokoris-Kogias, L.: Bullshark: DAG BFT Protocols Made Practical. In: CCS '22: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (2022)
44. Sui, X., Wang, X., Duan, S.: Signature-Free Atomic Broadcast with Optimal O(n2) Messages and O(n1) Expected Time (2023)
45. Team, T.T.: Tokio. `https://tokio.rs` (2024)
46. de Valence, H.: Ed25519 for consensus-critical contexts. `https://crates.io/crates/ed25519-consensus` (2024)
47. Yang, L., Park, S.J., Alizadeh, M., Kannan, S., Tse, D.: {DispersedLedger}:{High-Throughput} Byzantine Consensus on Variable Bandwidth Networks. In: 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22). pp. 493–512 (2022)
48. Zhang, H., Duan, S.: Pace: Fully parallelizable bft from reproposable byzantine agreement. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. pp. 3151–3164 (2022)

A   Mahi-Mahi Algorithms

This section presents the algorithms used in Mahi-Mahi in pseudocode format. If a high-level understanding of Mahi-Mahi is sufficient, then this section can be skipped.

Algorithm 1 Mahi-Mahi Main Function

1: waveLength   ▷ Set to at least 4 (see Section 3)
2: leadersPerRound   ▷ See Section 5 for details

   // Idempotent function called by the application layer to extend the commit sequence.
3: **procedure** EXTENDCOMMITSEQUENCE($r_{committed}, r_{highest}$)
4:     $L \leftarrow$ TRYDECIDE($r_{committed}, r_{highest}$)   ▷ See Line 11 below
5:     $L_{commit} \leftarrow [\,]$ ▷ Hold decided leader sequence
6:     **for** $status \in L$ **do**
7:         **if** $status = \perp$ **then break**   ▷ Stop at the first undecided leader
8:         **if** $status = $ commit(b_{leader}) **then**
9:             $L_{commit} \leftarrow L_{commit} \,||\, b_{leader}$
10:     **return** LINEARIZESUBDAGS(L_{commit})   ▷ See Line 20 of Algorithm 3

   // Try to decide as many proposals as possible, recursively, starting from the latest proposal.
11: **procedure** TRYDECIDE($r_{committed}, r_{highest}$)
12:     $L \leftarrow [\,]$ ▷ Hold decision of each leader
13:     **for** $r \in [r_{highest}$ down to $r_{committed} + 1]$ **do**
14:         **for** $l \in [$leadersPerRound $- 1$ down to $0]$ **do**   ▷ Loop over all possible leaders
15:             $i \leftarrow r \,\%\,$ waveLength
16:             D $\leftarrow$ Decider(waveLength, i, l)   ▷ Algorithm 2
17:             $w \leftarrow$ D.WAVENUMBER(r)
18:             **if** D.PROPOSEROUND(w) $\neq r$ **then continue**   ▷ Skip if not a leader
19:             $status \leftarrow$ D.TRYDIRECTDECIDE(w)   ▷ Apply direct decision rule
20:             **if** $status = \perp$ **then**
21:                 $status \leftarrow$ D.TRYINDIRECTDECIDE(w)   ▷ Apply indirect decision rule
22:             $L \leftarrow status \,||\, L$
23:     **return** L ▷ May still contain undecided leaders

In particular, Algorithm 1 specifies the Mahi-Mahi main algorithms, Algorithm 2 the Mahi-Mahi decider instance, and Algorithm 3 contains various DAG helper functions. As a reminder, Mahi-Mahi operates with a single type of message: a block whose validity is described in Section 2.3. Validators hold these blocks in a data structure called DAG. To access the block(s) of round r authored by validator v of the DAG, we write $DAG[r, v]$. If an equivocation happened at a slot v, then $DAG[r, v]$ may return multiple blocks. To access all blocks of a given round r, we write $DAG[r, *]$.

Then entry point is the procedure EXTENDCOMMITSEQUENCE($\cdot$) (Line 3 of Algorithm 1), which is called by the application layer to extend the commit sequence. This procedure is idempotent and is called by our implementation (Section 4) every time the validator receives a new block. This procedure calls TRYDECIDE($\cdot$) (Line 11 of Algorithm 1) to classify as many blocks as possible as either commit or skip. The TRYDECIDE($\cdot$) procedure iterates over all possible leaders and invokes the decider instance (Line 16 of Algorithm 1) to classify each leader slot. The decider instance is responsible for determining the leader of a given round, certifying blocks, and classifying leader slots. The decider instance uses various helper functions, such as IsVOTE($\cdot$) (Line 1 of Algorithm 3), and IsCERT($\cdot$) (Line 11 of Algorithm 3), that are generic utilities for working with the DAG.

B   Example of Mahi-Mahi Execution

This section completes Section 3 by guiding readers through the protocol execution, step-by-step protocol, using the example illustrated in Figure 2. This figure showcases

---

Algorithm 2 MAHI-MAHI Decider Instance

1: waveLength   ▷ Set to at least 4 (see Section 3)
2: waveOffset   ▷ Offset creating overlapping waves (Section 2)
3: leaderOffset   ▷ Each decider operates on a unique leader slot

4: **procedure** WAVENUMBER(r)
5:     **return** $(r - \text{waveOffset})/\text{waveLength}$

6: **procedure** PROPOSEROUND(w)
7:     **return** $w \cdot \text{waveLength} + \text{waveOffset}$ ▷ See Figure 1

8: **procedure** CERTIFYROUND(w)
9:     **return** $w \cdot \text{waveLength} + \text{waveLength} - 1 + \text{waveOffset}$ ▷ See Figure 1

10: **procedure** VOTEROUND(w)
11:     **return** Self.CERTIFYROUND(w) $- 1$ ▷ See Figure 1

12: **procedure** LEADERBLOCK(w)
13:     $r_{propose}, r_{certify} \leftarrow$ Self.PROPOSEROUND(w), Self.CERTIFYROUND(w)
14:     $c \leftarrow$ COMBINECOINSHARES($\{b.share \text{ s.t. } b \in DAG[r_{certify}, *]\}$)   ▷ Common coin
15:     $l \leftarrow c + \text{leaderOffset}$ ▷ Modulo committee size
16:     **return** $DAG[r_{propose}, l]$ ▷ May return more than one block in case of equivocations

17: **procedure** SKIPPEDLEADER(w, b_{leader})
18:     $r_{vote} \leftarrow$ Self.VOTEROUND(w)
19:     **return** $|\{\neg\text{ISVOTE}(b, b_{leader}) \text{ s.t. } b \in DAG[r_{vote}, *]\}| \geq 2f + 1$

20: **procedure** SUPPORTEDLEADER(w, b_{leader})
21:     $r_{certify} \leftarrow$ Self.CERTIFYROUND(w)
22:     **return** $|\{\text{ISCERT}(b, b_{leader}) \text{ s.t. } b \in DAG[r_{certify}, *]\}| \geq 2f + 1$

23: **procedure** TRYDIRECTDECIDE(w)
24:     **for** $b_{leader} \in$ Self.LEADERBLOCK(w) **do**   ▷ Loop over equivocations
25:         **if** Self.SKIPPEDLEADER(w, b_{leader}) **then return** skip(w)
26:         **if** Self.SUPPORTEDLEADER(w, b_{leader}) **then return** commit(b_{leader})
27:     **return** $\perp$

28: **procedure** TRYINDIRECTDECIDE(w, S)
29:     $s_{anchor} \leftarrow$ find first $s \in S$ s.t. $r_{certify} < s.round \wedge s \neq$ skip(w)
30:     **if** $s_{anchor} =$ commit(b_{anchor}) **then**
31:         **if** $\exists b_{leader} \in$ Self.LEADERBLOCK(w) s.t. ISCERTIFIEDLINK(b_{anchor}, b_{leader}) **then**
32:             **return** commit(b_{leader})
33:         **else**
34:             **return** skip(w)
35:     **return** $\perp$ ▷ The anchor is undecided or not found

---

MAHI-MAHI with four validators, labeled as (v_0, v_1, v_2, v_3), operating over a wave-length of five rounds, with two leader slots allocated per round. As mentioned in Section 3.2, all proposer slots are initially in the undecided state. The validator holds in memory the portion of the DAG depicted in Figure 2 and attempts to classify as many blocks in the leader slots as possible as either commit or skip.

The first step for the validator is to identify the leader slots by reconstructing the global perfect coin for each round. As described in Section 3.2 (step 1), the validator derives the following leader slots:

$$[L_{6b}, L_{6a}, L_{5b}, L_{5a}, L_{4b}, L_{4a}, L_{3b}, L_{3a}, L_{2b}, L_{2a}, L_{1b}, L_{1a}]$$

Algorithm 3 DAG Helper Functions

1: **procedure** IsVote(b_{vote}, b_{leader})
2:    **function** VotedBlock(b, id, r)
3:        **if** $r \geq b.round$ **then return** $\bot$
4:        **for** $b' \in b.parents$ **do**
5:            **if** $(b'.author, b'.round) = (id, r)$ **then return** b'
6:                $res \leftarrow$ VotedBlock(b', id, r)
7:                **if** $res \neq \bot$ **then return** res
8:        **return** $\bot$
9:    $(id, r) \leftarrow (b_{leader}.author, b_{leader}.round)$
10:    **return** VotedBlock(b_{vote}, id, r) $= b_{leader}$

11: **procedure** IsCert(b_{cert}, b_{leader})
12:    $res \leftarrow |\{b \in b_{cert}.parents : \text{IsVote}(b, b_{leader})\}|$
13:    **return** $res \geq 2f + 1$

14: **procedure** IsLink(b_{old}, b_{new})
15:    **return** exists a sequence of $k \in \mathbb{N}$ blocks $b_1, \ldots, b_k$ s.t. $b_1 = b_{old}, b_k = b_{new}$ and $\forall j \in [2, k] : b_j \in \bigcup_{r \geq 1} DAG[r, *] \wedge b_{j-1} \in b_j.parents$

16: **procedure** IsCertifiedLink(b_{anchor}, b_{leader})
17:    $w \leftarrow$ WaveNumber($b_{leader}.round$)
18:    $B \leftarrow$ GetDecisionBlocks(w)
19:    **return** $\exists b \in B$ s.t. IsCert(b, b_{leader}) $\wedge$ IsLink(b, b_{anchor})

20: **procedure** LinearizeSubDags(L)
21:    $O \leftarrow [\,]$ ▷ Hold output sequence
22:    **for** $b_{leader} \in L$ **do**
23:        $B \leftarrow \{b \in \bigcup_{r \geq 1} DAG[r, *]$ s.t. IsLink(b, b_{leader}) $\wedge b \notin O \wedge b$ not already output $\}$
24:        **for** $b \in B$ in any deterministic order **do**
25:            $O \leftarrow O \,\|\, b$
26:    **return** O

Next, the validator attempts to classify each leader slot as either `commit` or `skip` using the *direct decision rule* (step 2), starting with the highest slot, L_{6b}. As shown in Figure 6a, the validator classifies L_{6b} as `commit` since it is certified by $2f + 1$ blocks from round $R + 9$, specifically $B_{(v_0, R+9)}$, $B_{(v_1, R+9)}$, and $B_{(v_2, R+9)}$.

The validator then proceeds to L_{6a}. As illustrated in Figure 6b, it classifies this block as `skip` because $2f + 1$ blocks from round $R + 8$ ($B_{(v_1, R+8)}$, $B_{(v_2, R+8)}$, and $B_{(v_3, R+8)}$) do not vote for L_{6a}.

Next, the validator examines L_{5b}. In its local view of the DAG, it encounters two equivocations for this leader slot: L_{5b} and L'_{5b}. The validator then invokes the function IsVote($\cdot$) (shown in Line 1 of Algorithm 3, Appendix A) to determine which of these equivocations, if any, receives votes from the blocks of round $R + 7$. For each block of this round, the validator conducts a depth-first search starting from the block, following its hash references to see if it first encounters L_{5b} or L'_{5b}.

In this example, the validator first targets $B_{(v_0, R+7)}$ and finds it votes for L_{5b}. Upon targeting $B_{(v_1, R+7)}$, it discovers a vote for L'_{5b}. Continuing this process with $B_{(v_2, R+7)}$ and $B_{(v_3, R+7)}$ also leads to a vote for L'_{5b}. Consequently, since there are $2f + 1$ blocks from round $R + 7$ that do not vote for L_{5b}, the validator classifies it as `skip`. And since there are at least $2f+1$ blocks from round $R+8$ ($B_{(v_0, R+8)}$, $B_{(v_1, R+8)}$, $B_{(v_2, R+8)}$, and $B_{(v_3, R+8)}$) that certify L'_{5b}, the validator classifies it as `commit`.
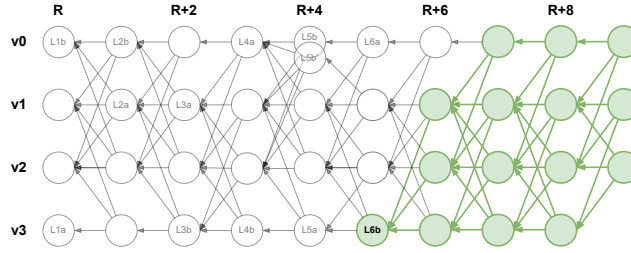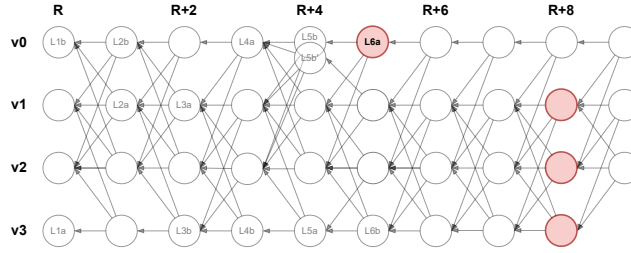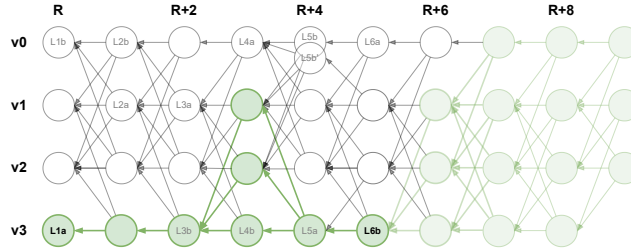
(a) Example of direct decision rule. L_{6b} is classified as commit.

(b) Example of direct decision rule. L_{6a} is classified as skip.

(c) Example of indirect decision rule. L_{1b} is classified as commit.

Fig. 6: Example of MAHI-MAHI execution with 4 validators, configured with a wave length of 5 rounds and 2 leader slots per round.

The validator then moves on to L_{5a}, classifying it as commit as it receives sufficient certification from blocks of round $R + 8$ ($B_{(v_0,R+8)}$, $B_{(v_1,R+8)}$, $B_{(v_2,R+8)}$, and $B_{(v_3,R+8)}$). Similarly, the validator classifies both L_{4b} and L_{4a} as commit since they are also certified by $B_{(v_0,R+7)}$, $B_{(v_1,R+7)}$, $B_{(v_2,R+7)}$, and $B_{(v_3,R+7)}$. This reasoning applies to the slots L_{3b} and L_{3a}, which are certified by blocks of round $R+6$, as well as to L_{2b} and L_{2a}, certified by blocks of round $R+5$. Finally, L_{1b} is certified by $B_{(v_0,R+4)}$ (through both L_{5b} and L'_{5b}), $B_{(v_1,R+4)}$, $B_{(v_2,R+4)}$, and $B_{(v_3,R+4)}$.

However, the direct commit rule fails to classify L_{1a}. There are neither $2f + 1$ blocks from round $R + 3$ that do not vote for it (which would classify it as skip) nor $2f + 1$ blocks from round $R+4$ that certify it (preventing its classification as commit). Therefore, the validator turns to the *indirect decision rule* (step 3) to classify L_{1a}.

As specified in Section 3.2, the validator first seeks the anchor for L_{1a}, which is L_{6b}. The anchor is defined as the first block with a round number $r' > r+5$ that is classified as either undecided or commit. Consequently, L_{6a} cannot serve as the anchor for L_{1b}, making L_{6b} its anchor. The validator then checks for an existing certificate over L_{1a} that is referenced by the causal history of its anchor, L_{6b}. As illustrated in Figure 6c, in this case, L_{1b} is certified by L_{5a}, which references $2f + 1$ votes for L_{1b} ($B_{(v_1,R+3)}$, $B_{(v_2,R+3)}$, and $B_{(v_3,R+3)}$). Thus, L_{1b} is classified as commit. Without such a certificate, the validator would have classified L_{1b} as skip.

C   Security Proofs

This section proves the correctness of MAHI-MAHI, by showing that MAHI-MAHI satisfies the properties of Byzantine Atomic Broadcast (BAB) from Section 2. We prove the correctness of both the 4-round and 5-round versions of MAHI-MAHI. We start with results that hold for both versions (these are mostly safety-related results) in Appendix C.1 and continue with version-specific results in Appendices C.2 and C.3.

C.1   Common Proofs for $w = 4$ and $w = 5$

We start by proving the Total Order and Integrity properties of BAB. A crucial intermediate result towards these properties is that all honest validators have consistent commit sequences, i.e., the committed sequence of one honest validator is a prefix of another's, or vice-versa. This is shown in Lemmas 5 and 6, which the following lemmas and observations build up to.

Lemma 1. *If in round r, $2f + 1$ blocks from distinct validators certify a block b, then all blocks at future rounds $r' > r$ will have a path to a certificate for b from round r.*

Proof. We prove the lemma by induction on r'. The base case is $r' = r + 1$. Let b' be a block at round r'. Since b' points to $2f + 1$ blocks at round r, by quorum intersection, b' must point to at least one of the certificates for b.

For the induction case, assume the lemma holds up to round r' and consider the case of round $r' + 1$. Let b' be a block at round $r' + 1$. By the induction hypothesis, $2f + 1$ blocks at round r' have paths to round-r certificates for b. Since b' points to $2f + 1$ blocks from round r', by quorum intersection, b' must point to at least one block that has a path to a round-r certificate for b.

Observation 1 *A block cannot vote for more than one block proposal from a given validator, in a given round.*

Proof. This is by construction. Honest validators interpret support in the DAG through deterministic depth-first traversal. So even if a block b in the vote round has paths to multiple leader round blocks from the same validator v (i.e., equivocating blocks), all honest validators will interpret b to vote for only one of v's blocks (the first block to appear in the depth-first traversal starting from b).

Lemma 2. *At most a single block per round from the same validator can be certified.*

Proof. Assume by contradiction that in a given round r, there exist two distinct blocks b and b' from the same validator v such that both b and b' are certified. This means that there exist round-$(r + w - 1)$ blocks c_b and $c_{b'}$ that certify b and b', respectively. c_b and $c_{b'}$ must point to $2f + 1$ votes for b and b', respectively. By quorum intersection, there exists an honest validator that has voted for both b and b' in the vote round. Since honest validators only produce a single block per round, this implies that there exists a block that votes for both b and b', contradicting Observation 1.

Observation 2 *If an honest validator v directly or indirectly commits a block b, then v's local DAG contains a certificate for b.*

Proof. This follows immediately from our direct and indirect commit rules.

Observation 3 *Honest validators agree on the sequence of leader slots.*

Proof. This follows immediately from the properties of the common coin, see Section 2.1.

Lemma 3. *If an honest validator v commits some block b in a slot s, then no other honest validator decides to directly skip the slot s.*

Proof. Assume by contradiction that some honest validator v' decides to directly skip s. Then it must be the case that in the local DAG of v', at least $2f + 1$ validators did not vote for b. However, since v commits b at s, by Observation 2, there must exist a certificate for b at s. So in v's local DAG there must be $2f + 1$ validators that vote for b. By quorum intersection, at least one honest validator both voted for b and did not vote for b. Since honest validators produce a single block in the vote round, this is a contradiction.

Lemma 4. *If an honest validator directly commits some block in a slot s, then no other honest validator decides to skip the slot s.*

Proof. Assume by contradiction that an honest validator v directly commits block b in slot s while another honest validator v' decides to skip s. By Lemma 3, v' cannot directly skip s; it must be the case therefore that v' skips s using the indirect decision rule. Let r be the round of s. Since v directly commits b, there exist $2f + 1$ certificates for b at s. Therefore, by Lemma 1, all blocks at rounds $r' > r + w - 1$, including the anchor of s, have a path to a certificate for b at s. Thus, v' cannot decide to skip s using the indirect decision rule. We have reached a contradiction.

Lemma 5. *If a slot s is committed at two honest validators, then s contains the same block at both validators.*

Proof. Let v and u be two honest validators and assume that v commits block b at slot s. We will show that if u commits slot s, then s contains b at s. Let w be the validator that produced block b. By Observation 2, for b to be committed at slot s at v, there must exist at least one certificate for b. By Observation 3, v and u agree that s must contain a block by w. By Lemma 2, at most a single block per round from w can be certified. So u cannot have a certificate for any other block than b at slot s.

We say that a slot is *decided* at a validator v if s is committed or skipped, *i.e.,* if it is categorized as `commit` or `skip`. Otherwise, s is *undecided*.

Lemma 6. *If a slot s is decided at two honest validators v and v', then either both validators commit s, or both validators skip s.*

Proof. Assume by contradiction that there exists a slot s such that v and v' decide differently at s. We consider a finite execution prefix and assume *wlog* that s is the highest slot at which v and v' decide differently (*). Further assume *wlog* that v commits s and v' skips s. By Lemma 3 and Lemma 4, neither v nor v' could have used the direct decision rule for s; they must both have used the indirect rule. Consider now the anchor of s: v and v' must agree on which slot is the anchor of s, since by our assumption (*) above, they make the same decisions for all slots higher than s, including the anchor of s. Let s' be the anchor of s; s' must be committed at both v and v'. Thus, by Lemma 5, v and v' commit the same block b' at s'. But then v and v' cannot reach different decisions about slot s using the indirect decision rule. We have reached a contradiction.

We have proven the consistency of honest validators' commit sequences: honest validators commit (or skip) the same leader blocks, in the same order. However, we are not done: we also need to prove that non-leader blocks are delivered in the same order by honest validators. We show this next.

Causal history & delivery conditions. Consider an honest validator v. We call the *causal history* of a block b in v's DAG, the transitive closure of all blocks referenced by b in v's DAG, including b itself. In MAHI-MAHI, a block b is delivered by an honest validator v if (1) there exists a committed leader block l in v's DAG such that b is in l's causal history (2) all slots up to l are decided in v's DAG and (3) b has not been delivered as part of a lower slot's causal history. In this case we say b is *delivered at* slot s, or *delivered with* block l.

Lemma 7. *If a block b is delivered by two honest validators v and v', then b is delivered at the same slot s, and b is delivered with the same leader block l, at both v and v'.*

Proof. Let s be the slot at which b is delivered at validator v, and l the corresponding leader block in s, also at validator v. Consider now the slot s' at which b is delivered at validator v', and l' the corresponding leader block. Assume by contradiction that $s' \neq s$. If $s' < s$, then v would have also delivered b at slot s', since by Lemma 5 must commit the same leader blocks in the same slots, so v could not have delivered b again at slot s; a contradiction. Similarly, if $s < s'$, then v' would have already delivered b at slot s, since by Lemma 5 v and v' must have committed the same block in slot s; contradiction. Thus it must be that $s = s'$, and by Lemma 5, $l = l'$.

We can now prove the main safety properties of MAHI-MAHI: Total Order and Integrity.

Theorem 1 (Total Order). *MAHI-MAHI satisfies the total order property of Byzantine Atomic Broadcast.*

Proof. This property follows immediately from Lemma 7 and the fact that honest validators order the causal histories of committed blocks using the same deterministic function, and deliver blocks in this order.

Theorem 2 (Integrity). *MAHI-MAHI satisfies the integrity property of Byzantine Atomic Broadcast.*

Proof. This is by construction: a block b is delivered as part of the causal history of a committed leader block only if b has not been delivered along with an earlier leader block (see "Causal history & delivery conditions" above). So an honest validator cannot deliver the same block twice.

We now turn to liveness properties. The following two lemmas establish that blocks broadcast by honest validators are eventually included in all honest validators' DAGs.

Lemma 8. *If a block b produced by an honest validator v references some block b', then b' will eventually be included in the local DAG of every honest validator.*

Proof. This is ensured by the synchronizer sub-component in each validator: if some validator w receives b from v, but does not have b' yet, w will request b' from v; since v is honest and the network links are reliable, v will eventually receive w's request, send b' to w, and w will eventually receive b'. The same is recursively true for any blocks from the causal history of b', so w will eventually receive all blocks from the causal history of b' and thus include b' in its local DAG.

Lemma 9. *If an honest validator v broadcasts a block b, then every correct validator will eventually include b in its local DAG.*

Proof. Since network links are reliable, all honest validators will eventually receive b from v. By Lemma 8, all honest validators will eventually receive all of b's causal history, and so will include b in their local DAG.

The following crucial lemma establishes that in any round r, there is at least one block b, called a common core, such that all blocks at round $r + 2$ have a path to b.

Lemma 10. *For any r, there is at least one block b from round r such that any valid block from round $r + 2$ has a path to b.*

Proof. Consider a set B of $2f + 1$ blocks in round $r + 1$ from honest validators. Using B, we create a table T, as follows: for blocks $b, c \in B$, let $T[b, c] = 1$ if b in $r + 1$ references c in r, $T[b, c] = 0$ otherwise. By quorum intersection, any b will reference at least $f + 1$ blocks in round r that are also in B, so each row of T has at least $f + 1$ entries equal to 1. Thus, T has at least $(2f+1)(f+1)$ entries equal to 1. By a counting argument, there is a block c^* in B that has a 1 entry in at least $f + 1$ rows, i.e., a block from round r which is referenced by $f + 1$ blocks from round $r + 1$. Let P' be the set of blocks from round $r + 1$ which reference c^*. Consider now any valid block b in round $r + 2$; b references $2f + 1$ blocks in $r + 1$, so by quorum intersection b references at least one block in B. Thus, b has a path to c^*.

C.2   Specific Proofs for $w = 5$

We continue with proofs that are specific to the liveness of the $w = 5$ version of MAHI-MAHI. We show that each wave has at least $2f + 1$ leader blocks that can be directly committed (Lemmas 11 and 12), and thus that each wave has a nonzero probability of directly committing at least one block (Lemma 13). We then show that each slot is eventually decided directly or indirectly (Lemma 14). Finally, we show that MAHI-MAHI satisfies the Validity and Agreement properties of BAB.

As a consequence of Lemma 10, we have the following:

Lemma 11. *For any r, there exists a set S of at least $2f + 1$ blocks from round r such that any valid block from round $r + 3$ is a vote for every block in S.*

Proof. Let $r' = r + 1$. By Lemma 10, there exists a block b in round $r' = r + 1$ such that any valid block from round $r' + 2 = r + 3$ has a path to b. Now let S be the set of blocks referenced by the block b. S must contain at least $2f + 1$ blocks from round r. Every block from round $r + 3$ has a path to b and thus, through b, to every block in S.

From this we can derive the following crucial lemma:

Lemma 12. *For any r, there exists a set S of at least $2f + 1$ blocks from round r such that every block in S has at least $2f + 1$ certificates in round $r + 4$.*

Proof. Take S to be the set from Lemma 11. There are at least $2f + 1$ blocks in $r + 4$. Any block b in round $r+4$ must reference $2f+1$ blocks from round $r+3$. By Lemma 11, every block from round $r + 3$ is a vote for every block in S, so b must be a certificate for every block in S.

We denote by $\ell \leq 3f + 1$ the number of leader slots per round.

Lemma 13. *Fix a round r. If $\ell > f$, then an honest validator directly commits at least one slot corresponding to round r. Otherwise, the probability that an honest validator directly commits at least one slot corresponding to round r is at least $p^\star = 1 - \frac{\binom{f}{\ell}}{\binom{3f+1}{\ell}} > 0$.*

Proof. By Lemma 12, at least $2f + 1$ blocks from round r can be directly committed, out of a maximum of $3f + 1$ blocks. When the common coin is released in round $r+4$, it selects uniformly at random ℓ round-r blocks as the ℓ slots of round r.

In the case $\ell > f$, by quorum intersection, there exists at least one slot selected by the common coin among the $2f + 1$ blocks that can be directly committed.

In the case $\ell \leq f$, we can model the number of directly committed slots in round r as a hypergeometric random variable, where a success event corresponds to selecting a slot that can be directly committed. The probability of 0 successes (i.e., not committing any slots directly) is therefore at most $\frac{\binom{f}{\ell}}{\binom{3f+1}{\ell}} < 1$.

Lemma 14. *Fix a slot s. Every honest validator eventually either commits or skips s, with probability 1.*

Proof. We prove the lemma by showing that the probability of s remaining undecided forever at some honest validator is 0. In order for s to remain undecided forever, s cannot be committed or skipped directly. Furthermore, s cannot be decided using the indirect rule. This means that the anchor s' of s must also remain undecided forever, and therefore the anchor s'' of s' must remain undecided forever, and so on. The probability of this occurring is at most equal to the probability of an infinite sequence of rounds with no directly committed slots, equal to $\lim_{t \to \infty} (1 - p^\star)^t = 0$, where $p^\star > 0$ is the probability from Lemma 13.

Theorem 3 (Validity). *MAHI-MAHI satisfies the validity property of Byzantine Atomic Broadcast.*

Proof. Let v be an honest validator and b a block broadcast by v. We show that, with probability 1, b is eventually delivered by every honest validator. By Lemma 9, b is eventually included in the local DAG of every honest validator. So every honest validator will eventually include a reference to b in at least one of its blocks. Let r be the highest round at which some honest validator includes a reference to b in one of its blocks. By Lemma 13, with probability 1, eventually some block b' at a round $r' > r$ will be directly committed. Block b' must reference at least $2f + 1$ blocks, thus at least $f + 1$ blocks from honest validators. Since all validators have b in their causal histories by round r, b' must therefore have a path to b. Lemma 14 guarantees that all slots before b' are eventually decided, so b' is eventually delivered. Thus, b will be delivered at all honest validators at the latest when b' is delivered along with its causal history.

Theorem 4 (Agreement). *MAHI-MAHI satisfies the agreement property of Byzantine Atomic Broadcast.*

Proof. Let v be an honest validator and b a block delivered by v. We show that, with probability 1, b is eventually delivered by every honest validator. Let l be the leader block with which b is delivered, and s the corresponding slot. By Lemma 14, all blocks up to and including s are eventually decided by all honest validators, with probability 1. By Lemma 5, all honest validators commit l in s. Therefore, all honest validators deliver b eventually.

C.3   Specific Proofs for $w = 4$

We now turn to the liveness of the $w = 4$ version of MAHI-MAHI. We first show that under an asynchronous network, liveness is guaranteed, albeit with a smaller probability of direct commit at each wave than in the $w = 5$ version. We later show that under a random network, MAHI-MAHI directly commits all valid leader blocks in each wave with high probability. Recall that in the random network model, a valid block from round $r + 1$ references a set of $2f + 1$ valid blocks from round r, sampled uniformly at random among all valid round-r blocks.

Lemma 15. *For any r, there exists a block b from round r such that b has at least $2f + 1$ certificates in round $r + 3$.*

Proof. By Lemma 10, there exists b in round r such that any block in round $r + 2$ has a path to b, and therefore is a vote for b. Since any block in round $r + 3$ must reference $2f + 1$ blocks in round $r + 2$, any block in round $r + 3$ must be a certificate for b. Since every honest validator publishes a block in round $r + 3$, there must exist at least $2f + 1$ certificates for b in round $r + 3$.

We again denote by $\ell \leq 3f + 1$ as the number of leader slots per round.

Lemma 16. *Assume the asynchronous network model and fix a round r. If $\ell = 3f + 1$, then an honest validator commits at least one slot corresponding to round r. Otherwise, the probability that an honest validator directly commits at least one slot corresponding to round r is at least $p^\star = \frac{\ell}{3f+1}$.*

Proof. By Lemma 15, there exists at least one block b in round r that can be directly committed. When the common coin is released in round $r + 3$, it selects uniformly at random ℓ round-r blocks as the ℓ slots of round r.

If $\ell = 3f + 1$, then all possible blocks of round r are included in the slots, and thus b is directly committed.

In the case $\ell < 3f + 1$, we can model the number of directly committed slots in round r as a hypergeometric random variable, where a success event corresponds to selecting the slot that can be directly committed. There are $3f + 1$ states in total, out of which only 1 is a success state; and there are ℓ draws. The probability of one success is therefore $\frac{\binom{1}{1}\binom{3f}{\ell-1}}{\binom{3f+1}{\ell}} = \frac{\ell}{3f+1}$.

Lemma 17. *In the random network model, with high probability, every block in round $r + 2$ is a vote for every block in round r.*

Proof. We prove the lemma by showing, through Markov's inequality, that the probability of any block in round r being unreachable from any block in round $r + 2$ approaches 0 exponentially in f.

Take a pair of blocks b_r and b_{r+2} in rounds r and $r + 2$, respectively. We compute the probability that there is no path from b_{r+2} to b_r. Block b_{r+2} must reference $2f + 1$ blocks from round $r + 1$; let b_{r+1} be one such block. The probability that b_{r+1} references b_r is at least $p = \frac{2f+1}{3f+1}$, since b_{r+1} references $2f + 1$ randomly selected blocks from round r. The probability that there is no path from b_{r+2} to b_r is therefore at most $q = (1 - p)^{2f+1}$.

We now compute the expected number of pairs of blocks b_r and b_{r+2} from rounds r and $r + 2$, respectively, such that b_r is not reachable from b_{r+2}. There are at most $(3f + 1)^2$ pairs, and each pair is not connected by a path with probability at most q, thus the expected number of unreachable pairs is at most $E = q(3f + 1)^2 = (3f + 1)^2(1 - p)^{2f+1}$.

Using Markov's inequality, the probability that there exists at least one unreachable pair is:

$$\Pr(\text{At least one unreachable pair}) \leq E = (3f + 1)^2(1 - p)^{2f+1}.$$

As f increases, this probability rapidly approaches 0 due to the exponential term.

Lemma 18. *Assume the random network model and fix a round r. With high probability, an honest validator directly commits every leader slot chosen by the common coin in round $r + 3$.*

Proof. By Lemma 17, every block in round $r + 2$ is a vote for every block in round r with high probability. Thus every block in round $r + 3$ is a certificate for every block in round r with high probability. So every honest validator sees $2f + 1$ certificates for every block in round r and thus sees also for the blocks chosen by the common coin at least $2f + 1$ certificates.

Lemma 19. *Fix a slot s. In both the asynchronous network model and the random network model, every honest validator eventually either commits or skips s with probability 1.*

Proof. The proof is analogous to the proof of Lemma 14. By Lemma 16 and Lemma 18, the probability of a honest validator directly committing any leader block in a given round is greater than 0, in both the asynchronous and the random network models. Thus, the probability of an infinite sequence of rounds without any directly committed blocks is 0. This implies that every slot that is not directly decided will eventually have a committed anchor and become decided itself.

Theorem 5 (Validity). *MAHI-MAHI satisfies the validity property of Byzantine Atomic Broadcast.*

Proof. The proof is similar to the proof of validity in the $w = 5$ case. Let v be an honest validator and b a block broadcast by v. We show that, with probability 1, b is eventually delivered by every honest validator. By Lemma 9, b is eventually included in the local DAG of every honest validator. So every honest validator will eventually include a reference to b in at least one of its blocks. Let r be the highest round at which some honest validator includes a reference to b in one of its blocks. By Lemma 16 and Lemma 18, with probability 1, eventually some block b' at a round $r' > r$ will be directly committed. Block b' must reference at least $2f + 1$ blocks, thus at least $f + 1$ blocks from honest validators. Since all validators have b in their causal histories by round r, b' must therefore have a path to b. Lemma 19 guarantees that all slots before b' are eventually decided, so b' is eventually delivered. Thus, b will be delivered at all honest validators at the latest when b' is delivered along with its causal history.

Theorem 6 (Agreement). *MAHI-MAHI satisfies the agreement property of Byzantine Atomic Broadcast.*

Proof. The proof is similar to the proof of agreement in the $w = 5$ case. Let v be an honest validator and b a block delivered by v. We show that, with probability 1, b is eventually delivered by every honest validator. Let l be the leader block with which b is delivered and s the corresponding slot. By Lemma 19, all blocks up to and including s are eventually decided by all honest validators, with probability 1. By Lemma 5, all honest validators commit l in s. Therefore, all honest validators deliver b eventually.
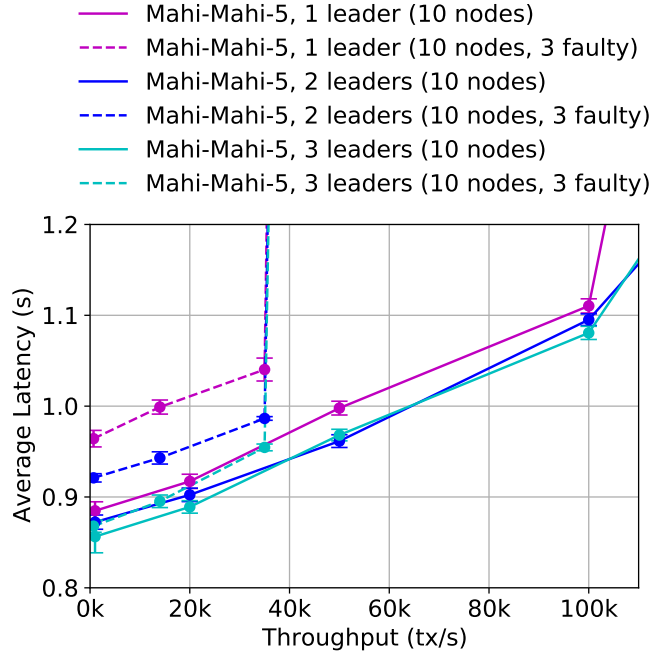
Fig. 7: Impact of the number of leaders per round in Mᴀʜɪ-Mᴀʜɪ. WAN measurements with 10 validators. Zero and three faults. 512B transaction size.

Note: Mᴀʜɪ-Mᴀʜɪ with $w = 3$**.** It is possible to configure Mᴀʜɪ-Mᴀʜɪ with 3-round waves, by removing all Boost rounds, and keeping the Propose round (r), Vote round ($r+1$) and Certify round ($r+2$). Such a protocol would still satisfy safety, as all results up to and including Theorem 2 hold when $w = 3$. However, this 3-round version of Mᴀʜɪ-Mᴀʜɪ would no longer satisfy liveness, because the common core approach in Lemma 10 can no longer be used to guarantee that at least one leader block can be directly committed in each wave.

D   Mᴀʜɪ-Mᴀʜɪ Impact of Multi-leader

This section completes Section 5 by presenting the impact of the number of leaders per round in Mᴀʜɪ-Mᴀʜɪ when implemented with 5 rounds per wave. Figure 7 illustrates how Mᴀʜɪ-Mᴀʜɪ configured with a wave length of 5 rounds performs with 1, 2, and 3 leaders per round under both normal conditions and scenarios involving 3 crash faults. We observe a latency reduction as the number of leaders increases similar to Figure 5 (Section 5). Specifically, when the number of leaders rises from 1 to 3, Mᴀʜɪ-Mᴀʜɪ's average latency decreases by approximately 40ms in ideal scenario, and by approximately 100ms in the crash failure scenario. Similarly to Figure 5, increasing the number of leaders beyond 3 did not further decrease latency.